Managing Multihoming Workers in the Gig Economy

Gad Allon
The Wharton School, University of Pennsylvania, Philadelphia, PA 19104
gadallon@wharton.upenn.edu

Maxime C. Cohen
Desautels Faculty of Management, McGill University, Montreal, QC H3A 1G5, Canada
maxime.cohen@mcgill.ca

Ken Moon
The Wharton School, University of Pennsylvania, Philadelphia, PA 19104
kenmoon@wharton.upenn.edu

Wichinpong Park Sinchaisri
Haas School of Business, University of California, Berkeley, CA 94720
parksinchaisri@haas.berkeley.edu

July 7, 2025

Gig economy platforms compete to source labor from common pools of workers, who multihome by dynamically allocating their services in real-time across multiple platforms. The question of how such workers choose between competing platforms has grown in salience. However, the unavailability of comprehensive data has limited our understanding of workers' dynamic multihoming decisions and their impact on the labor supply and operations of platforms. We address this gap by integrating a major ride-hailing platform's proprietary data regarding individual drivers' detailed trips with public data on the drivers' outside options. Using empirical methods that overcome the remaining data limitations using simulation and machine learning, we structurally estimate the perceived costs that motivate drivers' forward-looking decisions and behavior. Our analysis reveals that workers are short-sighted and value sequences of consistent rewards (e.g., jobs and pay) over highly variable ones. Based on these findings, we explore the impact of compensation and incentives on the platform's labor supply and operations. Our counterfactual analyses indicate that consistent pay dominates variable pay in retaining multihoming workers. For gig platforms seeking to maintain a stable platform workforce, they can further control their labor supply by rewarding uninterrupted work or delaying quits. For policymakers, our research gives insights regarding the design of gig economy incentive schemes and regulations, including specifically New York City's 2018 *Driver Income Rules*, and their impact on multihoming behavior.

Key words: gig economy, multihoming, structural estimation, applied generative adversarial networks, incentives, empirical operations, behavioral operations

1. Introduction

The gig economy forces platforms to compete for on-demand workers, who can flexibly choose where to work not only daily but also at every moment throughout the day. Increasingly and prevalently, such workers "multihome," or dynamically allocate their labor supply across various platforms in real-time. For example, ride-hailing drivers use multiple platforms to find and complete rides and provide additional services such as food delivery. In 2021, approximately 39% of gig workers in the United States actively work across two or more types of gig jobs, a substantial increase from 14% in 2016 (Anderson et al. 2021)

As the number of platforms has grown, each multihoming worker devotes less time and labor to any single platform. The competition for workers has further intensified in cities like New York and Seattle that have capped the total number of ride-hailing drivers who are eligible to work in order to ease traffic congestion and protect driver welfare (Fitzsimmons 2018, Wilson 2019). As such market conditions increasingly drive platforms to compete fiercely for the same common pools of multihoming workers, both platforms and policymakers must ask themselves: what rules and incentives can appropriately and desirably shape the multihoming labor supply decisions of gig workers?

At the same time, however, platforms must manage and compete for the uncertain demand of customers, who comparison-shop across different platforms based on factors such as availability and price. As a result, the prevalent model of platform operations aims to dynamically reconcile present levels of supply and demand through *pay-per-work* schemes that offer to pay workers dynamic, market-clearing wages for the tasks they complete. For instance, Uber compensates drivers for each trip they complete, at dynamic wages that depend on the time and distance, traffic conditions, and the current ratio of drivers over rider demand.

Whereas such dynamic pay-per-work mechanisms let workers flexibly decide their supply of labor in response to their individual preferences and market conditions (Chen et al. 2019), the implications for the platforms that use them are mixed. Historically, gig platforms have prioritized their operational flexibility in the face of demand uncertainty and therefore largely accepted the constantly evolving supply and wages of workers as a given. However, firms increasingly compete dynamically for a limited and multihoming labor supply, so firms are paying closer attention to how exactly workers select their gigs and choose to remain active on a given platform in order to assure that their labor supply is reliable when the demand for their services actually occurs. In fact, some ride-hailing and delivery platforms have begun compensating gig workers on an hourly basis, at dynamic rates that are guaranteed to workers as long as they remain available to immediately accept jobs. Several platforms add transparency or flexibility by announcing such hourly wage schedules in advance or allowing workers to choose from different rates themselves. Yet, the exact appeal of such pay policies is not well justified or understood in the literature. Do such rigid pay incentives respond effectively to escalation of multihoming behavior, and do they appeal to gig economy workers who are widely believed to value workplace flexibility?

Our study therefore addresses two primary research questions: (i) How do multihoming gig economy workers decide among the different work opportunities made available to them? In particular, how do they assess the *dynamic* and potentially *uncertain* evolution of both the tasks available to them and weigh the multiple, cross-platform incentives they face? Ultimately, their individual reactions to such dynamics collectively shape the gig economy's labor supply as a whole. (ii) In turn, given a multihoming labor supply, how should on-demand platforms compete for multihoming workers? What incentive scheme optimally balances a platform's response to demand uncertainty against maintaining adequate control over its available supply of labor?

To address these questions, we have collaborated with a ride-hailing platform based in New York City (NYC) and combine its proprietary data on driver activity with the public corpus of NYC ride-hailing trips reported by the Taxi and Limousine Commission (TLC). Using such data, we structurally model and estimate the preferences for work and pay that drive workers' intra-day labor supply decisions in the presence of multihoming opportunities. Our model considers the detailed time and location of drivers' sequential decisions, the pay incentives offered to individual drivers, and the contemporaneous availability and pay on the competing platform. Each driver then optimizes her dynamic decision-making. We estimate the structural model using indirect inference methods, which combine simulation with adversarial machine learning (Gourieroux et al. 1993, Kaji et al. 2023). These methods have been used previously to recover choice preferences from dynamic spatial trajectories in service networks (Moon 2023).

Using this structural empirical framework, we recover the heterogeneous discount factors and costs of working that explain drivers' forward-looking decisions. In particular, we find that drivers heavily discount their future earnings and therefore (over)react robustly, including by multihoming, to short-term income shocks. Because drivers focus heavily on immediate, short-term earnings and exhibit little patience for waiting to receive higher future earnings, 66.59% of the drivers on the partner platform multihome more often than not. Under their short-term focus, consistent pay rates give drivers a continual stream of reliable work and pay that greatly reduces their switching between platforms. In contrast, dynamic wages that clear the market may fluctuate to yield low availability or wages in the short run even if average effective wages stay high. Our partner platform offers drivers consistent wages in the form of preannounced wages that vary periodically but deterministically throughout the day. Nevertheless, we find that multihoming peaks around the scheduled times at which the dynamic pay rates change because workers are so sensitive to short-term pay.

We explore the prescriptive implications using counterfactual analyses. First, we evaluate the impact that pay schemes have on workers' multihoming decisions and on platform performance. Our main comparison pits dynamic, market-clearing wages against consistent pay rates in the form of a preannounced dynamic schedule of guaranteed wages by time of day. Offering guaranteed wages controls the multihoming that is otherwise caused by workers' short-run focus on avoiding even temporarily low effective wages; therefore, it offers gig platforms a cost-effective path to maintaining a reliable labor supply in the presence of multihoming. Second, we further evaluate two types of incentives that can enable a platform to control multihoming. Recently, several platforms have rewarded workers with "streak" or consecutive work bonuses when they maintain availability for work and complete all their assigned tasks within a pre-allotted time frame. We show that offering this type of reward plausibly discourages multihoming. Another platform policy is to delay active workers from quitting the platform. This could be achieved by introducing frictions (e.g., multiple buttons) in the quitting process or by requiring workers to signal their quits in advance. While this policy is appealingly non-monetary, we find that it could counterproductively motivate workers to switch to the competing platform earlier. Finally, we examine NYC's 2018 *Driver Income Rules*, which imposed minimum wage standards on ride-hailing platforms. We show that such minimum wage regulations can unintentionally harm multihoming workers. Effective regulations should take into account how platforms compensate workers and how workers make labor decisions.

2. Contributions and Related Work

Our work contributes to the growing stream of research in operations management on the gig economy (for a comprehensive review, see Benjaafar and Hu 2020) and the decentralized decision-making of gig workers. The majority of prior work provided analytical insights into gig economy worker decisions (e.g., Ibrahim 2018, Dong and Ibrahim 2020, Besbes et al. 2021, Benjaafar et al. 2022). Lobel et al. (2024) show that utilizing flexible contractors as workers offers key operational advantages over hiring them as employees, especially in the presence of substantial demand uncertainty. More recently, empirical researchers have studied factors that shape gig economy workers' labor decisions. Miao et al. (2022) analyze the effects of surge pricing on driver behavior using a quasi-experiment based on a ride-hailing company introducing surge pricing in two Chinese cities at different times. They find that surge pricing decreases drivers' daily earnings, encouraging more part-time drivers to work on more days compared to full-time drivers. In US-based ride-hailing, Allon et al. (2023b) empirically find that workers make platform labor decisions based on both the monetary value of incentives and their internal goals. Our work expands the study of gig worker labor supply decisions to the increasingly prevalent pattern of multihoming wherein workers choose when and where to provide services.

Multihoming behavior has been studied on consumer platforms that predate the gig economy. Recent surveys have shown that consumers typically "shop around" for similar products on other e-commerce sites before making a purchase and this is true across product categories (Critero 2017, Davies et al. 2022). The majority of academic work considers consumer multihoming in the advertising market (e.g., Athey et al. 2018, Park et al. 2021), suggesting that consumer switching leads to inefficiency in the matching of advertisers and consumers, lowering the demand and prices. A handful of prior works offer insights into provider-side multihoming. For instance, Landsman and Stremersch (2011) and Cennamo et al. (2018) show that provider multihoming in the video game industry is a result of variations in the complexity in adapting their product to different platforms.

The nature of the gig economy enables both consumers and providers to multihome across platforms and marketplaces. On the consumer side, Chitla et al. (2023) models the trade-offs that riders consider when they choose between ride-hailing platforms with potential loyalty effects. On the provider side, theory offers some guidance regarding gig workers' multihoming behavior. Specifically, by multihoming, workers can reduce the time they spend idly waiting for customers (Bryan and Gans 2019, Loginova et al. 2022). However, despite increasing their utilization and hours worked, multihoming gig workers can suffer declining earnings in equilibrium (Benjaafar et al. 2020, Bernstein et al. 2021). Tadepalli and Gupta (2020) further analyze how multihoming may alter the actions of platforms, leading to inefficient matching policies and decreasing welfare for all parties in the long term. Notwithstanding such predictions, cross-platform multihoming is empirically under-explored, and the broader literature offers a limited understanding of the ability of platforms to influence workers' decisions in the presence of multihoming opportunities.

We are among the first to model and empirically estimate multihoming behavior among workers in the gig economy. Our model captures temporal and spatial variations in demand and supply as well as competition across on-demand platforms. Closest to our work are the papers by Rosaia (2020) and Allon et al. (2023a). Rosaia (2020) presents a model of competing transportation platforms and uses New York City (NYC)'s high-frequency taxi ride data to estimate profit-maximizing allocations and prices and then go on to study policies for reducing traffic and incentivizing platforms to improve driver utilization. Unlike Rosaia (2020), our paper explores policies aimed at influencing gig workers' multihoming behavior. Using proprietary data from a driver analytics company, Allon et al. (2023a) empirically measure both multihoming and repositioning decisions among ride-hailing drivers and assess policies intended to limit multihoming or traffic congestion, whereas we focus solely on workers' multihoming behavior. We offer an in-depth study of the two pay schemes predominantly offered by gig economy platforms that compete. We additionally assess other monetary and non-monetary interventions for multihoming workers from both the platform's and policymakers' perspectives. Our structural estimation framework ultimately offers a response to the call for the integration of machine learning and behavioral science in operations management (Davis et al. 2024).

Finally, our prescriptive counterfactual analyses provide fresh insights into incentive design for gig economy platforms and policymakers. Offering a pay-per-task is generally more common in the gig economy, but a growing number of platforms have adopted guaranteed pay schemes (e.g., Browning 2023). Moving away from task-dependent compensation offers stability, which is highly valued by gig workers (Chen et al. 2019). Despite having full flexibility, gig workers tend to maintain a consistent schedule and stay on one platform over time (Allon et al. 2023b). Pay variability also leads to increased voluntary turnover, lower task completion rates and less motivated workers (Ikeda and Bernstein 2016, Cameron and Meuris 2022, Butschek et al. 2022). Platform incentive policies, in general, can be inefficient, affecting competition and not addressing workers' diverse preferences (Chen et al. 2022, Besbes et al. 2022). Recent works have proposed and demonstrated the effectiveness of non-monetary interventions, including providing feedback (Bolton et al. 2004, Athey et al. 2024), showing relative rankings (Choudhary et al. 2022, Shokoohyar and Katok 2022), limiting past performance information (Li and Zhu 2021), and introducing priority matching (Mai et al. 2023). Our analyses characterize the conditions under which income certainty could influence multihoming behavior and provide guidance for the design of optimal monetary and non-monetary incentive schemes as well as gig economy regulations.

3. Model of Gig Platforms and Worker Decisions

To study gig economy workers' multihoming decisions, we consider a market in which two platforms offer competing on-demand ride-hailing services fulfilled by gig workers. While the two platforms operate largely identically to meet customers' requests for rides, they compensate their workers differently. One platform offers guaranteed but dynamic pay to drivers based on their availability for work, while the other compensates drivers on a trip-by-trip basis with greater variability. Our model allows for spatially, differentiated, dynamic incentives, and captures how heterogeneous workers respond to these incentives depending on when and where they are working.

We denote the two competing platforms as firm A (the focal platform for our study) and firm B. Firm A guarantees drivers an hourly rate of pay while logged into firm A's platform, as long as they are available to drive customers. While drivers are logged on, Firm A instructs them where to drive even without passengers—e.g., along pickup-friendly routes. In contrast, Firm B pays drivers on a per-work basis (accounting for trip characteristics) and lets them move freely between rides. We further detail these differences when introducing our industry partner's background in Section 4.1.

The market is structured as follows. Spatially, it is defined over L regions in a single metropolitan area and covers all trips that originate and end within the city limits. Temporally, we discretize service hours into 20-minute intervals as decision epochs as the average ride-hailing trip duration is approximately 20 minutes. Effectively, we assume each ride lasts a multiple of 20 minutes. In between rides, on-demand drivers consider whether to continue to work and whether to switch platforms (multihome). Hence we assume that they do not abandon customers mid-ride.

