Driving the Drivers: Algorithmic Assignment in Ride-Hailing*

Yanyou Chen
University of Toronto

Yao Luo
University of Toronto

Zhe Yuan
Zhejiang University

This Version: *February 29, 2024*

Abstract

This paper investigates the impact of algorithmic assignment on worker behavior and welfare within the ride-hailing industry. We demonstrate how algorithms can impose a flexibility penalty on gig workers, despite their ostensible schedule autonomy. Utilizing rich transaction data from a leading ride-hailing company in Asia, we document a preferential assignment algorithm that favors drivers with longer working hours and consecutive hours during midday or late night. Drivers favored by the algorithm earn 8% more hourly than non-favored drivers. By constructing and estimating a two-sided market model, we quantify the welfare effects of such a preferential algorithm: Eliminating preferential assignment could raise ride fares by 7.79%, adversely affecting consumers and the platform. On the other hand, an additional 10% of drivers would switch to flexible schedules, leading to a 3.51% surplus gain, especially benefiting young, male, and local drivers.

Keywords: Two-Sided Market, Fair Pay, Work Schedule, Cross-Time Elasticity, Labor Supply, Market Power, Compensation Structure

*This paper was previously entitled "Driving the Drivers: Algorithmic Wage-Setting in Ride-Hailing." We are grateful to Victor Aguirregabiria, Gaurab Aryal, Robert Clark, Kory Kroft, Avi Goldfarb, Joshua Gans, Heski Bar-Isaac, El Hadi Caoui, and Chiara Farronato, as well as seminar participants at the Canadian Economic Association, China-VIOS, NBER SI Digitization, the Paris Conference on Digital Economics, IIOC, the BU Platform Strategy Symposium, CESI, CES-NA, NASM, ASSA, Digital Economics Conference at the Toulouse School of Economics, Peking University, Vanderbilt University, the University at Buffalo, HKUST, Shanghai Jiaotong University, the University of Toronto, Western University, and Zhejiang University for helpful discussions and comments. We also thank Duoyi Zhang, Anna Li, Mohaddeseh Heydari Nejad, Isaac Shiying Xi, Jinhan Liu, and Weiyue Zhang for excellent research assistance. All errors are our own.

1 Introduction

Recent years have seen the rapid acceleration of the gig economy—the ride-hailing market being a prominent example. Ride-hailing platforms offer riders an economical mode of transportation while allowing drivers the flexibility to tailor their work schedules to best fit their lives. Although gig work is usually thought of as piece-rate work that does not have the 'flexibility penalty' common to other jobs, such a penalty can still arise if the platform's algorithms introduce convexity into the earnings-hours relationship—for example, by prioritizing drivers who consistently work long hours. Specifically, algorithms may grant preferential order assignments to drivers based on their past work schedules, thereby influencing their hourly earnings. We refer to such phenomena as *algorithmic assignment.*

These assignment algorithms often restrict the work flexibility of gig workers, compelling them to undertake longer working hours. This issue is significant and not exclusive to ride-hailing; similar considerations apply to various other industries. For example, Uber Eats' algorithm gives preference to full-time over part-time workers when assigning orders. Similarly, DoorDash's algorithm discourages workers from strategically choosing orders; a worker who declines a long-distance delivery may stop receiving further delivery requests. Meanwhile, Instacart exercises significant control over the labor process, thereby restricting workers' autonomy over their time and the work they can undertake.[1] Each of these platforms employ a specific assignment algorithm to influence labor supply decisions. How do assignment algorithms affect the workers and the platform's performance? What are the welfare consequences? In this paper, we utilize rich transaction data from a leading ride-hailing company to explore these questions. We aim to provide the first empirical study of algorithmic assignment and its impact on worker behavior and welfare.[2]

First, we show that the ride-hailing company implements a preferential assignment algorithm, which limits drivers' utilization of schedule flexibility. In principle, a driver is free to choose whether or not to work for each hour of the day. However, a driver's hourly earnings depend not only on the specific hours during which they work, but also by their working hours in other parts of the day. For instance, at 7 AM, a driver who has also worked between 4 and 6 AM might have higher hourly earnings than a driver who starts work at 7 AM, as the algorithm may prioritize the former driver in the order of assignments at 7 AM.

[1]News reports for Uber Eats: How Uber got almost everything it wanted in Ontario's Working For Workers Act; DoorDash: Apps like Uber and DoorDash use AI to determine pay. Workers say this makes it impossible to predict wages. Instacart: At The Mercy Of An App: Workers Feel The Instacart Squeeze.

[2]There is a growing literature on algorithmic pricing, including works by Assad, Clark, Ershov and Xu (2023) for the retail gasoline market, and by Castillo (2023) and Gaineddenova (2022) for the ride-hailing industry. We provide the first empirical study on algorithmic order assignment in ride-hailing, complementing the existing literature on algorithmic pricing.

To understand why the platform's algorithm may prefer certain work schedules, we highlight one important channel the literature has thus far overlooked: *cross-time* labor supply elasticity. Most platforms apply surge pricing to balance demand and supply, which leverages *real-time* labor supply elasticity by increasing fares when demand exceeds supply. However, steeper fares discourage demand and reduce transactions if demand is overly elastic. In contrast, we focus on preferential assignment algorithms that adjust the labor supply in one time period by offering incentives in a different period, leveraging *cross-time* labor supply elasticity. In particular, such algorithms exploit the fact that drivers care about the total value of all hourly working decisions. Unlike the *real-time* reward of surge pricing, preferential algorithms may reward preferred drivers in any hour by prioritizing their trip assignments. Thus, even in hours when outside options are more attractive, some drivers may still prefer to work because they are rewarded in other hours. We provide a theoretical model in Section 2 to sharpen the intuition of when the adoption of a preferential algorithm is beneficial for the platform.

Second, we document significant wage differentials across workers due to work schedules. Using rich transaction data from a leading ride-hailing company in Asia, we show that three main factors drive the wage differential: high-performing drivers are given more ride requests per hour, wait fewer minutes for each request, and receive more requests from riders with lower cancellation rates. Here, it is important to note that the high- and low-performing labels do not reflect the drivers' efficiency or quality of work but merely indicate their performance from the platform's perspective. High-performing drivers typically have more committed work schedules and log more total hours on the platform. Due to the algorithmic assignment, drivers experience different hourly earnings, even if they are equally efficient and deliver the same quality of work. We also examine and rule out several alternative explanations documented in the literature about US ride-hailing markets (see, e.g., Cook, Diamond, Hall, List and Oyer, 2021), such as drivers strategically choosing where to work, strategically accepting or canceling orders, and driving faster. Moreover, we show the robustness of our finding to endogeneity concerns by employing instrumental variables: the rate of change in precipitation and the air quality index in drivers' hometown cities. The flexwork pay gap we identify is mainly due to algorithmic assignment, which penalizes low-performing drivers.

Third, to measure the impact of these preferential assignment algorithms on consumer and driver surplus, and to determine who gains and who loses from such a system, we construct and estimate a two-sided market model with time-varying demand and dynamic labor supply decisions. We propose a dynamic equilibrium model of a ride-hailing market and also incorporate the decisions of the platform in the two-sided market. The platform assigns orders and thus determines the hourly wage rate for drivers based on their overall work schedules. Our model accounts for riders' downward-sloping demand, drivers' dynamic labor supply with heterogeneous outside options, and the platform's fare and wage settings. Two sources of market power drive the platform's pricing decisions: the drivers face alternative time-varying outside options, and the riders have alternative modes of transportation. Drivers first choose to be high-performing or low-performing, and then choose their hourly work schedule by solving finite-horizon dynamic discrete choice problems. While drivers can set their own work schedules, the platform rewards high-performing drivers by assigning them more frequent and rewarding trips, leading to wage differentials between work schedules.

Our labor supply model with unobserved heterogeneity is point-identified using conditional choice probabilities in the drivers' dynamic labor supply. Regarding the estimation, we first estimate rider demand for service time for each hour of the day. We consider each hour a different market and aggregate our data to the day-hour level. We use the number of cars in competing ride-hailing companies on the given day as our supply-side instrumental variable. Then, we calculate the conditional choice probabilities of working for each hour of the day, based on drivers' observed work schedules. Together with the observed wage sequences, we then estimate the parameters in the labor supply model. Combining the estimated labor supply model and the rider demand model, we show how the platform leverages cross-time labor supply elasticity using the preferential algorithm. When ride fares are held fixed, eliminating the preferential algorithm would decrease labor supply, resulting in driver shortages for most hours. Our results show that the relation between wage differentials and labor shortages is not one-to-one. Instead, the platform smooths out the payment of high incentive wages by leveraging the variations in demand elasticity and the differing reservation values of drivers over time.

Next, we calculate the impact of eliminating the preferential assignment algorithm on consumer and driver surplus. Results show that both the platform and the riders would be worse off, but the drivers would be better off, enjoying more flexibility in choosing a work schedule under "fair" pay. Holding ride fares fixed while eliminating the preferential algorithm would result in platform revenues decreasing by 12.16% and total surplus decreasing by 7.16%. The proportion of high-performing drivers would decrease by 11.48%, as more drivers switch to being lower performing. Among the switchers, the driver surplus would increase by 3.51%. Allowing the platform to re-optimize ride fares after eliminating the preferential algorithm would raise ride fares to re-balance demand and supply. As a result, ride fares would increase by 7.79%. We also look at differential impacts across drivers. Namely, female and older drivers who choose to be high performing are more likely to suffer from this policy change. The effect for female drivers in general is ambiguous, because women are also more likely to prefer more flexible work schedules, and thus experience a larger welfare gain from the elimination of the preferential algorithm. Drivers who are not local residents tend to work long hours regardless of whether the platform has a preferential algorithm, and thus would suffer a welfare loss from the elimination of the preferential algorithm. Lastly, we investigate what factors determine the effectiveness of the preferential algorithm. We find that the platform benefits more from implementing a preferential algorithm when rider demand is more elastic or when warm-up cost is greater. Meanwhile, the loss of driver surplus with a preferential algorithm is smaller in the same conditions.

Related Literature

First, our paper contributes to the labor literature on the wage penalty of job flexibility. For example, Altonji and Paxson (1992) documents the constraints on hours and compensating differentials; Goldin (2014) uses the job flexibility penalty to explain the gender pay gap; and Aaronson and French (2004) examines the wage differentials between part-time and full-time employment. Economists are aware that workers value alternative work arrangements, as documented by Mas and Pallais (2017). One would assume that the gig economy, or any form of piece-rate work, would solve the issue of the flexwork wage penalty. However, our paper is the first to show that a flexibility penalty can still arise in subtle ways if platforms use algorithms that rely on historical work patterns to determine future work assignment and hence wage rate. Our findings contribute to the understanding of labor dynamics in the gig economy and the role of algorithmic management in shaping work conditions and welfare outcomes.

Second, our paper contributes to the rapidly growing literature on taxi (e.g. Frechette, Lizzeri and Salz, 2019; Buchholz, 2022) and ride-hailing markets (e.g. Chen, Rossi, Chevalier and Oehlsen, 2019; Liu, Wan and Yang, 2019; Chen, Ding, List and Mogstad, 2020; Rosaia, 2023). The existing literature primarily focuses on the crucial pricing aspect of algorithms used by ride-hailing platforms. For example, Castillo (2023) studies Uber's surge pricing using an empirical model of the two-sided market with riders, drivers, and the platform. Ming, Tunca, Xu and Zhu (2019) further demonstrates that surge pricing improves rider and driver surplus as well as platform revenues. Gaineddenova (2022) examines whether decentralizing the pricing mechanism improves market efficiency. Building on the foundation established by these studies, our paper highlights another important aspect of these algorithms: the preferential assignment of orders based on drivers' work history. We contribute to the literature by documenting how platforms balance demand and supply via the implementation of preferential assignment algorithms, and study the impacts of such algorithmic assignment on worker behavior and platform performance.

Third, our model builds on the literature on two-sided markets. See Rysman (2009) for a comprehensive survey. We take this view to the ride-hailing market, allowing for driver and rider outside options. While Rysman (2004) proposes a general setting with oligopolistic competition between platforms, we focus on one leading platform. This simplification approximates the industry structure well and allows us to incorporate important dynamics in drivers' labor supply. In estimating rider preferences, we employ an IV approach, similar to Kalouptsidi (2014), to deal with unobserved factors that may affect demand and rider fare schedules. In estimating drivers' preferences, we propose a GMM estimator that integrates the CCP estimator of Hotz and Miller (1993). We contribute to the estimation of ride-hailing drivers' outside option values by incorporating unobserved heterogeneity into these options. This approach enriches our welfare analysis and the understanding of distributional consequences when we examine the effects of eliminating preferential assignment algorithms.

The remainder of this paper is organized as follows. Section 2 elaborates on the details of the preferential assignment algorithm and introduces a theoretical model to explain why the platform has incentives to implement such an algorithm. Section 3 describes our data and the construction of key variables. Section 4 presents reduced-form evidence of the algorithmic assignment favoring high-performing drivers. Section 5 describes an equilibrium model with dynamic labor supply decisions. Section 6 discusses our identification argument and estimation results, and Section 7 discusses our counterfactual experiments. Finally, Section 8 concludes. The Appendix contains all omitted details.

2 Preferential Assignment Algorithm

Online platforms worldwide have been accused of implementing preferential assignment algorithms to restrict the work flexibility of gig workers. Our paper aims to understand why platforms employ such algorithms, as well as the resulting wage differential and its implications for consumer surplus, driver surplus, and platform profit. Specifically, we study one of the leading ride-hailing platforms in Asia, which we refer to as "Platform X" for confidentiality.[3] Platform X's algorithm grants preferential order assignment to drivers based on their total working time, particularly during incentivized hours. Below, we elaborate in detail on

[3]The leading ride-hailing platforms in Asia include Uber, Lyft, Didi, Grab, Gojek, and Ola, among others. With the development of Asia's residential travel demands, the number of ride-hailing users in Asia grew to 800 million by the end of 2020. Platform X has millions of ride-hailing drivers and serves over one hundred million people globally, having collected an annual revenue of over $10 billion USD in 2020. Platform X offers several tiers of operations: express, premium, and luxury. Our study focuses on its express service. Like UberX in the US, express service is the most popular way to travel on Platform X in Asia.

how Platform X's preferential algorithm works.

2.1 Preferential Algorithm in Ride-Hailing

Platform X typically distributes requests to drivers within three kilometers. Within this designated radius, Platform X gives priority to particular drivers based on their individual order assignment score. Drivers earn points based on the hours they spend shuttling riders and the time of day for which they work for the platform. In the city we are investigating (as of 2018), Platform X's fare schedules segment each workday into six intervals: (1) morning, 7:00–10:00; (2) midday, 10:00–16:00; (3) afternoon, 16:00–19:00; (4) evening, 19:00–22:00; (5) night, 22:00–00:00; and (6) early morning, 00:00–6:59 (next day). The points earned per hour for drivers vary depending on the specific time of day. Certain time intervals are designated as incentivized hours, during which drivers receive more points per hour of work. In the app, drivers have access to the point-earning formula, which provides them with precise information about the number of points they can earn at different times of the day. On average, drivers earn 0.3 points for each order they fulfill. A driver's order assignment score is computed by summing up all the points they earned over the previous thirty days.

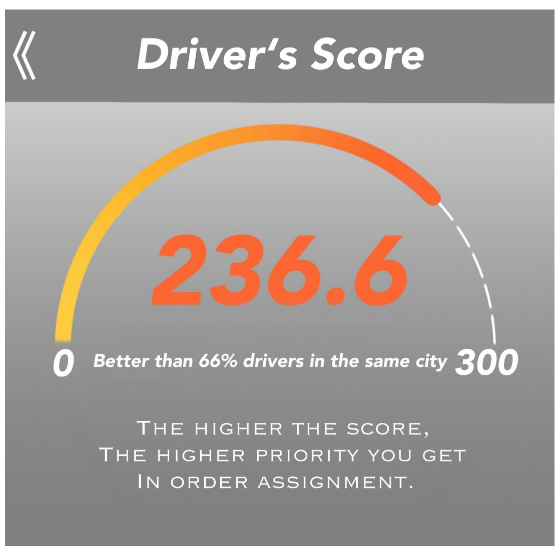

Figure 1: Information Displayed to Drivers

Figure 1 presents the information visible to the drivers.[4] The driver is presented with their current score, which in this case is 236.6. Next, there is a line indicating the percentile of their score, which in this case is "better than 66% of drivers in the same city." At the bottom of the screen, there is a line that explains the usage of the score to the drivers,

[4]For confidentiality reasons, we exclude all firm-identifying information in the graph.

stating that "the higher the score, the higher priority you will have in order assignment." To summarize, drivers are provided with the formula for earning points, possess full knowledge of their current score, and also understand that the score directly impacts their priority in order assignment.

Regarding fare schedules, all drivers face the same fare schedules on Platform X. Thus, hourly wage differentials across drivers mainly stem from systematic differences in their order assignment. Riders pay a 10 CCY base fare, 0.38 CCY per minute, and 1.9 CCY per mile for each Platform X Express trip. During the morning hours (7:00–10:00), the per-mile rate increases to 2.5 CCY, while during the afternoon (16:00–19:00), night (22:00–0:00), and early morning (0:00–7:00) hours, the per-mile rate is 2.4 CCY. Drivers receive 79.1% of the rider fare.[5]

2.2 Why Implement a Preferential Algorithm?

We will now elaborate on the platform's motivations behind introducing a preferential assignment algorithm. First, if the platform could apply first-degree price discrimination to both its consumers and drivers, it could maximize its profits to the fullest extent. In such a scenario, the platform would capture the entire surplus from both consumers and drivers, and introducing a preferential algorithm would not increase the platform's profit. The preferential algorithm is effective only in situations where the platform cannot achieve perfect price discrimination among drivers. This is because, while the platform may be able to extract consumer surplus through mechanisms like surge pricing, extracting the drivers' entire surplus is challenging due to factors such as labor laws and the design of the wage scheme. For instance, countries like France have regulations requiring a minimum payment for drivers per ride, resulting in a surplus for drivers.

Figure 2 presents a scenario in which the implementation of a preferential algorithm is profitable for the platform. Panels (a) and (b) represent two distinct time periods, t_1 and t_2. In both periods, the platform captures the entire consumer surplus by employing surge pricing, while providing drivers with a constant wage throughout each time period. The red lines represent the labor supply curve in each time period, while the blue lines represent the demand curve. Assuming without loss of generality that the drivers' outside option is zero at t_1. At time period t_1, the platform compensates drivers with a wage rate of w_1, leading to a surplus of B_1 for the drivers. The platform is unable to further decrease the wage rate at t_1 because of the minimum wage requirement. At t_1, there is an abundance of drivers willing to work at the wage rate w_1. However, in equilibrium, only L_1^* drivers are able to receive

[5] This information comes from Platform X's annual report.

orders and earn a surplus. The platform has the authority to select which drivers among the available pool will receive these orders. This ability to choose drivers necessitates the implementation of a preferential algorithm, which the platform utilizes to extract additional surplus from the drivers.

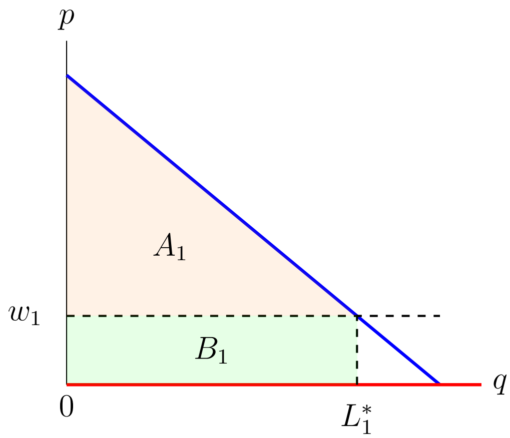

(a) t_1 with excessive labor supply

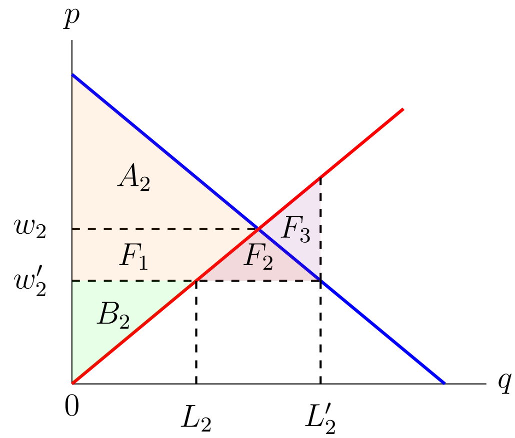

(b) t_2

Figure 2: Cross-time Labor Supply Elasticity

Notes: The red and blue lines represent the labor supply and demand curves in each time period, respectively. The light green area displays driver surplus, whereas the light yellow area displays platform surplus for each time period.

