

Subsidizing Drivers in the Ride-Sharing Market: A Full-Heterogeneity Supply Model

Job Market Paper of Chong Bo (Zack) Wang

With Qiyuan Wang, Tat Chan, Dennis Zhang

Current Version: 2024-07-21

*For the most updated version, please refer to [this link](#)

Abstract

The sharing economy allows suppliers full autonomy over when and how long they work, creating challenges for managing the supply in response to demand changes. In this study, we investigate how to subsidize suppliers to increase supply in a cost-effective way within the context of a ride-sharing company. We study the key economic and behavioral factors that determine suppliers' daily work decisions with a structural model that accounts for the full heterogeneity of work costs across drivers. We combine a field experiment with observational data to identify the income sensitivity and work costs of drivers. A novel nested iteration procedure is proposed to address the computational challenge due to the high dimensionality of the parameter space, as such the model estimation becomes scalable. Using the estimation results, we conduct a counterfactual analysis to explore the costs of different subsidization strategies. We show that a *time-based non-targeting* subsidization based on the time drivers work will incur loss to the platform because drivers are not very income-sensitive; however, an *individual-based targeting* subsidization based on the cost estimates of individual drivers can help the platform save 51-83% of the cost. Our findings highlight the importance of understanding and leveraging driver heterogeneity to improve the profitability of sharing platforms.

1. Introduction

Gig economy, particularly the ride-sharing sector, has experienced a significant growth in recent decades. According to a recent survey¹, the global market size of gig economy was \$91.63 billion in 2023 and is expected to reach \$418.53 billion by 2033. Unlike traditional industries, the gig economy offers suppliers much flexibility in supply decisions. As an illustration, different from taxi drivers who have to follow schedules, drivers on ride-sharing platforms can decide when and how long to work. Such flexibility however creates challenges for platforms to manage supply in response to demand changes. When the demand is high, for example, a low capacity of supply can lead to revenue loss and even affect future profit due to unsatisfactory experience on the demand side.

Various strategies have been employed to balance demand and supply. Platforms have typically relied on AI matching algorithms that dispatch drivers to satisfy rider demand.² Surge pricing has also been used in the time of high demand. While surge pricing helps increase the driver supply, riders may feel being exploited and thus have poor usage experience. Another common strategy is to subsidize drivers, offering them extra incentive to work during certain hours and in certain areas when demand surges. This strategy however is costly for ride-sharing platforms as they have to sacrifice their share of revenue from each order. It is important for platforms to optimize subsidies that can effectively motivate the supply and at the same time minimize the cost. This is the research focus of this paper.

An effective subsidization strategy in response to demand changes requires a comprehensive understanding of the underlying factors that drive the work decisions of drivers, as these factors will determine their responsiveness to the subsidy. To study this important managerial question, we use a rich dataset from a ride-sharing platform in Vancouver, Canada. Similar to other ride-sharing platforms such as Uber and Lyft, drivers on the platform differ in their overall tendency to work. While some drivers work for long hours in a day, others choose to work for significantly fewer hours and even not to work at all. Furthermore, some drivers prefer to work in the morning but some others prefer to work in the afternoon or at night. There is also a large heterogeneity in their work choices during weekdays and weekends.

¹ https://finance.yahoo.com/news/global-ride-sharing-market-size-110000245.html?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAANybyFh_xzwIelSlfA70SeE31Z_SBgv8t3KQeAj26aGH6KFSrbQpytzKaNTN-JQL-m2Jx0K8vw4Uka78Sf5qJaLTEX-LQKM0x9CWVHemWcNQulOFt0YLM-2Vs7xFX5QA9gV80aiaMDkHw8gfAu6GjuTjXNilluAPvJ44Hfl-ZxmE3, assessed on May 17, 2024

² For example, see “*Quantifying Efficiency in Ridesharing Marketplaces*” (by Alex Chin, <https://eng.lyft.com/quantifying-efficiency-in-ridesharing-marketplaces-affd53043db2>) that describes how the core system at Lyft works.

To have a full understanding of what leads to the heterogeneity in work choices, we develop and estimate from data a structural model that can quantify the different economic and behavioral factors driving drivers’ decisions. In our model, drivers evaluate not only the expected gain but also the cost of driving to decide the optimal daily work schedule that will maximize their payoff. To fully capture the heterogeneity in drivers’ work choices, our model employs a semi-parametric approach. While the payoff function is parametric, we allow each driver to have a unique marginal cost of work, and that the cost varies across morning, afternoon, and night, as well as between weekdays and weekends. Such a model specification allows us to profile drivers based on their driving costs in different time periods.

Our study faces two main challenges. First, to predict how drivers react to subsidies we have to estimate their income sensitivity. However, it is difficult to use observational data alone to separate income sensitivity from work cost. For example, observing many drivers working at nighttime could be due to a high income sensitivity (i.e. they are motivated by more orders during those hours) or a low cost (e.g. they do not work for full-time jobs at night). To tackle this challenge, we collaborate with the platform by running a field experiment for four weeks, during which we exogenously manipulate the level of subsidies across time periods. Such experimental data helps us identify the sensitivity of drivers for the expected income change. We combine the experimental data with observational data, in which drivers did not receive any subsidy, to further pin down the individual- and time-specific work costs.

The second challenge is that, since we estimate the work costs of each driver, the dimensionality of model parameters increases linearly with the number of drivers and, consequently, the computational burden grows exponentially. As the model is highly non-linear, the dimensionality issue may also lead to local optima during the model estimation. If the issue cannot be resolved, it implies that our model estimation is not *scalable* and, therefore, whatever solutions based on our study cannot be implemented for large businesses.

We propose a novel nested iteration estimation method to tackle this challenge. Specifically, we divide the model parameters into one set of parameters, Θ_1 , which are shared by all drivers, and another set, Θ_2 , that are individual-specific (i.e. each driver’s work costs). Only the dimensionality of Θ_2 increases with the number of drivers. Our estimation proceeds as follows: given trial Θ_1 , the inner loop estimates Θ_{2i} for each driver. In the outer loop, we search for the optimal Θ_1 that maximizes the full likelihood. The inner loop is repeated each time Θ_1 is updated at the outer loop. Since Θ_2 is estimated separately for each driver in the inner loop, while the dimensionality of Θ_1 remains fixed, the complexity of model estimation only increases linearly with the number of drivers. We use simulation studies to show that, compared with the standard method of jointly estimating Θ_1 and Θ_2 , the nested iterative estimation procedure can significantly reduce the model

estimation time, even when there is only a small number of drivers. Furthermore, the model estimation can be facilitated by large-scale parallel processing in the inner loop using the platform's computation resource. Therefore, business models that rely on the model estimation results are scalable.

The estimation results reveal substantial cost heterogeneity across drivers. Within each time period (e.g., weekday morning), there is a significant spread in the distribution of cost estimates, which generally follow a normal distribution. Additionally, we observe a high correlation between the estimated costs within a day for the same drivers, with higher correlations between adjacent time periods than between non-adjacent ones. For instance, the correlation between weekday morning and weekday afternoon is higher than that between weekday morning and weekend afternoon. Drivers also exhibit similar working preferences between weekdays and weekends: the correlation of costs for the same time periods on weekdays and weekends is higher than for different time periods across these days. Furthermore, we find that drivers are not income-sensitive, as 1% increase in income only results in a 0.06% increase in working time. Our analysis also indicates that drivers face startup costs, psychological flow, and fatigue when deciding when to work.

Using the model estimates, we conduct a counterfactual study to investigate how to subsidize drivers in response to demand changes in a cost-effective way. We assume that the platform faces a 5% increase in demand on a specific day and needs to increase drivers' work hours to balance this demand. In the base case, we assume no subsidy is offered and calculate how drivers are incentivized by the higher expected number of orders to increase their work hours, based on the model predictions. We find that even in the most ideal scenario, where the 5% increase in demand is fully reflected in drivers' expectations, the total working time increases by just around 0.035% and consequently the total number of orders can be fulfilled by drivers only increases by around 4%, suggesting that the supply does not match the increased demand without subsidization. This minimal increase is due to drivers' low income sensitivity. These results serve as a benchmark for the subsidization policy. The first counterfactual policy we consider is *time-based non-targeting* subsidization, under which the platform offers a unique subsidy in each time period (morning, afternoon and night). Drivers who work during the time period are offered the same amount of subsidy for each order they complete. This policy is similar to the second-degree price discrimination. In the second scenario, we assume that the platform offers *individual-based targeting* subsidization with the objective of minimizing the subsidization cost. Who to offer subsidies and how much are the subsidies are based on the individual cost estimates. This practice is similar to the third-degree price discrimination. In both scenarios, we impose the constraint that the increase in the supply capacity, i.e. the total number of orders that can be fulfilled by drivers

at work during the time period, cannot be smaller than 5%, as such the supply can balance the new demand.

Under the time-based non-targeting subsidization, we find that the platform must offer significant subsidies, which are generally highest in the morning (when supply is low) and lowest at night (when supply is the highest), to balance the increased demand. Consequently, the platform incurs a loss of 14% to 47% relative to what it could earn without offering subsidies. Under the individual-based targeting subsidization, we find that the best strategy relies on how the demand increase affects drivers' income expectation. If the income expectation fully reflects the higher demand, subsidization should be offered to drivers with higher work costs; otherwise, the company should target drivers with lower costs is more effective. In either case, subsidies should be offered to drivers with high work efficiency, i.e. those who can take more orders in an hour. The subsidization cost in this targeted approach is 51% to 83% lower than under the time-based non-targeting subsidization. Consequently, the platform can achieve 79% to 98% of the profit level when it does not offer subsidies. Although the profit is still lower than the benchmark, having more drivers on the street to fulfill the increased demand helps maintain customer satisfaction. Balancing demand and supply is critical for the platform's long-term growth.

Our study contributes to the marketing literature in two ways. Substantively, we show how a ride-sharing platform can target individual drivers with subsidization in a cost-effective way. In addition to AI matching algorithm and surge pricing, subsidization is an important instrument for ride-sharing platforms to manage driver supply in response to the demand change. Methodologically, we introduce a novel nested iteration procedure to make sure that, while our structural model captures full heterogeneity among drivers, the model estimation is still scalable. As such, the method can be used by large businesses in other empirical settings. As an example, ecommerce platforms can apply our method to estimate individual-specific product preferences and use the results to facilitate targeting promotions at the customer level.

The rest of the paper proceeds as follows: we first review the relevant literature and discuss our contribution. We then explain the empirical context and the dataset used in the paper and show reduced-form evidence regarding drivers' supply behavior. Building on the evidence, we develop our structural model and propose a novel nested iteration method for estimation. Next, we show the estimation results and the insights derived from the counterfactual analysis. Finally, we conclude the paper with discussion on implications and limitations.

2. Literature Review

Our paper is connected to several streams of research. First, our study contributes to the general literature on labor supply. As one of the most important economic decisions, labor supply is critical for understanding various micro and macroeconomic outcomes (Chetty et al., 2011; Keane and Rogerson, 2015). Previous literature investigates the labor supply choices at both the extensive margin (i.e., the choices about whether to participate in labor force) and intensive margin (i.e., the choices about hours of work) (Meyer 2002; Chetty et al., 2011; Attanasio et al., 2018). These decisions can be affected by various external and internal factors, such as tax (Keane 2011), expected wage (Attanasio et al., 2018), commuting (Monte et al., 2018), and household composition (Pabilonia and Ward-Batts, 2017). Our paper is particularly related to the labor supply decisions in the taxi industry. While taxi industry shares similarities with many other industries in terms of labor supply decisions, taxi drivers have the flexibility of when they can terminate their work shift for day. A large stream of literature therefore identifies whether drivers have an earnings target and how it influences their decision to stop working. The seminal work by Camerer et al. (1997) suggests a negative wage elasticity such that cumulative earnings within a day increase the likelihood of quitting work within the same day. They interpret this as evidence that drivers have an earnings target, and they are more likely to stop working after reaching this target. Following the work, several subsequent studies revisit the earnings target hypothesis using various samples and methodologies. For instance, Farber (2005) and Farber (2013) use alternative measurements of daily wages and larger samples to show that drivers respond positively to unanticipated as well as anticipated increases in earnings opportunities. Other studies formally develop structural models incorporating drivers' reference-dependent preferences and show that drivers may indeed have both earnings targets and supply hour targets (Farber 2008; Crawford and Meng 2011; Thakral and To 2021). This stream of literature primarily focuses on drivers' supply in the traditional taxi industry, where drivers usually have scheduled shifts. Therefore, drivers' primary supply decision is whether they stop working for the day. This is different from ride-sharing platforms, where drivers can flexibly choose when to work. For instance, drivers can work in the morning and night while taking a break in the afternoon. Our research complements this stream of literature by explicitly modeling such flexibility and providing novel insights into how drivers' heterogeneous costs influence their working decisions.