In order for a gig platform to maintain service performance, it must manage the dynamics of when multihoming workers switch platforms and quit for the day. We focus on platform A and its ability to retain drivers throughout the day once they join its platform at the start of their daily work schedules. Once a driver joins, she interacts with the platform in the following way during every subsequent epoch until she leaves. At the start of the epoch, she is matched to a customer with a probability R^A_{lt} which depends on her current location l and time t. We assume that drivers are required to accept a match and they do not relocate themselves when idling. If matched, she commences the ride and, at the end of the decision epoch in which she completes the ride, decides whether to continue working on A, switch to the competing firm B, or stop working for the day. Otherwise, she waits without a ride until the end of the current decision epoch and faces the same decision.

Once she switches to platform B, her trajectory follows a similar model—she is matched to a customer with a state-dependent probability $R^B_{(\cdot)}$, completes assigned trips, and then makes stop-or-continue decisions based on the new location and time. Empirical patterns show that very few drivers work multiple, disjointed shifts on a single platform, say the platform for firm A, within a single day; therefore, switches are permanent within a day, and we do not model a driver option to return to firm A on the same day once she has switched to firm B. Thus, once having switched to the competing platform B, a driver has only two options at the end of decision epochs—to stay on B or leave the market. Figure 1 illustrates our model.

Figure 1 Model overview

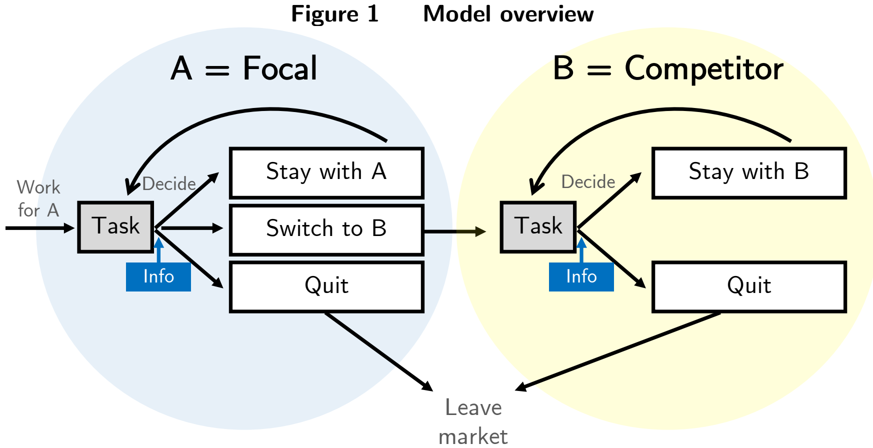

Note: A driver starts working on the focal firm A and, after completing each task, examines the time and location to decide whether to stay on firm A, switch to the competing firm B, or quit. After switching to firm B, she performs a task, examines the time and location, and decides whether to stay on firm B or quit.

Finally, drivers heterogeneously resolve their forward-looking decisions about whether to drive and whether to multihome based on their driver-specific costs to work, i.e., to stay on the road regardless of whether their time is spent waiting or driving a passenger. A driver i's average cost to work for a decision epoch, C_i, is Normally distributed with mean μ_C and variance σ^2_C over the truncated support $[\underline{\mathrm{C}}, \bar{\mathrm{C}}]$ for the overall population of drivers. In each decision epoch, an individual driver i's actual cost to drive is stochastic, $C_i + \epsilon_{ilt}^{(\cdot)}$, with the platform-specific contribution $\epsilon_{ilt}^{(\cdot)}$ of unobserved (to us) factors being i.i.d. extreme valued. Additionally, drivers' forward-looking decisions incorporate a homogenous time-discount factor $\beta < 1$.

The next subsections further detail drivers' forward-looking decisions on each platform.

3.1. Work Decisions on Focal Firm A

Except when she is mid-ride, a driver i active on platform A decides in each decision epoch whether to continue working on A, switch to B, or exit the service market entirely by quitting for the day.

Given that driver i finds herself at location l at the time t of her decision-making, we introduce the following relevant state variables and parameters. As prefaced, she accounts for the probability R_{lt}^A that she matches to a customer within the current epoch. Additionally, based on the location and time, she considers the characteristics of the trips she could receive. First, she anticipates the probability π_{lkt}^A that a customer departing l requests a trip to destination k. Then, for each potential destination k, she considers the average trip duration, in epochs, τ_{lk}^A, of trips made to k when departing l at time epoch t.

Accordingly, if she stays and is matched under our model, the driver takes her customer to the destination k with probability π_{lkt}^A, and such trip will take τ_{lk}^A time epochs. We ignore as negligible the driver's travel within l to the customer's exact pickup location. Whether she is matched or waiting, the driver i earns a flat hourly rate of w_{it} as long as she remains active on firm A's platform. Section 4.1 provides more information about drivers' hourly compensation rates on A. If active but unmatched, the driver remains at location l and revisits her work decision in the next decision epoch $t+1$. If the driver chooses to quit, she earns her outside option value of $V_i(Quit; l, t)$ which we normalize to 0.

Therefore, when excluding the stochastic component ϵ_{ilt}^{Stay} of her cost to drive, the driver's expected discounted value, $V_i(Stay; A, l, t)$, of continuing to work for A accounts for her wage rate per epoch, $\dfrac{w_{it}}{3}$, her cost to drive, C_i, and the best choice she could make at her next opportunity to decide whether to work, switch, or quit:

$$
\begin{aligned}
V_i(Stay; A, l, t) = &\frac{w_{it}}{3} - C_i + R_{lt}^A \cdot \left(\sum_{k=1}^{L} \pi_{lkt}^A \cdot \beta^{\tau_{lk}^A} \cdot \left((\tau_{lk}^A - 1)(w_{it}/3 - C_i) + V_i^*(A, k, t + \tau_{lk}^A) \right) \right) \\
&+ \beta \cdot (1 - R_{lt}^A) \cdot V_i^*(A, l, t+1), \qquad (1)
\end{aligned}
$$

where $V_i^*(A, l', t')$ represents her optimal future decision at location l' and time t':

$$
\begin{aligned}
V_i^*(A, l', t') &:= \max_{a \in \{Stay, Switch, Quit\}} V_i(a; A, l', t') + \epsilon_{il't'}^a \qquad (2) \\
&= \log \left(\sum_{a \in \{Stay, Switch, Quit\}} \exp \left(V_i(a; A, l', t') \right) \right). \qquad (3)
\end{aligned}
$$

The value of switching to the competing firm B will be derived in the next subsection, as equal to the value of continuing to work once on B: $V_i(Switch; A, l', t') = V_i(Stay; B, l', t')$. In each case, V_i excludes the stochastic component of the driver's current-period driving cost, which is based on factors unobserved to us as researchers. Once these are included, the utility-maximizing driver chooses to continue working on at least one platform as long as $\max_{a\in\{Stay,Switch\}} V_i(a; A, l, t) + \epsilon^a_{ilt} \geq V_i(Quit; A, l, t) + \epsilon^{Quit}_{ilt} = V_i(Quit; l, t) + \epsilon^{Quit}_{ilt} = \epsilon^{Quit}_{ilt}$.

3.2. Work Decisions on Competing Firm B

Similarly, except when mid-ride, each driver active for firm B decides at the end of each decision epoch whether to stay on for the next epoch or quit working for the day.

We use notation analogous to firm A's parameters to analyze the choice of a platform B driver at location l and time t. If she continues working for B, she is matched to a passenger with the probability R^B_{lt}. If matched, the requested destination is location k with probability π^B_{lkt}, and the ride occupies τ^B_{lk} epochs in work duration. For each potential destination k, the driver further considers the average fare f^B_{lkt} she would receive. If active but unmatched, the driver remains at location l and revisits her work decision in the next decision epoch $t+1$.

Therefore, the driver's expected discounted value, $V_i(Stay; B, l, t)$, of continuing on B is:

$$\begin{aligned} V_i(Stay; B, l, t) = &- C_i + R^B_{lt}\left(\sum_{k=1}^{L} \pi^B_{lkt}\beta^{\tau^B_{lk}}\left(f^B_{lkt} - (\tau^B_{lk} - 1)C_i + \max_{a\in\{Stay,Quit\}} V_i(a; B, l, t+\tau_{lk})\right)\right) \\ &+ (1 - R^B_{lt})\left(\beta\cdot \max_{a\in\{Stay,Quit\}} V_i(a; B, l, t+1)\right). \end{aligned} \tag{4}$$

The utility-maximizing driver chooses to continue working for B as long as $V_i(Stay; B, l, t) \geq V_i(Quit; B, l, t) = V_i(Quit; l, t) = 0$.

4. Data: Ride-hailing Trips in New York City

To estimate the model introduced in Section 3, we combine two sets of data on ride-hailing trips in New York City (NYC). First, we obtain proprietary data from our industry partner, a U.S. ride-hailing platform, that contain information regarding driver-level incentives and work decisions as well as trips within drivers' daily work sessions. Second, we leverage the trip record data collected by NYC's Taxi and Limousine Commission (TLC), which report the time and location of each trip fulfilled by a ride-hailing platform, including those provided by our industry partner. Together, these datasets allow us to examine the focal platform's ride-hailing drivers' decisions under real-time market conditions.

4.1. Proprietary Data from a Ride-hailing Industry Partner

Our industry partner is a ride-hailing platform, referred to as platform A, that operates in several U.S. cities. We focus only on its operations within NYC in this paper, in order to incorporate accurate data regarding drivers' outside options. The focal platform mainly offers shared rides and the majority of its customers are regular commuters, who have relatively consistent demand for their work commutes and require reliable service.

Another key feature of the partner's platform is that drivers are paid at a guaranteed hourly wage rate regardless of the number of customer trips they make as long as they are online and follow the navigation instructions supplied by the platform.[1] Hourly rates, referred to as *offers*, offered to drivers, are determined at the granularity of day of the week and time interval within each day (e.g., the morning rush hours of a given day, or the late night hours of the same day) and are communicated in advance to the driver through text messages. On this platform, a *shift* refers to a set time interval during which the hourly offer rate is fixed at a single value. The shifts studied in this paper are as follows: morning rush hours from 7 to 9 a.m. *(AM peak)*, from 9 a.m. to 5 p.m. *(midday)*, afternoon rush hours from 5 to 8 p.m. *(PM peak)*, evening non-rush hours from 8 to 9 p.m. *(PM off-peak)*, and from 9 p.m. to midnight *(late night)*. Note that drivers do not need to work an entire shift's duration in order to be compensated; instead, they are paid pro rata at the offer rate times the actual time they work (i.e., are online and following navigation instructions). Most rides are shared by multiple passengers, and the platform utilizes its drivers at higher rates than other ride-hailing platforms. Consequently, drivers rarely need to decide to relocate, justifying our modeling assumption that drivers stay in their current regions when not matched with passengers.

We leverage two sets of proprietary data provided by our industry partner for their NYC operations in July 2017. The first dataset contains, for every driver's active shift, information regarding her driving activities and financial incentives including offers. The data include each driver's anonymized ID and vehicle type, her experience with the platform, the duration she worked during the shift, and the shift's hourly offer rate. We merge these with another comprehensive dataset from our industry partner that reports a driver's first pick-up and last drop-off during every work session. A work session is defined as an uninterrupted span of time during which the driver remains actively logged into the platform. For each work session, we observe the driver ID and the exact timestamps and GPS coordinates of the driver's first and last trips. Effectively, we observe, on a daily basis, when and where each driver starts and stops working. In total, these data cover 45,208 work sessions performed by 2,875 drivers.

[1] When not matched with any customers, drivers are instructed to relocate to specific locations and their real-time location is tracked by GPS to confirm that they are still working for the platform.

4.2. Public Trip Record Data from TLC

We also leverage the large, public, trip record dataset collected by NYC's Taxi and Limousine Commission (TLC) that reports the pick-up and drop-off information for every ride-hailing trip taken in the city across all platforms.[2] For each trip, we observe the timestamps and neighborhood-level locations of the pick-up and drop-off, as well as the dispatching base information, which we use to identify the specific ride-hailing platform operating each reported trip. Despite knowing trip times, locations, and destinations by platform, we cannot identify and trace individual drivers across trips.

We use this data in two ways. First, we use trip-level information to approximate the path taken by each active driver on the focal platform during each of her work sessions. As we could observe only the first pick-up and the last drop-off of each of work session, the TLC dataset allows us to estimate origin-destination transition probabilities when a driver is matched to a customer at a given pickup time and location. Second, we use the data to reconstruct the market conditions and incentives offered to our partner's drivers by competing ride-hailing platforms. During the summer of 2017, four major ride-hailing platforms were operating in the city, including our industry partner. We choose the largest competing platform to serve as the representative competing firm, referred to as platform B, in our model and estimation. From the TLC data, we obtain the average trip durations and transition probabilities as well as the overall volume of trips by location and time on the competing platform B; we use these values in our model. We lastly augment these data with trip fares calculated using a duration-and-distance-based formula posted on platform B's website, allowing us to reconstruct the real-time market conditions and potential outside opportunities offered to drivers by platform B.

4.3. Preliminary Analysis of Gig Work

The proprietary data from our industry partner provides preliminary insight into gig economy workers' behaviors and their work settings. We analyze 45,208 work sessions performed by 2,875 drivers on the focal platform during the month of July 2017. As stated earlier, as long as they are logged into the system and stay active, drivers earn guaranteed hourly offer rates that are determined and communicated to them ahead of time (e.g., typically the evening prior). These rates are driver-, day-, and shift-specific. Figure 2 shows the distributions of hourly offer rates by operating shift for work sessions on weekdays and on weekends. On average, the rates are at their daily high during the midday shift (labeled as 2 in the figure) and at their daily low during the late night shift (labeled as 5 in the figure).

[2] See `https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page`

Figure 2 Hourly guaranteed wage rates offered to drivers, by shift, in the combined dataset

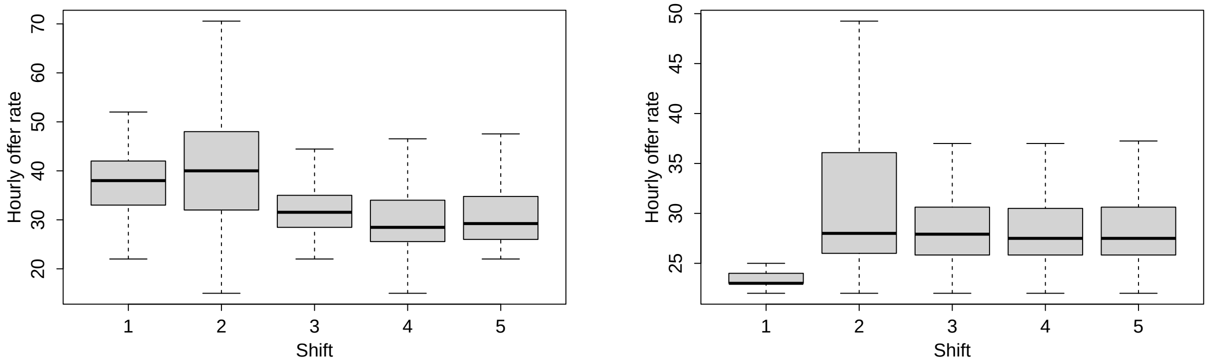

(a) Weekdays (b) Weekends

Note: The shifts, defined by the focal platform, from 1 to 5 are morning rush hours (7-9am), midday (9am-5pm), afternoon rush hours (5-8pm), afternoon non-rush hours (8-9pm), and late night (9pm-12am).