Without a preferential algorithm, the equilibrium wage rate in t_2 is w_2, determined by the point of intersection between the demand and supply curves. Given that platform surplus is A_1 in period t_1 and A_2 in period t_2, total platform surplus (profit) in absence of a preferential algorithm is $A_1 + A_2$, while driver surplus amounts to $B_1 + B_2 + F_1$.

With a preferential algorithm, the platform communicates to drivers that if they work during time period t_2, they will also be given priority should they choose to work during time period t_1. Hence, the platform can motivate drivers to work during time period t_2 without offering substantial incentive wages. This allows the platform to lower the wage rate to w_2' during period t_2 and still sustain the desired level of labor supply, denoted as L_2'. Therefore, with the implementation of a preferential algorithm, total platform surplus becomes $A_1 + A_2 + F_1 + F_2$, while driver surplus becomes $B_1 + B_2 - F_2 - F_3$.

Typically, the platform needs to offer high incentive wages to incentivize drivers to supply more labor. For instance, at wage rate w_2' in t_2, if the platform intends to increase its labor supply from L_2 to L_2', it must provide $F_2 + F_3$ in additional compensation to induce more drivers to work. Nevertheless, with the implementation of a preferential algorithm, such incentive wages are no longer necessary, and the platform can instead simply offer to prioritize drivers who work in period t_2 for order assignment in period t_1. This is because drivers who choose to work in t_2 are now compensated by earning a surplus in t_1, as the preferential algorithm creates a wage differential between drivers who work exclusively in t_1 and drivers who work in both t_1 and t_2. The power of implementing a preferential algorithm arises from the excess supply of drivers in period t_1 and the platform's ability to select which drivers will receive orders in such situations. The effectiveness of the preferential algorithm depends on the disparity in demand elasticity and drivers' reservation values across different time periods.

This simple theoretical model highlights the incentive for using a preferential algorithm. The preferential algorithm allows the platform to extract additional driver surplus in situations where it is unable to conduct first-price discrimination for drivers. The preferential algorithm leverages cross-time labor supply elasticity to extract additional driver surplus, while surge pricing utilizes real-time elasticity to maximize the platform's profit. Despite their different mechanisms, the platform achieves higher profits from utilizing both surge pricing and the preferential algorithm. In Appendix A, we use our theoretical model to illustrate the contrasting ways surge pricing and the preferential algorithm operate, as well as to highlight their potential complementarity. Specifically, we compare the equilibrium outcomes in four scenarios: without surge pricing and the preferential algorithm, with only surge pricing, with only the preferential algorithm, and with both surge pricing and the preferential algorithm.

3 Data

To study the wage differential resulting from the preferential algorithm, we acquire transaction-level data from the Transportation Bureau of a major city in Asia. We observe the universe of completed transactions by all ride-hailing platforms in December 2018 for that city.[6] We also observe drivers' attributes, such as age, gender, and place of birth.

For each transaction, we observe the trip's origin, destination, and distance, as well as the duration spent on passenger pickup and transportation, and the corresponding payment the driver receives. The transaction-level data allow us to observe the drivers' detailed work schedules and the characteristics of the orders they receive. In particular, the order details—encompassing origin, destination, wait time, pickup time, and drive time—enable us to assess the quantity and quality of orders received by different drivers. Therefore, we can investigate the underlying factors contributing to drivers' wage differentials. While we have

[6]The city we study has a population of around 11 million.

data on completed transactions from all ride-hailing platforms, our main analysis focuses exclusively on Platform X for two key reasons. First, Platform X holds a dominant position in the city under study, accounting for more than 90% of the market share. Second, our data indicate that drivers rarely multi-home or switch between different platforms, suggesting that platform competition is almost negligible in our city of study.[7]

Table 1 summarizes our data set. The unit of observation is at the driver-hour level.[8] A driver fulfills, on average, 1.9 orders per hour and earns 50 CCY. The number of fulfilled orders per hour ranges from 1 to 9 between the 25$^{\text{th}}$ percentile and the maximum value, demonstrating significant variation. For each hour worked, drivers generally only spend half the time transporting riders, devoting on average 10 minutes of their time to picking up riders and another 19 minutes to waiting for subsequent orders. Considering that these latter two (non-billable) tasks take up a substantial proportion of drivers' work schedules, having a higher priority in order assignment plays a significant role in improving a driver's hourly wage.

Table 1: Summary Statistics (Driver-Hour)

	Mean	Std. Dev.	Min	25 Pctl	Median	75 Pctl	Max
Hourly Wage (CCY)	49.98	24.52	0	32.83	47.42	62.74	286.86
Earning Time (minutes)	30.60	12.01	0	21	31	40	60
Pickup time (minutes)	10.62	6.67	0	6	10	15	60
Order wait time (minutes)	18.78	14.32	0	6	17	29	60
Number of Orders	1.89	1.11	0	1	2	3	9
Distance (km)	14.11	7.41	0	8.78	13.1	18.2	94.13
Number of Observations				4,182,318			

Table 2 summarizes the driver characteristics. There are $40,104$ unique drivers in our data, of which 2.7% are female and 37.4% are local drivers. We define local drivers as those with local household registration permits.[9] Household registration permits have a profound impact on residents' eligibility to purchase property and access schools and childcare facilities, thereby influencing their job market opportunities. Table 2 shows that drivers work a median

[7]Online Appendix B shows a detailed analysis of multi-homing.

[8]The unit of observation in our raw database is at the driver-rider-order level. We primarily concentrate on weekdays (a total of 21 days) due to variations in supply, demand, and fee schedules between weekdays and weekends. Following the literature, we conduct our main demand estimation and counterfactual analysis at the driver-hour level. For more details on how we derive the driver-hour level data from the raw data, please refer to Online Appendix C.

[9]Household registration permits are issued by the government and indicate the particular region a person is from; in this region, the registrant is entitled to benefits such as hospitals, schools, or land-purchasing rights.

of 13 out of the 21 workdays. There is considerable heterogeneity across the number of days each driver works, ranging from 5 to 19 days between the 25th and 75th percentiles. Moreover, there is also substantial variation in the number of work hours per day, with the 25th percentile driver working 4.8 hours, while the 75th percentile driver works 10.5 hours.

Table 2: Summary Statistics of Driver Characteristics

Characteristic	Mean	Std. Dev.	Min	25 Pctl	Median	75 Pctl	Max
Age	37.29	8.24	21	31	37	43	61
Work Days	12.02	7.03	1	5	13	19	21
Daily Work Hours	7.61	3.59	1	4.75	8.09	10.47	18

The summary statistics of drivers working for Platform X in our city of study are substantially different from Uber data from the US market. For instance, according to Cook, Diamond, Hall, List and Oyer (2021), 27.3% of Uber drivers are women, whereas just 2.7% of the drivers in our data are female. Additionally, our drivers dedicate significantly more time to their work, averaging around 7.6 hours per day, while Uber drivers usually work approximately 3 hours per day. The substantial difference in working hours may be attributable to the preferential algorithm discussed in this study.

High-Performing and Low-Performing Drivers

Section 2.1 provides descriptions of the preferential algorithm used by Platform X. In this subsection, we explain the key variables that we utilize to characterize drivers and assess the algorithm's impact in our main analysis. First, we examine whether the algorithm exhibits any preference towards drivers based on their past work schedules. Then, we elaborate on the construction of the key variables designed to capture drivers' past work schedules for our main analysis.

First, we verify whether working longer periods, particularly during incentivized hours, leads to higher priority of order assignment and subsequently higher hourly wages, as explained in Section 2.1. We regress the hourly wage of a driver on the total number of hours worked in a month and the percentage of incentivized hours worked, controlling for day, hour, and operation area fixed effects. According to interviews with drivers and engineers at Platform X, midday (10 AM to 4 PM) and night hours (starting from 7 PM) were identified as incentivized hours during our study period. Table 3 shows the results. Generally, drivers who work more, especially during incentivized hours, earn a higher hourly wage than other drivers. Column (2) of Table 3 shows that working one additional hour in a month increases a driver's hourly wage by 0.3 cents. Given that the 25th to 75th percentile drivers work 27 to 172 hours, their hourly wage gap is 0.435 CCY or 0.87% of the average hourly wage. A more important feature of the higher hourly wages is the percentage of incentivized hours worked. Allocating 1% more work time to incentivized hours increases the hourly wage by 0.187 CCY. Given that the 25$^{\text{th}}$ to 75$^{\text{th}}$ percentile drivers spend 55% to 72% of their work time on incentivized hours, respectively, their hourly wage gap is 3.2 CCY or 6.4% of the average hourly wage. These findings indicate that the preferential algorithm employed by Platform X aligns with its description: Drivers who put in longer hours and work more during incentivized hours are indeed given priority in order assignments, resulting in a higher hourly wage. In Section 4, we thoroughly investigate how the preferential algorithm drives this wage differential and eliminate alternative explanations.

Table 3: Factors Correlated with Hourly Wage

Hourly Wage	(1)	(2)
# of Work Hours in month	0.003***	0.003***
	(0.000)	(0.000)
% Incentivized Hours		18.724***
		(0.170)
Constant	54.918***	39.201***
	(0.126)	(0.190)
Observations	4,182,318	4,182,318
R-squared	0.040	0.043

Notes: We control for day-hour fixed effects, origin district fixed effects, and destination district fixed effects. Standard errors are in parentheses. *** p<0.01.

Next, we aim to characterize drivers based on their past work schedules. A computational challenge arises from the fact that drivers have the option to choose whether to work each hour, leading to potentially 2^{24} different driver work schedules on any given day. Given the computational impracticality of tracking this vast number of possibilities in our analysis, we explore clustering drivers into discrete types based on their historical work schedules. Discretizing the driver types in this manner also makes our dynamic labor supply model more tractable.

We implement machine learning algorithms to cluster drivers based on their past work schedule. Our findings indicate that we can primarily classify drivers into two distinct types, and this result is robust to various specifications of clustering. The first type comprises high-performing drivers, who can also be perceived as committed or full-time workers. The second type are low-performing drivers, who can also be perceived as uncommitted or part-time workers. We formally classify the two types of drivers as follows: A driver is considered high-performing if they work for at least two consecutive hours during incentivized hours (midday or night) on a minimum of 8 out of the 21 workdays. Conversely, a driver is classified as low-performing if they do not meet these criteria. Throughout our subsequent analyses, we use "H-type" and "L-type" to represent the status of drivers as high-performing and low-performing, respectively.[10]

Table 4: High/Low-performing Driver Characteristics

	High-performing (1)	Low-performing (2)
Panel I: Driver/Vehicle Characteristics		
% female	2.2%	3.5%
% non-local	69%	53%
Age	37.2	37.4
Panel II: Performance (in a month)		
Work Days	17	5
Work Hours	159	26
# of orders	301	46
Monthly Revenue	7,985	1,202
Panel III: Performance (in an hour)		
Work Time	30.7	29.3
Pickup time	10.7	10.2
Idle Time	18.6	20.4
# of orders	1.90	1.76
Hourly Revenue	50.4	46.5
# of drivers	23,712	16,392
Share of Drivers	59.1%	40.9%

Table 4 summarizes the characteristics of high- and low-performing drivers. There are $23,712$ high-performing drivers and $16,392$ low-performing drivers. Panel I reports the drivers' characteristics. High-performing drivers are more likely to be non-local and male. Women account for 2.2% of the high-performing drivers and 3.5% of the low-performing ones. Non-locals account for 69% of the high-performing drivers and only 53% of the low-

[10]Online Appendix D describes the machine learning algorithm we use to cluster drivers. We performed various robustness checks by changing the threshold for identification as a high-performing driver. For example, we changed the required number of days from 8 to 9, 10, 11, etc., out of 21 workdays, and we put further restrictions on the total number of hours worked per month at various levels. The reduced-form results in Section 4 are robust to these definitions. As explained in the main context, the key feature of high-performing drivers is the percentage of hours worked consecutively in incentivized hours. Because of the high fixed cost of starting a shift, the number of consecutively worked hours during incentivized hours is highly correlated with total number of hours worked. This may help explain why the two criteria we use in the main context are robust to all the variations mentioned here.

performing ones. The average age is comparable between high- and low-performing drivers. Panels II and III report driver performance. On average, high-performing drivers work more, averaging 17 out of 21 workdays, while low-performing drivers, on average, work 5 out of 21 workdays. In any given hour, conditional on working, high-performing drivers have more passenger-service time (30.7 minutes versus 29.3 minutes) and spend less time waiting for orders (18.6 minutes versus 20.4 minutes). High-performing drivers also complete more orders (1.9 orders versus 1.74 orders) and earn more (50.4 CCY versus 46.5 CCY per hour) than low-performing drivers.

Last, to further validate the categorization into two driver types in our analysis, we compare results obtained from this categorization with those obtained from imputed assignment scores. While the platform's assigned driver scores are not directly observable in our data, we observe the number of completed transactions and hours of operations for all drivers. This allows us to compute an assignment score that mimics the actual score calculation process, which we refer to as the imputed assignment score. We present the results based on imputed assignment scores in Online Appendix E.

We obtain two primary findings using imputed assignment scores. First, there exists a significant correlation between a driver's hourly wage and their assignment score, with drivers possessing higher assignment scores typically earning a higher hourly wage than their lower-scoring counterparts. Second, we observed a significant difference in scores between drivers classified as high-performing versus those categorized as low-performing, and both the imputed driver score and our clustering method reflect similar aspects of drivers' historical work schedules. These findings confirm the validity of our approach to cluster drivers into high- and low-performing groups based on their work schedules.

4 Reduced-Form Evidence

In this section, we provide evidence that high-performing drivers earn a higher hourly wage. Subsequently, we investigate what factors contribute to this wage disparity. Last, we rule out alternative explanations for the observed wage differential between high- and low-performing drivers, including strategically choosing where to work, strategically selecting and canceling orders, and driving faster.

4.1 Wage Differential

First, conditional on working in the same hour, we test whether high-performing drivers earn more than low-performing ones. We regress the hourly wage of a driver on an indicator of being high-performing and control for day-hour, origin, and destination fixed effects. Table 5 shows a significant difference in hourly wage between high- and low-performing drivers. High-performing drivers earn 3.8 CCY, or 8.2% more per hour, than their low-performing counterparts. The result is very robust, with or without controlling for various fixed effects.

Table 5: Wage Differential: High-performing versus Low-performing

Dependent Variables	Hourly Wage		
	(1)	(2)	(3)
High-performing	3.886***	3.794***	3.851***
	(0.0397)	(0.0393)	(0.0391)
Constant	46.49***	46.57***	47.24***
	(0.0376)	(0.0372)	(0.0701)
Day-Hour FE		Y	Y
Origin FE			Y
Destination FE			Y
Observations	4,182,318	4,182,318	4,182,318
R-squared	0.002	0.039	0.050

Notes: Standard errors in parentheses. *** p<0.01

Given that high-performing drivers earn significantly higher hourly wages, we investigate what factors drive this wage differential. To this end, we study the characteristics of the orders that high- and low-performing drivers receive. For example, we evaluate the number of orders they receive and how often a rider cancels their order. We also compare the the amount of idle time and the time spent serving the customer for the two types of drivers. Table 6 shows the results. Column (1) shows that conditional on working in the same hour, high-performing drivers receive more orders than low-performing ones. On average, high-performing drivers receive 0.125 more orders or 7.1% more every hour. Second, orders assigned to high-performing drivers are 2.8% less likely to be canceled by riders (column 2).[11] Because high-performing drivers get assigned more orders every hour, they also drive 0.748 more kilometers and spend 5.4% more time carrying riders in an hour (column 4).[12] More importantly, high-performing drivers spend 10.5% less time waiting for orders (column 5). This result is consistent with the descriptions in Section 2 of how the algorithm prioritizes

[11]Our main analysis throughout this paper uses data on completed transactions. Our data include information on canceled orders in the first ten days (from December 1 to December 10, 2018). We use data on completed transactions and canceled orders for all regressions involving cancellation rates. Therefore, the number of observations differs from that of other regressions.

[12]Gaineddenova (2022) shows that drivers prefer more expensive trips with a shorter pickup distance, using data from a decentralized ride-hailing platform.

high-performing drivers for better order assignments.

Table 6: Driving Forces of Wage Differential

Dependent Variables	# Orders	Cancellation Rate (Rider)	Drive Dist	Earning Time	Idle Time
	(1)	(2)	(3)	(4)	(5)
High-performing	0.125***	-0.0023***	0.748***	1.579***	-2.140***
	(0.0018)	(0.0004)	(0.0003)	(0.0187)	(0.0221)
Constant	1.468***	0.0894***	12.85***	32.35***	17.04***
	(0.00313)	(0.0005)	(0.0212)	(0.0334)	(0.0395)
Mean of Low-performing	1.76 (orders)	8.2%	13.4 (km)	29.3 (min)	20.4 (min)
High-performing compared to Low-performing	7.1%	-2.8%	5.6%	5.4%	-10.5%
Observations	4,182,318	4,815,026	4,182,318	4,182,318	4,182,318
R-squared	0.080	0.006	0.045	0.100	0.115

Notes: In all columns except column (2), our analysis uses completed transactions, which are available from Dec. 1, 2018, to Dec. 31, 2018. In column (2), we also include canceled orders to compute rider cancellation rates. Information on canceled order is available from Dec. 1, 2018, to Dec. 10, 2018. Standard errors are in parentheses. All specifications control for day-hour fixed effects, origin district fixed effects, and destination district fixed effects. *** p<0.01

In summary, Table 6 shows that three main factors are driving the wage differentials between high- and low-performing drivers. The former are given more rides from the platform, waste less idle time waiting for orders, and receive more orders from higher quality riders (those with a lower probability of rider-initiated cancellation). As the platform's algorithm determines the assignment of orders, we hereafter term the systematic difference in the quantity and quality of order assignments based on work schedule (high-performing versus low-performing) algorithmic preferential assignment.

4.2 Rule Out Alternative Explanations

There could, of course, be alternative explanations for the wage differentials between high- and low-performing drivers. Rather than having the algorithm prioritize different work schedules when assigning orders, some may argue that drivers make decisions endogenously, resulting in the observed wage difference. For example, Cook, Diamond, Hall, List and Oyer (2021) find that the gender earnings gap amongst drivers can be entirely attributed to three factors: experience on the platform (learning-by-doing), preferences and constraints about where to work (driven largely by where drivers live and, to a lesser extent, by safety), and preferences about driving speed. To provide a robustness check for our findings, we consider three alternative explanations and use our data to prove that such alternative explanations are unlikely to be true in our context. First, high-performing drivers may have better knowledge of the popular rider pickup areas (hot spots) and get more orders. Second, high-performing drivers may learn how to reject and cancel rides strategically. Third, high-performing drivers may drive faster than others and earn a higher hourly rate.

High-Performing Drivers Strategically Choose Where to Work

First, we explore whether high-performing drivers have better knowledge of hot spots, and hence are strategically choosing where to work in order to earn more hourly.[13] There are eight districts in the city we study. We first examine where the high- and low-performing drivers work and whether they tend to pick up or drop off clients in different areas. Table 7 suggests no substantial difference between the origin or destination districts where high- and low-performing drivers work.