Our paper is also related to the research examining users' participation on sharing economy platforms. Previous studies suggest that the gig economy provides a flexible channel for people to find employment opportunities. Therefore, gig economy platforms are particularly relevant for unemployed and underemployed workforces (Burtch et al., 2018). For instance, Huang et al. (2020) show that unemployment is positively associated with participation in the online labor market,

particularly in areas with better internet access and younger, more educated populations. Similarly, using data from South Korea, Liu et al. (2023) demonstrate that the entry of online food delivery platforms can significantly boost the employment rate for female workers who can better use their time beyond their housework. Specifically for the ride-sharing platforms, Chen et al. (2022) examine driver supply on DiDi, the largest ride-sharing platform. They similarly develop a structural model to capture drivers' supply decisions but do not allow for full heterogeneity among drivers. Most relatedly, Chen et al. (2019) investigate the value of flexible work on Uber and estimate drivers' reservation wages for each hour using a Bayesian method. Therefore, drivers' heterogeneity in their model is captured in their reservation wages. Building on prior literature, we capture drivers' heterogeneity through their working costs. The advantage of this approach is that we can incorporate drivers' temporal preferences for working (e.g., working on weekday mornings rather than weekend mornings) and interdependency over time, such as how working in the morning influences working in the afternoon within the same day. Therefore, our cost specification captures more nuances of driver supply on the platform.

Finally, the insights from our empirical and counterfactual analysis also contributes to the literature on subsidy. Firms and governments routinely use subsidization to nudge people's behavior towards the desirable direction. For instance, the famous U.S. federal subsidy program, the Special Supplemental Nutrition Program for Women, Infants, and Children (also known as WIC), provides subsidy to purchasing food and aims to safeguard the nutrition in the diet for vulnerable population groups (Owen and Owen 1997). Previous study demonstrate that the subsidy can influence subsidized consumers' food consumption (Hinnosaar 2023), birth (Hoynes et al., 2011), and long-term lifetime outcomes (Jackson 2015). Prior papers have studied several other programs, such as subsidizing rural consumers to purchase household appliances (Xiao et al, 2020), agricultural production subsidy (Fan et al., 2024), and subsidy for medical innovation (Olsder et al., 2022). The subsidization program studied in prior literature is usually uniform for eligible recipients. In contrast, we leverage the information on drivers' heterogeneous preferences and customize the targeted subsidy at the individual level. We demonstrate that such subsidy design can significantly reduce the cost for the platform and achieve better outcomes.

3. Empirical Setting and Data

3.1. The Platform

Our study is conducted in collaboration with a local ride-sharing platform in Vancouver, Canada, which resembles well-known services like Uber and Lyft. This platform pairs available drivers with riders seeking transportation. The matching process begins when drivers log onto the driver-

side application, indicating their readiness to accept rides. When a ride request is made within a certain proximity, the driver receives a notification to either accept or decline the ride. Before accepting, drivers are not informed of ride details such as the final destination, fare, or rider information. Once a ride is accepted, the platform reveals the pickup location and the rider's contact details, but the final destination is disclosed only when the driver picks up the rider. This protocol is designed to prevent selective acceptance of rides based on destination, promoting equitable ride opportunities across the platform.

Like many other ride-sharing platforms, the platform we study allows drivers to decide their work schedule, indicating the voluntary nature of driver participation. This policy ensures efficient management of driver availability and system resources while respecting driver autonomy, enabling them to adjust their work hours to their personal needs. We measure drivers' working status after they log on to the platform's app because it indicates that they are ready to take orders. Since we also know each order's pick-up and drop-off time, we can use this information to infer whether the car is vacant or occupied when the driver is working.

After a completed trip, riders pay a fee to the platform, calculated based on the trip's distance, duration, and timing. For drivers, the platform collects a fixed initiation fee from each trip. Generally, drivers earn 80% of the remaining trip fee as their income, while the platform retains the remaining 20%. To incentivize driver supply during peak demand, the company regularly implements surge pricing, multiplying the total order price by a factor dependent on market conditions. While effective in increasing driver earnings, surge pricing can lead to rider dissatisfaction due to higher costs. Alternatively, the platform also offers subsidies, providing drivers with an additional amount for each completed order within specified periods. This strategy avoids high rider charges while motivating drivers to accept more orders. Compensation for drivers is distributed weekly and is derived from both the trip fee and subsidies.

3.2. Subsidy Experiment

In this paper, we study drivers' participation decisions on the platform. These decisions are largely based on two main factors: drivers' expected income and working costs (Farber 2008; Crawford and Meng 2011; Chen et al. 2019). Without exogenous variation, it can be difficult to identify these two factors separately. For example, observing many drivers working at a certain time period could be attributed to either a high-income sensitivity or a low cost. To better distinguish expected income from working costs and separately identify their impacts on drivers' participation decisions, we combine a subsidy experiment with observational records to assess drivers' income sensitivity.

We collaborated with the platform to conduct the field experiment (more details can be found in Wang et al. 2023). The experiment exogenously varied the subsidy provided to drivers to influence

their supply behavior. The value of the experiment to our study lies in the subsidy manipulation, which results in exogenous variation of drivers' expected income, uncorrelated with drivers' working costs. This helps us disentangle how expected income and working costs distinctively influence drivers' decisions to work on the platform.

The experiment lasted for four weeks, from June 20 to July 17, 2021. Following the platform's classification, time slots are divided into afternoon (12:00 PM to 5:59 PM) and night sessions (6:00 PM to 11:59 PM). We exogenously set the subsidy at \$2, \$4, or \$6 per trip for afternoon sessions and \$0, \$3, or \$5 per trip for night sessions. These subsidies were applied differently across the week, divided into three-day groups: Monday to Thursday, Friday to Saturday, and Sunday, with each group receiving rotating subsidies. Each subsidy level was offered once per week in each time slot. The initial three weeks followed a detailed and balanced subsidy schedule, while the fourth week's schedule was simplified, offering only \$0, \$4, and \$5 subsidies across all time slots. Details of the subsidy schedule is in Table 1.

Table 1. Subsidy Schedule (\$ Per Completed Trip)

Week 1			Week 2		
Mon.-Thu.	Fri.-Sat.	Sun.	Mon.-Thu.	Fri.-Sat.	Sun.
6	2	4	2	4	6
0	5	3	3	0	5
Week 3			Week 4		
Mon.-Thu.	Fri.-Sat.	Sun.	Mon.-Thu.	Fri.-Sat.	Sun.
4	6	2	4	5	5
5	3	0	0	0	5

3.3. Data

The platform provides a detailed dataset with comprehensive demand and supply information for all drivers and riders. This dataset not only covers the experiment period (June 20 to July 17, 2021) but also extends from January 1 to July 30, 2022, during which no subsidy was offered. The additional information helps enhance the identification of model parameters. We intentionally excluded data from July 18 to December 31, 2021, due to several field experiments the ride-sharing company conducted, which could skew the drivers' earning behaviors.³

³ For example, one experiment during this period featured a competition where drivers could earn monetary rewards by completing a designated number of trips within a specific timeframe.

The demand-side data is organized at the order level, where an order represents a ride request made by a rider. For each order, we have the rider ID, driver ID, and timestamps for order requests, driver acceptance, pickup, and drop-off. Additionally, we know the key financial details for each order, including the price of the trip, any surge pricing applied, subsidies provided, and the total earnings accrued by the driver. This rich dataset enables us to analyze and estimate drivers' expected income for specific periods during the week, providing insights into how earnings fluctuate based on time and day.

On the supply side, the observation unit is at the driver-shift level. For each shift, we observe the driver ID and the driver's logon and logoff timestamps, allowing us to determine when drivers begin and end their work in a day. It should be noted that a driver's car can be either vacant or occupied during a shift. We merge the shift data with the trip data to ascertain the status of the driver (i.e., vacant or occupied).

This comprehensive dataset allows us to conduct a detailed analysis of drivers' participation decisions by examining their expected income and working costs across different periods and conditions.

3.4. Variable Construction

Time Periods. As discussed previously, drivers make their work decisions primarily considering the expected income and working cost, which depend on the time period. To facilitate the analysis, we treat a week as one complete temporal cycle and aggregate days and hours into different time periods. Specifically, we divide a week into six time periods: weekday morning, weekday afternoon, weekday evening, weekend morning, weekend afternoon, and weekend evening. Weekdays are Monday to Thursday, whereas weekends are Friday to Saturday. Mornings are from 6:00 AM to 11:59 AM, afternoons are from 12:00 PM to 5:59 PM, and evenings are from 6:00 PM to 11:59 PM. We exclude early mornings (i.e., 12:00 AM to 5:59 AM) from our analysis because of the very low level of activity on the platform.

We make this time aggregation for two reasons. First, the ride-sharing platform uses the same time classification during the experimental period. Our time aggregation is aligned with the platform operation. Second, the time aggregation is also consistent with temporal patterns of the driver's working behavior on the platform. As shown in Figure 1, the average percentage of online drivers increases in the morning, stabilizes at a high level in the afternoon, and gradually decreases at night. Similarly, we show the average number of working hours by the day of the week in Figure 2. We find that drivers generally work more on weekends than weekdays, with the exception of Sunday. However, because Sunday has similar characteristics to Saturday—people generally have the whole day to plan without regular day jobs—we still include it in the weekend category.

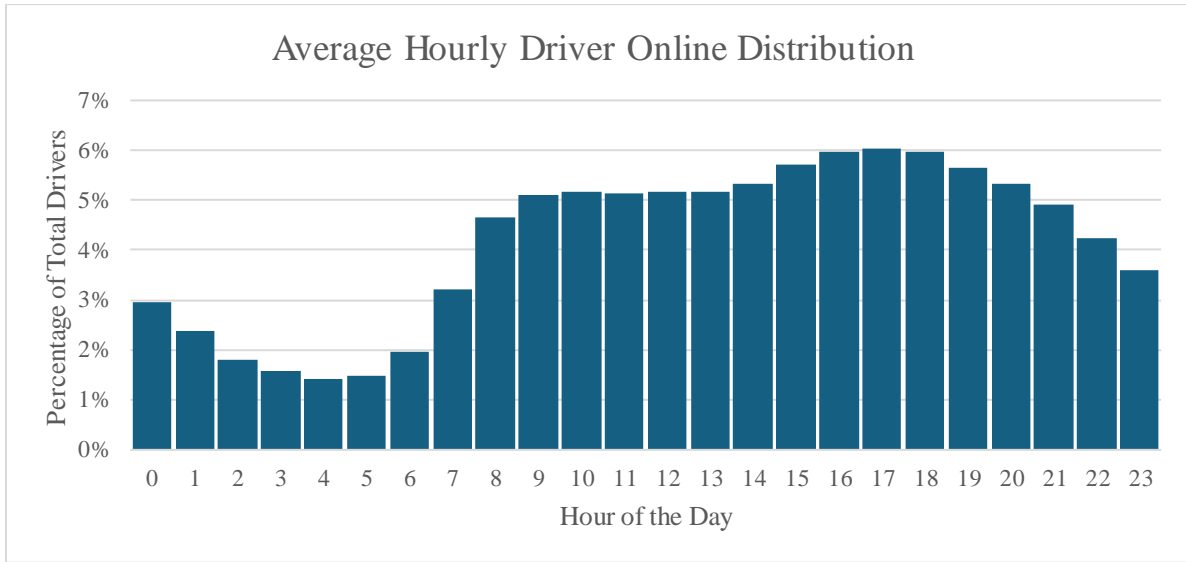


Figure 1. Average Hourly Driver Distribution




Figure 2. Average Number of Working Hours by Day of the Week

Expected income. To model drivers' working decisions, we need to measure their expected income. Directly asking drivers' beliefs is challenging and may be biased. Following prior

literature (Farber 2008; Crawford and Meng 2011; Thakral and To 2021), we use the observed order data to approximate drivers' beliefs. This is justifiable because drivers should also form their expectations based on their real earnings, which are reflected in the demand data.

We construct the belief on the expected income at the driver and hour level. This is because, similar to drivers' working behavior, we find substantial cross-driver and temporal heterogeneity in drivers' income. Specifically, for temporal heterogeneity, the total number of orders varies substantially within and across days. The total demand increases from morning to afternoon, peaks at 6 PM, and then gradually declines. The total demand is also higher around weekends (i.e., Friday to Sunday) than weekdays (i.e., Monday to Thursday). This temporal heterogeneity appears to be correlated with people's mobility needs due to work or leisure activities. As for cross-driver heterogeneity, given the same time slot, the number of orders each driver takes differs significantly. This may be attributed to drivers' heterogeneous skills in "finding" orders. For example, given that riders' requests will be dispatched to nearby drivers, cruising in a neighborhood with high demand can help drivers get more orders. Therefore, more experienced drivers may be more proficient in identifying such high-demand areas and eventually have more orders.