The duration of a work session has a mean of 4.50 hours, a median of 3.03 hours, and the middle quartiles cover between 1.10 hours and 6.66 hours. For each driver, the time between two consecutive work sessions has a mean of 36.83 hours and a median of 12.51 hours, and the middle quartiles cover between 4.43 hours and 27.58 hours. The average number of daily work sessions for an individual driver ranges from one session to five sessions, with a mean of 1.33 sessions and a median of 1.24 sessions. Therefore, it is safe to assume that most drivers only drive one long work session per day on the focal platform.

Figure 3 Daily work outcomes conditional on working

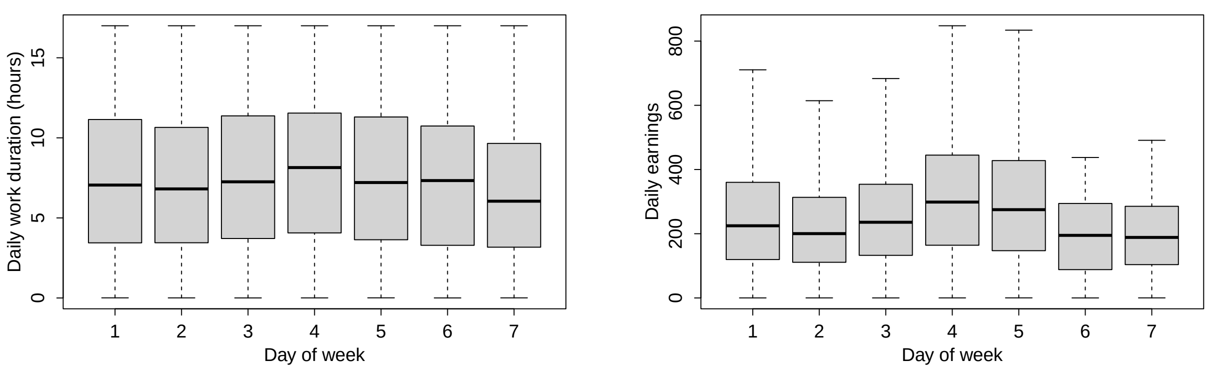

(a) Daily work duration (hours) (b) Daily earnings

Summing up across same-day work sessions, the daily work duration on the focal platform has a mean of 6.34 hours and a median of 5.55 hours, and the middle 50 percent range spans are from 2.28 hours to 10.32 hours. A comparison of our proprietary data and the aggregated reports published by TLC reveals that the daily work duration on the focal platform approximately matches the average daily duration of 65,178 ride-hailing drivers across all platforms in July 2017 (6 hours), suggesting that the behavior on the focal platform is not significantly different from the other platforms. Figure 3a depicts the distribution of the daily work duration for each of the weekdays. Similarly, Figure 3b illustrates the distribution of workers' earnings for each day of the week. Daily earnings on the focal platform have a mean of \$230.90 and a median of \$211.35, and the middle quartiles are from \$116.67 to \$322.77. Dividing these daily earnings by the average daily work duration, a driver would earn on average \$36.42 per hour, which is much higher than NYC's minimum wage of \$11.00 per hour in 2017. Potentially, the higher pay could encourage drivers to stay on the focal platform rather than switch to the competitor.

Under the gig economy's unique flexibility, workers can start working at any time and location. Here, we compare each work session's starting location as well as its ending location and time. We find that 91.86% of work sessions started and ended on the same day, suggesting that most drivers do not work overnight on this focal platform, thus justifying our model assumption. Figure 4 shows the distribution of drivers' decisions of when to start and end their work sessions. We observe a spike in drivers starting new sessions right before the morning rush hour begin (5-6 a.m.), at the beginning of the midday shift (9 a.m.), before the afternoon rush hour (3-4 p.m.), and right after midnight. Such patterns suggest that drivers strategically maximize the benefit of these shifts and their offers by starting early. We also find that drivers tend to end their sessions right before the end of the morning rush hours (8-9 a.m.) and after the afternoon rush hours (after 5 p.m.). These clustered ending times suggest that after rush hour the drivers find better value elsewhere.

Figure 4 Distributions of the hours of the first and last trips in drivers' work sessions

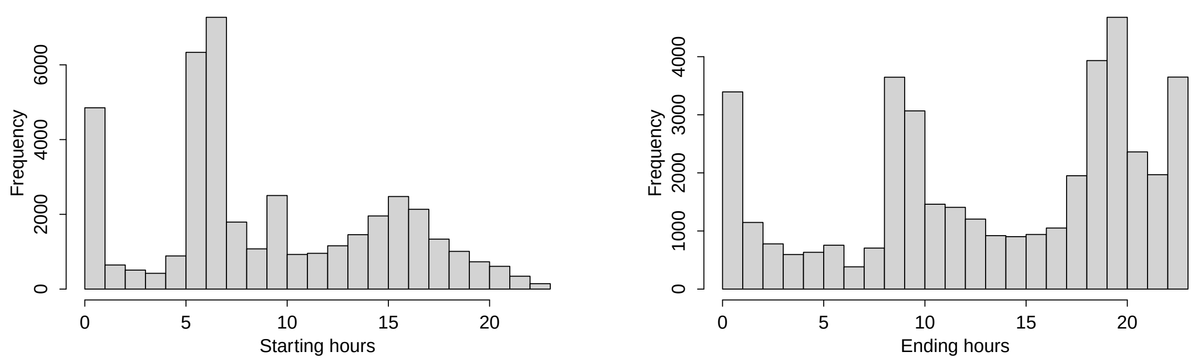

(a) Hour of the first trip in the session

(b) Hour of the last trip in the sessions

In terms of individual drivers' location decisions, 74.62% of sessions started and ended in the same borough in New York City; however, only 2.06% started and ended in the same neighborhood, and 9.66% started and ended in the same region defined by us. These findings suggest that drivers tend to stay in their focal borough, but within it, they explore multiple neighborhoods based on where their trips can lead.

These preliminary analyses provide early insights that drivers on the focal platform are sensitive to real-time monetary incentives as well as potential outside options happening at the same time across the city.

5. Estimation Strategy and Implementation

In this section, we estimate the model we developed in Section 3 using the proprietary data from our industry partner and the publicly available New York City (NYC) data as described in Section 4.

We group adjacent NYC neighborhoods into $L = 20$ regions: the Bronx, Brooklyn, Newark Liberty International Airport, Central Park, Chelsea, Downtown, Governors Island, Gramercy, Harlem, Lower East Side, Lower West Side, Midtown, Morningside Heights, Upper East Side, Upper West Side, Upper Manhattan, JFK International Airport, LaGuardia Airport, Queens, and Staten Island. Along the time dimension, we partition each operating day into seven blocks to match the shifts defined by our industry partner: 7-9am (AM peak), 9am-12pm (Late AM), 12-2pm (Midday), 2-5pm (Early PM), 5-8pm (PM peak), 8-9pm (PM off-peak), and 9pm-midnight (Late night).[3] Therefore, each shift spans multiple decision epochs. Within each decision epoch and region, we assume that available matches are made instantly between the active driver and her customer such that there is no waiting time incurred within the epoch when matched.

Our goal is to estimate three population parameters: β, μ_C, σ_C. Recall that $\beta \in [0,1]$ is the common time-discount factor, whereas the mean μ_C and variance σ_C^2 govern the truncated Normal distribution of driver-specific average cost to work, C_i for driver i, per time epoch.

5.1. Pre-Computation

Our estimation involves computing the values of different work options for each driver at each time and location for each day. To reduce the computational effort, we pre-compute several key values from our data prior to estimating for the key population parameters. First, to determine a reasonable range for our cost parameters, we consult a 2014 report by the American Automobile Association (AAA). The cost per mile is estimated to be 59.2 cents for an average sedan and 73.6 cents for an SUV (Hunter 2014). The average NYC citywide bus speeds in 2017 were 7.44 miles per hour (New York City Department of Transportation 2018).[4] The hourly cost is then \$4.40 for a sedan and \$5.48 for an SUV. Since we consider a 20-minute interval in our model, μ_C is then expected to be around \$1.47–\$1.83 per interval. To be conservative, we set $\underline{C} = 0$ and $\bar{C} = 5$ as the support for the drivers' cost distribution.

[3] Our industry partner groups the midday and early PM blocks together into a single midday shift. We break the shift down into two blocks to make all blocks more comparable in terms of duration.

[4] We use the average citywide bus speeds as a proxy for travel speeds in the city as it is the only statistic reported at the city level. The only other available statistic was the average taxi speeds in the Central Business District (Manhattan south of 60$^{\text{th}}$ Street) which were 7.1 miles per hour.

We also pre-compute the service parameters specific to each platform from our datasets. The probability of getting matched with a passenger at each location l and time t on platform j, R^j_{lt}, is computed by first summing up the number of trips operated by j starting from location l at time t (from TLC trip record data) and then dividing it by the number of unique drivers on the focal platform who were online during the same time in all locations (from our proprietary data). The transition probabilities π^j_{lkt} are computed by dividing the number of trips going from location l to location k at time t by the number of trips leaving location l at time t. The average trip duration τ^j_{lk} is computed by averaging the durations of all trips from l to k on platform j. We then assume that these values are fixed as the averages within each location, time block, and day of week.

5.1.1. Expected values of working for a competing firm B. We consider a parameter grid of a heterogeneous driver's average cost of working during a decision epoch, $C_i \in [0,5]$ with an increment of 0.05, and a homogeneous forward-looking discount factor per decision epoch, $\beta \in \{0.8, 0.825, 0.85, 0.875, 0.9, 0.925, 0.95, 0.975, 0.98\}$. For each combination of (C_i, β), we compute the value of working for B at location l and time block b either on a weekday or a weekend: $V^B_{l,b,weekday}(C_i,\beta)$ or $V^B_{l,t,weekend}(C_i,\beta)$. Recall that for an active driver on B, we assume that they can choose between two options: (i) continue working or (ii) leave the market. The value of leaving the market is normalized to zero for any time block and any location. We omit the subscript for day of week for simplicity. For each day type, we solve for the values backward from the last (7^{th}) time block in which all active drivers are expected to quit at the end of the block (i.e., midnight). The terminal condition is $V^B_{l,7,midnight}(\cdot) = 0$ for all $l = 1,...,L$. Then, we compute the values of working for B during the earlier time intervals within the same time block based on the following equation:

$$\begin{aligned} V^B_{l,7,t}(C_i,\beta) = R^B_{l,7} &\left(\sum_{k=1}^{L} \pi^B_{l,k,7} \beta^{\tau^B_{l,k,7}} \left(f^B_{l,k,7} - (\tau^B_{l,k,7} - 1)C_i + \max\left(V^B_{k,7,(t+\tau^B_{l,k,7})}(C_i,\beta), 0\right)\right)\right) \\ &+ (1 - R^B_{l,7})\left(\beta \max\left(V^B_{l,7,(t+1)}(C_i,\beta), 0\right)\right) - C_i, \end{aligned}$$

until we obtain convergence: $V^B_{l,7,t}(\cdot) = V^B_{l,7,t+1}(\cdot)$ for all l. The $L \times 1$ vector of converged values reflects the stationary expected value of working for platform B during the focal time block: $\mathbf{V^B_7} = (V^B_{1,7}, ..., V^B_{L,7})$. This vector is then used as a terminal condition for the earlier time block: $V^B_{l,6,T}(\cdot) = V^B_{l,7}(\cdot)$ for all $l = 1,...,L$. In the same fashion, for any time block $b < 7$, we compute intermediate values of working for B by

$$\begin{aligned} V^B_{l,b,t}(C_i,\beta) = R^B_{l,b} &\left(\sum_{k=1}^{L} \pi^B_{l,k,b} \beta^{\tau^B_{l,k,b}} \left(f^B_{l,k,b} - (\tau^B_{l,k,b} - 1)C_i + \max\left(V^B_{k,b,(t+\tau^B_{l,k,b})}(C_i,\beta), 0\right)\right)\right) \\ &+ (1 - R^B_{l,b})\left(\beta \max\left(V^B_{l,b,(t+1)}(C_i,\beta), 0\right)\right) - C_i. \end{aligned}$$

The converged values are then the stationary values for the time block to be used as a terminal condition for the next (earlier) time block: $V^B_{l,b,T}(\cdot) = V^B_{l,b+1}(\cdot)$ for all $l = 1,...,L; b = 1,..,6$.

5.1.2. Expected values of working for the focal firm A. For each driver i, we observe whether they worked for the focal platform A on a particular date d, $Work_{i,d}$, and if so, when and where they started (t_i^s, l_i^s) and ended (t_i^q, l_i^q) their work session for the day. Similar to the computation of the B-related values, we assume that the probability of being matched to a passenger to provide a ride R^A, the transition probability π^A, and the trip duration τ^A, as well as the values of switching over to platform B, V^B, are fixed within each location, time block, and day of week. However, we will estimate the *driver-* and *date-specific* value of working for platform A for every 20-minute interval starting from the time the driver started working until midnight for each date that the driver worked for platform A. For weekdays in July 2017, there are 11,109 driver-date pairs.

Due to the computational burden, we adopt a sparser grid to compute platform A-related values for each driver: 13 values of C (0, 0.5, ..., 5) $\times$ 6 values ofβ (0.825, 0.85, ..., 0.975). For each combination of (C, β), we compute the value of working for platform A, $V^A_{i,l,d,t}(C,\beta)$, for each driver i at location l on date d and time interval t, where d is taken from the set of dates when the driver i worked for platform A, and $t \in [t_i^s + 1, \text{midnight}]$. The terminal condition at midnight is that the driver quits: $V^A_{i,l,d,midnight}(\cdot) = 0$ for all $l = 1, ..., L$.