Table 7: Active Area for High-performing and Low-performing Drivers

District	Origin		Destination	
	Low-performing	High-performing	Low-performing	High-performing
1	7%	7%	7%	7%
2	9%	8%	9%	8%
3	20%	22%	21%	23%
4	7%	7%	7%	7%
5	16%	15%	15%	14%
6	10%	11%	10%	11%
7	16%	15%	16%	15%
8	16%	15%	15%	13%
Total	100%	100%	100%	100%

To better control for location-fixed effects, we manually divide the eight districts into even finer 1km $\times$ 1km grids. Because we observe the coordinates of each pickup and drop-off location, we can accurately place trip origins and destinations into one square on the fine grids. Column (2) of Table 8 reports the results of re-running our main regression with day-hour and grid fixed effects, which are close to the benchmark result in column (1).[14] This shows that the wage differential between high- and low-performing drivers cannot be explained by high-performing drivers picking up or dropping off passengers from certain locations or neighbourhoods. We further divide each hour into four 15-minute intervals as a

[13] For example, Haggag, McManus and Paci (2017) find that New York taxi drivers accumulate neighborhood-specific experience, which helps them find riders.

[14] Column (1) of Table 8 has the same specification as column (3) of Table 5, where we control for origin, destination, and day-hour fixed effects. The two results differ because we restrict our sample to drivers who were active in the last hour in Column (1) of Table 8.

robustness check. Instead of controlling for day-hour fixed effects, we now control for day-hour-15-minute fixed effects. With a finer measure of both location- and time-fixed effects, we are essentially comparing drivers who work in the same location at the same time. The only difference between the drivers is their performance level, which depends on their past work schedules. Column (3) reports the result after controlling for day-hour-15-minute and grid fixed effects, and column (4) reports the result controlling for day-hour-15-minute-grid fixed effects. The results in all robustness checks are close to our benchmark result in column (1). Thus, a knowledge of hot spots and strategically choosing where to work are unlikely explanations for the wage differential between high- and low-performing drivers.

Table 8: Wage Differentials with Finer Grids

Dependent Variables	Hourly Wage (OLS)				IV
	(1)	(2)	(3)	(4)	(5)
High-Performing	2.742***	2.704***	2.705***	2.731***	8.99***
	(0.0391)	(0.0453)	(0.0448)	(0.0448)	(0.9614)
Constant	47.98***	21.38	23.90	47.56***	41.85***
	(0.0701)	(22.75)	(22.51)	(0.0427)	(0.8755)
Time Controls:					
Day-Hour	Y	Y			
15Minute			Y		
Location Controls:					
Origin/Destination	Y				
Grid		Y	Y		
Grid-15Minute				Y	Y
Observations	3,160,528	3,160,528	3,160,528	3,160,528	3,160,528
R-squared	0.053	0.075	0.094	0.097	(omitted)

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

We restrict our sample to drivers who were active in the last hour.

To mitigate potential biases in driver selection based on unobservable characteristics, we have employed instrumental dummy variables (IVs): the rate of change in precipitation and air quality index (AQI) in the driver's hometown city between 2017 and 2018. These weather variables satisfy the two conditions required for valid IVs: first, the occurrence of precipitation and changes in air pollution may be correlated with a driver's decision to become a high-performing driver. For instance, alterations in precipitation and air pollution might affect the benefits of farming, potentially motivate more drivers to leave their hometowns and become full-time drivers in the focal city under study, as suggested by (Miguel et al., 2004). Second, the variation in weather conditions in a driver's hometown should not directly impact the driver's hourly rate or order distribution in the city being studied. The IV results are presented in column (5), revealing a more significant wage differential between high- and low-performing drivers. We include additional IV results in Online Appendix F.

High-Performing Drivers Strategically Cancel Orders

Second, the literature shows that more experienced drivers can learn how to strategically reject and cancel rides, hence earning more. To examine whether such a mechanism exists in our data, we regress the probability of a driver canceling an order on driver type and control for day-hour, origin, and destination fixed effects. The results in column (1) of Table 9 show that, if anything, high-performing drivers have a lower cancellation rate than low-performing ones in our data. The robustness check results, controlling for finer location and time-fixed effects, remain consistent with these findings. It is very difficult for drivers to cancel an order on Platform X , which may help explain why Platform X drivers behave differently from the Uber drivers described in Cook, Diamond, Hall, List and Oyer (2021). Because high-performing drivers are less prone to cancelling an order, it is unlikely, in our case, that the higher hourly wage of high-performing drivers is caused by drivers strategically rejecting and choosing rides.

Table 9: Driver Cancellation and Driver Speed

Dependent Variables	Probability of Cancellation (by Driver)	Driving Speed
	(1)	(2)
High-performing	-0.0062***	0.1313***
	(0.0002)	(0.0194)
Constant	0.0365***	0.410***
	(0.0003)	(0.0006)
Mean of Low-performing	3.4%	24.63 (km/h)
High-performing compared to Low-performing	-18.2%	0.5%
Observations	4,815,026	4,168,889
R-squared	0.004	0.089

Notes: Standard errors are in parentheses. We control for day-hour fixed effects, origin district fixed effects, and destination fixed effects. *** p<0.01

High-Performing Drivers Drive Faster

Third, some drivers may drive faster than others, hence completing more trips and earning a higher hourly wage. We examine whether high-performing drivers drive faster than their low-performing counterparts by regressing the average driving speed per hour on an indicator of being high-performing. We continue to control for day-hour, origin, and destination fixed effects. Column (2) of Table 9 shows the results. While we do find that high-performing drivers drive slightly faster (0.5%) than low-performing ones, the 0.5% faster driving speed is insufficient to explain the 8% wage differential we find in our main analysis. This 0.5% faster driving speed only converts into an extra 0.24 CCY per hour,[15] thus explaining very little of the 3.8 CCY (or 8%) wage differential between high- and low-performing drivers.

To summarize, we examine three alternative explanations for the wage differential between high- and low-performing drivers: that high-performing drivers may strategically choose where to work, strategically select and cancel orders, or drive faster. However, upon more in-depth analysis of our data, we can rule out all three of these as likely explanations of the observed wage differential between high- and low-performing drivers.

5 Model

Given that the preferential algorithm prioritizes specific drivers based on their work schedule for order assignments, the labor supply decision is now subject to the rules specified by the preferential algorithm. To understand who benefits and who loses under such a preferential algorithm, we propose a dynamic equilibrium model of a ride-hailing market. In our model, each driver decides when and how long to work, depending on the wage rates and reservation values.

We model the decisions of market participants for one day. At each hour of the day, there is a demand curve for rides, $D_t(P_t)$. Given this demand curve, the platform makes two types of decision. First, it determines the price to charge riders, P_t. We allow for dynamic pricing and thus let prices vary across different times of the day. Second, the platform's algorithm allocates ride orders to each driver. The algorithm distinguishes between two types of drivers: high-performing and low-performing. The former commit to working consecutively for at least 2 hours during incentivized hours, 10 AM–4 PM and 7 PM–6 AM the next day. The latter make no work schedule commitments. Driver i first decides whether to be high

[15] The average driving speed for a low-performing driver is 24.63 km/h. Thus, by driving 0.5% faster, high-performing drivers drive 0.12 km more per hour. The average ride fare is about 2 CCY/km. Therefore, assuming that the extra 0.12 km is entirely used in carrying a rider, without any time wasted to wait for and pick up customers, then $0.12 * 2 = 0.24$ CCY. Therefore, this 0.5% faster driving speed only converts into an extra 0.24 CCY per hour.

performing or low performing $\tau \in \{H, L\}$. We assume that drivers choose their type (H or L) at the start of the day, and drivers cannot change their worker type throughout the day. Conditioning on the choice of being an H-type or L-type, each driver chooses whether to work for each hour of the day. The problem is dynamic, because whenever a driver starts working or resumes working after a break, there is a fixed "warm-up" cost. If the driver chooses to be high performing, the dynamic problem is under the constraint that working hours need to satisfy certain work schedules. Otherwise, the problem is unconstrained.

Our model prioritizes within-day dynamics over day-to-day dynamics to emphasize the major trade-off involved in drivers' standard decision-making process. According to our interviews with drivers, they tend to maintain consistent working habits from day to day. This may be attributed to the fact that the initial introduction of ride-hailing in the focal market took place as early as 2014, allowing drivers to have established their daily patterns by the time of our study in 2018. Our data also validate this consistency in hourly driving patterns. In Online Appendix G we analyze the labor supply patterns of drivers for Platform X in our data, especially across-day patterns. We show that across-day patterns are very stable, especially for High-performing drivers. As per our interviews with drivers, day-to-day dynamics are typically established within a short period. For individuals who have not yet settled into a consistent routine, choosing full-time ride-hailing work requires making arrangements with family members, such as finding suitable daycare for their children. Once drivers have settled into their day-to-day dynamics, they revert to the within-day dynamics, as described in our study.[16]

We use a bold typeface to denote vectors containing values for each hour of the day. For example, $\boldsymbol{P}$ denotes prices for all $t = 1, \cdots, 24$. The sequence of wage rates is $\boldsymbol{W}^{\boldsymbol{\tau}}$, which is determined by the platform's pricing decisions $\boldsymbol{P}$ and the algorithm deciding which driver receives an order.

5.1 Drivers' Dynamic Labor Supply

Drivers first choose to be either high performing or low performing. Low-performing drivers solve an unconstrained dynamic discrete choice problem of when to work. As per our discussion in Section 3, high-performing drivers are required to work consecutively for at least 2 hours during incentivized hours, between 10 AM–4 PM and 7 PM–6 AM the next day. Besides fulfilling the required working hours, a high-performing driver also makes an hourly

[16] While it is technically feasible to incorporate day-to-day dynamics into our existing one-day model, doing so presents significant challenges due to data limitations and computational complexities. As a result, we leave this to future research. Investigating day-to-day dynamics could offer interesting insights into how drivers make choices among platforms and whether they opt to become ride-hailing drivers.

choice of whether to work. A driver who chooses to be high performing must choose which minimum requirement to satisfy in advance. For example, driver A may choose to be a high-performing driver by committing to work between 10 AM and 12 PM. Between 10 AM and 12 PM, driver A will be active on the app with probability 1; at any other time of the day, driver A can freely choose whether to work. We assume that drivers choose their work type (H or L) at the start of the day, and that drivers cannot change their type during that day. A low-performing driver B does not commit to any work schedule. *Ex post*, even if driver B ends up working long hours, including from 10 AM to 12 PM, they would still be considered a low-performing type.

Sixteen possible work schedules satisfy the high-performing requirement.[17] Work status is summarized by different work schedules, $\mathcal{L} \equiv \{0\}$ and $\mathcal{H} \equiv \{1, \cdots, 16\}$. The choice of work schedule is a simple logit model,

$$N^j = N \cdot \frac{\exp(EV^j)}{\sum_{k=0}^{16} \exp(EV^k)},$$

where N is the total number of potential drivers and EV^j represents the expected value of choosing work schedule j.[18] Therefore, the total number of high-performing drivers is $N^H = \sum_{k=1}^{16} N^k$, and the total number of low-performing drivers is $N^L = N^0$.

After deciding whether to be H-type, drivers then find the optimal solution to their dynamic discrete choice problem by choosing whether to work at each time t. Drivers observe the warm-up cost κ, sequence of hourly wages $\boldsymbol{W^\tau}$, and reservation values $\boldsymbol{O}$. At each time t, low-performing drivers compare the hourly wage plus the difference in expected future values to the value of their outside option before deciding whether to work at time t. This is a dynamic problem because if the driver chooses to work at time t and continues working at $t+1$, they would not need to pay an extra warm-up cost at $t+1$. Hence, the expected future value at time t is higher if the driver chooses to work than if they choose not to work at time t. High-performing drivers have to work with probability 1 during committed hours. At any other time of day, high-performing drivers solve the same dynamic discrete choice problem by comparing their hourly wage plus the difference in expected future values to their outside option, then decide whether to work at each time t.

Specifically, at the beginning of hour t, a driver receives a random draw from the wage

[17] For example, if a driver chooses to satisfy the high-performing requirement by working 10 AM–12 PM, then they are categorized as schedule 1. If a driver chooses to satisfy the high-performing requirement by working 11 AM–1 PM, then they are categorized as schedule 2, etc.

[18] We use the number of unique drivers in the 21 workdays as the number of potential drivers in our model.

distribution and another draw from the outside option:

$$
\begin{aligned}
U_{1t}^{\tau} &= \underbrace{W_t^{\tau}}_{\text{preferential wage rate}} + \sigma \cdot \epsilon_{1t}, \\
U_{0t}^{\tau} &= \underbrace{O_t^{\iota}}_{\text{outside option value}} + \sigma \cdot \epsilon_{0t},
\end{aligned} \tag{1}
$$

where U_{1t}^{τ} corresponds to the utility associated with working during time period t, while U_{0t}^{τ} represents the utility of not working during the same period. O_t^{ι} represents the reservation value from working on something else, and $\epsilon_{\cdot t}$ represents the error term, which is Type-I extreme value distributed. Differences in reservation values among various groups of drivers may impact their labor supply decisions. To account for this, we consider the presence of unobserved heterogeneity in drivers' reservation values. Specifically, we have $O_t^{\iota} = O_t + \eta_{\iota,t}$, where O_t is the average reservation value per hour, $\eta_{\iota,t}$ is driver-specific unobserved heterogeneity representing their preference for certain parts of the day, and ι is the unobserved heterogeneity type. According to our survey of drivers, these preferred parts often closely align with the intervals on Platform X's fare schedule.[19] Therefore, we consider seven unobserved heterogeneity types that align with these intervals. For the benchmark driver group 0, the unobserved heterogeneity term η_{0t} is set to 0 for all time periods t. In contrast, driver group 1 has an unobserved heterogeneity term η_{1t}, equal to η_1, during the time period from 7 AM to 10 AM and 0 for all other time periods. Similarly, driver group 2 shows an unobserved heterogeneity term η_{2t}, equal to η_2, during the period from 10 AM to 4 PM.

If the driver took the outside option in the previous hour, there is a fixed warm-up cost $\kappa > 0$ to start work. This is to rationalize the fact that drivers often drive for consecutive hours. The value function for the low-performing driver at any time t that is not the first or the last period is

$$
V_t^L = \begin{cases}
W_t^L + \sigma \cdot \epsilon_{1t} + \beta E V_{1t+1}^L & \text{if } a_t = 1 \;\&\; a_{t-1} = 1, \\
W_t^L - \kappa + \sigma \cdot \epsilon_{1t} + \beta E V_{1t+1}^L & \text{if } a_t = 1 \;\&\; a_{t-1} = 0, \\
O_t^{\iota} + \sigma \cdot \epsilon_{0t} + \beta E V_{0t+1}^L & \text{if } a_t = 0.
\end{cases}
$$

Here, W_t^L represents the wage rate for the low-performing driver at time t, κ is the warm-up cost, and σ is the scale parameter. The terms EV_{1t+1}^L and EV_{0t+1}^L denote the expected values if the driver chooses to work or not to work, respectively, at time t. The labor supply decision of the driver is denoted by a_t, where $a_t = 0$ indicates not working at time t.

[19] Section 2.1 offers details of the fare schedules.

The value functions for all time periods, for both low- and high-performing drivers, appear in Online Appendix H. We solve the dynamic discrete choice problems through backward induction.

Individual driver choices, in turn, generate the aggregate labor supply for each hour by driver type:

$$\begin{aligned} N_t^H &= \sum_{j=1}^{16} N^j \times \Pr(\text{work in hour t}|\text{work schedule j}), \\ N_t^L &= N^0 \times \Pr(\text{work in hour t}|\text{work schedule 0}), \end{aligned}$$

where the conditional choice probabilities $\Pr(\cdot|\cdot)$ are the solutions to the above-mentioned dynamic discrete choice problems. We denote the type-specific labor supply as

$$\begin{aligned} N_t^H &= \mathcal{N}_t^H(\boldsymbol{W^H};\boldsymbol{\theta}) = \mathcal{N}_t^H(\boldsymbol{P},\boldsymbol{s};\boldsymbol{\theta}), \\ N_t^L &= \mathcal{N}_t^L(\boldsymbol{W^L};\boldsymbol{\theta}) = \mathcal{N}_t^L(\boldsymbol{P},\boldsymbol{s};\boldsymbol{\theta}). \end{aligned}$$

5.2 Demand for Rides and the Platform's Problem

Riders only demand driver-earning hours. The number of earning hours they demand is $D_t(P_t)$, where P_t is the hourly serving rate that the platform posts at hour t. For simplicity, we assume that the demand for rides is downward-sloping and iso-elastic:

$$Q_t = D_t(P_t) = \delta_t P_t^{-\epsilon}, \tag{2}$$

where ϵ is the constant demand elasticity. The demand shifter δ_t includes daily weather indices, such as precipitation and temperature.

The platform takes demand shifter δ_t and demand elasticity ϵ as given and chooses prices and assignments to balance the demand and supply of rides so as to maximize platform profit. Let s_t be the percentage of orders assigned to high-performing drivers at time t, where $s_t \in [0,1]$. The platform's choice of $(\boldsymbol{P},\boldsymbol{s})$ maximizes its own payoff:

$$\begin{aligned} \max_{(\boldsymbol{P},\boldsymbol{s})} \quad & r\sum_t P_t D_t(P_t) \\ \text{s.t.} \quad & D_t(P_t)s_t \leq \lambda_t^H \mathcal{N}_t^H(\boldsymbol{P},\boldsymbol{s};\boldsymbol{\theta}) \\ & D_t(P_t)(1-s_t) \leq \lambda_t^L \mathcal{N}_t^L(\boldsymbol{P},\boldsymbol{s};\boldsymbol{\theta}). \end{aligned} \tag{3}$$

Here, r represents the platform's commission rate, while drivers receive a $1-r$ portion of the ride fare. $D_t(P_t)$ is the demand for rides, measured in earning hours, and $\mathcal{N}_t^\tau$ represents the total number of working hours (active app hours) that the drivers provide. λ_t^τ is the technological constraint restricting the relationship between working hours and earning hours, where $\lambda_t^\tau \in [0,1]$. For example, $\lambda_t^\tau = 0.5$ means that for every 15 minutes driving with a rider, a typical driver spends another 15 minutes on pickup, payment, etc. If $\lambda_t^\tau = 1$, there is no time spent on pickup.[20] We have idle drivers waiting for trip requests when one of the two inequalities is unbounded. In our empirical analysis, we set the commission rate r equal to 20%.[21]

Given the choice of prices and assignments $(\boldsymbol{P}, \boldsymbol{s})$, the platform effectively determines the sequence of wages $(\boldsymbol{W^H}, \boldsymbol{W^L})$. Each high- or low-performing driver expects to receive a wage rate

$$\begin{aligned} W_t^H &= (1-r)P_tD_t(P_t)s_t\frac{1}{N_t^H}, \\ W_t^L &= (1-r)P_tD_t(P_t)(1-s_t)\frac{1}{N_t^L}, \end{aligned} \tag{4}$$

where $1-r$ is the revenue share that the driver receives, while s_t represents how the algorithm favors high-performing drivers (the proportion of orders assigned to high-performing drivers).

6 Estimation

6.1 Demand Estimation

We first estimate rider demand for service time for each hour h. We consider each hour a different market and aggregate our data to the day-hour level. Based on the logarithm of total earning time (Q_t) and the logarithm of average hourly ride fare (P_t), we estimate demand parameters as follows:

$$\log Q_t = \log \delta_h - \epsilon \log P_t + e_h. \tag{5}$$

Our demand estimation suffers from classic supply-demand endogeneity: the platform may set a higher price when there is a higher demand shock in the market, so our OLS estimates may be biased. Similar to Kalouptsidi (2014), we use the number of cars in competing ride-hailing companies on a given day as our supply-side instrumental variable. Suppose the hourly demand shock e_h is instantaneous, with an expected value of zero *ex ante*. In

[20] We obtain the value of the technological constraint from the data. We compute the driving time as a percentage of driver work time in each day-hour for both high- and low-performing drivers. Then, we calculate the maximum as the technological restriction.