To incorporate such cross-driver and temporal heterogeneity, we decompose the expected income for driver i in hour t as the product of the expected number of orders and the expected earnings per order:

$$\text{Expected Income Per Hour}_{it} = \text{Expected Orders Per Hour}_{it} * \text{Expected Earnings Per Order}_t \quad (1)$$

For *Expected Orders Per Hour*_{it}, we first calculate, in each hour (e.g., Monday 10 AM), the average number of orders a driver actually took during our sampling period, conditional on the driver working in this hour slot. Since we group hour slots into time periods, we pool all observations within a time period to calculate the average. One limitation of this calculation is that drivers should work at least once for a given time slot. However, it is likely that some drivers may never work in a time period (e.g., weekday morning), leading to a missing data problem. To address this limitation, we impute the expected number of orders for the time period that a driver never works by multiplying a factor with the average expected number of orders for all other drivers. We obtain this factor by comparing the average expected number of orders for this driver and all other drivers in other time slots. The basic idea is that if a driver takes many (or few) orders relative to the platform average level in general, he should similarly take many (or few) orders if he had ever worked in the time period (i.e., weekday morning).

The expected earnings per order is the sum of the expected fee paid to the driver and the subsidy per order offered by the platform in hour t :

$$\text{Expected Earnings Per Order}_t = \text{Expected Fee Per Order} + \text{Subsidy Per Order}_t \quad (2)$$

While the subsidy per order is directly controlled by the platform and observed by us as researchers, we need to measure the expected fee per order paid by the rider. We calculate the average fee across all orders, which has no temporal and cross-driver variation. We make this simplification for two reasons. First, while the incidence of orders may be highly dependent on time, the duration of the trip, a key factor that determines the fee, should be much less dependent on time, resulting in limited temporal variation. For instance, there is no strong reason to believe that riders tend to take significantly longer trips in the morning than in the afternoon. Second, drivers do not know the order details until they pick up the order. Therefore, order selection may be difficult for drivers, leaving much less room for them to influence the order fee with their skills. This limits cross-driver variation. It should be noted that our model does not rely on the simplification of the expected fee per order. We can easily accommodate heterogeneous expected fees if needed.

Supply. We construct drivers' work hours using the driver shift data. We segment each driver's shift into hourly intervals and define that a driver is working in an hour if they are logged into the platform app for at least 10 minutes (Chen et al., 2019). Therefore, our main outcome variable for measuring drivers' supply is the number of hours they worked.

3.5. Descriptive Evidence

Using our datasets, we present descriptive evidence to understand the underlying factors that influence drivers' work behavior.

The Heterogeneity in Driver Supply. As previously discussed, the ride-sharing platform allows for supply flexibility, enabling drivers to choose when and how long they work. Due to varying beliefs about income and work costs, observed supply behavior can differ significantly across time and drivers.

To understand when drivers work, we first examine the distribution of their starting and ending times each day. For analytical purposes, we only focus on the duration from 6 AM to 12 AM each day. The starting time is defined as the earliest time drivers log on to the platform app each day, and the ending time is the latest time they log off. As shown in Figure 3, the starting and ending times vary substantially across drivers. Around 55% of drivers start working after 9 AM, whereas nearly 30% stop working before 5 PM. This differs significantly from the common working schedule of employees. It is important to note that this does not mean drivers work continuously between these times, as they may take breaks.




Figure 3. Drivers’ Work Starting and Ending Time Distribution

Drivers also differ substantially in the length of their work periods. We first show the distribution of total working hours per week in Figure 4. As suggested by Figure 4, drivers exhibit different levels of activity on the platform. For instance, about 20% of drivers work over 40 hours, while the majority work much less, primarily because most drivers are part-time.




Figure 4. Weekly Working Hours Distribution

We further explore the heterogeneity in drivers' working choices by examining time periods. For each driver, we calculate the percentage of working hours in each time period during our sample period and plot the distribution of these percentages by different time slots in Figure 5. It shows a large variation in preferences for time slots. For example, there are significant proportions of drivers whose working time is above 50% in the afternoon and at night (see the right side of the figure). The morning slot is the least popular choice of working time, as it exhibits the highest number of drivers working less than 10% of their total time in that slot (see the left of the figure).




Figure 5. Working Hours Choice in Different Time Slot Distribution

We also compare the differences between weekdays and weekends. Specifically, we calculate the total weekday and weekend working time over the total working time in a week for each driver and plot this in a histogram. The histograms shows that, while a large proportion of drivers split their working time equally on weekdays or weekends (see the mode at 50-60% working time in the figure), a significant number of drivers have a very strong preference for working on either weekdays or weekends, as they spend above 60% of working time on either weekdays or weekends (see the right side of the figure).

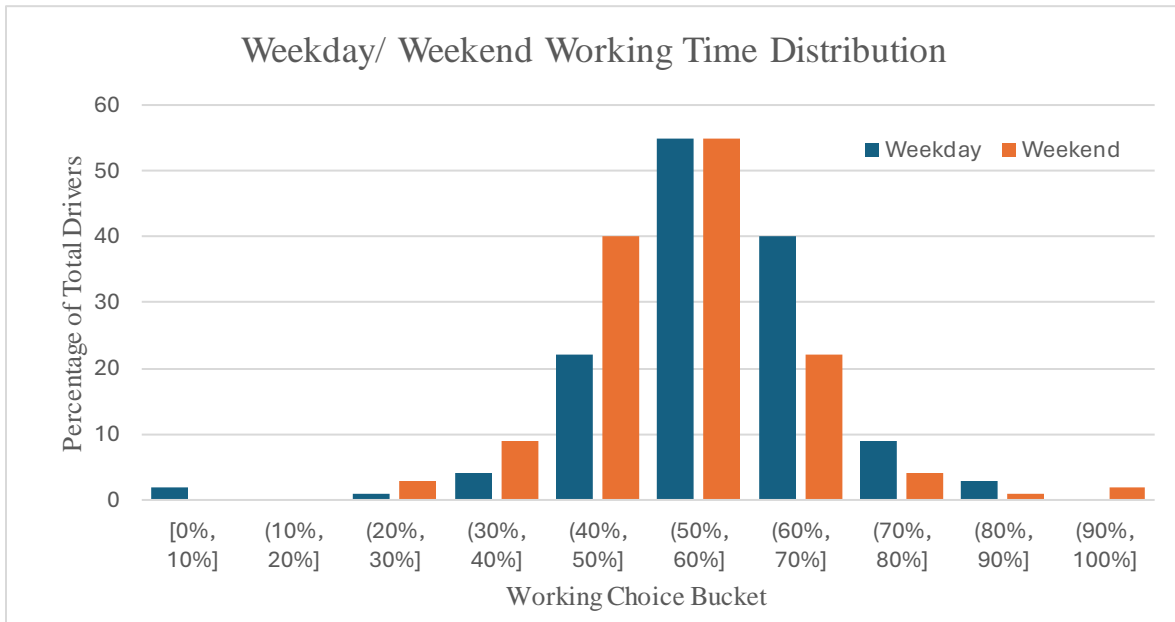


Figure 6. Working Hours Choice in Weekday/Weekend Distribution

In summary, drivers' preferences for both working times and workday choices vary significantly across individuals. This variability is likely to be driven by not only drivers' working time preference but also their outside activities, in particular whether they have to work for full time jobs. For instance, if drivers need to work for a regular job in weekday morning, it would be more difficult for them to work for the ride-sharing platform. We capture such differential value of time across drivers by using a flexible cost function. Understanding such cost heterogeneity at the individual level is important for more effectively incentivizing them. For instance, subsidizing a driver who will never choose to work at weekday morning might not be effective, as the driver may have other work obligations during that time.

The Impact of Income on Driver Supply. Since our major goal is to investigate how to use subsidies to encourage drivers to work, we further examine how drivers respond to subsidies. Using the experimental data, we regress whether drivers work during a specific hour on the exogenously manipulated subsidy. As discussed in the variable construction, we define a driver as working for an hour if they log into the platform app for at least 10 minutes. We control for driver fixed effects, day-of-week fixed effects, and hour-of-day fixed effects. The results in Column (1) of Table 1 suggest that increasing the subsidy encourages drivers to work. Since the subsidy is exogenously determined by the platform, we can interpret the estimate as the causal positive impact of subsidy on work, indicating that drivers' working decisions are sensitive to income variation.

Since drivers decide both whether to work and how long to work, we further examine the impact of subsidy on drivers' working duration. The results in Column (2) of Table 1 show consistent findings, indicating that drivers are more likely to work for longer time with higher subsidies.

Although the estimated coefficients for whether to work and how long to work are statistically significant, the magnitudes are relatively small, suggesting that the supply is not very elastic toward subsidy changes. This may imply that subsidies are not effective in encouraging supply. To investigate whether the platform can offer a more cost effective subsidization plan, which is the main goal of this study, we develop a structural model to fully capture the heterogeneity in drivers' supply behaviors.

Table 2. The Impact of Subsidy on Driver Supply (Driver Hour Level Analysis)

	DV: Whether Working	DV: Working Duration (Mins)
Log(Subsidy+1)	0.001*** (0.0004)	0.052*** (0.020)
Driver Fixed Effect	Yes	Yes
Day of Week Fixed Effect	Yes	Yes
Hour of Day Fixed Effect	Yes	Yes
Observations	195,648	195,648
R ²	0.268	0.259
Adjusted R ²	0.268	0.258

4. A Structural Model of Driver Supply

4.1. Model Overview

We develop a structural model to characterize drivers' supply decisions on the ride-sharing platform. The supply decision involves determining how many hours a driver i will work on day d . We consider the decision at the time period level $p \in (1, 2, 3)$, corresponding to morning ($p = 1$), afternoon ($p = 2$), and night ($p = 3$). Drivers will choose the number of work hours (from 0 to 6 hours) in each time period, namely $h_{idp} \in (0, 1, 2, \dots, 6)$. We provide a summary of notation in our model in Table 3.

Table 3. Notation Explanation

Notation	Meaning
i	Index for driver
d	Index for calendar day
p	Index for time period for the day with $p = 1$ for morning (6:00 AM- 11:59 AM), $p = 2$ for afternoon (12:00 PM to 17:59 PM), $p = 3$ (18: 00 PM to 23: 59 PM) for night
w_d	Index for day of the week (Monday - Sunday) for day d
h_{idp}	The working hour driver i has at time period p on day d
\mathbf{h}_{id}	The vector of working hour driver i has on day d
$Earning\ Per\ Hour_{idp}$	The expected earning hour driver i has at time period p on day d
$\overline{\# Order\ Per\ Hour}_{iw_d p}$	The expected number of orders driver i has at time period p on day of week w_d
$\overline{Earning\ Per\ Order}$	The expected earnings per order
$Subsidy_{dp}$	The subsidy per order offered by the platform at time period p on day d
ξ_{idp}	The unobserved demand and cost shocks for driver i at time period p on day d
α	The income sensitivity shared by all drivers
$c_{1iw_d p}$	The marginal cost for driver i at time period p on day of week w_d
c_2	The cost parameter that captures the benefit from work flow
c_3	The cost parameter that captures the benefit from work continuity
c_4	The cost parameter that captures the work fatigue
Σ_ξ	The variance-covariance matrix for the distribution of the demand and cost shocks ξ_{idp}

Drivers decide their working hours by evaluating their gain from the expected income and working costs. We specify the driver's utility $u(\mathbf{h}_{id})$ as following:

$$u(\mathbf{h}_{id}) = \text{Gain}(\mathbf{h}_{id}) - \text{Costs}(\mathbf{h}_{id}), \quad (3)$$

where $\text{Gain}(\mathbf{h}_{id})$ and $\text{Costs}(\mathbf{h}_{id})$ are associated with the expected income and costs conditional on working hours \mathbf{h}_{id} which consists of supply decision throughout the day, namely $\mathbf{h}_{id} = (h_{id1}, h_{id2}, h_{id3})$. Drivers thus choose the working hours \mathbf{h}_{id}^* for the whole day that maximize their total utility:

$$\mathbf{h}_{id}^* = \underset{\mathbf{h}_{id}}{\operatorname{argmax}} u(\mathbf{h}_{id}). \quad (4)$$

We consider two important features of drivers' supply behavior when we specify the utility function below. First, drivers' utility across different time periods may be correlated within a day. Such correlation can be caused by both *observable* factors and *unobservable* factors. For instance, for observable factors, working longer in the morning may inevitably influence drivers' state to work in the afternoon either by increasing or decreasing their working costs. For unobservable factors, for example, road repairs, which we do not observe in our data, can cause a negative shock to drivers' expected income and costs throughout the whole day. We account for the role of both observable and unobservable factors in drivers' utility specification.

Second, drivers may have heterogeneous expected income and working costs conditional on time. As shown in the descriptive evidence, drivers vary substantially in terms of when and how long they work. Such heterogeneity can be attributed to differences in both their expected income and working costs. To accommodate such differences, we compute the heterogeneous expected income at the driver level to reflect that drivers' earnings can be dispersed even if they work in the same hour. Furthermore, we allow for the full heterogeneity in drivers' costs at the individual level. However, such model specification leads to very challenging estimation because we have a very large set of parameters to estimate. To address this challenge, we propose a novel estimation method that allows us to recover individual cost parameters in a tractable way.