Similar to the computation of platform B-related values, we solve backward from midnight, for every 20 minutes, until we reach 20 minutes after the actual time the driver started t_i^s. The hourly offer for driver i at time t of day d is given by $w_{i,d,t}$. Define $b(t)$ as the time block that the time interval t is in and $wk(d)$ as an indicator whether the date d is a weekday or weekend. We omit the subscript for time block and day of week for simplicity.

$$
\begin{aligned}
V^A_{i,l,d,t}(C,\beta) &= \frac{w_{i,d,t}}{3} - C + \beta(1-R_l^A)V_{i,l,d,t+1} + R_l^A \sum_{k=1}^{L} \pi^A_{l,k}\beta^{\tau^A_{l,k}} \left((\tau^A_{l,k}-1)\left(\frac{w_{i,d,t}}{3} - C\right) + V_{i,k,d,t+\tau^A_{l,k}} \right) \\
&= \frac{w_{i,d,t}}{3} - C + \beta(1-R_l^A) \log\left(\sum_{j\in(w,s,q)} \exp\left(V^j_{i,l,d,t+1}\right) \right) \\
&\quad + R_l^A \sum_{k=1}^{L} \pi^A_{l,k}\beta^{\tau^A_{l,k}} \left((\tau^A_{l,k}-1)\left(\frac{w_{i,d,t}}{3} - C\right) + \log\left(\sum_{j\in(w,s,q)} \exp\left(V^j_{i,k,d,t+\tau^A_{l,k}}\right) \right) \right),
\end{aligned}
$$

where $V^w_{i,l,d,t}(\cdot) = V^A_{i,l,d,t}(\cdot), V^s_{i,l,d,t}(\cdot) = V^B_{l,b(t),wk(d)}(\cdot)$, and $V^q_{i,l,d,t}(\cdot) = 0$.

5.2. Adversarial Estimation for Population Parameters

Finally, we estimate our population parameters. Our methodology addresses a basic empirical challenge. Whereas, we observe a large number of heterogeneous drivers' start and end times and locations, their intermediate paths are unobserved. Although such data still allow us to identify location- and time-specific worker outflows that depend upon granular changes in the types of trips and fares that are available, a maximum likelihood approach would formally entail the computationally daunting task of simulating and integrating out the possible intermediate paths at an individual driver level. Therefore, we adopt a machine-learning-based estimation method called *adversarial estimation* recently proposed by Kaji et al. (2023) in order to estimate parameters that describe drivers' characteristics at an aggregated level.

Outcomes of interest. Since we do not directly observe drivers' intermediate locations during their shifts, we perform the estimation based on the aggregated outflows of workers at specific times and locations. The daily outcome of interest is the array X that captures, at each hour and location, the fraction of ride-hailing drivers leaving the focal platform A. Let $f^d_{l,h}$ be the daily fraction of drivers who worked for platform A on day d and left the platform at location l and hour h, such that $\sum_{l\in L}\sum_h f^d_{l,h}=1$ for any d. For example, a row in X contains the day of week (e.g., Wednesday), the hour (e.g., 9-10am), and, across 20 columns, the fraction of drivers quitting at each of the 20 locations in New York City.

Using our model, we can also simulate such outflows as outcomes. For a set of parameters $\theta=(\beta,\mu_C,\sigma_C)$, we first draw each driver i's cost of working, C_i, from the truncated Normal(μ_C,σ_C) distribution and then simulate her outcomes X^θ_i that indicates when and from which location she left the focal platform, aggregated over all days that the driver i worked on the focal platform. We then obtain the overall outflows X^θ by aggregating X^θ_i across all drivers.

Adversarial estimation. To identify the parameter set θ that generates the simulated outflows X^θ that match the observed data X as closely as possible, we employ the estimation method proposed by Kaji et al. (2023) that is inspired by a machine learning algorithm called a *generative adversarial network (GAN)* (Goodfellow et al. 2014). GANs are often employed as models that can be trained and used to generate artificial data (e.g., images) that look real (e.g., have characteristics that match the real data). In a similar manner, we use GANs to obtain estimates of our structural model that generate platform outflows that match the patterns produced by real human decision-makers.

The idea of GANs and adversarial estimation is the result of a two-player minimax game. The first player is a discriminator that evaluates a given row of data and determines whether the data are real or simulated. The second player is a generator that produces the simulated data from the model. In each iteration, the model's parameters are updated to maximally increase the error rate of the discriminator. As the discriminator, we use a neural network that acts as a classifier of real versus simulated data using accurate labels.

Ultimately, the adversarial estimation framework consists of two levels of estimation. The objective of the inner maximization problem is the negative cross-entropy loss function of the discriminator model. The inner maximization can be written as follows:

$$\max_{D\in\mathcal{D}} \frac{1}{N}\sum_{i=1}^{N}\log(D(X_i)) + \frac{1}{H}\sum_{i=1}^{H}\log(1-D(X_i^{\theta})),$$

where $D(X)$ is a binary classifier function that decides whether an observed outcome X is from the actual data ($D=1$) or from the simulation ($D=0$), N is the number of actual data observations, and H is the number of simulation trials. Then, the outer minimization problem is the search for the set of parameter values, θ, that minimizes the discriminator's level of classification accuracy. Ideally, these are parameter values under which the real and simulated data are indistinguishable.

$$\min_{\theta}\max_{D\in\mathcal{D}} \frac{1}{N}\sum_{i=1}^{N}\log(D(X_i)) + \frac{1}{H}\sum_{i=1}^{H}\log(1-D(X_i^{\theta})).$$

Kaji et al. (2023) provide theoretical guarantees and characterize the statistical properties of this GAN-based estimator.

Implementation. We first combine the simulated data X^{θ} with the actual data X and randomize the order of the rows. For each row, we set the label Y to be 1 if the row is from the real data or 0 if it is from the simulated data. We use a feed-forward neural network as our discriminator D as it is known to perform well in terms of classification accuracy under substantial data heterogeneity.

Our neural network is specified with a three-unit dense layer followed by a batch normalization and the `leakyReLu` activation function. We use a dropout rate of 40% between the first and the second block of dense. The dense, batch normalization, `leakyReLu`, and dropout pattern are repeated for three hidden layers before the last sigmoid output later. These are trained with a batch size of 32. To train the generator network, we use the gradient from the discriminator model to back-propagate through the cost function, and this allows us to calculate the gradient with respect to our parameters at the end of each iteration. Additionally, a soft labeling technique is implemented, to increase the stability of our GAN adaptation. To prevent the discriminator from overfitting, its weights are reinitialized when a gradient explosion occurs. We use the `tensorflow` and `keras` packages in R and the default Adam optimization algorithm for our implementation. After we train $\hat{D}$ using X and X^{θ}, we then fit it back to the same inputs to predict the labels $\hat{Y}$. We consider a cross-entropy loss function:

$$Loss = \sum_{i \text{ s.t. } Y_i=1}(-\log(\hat{Y}_i)) + \sum_{i \text{ s.t.} Y_i=0}(-\log(1-\hat{Y}_i))$$

The advantage of using the off-the-shelf neural network is that we can obtain the gradients using automatic differentiation. Given $\hat{D}$, we compute numerical gradient of the loss function with respect to the input data $\{X, X^{\theta}\}$. Using the chain rule, we obtain the gradient with respect to each of the parameters $\Delta(\theta)$. This step allows us to update $\theta^{(t)}$ to $\theta^{(t+1)}$ with one-step gradient descent:

$\theta^{(t+1)} = \theta^{(t)} - \xi\Delta(\theta^{(t)})$, where $\xi > 0$ is a learning rate. The derivation of our numerical gradient computation is provided in Appendix A.1. We repeat the procedure until we obtain convergence (e.g., $\Delta(\theta) \approx 0$).

6. Estimation Results

We present the structural estimates we obtain.

6.1. Population-level and Individual-level Parameters

Table 1 reports the population-level paramater estimates with standard errors computed from 10,000 bootstrap runs. As a key empirical finding, we estimate the discount factor to be 0.94985 per a 20-minute decision epoch, which indicates that although forward-looking, drivers are strongly myopic in prioritizing near-term earnings and costs. For example, drivers treat each US dollar of net earnings obtainable in 2 hours (4 hours) as though it is worth 73 cents (54 cents) at present. Under such myopic discounting, drivers are prone to switch platforms (multihome) in response to short-term differentials in platforms' immediately available fares, even when the current platform's fares will quickly recover so that remaining on the current platform yields better daily earnings overall.

Over the population of drivers, their per-epoch costs to drive are distributed under the truncated Normal with mean $\mu_C = \$0.55358$ and standard variation $\sigma_C = \$0.664725$. These estimated costs to drive include, at the lower end of the distribution, some values below the driving cost figures calculated by AAA,[5] which suggests that at least some drivers underestimate their own costs to drive even after factors such as automobile depreciation are correctly taken into account.

Table 1 Population-level parameter estimates

Discount factor β	Mean cost of working μ_C	Std. Dev. cost of working σ_C
0.94985 (0.00187)	\$0.55358 (\$0.01145)	\$0.664725 (\$0.01197)

Note: Standard errors in parentheses are obtained via bootstrap with 10,000 runs.

Given the population-level distribution of costs, we can determine a posterior estimate of the perceived costs borne by each particular driver based on her observed choices; each driver's posterior type is obtained by Bayesian updating from the population distribution of cost types. Figure 5 plots the distribution of the population's individual-level posterior cost estimates conditional on drivers' recorded decisions.

[5] Since 1950, the American Automobile Association (AAA) has published its annual *Your Driving Costs* study that breaks down the cost of owning and driving a car. The 2014 report suggests the hourly cost for a sedan (SUV) to be \$4.4 (\$5.48), translating into \$1.47 (\$1.83) for a 20-minute interval.

Figure 5 Distribution of individual-level cost estimates

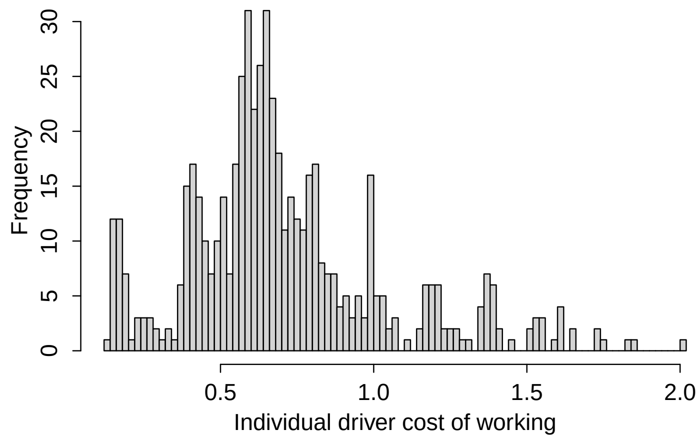

Model fit: Having recovered the operating costs of the individual drivers, we simulate their decisions to generate an informal comparison of model fit against the data. Under simulation with the estimated parameters, the average daily work duration on platform A is 7.3 hours, with a median of 6.3 hours. The simulated daily work hours span from a minimum of 0.3 to a maximum of 17.3 hours, with a standard deviation of 4.44 hours. These statistics are comparable to those of the actual work hours observed in the data.

Figure 6 shows the average simulated utilization rate of drivers on the focal platform throughout an average weekday. We observe that the utilization decreases as the morning proceeds, spikes up at noon, and remains high until late in the evening. Figure A.1 further illustrates variation across days.

Figure 6 Simulated utilization rates by the day of week and time of day

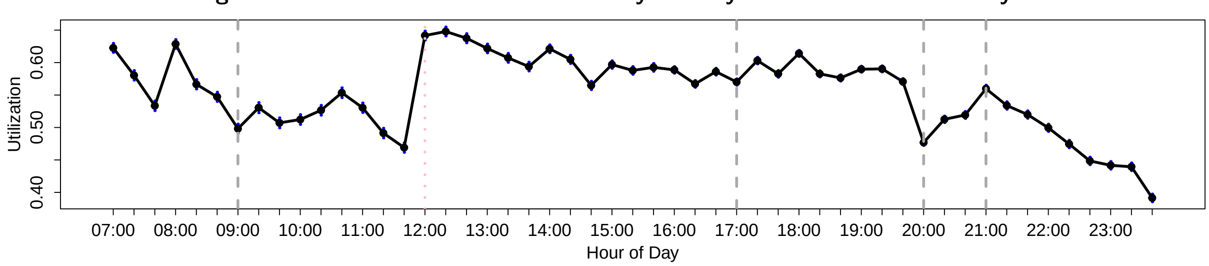

Note: Grey dashed lines represent shift boundaries on the focal platform; the pink dotted line represents noon.

6.2. Multihoming Behavior

In general, drivers cannot be tracked across platforms, and multihoming behavior must therefore be imputed. Based on our estimates which account for the competing platform, 66.59% of platform A's drivers multihome during the majority of the days they work. Of these, 42.23% of all drivers

Figure 7 **Rates of drivers switching to the competing platform by the day of week and time of day**

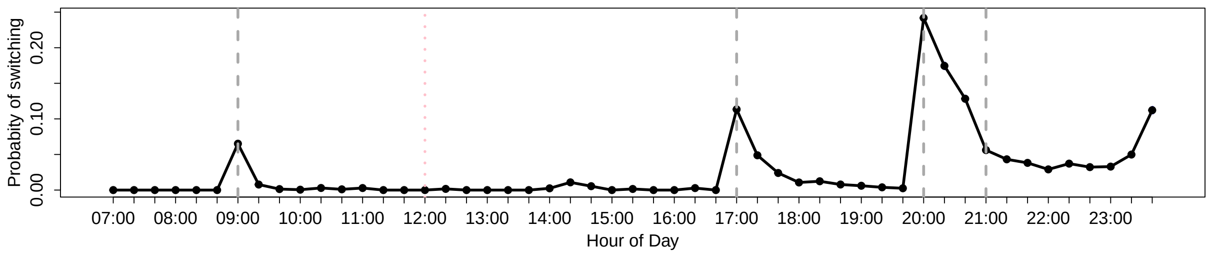

Note: Grey dashed lines represent shift boundaries on the focal platform; the pink dotted line represents noon.

multihome every day they work, whereas a separate 16.01% of all drivers never multihome. Figure 7 depicts the timing of multihoming switches on an average weekday. Multihoming transitions out of platform A typically occur at higher rates during the first hour of a new shift. At these times, drivers' hourly offer rates shift (as illustrated in Figure 2), causing both inflows and outflows of drivers, and for certain of these transitions (i.e., 9am, 5pm, and 8pm) they coincide with the arrival of significant daily shifts in traffic and ride demand patterns. The post-afternoon peak hours after 8pm feature a continuous outflow of drivers who multihome. In the evenings when offer rates decline, the platform's driver utilization rates decrease (see Figure 6). Additionally, the competing ride-hailing platforms offer ride services that more closely cater to nightlife events in NYC. These observations highlight how rapidly drivers flow in and out in response to shifts, due to their short-term focus.