[21] According to Platform X's IPO document, the national average commission rate is 20.9%. In our survey, most drivers suggest that the commission rate is about 20%. Therefore, we use $r = 0.2$ in our empirical analysis.

this case, the number of cars in competing ride-hailing companies is not correlated with hour-level demand shocks. On the other hand, the number of cars competitors operate is negatively correlated with the ride fare that the platform can charge. Therefore, the number of cars in competing ride-hailing companies is a valid instrument.

Table 10: Demand Estimation

Dependent Variables	ln(Service Hours)			
	(1) OLS	(2) OLS	(3) OLS	(4) IV
ln(Price)	-5.151***	-5.158***	-0.767***	-1.186**
	(0.0743)	(0.0737)	(0.152)	(0.553)
Rain		-0.0020	-0.0005	-0.0006
		(0.002)	(0.0007)	(0.0007)
Temperature		0.0127***	0.0094***	0.0098***
		(0.0033)	(0.0011)	(0.0012)
Constant	32.12***	32.06***	10.62***	12.69***
	(0.350)	(0.348)	(0.752)	(2.736)
Hour FE			Y	Y
Day of Week FE			Y	Y
Observations	744	744	744	744
R-squared	0.866	0.869	0.988	0.988

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 10 reports demand estimates for the city of study. Column (1) reports the estimates without fixed effects. Column (2) reports estimates with the weather as a demand shifter. Column (3) further includes day and hour fixed effects. Column (4) reports our IV estimates. After controlling for hourly fixed effects and day-of-the-week fixed effects, column (3) reports a demand elasticity of -0.767. The estimated demand elasticity is much smaller with fixed effects than the estimates in the naïve OLS regression. Our IV estimates in column (4) are similar to those with fixed effects in column (3), and show that when the hourly ride fare increases by 1%, the total demand for service time decreases by 1.2%. We use these IV estimates as the demand elasticity in our counterfactual analysis. Our estimated demand elasticity of -1.186 is comparable to the values estimated in the literature. For example, Frechette, Lizzeri and Salz (2019) estimate an elasticity of -1.225 for New York City's taxi market, while Cohen, Hahn, Hall, Levitt and Metcalfe (2016) rely on the surge pricing algorithm and estimate a smaller price elasticity for UberX (between -0.4 and -0.6).

6.2 Identification and Estimation of Supply Parameters

Our model with unobserved heterogeneity is point-identified using conditional choice probabilities in drivers' dynamic labor supply. Online Appendix I.1 contains the details of our identification arguments. We follow our identification argument closely in estimating the model and estimate the following structural parameters: $\theta \equiv (\{O_t^{\iota}\}, \kappa, \sigma)$. $\{O_t^{\iota}\}$ is the reservation value at each time t, which includes the average reservation value O_t and the unobserved heterogeneity $\eta_{\iota,t}$. κ is the warm-up cost of starting to work, and σ is the normalization term of the EVT1 errors (the scale parameter). We explain the details of our estimation procedure in Online Appendix I.2.

Table 11 shows the estimation results. The first row shows the estimated population density of each driver group. Three main driver groups clearly dominate: group 3 with probability 0.42, group 2 with probability 0.18, and group 4 with probability 0.18. The second row of Table 11 shows the probability of being high-performing for each driver group. Driver groups 2 and 4 are 96.5% and 93.4% likely to be high performing, respectively, while driver groups 2 and 4 have the lowest average reservation values.

Table 11: Estimation Results for Unobserved Heterogeneity

	Group 0	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Population density of each group	0.07	0.06	0.18	0.42	0.18	0.04	0.05
Probability of H-Type	76.7%	78.7%	96.5%	49.6%	93.4%	82.8%	81.0%
Average Reservation Value	46.2	45.6	36.5	50.9	40.6	45.1	44.8

Figure 3 shows the estimated reservation values with unobserved driver heterogeneity. The average estimated reservation value is 49 CCY. The black line shows the estimated reservation values for group 0 drivers. Reservation values are lowest during morning hours, around 25 CCY, and highest at late night, around 68 CCY. For context, the minimum hourly wage in the city under study was 18.5 CCY in 2018. From the estimated results, we can see that drivers have higher reservation values during incentivized hours between 10 AM—2 PM and 7 PM—5 AM. This helps explain why the ride-hailing platform wants to implement algorithmic preferential wage-setting and convince drivers to work more during those incentivized hours. In terms of driver heterogeneity, driver groups 2 and 4 exhibit low reservation values during midday and early night periods, respectively. The estimated warm-up cost is 124 CCY, or around 2.5 times the average hourly reservation value. The high warm-up cost helps explain why drivers usually choose to drive consecutive hours.

Table 11 shows that three main driver groups dominate: the benchmark drivers (group 3), drivers with low midday reservation values (group 2), and those with low early night

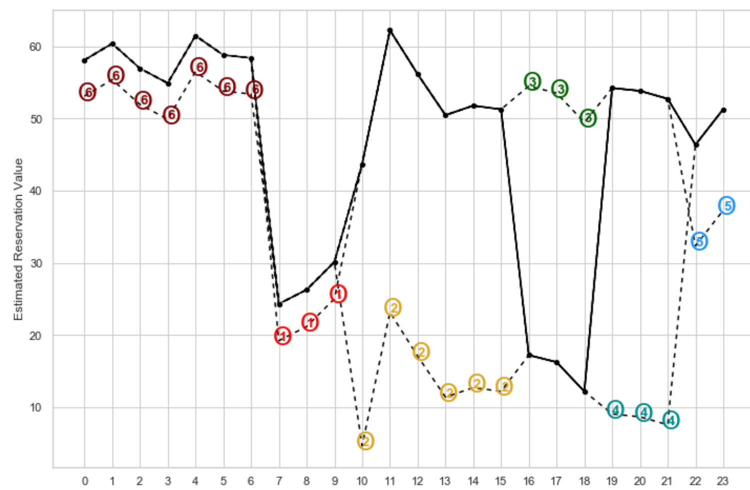

Figure 3: Estimated Reservation Values with Unobserved Driver Heterogeneity

reservation values (group 4). To better understand the estimated driver groups, we associate the observed driver characteristics with the respective driver groups. Because driver group 2 has a low midday reservation value and a high probability of being high-performing, observed high-performing drivers who choose to work midday are more likely to be in driver group 2. Similarly, observed high-performing drivers who work early nights are more likely to be in driver group 4. Based on this definition, we divide observed high-performing drivers into several groups according to their working hours. The *OnlyDay* group comprises drivers who work at least 2 consecutive hours in the daytime for at least 8 of 21 workdays.[22] The *OnlyNight* group refers to drivers working at least 2 consecutive hours at night for at least 8 of 21 workdays. Drivers who satisfy both criteria belong to the *BothDay&Night* group. The three groups (*OnlyDay*, *OnlyNight*, *BothDay&Night*) are mutually exclusive. The remaining high-performing drivers are grouped into the *Rest* group.

Table 12 compares the driver characteristics of different driver groups in the data, with several interesting findings. First, the *OnlyDay* group has a higher proportion of female drivers (3.5%) than the *OnlyNight* group (1.2%). Second, the average age in the *OnlyDay* group (38.3) is higher than that of the *OnlyNight* group (36.5). Third, non-locals are more likely to be high-performing drivers. For instance, 76% of the drivers in the *BothDay&Night* group are non-local, as compared to 62% of those in the *OnlyNight* or *OnlyDay* group. Therefore, the results suggest that driver groups 2 and 4 are more likely to consist of older, non-local, and female drivers. On the other hand, driver group 3 is more likely to include

[22]This is consistent with our definition of high-performing drivers, who we require to satisfy the condition in at least 8 of the 21 workdays.

younger, local, and male drivers. By associating observed driver demographics with estimated driver groups through unobserved heterogeneity, the counterfactual analysis allows us to better understand which individuals may benefit or suffer from the implementation of a preferential algorithm.

Table 12: Observed Driver Characteristics

Type	Low-Performing	High-Performing			
Group		Only Night	Only Day	Both Day & Night	Rest
Female	3.5%	1.2%	3.5%	1.7%	1.6%
Age	37.4	36.5	38.3	36.8	36.4
Non-local	53%	62%	63%	76%	62%
# of Drivers	16,392	3,073	6,659	11,939	2,041

Lastly, we validate our model by checking its goodness of fit. Specifically, we examine whether the simulated values fit the observed CCPs well. Online Appendix J shows the model's goodness of fit. Figure J.1 shows that overall, the simulated values do indeed fit the observed CCPs well.

7 Counterfactual Analysis

We conduct two main counterfactual experiments. First, we show the welfare effects of eliminating the preferential algorithm. To this end, we break down how using a preferential algorithm affects the balance between demand and supply over time, as well as how it influences passenger fares. Second, we investigate what factors determine the effectiveness of the preferential algorithm.

7.1 Elimination of Preferential Algorithm ("Fair" Pay)

In the first counterfactual analysis, we study changes in welfare if the preferential algorithm based on work schedules is eliminated. In this case, orders would be randomly assigned to nearby active workers. Effectively, the hourly wage each driver earns will become

$$\widetilde{W}_t = \frac{\eta P_t D_t(P_t)}{N_t}.$$

Given the new sequence of hourly wages $\{\widetilde{W}_t\}$, drivers solve the unconstrained dynamic discrete choice problem for each hour t:

$$
\begin{aligned}
U_{1t} &= \underbrace{\widetilde{W}_t}_{\substack{\text{non-preferential}\\ \text{wage rate}}} + \sigma \cdot \epsilon_{1t}, \\
U_{0t} &= \underbrace{O_t^{\iota}}_{\text{outside option value}} + \sigma \cdot \epsilon_{0t},
\end{aligned}
$$

where we have replaced the preferential wage rates W_t^H and W_t^L by the "fair" rate $\widetilde{W}_t$.

First, we show how the platform leverages cross-time labor supply elasticity using the preferential algorithm. We maintain the ride fares at the same level as when using a preferential algorithm, but we eliminate the wage differential among drivers. While keeping the ride fares unchanged, eliminating the preferential algorithm will decrease labor supply, resulting in labor shortages for most hours. Panel (a) of Figure 4 shows the level of labor shortage.[23] Upon eliminating the preferential algorithm, we can see a severe labor shortage during midday and in the late afternoon. Panel (b) shows the wage differential between high- and low-performing drivers when the preferential algorithm is present. A high wage differential in a particular hour means a high incentive wage for that hour. These results show that the relation between wage differentials and labor shortages is not one-to-one. For example, there is a severe labor shortage at 1 PM and 2 PM, even though the platform does not directly provide high incentive wages specifically at 1 PM and 2 PM, instead doing so from 5 AM to 8 AM. Panel (a) shows that the labor shortage is very mild in the early morning from 5 AM to 7 AM. Therefore, the platform does not necessarily provide direct high incentive wages to mitigate labor shortages in a given hour, but instead smooths out the payment of high incentive wages by leveraging the variations in demand elasticity and the differing reservation values of drivers over time.

To further illustrate the idea of cross-time labor supply elasticity, we eliminate the wage differential between high- and low-performing drivers in only one hour (the treatment hour) and study the implied elasticity of labor supply. More precisely, we calculate the elasticity as

$$
\mathcal{E}_t^{\tau}(h) = \frac{\left(N_t^{\tau}(\widetilde{\boldsymbol{W}}^H, \widetilde{\boldsymbol{W}}^L) - N_t^{\tau}(\boldsymbol{W}^H, \boldsymbol{W}^L)\right) / N_t^{\tau}(\boldsymbol{W}^H, \boldsymbol{W}^L)}{(\widetilde{W}_h^{\tau} - W_h^{\tau})/W_h^{\tau}},
$$

where h is the chosen hour during which we eliminate the wage differential between high-

[23]To better illustrate the results, we normalize the maximum labor shortage and the maximum wage differential to 1 in Figure 4.

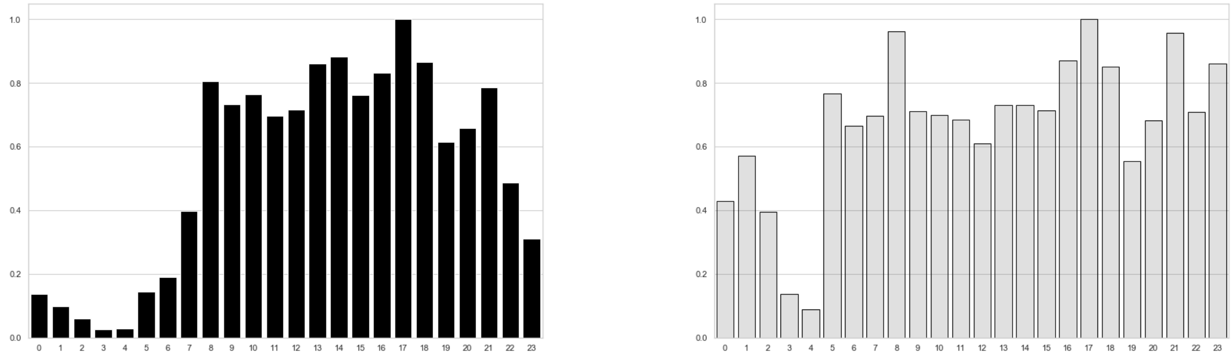

(a) Labor Shortage

(b) Wage Differential (W_t^H - W_t^L)

Figure 4: Illustration of Leveraging Cross-time Labor Supply Elasticity

and low-performing drivers.

Figure 5 shows the absolute value of the elasticity of labor supply corresponding to the elimination of the wage differential at 12 PM. The blue line represents the low-performing drivers, while the red line represents the high-performing ones. The former are much more responsive to the elimination of the wage differential than the latter. This is because, on the one hand, high-performing drivers' labor supply is inelastic in all hours (less than 0.7). On the other, low-performing drivers' labor supply elasticities are higher than 0.9 in all hours and even higher than 1 in the hours near the treated hour. The absolute elasticity value generally decreases for hours further away from the treatment hour. This is because there is a high warm-up cost for starting to work, and adjacent hours of the treatment hour will thus be affected more. However, the absolute elasticity value does not monotonically decrease with respect to the distance to the treatment hour because of the variation in reservation values across the different hours of the day. Given that multiple-hour labor supply responds to the wage differential at one particular hour, the platform can strategically choose when to provide high incentive wages. In Online Appendix K, we replicate this exercise by changing the treatment hour from 7 AM to 6 PM.

Next, we show the welfare effects of eliminating the preferential algorithm. We study two scenarios: first, we eliminate the wage differential among drivers but maintain the ride fares at the same level as when using a preferential algorithm. Second, we allow the platform to re-optimize its pricing strategy and change ride fares when the preferential algorithm is

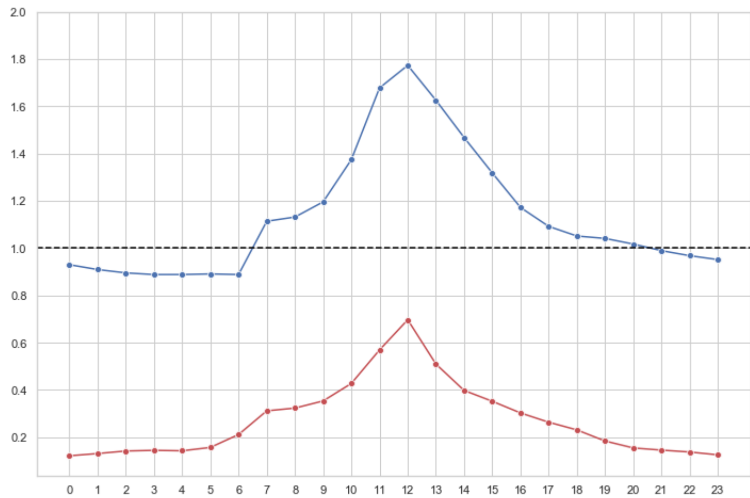

Figure 5: Elasticity of Labor Supply When Eliminating Wage Differential at 12 PM

eliminated:[24]

$$\begin{aligned} \max_{\vec{P}} \quad & (1-\eta)\sum_t P_t D_t(P_t) \\ \text{s.t.} \quad & D_t(P_t) \leq \widetilde{\lambda}_t \mathcal{N}_t(\boldsymbol{W_t};\boldsymbol{\theta}) \end{aligned} \tag{6}$$

Using the estimated parameters, we solve for the new equilibrium outcome if the platform can no longer implement algorithmic preferential wage-setting based on the work schedule. Then, we calculate the changes in platform revenue, consumer surplus, and driver surplus by comparing the outcome without the preferential algorithm to the outcome with it. We calculate consumer surplus as $\sum_t \int_{P_t}^{\infty} \delta_t x^{-\epsilon} dx$ and the driver surplus of each schedule j as

$$EV_0^j = \sigma \left[\ln\left(\exp((\widetilde{W}_1 - \kappa + \beta EV_{12})/\sigma) + \exp((O_1 + \beta EV_{02})/\sigma) \right) + \gamma \right],$$

where EV_0^j represents the expected value from choosing each work schedule type j.[25]

Table 13 shows the results. In the first scenario, on the one hand, eliminating the preferential algorithm will result in a massive loss for both the platform and the rider because of a driver shortage. On the other, drivers enjoy more flexibility in choosing a work schedule under "fair" pay. Hence, there will be a 0.14% increase in driver surplus. High-performing drivers suffer a loss of 0.63% because there is no longer a bonus for being high-performing.

[24]Note that the two feasibility constraints in equation 3 become one because all drivers have the same likelihood of receiving a task; $\widetilde{\lambda}_t$ is the technology restriction without algorithmic preferential wage-setting. Under "fair" pay, there is only one group, so $s = 1$.

[25]Note that consumer surplus $\sum_t \int_{P_t}^{\infty} \delta_t x^{-\epsilon} dx = \sum_t \frac{\delta_t}{\epsilon-1}(P_t)^{1-\epsilon} = \frac{1}{(\epsilon-1)(1-\eta)} \times$ platform revenue. When ride fares are unchanged, the total number of riders served equals $\min\{D_t(P_t), \widetilde{\lambda}_t \mathcal{N}_t(\boldsymbol{W_t};\boldsymbol{\theta})\}$.

Low-performing drivers see an increase in hourly wage, hence a 0.69% increase in surplus. In aggregate, the total surplus will decrease by 7.16% if we eliminate the preferential algorithm.

Table 13: Changes in Welfare

Changes in	The first scenario: Maintain ride fares unchanged	The second scenario: The platform re-optimizes ride fares
Platform revenue	-12.16%	-1.42%
Consumer surplus	-12.16%	-1.42%
Driver surplus	0.14%	0.49%
Total surplus	-7.16%	-0.64%
Decomposition of Per-Driver Surplus		
High-performing driver (non-switcher)	-0.63%	-0.16%
Low-performing driver (non-switcher)	0.69%	0.99%
Switcher (from H-type to L-type)	3.51%	3.81%
Change in Probability of being H-type	-11.48%	-9.98%

Notes: We calculate changes in welfare by measuring the results without a preferential algorithm minus the results with a preferential algorithm. In the first scenario, ride fares without a preferential algorithm are held the same as ride fares with a preferential algorithm. In the second scenario, the platform re-optimizes its pricing strategy without a preferential algorithm.

In the second scenario, the platform will re-optimize its pricing strategy, increasing ride fares to reduce the driver shortage that we see in the first scenario after eliminating the preferential algorithm. Ride fares, on average, increase by 7.79% across different hours of the day, with the highest increase being 12.16%. This increase in fares helps alleviate the driver shortage. As a result, the losses of the platform and riders will be smaller than in the first scenario, resulting in a total 1.42% decrease in surplus. Driver surplus will increase further, because drivers benefit in the second scenario from the increased ride fares. Total driver surplus will hence increase by 0.49% if we eliminate the preferential algorithm. Low-performing drivers have a 1% increase in surplus because they benefit from more flexibility in working and a higher ride fare. Regarding the extensive margin, the probability of being high-performing decreases by 11.48 percentage points in the first scenario and 9.98 percentage points in the second. After we eliminate the preferential algorithm, the probability of being a high-performing driver slightly increases in the second scenario, as compared to the first. The total surplus will decrease by 0.64% if we eliminate the preferential algorithm.