4.2. The Gain Function

First, we discuss drivers' gain from expected income in the utility function $u(\mathbf{h}_{id})$. With working hour h_{idp} , driver i 's gain from the expected income in time period p on day d can be expressed as:

$$\text{Gain}_{idp}(h_{idp}) = \alpha * E(\text{Earning Per Hour}_{idp}) * h_{idp}, \quad (5)$$

where α denotes the income sensitivity shared among all drivers, indicating how changes in potential earnings influence their decisions to work additional hours. This parameter is crucial for modeling their responsiveness to varying hourly earnings. $E(\text{Earning Per Hour}_{idp})$ is the expected hourly earnings for the driver in time period p on day d .

We approximate drivers' expected earning per hour using the following formula:

$$E(\text{Earning Per Hour}_{idp}) = \overline{\# \text{ Order Per Hour}_{iw_{dp}}} * (\overline{\text{Earning Per Order}} + \text{Subsidy}_{dp}) + \xi_{idp}^e. \quad (6)$$

In this equation, $\overline{\# \text{ Order Per Hour}_{iw_{dp}}}$ denotes the average number of hourly trips undertaken by the driver i per hour in period p on day d that is on the day of week w_d . $\overline{\text{Earning per Order}}$ is the average earning from each trip. Due to the design of the platform, drivers do not know the destination of the trip before they pick up the rider. Therefore, it is very difficult for drivers to selectively choose orders. We therefore assume drivers have the same expectation regarding how much they can earn from each trip. See the discussion in the previous section how we construct these two variables. Subsidy_{dp} is the actual subsidy provided by the platform in time period p on day d . It is important to note that during the non-experimental period from January 1 to July 30, 2022, no subsidies were provided, thus this variable is set to zero for that timeframe. Finally, ξ_{idp}^e is a stochastic term that captures demand fluctuations known to the driver beforehand but not observable to researchers. This term accounts for any additional, unanticipated variability in driver earnings. Importantly, ξ_{idp}^e can be correlated across different time periods within a day.

Driver i 's gain from the expected income on day d can thus be specified as:

$$\begin{aligned} \text{Gain}_{id}(\mathbf{h}_{id}) &= \sum_{p=1}^3 \alpha * E(\text{Earning Per Hour}_{idp}) * h_{idp} \\ &= \sum_{p=1}^3 \alpha * (\overline{\# \text{ Order Per Hour}_{iw_{dp}}} * (\overline{\text{Earning Per Order}} + \text{Subsidy}_{dp}) + \xi_{idp}^e) * h_{idp}. \quad (7) \end{aligned}$$

4.3. The Cost Function

We model drivers' costs by considering several factors. First, we assume an individual-specific marginal cost of work, denoted as $c_{1iw_{dp}}$. This captures the "opportunity cost" of using time for other activities, such as working a full-time job. Note that the overall utility function is parametric, including both the gain and cost functions. However, we model $c_{1iw_{dp}}$ in a non-parametric way to capture the full heterogeneity of work behaviors across drivers. Specifically, we estimate the cost for each driver that varies across weekdays and weekends, w_d , and across time periods, p . Specifying such granular parameters for costs is necessary because we have shown that drivers differ from each other in terms of not only their overall tendency to work but also when and how long they choose to work across time periods. As a result, driver-specific and time-variant marginal costs allow us to rationalize such rich heterogeneity. This is also consistent with the reality of ride-sharing platforms where drivers usually take orders as part-time work. Therefore, their preference for working is highly dependent on daily work and life schedules that vary substantially across individuals.

To account for unobserved and short-term factors affecting this cost, we further introduce a stochastic term, ξ_{iwdp}^c , which is known to the driver but not to researchers. This term reflects individual, day, and time period-specific cost shocks. For example, the cost of driving may be higher in good weather because the driver might prefer to spend time with family or friends. This component can also be correlated across different time periods within a day. In short, the driver's individual-specific marginal cost is represented by:

$$(c_{1iwdp} + \xi_{iwdp}^c) * h_{idp}.$$

The cost function further allows for psychological and physiological factors that may affect the work cost, including psychological flow, cost benefits from continuous work, and fatigue. These factors help us capture the interdependency between different time periods. Psychological flow (Norsworthy 2021) indicates that drivers are more likely to continue working if they are already in the mood. To capture this behavior, we calculate the cumulative working hours $\sum_{p' < p} h_{idp'}$ and use c_2 to incorporate its impact on subsequent costs within the same day. Therefore, this component is captured by:

$$c_2 * \sum_{p' < p} h_{idp'} * h_{idp}.$$

Additionally, we use c_3 to denote the cost benefits from continuous work, where drivers who have already been active in a previous period might save on commuting costs, encouraging more efficient work distribution and reducing downtime. That is, if driver i works in the previous time period, i.e. $I\{h_{idp-1} > 0\} = 1$, the cost of working will possibly be lower in the current period than the case when the driver does not work in the previous period. This is represented by:

$$c_3 * 1\{h_{idp-1} > 0\} * 1\{h_{idp} > 0\}.$$

That is, if the driver has worked in the previous period and continue to work in the current period, his/her cost will drop (assuming c_3 is negative).

Finally, we use c_4 to capture the fatigue of the driver. This means that drivers will incur a penalty if their total working hours for the day, i.e. the sum of h_{id1} , h_{id2} , h_{id3} , is too high. We assume that this penalty is convex, meaning that the marginal cost increases as working hours accumulate. Specifically, we include the following components:

$$\frac{c_4}{2} * (h_{id1} + h_{id2} + h_{id3})^2.$$

Therefore, the total costs for the driver i on day d can be expressed as:

$$Cost_{id}(\mathbf{h}_{id}) = \sum_{p=1}^3 \left((c_{1iwdp} + \xi_{iwdp}^c + c_2 * \sum_{p' < p} h_{idp'}) * h_{idp} + c_3 * 1\{h_{idp-1} > 0\} * 1\{h_{idp} > 0\} \right) + \frac{c_4}{2} * (h_{id1} + h_{id2} + h_{id3})^2. \quad (8)$$

4.4. Total Utility

Combining drivers' gain and costs, the utility for driver i on day d can be summarized as follows:

$$u(\mathbf{h}_{id}) = \sum_{p=1}^3 \left[\left(\alpha * \overline{\# Order Per Hour}_{iwdp} * \left(\overline{Earning Per Order} + \overline{Subsidy}_{dp} \right) - c_{1iwdp} - c_2 * \sum_{p' < p} h_{idp'} + \alpha * \xi_{idp}^e - \xi_{idp}^c \right) * h_{idp} - c_3 * 1\{h_{idp-1} > 0\} * 1\{h_{idp} > 0\} \right] - \frac{c_4}{2} * (h_{iwd1} + h_{iwd2} + h_{iwd3})^2. \quad (9)$$

We redefine $\xi_{idp} = \alpha * \xi_{idp}^e - \xi_{idp}^c$ to consider both unobserved demand and cost shocks that may

affect specific to each driver. We assume $\boldsymbol{\xi}_{id} = \begin{pmatrix} \xi_{id1} \\ \xi_{id2} \\ \xi_{id3} \end{pmatrix}$ follows a multivariate normal distribution

with the CDF as $F(0, \boldsymbol{\Sigma}_{\xi})$. $\boldsymbol{\Sigma}_{\xi}$ is the 3*3 variance-covariance matrix that is specified as:

$$\boldsymbol{\Sigma}_{\xi} = \begin{bmatrix} \sigma_{11}^2 & s_{21} & s_{31} \\ s_{21} & \sigma_{22}^2 & s_{32} \\ s_{31} & s_{32} & \sigma_{33}^2 \end{bmatrix}. \quad (10)$$

To further “smooth” the probability distribution across different combinations of work hours, \mathbf{h}_{id} , we introduce $\varepsilon_{h_{id}}$ to the utility function:

$$U(\mathbf{h}_{id}) = u(\mathbf{h}_{id}) + \varepsilon_{h_{id}}, \quad (11)$$

where $\varepsilon_{h_{id}}$ is independent across individuals and work hour choices $\mathbf{h}_{id} = (h_{id1}, h_{id2}, h_{id3})$ (see details below) and it follows a Type-1 extreme value distribution.

4.5. Likelihood Function

Given the utility specification and the distribution for $\varepsilon_{h_{id}}$, we can write the probability for drivers' supply choice \mathbf{h}_{id}^* as:

$$P_{id}(\mathbf{h}_{id}^*) = \int \frac{\exp(u(\mathbf{h}_{id}^*; \boldsymbol{\xi}_{id}))}{\sum_{\mathbf{h}_{id}} \exp(u(\mathbf{h}_{id}; \boldsymbol{\xi}_{id}))} f_{\xi}(\boldsymbol{\xi}_{id}) d\boldsymbol{\xi}_{id}, \quad (12)$$

where $f_{\xi}(\boldsymbol{\xi}_{id})$ is the PDF for $\boldsymbol{\xi}_{id}$. Note that \mathbf{h}_{id}^* is the observed work hour choice of driver i on day d , which is a vector of three time periods. This indicates that each driver faces a total of 343 potential work choice combinations daily, that is $(0, 1, 2, \dots, 6) \times (0, 1, 2, \dots, 6) \times (0, 1, 2, \dots, 6)$. The probability function is independent across days within the same person after

controlling for the individual marginal cost c_{1iwdp} , based on the assumption that ξ 's and ε 's are independent across days. We can then write the likelihood as follows:

$$l = \prod_i \prod_d P_{id}(\mathbf{h}_{id}^*). \quad (13)$$

5. Model Estimate and Identification

5.1. Simulated Maximum Likelihood Estimation

Given that we do not have the closed-form expression for the choice probability, we opt for a simulated likelihood method. To simulate ξ_{id}^s , where $s = 1, 2, \dots, NS$, we use the Cholesky decomposition method by first decomposing the variance-covariance matrix Σ_ξ . That is,

$$\Sigma_\xi = \begin{bmatrix} \sigma_{11}^2 & s_{21} & s_{31} \\ s_{21} & \sigma_{22}^2 & s_{32} \\ s_{31} & s_{32} & \sigma_{33}^2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ S21 & S22 & 0 \\ S31 & S32 & S33 \end{bmatrix} \begin{bmatrix} 1 & S21 & S31 \\ 0 & S22 & S32 \\ 0 & 0 & S33 \end{bmatrix}. \quad (14)$$

where parameters $\{S21, S22, S31, S32, S33\}$ are what we actually estimate in the model estimation. We then draw NS samples of \mathbf{v}_{id}^s from the standard normal distribution, and multiply \mathbf{v}_{id}^s with the lower-triangular matrix on the right-side of equation (14), given $\{S21, S22, S31, S32, S33\}$. This gives us ξ_{id}^s . Therefore, the probability can be rewritten as:

$$P_{id}(\mathbf{h}_{id}^*) = \frac{1}{NS} \sum_{ns=1}^{NS} \frac{\exp(u(\mathbf{h}_{id}^*; \xi_{id}^{ns}))}{\sum_{\mathbf{h}_{id}} \exp(u(\mathbf{h}_{id}; \xi_{id}^{ns}))}. \quad (15)$$

Let $\Theta_1 = \{\alpha, c_2, c_3, c_4, \Sigma_\xi\}$ consist of all the common shared parameters and $\Theta_2 = \{c_{1iwdp}; i = 1, \dots, N\}$ represent each driver's time period-specific costs. We can rewrite our log-likelihood function as:

$$\max_{\Theta_1, \Theta_2} ll = \max_{\Theta_1, \Theta_2} \sum_i \sum_d \ln \left(\frac{1}{NS} \sum_{ns} \frac{\exp(u(\mathbf{h}_{id}^*; \xi_{id}^{ns}, \Theta_1, \Theta_2))}{\sum_{\mathbf{h}_{id} \in H} \exp(u(\mathbf{h}_{id}; \xi_{id}^{ns}, \Theta_1, \Theta_2))} \right). \quad (16)$$

Although using the simulated likelihood method allows us to approximate the probability of each driver's work hour choices by integrating over the distribution of unobserved heterogeneity, the model's large number of parameters poses a computation challenge. The dimension of Θ_1 is 16 in total, whereas the dimension of Θ_2 is $N \times 6$ with N as the number of drivers, which is far larger than the dimension of Θ_1 . Note that, Θ_2 is associated with the number of drivers, which increases linearly with the number of drivers; however, Θ_1 does not. This characteristic makes the problem non-scalable. For instance, if we have a million drivers in a certain city, we will need to estimate over six million parameters in total. Thus, when estimating the model, we face the critical

challenge with the simulated maximum likelihood procedure. That is, when the dimensionality of Θ_2 is large, the estimation time will be too long because we are searching over a very large space to optimize the objective function.