7. Counterfactual Analyses

Our estimation results suggest that drivers are relatively myopic in their intra-day behavior. If they were less myopic, they would likely do a better job of balancing overall supply and demand. Then, there is potentially an efficiency argument to making platform choices "stickier" so as to keep the workers longer on platforms. Although at times platforms may want their active workers to switch to accommodate persistent shifts in supply and demand, getting them to flow out quickly is not hard. Using counterfactual analyses, we assess the design and the impact of policies employed by the gig platforms and the government on both the workers' decisions and the platform operations.

7.1. Managing Long-Term Capacity: Pay-per-work and Dynamic Guaranteed Pay

Gig economy platforms need to constantly manage their service capacity, which depends on an unpredictable and decentralized labor supply. As the magnitude of the gig economy has drastically grown, its on-demand services no longer cater only to short-term demand and supply. For example, the majority of the customers of the focal firm are regular commuters, who have consistent demand for their work commutes and require reliable service.

A carefully designed compensation scheme critically underlies the platform's ability to deliver consistent operational performance. The platform needs to retain a steady pool of workers and at the same time compete for temporary workers when unpredictable demand occurs. In practice, gig platforms use one of two primary schemes to compensate workers: pay-per-work and dynamic guaranteed wage rates. Both schemes have their advantages and shortcomings.

Ride-hailing platforms Uber and Lyft, as well as the delivery platforms Instacart and Grubhub are known to employ the pay-per-work scheme that compensates workers only when they perform jobs on the platform. This policy transparently reflects the market conditions of supply and demand. Because workers are compensated based on the work that they actually perform, they are attracted when jobs are available and not when work is scarce. In theory, this allows the firm to respond quickly and fluidly to real-time changes in demand and supply. Yet, such advantages depend on workers constantly checking the availability of jobs, and the platform may fail to attract workers in time to meet sudden surges in demand. At the same time, workers bear the opportunity costs of idly waiting for jobs when utilization is low, and they are prone to respond by multihoming or finding better alternatives.

The dynamic guaranteed wage scheme guarantees workers a fixed rate of hourly compensation during their shifts as long as they are actively available for work on the platform. Our focal company as well as the grocery delivery platforms GoPuff and Postmates are examples of platforms that offer a dynamic guaranteed hourly wage. In June 2023, DoorDash started offering options to Dashers to choose whether they earn a guaranteed hourly rate while delivering or receive compensation for each order (Browning 2023). The guaranteed pay option allows the platform greater certainty about its costs while attracting workers who may value income certainty and stable workflows. Research in organizational behavior such as DeVoe and Pfeffer (2007) has empirically shown that greater certainty in incentives (e.g., through guaranteed hourly wages) can increase worker productivity and induce longer work durations. Such a scheme is therefore appealing to firms that seek to deliver consistent service levels to meet regular demand levels (e.g., daily commuter flows).

When demand is unexpectedly low, however, the platform's underutilized labor supply may fail to adjust appropriately. Workers who do not serve customers nonetheless receive their guaranteed compensation rates by staying active. Unlike under the pay-per-work regimes, the firm bears the risk, through possibly close-to-fixed labor costs, of mismatched supply. Throughout the relatively short history of the gig economy, the larger platforms have predominantly utilized the pay-per-work scheme; nonetheless, as competition for multihoming workers grows between increasingly many gig platforms, it is an open empirical question whether guaranteed wages' advantages at retaining such workers outweighs their relative inability to alleviate real-time demand and supply mismatches.

We therefore ask: what are the operational effects of offering gig workers guaranteed pay rates, and under what conditions should a platform adopt this scheme?

Implementation. For this analysis, we evaluate how a platform's adoption of a pay-per-work scheme affects its operations, particularly: its capacity in terms of the average number of daily hours worked on the platform; and its labor supply in terms of the average number of daily trips a driver completes and the rates of multihoming across platforms. Using the estimated model of driver behavior, we generate counterfactual outcomes under pay-per-work compensation that we compare to those under the platform's practice of using guaranteed wages.

First, for each driver and hour on each day, we determine the pay-per-work rate r_{it} that is equivalent, on average, to the hourly guaranteed wage rate. That is, during a given hour t, driver i is paid a per-work rate r_{it} that is related to her hourly guaranteed rate w_{it} through the following equation:

$$w_{it} = r_{it} R_t \cdot \frac{60}{\tau_t},$$

where, during the hour t, R_t is the probability that drivers get a ride (i.e., the rate of getting matched per decision epoch) and τ_t is the average trip duration in minutes. We therefore personalize the per-job pay rates to the individual drivers by the time of day and day of week. However, we do not tailor the pay rates to the spatial locations of the drivers. Starting from this "equivalent" pay-per-work scheme, we consider a range of pay-per-work policies that offer pay-per-work rates at multiples of 0.5×, 0.75×, 1× (the pay-per-work scheme most comparable to the status quo), 1.25×, 1.5×, 2×, 2.5×, 3×, 3.5×, and 4×. Against these, we compare the focal platform's current use of guaranteed hourly wages as a base policy. Under each policy, we simulate every driver's decisions 100 times for each day in the data in order to study the platform's operational outcomes.

Results. We first investigate the the two types of wages' impact on the numbers of hours worked and trips completed on the focal platform. Figure 8 illustrates the impact as it varies by the day of week.

We find that in order to maintain equivalent levels of service capacity, pay-per-work wages incur higher labor costs than guaranteed pay rates. Compared to the guaranteed pay rates, the pay-per-work pay policy (1×) that incurs costs at equivalent rates diminishes the average number of hours worked and trips completed. The degree of the cost disadvantage varies depending on the time of day. Overall, the platform must pay a pay-per-work rate that is effectively 1.25 to 1.5 times greater than the guaranteed pay rate in order to keep the same level of supply. Guaranteed wage rates are thus clearly cost-effective.

The choice of pay policy further impacts multihoming behavior. Based on drivers' estimated discount factors, drivers' decisions regarding whether to continue working robustly discount their potential earnings received further into the future. Under the guaranteed pay rates, drivers receive a steady stream of earnings. Yet, despite compensating each driver at the same effective rate

Figure 8 Daily hours worked and trips completed, by the day of week, under the guaranteed pay and pay-per-work wage policies

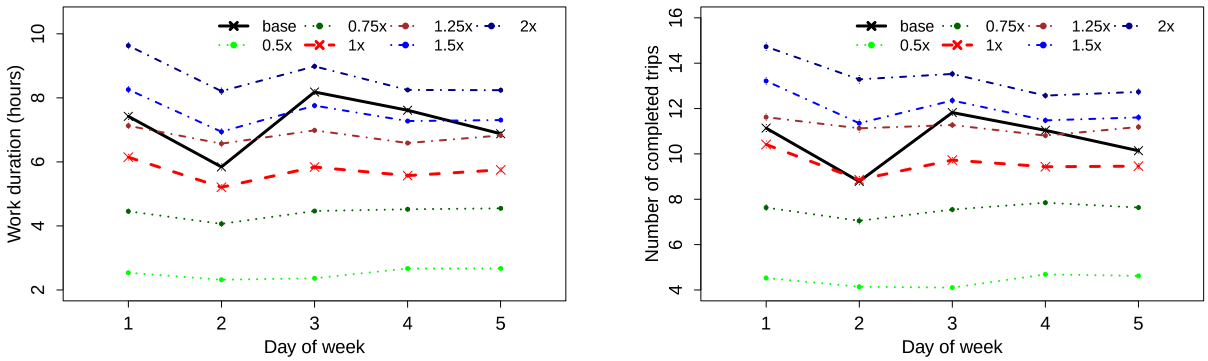

(a) Daily hours worked (b) Number of trips completed

Note: The outcomes are compared across the current guaranteed pay policy ("base") and pay-per-work policies of different scaling factors across weekdays. $\cdot\times$ represents the pay-per-work policy with an effective wage rate that is $\cdot$ times the guaranteed pay rate under the base policy. The horizon axis represents each week day (1: Monday, 5: Friday).

Figure 9 Distribution of daily hours worked and trips completed under the guaranteed pay and pay-per-work wage policies

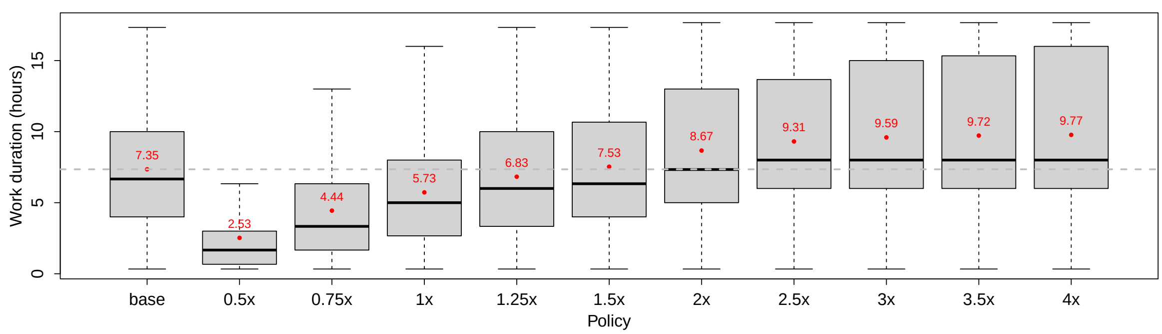

(a) Daily work duration (hours)

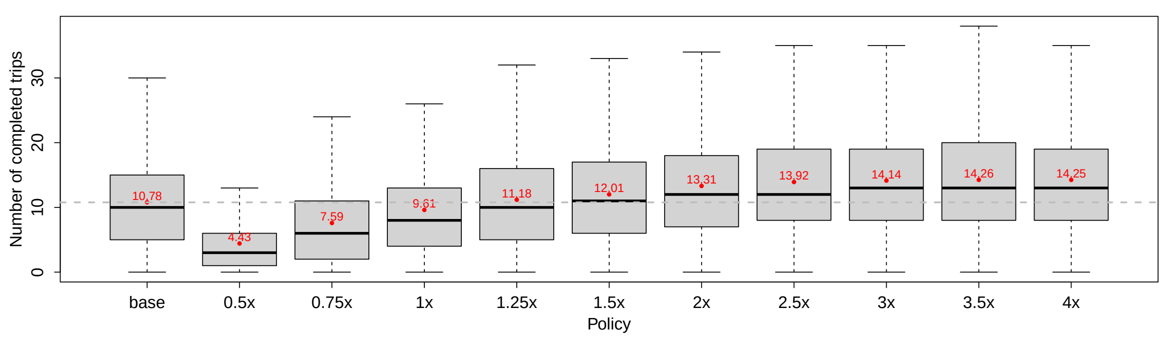

(b) Number of completed trips

Note: The outcomes are compared across the current guaranteed pay policy ("base") and pay-per-work policies of different scaling factors across weekdays. $\cdot\times$ represents the pay-per-work policy with an effective wage rate that is $\cdot$ times the guaranteed pay rate under the base policy. Red dots denote mean levels of a particular outcome. Grey dashed lines are the average fractions under the current guaranteed pay policy.

on average, the 1× pay-per-work pay policy delivers earnings only after each trip is completed and with uncertainty regarding when trips are matched—thus potentially delaying pay into the future. By reducing the expected *discounted* value of continuing to work for the focal platform, the pay-per-work policy raises the likelihood that drivers switch to the competing platform.

To shed further light on multihoming, we classify drivers into three types based on the multihoming behavior we observe over time. The first group, which we call *"never-switchers"*, does not switch to the competing platform after working for the focal platform. As a group, they are the most loyal to the focal platform. The second group, *"mostly-switchers"*, multihome frequently by switching from the focal platform to the competing platform on at least 50% of their workdays. The third and final group, *"always-switchers"*, switch to the competing platform on every workday started on the focal platform. Each plot in Figure 10 illustrates the fractions of drivers belonging to each of the three groups under the current guaranteed pay rates (labeled as 0) and under the pay-per-work policy at various multipliers of the equivalent rates of effective pay (labeled 0.5 through 4).

Figure 10 Driver multihoming behavior, by group, under guaranteed pay and pay-per-work wage policies

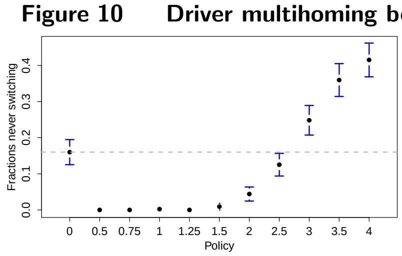

(a) Fraction of never-switchers

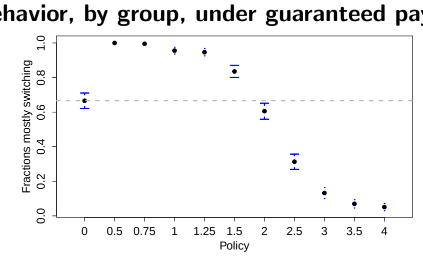

(b) Fraction of mostly-switchers

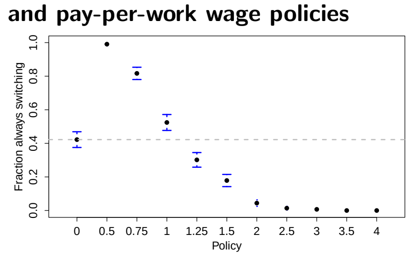

(c) Fraction of always-switchers

Note: The fractions are compared across the current guaranteed pay policy ("base") and pay-per-work policies of different scaling factors across weekdays. $\cdot\times$ represents the pay-per-work policy with an effective wage rate that is $\cdot$ times the guaranteed pay rate under the base policy. Grey dashed lines are the average fractions under the current guaranteed pay policy. "Never-switchers" refers to drivers who never switch to the competing platform after working on the focal platform. "Mostly-switchers" refers to drivers who switch to the competing platform after working on the focal platform at least 50% of the days that they work. "Always-switchers" refers to drivers who always switch to the competing platform after working on the focal platform.

In order to retain loyal drivers as effectively as under guaranteed pay rates, we find that the pay-per-work rates must be between 2.5× and 3× the current hourly offer (see Figure 10a). In other words, a guaranteed pay rates save the platform up to 75% of the per-unit labor cost to maintain the same pool of loyal drivers. Similarly, to keep the same fraction of drivers who switch most of the time, the pay-per-work rate has to be between 1.5× and 2× (see Figure 10b). Finally, to keep the same fraction of drivers who always switch, the wage multiple should be between 1× and 1.25× (see Figure 10c).

These findings suggest that to determine the pay-per-work rate the platform needs to decide the targeted fraction of loyal workers needed for their service capacity. This segment of the never-switchers is most affected by the change in the pay-per-work rate multiplier. Managing the pay impact on loyal workers is therefore critical for the platform to maintain its desired service capacity. Again, these effects vary by the day of the week as illustrated in Figure 11.