Lastly, we look at how different groups of drivers are affected if we eliminate the preferential algorithm. We characterize driver groups by the unobserved heterogeneity types described in Section 6.2.

Table 14 shows the results. First, we can see both winners and losers from eliminating the preferential algorithm. From the elimination of the preferential algorithm, driver groups 2 and 4 experience a decrease in their surpluses of 0.36% and 0.22%, respectively, in the first scenario, while all other driver groups experience an increase. The welfare loss of driver groups 2 and 4 occurs because they are more likely to be high-performing, and high-performing drivers will no longer earn extra hourly wages without the preferential algorithm. Previous results showed that drivers in groups 2 and 4 are likely to be high-performing with probabilities of 96.5% and 93.4%, respectively. Previous results also showed that among high performers, female or older drivers are more likely to fall into groups 2 and 4. Therefore, the counterfactual results indicate that women and older drivers who choose to be high-performing are more likely to suffer from eliminating the preferential algorithm. The general effect for female drivers is ambiguous, because women are also more likely to be low-performing, with a larger welfare gain from the elimination of the preferential algorithm. Non-locals are more likely to suffer a welfare loss if we eliminate the preferential algorithm because they are more likely to be high-performing. All other driver groups (younger, local, male) will benefit from its elimination.

Table 14: Change in Driver Surplus, by Groups of Drivers

Changes in Driver Surplus	Driver Group						
	Group 0	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Panel I: Maintain ride fares unchanged							
Total	0.08%	0.05%	-0.36%	0.20%	-0.22%	0.00%	0.07%
H-Schedule	-0.41%	-0.43%	-0.50%	-0.14%	-0.44%	-0.38%	-0.42%
L-Schedule	0.35%	0.38%	0.86%	0.12%	0.57%	0.37%	0.36%
Panel II: The platform re-optimizes ride fares							
Total	0.29%	0.28%	-0.02%	0.22%	0.08%	0.23%	0.29%
H-Schedule	-0.14%	-0.15%	-0.17%	-0.04%	-0.13%	-0.12%	-0.14%
L-Schedule	0.54%	0.58%	1.19%	0.16%	0.86%	0.57%	0.56%

Notes: We calculate changes in welfare by results without a preferential algorithm minus results with a preferential algorithm. We characterize driver groups by unobserved heterogeneity.

In summary, the platform benefits from implementing a preferential algorithm by leveraging the cross-time labor supply elasticity. Without a change in ride fares, eliminating the preferential algorithm results in a significant welfare loss for both the platform and riders due to driver shortages. Drivers, on the other hand, experience an increase in surplus because of the increased flexibility when choosing their work schedule. If we allow the platform to re-optimize pricing, ride fares will significantly increase without a preferential algorithm. However, the driver shortage will also be mitigated, and welfare loss will be smaller for both the platform and riders. Drivers, meanwhile, will have an even greater increase in welfare because of increased ride fares. Among the different groups of drivers, those who are male, young, and local are more likely to benefit from the elimination of the preferential algorithm. Older drivers, conversely, are likely to experience a welfare loss. The net effect for female drivers is ambiguous, with a welfare loss for high-performing female drivers and a welfare gain for low-performing ones.

7.2 Factors Determining Preferential Algorithm Effectiveness

To further investigate what factors determine the effectiveness of the preferential algorithm, we conduct counterfactuals by alternating key structural parameters. We specifically focus on the demand elasticity ϵ and warm-up cost κ. Table 15 shows the results when we alter the value of demand elasticity. When demand is more elastic, the platform benefits more from utilizing the cross-time labor supply elasticity by implementing the preferential algorithm. Therefore, in column (1) of Table 15, we see a larger increase in platform revenue, from 1.44% to 2.89%, if the platform implements the preferential algorithm. On the other hand, drivers suffer less from the preferential algorithm if demand elasticity increases. Total driver surplus will decrease by 0.32% when demand is more elastic, compared to a decrease of 0.49% when demand is less elastic. Our intuition is that when demand is very elastic, the platform is less willing to incentivize labor supply by increasing ride fares; otherwise, there would be a large decrease in rider demand. Therefore, drivers will experience a smaller increase in wage rate when the platform eliminates the preferential algorithm. Equivalently, this means that drivers will experience a smaller decrease in wage rate, and hence driver surplus, when the platform implements a preferential algorithm. Column (4) of Table 15 confirms this intuition by showing that the average decrease in wage is smaller (5.03% versus 7.26%) when demand is more elastic. As a result, the loss of low-performing drivers decreases from 0.98% to 0.17%.

Next, we examine the effect of the warm-up cost κ. Table 16 shows the results when we vary the value of this factor. When the warm-up cost is higher, the platform must pay higher wages to incentivize drivers to work. Hence, it is more profitable for the platform to avoid paying such high incentive wages by implementing the preferential algorithm. On the other hand, saving these high incentive wages reduces the ride fare, and hence more riders can be served. Serving more riders also generates more hourly revenues for the drivers. As a result, the loss in driver surplus from the preferential algorithm will be smaller when the

Table 15: Varying the Value of Demand Elasticity ϵ

Demand Elasticity	Changes in (With - Without)			
	Platform Revenue/ Consumer Surplus	Driver Surplus	Driver Surplus (Low-performing)	Average Wage
Benchmark	1.44%	-0.49%	-0.98%	-7.26%
$\epsilon \times 1.1$	2.13%	-0.47%	-0.52%	-6.55%
$\epsilon \times 1.2$	2.60%	-0.40%	-0.29%	-5.78%
$\epsilon \times 1.3$	2.89%	-0.32%	-0.17%	-5.03%

warm-up cost is larger. Column (4) of Table 15 confirms this intuition by showing that the change in the number of served riders is greater (9.76% versus 9.55%) when the warm-up cost is larger.[26]

Table 16: Varying the Value of Warm-up Cost κ

Warm-up Cost	Changes in (With - Without)			
	Platform Revenue/ Consumer Surplus	Driver Surplus	Driver Surplus (Low-performing)	Consumers Served
Benchmark	1.44%	-0.49%	-0.98%	9.55%
$\kappa \times 1.1$	1.45%	-0.49%	-0.93%	9.64%
$\kappa \times 1.2$	1.46%	-0.49%	-0.86%	9.71%
$\kappa \times 1.3$	1.47%	-0.48%	-0.79%	9.76%

To summarize, the platform benefits more from implementing a preferential algorithm when the demand is more elastic or when the warm-up cost is greater. Meanwhile, the loss of driver surplus with a preferential algorithm is also smaller given these two conditions.

8 Conclusion

In recent years, the gig economy has rapidly accelerated, with the ride-hailing market standing out as a prominent example. It is often assumed that this sector, or any form of piece-rate

[26]In the second case, when we alter the value of the warm-up cost κ, showing the change in average wage will not directly reveal how driver surplus changes, because driver utility depends on both the warm-up cost and the average wage. Instead, when demand elasticity is fixed in this case, the surplus of low-performing drivers will monotonically increase with respect to the number of riders served. Similarly, in the first case, when we alter the value of demand elasticity η, showing the change in the number of riders served will not directly reveal how driver surplus changes, because the number of riders served is determined by both the labor supply decision and demand elasticity. In this case, when the warm-up cost is fixed, the surplus of low-performing drivers will monotonically increase with respect to the average wage rate. This is why we report different variables in the last columns of Tables 15 and 16.

work, would address the issue of job flexibility penalties. However, such penalties can still arise when platforms' algorithms preferentially assign orders to drivers based on their historical work patterns. Our paper is the first to empirically study such an algorithmic assignment and its impact on worker behavior and welfare. Using rich transaction data from a leading ride-hailing company in Asia, we first document significant wage differentials due to work schedules between high-performing drivers who work long and consecutive hours and their low-performing counterparts. We show that three main factors drive the wage differential: high-performing drivers are given more ride requests per hour, wait fewer minutes for each request, and receive more requests from riders with lower cancellation rates. Next, we exclude alternative explanations for the wage differentials, such as drivers strategically choosing where to work, strategically selecting and canceling orders, and driving faster. The large wage differential we identify is mainly due to algorithmic assignment, which penalizes low-performing drivers.

We then propose a dynamic equilibrium model of a ride-hailing market to quantify the welfare effects of such a preferential assignment algorithm. Results show that platform revenues will decrease by 12.16%, while the total surplus will decrease by 7.16%, if ride fares are held constant when we eliminate the preferential assignment algorithm. The probability of drivers being high-performing will decrease by 11.48% without a preferential algorithm. For the switchers, driver surplus will increase by 3.51%. If we allow the platform to re-optimize ride fares after eliminating the preferential algorithm, it will raise rider fares to re-balance demand and supply, resulting in minimal welfare loss. Moreover, an additional 10% of drivers would switch to flexible schedules. Among drivers, those who are young, male, and local benefit more from the elimination of the preferential algorithm. Lastly, our simulations show preferential algorithms benefit the platform more and hurt drivers less when rider demand is more elastic or when the warm-up cost is higher.

References

Aaronson, Daniel and Eric French, "The effect of part-time work on wages: Evidence from the social security rules," *Journal of Labor Economics*, 2004, *22* (2), 329–252.

Altonji, Joseph G and Christina H Paxson, "Labor supply, hours constraints, and job mobility," *Journal of Human Resources*, 1992, *27* (2), 256–278.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market," 2023. *Journal of Political Economy (forthcoming).*

Buchholz, Nicholas, "Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry," *The Review of Economic Studies*, 2022, *89* (2), 556–591.

Castillo, Juan Camilo, "Who benefits from surge pricing?," *Available at SSRN 3245533*, 2023.

Chen, Kuan-Ming, Claire Ding, John A List, and Magne Mogstad, "Reservation wages and workers' valuation of job flexibility: Evidence from a natural field experiment," Technical Report, National Bureau of Economic Research 2020.

Chen, M Keith, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen, "The value of flexible work: Evidence from Uber drivers," *Journal of Political Economy*, 2019, *127* (6), 2735–2794.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, "Using big data to estimate consumer surplus: The case of Uber," Technical Report, National Bureau of Economic Research 2016.

Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer, "The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers," *The Review of Economic Studies*, 2021, *88* (5), 2210–2238.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz, "Frictions in a competitive, regulated market: Evidence from taxis," *American Economic Review*, 2019, *109* (8), 2954–92.

Gaineddenova, Renata, "Pricing and efficiency in a decentralized ride-hailing platform," *Unpublished working paper*, 2022.

Goldin, Claudia, "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Haggag, Kareem, Brian McManus, and Giovanni Paci, "Learning by driving: Productivity improvements by New York City taxi drivers," *American Economic Journal: Applied Economics*, 2017, *9* (1), 70–95.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Kalouptsidi, Myrto, "Time to build and fluctuations in bulk shipping," *American Economic Review*, 2014, *104* (2), 564–608.

Liu, Tracy, Zhixi Wan, and Chenyu Yang, "The efficiency of a dynamic decentralized two-sided matching market," 2019. Available at SSRN 3339394.

Mas, Alexandre and Amanda Pallais, "Valuing alternative work arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti, "Economic shocks and civil conflict: An instrumental variables approach," *Journal of Political Economy*, 2004, *112* (4), 725–753.

Ming, Liu, Tunay I Tunca, Yi Xu, and Weiming Zhu, "An empirical analysis of market formation, pricing, and revenue sharing in ride-hailing services," 2019. Available at SSRN 3338088.

Rosaia, Nicola, "Competing platforms and transport equilibrium," *Unpublished working paper*, 2023.

Rysman, Marc, "Competition between networks: A study of the market for Yellow Pages," *The Review of Economic Studies*, 2004, *71* (2), 483–512.

_ , "The economics of two-sided markets," *Journal of Economic Perspectives*, 2009, *23* (3), 125–43.

Appendix

A Preferential Algorithm and Surge Pricing

The platform benefits from both surge pricing and a preferential algorithm, although their mechanisms differ. To illustrate the contrasting ways surge pricing and the preferential algorithm operate and to showcase their potential complementarity, we use our theoretical model to solve equilibrium outcomes in the following four scenarios:

1. Without surge pricing or a preferential algorithm
2. With only surge pricing
3. With only the preferential algorithm
4. With both surge pricing and the preferential algorithm

As discussed in section 2, we consider a scenario with two time periods, denoted as t_1 and t_2. For both periods, we the demand to be $P_t^d = 10 - q$. During time period t_1, drivers have a reservation value of 0. During time period t_2, drivers have positive and heterogeneous reservation values, and the supply curve is defined as $P_{t_2}^s = q$.

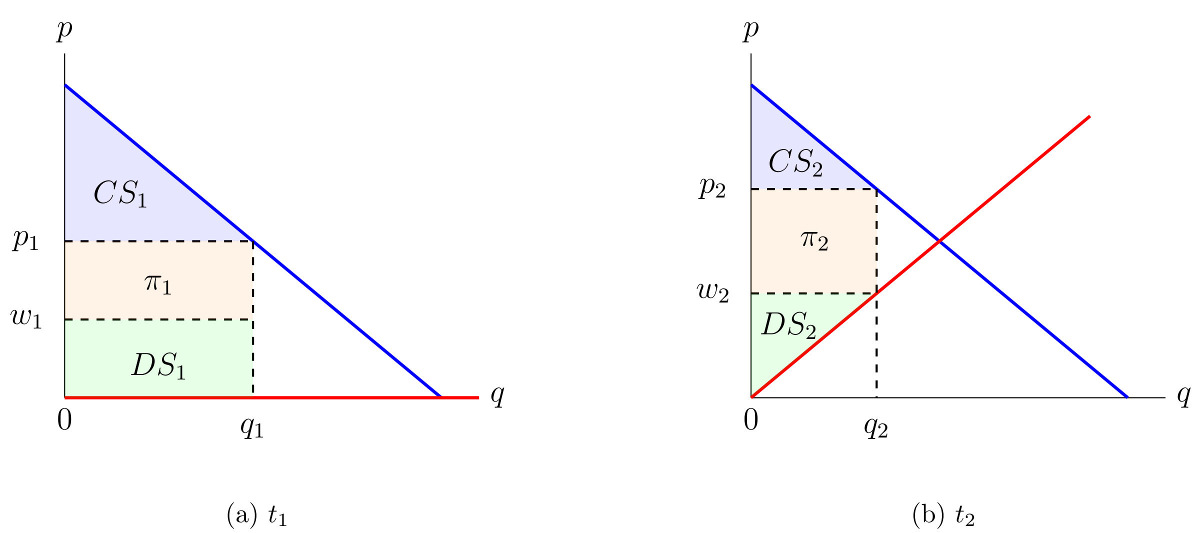

Figure A.1: Case 1, without surge pricing and preferential algorithm

In the baseline case, with neither surge pricing nor a preferential algorithm, the platform sets a price p_t for time period t, and drivers earn $w_t = (1 - \eta) * p_t$, where η represents the fractional commission fee. The platform's earnings in this case are $\eta * p_t * q_t$. For this numerical example, we assume $\eta = 0.5$. Figure A.1 shows the results for case 1. In both periods, the platform chooses the optimal ride fare p_t to maximize its profit. For the numerical example, the platform sets ride fares at $p_1^* = 5$ and $p_2^* = 6.67$, leading to wage rates of $w_1^* = 2.5$ and $w_2^* = 3.33$ for the drivers. The blue area in Figure A.1 represents the consumer surplus, the orange area represents the platform's profit, and the green area shows the drivers' surplus.

Figure A.2 shows the results for case 2, in which only surge pricing is implemented. With surge pricing, the platform captures the entire consumer surplus while offering a constant wage rate to the drivers in each period.[27] Consequently, in this scenario, the platform optimizes the wage rate to pay the drivers. In the given numerical example, the optimal wage rate is set at $w_1^* = w_{min} = 2$,[28] and $w_2 = 3.33$. The implementation of surge pricing results in a complete elimination of consumer surplus in both periods.

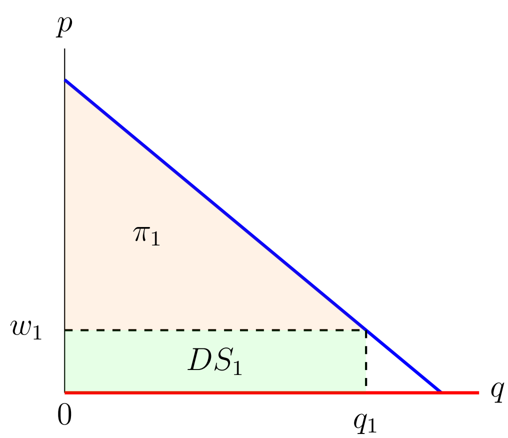

(a) t_1

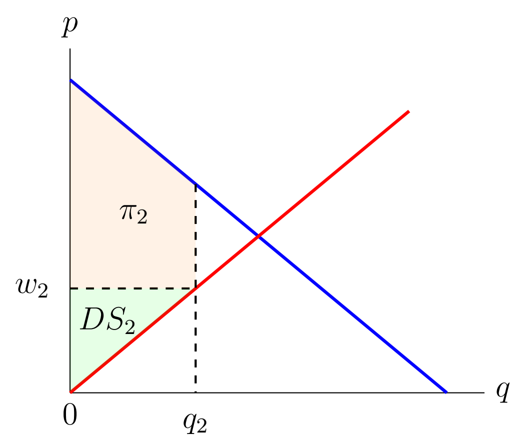

(b) t_2

Figure A.2: Case 2, with only surge pricing

Figure A.3 shows the results for case 3, in which only the preferential algorithm is implemented. With a preferential algorithm, the platform communicates to drivers that if they work during time period t_2, they will be given priority and receive an order during time period t_1. Hence, the platform can motivate drivers to work during time period t_2 without

[27]Alternatively, we can use the same wage scheme as in case 1, where drivers receive a constant percentage of the ride fare as their payment. This would lead to a higher surplus for the drivers. However, to highlight the key trade-off and to simplify the model here, we opt for a constant wage rate in this context.

[28]The platform is unable to further decrease the wage rate at t_1 because of the minimum wage requirement, as explained in 2. Otherwise, the platform would charge $w_1^* = 0$ at t_1, resulting in no driver surplus during that period. Consequently, the implementation of the preferential algorithm on top of surge pricing would no longer be feasible. In this particular scenario, to enable a meaningful comparison between case 2 and case 4, it is necessary to ensure a positive drivers' surplus at t_1.

offering extra incentive wages. As a result, the labor supply curve is shifted outwards at t_2, as illustrated by the dashed red line. In this scenario, the optimal prices are $p_1^* = 5$ and $p_2^* = 5$, leading to $w_1^* = 2.5$ and $w_2^* = 2.5$, with the number of orders served being $q_1^* = 5$ and $q_2^* = 5$. However, with a wage rate of $w_2^* = 2.5$, some drivers actually earn a negative surplus at t_2. For drivers falling within the range of q_2' to q_2, the wage at t_2 is insufficient to cover their reservation values. Despite this, they are willing to work during this time period because they anticipate earning a positive surplus at t_1 with prioritized order assignment. The platform's profit at t_2 is represented by the orange rectangle π_2 in panel (b) of Figure A.3, which is $(p_2 - w_2) * q_2$. The driver surplus is equivalent to the green area $DS_2 - DS_2'$. The preferential algorithm serves to assist the platform in extracting additional driver surplus.