To address these challenges, we propose a nested iteration procedure to manage the large dimensionality in the model estimation. This novel method is based on the observation that equation (16) can be rewritten as:

$$\max_{\Theta_1, \Theta_2} ll = \max_{\Theta_1} \sum_i \max_{\Theta_{2i}} \sum_d \ln \left(\frac{1}{NS} \sum_s \frac{\exp(u(\mathbf{h}_{id}^*; \xi_{id}^s, \Theta_{2i} | \Theta_1))}{\sum_{\mathbf{h}_{id} \in H} \exp(u(\mathbf{h}_{id}; \xi_{id}^s, \Theta_{2i} | \Theta_1))} \right). \quad (17)$$

where $\Theta_{2i} = \{c_{1iwdp}\}$ is the cost estimates for driver i . That is, conditional on Θ_1 , Θ_{2i} can be obtained by maximizing the likelihood of the observed work hour choices of the driver across days. Note that the likelihood can be rewritten in this way because the likelihood functions across drivers are independent from each other given the shared parameters Θ_1 .⁴

The nested iteration procedure can work because of equation (17). Details of the procedure are as follows: (1) At the outside of the estimation, we simulate \mathbf{v}_{id}^s from standard normal distribution for every driver-day. (2) Given the trial of Θ_1^* , we compute ξ_{id}^s . We then obtain the estimate $c_{1i, \widehat{\text{Weekday}, p}}$ for driver i 's weekday period-specific cost by

$$\max_{c_{1i, \text{Weekday}, p}} l_{1i} = \sum_{d \in \{\text{weekday}\}} \frac{1}{NS} \sum_s \frac{\exp(u(\mathbf{h}_{id}^*; \xi_{id}^s, c_{1i, \text{Weekday}, p} | \Theta_1^*))}{\sum_{\mathbf{h}_{id}} \exp(u(\mathbf{h}_{id}; \xi_{id}^s, c_{1i, \text{Weekday}, p} | \Theta_1^*))}. \quad (18)$$

We can similarly obtain the likelihood function for the weekend $l_{1,i}$ and calculate $c_{1i, \widehat{\text{Weekend}, p}}$. We repeat this process for all drivers. (3) Given the estimated $\widehat{\Theta}_2$, we can evaluate the full likelihood function:

$$l = \sum_i \sum_d \ln \left(\frac{1}{NS} \sum_s \frac{\exp(u(\mathbf{h}_{id}^*; \xi_{id}^s, \Theta_1^*, \widehat{\Theta}_2))}{\sum_{\mathbf{h}_{id}} \exp(u(\mathbf{h}_{id}; \xi_{id}^s, \Theta_1^*, \widehat{\Theta}_2))} \right). \quad (19)$$

Steps (2) and (3) describe the inner loop of the nested iteration procedure. Finally, (4) at the outer loop, we search over Θ_1 to maximize the log-likelihood. We repeat steps (2) and (3) each time Θ_1 is updated at the outer loop. We can visualize our estimation steps as below.

⁴ If Θ_{2i} affects the likelihood function of another driver i' , the proposed estimation method will not work. This can happen when, for example, there are spillover or peer effects in drivers' working choices. The choice of driver i' can be a function of the choice of i in this case.

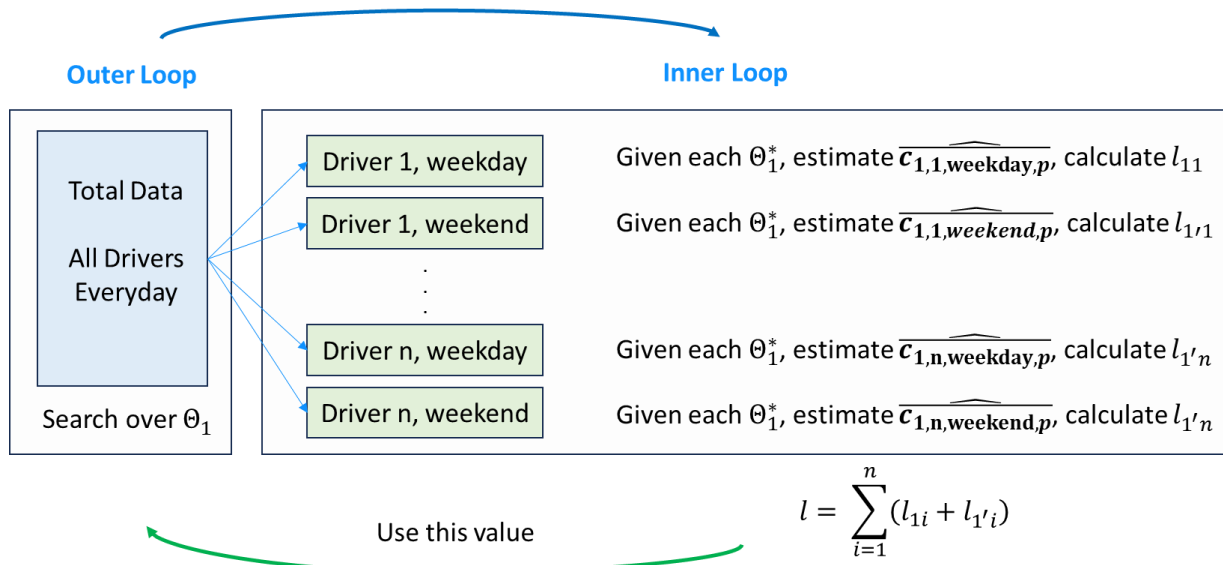


Figure 7. Nested Iteration Estimation

Compared with the joint estimation of Θ_1 and Θ_2 , the key advantage of the proposed method is that we can separately estimate Θ_2 for each individual driver. Note that only the dimensionality of Θ_2 increases with the number of drivers, and the estimation of Θ_2 in the inner loop is independent for each driver. This suggests that the computation burden of this procedure only increases linearly as the number of drivers grows. This will significantly reduce the time of model estimation. We can even decompose the estimation further by separating Θ_2 for weekdays and weekends within a driver. This is because while drivers' utility is correlated across time periods within a day, it is uncorrelated across days. Since we specify Θ_2 for weekday $c_{1i,Weekday,p}$ and weekend $c_{1i,Weekend,p}$, we can use the corresponding observations to separately estimate these parameters.

For large platforms such as Uber or Lyft that hire millions of drivers, researchers have to estimate tens of millions of individual-level Θ_2 . Even though the computation burden in the model estimation is linear in the number of drivers, estimating so many cost parameters is still challenging. However, since the estimation of Θ_2 in the inner loop is independent for each driver, platforms can employ large-scale parallel programming to allocate the computation of Θ_{2i} for each driver. The model estimation will reduce proportionately with the number of processors platforms can allocate to the problem. Since the nested iteration procedure can continue to function well with larger parameter space, the model estimation problem is *scalable*.

Our proposed method shares a similar nested iteration structure with Chan et al. (2014) and Chan et al. (2014). The key difference is that, while the referenced papers can recover Θ_2 conditional on

Θ_1 in a closed form, the inner loop of our method does not have this property. Therefore, while they can use least squares regressions in the inner loop, our estimation has to rely on maximum likelihood. The advantage of our proposed method is that the procedure can be applied to general structural models that are typically non-linear (even conditional on Θ_1).

5.2. Identification

Our model includes drivers' sensitivity to income α , cost parameters $\{c_{1iwdp}, c_2, c_3, c_4\}$, and the variance-covariance matrix for the unobserved factors Σ_ξ within a day. We use the experimental data to pin down α from c_{1iwdp} . This is because, during the experiment, we exogenously manipulate the level of subsidies across time periods. As c_{1iwdp} remains constant, change of driver supply can inform us of the value of α . Given the identification of α , c_{1iwdp} can also be identified from a driver's work choices across different time periods outside the experimental period.

Note that, if the outcome is only discrete choice (e.g. work or not work in a time period), α and c_{1iwdp} cannot be both identified, as one of the parameters has to be normalized. However, although the likelihood function in our model follows a multinomial logit structure, the outcome variables are the number of hours (from 0 to 6) in a time period. Therefore, it should be treated as a multiple-unit choice model, making both parameters identifiable.

The identification of $\{c_2, c_3, c_4\}$ is related to the general working behavior across all drivers. The identification of c_2 comes from the relationship between previous cumulative working hours and subsequent working hours. If drivers work more when they have worked more in prior time periods, we will have a negative estimate for c_2 . That is, the momentum of working reduces the cost of continuing to work. Note that we have controlled the drivers' average tendency to work using c_{1iwdp} . Similarly, the identification of c_3 comes from the relationship between whether drivers work in the adjacent time period and the subsequent working hour. Finally, the identification of c_4 comes from the average tendency to choose specific working hours for the whole day. In our model, drivers can work for a minimum of 0 hours and a maximum of 18 hours. If we observe that drivers are more likely to work shorter hours than longer hours, we will obtain a positive estimate for c_4 .

We capture the correlation of working hours across time periods within a day using c_2 and c_3 . However, these two components are unlikely to control for all factors. After controlling for c_2 and c_3 , any remaining correlations will be reflected in the variance-covariance matrix Σ_ξ . For instance, suppose in the data we observe that if a driver drives more in the morning on a given day, they are also more likely to drive in the afternoon; conversely, if they drive less in the morning, they are also less likely to work in the afternoon. Then, the covariance of the two time periods will be positive.

5.3. Numerical Experiments

To understand the performance of the proposed nested iteration method compared with the traditional joint estimation approach, we conduct simulation studies to compare the two methods. In these simulations, we assessed the performance and accuracy of our approach by simulating data for 20 drivers over 24 weeks. For each driver-day, we drew 100 random errors ξ_{id}^s from the multivariate normal distribution (NS = 100). We then estimated the results using various sets of starting values, including randomly generated values (ranging from 0 to 1) and fixed values of 1, 2, and 3. We use two estimation methods: the proposed nested iteration method and traditional joint estimation. We conduct the estimation on the desktop with 64 GB RAM and a 12th Gen Intel(R) Core(TM) i9-12900KF 3.19 GHz processor. We estimate a model in MATLAB 2022a. One advantage of the proposed nested iteration method is that we can use parallel programming for the inner loop, as the estimation for each driver is separable and independent conditional on the outer loop parameters. To leverage this advantage, we use 10 parallel CPU cores to estimate the model in the inner loop.

We first show the recovery accuracy for the homogenous parameters $\{\alpha, c_2, c_3, c_4\}$ in Table 4 by trying three different sets of model parameters. The results suggest that the nested iteration method consistently recovers the different true parameters with high accuracy. We further show the recovery of the individual parameters $c_{1i, \text{weekday}/\text{end}, p}$ in Figure 8, using one trial as an example. Consistent with the results of homogenous parameters, we find that the proposed method can also identify the individual-level parameters accurately. Taken together, these results suggest that the nested iteration method can achieve high accuracy in parameter recovery, both for homogeneous and individual-specific parameters.

Table 4. Recovery Test

Parameter	Trial 1		Trial 2		Trial 3	
	True Value	Recovered Value	True Value	Recovered Value	True Value	Recovered Value
Alpha	0.87	0.87	0.38	0.38	0.11	0.11
C2	-0.46	-0.43	-0.96	-1.04	0.70	0.72
C3	0.19	0.18	-0.72	-0.75	0.64	0.69
C4	0.20	0.21	0.41	0.42	0.40	0.39
Log-likelihood Function Value	-12332		-9355		-6352	



Figure 8. Recovery of Individual Cost Parameters

Next, we show that the proposed method can achieve faster estimation speeds than the traditional joint estimation method. Specifically, we compared the estimation times for different numbers of simulated drivers. Using the same set of parameter values, we simulated scenarios with 10, 20, 30, 40, and 50 drivers and compared the computational time required for each method. As shown in Table 5, while the traditional joint estimation method provided faster estimation times with fewer simulated drivers (i.e., 10 drivers), the computational efficiency of our proposed nested iteration method became apparent as the number of drivers increased. For example, with 50 drivers, the running time was significantly reduced from 18.05 hours with the traditional method to 9.87 hours with our proposed method. These findings illustrate that our nested iteration method not only maintains accuracy but also offers substantial computational advantages, especially as the dataset size grows.

Note that this exercise involves only 50 drivers at the maximum. For most of the ride-sharing platforms, the number will be tens or hundreds of thousands larger. The advantage of the proposed method will be much bigger in those cases. Furthermore, we restrict the number of CPU cores to ten for all studies. Big platforms can employ more processors in the parallel programming, which can vastly reduce the estimation time.

Table 5. The Estimation Time

N	Time (Hours)	
	Nested Iteration	Joint
10	2.04	0.57
20	2.75	3.00
30	7.62	7.56
40	6.41	9.34
50	9.87	18.05

6. Main Results

6.1. Estimation Results

Based on our dataset, we estimate the structural model using our proposed method. For the empirical application, since our panel covers a relatively long period, we include time-fixed effects in the model. Specifically, we adapt the original linear cost part $(c_{1iw_{dp}} + c_2 * \sum_{p' < p} h_{idp'}) * h_{idp}$ to $(c_{1iw_{dp}} + c_2 * \sum_{p' < p} h_{idp'} + Month_d) * h_{idp}$ where $Month_d$ denotes in which month day d is. We use the month of the field experiment as the baseline. The results are shown in Table 6.