Figure 11 Driver multihoming behavior, by group and weekday, under guaranteed pay and pay-per-work wage policies

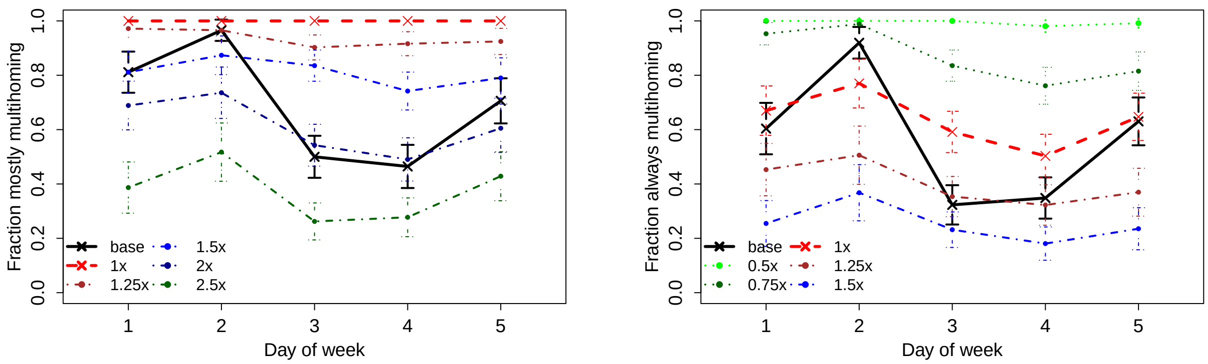

(a) Fraction of drivers that are mostly-switchers (b) Fraction of drivers that are always-switchers

Note: The fractions are compared across the current guaranteed pay policy (“base”) and pay-per-work policies of different scaling factors across weekdays. $\cdot\times$ represents the pay-per-work policy with an effective wage rate that is $\cdot$ times the guaranteed pay rate under the base policy. The horizon axis represents each week day (1: Monday, 5: Friday).

In general, driver outcomes follow the platform’s earnings. We calculate the pay-per-work rate r_{it} ($1\times$) that leads to the drivers earning a comparable amount of income to what they would have earned under the guaranteed wage policy. Figure 12 compares the simulated earnings across policies. Compensating workers per trip at a $1\times$ rate translates into fairly higher (30.16%) hourly earnings compared to the guaranteed wage policy as seen in Figure 12a. However, aggregating within each day, the $1\times$ pay-per-work rate generally maintains the level of daily earnings on the focal platform as seen in Figure 12b.

In summary, we observe that a schedule of dynamic guaranteed pay rates outperforms the pay-per-work scheme. Based on estimating drivers’ individual costs to work, we find that such a scheme better motivates loyal workers who prioritize income certainty in the short run. The retention of these workers is important to the platform, which matches their labor to handle the majority of demand. In part, this finding suggests that being able to reliably cover the regular demand patterns of customers requires aligning the wage policy to maintain a stable and consistent labor supply of available drivers. By accommodating drivers’ preferences for consistent workload and pay, the platform can deliver its customers consistent service levels, reasonable wait times, and lower prices based on reduced labor costs.

Figure 12 Comparison of simulated driver earnings (a: hourly, b: daily) from the current pricing scheme ("base") and pay-per-work policies of different scaling factors.

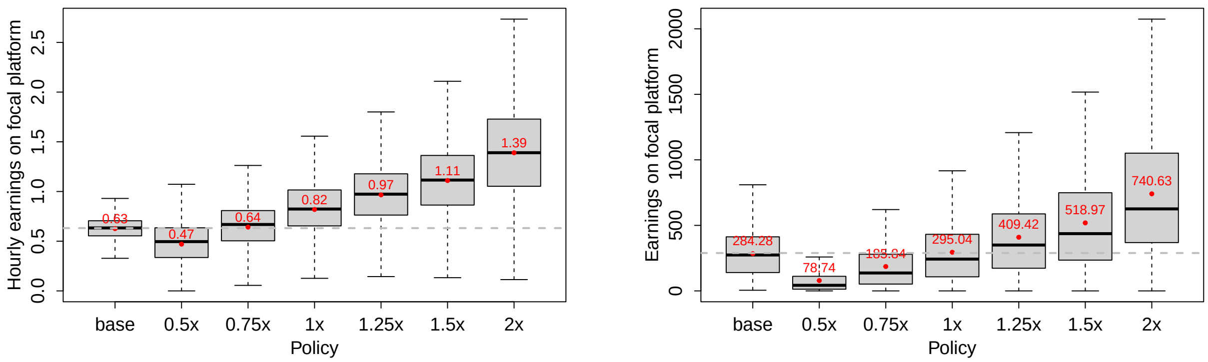

(a) Hourly earnings on focal platform

(b) Daily earnings on focal platform

Note: The outcomes are compared across the current guaranteed pay policy ("base") and pay-per-work policies of different scaling factors across weekdays. $\cdot\times$ represents the pay-per-work policy with an effective wage rate that is $\cdot$ times the guaranteed pay rate under the base policy. Grey dashed lines are the average outcomes under the current guaranteed pay policy.

7.2. Managing Short-Term Capacity: Bonuses and Delays

Beyond their everyday pay schemes, platforms must react to demand and supply shocks in real time. Perhaps most critically, at times of relatively high demand, platforms ought to retain the workers they have available to them instead of losing them to competitors or letting them quit. A key leverage point is that workers on a platform are less likely to switch over to a platform's competitors when kept active. Thus, a platform can hinder multihoming by keeping workers active; although high utilization naturally helps and arises in situations of high demand, the platform may use additional means.

In this section, we investigate two common levers that received attention in industry and academia: offering bonuses for continued work and imposing procedural frictions that delay workers from leaving. We examine the impact of these policies on platform operations and on drivers' multihoming decisions.

Real-time bonuses are meant to incentivize higher supply to match demand in different locations or time periods. One particular type of bonuses aims to keep workers working on a platform by awarding bonuses when they accept and complete a required number of consecutive jobs or stay active for a minimum duration. For example, Lyft awards a "Ride Streak" bonus to drivers who accept a ride within the set offer period and complete all required rides (e.g., meeting the quantity threshold) before the period is over. The dollar value of the Streak offers and the times they will be available are communicated to drivers ahead of time, and drivers could receive multiple bonuses during an offer period. DoorDash offers a similar promotion ("Guaranteed Earnings"); for example, delivery workers could earn at least \$500 in total earnings if they complete 50 deliveries within a week.

Such quantity-based threshold or consecutive work bonuses work well for two reasons. First, the marginal benefit of completing an additional job on the platform becomes higher than that of accepting and completing work elsewhere. In a ride-hailing context, threshold incentives have been found to be more effective at motivating drivers than linear incentives (Kabra et al. 2017), and the positive impact is more pronounced for full-time drivers compared to part-time drivers (Liu et al. 2023). Second, it presents workers with a goal; behavioral science research shows that having a clear goal is a strong motivator (e.g., Locke 1996). The main downside of this type of policy is that it requires platforms to commit to these relatively longer-term incentives with specific threshold requirements compared to real-time, on-off, surge pricing. Thus, if market conditions change, such bonus schemes may not elicit the desired response in a timely manner.

On the other hand, rather than paying more to workers for staying, platforms can impose costs or frictions on workers who consider leaving. For instance, Uber has experimented with psychological nudges in the form of short time delays to keep active drivers from leaving (Scheiber 2017). In addition, ride-hailing platforms often assign drivers new trips before they complete their previous trips in order to lock them in. Appealingly, this policy comes at no direct cost to the platform. However, it could conceivably backfire. Over time, it may induce workers to become highly anticipatory in terms of actively surveying and comparing their work options and outside options in order to avoid being locked in. The additional time between when the workers might have wanted to leave and the actual time they may result in a negative experience, hampering their intention to work on the platform in the future. This is similar to the finding in Tadepalli and Gupta (2020) on the negative impact of driver latency in a single-homing setting.

Implementation. We study both threshold bonuses and time-delay frictions in our counterfactual analysis. Specifically, the *time delay* policy impedes drivers from leaving the platform immediately at will. Under this policy, after a driver opts to leave the platform, she faces a short time delay of one decision epoch (20 minutes in our setting) until she is actually logged off. During the delay, the driver still receives her base pay for her time, and she may be assigned an additional trip. If the additional trip outlasts the time delay of 20 minutes, the driver is logged off immediately after she drops off all passengers.

Second, we study a *consecutive work bonus* policy that awards drivers hourly bonuses whenever they work for an entire regular hour of the day (e.g., for the entire 60 minutes starting at 9am and ending at 10am). We choose one hour as our threshold as it is the minimum threshold that is commonly used by gig economy firms in practice. Our robustness checks that use different durations as thresholds yield qualitatively consistent insights.

Figure 13 **Hours worked and multihoming rates under the time delay and consecutive work bonus policies**

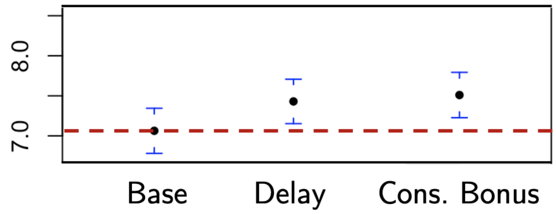

(a) Number of hours worked on the focal platform

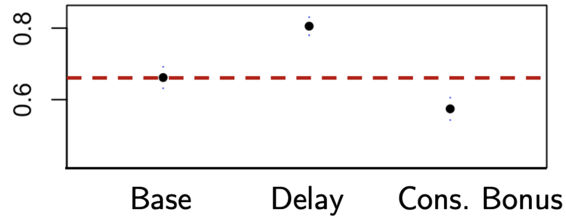

(b) Multihoming rates

Note: The outcomes are compared across the current guaranteed pay policy ("Base"), the time delay policy in which drivers have to wait at least one decision epoch before quitting ("Delay"), and the consecutive work bonus policy in which drivers are rewarded for staying active for an hour ("Cons. Bonus"). Each red dashed line represents the mean level of the corresponding outcome under the base policy.

Results. We compare the two policies of bonuses and time delays against a baseline scheme of guaranteed wages without either of them. We are interested in four key operational outcomes: the rates of utilization and multihoming among drivers active on the platform, respectively, the drivers' daily earnings, and their daily hours worked.

In certain respects, we find surprisingly little divergence in the key outcomes across the policies. First, compared to the baseline, we find that each of the policies appears to decrease utilization rates but not at a statistically significant level. Drivers' daily earnings increase compared to the baseline, but again not at a statistically significant level.

In contrast, the policies do impact drivers' hours worked and multihoming, yet they do so with considerable differences between them. First, both the time delay and consecutive work bonuses induce workers to work significantly longer hours. Figure 13a reports the number of hours worked per driver on the focal platform under each policy and the baseline without them. Drivers who are offered consecutive work bonuses work slightly longer than those facing time delays. Figure 13b illustrates drivers' intra-day multihoming rates of switching over to work for the competitor firm. Under the baseline policy, we find that multihoming rates are high at close to 70%. Under the time delay policy, multihoming rates rise significantly to over 80%. In contrast, under the consecutive work bonus policy, multihoming rates fall significantly. Thus, whereas both policies increase the number of hours worked, the use of time delays causes workers to opt into switching platforms anticipatorily to avoid getting locked in.

Our results produce interesting prescriptive insights for different demand scenarios. During peak hours, the platform needs to expand its service capacity and encourage workers to stay reliably engaged. By offering drivers temporary additional income for a specified duration on top of the certainty of the guaranteed wage rates, the consecutive work bonus helps address this need for additional labor supply. Several on-demand platforms have incorporated the consecutive work bonuses and similar forms of gamification into their worker retention efforts; examples include

Uber's *Consecutive Trips* and Lyft's *Ride Streaks* promotions. For the consecutive work bonus to be effective, however, platforms need a consistent stream of work to justify the income certainty that the policy provides.

Unlike bonuses, incorporating time delays does not deter multihoming. In fact, imposing time delays can be an operationally effective way to encourage strategic multihoming. By forcing workers to wait before being able to leave, it triggers them to plan ahead. Empirically, drivers appear prone to multihoming when facing low utilization, and they do not want to be stuck in low-demand locations or be forced to be idle for longer due to the time delay. Thus, perhaps counterintuitively, time delays can help gig economy platforms encourage timely and efficient multihoming outflows from the least utilized portions of their service networks.

7.3. Designing Regulations: *NYC's Driver Income Rules* Case Study

In the gig economy, workers lack traditional employment protections, and platforms' decisions regarding wages and compensation can impact worker welfare. Low pay is a major public concern: 14% (29%) of gig economy workers earned less than the federal (state) minimum hourly wage (Zipperer et al. 2022). For policymakers and regulators, a natural remedy is to enforce a minimum wage for gig work. However, theoretical analyses indicate that enforcing minimum wages without accompanying regulations and platform actions may lead to negative welfare effects, such as through flooding the gig market with drivers competing for too few jobs. Asadpour et al. (2023) offers a theoretical prediction that a minimum wage-type policy is only feasible for loose labor markets. If the minimum wage is too high, the optimal response for the platform is to cap its supply, i.e., only allowing a fixed number of workers to work for the platform. The limited access to gig work could then negatively affect worker welfare (Lao 2017).

In this section, we investigate whether policymakers should account for platforms' compensation schemes when evaluating new minimum-wage regulations. We take a minimum-wage policy for ride-hailing drivers instituted in New York City as our case study. In December 2018, New York City's TLC promulgated its *"Driver Income Rules"*, requiring that ride-hailing drivers be compensated at wage rates exceeding $17.22 per hour of work, and that the rates must fairly reflect the times and distances involved in the work.[6] In early 2019, the rules became enforced for most major ride-hailing platforms; our industry partner was exempted due to its different compensation scheme.