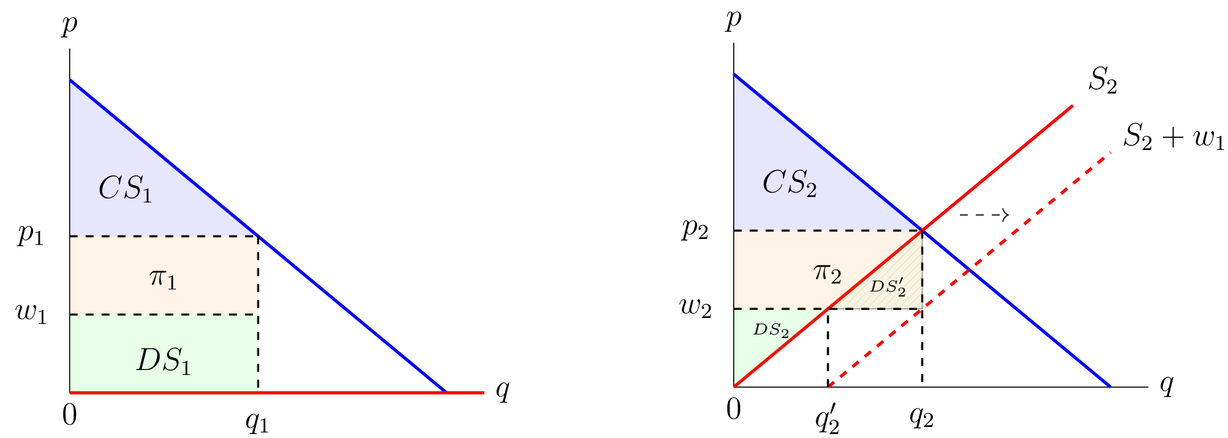

(a) t_1 (b) t_2

Figure A.3: Case 3, with only preferential algorithm

Figure A.4 shows the results for case 4, with both surge pricing and the preferential algorithm. First, similar to case 2, surge pricing allows the platform to capture the entire consumer surplus. Additionally, in case 4, the platform further leverages the preferential algorithm to incentivize drivers to work more. In Figure A.4, the platform's profit is denoted as π_1 and π_2, respectively, for time periods t_1 and t_2, while driver surplus is DS_1 during time period t_1 and $DS_2 - DS_2'$ during time period t_2.

We proceed to compare the equilibrium outcomes in each scenario. Table A.1 provides a summary of the equilibrium profit of the platform, the ride fare, the wage rate, and the quantity served for each scenario. On the other hand, Table A.2 shows the consumer surplus, driver surplus, and total surplus for the respective scenarios. We refer to the scenario with neither the preferential algorithm nor surge pricing as the benchmark case (case 1). When comparing the results of case 2 with the benchmark case, we can see that by implementing

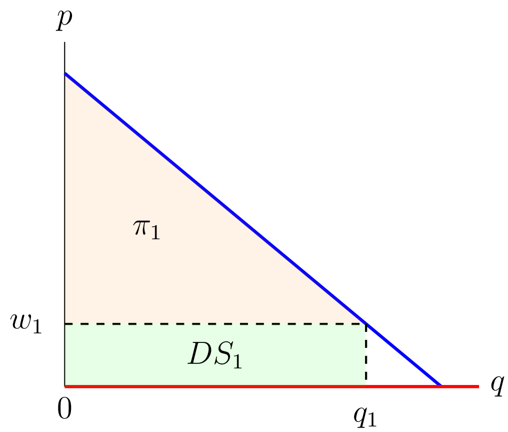

(a) t_1

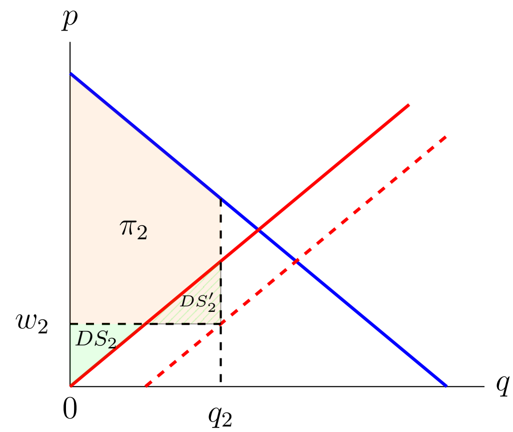

(b) t_2

Figure A.4: Case 4, with both surge pricing and preferential algorithm

surge pricing, the platform serves more consumers during t_1. While drivers experience a lower wage rate at t_1, the total drivers' surplus increases from 12.5 to 16 due to the higher number of orders served. Consequently, with the implementation of surge pricing, although the entire consumer surplus is captured by the platform, drivers also benefit from a higher driver surplus.

Table A.1: Equilibrium Outcomes

	Platform's profit		Price		Wage rate		Quantity	
	t_1	t_2	t_1	t_2	t_1	t_2	t_1	t_2
Case 1: No algorithm, no surge	12.50	11.11	5.00	6.67	2.50	3.33	5.00	3.33
Case 2: Only surge pricing	32.00	16.67	/	/	2.00	3.33	8.00	3.33
Case 3: Only algorithm	12.50	12.50	5.00	5.00	2.50	2.50	5.00	5.00
Case 4: Both surge and algorithm	32.00	24.00	/	/	2.00	2.00	8.00	4.00

Table A.2: Comparing Surplus

	Driver Surplus		Consumer Surplus		Total Surplus
	t_1	t_2	t_1	t_2	$t_1 + t_2$
Case 1: No algorithm, no surge	12.50	5.56	12.50	5.56	59.72
Case 2: Only surge pricing	16.00	5.56	0.00	0.00	70.22
Case 3: Only algorithm	12.50	0.00	12.50	12.50	62.50
Case 4: Both surge and algorithm	16.00	0.00	0.00	0.00	72.00

When comparing case 3, where only the preferential algorithm is implemented, with the benchmark case, we can see that the platform now charges lower ride fares and pays lower wage rates during t_2. Consequently, the driver surplus is reduced upon the introduction of the preferential algorithm. However, consumers also benefit from the lower ride fares during t_2, resulting in an increase in consumer surplus.

Lastly, in case 4, we examine the outcomes when both surge pricing and a preferential algorithm are implemented. When comparing case 4 with case 2, we observe that the platform not only captures the entire consumer surplus but also further extracts driver surplus by introducing the preferential algorithm. As a consequence, the total driver surplus decreases from 21.56 to 16. On the other hand, when comparing case 4 with case 3, we observe that after implementing the preferential algorithm, the additional introduction of surge pricing actually leads to an increase in driver surplus. This is because the platform serves more consumers, leading to drivers benefiting from the increased demand. Consequently, the total driver surplus increases from 12.5 to 16 between case 3 and case 4.

In summary, the findings show that both surge pricing and the preferential algorithm help increase the platform's profit, but they operate through distinct mechanisms. By implementing surge pricing, the consumer surplus is reduced as compared to the baseline scenario; however, both the platform's profits and the drivers benefit from this strategy. On the other hand, implementing the preferential algorithm may decrease driver surplus as compared to the baseline scenario, but it also leads to improved profitability for the platform and greater benefits for consumers. The results from case 4 demonstrate the complementarity between surge pricing and the preferential algorithm. When both methods are implemented, the platform's profit is the highest, and the total surplus is also maximized. This highlights the strong synergy between surge pricing and the preferential algorithm in achieving the best overall outcomes for the platform. However, this combination also results in a significant distributional effect, leading to a reduction in both consumer and driver surplus when compared to the baseline scenario.

Online Appendices (Not for Publication)

DRIVING THE DRIVERS: ALGORITHMIC ASSIGNMENT IN RIDE-HAILING

by Yanyou Chen, Yao Luo and Zhe Yuan

B Multi-Homing versus Single-Homing

In our main analysis, we focus on Platform X, because it accounts for more than 90% of China's mainland ride-hailing market. Nonetheless, there is a concern that drivers may switch between working for different platforms if they pay different hourly wages. To address such concerns, we document the number of vehicles/drivers that are multi-homed versus those that are single-homed in our data. First, we look at the number of vehicles that are multi-homed based on registration data. Panel (a) of Table B.1 shows that 85% of vehicles are registered to only one platform, while only 1.8% of vehicles are registered to more than two platforms. Therefore, multi-homing is not very common, based on vehicle registration information. Then, we look at how common multi-homing is directly from actual transactions. Panel (b) of Table B.1 shows that among all the vehicles that conducted business in December 2018, 92.5% used a single platform and never switched to another platform within the month. Only 0.3% of vehicles used more than two platforms in this month. The evidence shows that the majority of vehicles/drivers are single-homed.

Table B.1: Multi-Homing versus Single-Homing

Number of Registered Platforms	Number of Vehicles	Percent
1	86,422	84.6%
2	13,838	13.5%
3	1,866	1.8%

(a) Based on Vehicle Registration Data

Number of Used Platforms	Number of Vehicles	Percent
1	49,213	92.5%
2	3,836	7.2%
3	141	0.3%

(b) Based on Transactional Data

Among the multi-homed drivers, we further study how these drivers switch between different ride-sharing platforms, as well as calculating the number of multi-homed and single-homed drivers within a whole day, based on actual transactions. In any given day of December, only about 1% of drivers used more than one platform. This suggests that drivers in our data are mostly single-homed and rarely switch between platforms.

C Additional Data Description

For the working dataset, we are interested in driver operation and wage information, and we construct several important variables for each driver-hour:

- **Earning time** is the trip duration, measured as the amount of time a driver spends with the rider. A driver can transport riders and collect revenue only during their earning time.
- **Drive Distance** measures the distance over which a driver serves a rider in an hour.
- **Driver's Hourly Wage** measures the revenue of a driver in an hour.[29] Given that the platform fee is around 20% of revenue, the driver income is roughly 80% of the ride fare.
- **Pickup Time** measures the time a driver spends on the way to pick up riders.
- **Idle Time** is the time a driver spends waiting for orders in an hour, given by the following relationship: Idle time = 60 - Work time - Pickup time.
- **Number of Orders** measures the number of orders a driver receives in an hour.[30]

Next, we discuss how we construct a driver-hour level dataset from the driver-rider-order level dataset. First, We keep all orders with a departure and arrival in the urban area (eight districts) within the city, and drop orders with a price of zero, those with a price above 200, or those that span over four hours. In total, we drop less than 0.5% of the observations. Second, following Chen et al. (2019), we define a driver as working during an hour t if he works at least ten minutes out of that hour. At night (22PM–6AM), when orders are sparse, we define a driver as working during hour t if they work at hour $t-1$ and hour $t+1$. All a driver's working hours comprise their work schedule. Third, we match order to Hour. Suppose an order spans x hours. We divide this order into x sub-orders, with each sub-order corresponding to an hour. The hourly wage rate and driving distance are defined as proportional to each hour. For instance, suppose an order starts at 8:50 and finishes at 9:20, yielding a revenue of 60 CCY. We say that $\frac{10}{10+20} = \frac{1}{3}$ of the order belongs to the 8 AM operations, and the rest contributes to the 9 AM operations. By doing so, we divide this order into two sub-order operations: The driver drives 10 minutes and earns 20 CCY at 8

[29] Our definition is different from that of Chen et al. (2019), who define "wage rate" as a driver's total earnings in an hour, divided by minutes worked, multiplied by sixty. In other words, they study the wage rate when the driver is driving a rider, and we focus on the wage rate when the driver is active on the platform.

[30] In rare cases, an order may span several hours, which we attribute to the hour of departure.

AM and drives 20 minutes (10 miles) and makes 40 CCY at 9 AM. After matching orders to hours, we aggregate all sub-orders in an hour and obtain the given driver's earning time, ride prices, pickup time, idle time, and number of orders during this hour.

Table C.1 summarizes orders and transactions; the unit of observation is at the order level. There are around 15 million total order transactions in our sample period, with an average route length of 6.9 km and drive time of 17 minutes. The average price per order is 25.31 CCY (about $4 USD).

Table C.1: Summary Statistics: Orders and Transactions

Variable	Mean	Std. Dev.	Min	Max
Price	25.31	26.44	0	3,387
Drive Distance (km)	6.92	6.85	0	727
Drive Time (minutes)	17.36	13.14	0	1,458
Number of Observations		14,471,573		

Orders, Transactions, Precipitation, and Temperature. Figure C.1 reports the daily number of orders and transactions during our sample period (10 days of order data and 31 days of transaction data). We compare these with daily precipitation and average temperatures. From December 6 to 10, the precipitation increases and the temperature decreases, resulting in more ride orders (customer demand). However, the number of completed transactions across days remains the same throughout our sample period. We use information about precipitation and temperature in our demand estimation.

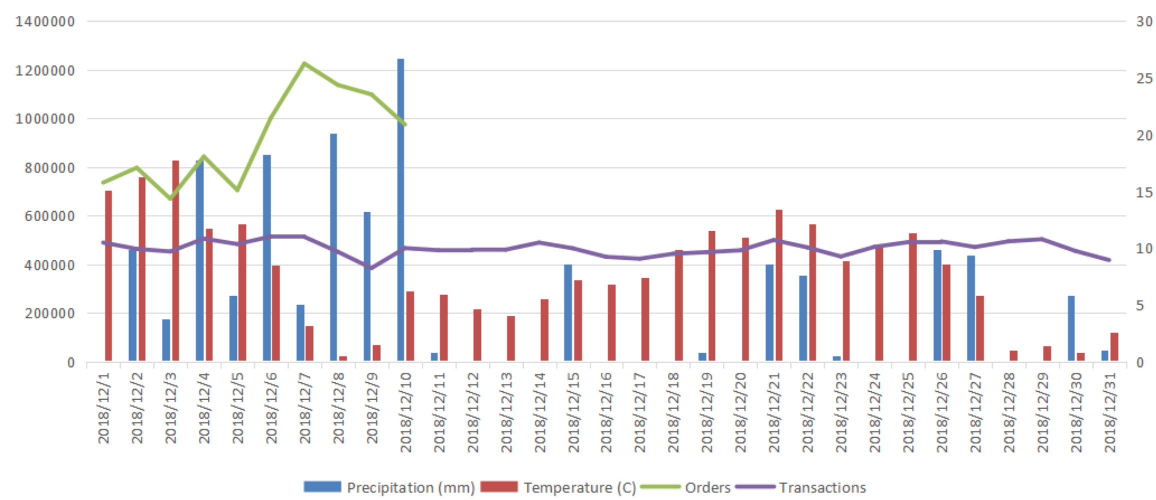

Figure C.1: Orders, Transactions, Precipitation, and Temperature across Days

D Cluster Schedules

This online appendix describes how we cluster drivers based on their past work schedules. Our sample includes 40,104 unique drivers and 470,733 driver-workdays (observations). We construct the following variables to measure a driver's work schedule for a given day and include them in our study:

- **totHour**: driver's total working hours in a day.
- **consecutive**: the longest period a driver has worked consecutively in a day.
- **start**: the hour during which the driver starts working.
- **earlyMorning**: a dummy variable indicates whether the driver is working during the early morning hours of 0 AM to 7 AM.
- **morning**: a dummy variable indicates whether the driver is working during the morning peak hours of 7 AM to 10 AM.
- **midday**: a dummy variable indicates whether the driver is working during the midday hours of 10 AM to 4 PM.
- **afternoon**: a dummy variable indicates whether the driver is working during the afternoon peak hours of 4 PM to 7 PM.
- **evening**: a dummy variable indicates whether the driver is working during the evening hours of 7 PM to 10 PM.
- **night**: a dummy variable indicates whether the driver is working during the night hours of 10 PM to 12 AM.
- **morningCon/afternoonCon/eveningCon**: the longest consecutive working periods a driver has in the morning, afternoon, and evening, respectively.

The first two variables capture the total working hours and the longest consecutive working hours of a driver in a workday. The variable *start* captures the time a driver begins working on a given day.[31] Next, we use six dummy variables to capture whether the driver is working during the early morning, morning, midday, afternoon, evening, and/or night on

[31] Some drivers start working in the late afternoon or early evening and continue until the early morning of the next day. To accurately capture the start working hours for these drivers, we define the start working hour as the earliest hour of the day, with the condition that the driver has not worked for at least six consecutive hours beforehand.

a given day. Last, we calculate the longest consecutive working periods a driver has in the morning, afternoon, and evening hours on a given day.

We apply the k-means method to cluster the drivers in our dataset, aiming to categorize various work schedules into discrete types. The k-means clustering method divides observations into a specific number of groups based on their similarity. We do not know the optimal number of groups to define *ex ante*. Therefore, we test with $k = 2$, 3, and 4 clusters.

Table D.1 shows the results of clustering drivers using the k-means method. The results clearly show that drivers can primarily be classified into two types. The first type comprises committed or full-time drivers, which we refer to as H-type drivers. The second type consists of uncommitted or part-time drivers, referred to as L-type drivers.

Panel I of Table D.1 shows that H-type drivers (type 1 in Panel I) work a total of approximately 12 hours, including 8 consecutive hours. These drivers typically start working at 7 AM. The result remains robust across different specifications when we change the number of clusters from 2 to 3 and 4: we consistently find the existence of such H-type drivers, with around 12 total working hours, 8 hours of consecutive work, and a start time at 7 AM. Moreover, in all specifications, H-type drivers account for about half of the total number of observations.

Regarding L-type drivers, Panel I of Table D.1 shows that these drivers typically work a total of 5 hours, including 3 consecutive hours, and on average start working at 11 AM. When we increase the number of clusters, we find L-type drivers with varying work habits. For example, Panel III of Table D.1 reveals L-type drivers (type 2 in Panel III) who start working in the morning around 8 AM and finish in the early afternoon. Additionally, there are L-type drivers (type 3 in Panel III) who begin in the afternoon around 5 PM and finish at night.

These findings are also consistent with our subsequent analysis of driver labor supply patterns in Online Appendix G. There, we discover that H-type drivers are significantly more committed, showing a tendency to work a similar number of total hours both daily and weekly. They establish a stable work routine, operating as full-time drivers. In contrast, L-type drivers log fewer total working hours and exhibit much greater variations from day to day and week to week.

Table D.1: Results of Clustering Drivers Using K-means

Types	Total Hrs	Consecutive Hours	Start	Early-morning	Morning	Midday	Afternoon	Evening	Night	morning-Con	afternoon-Con	evening-Con	N
Panel I: K-means with 2 types of drivers													
1	11.81	8.09	7	0.22	0.83	0.97	0.94	0.82	0.52	2.02	2.44	1.98	262,268
2	4.83	3.44	11	0.16	0.39	0.58	0.50	0.45	0.29	0.77	0.98	0.94	208,465
Panel II: K-means with 3 types of drivers													
1	12.42	8.68	7	0.21	0.81	0.98	0.97	0.87	0.57	2.00	2.61	2.16	218,533
2	5.13	3.82	17	0.23	0.04	0.36	0.65	0.72	0.53	0.10	1.32	1.56	94,141
3	5.74	3.69	8	0.15	0.74	0.80	0.47	0.31	0.14	1.55	0.96	0.61	158,059
Panel III: K-means with 4 types of drivers													
1	12.70	8.82	7	0.20	0.95	0.99	0.97	0.86	0.53	2.33	2.60	2.11	184,259
2	5.63	3.63	8	0.16	0.81	0.78	0.45	0.28	0.12	1.69	0.90	0.55	145,470
3	3.58	2.82	17	0.15	0.02	0.29	0.53	0.60	0.42	0.04	0.93	1.21	68,549
4	9.67	6.80	13	0.30	0.08	0.83	0.94	0.93	0.72	0.16	2.36	2.27	72,455

E Results from Driver Score

As outlined in Section 2.1, Platform X assigns orders to drivers based on their order assignment scores. For a robustness check, we have created an imputed assignment score that mimics the actual score calculation process.

First, we examine whether higher imputed driver scores correlate with higher hourly wages. Given that we have one month of data, we first calculate the score using the first three weeks of data. Then, we explore how different scores affect driver hourly wages in the last week. Results in Table E.1 show that each unit difference in score corresponds to a significant difference in hourly wage, amounting to 0.0149 CNY per score point.