First, we find that the income sensitivity parameter α is significantly positive, suggesting that drivers are more likely to work and work longer when their expected income increases. This is consistent with our descriptive evidence. However, the small value also suggests that drivers are not very income-sensitive, with an average income elasticity of only 0.008, calculated based on our data and the model estimates. Therefore, providing a uniform subsidy to every driver may be inefficient, as we will demonstrate in the counterfactual analysis.

Second, interestingly, we find that the estimate for c_2 is negative, suggesting that when drivers' previous cumulative working hours are longer, their working costs are smaller. This aligns with the psychological concept of flow (Norsworthy 2021), where individuals become more absorbed and engaged in an activity, making them more likely to continue it. For instance, a driver who has been driving for several hours may find it easier and more rewarding to keep driving rather than stop and start again later. This continuous engagement can lead to a more seamless and enjoyable work experience, reducing the perceived effort required to maintain the activity.

Table 6. Structural Model Estimates

	Coefficient	SE
α	0.06***	0.009
c_2	-0.45***	0.021
c_3	-0.07***	0.007
c_4	0.13***	0.006
S21	1.27***	0.020
S22	0.53***	0.031
S31	0.78***	0.033
S32	0.73***	0.036
S33	-0.01	0.023
Jan 2022	0.06**	0.038
Feb 2022	0.02	0.033
Mar 2022	-0.26***	0.033
Apr 2022	-0.32***	0.033
May 2022	-0.28***	0.035
June 2022	-0.24***	0.033
July 2022	-0.20***	0.033

Similarly, the estimate for c_3 is also negative. This suggests that when drivers work in the prior adjacent time period, they are more likely to work in the subsequent time period. Since we operationalize this prior working state as a binary variable, it captures the starting cost. For instance, when drivers start working on the platform, they may need to physically set up their cars to take riders and adjust their phones to interact with the platform's dispatching system. All these factors contribute to an initial starting cost, which can be avoided if drivers continue their work. Note that the difference between c_2 and c_3 is that c_2 captures the cumulative working hours in all prior time periods, whereas c_3 only reflects the working status in the closest time periods. It is likely that

drivers may have the same cumulative working hours but different working statuses in prior time periods.

In contrast, we find the estimate for c_4 , which measures the fatigue of the drivers, to be positive. A key feature in c_4 is that it suggests that drivers' costs are convex in working hours such that the marginal cost is increasing. It is important to note that while c_2 suggests a momentum effect such that higher cumulative working hours result in lower working costs, this momentum effect is linear. In comparison, c_4 is quadratic in total working hours. Taken together, this suggests that when the total working hours are not large, the momentum effect may be influential in shaping drivers' supply decisions. However, when the total working hours are large, the convex costs may play a more dominant role.

Apart from the costs, we also estimate the variance-covariance matrix for the unobserved factors, using Cholesky decomposition with $\{S21, S22, S31, S32, S33\}$. When we calculate the correlation matrix based on this, we get the following results:

$$\Sigma_{\xi} = \begin{bmatrix} 1 & 0.92 & 0.73 \\ 0.92 & 1 & 0.94 \\ 0.73 & 0.94 & 1 \end{bmatrix}$$

These results align with our expectations that periods closer together have higher correlations and *vice versa* for periods further apart. For example, the correlation between morning and afternoon is high at 0.92, whereas the correlation between morning and night is lower at 0.73. Still, the high correlations between time slots suggests that there might be some unobserved factors that influence a driver's behavior throughout the day. For instance, if a driver works late with their regular job the previous night, they may have a very high cost of working throughout the day. This correlation pattern helps us understand how such unobserved factors can have a lingering effect throughout the day, influencing the driver's willingness and ability to work.

Finally, the month fixed effects also align with our expectations. We set the experiment month as the baseline. During the winter months (January and February), drivers tend to have higher costs, likely due to increased difficulties in cold weather. As the weather starts to warm up, the costs for drivers decrease, as seen in the lower costs from March to July.

6.2. Understanding Individual-Level Costs

An important feature of our model is that we estimate individual driver-level costs for different time periods. It is important to note that some drivers have never worked during certain time periods. For these drivers, we lack data to estimate their costs directly. Consequently, we estimated their costs to be the highest among all drivers who worked during that time period. This serves as the lower bound for their costs.

We first show the distribution of these costs for each time period in Figure 9, with the costs of drivers who have never worked during the time slot highlighted in red.

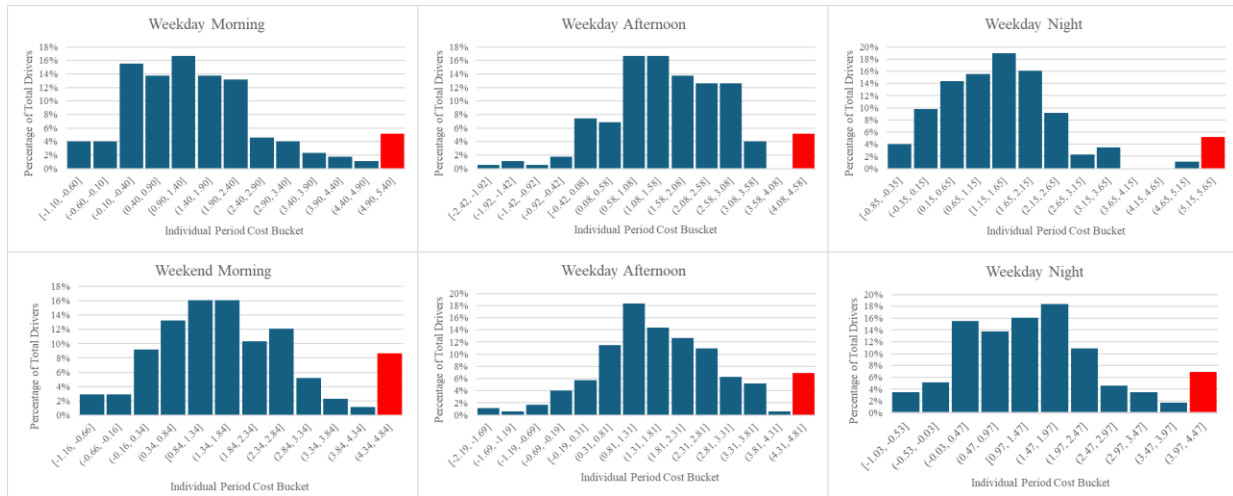


Figure 9. Estimated Individual Cost Distribution

The cost distribution for all drivers in each time period generally follows a normal distribution shape. For a very small number of drivers with negative costs, this implies that working during that period provides them with additional utility beyond their earnings, motivating them to work regardless of the financial incentive. This could be due to various factors such as personal satisfaction, social interactions, or other non-monetary benefits they derive from working. However, the negative costs do not mean that drivers work as much as they can. This is because drivers’ total costs also contain a quadratic term, which prevents them from working very long hours.

Apart from the cost heterogeneity across drivers, we are able to examine the cost correlation within each driver because we have six cost estimates for each driver (i.e., weekday morning/afternoon/night and weekend morning/afternoon/night). Such correlation patterns are informative of the temporal dynamics of drivers’ costs. We show the correlation among these six dimensions in Table 7.

The correlation patterns show the high face validity of our cost estimates. We generally observe a higher correlation between time periods that are adjacent. For instance, within weekdays or weekends, the cost correlation between mornings and afternoons is higher than between mornings and nights. Across weekdays and weekends, the cost correlation between mornings and mornings is higher than between mornings and afternoons. This is likely due to the fact that drivers may be more likely to have similar offline arrangements in adjacent time periods than in distant time periods. For instance, if they have a regular job in the morning, they are more likely to work the

same job in the afternoon than at night. Similarly, if drivers have an exercise routine on weekday mornings, they may be more likely to have a similar routine on weekend mornings than on weekend afternoons. Such behavior consistency outside the ride-sharing platforms may result in the correlation patterns we observe.

Table 7. Individual Cost Correlation

	Weekday Morning	Weekday Afternoon	Weekday Night	Weekend Morning	Weekend Afternoon	Weekend Night
Weekday Morning	1					
Weekday Afternoon	0.80	1				
Weekday Night	0.70	0.74	1			
Weekend Morning	0.67	0.59	0.53	1		
Weekend Afternoon	0.52	0.69	0.54	0.83	1	
Weekend Night	0.45	0.52	0.61	0.79	0.84	1

Given the high degree of cost differences between drivers and correlation within drivers, it is possible for us to classify drivers into different categories based on distinctive patterns. To understand this issue, we conducted a cluster analysis to determine whether drivers exhibit similar behaviors. In this analysis, we compared clustering results for group sizes ranging from 1 to 7 and found that the 6-group model provided the most accurate segmentation. The average costs for drivers in different time periods are shown in Table 8, which represents the weekday cost clustering result. We use this as an example, as the weekend distribution is very similar.

The table reveals distinct patterns among the clusters. Group 6 consists of 5% of drivers who never worked during weekdays, exhibiting the highest costs across all time periods. These drivers are less willing to work unless they receive higher compensation. In contrast, Group 5 includes the most active drivers with the lowest costs, suggesting a greater willingness to work under current compensation rates.

Group 1 is characterized by late workers, who experience the lowest costs at night compared to the morning and afternoon, implying their reservation wage is met more easily in the evening. Conversely, Group 4 comprises morning workers, showing the lowest costs in the morning relative to other periods. Groups 2 and 3 display similar cost patterns throughout the day, indicating indifference to the time of day. However, Group 3 is more active than Group 2, as evidenced by their lower costs.

Table 8. Individual Cost Clustering (Weekday)

Cluster	Morning	Afternoon	Night	Percentage of Drives
1	3.62	2.69	2.06	11%
2	1.77	2.16	1.80	32%
3	0.75	1.09	0.80	33%
4	1.42	1.99	3.96	4%
5	-0.26	-0.31	-0.02	15%
6	5.20	4.26	5.47	5%

7. Counterfactual Analysis

Similar to the flexibility on the supply side of ride-sharing platforms, the demand side is also characterized by significant fluctuations. When facing a sudden increase in rider demand, the platform faces the challenge of motivating drivers to work. When demand increases without a corresponding supply boost, high cancellation rates become a concern for the platform. High cancellation rates cause short-term revenue loss and lead to rider churn due to long waiting times and unsatisfactory experiences. Therefore, incentivizing drivers' supply to meet riders' demand is critical for platform growth and sustainability.

One popular strategy for ride-sharing platforms is to provide subsidies to motivate drivers' supply. For instance, Lyft provides a healthcare subsidy for drivers in California with specific working hours.⁵ For drivers who average between 15 and 25 hours per week of booked time, Lyft will make 50% of the average Affordable Care Act (ACA) contribution for the driver. If drivers average above 25 hours, Lyft's contribution increases to 100%. As another example, Uber offers its drivers the Quest program based on how much a driver can make. When drivers complete a certain amount of earnings within a period of time, Uber offers monetary rewards to the driver.⁶ Note that none of these subsidies are individually customized, as they are offered to all drivers. Furthermore, these subsidies can incur significant financial costs for platforms, while the effectiveness of stimulating drivers to work more is questionable.

⁵ <https://help.lyft.com/hc/en-us/all/articles/360061623553-Healthcare-subsidy>

⁶ <https://therideshareguy.com/the-best-strategies-and-hacks-for-uber-quest/>,
<https://www.nytimes.com/interactive/2017/04/02/technology/uber-drivers-psychological-tricks.html>

To investigate whether the focal platform can offer an effective subsidization plan that can incentivize drivers' work supply in a cost-effective manner, we consider a hypothetical scenario where rider demand in different time periods throughout a day exogenously increases by 5%. We first create a base case to see how drivers' working hours and orders taken would change regarding the demand increase without subsidies. We then consider two counterfactual policies—*time-based non-targeting* subsidization and *individual-based targeting* subsidization—and compare their performance and cost to the platform.

7.1. Base Case: No Subsidy Offers

Driver supply is influenced by demand, as higher demand naturally increases drivers' actual income. This increase in actual income will also be reflected in drivers' expected income, leading to an increase in their working hours. However, the impact can vary among drivers. In extreme cases, if a driver is already fully occupied, the increase in demand will likely not affect their actual and expected income.

We consider two extreme cases in all counterfactual scenarios: the most ideal and worst case. For the most ideal scenario, we first assume that the change in drivers' expected income mirrors the change in their actual income, i.e., drivers have rational expectations. We then impose the restriction that increased demand will not affect drivers' hourly orders if their current number of orders is 1.2 or above. This assumption is based on the observation that, on average, it takes about 30 minutes for a driver to complete an order, from assignment to drop-off. Additionally, 1.2 orders per hour represent the top 1 percentile of hourly orders among drivers. Thus, we assume that drivers already operating at this level will not see an increase in their hourly orders, and their expected income will remain unchanged. The remaining drivers' hourly orders will be adjusted accordingly, capped at 1.2. Additionally, we assume that the actual hourly orders taken align with the expected hourly order change, which is increased by 5%.