Because our data from 2016–2017 precede the enforcement of this regulation, we use counterfactual simulations to compare different implementations of the rules and their consequences for platform operations and worker welfare. We continue to assume a labor market with two competing, on-demand platforms that source labor from multihoming workers: our industry partner and

[6] The approved rules can be found at `https://nyc.gov/assets/tlc/downloads/pdf/driver_income_rules_12_04_2018.pdf`.

a representative competitor. The competitor compensates workers (pay-per-work) in line with the two platforms that held the largest market shares. We further assume that customer fares are not affected by the new rules. This assumption is supported by the findings of Parrott and Reich (2018) that the fare increase after the enforcement was not significant—e.g., only three to five percent higher than prior to enforcement. We compare the following four scenarios:

a. *Pre-2019 baseline of no enforced minimum wage for gig platforms.* The baseline scenario captures the platforms' decisions and outcomes based on the market conditions prior to the the minimum wage rules being enforced in 2019. The scenario is evaluated using the estimated driver characteristics from our earlier analyses as well as the market conditions obtained from the public data.

b. *Enforced on pay-per-work gigs only.* The second scenario reflects the minimum wage rules as they were implemented in early 2019. These rules were only enforced on the pay-per-work competitors to our industry partner's platform. In our counterfactual analysis, we therefore simulate the policy's effects on the drivers' outside option of multihoming (i.e., their value of working for the competing platform), before simulating and analyzing the follow-on effects on our industry partner's platform operations and driver decisions.

c. *Enforced on guaranteed wage gigs only.* The third scenario implements a hypothetical reversal of the actual enforcement regime by applying the minimum wage rules only to our industry partner but not to its pay-per-work competition.

d. *Universal enforcement.* The fourth and final scenario hypothetically applies the minimum wage rules universally to all platforms.

Results. We compare each of the latter three scenarios against the baseline without enforced minimum wages. Under the actually implemented second scenario, we find that workers earn 3.5% less and spend 2.1% more time idling. Our simulated findings are consistent with the regulation's actual outcomes in practice: drivers generally failed to earn more and spent more time waiting for work assignments due to an influx of drivers attracted by the new minimum wage. In response, ride-hailing platforms ultimately capped the number of drivers at certain times of day to avoid leaving them idle.

Under the second, hypothetical scenario, the rules were only enforced on guaranteed wage gigs. As a result, workers earned 3% more but also spent 3.3% more time idling compared to the pre-2019 baseline. In contrast, the universal enforcement scenario results in workers earning 1.2% less and spending 7% more time idling than pre-2019. These results indicate that enforcing minimum wages on pay-per-work platforms in particular consistently harms workers' earnings and utilization rates. Whereas such results alone do not negate the public interest in minimum wage standards, they suggest that additional regulations or steps are required (e.g., capping the number of drivers at certain hours) in order to raise worker welfare and avoid inefficient congestion and competition among idling workers. The guaranteed gigs respond to minimum wage standards more in line with expected outcomes at traditional employers. Lastly, it is interesting to compare the effect of universal enforcement against that of enforcement on pay-per-work gigs only—i.e., the former additionally brings guaranteed gigs up to the minimum wage. With the influx of labor supply, the overall earnings effects are still negative but less so than before, since unlike pay-per-work gigs, guaranteed wages are not diluted by low utilization. However, drivers still suffer, relative to pre-2019, from a diminished outside option at the competitor's platform, and utilization naturally falls even further under the broadened minimum wage that attracts drivers.

For policymakers, our findings highlight the importance of considering platforms' compensation schemes when designing minimum wage rules and other gig labor regulations. Similar analyses can help reveal which and when platforms will react more similarly to traditional employment markets and when additional regulations or steps should be deployed in tandem in order to achieve policy goals such as improving worker welfare.

8. Concluding Remarks

By exercising their flexibility to choose gigs, gig economy workers often work for multiple platforms. A growing number of platforms competes ever more fiercely to win over a limited pool of workers. Therefore, understanding how workers respond to multihoming incentives is becoming critical, but empirical studies have been limited by data that cannot directly capture workers' other options.

In this paper, we combine proprietary data from a ride-hailing industry partner and publicly available data documenting trips on all platforms in order to estimate a structural model of gig workers' sequential dynamic decisions that does account for the availability of work opportunities on other platforms. Methodologically, we contribute the modeling and structural estimation of gig workers' dynamic labor decisions using the aforementioned combined data and econometric techniques that leverage machine learning. Based on the estimated population parameters, our counterfactual analyses provide novel prescriptive insights for on-demand platforms and policy-makers.

Specifically, our analysis shows that workers are short-sighted and value consistent rewards over variable rewards. As a result, we find that consistent pay dominates variable pay-per-work at retaining multihoming workers in order to maintain a stable platform workforce. Platforms can further control their labor supply by rewarding uninterrupted work or delaying quits, where the latter can, counterintuitively, encourage multihoming. More generally, our research yields insights into multihoming behavior that can guide the design of optimal incentive schemes and gig economy regulations.

Our study offers a first empirical step towards understanding prevalent multihoming behavior. As more data become available, researchers could improve the approximation and estimation of a wider range of work options available to workers. In this study, we use the varying levels of imbalance between pick-ups and drop-offs in each region across the city as a proxy for time-specific and location-specific outside promotions and incentives. Although we do not directly observe the actual incentives and workload on other platforms, we argue that these approximated options could serve as a lower bound on the values of potential outside options. In addition, our counterfactual analyses focus on simple implementations of select interventions. Future researchers could delve more deeply into studying the variations of these policies, as well as other policies not explored in this study. Our estimation framework could also be extended to other settings involving human agents' sequential decisions, where only aggregated outcomes are observed, and traditional maximum likelihood estimation methods are infeasible. Finally, our study highlights that while gig economy workers are drawn to flexibility, they largely prefer certainty when it comes to earnings. Our industry partner is one of several firms that offer guaranteed pay to their workers. Therefore, we believe that our findings are not specific to this particular company's operations, but instead reflect broader patterns in the gig economy. Future research could investigate the interplay between competing platforms that offer different pay schemes and what the equilibrium mix of payment schemes across competing platforms should be.

By shedding light on the dynamic decision-making processes of gig workers and how incentives affect their decisions, we provide valuable insights into the challenges faced by platforms and policymakers. We believe that our findings have the potential to inform more effective strategies for managing multihoming workers in the gig economy, thereby improving their well-being.

References

Allon G, Chen D, Moon K (2023a) Measuring strategic behavior by gig economy workers: Multihoming and repositioning. *Available at SSRN 4411974* .

Allon G, Cohen MC, Sinchaisri WP (2023b) The impact of behavioral and economic drivers on gig economy workers. *Manufacturing & Service Operations Management* 25(4):1376–1393.

Anderson M, McClain C, Faverio M, Gelles-Watnick R (2021) The state of gig work in 2021 .

Asadpour A, Lobel I, van Ryzin G (2023) Minimum earnings regulation and the stability of marketplaces. *Manufacturing & Service Operations Management* 25(1):254–265.

Athey S, Calvano E, Gans JS (2018) The impact of consumer multi-homing on advertising markets and media competition. *Management cience* 64(4):1574–1590.

Athey S, Castillo JC, Chandar B (2024) Service quality on online platforms: Empirical evidence about driving quality at uber. Technical report, National Bureau of Economic Research.

Benjaafar S, Ding JY, Kong G, Taylor T (2022) Labor welfare in on-demand service platforms. *Manufacturing & Service Operations Management* 24(1):110–124.

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.

Benjaafar S, Xiao S, Yang X (2020) Do workers and customers benefit from competition between on-demand service platforms? *Available at SSRN 3645882* .

Bernstein F, DeCroix GA, Keskin NB (2021) Competition between two-sided platforms under demand and supply congestion effects. *Manufacturing & Service Operations Management* 23(5):1043–1061.

Besbes O, Castro F, Lobel I (2021) Surge pricing and its spatial supply response. *Management Science* 67(3):1350–1367.

Besbes O, Goyal V, Iyengar G, Singal R (2022) Effective wages under workforce scheduling with heterogeneous time preferences. *Available at SSRN* .

Bolton GE, Katok E, Ockenfels A (2004) How effective are electronic reputation mechanisms? an experimental investigation. *Management science* 50(11):1587–1602.

Browning K (2023) Doordash, shifting business model, will offer drivers hourly pay. URL `https://www.nytimes.com/2023/06/28/business/doordash-hourly-wage-option.html`.

Bryan KA, Gans JS (2019) A theory of multihoming in rideshare competition. *Journal of Economics & Management Strategy* 28(1):89–96.

Butschek S, Amor RG, Kampkötter P, Sliwka D (2022) Motivating gig workers–evidence from a field experiment. *Labour Economics* 75:102105.

Cameron L, Meuris J (2022) The peril of paycheck dispersion: When fluctuations in compensation jeopardize worker retention. *Academy of Management Proceedings*, volume 2022, 15162 (Academy of Management Briarcliff Manor, NY 10510).

Cennamo C, Ozalp H, Kretschmer T (2018) Platform architecture and quality trade-offs of multihoming complements. *Information Systems Research* 29(2):461–478.

Chen L, Cui Y, Liu J, Liu X (2022) Bonus competition in the gig economy. *Available at SSRN 3392700* .

Chen MK, Rossi PE, Chevalier JA, Oehlsen E (2019) The value of flexible work: Evidence from uber drivers. *Journal of Political Economy* 127(6):2735–2794.

Chitla S, Cohen M, Jagabathula S, Mitrofanov D (2023) Customers' multihoming behavior in ride-hailing: Empirical evidence using a structural model. *Working Paper* .

Choudhary V, Shunko M, Netessine S, Koo S (2022) Nudging drivers to safety: Evidence from a field experiment. *Management Science* 68(6):4196–4214.

Critero (2017) URL `https://www.criteo.com/wp-content/uploads/2017/10/TheShopperStory_US_Final.pdf`.

Davies J, Khodjamirian S, Giallombardo F, Aletti P (2022) Survey evidence on multi-homing in online retail business. URL `https://www.compasslexecon.com/the-analysis/survey-evidence-on-multi-homing-in-online-retail-businesses/11-23-2022/`.

Davis AM, Mankad S, Corbett CJ, Katok E (2024) Om forum—the best of both worlds: Machine learning and behavioral science in operations management. *Manufacturing & Service Operations Management* 26(5):1605–1621.

DeVoe SE, Pfeffer J (2007) When time is money: The effect of hourly payment on the evaluation of time. *Organizational Behavior and Human Decision Processes* 104(1):1–13.

Dong J, Ibrahim R (2020) Managing supply in the on-demand economy: Flexible workers, full-time employees, or both? *Operations Research* .

Fitzsimmons EG (2018) Uber hit with cap as new york city takes lead in crackdown. URL `https://www.nytimes.com/2018/08/08/nyregion/uber-vote-city-council-cap.html`.

Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. *Advances in neural information processing systems* 27.

Gourieroux C, Monfort A, Renault E (1993) Indirect inference. *J. Appl. Econometrics* 8(S1):S85–S118.

Hunter H (2014) Owning and operating your vehicle just got a little cheaper according to aaa's 2014 'your driving costs' study. *AAA NewsRoom* .

Ibrahim R (2018) Managing queueing systems where capacity is random and customers are impatient. *Production and Operations Management* 27(2):234–250.

Ikeda K, Bernstein MS (2016) Pay it backward: Per-task payments on crowdsourcing platforms reduce productivity. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 4111–4121.

Kabra A, Elena B, Karan G (2017) The efficacy of incentives in scaling marketplaces, working paper.

Kaji T, Manresa E, Pouliot G (2023) An adversarial approach to structural estimation. *Econometrica* 91(6):2041–2063.

Landsman V, Stremersch S (2011) Multihoming in two-sided markets: An empirical inquiry in the video game console industry. *Journal of Marketing* 75(6):39–54.

Lao M (2017) Workers in the gig economy: The case for extending the antitrust labor exemption. *UCDL Rev.* 51:1543.

Li H, Zhu F (2021) Information transparency, multihoming, and platform competition: A natural experiment in the daily deals market. *Management Science* 67(7):4384–4407.

Liu T, Xu Z, Vignon D, Yin Y, Li Q, Qin Z (2023) Effects of threshold-based incentives on drivers' labor supply behavior. *Transportation Research Part C: Emerging Technologies* 152:104140.

Lobel I, Martin S, Song H (2024) Frontiers in operations: Employees vs. contractors: An operational perspective. *Manufacturing & Service Operations Management* 26(4):1306–1322.

Locke EA (1996) Motivation through conscious goal setting. *Applied and preventive psychology* 5(2):117–124.

Loginova O, Wang XH, Liu Q (2022) The impact of multi-homing in a ride-sharing market. *The Annals of Regional Science* 69(1):239–254.

Mai Y, Hu B, Pekeč S (2023) Courteous or crude? managing user conduct to improve on-demand service platform performance. *Management Science* 69(2):996–1016.

Miao W, Deng Y, Wang W, Liu Y, Tang CS (2022) The effects of surge pricing on driver behavior in the ride-sharing market: Evidence from a quasi-experiment. *Journal of Operations Management* .

Moon K (2023) Strategic path selection in service networks: Leveraging machine learning to estimate combinatorially complex preferences for consumption and waiting. *Available at SSRN 3819117* .

New York City Department of Transportation (2018) New york city mobility report .

Park KF, Seamans R, Zhu F (2021) Homing and platform responses to entry: Historical evidence from the us newspaper industry. *Strategic Management Journal* 42(4):684–709.

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* .

Rosaia N (2020) Competing platforms and transport equilibrium: Evidence from new york city. Technical report, mimeo, Harvard University.

Scheiber N (2017) How uber uses psychological tricks to push its drivers' buttons. *Ethics of Data and Analytics*, 362–371 (Auerbach Publications).

Shokoohyar S, Katok E (2022) Incentivizing suppliers using scorecard: A behavioral study. *Journal of the Operational Research Society* 73(5):1105–1126.

Tadepalli K, Gupta A (2020) Multihoming in ridesharing markets: Welfare and investment. *Working Paper* .

Wilson R (2019) Seattle becomes first city to cap uber, lyft vehicles. URL `https://www.washingtonpost.com/blogs/govbeat/wp/2014/03/18/seattle-becomes-first-city-to-cap-uber-lyft-vehicles/`.

Zipperer B, McNicholas C, Poydock M, Schneider D, Harknett K (2022) National survey of gig workers paints a picture of poor working conditions, low pay .