Based on our classification of drivers into L-type (low-performing) and H-type (high-performing), we found that low-performing drivers have an average driver score of 30.1, whereas high-performing drivers have an average score of 154.8 across 15 workdays. Consequently, this difference in driver scores translates to an hourly wage differential of 3.72 CNY based on the results of Table E.1.[32] Notably, this wage discrepancy is closely aligned with the hourly wage gap found in our main reduced-form analysis, as detailed in Table 5.

Table E.1: Wage Differentials with Imputed Driver Scores

Dependent Variables	Hourly Wage		
	(1)	(2)	(3)
Driver Score	0.0157***	0.0159***	0.0149***
	(0.000281)	(0.000276)	(0.000278)
Constant	47.93***	47.89***	48.93***
	(0.0472)	(0.0464)	(0.132)
Day-Hour FE		Y	Y
Origin FE			Y
Destination FE			Y
Observations	1,019,126	1,019,126	1,019,126
R-squared	0.003	0.039	0.050

Notes: Standard errors in parentheses. *** p<0.01

Next, we assess the validity of our classification of drivers into two types using the imputed driver score. Table E.2 presents the correlation between imputed driver scores and attributes of their past work schedules. The findings indicate that drivers with more total and

[32] According to the platform, the driver score is calculated based on performance over the last 30 days. On average, the monthly driver score disparity between H-type and L-type drivers is 249.4. Therefore, given the estimated coefficient in Table E.1, the hourly wage differential is calculated as $249.4 \times 0.0149 = 3.72$.

consecutive hours tend to have higher imputed driver scores, aligning with the criteria used for clustering drivers. Furthermore, drivers who work more hours during incentivized periods (midday and late night) typically exhibit higher imputed driver scores. These observations suggest that both the imputed driver score and our clustering method reflect similar aspects of drivers' historical work schedules.

Table E.2: Correlation Between Imputed Driver Scores and Past Work Schedules

Dependent Variable	Imputed Driver Score
	(1)
Total Daily Hours	3.657***
	(0.091)
Consecutive Hours	0.485***
	(0.064)
Morning	-3.396***
	(0.333)
Midday	20.983***
	(0.389)
Afternoon	6.653***
	(0.364)
Evening	2.104***
	(0.355)
Night	-3.957***
	(0.322)
Late Night	13.804***
	(0.336)
Observations	470,733
R-squared	0.085

Notes: Standard errors in parentheses. *** $p<0.01$

In conclusion, our robustness check yielded two key findings. First, there exists a significant correlation between a driver's hourly wage and their assignment score, with drivers possessing higher assignment scores typically earning a higher hourly wage than their lower-scoring counterparts. Second, using imputed driver scores validates our clustering of drivers into two types. We observed a significant difference in scores between drivers classified as high-performing versus those categorized as low-performing, and both the imputed driver score and our clustering method reflect similar aspects of drivers' historical work schedules.

F Control For Unobservable Selection: IV results

To mitigate potential biases in driver selection based on unobservable characteristics, we employ instrumental dummy variables: the rate of change in precipitation and air quality index (AQI) in the driver's hometown city between 2017 and 2018. The selection of drivers based on these variables is not influenced by the drivers' unobserved characteristics.

These weather variables satisfy the two conditions required for valid instrumental variables (IVs). First, the occurrence of precipitation and changes in air pollution may be correlated with a driver's decision to become a high-performing driver, as suggested by (Miguel et al., 2004). For instance, alterations in precipitation and air pollution might motivate more drivers to leave their hometowns and become high-performing drivers in the city under study. Second, the variation in weather conditions in a driver's hometown should not directly impact the driver's hourly rate or order distribution in the city being studied. The IV results appear in column (5), revealing a more significant wage differential between high- and low-performing drivers.

Table F.1: Wage Differential: First Stage

	High-performing					
	(1)	(2)	(3)	(4)	(5)	(6)
Change in Precipitation	-0.157***	-0.155***	-0.154***	-0.159***	-0.157***	-0.155***
	(0.00131)	(0.00129)	(0.00130)	(0.00134)	(0.00132)	(0.00132)
Change in AQI				-0.0270***	-0.0231***	-0.0206***
				(0.00322)	(0.00318)	(0.00318)
Constant	0.898***	0.898***	0.886***	0.895***	0.896***	0.884***
	(0.000148)	(0.000146)	(0.000450)	(0.000327)	(0.000324)	(0.000529)
Controls:						
Day-Hour FE		Y	Y		Y	Y
Origin/Dest FE			Y			Y
Observations	4,182,318	4,182,318	4,182,318	4,182,318	4,182,318	4,182,318
R-squared	0.003	0.023	0.024	0.003	0.023	0.024

Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

The first stage results appear in Table F.1. We regress the High-performing dummy on

Table F.2: Wage Differential: IV results

	Hourly Wage					
	(1)	(2)	(3)	(4)	(5)	(6)
High-performing	8.908***	9.480***	5.599***	9.268***	9.813***	5.759***
	(0.679)	(0.676)	(0.677)	(0.678)	(0.675)	(0.676)
Constant	41.98***	47.28***	49.47***	41.65***	46.97***	49.32***
	(0.611)	(1.192)	(1.186)	(0.609)	(1.191)	(1.185)
Day-Hour FE		Y	Y		Y	Y
Origin FE			Y			Y
Destination FE			Y			Y
IV (2017-2018)						
Change in Precipitation	Y	Y	Y	Y	Y	Y
Change in AQI				Y	Y	Y
Observations	4,182,318	4,182,318	4,182,318	4,182,318	4,182,318	4,182,318
R-squared	0.024	0.034	0.050	0.024	0.034	0.050

Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

the weather conditions, with F-values exceeding 1000 for all specifications. In column (1), we regress the High-performing dummy on the precipitation change without any fixed effects. In column (2), we include day-hour fixed effects, and in column (3), we include day-hour, origin, and destination fixed effects. Columns (4)–(6) are replicates of (1)–(3), except that we include both precipitation change and AQI change as explanatory variables.

Our analysis reveals a robust negative correlation between precipitation and driver performance. Specifically, we find a one percent increase in precipitation in a driver's hometown to be associated with a 0.154% decrease in the probability of a driver attaining high-performance status. More precipitation is usually associated with higher agricultural output (Miguel et al., 2004) and may inadvertently diminish drivers' incentives to move to the city we study and become a high-performing driver. Moreover, we identify a negative impact of air quality on driver performance. A one percent increase in the Air Quality Index (AQI), reflecting elevated pollution levels, corresponds to a 0.02% decrease in the likelihood of drivers achieving high-performance status. This observation aligns with the hypothesis that heightened pollution levels, often linked to increased industrial output in drivers' hometowns, may undermine their motivation to move to the city we study and become high-performing drivers.

The instrumental variable (IV) results appear in Table F.2. Column (1) presents the results without any fixed effects, while column (2) includes day-hour fixed effects. In column (3), we include day-hour fixed effects, origin fixed effects, and destination fixed effects. Columns (4)–(6) are replicates of (1)–(3), except that we include both precipitation change and AQI change as instruments.

Our findings demonstrate that high-performing drivers outearn their low-performing counterparts. Specifically, when examining column (3)/(6), we observe that high-performing drivers earn 5.6/5.8 CCY more per hour than low-performing drivers. This wage differential is notably higher than that indicated by our main ordinary least squares (OLS) findings.

We conjecture that low-performing drivers exhibit a higher level of strategic behavior. For instance, they may strategically select optimal times and locations to work, which contributes to their comparatively lower number of working hours. This aligns with the fact that their working decisions are more discerning and selective. In summary, our OLS estimates underestimate the wage disparity due to algorithmic preferences.

G Driver Labor Supply Patterns

In this section we document the labor supply patterns of drivers for Platform X, especially across-day patterns. We follow Chen et al. (2019) in documenting the distribution of active drivers for Platform X, and the transition matrix of hours worked in contiguous weeks for drivers who meet our active driver criterion.

Table G.1 provides the distribution of total hours worked by week supplied by drivers in the full sample. The table displays the share of the drivers who were active on the system for various time bins. The vast majority of H-type drivers work more than 20 hours a week, while the majority of L-type drivers work less than 20 hours. This is consitent with our earlier claim that H-type drivers are more likely to be full-time drivers while L-type drivers are part-time drivers. This is different from the patterns of Uber drivers in the US as reported by Chen et al. (2019). For the Uber drivers in the US, the overwhelming majority of Uber drivers are working part-time hours with fewer than 12 hours per week.

Table G.1: Distribution of Active Hours Per Week for Drivers

Total Hours	Share of Driver Weeks (%)	
	H-type	L-type
0	7%	54%
1–10	5%	27%
11–20	9%	10%
21–30	13%	5%
31–40	16%	3%
41–50	19%	1%
51–60	18%	1%
> 60	13%	0%

Table G.2 and Table G.3 illustrate the extent to which a driver's total activity varies from week to week. The stub column shows bins of hours of driver activity in a week, and the remaining columns show the share of the drivers who are active on the system for various time bins during the subsequent week. For example, of the drivers active from 21 to 30 hours in a week, 26 percent H-type drivers and 12 percent L-type drivers fall into the same time-supplied bin in the subsequent week.

Table G.2 show that H-type drivers have strong tendency to work a similar number of total hours from week to week. For example, for drivers who work more than 60 hours per week, the probability of doing the same in the contiguous week is 48%, and the probability of those drivers working more than 40 hours is 87%. Only 1% of drivers in the group will not work in the contiguous week.

Table G.2: Transition Matrix of Hours Worked in Contiguous Weeks (H-type drivers)

	t+1							
t	0	1–10	11–20	21–30	31–40	41–50	51–60	> 60
0	29%	8%	12%	12%	13%	12%	9%	5%
1–10	18%	15%	21%	15%	12%	9%	6%	5%
11–20	6%	12%	27%	22%	12%	9%	7%	4%
21–30	4%	7%	18%	26%	20%	13%	7%	4%
31–40	4%	5%	8%	18%	27%	23%	11%	5%
41–50	3%	3%	5%	10%	20%	31%	22%	7%
51–60	2%	2%	3%	5%	12%	26%	33%	17%
> 60	1%	1%	2%	3%	6%	12%	27%	48%

Table G.3 shows that for L-type drivers, they are much less likely to work long hours for the platform. Moreover, week-to-week variation is much greater for the L-type drivers. The probability of not working is much greater for the L-type drivers. Compared to H-type drivers, L-type drivers seem to much less committed to the total hours of working from week to week.

Table G.3: Transition Matrix of Hours Worked in Contiguous Weeks (L-type drivers)

	t+1							
t	0	1–10	11–20	21–30	31–40	41–50	51–60	> 60
0	73%	17%	5%	2%	1%	1%	1%	0%
1–10	46%	41%	9%	2%	1%	1%	0%	0%
11–20	34%	29%	25%	7%	2%	1%	1%	1%
21–30	44%	20%	17%	12%	5%	2%	0%	0%
31–40	52%	17%	11%	10%	8%	2%	0%	0%
41–50	56%	16%	12%	8%	6%	2%	0%	0%
51–60	52%	20%	18%	8%	2%	1%	0%	0%
> 60	36%	26%	27%	10%	1%	0%	0%	0%

Next, we perform similar analysis for day-to-day work patterns. Table G.4 provides the distribution of total hours worked by day in the full sample. About 70% of H-type drivers work more than 5 hours per day, while only 9% of L-type drivers do so.

Table G.4: Distribution of Active Hours Per Day for Drivers

Total Hours	Share of Driver Days (%)	
	H-type	L-type
0	20%	76%
1–4	12%	14%
5–8	20%	5%
9–12	29%	3%
13+	19%	1%

When we examine the day-to-day work patterns, the results are similar to those we found for week-to-week patterns. Table G.5 and Table G.6 illustrate the extent to which a driver's total activity varies from day to day. H-type drivers have strong tendency to work a similar number of total hours from day to day. For example, for drivers who work about 5–8 hours per day, the probability of doing the same in the contiguous day is 31%, while the probability of working more than 5 hours is 73%.

Table G.5: Transition Matrix of Hours Worked in Contiguous Days (H-type drivers)

	t+1					
t	0	1–4	5–8	9–12	13–16	17–20
0	63%	10%	11%	10%	5%	0%
1–4	19%	26%	27%	20%	9%	0%
5–8	10%	16%	31%	30%	12%	0%
9–12	5%	8%	19%	44%	23%	1%
13–16	3%	5%	12%	33%	45%	2%
17–20	2%	4%	8%	21%	54%	10%

In comparison, Table G.6 shows that for L-type drivers, the day-to-day variation is much greater. The probability of not working is much greater for the L-type drivers. Among all the groups for L-type drivers, their probability of not working in the contiguous day is significantly greater than those for H-type drivers.

Table G.6: Transition Matrix of Hours Worked in Contiguous Days (L-type drivers)

	t+1					
t	0	1–4	5–8	9–12	13–16	17–20
0	89%	8%	2%	1%	0%	0%
1–4	43%	42%	10%	3%	1%	0%
5–8	28%	27%	27%	13%	4%	0%
9–12	21%	15%	21%	30%	13%	0%
13–16	11%	10%	15%	29%	34%	2%
17–20	5%	10%	12%	18%	47%	9%

In general, for H-type drivers, they are more likely to work a similar number of total hours from day-to-day and week-to-week. They build a stable work routine across days as they work as full-time drivers. For L-type drivers, they work less total working hours and have much greater week-to-week and day-to-day variations. Therefore, H-type drivers work longer and seem to be committed to working similar hours, as predicted in our results. Also, this highlights that focusing on within-day dynamics is sufficient to capture the main decisions of the drivers.

H Drivers' Finite-Horizon Dynamic Problem

This appendix describes in detail drivers' finite-horizon dynamic choices. For each hour t, the utility of working and not working are specified as

$$\begin{aligned} U_{1t}^{\tau} &= \underbrace{W_t^{\tau}}_{\text{preferential wage rate}} + \sigma \cdot \epsilon_{1t}, \\ U_{0t}^{\tau} &= \underbrace{O_t^{\iota}}_{\substack{\text{outside option value,} \\ O_t^{\iota} = O_t + \eta_{\iota,t}}} + \sigma \cdot \epsilon_{0t}, \end{aligned} \tag{H.1}$$

Drivers first observe random shocks ϵ, then decide whether to work. To keep things straightforward, we leave out the hyperscript ι when referring to O_t^{ι} in the derivations below.

H.1 Low-Performing Drivers

For the final period, $t = T$,

$$V_T^L = \begin{cases} W_T^L + \sigma \cdot \epsilon_{1T} & \text{if } a_T = 1 \ \& \ a_{T-1} = 1, \\ W_T^L - \kappa + \sigma \cdot \epsilon_{1T} & \text{if } a_T = 1 \ \& \ a_{T-1} = 0, \\ O_T + \sigma \cdot \epsilon_{0T} & \text{if } a_T = 0. \end{cases}$$

So, the expected utility for the last period T is

$$\begin{aligned} EV_{1T}^L &= \sigma \left[\ln \Big(\exp(W_T^L/\sigma) + \exp(O_T/\sigma) \Big) + \gamma \right], \\ EV_{0T}^L &= \sigma \left[\ln \Big(\exp((W_T^L - \kappa)/\sigma) + \exp(O_T/\sigma) \Big) + \gamma \right]. \end{aligned}$$

Throughout our model, EV's subscript 1 represents $a_{t-1} = 1$. In this case, EV_{1T}^L represents the expected value of a low-performing driver at time T if $a_{T-1} = 1$.

At any time $t \in [T-1, 2]$,

$$V_t^L = \begin{cases} W_t^L + \sigma \cdot \epsilon_{1t} + \beta EV_{1t+1}^L & \text{if } a_t = 1 \ \& \ a_{t-1} = 1, \\ W_t^L - \kappa + \sigma \cdot \epsilon_{1t} + \beta EV_{1t+1}^L & \text{if } a_t = 1 \ \& \ a_{t-1} = 0, \\ O_t + \sigma \cdot \epsilon_{0t} + \beta EV_{0t+1}^L & \text{if } a_t = 0. \end{cases}$$

So, the expected utility of period t is

$$EV_{1t}^L = \sigma\left[\ln\left(\exp((W_t^L + \beta EV_{1t+1}^L)/\sigma) + \exp((O_t + \beta EV_{0t+1}^L)/\sigma)\right) + \gamma\right],$$
$$EV_{0t}^L = \sigma\left[\ln\left(\exp((W_t^L - \kappa + \beta EV_{1t+1}^L)/\sigma) + \exp((O_t + \beta EV_{0t+1}^L)/\sigma)\right) + \gamma\right].$$

For the first period, $t = 1$,

$$V_1^L = \begin{cases} W_1^L - \kappa + \sigma \cdot \epsilon_{11} + \beta EV_{12}^L & \text{if } a_1 = 1, \\ O_1 + \sigma \cdot \epsilon_{01} + \beta EV_{02}^L & \text{if } a_1 = 0. \end{cases}$$

The expected value of being a low-performing driver is then

$$EV^L = \sigma\left[\ln\left(\exp((W_1^L - \kappa + \beta EV_{12}^L)/\sigma) + \exp((O_1 + \beta EV_{02}^L)/\sigma)\right) + \gamma\right]. \tag{H.2}$$

H.2 High-Performing Drivers

High-performing drivers are required to work at T_0 and for at least 2 consecutive hours. T_0 can be any hour between 10AM–2PM or 7PM–5AM. There are 16 possible work schedules to choose from. For schedule $j \in \{1, \cdots, 16\}$, with committed working hours $[T_0, T_0+1]$,

if $T_0 + 2 < T$, then for the last period T,

$$V_T^j = \begin{cases} W_T^H + \sigma \cdot \epsilon_{1T} & \text{if } a_T = 1 \ \&\ a_{T-1} = 1, \\ W_T^H - \kappa + \sigma \cdot \epsilon_{1T} & \text{if } a_T = 1 \ \&\ a_{T-1} = 0, \\ O_T + \sigma \cdot \epsilon_{0T} & \text{if } a_T = 0. \end{cases}$$

The expected utility of period T is

$$EV_{1T}^j = \sigma\left[\ln\left(\exp(W_T^H/\sigma) + \exp(O_T/\sigma)\right) + \gamma\right],$$
$$EV_{0T}^j = \sigma\left[\ln\left(\exp((W_T^H - \kappa)/\sigma) + \exp(O_T/\sigma)\right) + \gamma\right].$$

At $t \in [T_0 + 3, T - 1]$,

$$V_t^j = \begin{cases} W_t^H + \sigma \cdot \epsilon_{1t} + \beta EV_{1t+1}^j & \text{if } a_t = 1 \ \&\ a_{t-1} = 1, \\ W_t^H - \kappa + \sigma \cdot \epsilon_{1t} + \beta EV_{1t+1}^j & \text{if } a_t = 1 \ \&\ a_{t-1} = 0, \\ O_t + \sigma \cdot \epsilon_{0t} + \beta EV_{0t+1}^j & \text{if } a_t = 0. \end{cases}$$

The expected utility of period $t \in [T_0+3, T-1]$ is

$$EV^j_{1t} = \sigma \left[\ln \Big(\exp((W^H_t + \beta EV^j_{1t+1})/\sigma) + \exp((O_t + \beta EV^j_{0t+1})/\sigma) \Big) + \gamma \right],$$
$$EV^j_{0t} = \sigma \left[\ln \Big(\exp((W^H_t - \kappa + \beta EV^j_{1t+1})/\sigma) + \exp((O_t + \beta EV^j_{0t+1})/\sigma) \Big) + \gamma \right].$$

At time T_0+2, because the driver commits to work at T_0 and T_0+1, $a_{T0+1} = 1$ with probability 1,

$$V^j_{T_0+2} = \begin{cases} W^H_{T_0+2} + \sigma \cdot \epsilon_{1T_0+2} + \beta EV^j_{1T_0+3} & \text{if } a_{T_0+2} = 1, \\ O_{T_0+2} + \sigma \cdot \epsilon_{0T_0+2} + \beta EV^j_{0T_0+3} & \text{if } a_{T_0+2} = 0. \end{cases}$$