In the worst-case scenario, we consider the possibility that drivers' expected income does not change despite the increase in demand. This scenario accounts for potential friction in matching riders and drivers, which may prevent the full conversion of increased demand into fulfilled orders. Therefore, the expected income of drivers remains unchanged even with a 5% increase in demand, and there is no change in the actual hourly number of orders taken. This approach allows us to evaluate the potential range of outcomes and better understand how different levels of demand realization impact driver behavior and platform performance.

Table 9 represents the benchmark scenario where no subsidy is offered, illustrating how supply would change in the most ideal case where the number of orders fully reflects the 5% increase in demand. We separate the analysis for weekdays and weekends. Even in this ideal scenario, the

increase in working time is very small, at just 0.035%. Consequently, the total number of orders drivers take cannot fulfill the increased demand, achieving only around a 4% increase. The demand fulfillment gap is largely because the 1% drivers do not incorporate the demand increase into their expectation and they account for a significant portion of orders on the platform (20%). In addition, this small increase in working time is due to drivers' low-income sensitivity. The results suggest that to fulfill the total demand increase without generating additional costs for the riders, the platform needs to offer subsidies. Otherwise, roughly 20% of the increased orders will remain unfulfilled, potentially leading to negative consequences, as previously discussed.⁷

In the worst case, if there is no increase in both the expected and actual hourly number of orders, the working time of all drivers will remain unchanged, and there will be a 100% revenue loss due to the lack of a supply of drivers. In reality, the revenue loss of the platform should be between 20% and 100%.

Table 9. Base Case: Order and Work Time Change

	Expected Hourly Order Increase by 5%	
	Order Increase	Work Time Increase
Weekday	4.15%	0.034%
Weekend	3.95%	0.037%

7.2. Policy 1: Time-Based Non-Targeting Subsidization

We examine the effects of the *time-based non-targeting* subsidization strategy. This strategy involves offering uniform subsidies to all drivers for every order completed in a time period and letting drivers decide whether to increase the working time during the period, which is similar to second-degree price discrimination. We assume that the subsidy is provided at integer levels, ranging from \$0 to \$9⁸. By design, this subsidy structure means the platform will lose some profit because it will also subsidize orders drivers would have taken without the subsidy. However, the benefit is that the subsidy is provided equally to all drivers, avoiding discrimination.

⁷ The increase in orders is larger than the increase in working hours because our model measures supply by the hour. The number of orders is calculated by multiplying the supply hours by the expected hourly orders. Therefore, even if the supply hours do not change, an increase in the expected hourly orders results in a change in the total number of orders taken.

⁸ The platform would already incur a loss when the subsidy exceeds \$6; therefore, we assume the subsidy is at most \$9. Note that \$6 is just a rough number, not the exact number, to avoid revealing sensitive information about the platform.

We first examine cases when the demand increase occurs in only one period of the day. We compare the results for two scenarios: the worst-case scenario, where 0% of the demand increase is reflected in drivers' actual and expected change of orders, and the most ideal case, where the demand increase is fully reflected in the order changes. We compute *the least-cost subsidy level* in each time period under the constraint that there is at least a 5% increase in the orders drivers will take, matching the 5% increase to match the demand increase. Table 10 reports the results.

The first observation is that the increase in orders taken and working time during the subsidized period negatively affects the rest of the periods, as seen in the "order taken change" and "working time change" sections of the table. This indicates that drivers spend more time working during the subsidized period, sacrificing their working time in other periods. We also noticed that the afternoon generally results in higher subsidy costs, around 16% and 8% for the most ideal and worst case scenarios, respectively, compared to the morning and night, which are around 10% and 5% for the most ideal and worst case scenarios. This is because there tend to be more orders in the afternoon. To incentivize drivers to work more during the afternoon, the platform would have to subsidize all the previous orders, resulting in higher costs.

Another interesting observation is that in the most ideal scenario, for three time periods—weekday morning, weekday night, and weekend morning—the platform does not need to provide any subsidies to the drivers. This is because demand has been low during these periods, and no drivers working during these time periods exceed 1.2 orders per hour, even with a 5% increase. That said, every driver's expected hourly order is fully increased by 5%. Therefore, even though there is no subsidy to incentivize the drivers, i.e., with no working time change, the orders taken still meet the 5% increase.

We then examine how the non-targeting subsidy would change driver performance if the increase in demand occurs throughout the entire day instead of just one period. Using the same criteria—ensuring the lowest subsidy that can fulfill the demand in each period by at least 5%—we observe that the subsidy generates a profit loss ranging from 14% to 47% compared to no subsidies (see Table 11). Additionally, the platform tends to offer the highest subsidy in the morning to balance the increased demand. This strategy is likely due to the need to incentivize drivers early in the day to ensure adequate supply, which then influences their availability and work patterns for the rest of the day.

7.3. Policy 2: Individual-Based Targeting Subsidization

We then examine the effects of *individual-based targeting* subsidies, i.e. offering individually customized subsidies to each driver during each time period, and their costs to the platform. To conduct individual-based targeting, we use linear programming to optimize driver-specific

subsidies. Linear programming is suitable for this purpose because the subsidies are designed at the individual level, while the total supply constraint is applied at the aggregate level. Linear programming allows us to effectively search for the optimal combination of individual subsidies while ensuring the aggregate constraint is met.

Let $Subsidy_{ip}$ be the subsidy provided to driver i at time period p , $\overline{\# Order Per Hour}_{ip}$ be the expected number of orders per hour after the adjustment, and $h_{ip}(Subsidy_{ip})$ be the working hours supplied by driver i at time period p , which depends on the subsidy offered. The exact dependency is based on the structural model we have estimated. We assume that the platform aims to maximize the revenue it receives from the orders. Therefore, we have the following constrained optimization problem:

$$\max_{Subsidy_{ip}} \sum_i \sum_p \overline{\# Order Per Hour}_{ip} * (Profit Per Order - Subsidy_{ip}) * h_{ip}(Subsidy_{ip}),$$

Subject to:

$$\sum_i \sum_p h_{ip}(Subsidy_{ip}) * \overline{\# Order Per Hour}_{ip} \geq \sum_i \sum_p h_{ip}(0) * \overline{\# Order Per Hour}_{ip} * 1.05, \quad (20)$$

where $h_{ip}(0)$ is drivers' supply hour without any subsidy. The optimization constraint is that the total supply hours must increase by at least 5%. Consistent with managerial practice on the platform, the optimization constraint is that the total supply hours must increase by at least 5%. We also assume that the subsidy is provided at integer levels, ranging from \$0 to \$9.

Table 12 presents the results showing how individual-based targeting subsidization can encourage drivers to work and meet a 5% demand increase during a certain period. Similar to the non-targeted case, in the most ideal scenario, weekday morning, weekday night, and weekend morning do not require any subsidy because the natural increase in orders is sufficient to cover the demand increase. However, for the remaining periods, in both the worst and most ideal cases, the overall subsidy cost is reduced by at least nearly 50% compared to the non-targeting case. This significant reduction suggests that the targeting method is highly effective in minimizing subsidy costs while still meeting the increased demand.

We also compare the performance of the targeted approach with non-targeted approaches when demand increases by 5% for each period throughout the entire day. As shown in Table 13, the subsidization cost in the targeted approach is 51% to 83% lower than in the non-targeted approach. Furthermore, the platform can achieve 79% to 98% of the profit level it would have without providing any subsidies.

To understand how the targeting approach distributes subsidies, we further analyze the breakdown of each subsidy level for every period. Using weekdays as an example (weekends exhibit very similar features), Tables 14 and 15 reveal that the targeting subsidy approach tends to assign subsidies to a relatively small portion of drivers, ranging from 25% to 2%, depending on the scenario and time period.

For the worst-case scenario with a 0% increase in expected orders, the targeting approach tends to provide more subsidies to drivers with the highest hourly orders and lower individual period costs. This strategy aims to maximize the immediate impact by encouraging already active and efficient drivers to work more. Conversely, in the most ideal scenario where there is a full 5% increase in expected orders, the algorithm targets drivers with lower hourly expected orders and higher individual period costs. This is because, in the most ideal scenario, the majority of the increased orders come from active drivers. Therefore, the algorithm focuses on incentivizing less active drivers to come online and work, thus balancing the supply more effectively across all drivers.

This differential approach ensures that the subsidies are utilized in the most cost-effective manner, optimizing the supply to meet the increased demand. By strategically targeting drivers based on their efficiency and activity levels, the platform can achieve a higher level of service with significantly lower subsidization costs.

Table 10. Policy 1- Time-Based Non-Targeting Subsidization: One Time Period Subsidy

		Subsidy			Order Taken Change			Working Time Change			Subsidy Cost Rate
		Morning	Afternoon	Night	Morning	Afternoon	Night	Morning	Afternoon	Night	
0% Expected Order Increases											
Weekday	Morning	3	0	0	5.4%	-0.4%	-0.5%	2.2%	-0.2%	-0.1%	10.5%
	Afternoon	0	2	0	-0.5%	5.0%	-0.4%	-0.2%	2.3%	-0.2%	16.9%
	Night	0	0	3	-0.6%	-0.4%	6.5%	-0.2%	2.6%	-0.6%	12.0%
Weekend	Morning	4	0	0	5.8%	-0.3%	-0.5%	2.6%	-0.1%	-0.2%	10.1%
	Afternoon	0	2	0	-0.8%	6.1%	-0.5%	-0.3%	2.6%	-0.2%	15.9%
	Night	0	0	2	-0.8%	-0.4%	6.1%	-0.3%	-0.2%	2.7%	11.1%
5% Expected Order Increases											
Weekday	Morning	0	0	0	5.1%	-0.1%	0.0%	0.0%	0.0%	0.1%	0.0%
	Afternoon	0	1	0	-0.3%	6.0%	-0.3%	-0.1%	1.2%	-0.1%	8.5%
	Night	0	0	0	0.0%	-0.1%	5.1%	0.0%	0.0%	0.0%	0.0%
Weekend	Morning	0	0	0	5.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	Afternoon	0	1	0	-0.4%	6.6%	-0.3%	-0.2%	1.4%	-0.1%	8.0%
	Night	0	0	1	-0.4%	-0.2%	7.6%	-0.2%	-0.1%	1.5%	5.6%

Table 11. Policy 1- Time-Based Non-Targeting Subsidization: Full Day Subsidy

Subsidy			Order Increase			Working Time increase			Subsidy Cost Rate	Profit Lose Comparing to No Subsidy	
Morning	Afternoon	Night	Morning	Afternoon	Night	Morning	Afternoon	Night			
0% Expected Order Increases											
Weekday	4	3	3	5.8%	6.7%	5.1%	2.4%	3.2%	2.1%	50.5%	-46.4%
Weekend	5	3	3	5.5%	8.0%	7.5%	2.5%	3.5%	3.4%	49.1%	-47.3%
5% Expected Order Increases											
Weekday	1	1	1	6.5%	5.7%	7.0%	0.6%	1.2%	0.8%	15.4%	-13.7%
Weekend	1	1	1	6.0%	6.3%	7.1%	0.4%	1.2%	1.3%	15.4%	-13.3%

Table 12. Policy 2-Individual-Based Targeting Subsidization: One Time Period Subsidy

		Subsidy			Order Taken Change			Working Time Change			Subsidy Cost Rate
		Morning	Afternoon	Night	Morning	Afternoon	Night	Morning	Afternoon	Night	
0% Expected Order Increases											
Weekday	Morning	2.4	0	0	5.0%	-0.3%	-0.5%	1.2%	-0.1%	-0.1%	5.4%
	Afternoon	0	2.3	0	-0.5%	5.0%	-0.5%	-0.1%	1.5%	-0.1%	10.0%
	Night	0	0	2.2	-0.4%	-0.2%	5.0%	-0.1%	-0.1%	1.1%	5.2%
Weekend	Morning	3.4	0	0	5.0%	-0.3%	-0.4%	1.4%	-0.1%	-0.1%	5.0%
	Afternoon	0	2.3	0	-0.5%	0.1%	-0.5%	-0.1%	1.2%	0.1%	7.9%
	Night	0	0	1.7	-0.5%	-0.3%	5.0%	-0.1%	-0.1%	1.3%	5.0%
5% Expected Order Increases											
Weekday	Morning	0	0	0	5.1%	-0.1%	0.0%	0.0%	0.0%	0.1%	0.0%
	Afternoon	0	4.0	0	-0.2%	5.0%	-0.2%	0.0%	0.4%	0.0%	2.6%
	Night	0	0	0	0.0%	-0.1%	5.1%	0.0%	0.0%	0.0%	0.0%
Weekend	Morning	0	0	0	5.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	Afternoon	0	2.5	0	-0.2%	5.0%	0.0%	0.0%	0.4%	0.0%	2.1%
	Night	0	0	1.0	-0.1%	0.0%	5.0%	-0.1%	0.0%	0.3%	0.5%