Appendix A: Model Derivations

A.1. Numerical Gradients Computation

We use the following loss function:

$$Loss = \sum_{i \text{ s.t. } Y_i=1} (-\log(\hat{Y}_i)) + \sum_{i \text{ s.t.} Y_i=0} (-\log(1-\hat{Y}_i)).$$

Our input X has a $(D \cdot H) \times L$ dimension, where D is the number of days in the data, H is 17 (e.g., 17 hours from 7am to 11pm), and L is 20 (e.g., 20 regions in NYC). Keras automatic differentiation provides us with $\partial Loss/\partial X$,

$$\frac{\partial Loss}{\partial \theta} = \frac{\partial Loss}{\partial X} \cdot \frac{\partial X}{\partial \theta}.$$

We need

$$\frac{\partial X}{\partial \theta} = \begin{bmatrix} \frac{\partial X}{\partial \beta} & \frac{\partial X}{\partial \mu} & \frac{\partial X}{\partial \sigma} \end{bmatrix}.$$

Formulation of X: Each cell is a fraction of drivers quitting at location l at time h on day d.

$$X_{l,h,d} = \frac{1}{|Drove_d|} \sum_{i \in Drove_d} q^i_{l,d,h},$$

where $Drove_d$ is the set of drivers who worked on day d and $q^i_{l,d,h}$ is an average quitting decision for driver i at location l at time h on day d.

$$\frac{\partial X_{l,h,d}}{\partial \theta} = \frac{1}{|Drove_d|} \sum_{i \in Drove_d} \frac{\partial q^i_{l,d,h}}{\partial \theta}$$

Aggregation across drivers: For each driver i, their cost c_i is drawn from $TruncatedNormal(\mu, \sigma^2)$. We also set $\underline{c} = 0$ and $\bar{c} = 5$ as the range of possible costs. Technically, we draw the error term for each driver: ϵ_i from $Uniform[0,1]$. Then,

$$c_i = \Phi^{-1}(\Phi(\underline{c}) + \epsilon_i(\Phi(\bar{c}) - \Phi(\underline{c})))\sigma + \mu.$$

So we have

$$\frac{\partial c_i}{\partial \beta} = 0, \quad \frac{\partial c_i}{\partial \mu} = 1, \quad \frac{\partial c_i}{\partial \sigma} = \Phi^{-1}(\Phi(\underline{c}) + \epsilon_i(\Phi(\bar{c}) - \Phi(\underline{c}))).$$

In the pre-computation, we have already computed $q^i_{l,d,h}(\beta, c)$ for a specific β and c. We do bi-linear interpolation using the four nearest points: (β_1, c_1), (β_1, c_2), (β_2, c_1), (β_2, c_2), such that $\beta_1 < \beta < \beta_2$ and $c_1 < c < c_2$.

$$\begin{aligned} q^i_{l,d,h}(\beta, c_i) = {} & \frac{[\beta_2 - \beta \quad \beta - \beta_1]}{(\beta_2 - \beta_1)(c_2 - c_1)} \left[q^i_{l,d,h}(\beta_1, c_1) \quad q^i_{l,d,h}(\beta_1, c_2); \quad q^i_{l,d,h}(\beta_2, c_1) \quad q^i_{l,d,h}(\beta_2, c_2) \right] [c_2 - c_i; \quad c_i - c_1] \\ = {} & \frac{1}{(\beta_2 - \beta_1)(c_2 - c_1)} \cdot [q^i_{l,d,h}(\beta_1, c_1)(\beta_2 - \beta)(c_2 - c_i) + q^i_{l,d,h}(\beta_2, c_1)(\beta - \beta_1)(c_2 - c_i) \\ & + q^i_{l,d,h}(\beta_1, c_2)(\beta_2 - \beta)(c_i - c_1) + q^i_{l,d,h}(\beta_2, c_2)(\beta - \beta_1)(c_i - c_1)] \end{aligned}$$

$$\begin{aligned} \frac{\partial q^i_{l,d,h}(\beta, c_i)}{\partial \beta} = {} & \frac{1}{(\beta_2 - \beta_1)(c_2 - c_1)} \cdot [-q^i_{l,d,h}(\beta_1, c_1)(c_2 - c_i) + q^i_{l,d,h}(\beta_2, c_1)(c_2 - c_i) \\ & - q^i_{l,d,h}(\beta_1, c_2)(c_i - c_1) + q^i_{l,d,h}(\beta_2, c_2)(c_i - c_1)] \end{aligned}$$

$$\begin{aligned}\frac{\partial q^i_{l,d,h}(\beta,c_i)}{\partial\mu} &= \frac{1}{(\beta_2-\beta_1)(c_2-c_1)}\cdot[-q^i_{l,d,h}(\beta_1,c_1)(\beta_2-\beta)\frac{\partial c_i}{\mu}-q^i_{l,d,h}(\beta_2,c_1)(\beta-\beta_1)\frac{\partial c_i}{\mu}\\ &+q^i_{l,d,h}(\beta_1,c_2)(\beta_2-\beta)\frac{\partial c_i}{\mu}+q^i_{l,d,h}(\beta_2,c_2)(\beta-\beta_1)\frac{\partial c_i}{\mu}]\\ &= \frac{1}{(\beta_2-\beta_1)(c_2-c_1)}\cdot[-q^i_{l,d,h}(\beta_1,c_1)(\beta_2-\beta)-q^i_{l,d,h}(\beta_2,c_1)(\beta-\beta_1)\\ &+q^i_{l,d,h}(\beta_1,c_2)(\beta_2-\beta)+q^i_{l,d,h}(\beta_2,c_2)(\beta-\beta_1)]\end{aligned}$$

$$\begin{aligned}\frac{\partial q^i_{l,d,h}(\beta,c_i)}{\partial\sigma} &= \frac{1}{(\beta_2-\beta_1)(c_2-c_1)}\cdot[-q^i_{l,d,h}(\beta_1,c_1)(\beta_2-\beta)\frac{\partial c_i}{\sigma}-q^i_{l,d,h}(\beta_2,c_1)(\beta-\beta_1)\frac{\partial c_i}{\sigma}\\ &+q^i_{l,d,h}(\beta_1,c_2)(\beta_2-\beta)\frac{\partial c_i}{\sigma}+q^i_{l,d,h}(\beta_2,c_2)(\beta-\beta_1)\frac{\partial c_i}{\sigma}]\end{aligned}$$

Appendix B: Additional Results

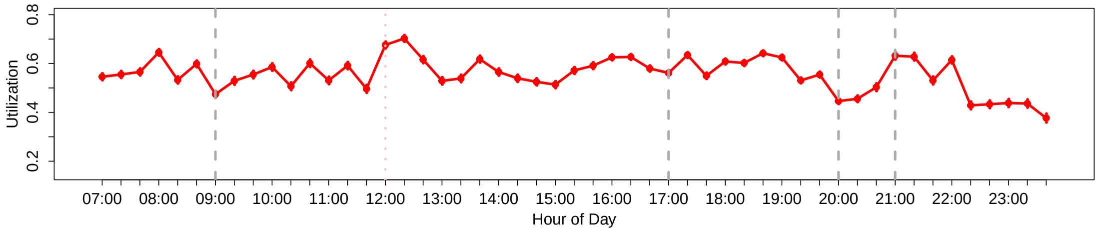

(a) Mondays

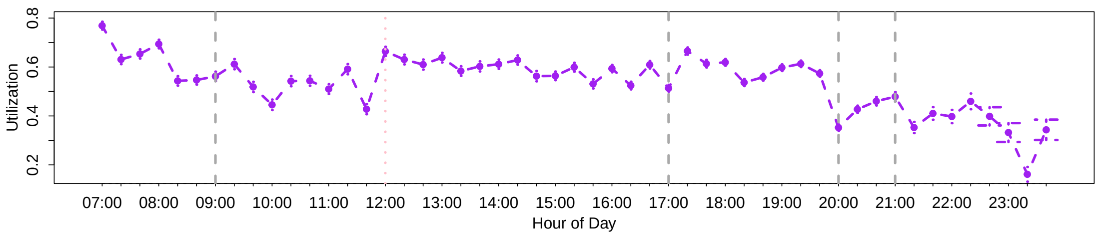

(b) Tuesdays

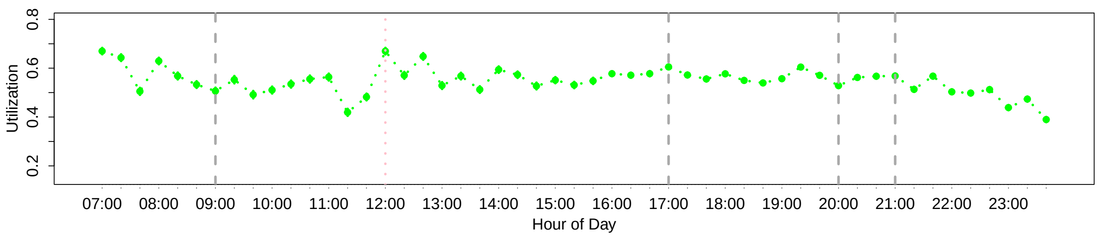

(c) Wednesdays

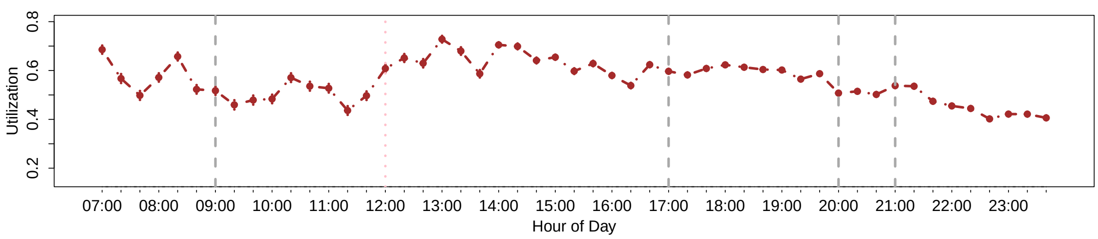

(d) Thursdays

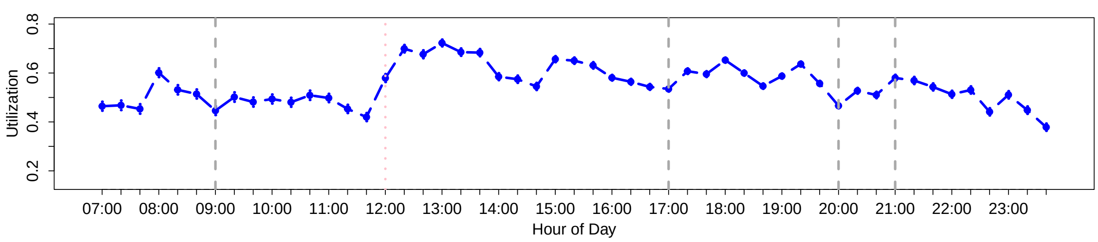

(e) Fridays

Figure A.1 Utilization rates by day of week and time of day. Grey dashed lines represent shift boundaries on the focal platform and the pink dotted line represents noon.

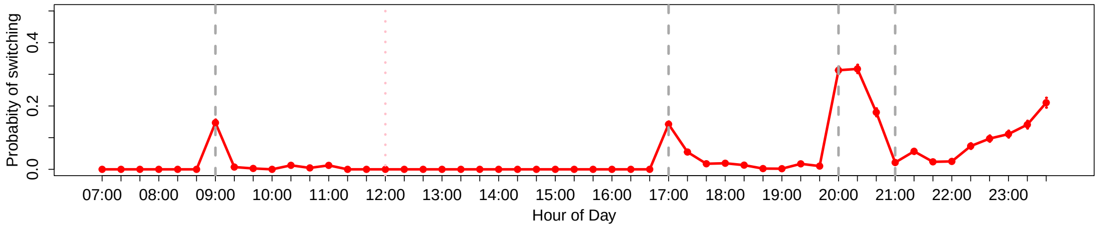

(a) Mondays

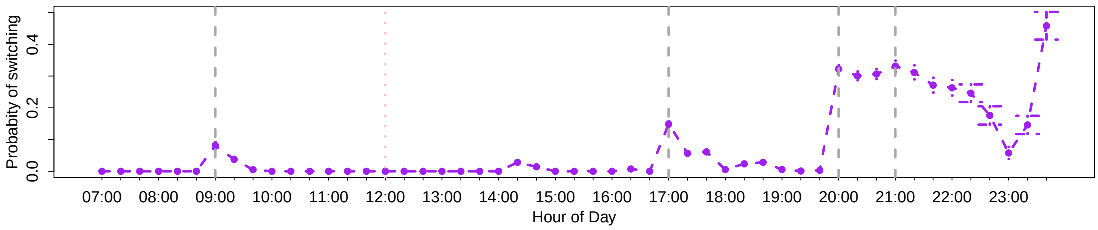

(b) Tuesdays

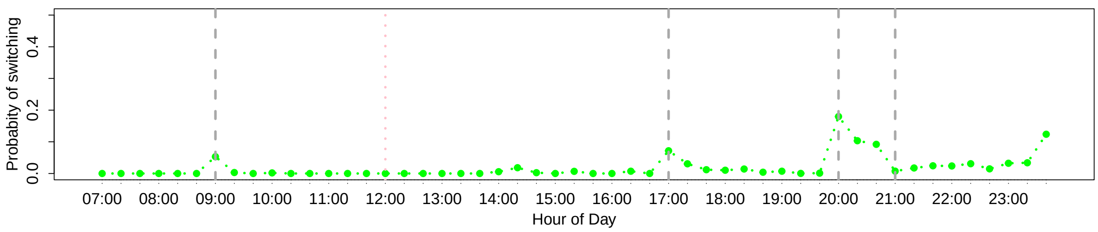

(c) Wednesdays

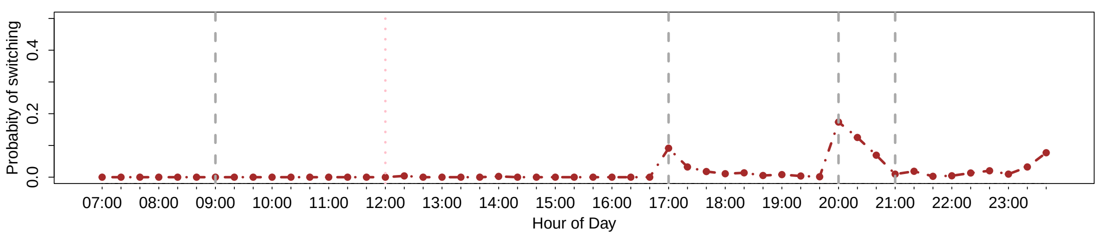

(d) Thursdays

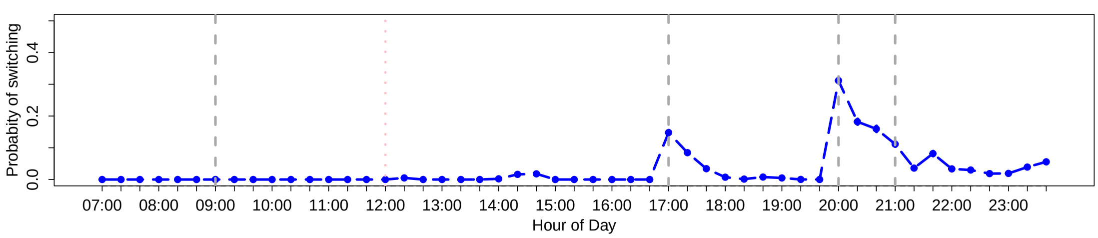

(e) Fridays

Figure A.2 Probability of switching to the competing platform by day of week and time of day. Grey dashed lines represent shift boundaries on the focal platform and the pink dotted line represents noon.