Table 13. Policy 2-Individual-Based Targeting Subsidization: Full Day Subsidy

Subsidy			Order Increase			Working Time increase			Subsidy Cost Rate	Profit Lose Comparing to No Subsidy	
Morning	Afternoon	Night	Morning	Afternoon	Night	Morning	Afternoon	Night			
0% Expected Order Increases											
Weekday	2.8	2.8	3.3	5.0%	5.0%	5.0%	1.3%	1.6%	1.1%	24.4%	-20.6%
Weekend	2.2	2.2	2.1	5.0%	5.0%	5.0%	1.6%	1.2%	1.2%	20.9%	-16.9%
5% Expected Order Increases											
Weekday	2.4	3.3	1.7	5.0%	5.0%	5.0%	0.0%	0.0%	0.0%	2.7%	-1.9%
Weekend	1.3	1.7	1.1	5.0%	5.0%	5.0%	0.0%	0.0%	0.0%	2.9%	-1.9%

Table 14. Policy 2-Subsidy Breakdown (Weekday): Full Day Subsidy with 0% Expected Order Increase

Subsidy	Morning					Afternoon					Night				
	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change
0	75%	0.05	1.2	-1.0%	0.3%	92%	0.15	1.6	-4.6%	-1.1%	89%	0.14	1.3	-8.2%	-7.9%
1	10%	0.07	2.1	2.0%	4.9%	2%	0.27	2.5	1.8%	2.3%	5%	0.02	3.3	0.0%	0.5%
2	5%	0.15	2.9	10.4%	11.1%	2%	0.96	0.7	16.4%	17.8%	1%	0.46	2.8	3.4%	4.6%
3	3%	0.23	1.8	9.0%	11.0%	2%	0.95	0.2	23.0%	29.2%	2%	0.37	3.0	16.1%	15.3%
4	2%	0.28	1.8	12.7%	14.9%	0%		NA			1%	0.00	5.5	0.0%	0.0%
5	2%	0.13	3.0	6.4%	9.1%	0%		NA			0%		NA		
6	2%	0.04	4.8	0.0%	0.0%	1%	1.29	0.1	24.4%	22.7%	1%	0.20	5.5	0.0%	0.0%
7	0%		NA			0%		NA			1%	0.01	5.5	0.0%	0.0%
8	0%		NA			0%		NA			1%	1.14	0.5	52.6%	52.0%
9	2%	0.44	1.8	60.5%	48.6%	1%	1.57	1.1	38.9%	29.2%	1%	1.24	0.3	36.1%	35.3%

Table 15. Policy 2-Subsidy Breakdown (Weekday): Full Day Subsidy with 5% Expected Order Increase

Subsidy	Morning					Afternoon					Night				
	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change	% Drivers	Avg Hourly Order	Avg Individual Period Cost	% to Total Order Change	% to Total Working Time Change
0	87%	0.08	1.2	94.8%	11.9%	98%	0.19	1.6	66.6%	12.7%	96%	0.16	1.4	97.9%	65.0%
1	7%	0.05	2.9	3.5%	53.6%	1%	0.07	3.2	0.0%	0.2%	2%	0.16	3.2	2.0%	28.6%
2	3%	0.05	3.8	1.4%	28.2%	1%	0.59	2.0	1.1%	4.9%	1%	0.20	5.5	0.0%	0.3%
3	1%	0.22	2.7	0.3%	5.6%	0%		NA			1%	0.08	2.7	0.1%	6.1%
4	0%		NA			0%		NA			0%		NA		
5	2%	0.04	4.6	0.0%	0.7%	0%		NA			0%		NA		
6	1%	0.03	3.9	0.0%	0.0%	0%		NA			0%		NA		
7	0%		NA			1%	1.57	1.08	32.2%	82.1%	0%		NA		
8	0%		NA			0%		NA			0%		NA		
9	1%	0.06	4.7	0.0%	-0.1%	0%		NA			0%		NA		

8. Conclusion

Ride-sharing platforms have revolutionized the matching of drivers and riders, greatly improving the efficiency of the taxi industry. While these platforms provide significant benefits to both drivers and riders, they also face the critical challenge of motivating driver supply and ensuring a satisfactory rider experience. This challenge arises from the fact that drivers have complete flexibility over their work schedules. Since the platform cannot compel drivers to be at specific locations and times, this flexibility can undermine the goal of seamlessly transporting passengers whenever and wherever needed.

To tackle this challenge, our study aims to understand the underlying factors that influence drivers' supply behavior and to shed light on how platforms can design effective subsidy strategies to motivate drivers and meet riders' demand. To achieve this goal, we built a structural model that accounts for the heterogeneous costs and income sensitivities of drivers. A unique feature of our model is that it allows the working cost to vary for each driver. This rich heterogeneity is valuable for rationalizing drivers' idiosyncratic working behavior, as different drivers choose to work at different times and for varying durations. While the model provides deep insights into drivers' working decisions, it also introduces significant computational burdens in estimation. To overcome this challenge, we developed a novel nested iteration estimation method. This method achieves comparable accuracy to traditional joint estimation approaches while offering significant computational advantages, especially with large-scale datasets involving thousands of parameters. This methodological advancement provides a powerful tool for analyzing complex behavioral models in extensive datasets.

Applying the model to a dataset from a ride-sharing platform in Canada, we find substantial heterogeneity in drivers' costs, which aligns with the observed variation in their working behavior. More importantly, we use the model estimates to conduct counterfactual analyses and develop optimal subsidy strategies to motivate driver effort. We consider a hypothetical scenario in which the platform faces a 5% increase in demand and needs to incentivize drivers to meet this demand. We evaluate three counterfactual subsidy strategies. In the first baseline case, no subsidy is offered, as drivers may increase their work expecting higher earnings due to the demand increase. In the second and third cases, we offer subsidies to drivers. In the second case (*time-based non-targeting*), the subsidy varies only by time periods, while in the third case (*individual-based targeting*), it varies by both time periods and individual drivers. This third strategy is a targeted subsidization approach based on drivers' heterogeneous working costs.

Our findings indicate that because drivers are not highly sensitive to subsidies, the platform incurs significant loss by offering *time-based non-targeting* subsidization relative to the baseline case.

However, by customizing the subsidy for each individual driver, the *individual-based targeting* subsidization can substantial costs compared with the *time-based non-targeting* subsidization and achieve the majority level of the profit level it would obtain without subsidies. This suggests that targeted subsidies can effectively incentivize drivers while saving costs for the platform.

Our paper makes significant contributions to both academic literature and managerial practices. Academically, we develop a structural model that incorporates heterogeneous working costs and rationalizes drivers' working choices. Additionally, we introduce a novel nested iteration estimation method that accommodates the high dimensionality of the parameter space, significantly improving estimation efficiency while maintaining high precision. On the managerial front, our empirical framework and counterfactual analysis provide actionable guidance for platforms to understand drivers' working behavior and design effective subsidization strategies. Our framework can be easily applied to similar ride-sharing platforms. The nested iteration estimation approach ensures that the computational burden remains manageable even as the number of drivers increases. We demonstrate that platforms can leverage estimates of individual drivers' costs to design targeted subsidy strategies that effectively induce driver effort while maintaining reasonable subsidization costs.

There are several limitations related to our study that warrant future research. First, while our structural model considers rich heterogeneity in drivers' supply-side behavior, we do not endogenize riders' demand relative to different levels of supply. We made this choice because our main goal is to understand the underlying factors driving drivers' working decisions. We construct drivers' belief of demand by using the observed demand data, which is consistent with drivers' belief formation and empirical approaches adopted by previous studies (Crawford and Meng, 2011; Chen et al., 2019; Thakral and Tô, 2021). However, future studies can further consider riders' demand and develop an integrated model to better understand the platform. Second, our model allows for dynamics within a day but not across days. Specifically, within a day, drivers' working behavior in prior time periods can influence their costs in subsequent time periods. Such dynamics influence drivers' working decision optimization. However, we do not consider cross-day dynamics where drivers' working behavior in one day could shift their costs in the following days. While considering such cross-day dynamics may bring the model closer to reality, it can significantly complicate the estimation. Third, the subsidy design in the counterfactual analysis may be specific to the platform we study. We consider the subsidy for morning, afternoon, and night given that this is the temporal segmentation implemented on the platform. However, other platforms may offer different subsidy plans, such as varying subsidies by hour. Therefore, when

applying our method to alternative platforms, researchers and managers should consider the specific practices and adapt accordingly.

References

- Attanasio, O., Levell, P., Low, H., & Sánchez-Marcos, V. (2018). Aggregating elasticities: intensive and extensive margins of women's labor supply. *Econometrica*, 86(6), 2049-2082.
- Bäro, A., Toepler, F., Meynhardt, T., & Velamuri, V. K. (2022). Participating in the sharing economy: The role of individual characteristics. *Managerial and Decision Economics*, 43(8), 3715-3735.
- Burtch, G., Carnahan, S., & Greenwood, B. N. (2018). Can you gig it? An empirical examination of the gig economy and entrepreneurial activity. *Management Science*, 64(12), 5497-5520.
- Camerer, C., Babcock, L., Loewenstein, G., & Thaler, R. (1997). Labor supply of New York City cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2), 407-441.
- Chan, T. Y., Li, J., & Pierce, L. (2014). Learning from peers: Knowledge transfer and sales force productivity growth. *Marketing Science*, 33(4), 463-484.
- Chan, T. Y., Li, J., & Pierce, L. (2014). Compensation and peer effects in competing sales teams. *Management Science*, 60(8), 1965-1984.
- Chen, M. K., Rossi, P. E., Chevalier, J. A., & Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of Political Economy*, 127(6), 2735-2794
- Chen, X., Li, Z., Ming, L., & Zhu, W. (2022). The incentive game under target effects in ridesharing: A structural econometric analysis. *Manufacturing & Service Operations Management*, 24(2), 972-992.
- Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3), 471-475.
- Crawford, V. P., & Meng, J. (2011). New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101(5), 1912-1932.
- Fan, Tijun, et al. "Output-oriented agricultural subsidy design." *Management Science* 70.3 (2024): 1448-1464.
- Farber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of political Economy*, 113(1), 46-82.
- Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *American Economic Review*, 98(3), 1069-1082.

- Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, 130(4), 1975-2026.
- Hinnosaar, M. (2023). The persistence of healthy behaviors in food purchasing. *Marketing Science*, 42(3), 521-537.
- Hoynes, H., Page, M., & Stevens, A. H. (2011). Can targeted transfers improve birth outcomes?: Evidence from the introduction of the WIC program. *Journal of Public Economics*, 95(7-8), 813-827.
- Huang, N., Burtch, G., Hong, Y., & Pavlou, P. A. (2020). Unemployment and worker participation in the gig economy: Evidence from an online labor market. *Information Systems Research*, 31(2), 431-448.
- Liu, J., Pei, S., & Zhang, X. (2023). Online Food Delivery Platforms and Female Labor Force Participation. *Information Systems Research*.
- Jackson, M. I. (2015). Early childhood WIC participation, cognitive development and academic achievement. *Social science & medicine*, 126, 145-153.
- Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, 49(4), 961-1075.
- Monte, F., Redding, S. J., & Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12), 3855-3890.
- Keane, M., & Rogerson, R. (2015). Reconciling micro and macro labor supply elasticities: A structural perspective. *Annu. Rev. Econ.*, 7(1), 89-117.
- Meyer, B. D. (2002). Labor supply at the extensive and intensive margins: The EITC, welfare, and hours worked. *American economic review*, 92(2), 373-379.
- Norsworthy, C., Jackson, B., & Dimmock, J. A. (2021). Advancing our understanding of psychological flow: A scoping review of conceptualizations, measurements, and applications. *Psychological bulletin*, 147(8), 806.
- Olsder, W., Martagan, T., & Tang, C. S. (2023). Improving access to rare disease treatments: Subsidy, pricing, and payment schemes. *Management Science*, 69(9), 5256-5274.
- Owen, A. L., & Owen, G. M. (1997). Twenty years of WIC: a review of some effects of the program. *Journal of the American Dietetic Association*, 97(7), 777-782.

Pabilonia, S. W., & Ward-Batts, J. (2007). The effect of child gender on parents' labor supply: An examination of natives, immigrants, and their children. *American Economic Review*, 97(2), 402-406.

Thakral, N., & Tô, L. T. (2021). Daily labor supply and adaptive reference points. *American Economic Review*, 111(8), 2417-2443.

Xiao, P., Xiao, R., Liang, Y., Chen, X., & Lu, W. (2020). The effects of a government's subsidy program: Accessibility beyond affordability. *Management Science*, 66(7), 3211-3233.