At T_0+1, the high-performing driver has to work. The expected value at any T_0+1 is

$$EV^j_{1T_0+1} = W^H_{T_0+1} + \beta EV^j_{1T_0+2} + \sigma\gamma.$$

At period T_0, the expected value is

$$EV^j_{1T_0} = W^H_{T_0} + \beta EV^j_{1T_0+1} + \sigma\gamma,$$
$$EV^j_{0T_0} = W^H_{T_0} - \kappa + \beta EV^j_{1T_0+1} + \sigma\gamma.$$

At any time before T_0, $t \in [2, T_0-1]$, the expected utility is

$$EV^j_{1t} = \sigma \left[\ln \Big(\exp((W^H_t + \beta EV^j_{1t+1})/\sigma) + \exp((O_t + \beta EV^j_{0t+1})/\sigma) \Big) + \gamma \right],$$
$$EV^j_{0t} = \sigma \left[\ln \Big(\exp((W^H_t - \kappa + \beta EV^j_{1t+1})/\sigma) + \exp((O_t + \beta EV^j_{0t+1})/\sigma) \Big) + \gamma \right].$$

For period 1,

$$V^j_1 = \begin{cases} W^H_1 - \kappa + \sigma \cdot \epsilon_{11} + \beta EV^j_{12} & \text{if } a_1 = 1, \\ O_1 + \sigma \cdot \epsilon_{01} + \beta EV^j_{02} & \text{if } a_1 = 0. \end{cases}$$

The expected value of being a high-performing driver is then

$$EV^j = \sigma \left[\ln \Big(\exp((W^H_1 - \kappa + \beta EV^j_{12})/\sigma) + \exp((O_1 + \beta EV^j_{02})/\sigma) \Big) + \gamma \right]. \tag{H.3}$$

I Identification and Estimation

I.1 Identification

We start with the case without UH, so that $O_t^\iota = O_t$. Denote $P_t^\tau(a_T = a|a_{T-1} = b) = P_t^\tau(a|b)$, where $a, b = 0, 1$ and $\tau = L, H$. We have

$$\log P_T^L(1|1) - \log P_T^L(0|1) = \frac{W_T^L - O_T}{\sigma}, \tag{I.1}$$

$$\log P_T^L(1|0) - \log P_T^L(0|0) = \frac{W_T^L - O_T - \kappa}{\sigma}, \tag{I.2}$$

which implies that

$$\frac{\kappa}{\sigma} = [\log P_T^L(1|1) - \log P_T^L(0|1)] - [\log P_T^L(1|0) - \log P_T^L(0|0)]. \tag{I.3}$$

Similarly, H-type drivers have

$$\log P_T^H(1|1) - \log P_T^H(0|1) = \frac{W_T^H - O_T}{\sigma}, \tag{I.4}$$

$$\log P_T^H(1|0) - \log P_T^H(0|0) = \frac{W_T^H - O_T - \kappa}{\sigma}. \tag{I.5}$$

Combining (I.1) and (I.4) gives

$$\sigma = \frac{W_T^H - W_T^L}{[\log P_T^H(1|1) - \log P_T^H(0|1)] - [\log P_T^L(1|1) - \log P_T^L(0|1)]}, \tag{I.6}$$

which implies that κ is identified by following (I.3)

$$\kappa = (W_T^H - W_T^L)\frac{[\log P_T^L(1|1) - \log P_T^L(0|1)] - [\log P_T^L(1|0) - \log P_T^L(0|0)]}{[\log P_T^H(1|1) - \log P_T^H(0|1)] - [\log P_T^L(1|1) - \log P_T^L(0|1)]}$$

and O_T is identified by following (I.1) or (I.4)

$$O_T = W_T^L - (W_T^H - W_T^L)\frac{\log P_T^L(1|1) - \log P_T^L(0|1)}{[\log P_T^H(1|1) - \log P_T^H(0|1)] - [\log P_T^L(1|1) - \log P_T^L(0|1)]}.$$

When there is no UH, we have three unknown parameters O_T, σ, κ and four equations that capture the observed CCPs $P_T^\tau(1|a_{T-1})$, where $a_{T-1} = 0, 1$ and $\tau = L, H$. Note that $P_T^\tau(1|a_{T-1})$ and the above defined odd ratios $P_T^\tau(1|a_{T-1})/P_T^\tau(0|a_{T-1})$ capture the same amount of identifying information because $P_T^\tau(1|a_{T-1}) + P_T^\tau(0|a_{T-1}) = 1$. The system is overidentified using just the last period, which is clear because our identification steps do not involve (I.5).

When there is UH, we assume that

$$O_t^\iota = O_t + \eta_{\iota,t},$$

where $\eta_{\iota,t}$ represents time-interval-specific preference. Considering data from the last period, we have five unknown parameters $O_T, \eta_{(5)}, p_{(5)}, \sigma, \kappa$, where (5) denotes the 5-th UH type, and again four equations that capture the observed CCPs $P_T^\tau(1|a_{T-1})$, where $a_{T-1} = 0, 1$ and $\tau = L, H$. Obviously, data from the last period are insufficient for point identification. Combining the last two periods, we have six unknown parameters $O_T, O_{T-1}, \eta_{(5)}, p_{(5)}, \sigma, \kappa$ and eight equations. More specifically, the eight equations represent the observed CCPs $\overline{P}_t^\tau(1|a_{t-1}) = (1 - p_{(5)})P_t^\tau(1|a_{t-1}) + p_{(5)}P_{t,(5)}^\tau(1|a_{t-1})$, where $a_{t-1} = 0, 1$, $t = T-1, T$, and $\tau = L, H$, relate to the unknown parameters. In particular,

$$\overline{P}_T^\tau(1|a_{T-1}) =(1 - p_{(5)})\frac{\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma})}{\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma}) + \exp(\frac{O_T}{\sigma})} \tag{I.7}$$

$$+ p_{(5)}\frac{\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma})}{\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma}) + \exp(\frac{O_T + \eta_{(5)}}{\sigma})}, \tag{I.8}$$

$$\overline{P}_{T-1}^\tau(1|a_{T-2}) =(1 - p_{(5)})\frac{\exp(\frac{W_{T-1}^\tau - \kappa 1(a_{T-2}=0) + \beta EV_{1T}^\tau}{\sigma})}{\exp(\frac{W_{T-1}^\tau - \kappa 1(a_{T-2}=0) + \beta EV_{1T}^\tau}{\sigma}) + \exp(\frac{O_{T-1} + \beta EV_{0T}^\tau}{\sigma})} \tag{I.9}$$

$$+ p_{(5)}\frac{\exp(\frac{W_{T-1}^\tau - \kappa 1(a_{T-2}=0) + \beta EV_{1T}^{\tau,(5)}}{\sigma})}{\exp(\frac{W_{T-1}^\tau - \kappa 1(a_{T-2}=0) + \beta EV_{1T}^{\tau,(5)}}{\sigma}) + \exp(\frac{O_{T-1} + \eta_{(5)} + \beta EV_{0T}^{\tau,(5)}}{\sigma})}, \tag{I.10}$$

where the expected utility for the last period T is given by

$$EV_{a_{T-1}T}^\tau = \sigma\left[\ln\Big(\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma}) + \exp(\frac{O_T}{\sigma})\Big) + \gamma\right],$$
$$EV_{a_{T-1}T}^{\tau,(5)} = \sigma\left[\ln\Big(\exp(\frac{W_T^\tau - \kappa 1(a_{T-1}=0)}{\sigma}) + \exp(\frac{O_T + \eta_{(5)}}{\sigma})\Big) + \gamma\right].$$

Note that $\overline{P}_t^\tau(0|a_{t-1}) = 1 - \overline{P}_t^\tau(1|a_{t-1})$ does not provide additional identification power.

We can continue the identification process backward and identify all the remaining parameters. In summary, we can identify the model with UH as long as each type involves at least two periods.

I.2 Estimation

We follow our identification argument closely in estimating the model. For each hour t, the utility of working and not working are

$$\begin{aligned}
U_{1t}^{\tau} &= W_t^{\tau} + \sigma \cdot \epsilon_{1t}, \\
U_{0t}^{\tau} &= O_t^{\iota} + \sigma \cdot \epsilon_{0t} \\
&= O_t + \eta_{\iota,t} + \sigma \cdot \epsilon_{0t}.
\end{aligned}$$

We use π_ι to denote the probability of individual driver i being the unobserved type ι. We make use of observed conditional choice probabilities to estimate the structural parameters. First, we derive the conditional choice probability of working for each type of driver. To keep things straightforward, we leave out the hyperscript ι when referring to O_t^ι in the derivations below.

Low-performing Drivers

For the final period T, the conditional probability of working for each unobserved type ι is

$$\begin{aligned}
P^L(a_T = 1 | a_{T-1} = 1, \iota) &= \frac{\exp(W_t^L/\sigma)}{\exp(W_t^L/\sigma) + \exp(O_T/\sigma)}, \\
P^L(a_T = 1 | a_{T-1} = 0, \iota) &= \frac{\exp((W_t^L - \kappa)/\sigma)}{\exp((W_t^L - \kappa)/\sigma) + \exp(O_T/\sigma)}.
\end{aligned}$$

For any $t \in [2, T-1]$,

$$\begin{aligned}
P^L(a_t = 1 | a_{t-1} = 1, \iota) &= \frac{\exp((W_t^L + \beta EV_{1t+1}^L)/\sigma)}{\exp((W_t^L + \beta EV_{1t+1}^L)/\sigma) + \exp((O_t + \beta EV_{0t+1}^L)/\sigma)}, \\
P^L(a_t = 1 | a_{t-1} = 0, \iota) &= \frac{\exp((W_t^L - \kappa + \beta EV_{1t+1}^L)/\sigma)}{\exp((W_t^L - \kappa + \beta EV_{1t+1}^L)/\sigma) + \exp((O_t + \beta EV_{0t+1}^L)/\sigma)}.
\end{aligned}$$

For $t = 1$,

$$P^L(a_1 = 1, \iota) = \frac{\exp((W_1^L - \kappa + \beta EV_{12}^L)/\sigma)}{\exp((W_1^L - \kappa + \beta EV_{12}^L)/\sigma) + \exp((O_1 + \beta EV_{02}^L)/\sigma)}.$$

Therefore, at any t, the conditional probability for low-performing drivers is

$$
\begin{aligned}
P^L(a_t = 1|a_{t-1} = 0) &= \sum_{\iota} \pi_\iota \cdot P^L(a_t = 1|a_{t-1} = 0, \iota), \\
P^L(a_t = 1|a_{t-1} = 1) &= \sum_{\iota} \pi_\iota \cdot P^L(a_t = 1|a_{t-1} = 1, \iota).
\end{aligned} \tag{I.11}
$$

High-performing Drivers

For any schedule $j \in \{1, \cdots, 16\}$, the conditional probability of working in the final period T is

$$
\begin{aligned}
P^j(a_T = 1|a_{T-1} = 1, \iota) &= \frac{\exp(W_t^H/\sigma)}{\exp(W_t^H/\sigma) + \exp(O_T/\sigma)}, \\
P^j(a_T = 1|a_{T-1} = 0, \iota) &= \frac{\exp((W_t^H - \kappa)/\sigma)}{\exp((W_t^H - \kappa)/\sigma) + \exp(O_T/\sigma)}.
\end{aligned}
$$

For any $t \in [T_0 + 3, T - 1]$, we have

$$
\begin{aligned}
P^j(a_t = 1|a_{t-1} = 1, \iota) &= \frac{\exp((W_t^H + \beta EV_{1t+1}^j)/\sigma)}{exp((W_t^H + \beta EV_{1t+1}^j)/\sigma) + \exp((O_t + \beta EV_{0t+1}^j)/\sigma)}, \\
P^j(a_t = 1|a_{t-1} = 0, \iota) &= \frac{\exp((W_t^H - \kappa + \beta EV_{1t+1}^j)/\sigma)}{\exp((W_t^H - \kappa + \beta EV_{1t+1}^j)/\sigma) + \exp((O_t + \beta EV_{0t+1}^j)/\sigma)}.
\end{aligned}
$$

At $t = T_0 + 2$, we have

$$
P^j(a_t = 1|a_{t-1} = 1, \iota) = \frac{\exp((W_t^H + \beta EV_{1t+1}^j)/\sigma)}{\exp((W_t^H + \beta EV_{1t+1}^j)/\sigma) + \exp((O_t + \beta EV_{0t+1}^j)/\sigma)}.
$$

At $t = T_0 + 1$, we have

$$
P^j(a_t = 1|a_{t-1} = 1, \iota) = 1.
$$

At $t = T_0$, we have

$$
\begin{aligned}
P^j(a_t = 1|a_{t-1} = 1, \iota) &= 1, \\
P^j(a_t = 1|a_{t-1} = 0, \iota) &= 1.
\end{aligned}
$$

For any $t \in [2, T_0 - 1]$, we have

$$P^j(a_t = 1 | a_{t-1} = 1, \iota) = \frac{\exp((W_t^H + \beta E V_{1t+1}^j)/\sigma)}{\exp((W_t^H + \beta E V_{1t+1}^j)/\sigma) + \exp((O_t + \beta E V_{0t+1}^j)/\sigma)},$$
$$P^j(a_t = 1 | a_{t-1} = 0, \iota) = \frac{\exp((W_t^H - \kappa + \beta E V_{1t+1}^j)/\sigma)}{\exp((W_t^H - \kappa + \beta E V_{1t+1}^j)/\sigma) + \exp((O_t + \beta E V_{0t+1}^j)/\sigma)}.$$

At $t = 1$, we have

$$P^j(a_1 = 1, \iota) = \frac{\exp((W_1^H - \kappa + \beta E V_{12}^j)/\sigma)}{\exp((W_1^H - \kappa + \beta E V_{12}^j)/\sigma) + \exp((O_1 + \beta E V_{02}^j)/\sigma)}.$$

Therefore, at any t, the conditional probability for high-performing drivers is

$$\begin{aligned} P^H(a_t = 1 | a_{t-1} = 0) &= \sum_{\iota} \sum_{j=1}^{16} \pi_\iota \cdot \widetilde{P}^j(\iota) \cdot P^j(a_t = 1 | a_{t-1} = 0, \iota), \\ P^H(a_t = 1 | a_{t-1} = 1) &= \sum_{\iota} \sum_{j=1}^{16} \pi_\iota \cdot \widetilde{P}^j(\iota) \cdot P^j(a_t = 1 | a_{t-1} = 1, \iota), \end{aligned} \tag{I.12}$$

where $\widetilde{P}^j$ is the probability of choosing each high-performing schedule, and

$$\widetilde{P}^j(\iota) = \frac{\exp(EV^j(\iota))}{\sum_{k=1}^{16} \exp(EV^k(\iota))}.$$

We estimate the model by minimizing the weighted distance between the data moments and the simulated moments of the finite mixture model:

$$\{\hat{\theta}, \hat{\pi}\} = arg \min_{\theta, \pi} \quad [\boldsymbol{P^\tau} - \widehat{\boldsymbol{P}}^{\boldsymbol{\tau}}(\theta, \pi)]' W [\boldsymbol{P^\tau} - \widehat{\boldsymbol{P}}^{\boldsymbol{\tau}}(\theta, \pi)].$$

Here, W represents a positive definite matrix. The vector $\boldsymbol{P^\tau}$ contains the actual conditional choice probabilities (CCPs) derived from the data. This vector is of dimension 96-by-1 and includes the conditional probabilities of the different choices for each time period t. Specifically, $P^H(a_t = 1 | a_{t-1} = 0), P^H(a_t = 1 | a_{t-1} = 1), P^L(a_t = 1 | a_{t-1} = 0), P^H(a_t = 1 | a_{t-1} = 0)$. Additionally, $\widehat{\boldsymbol{P}}^{\boldsymbol{\tau}}$ denotes a vector containing simulated conditional choice probabilities, calculated based on equations I.11 and I.12.

J Model Validation

Figure J.1 illustrates the model's goodness of fit. The simulated conditional choice probabilities (CCPs) reasonably align with the observed CCPs for both high- and low-performing drivers. However, there is a slight discrepancy in the fit of high-performing drivers at $a_{t-1} = 0$ during early morning hours. This discrepancy could be attributable to the relatively low number of transactions occurring during this period, as compared to other working hours.

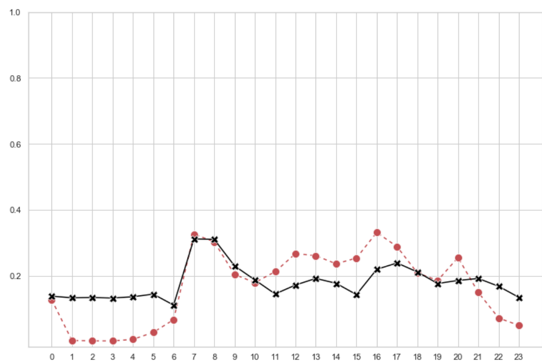

(a) High-performing Drivers, State = 0

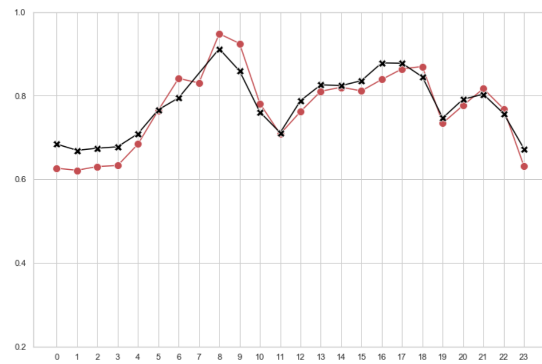

(b) High-performing Drivers, State = 1

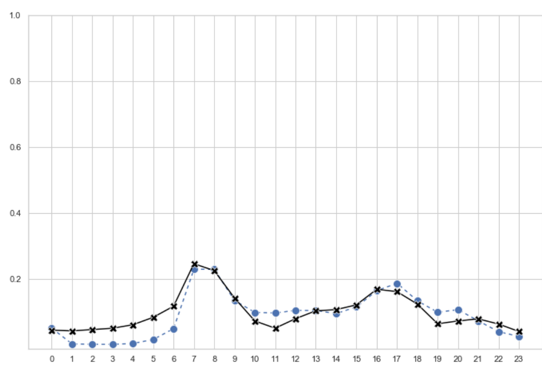

(c) Low-performing Drivers, State = 0

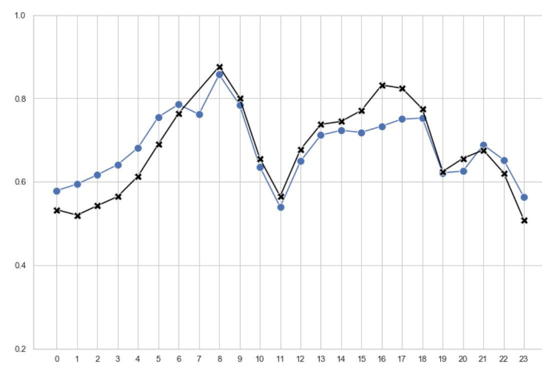

(d) Low-performing Drivers, State = 1

Figure J.1: Model Goodness of Fit

Note: Figure J.1 shows the model's simulated values against the empirically observed CCPs. The black lines represent the model's simulated values.

K Results from Eliminating the Wage Differential

To understand the effect of eliminating wage differentials between high- and low-performing drivers, Figure K.1 shows the equilibrium labor supply decision in panel (a) and the equilibrium wage rate in panel (b). In this scenario, we maintain the ride fares at the same level as when the platform uses a preferential algorithm. When we eliminate the wage differential between high- and low-performing drivers, drivers will switch from being high performing because it no longer provides any bonus. Because we fix the ride fares, and hence rider demand, there will be a labor shortage because of the lack of high-performing drivers. As a result of the excess demand, the equilibrium wage rate without a preferential algorithm will be higher than the wage rate of low-performing drivers when there is a preferential algorithm. The equilibrium wage rate without a preferential algorithm lies between the former wage rates of the high- and low-performing drivers.

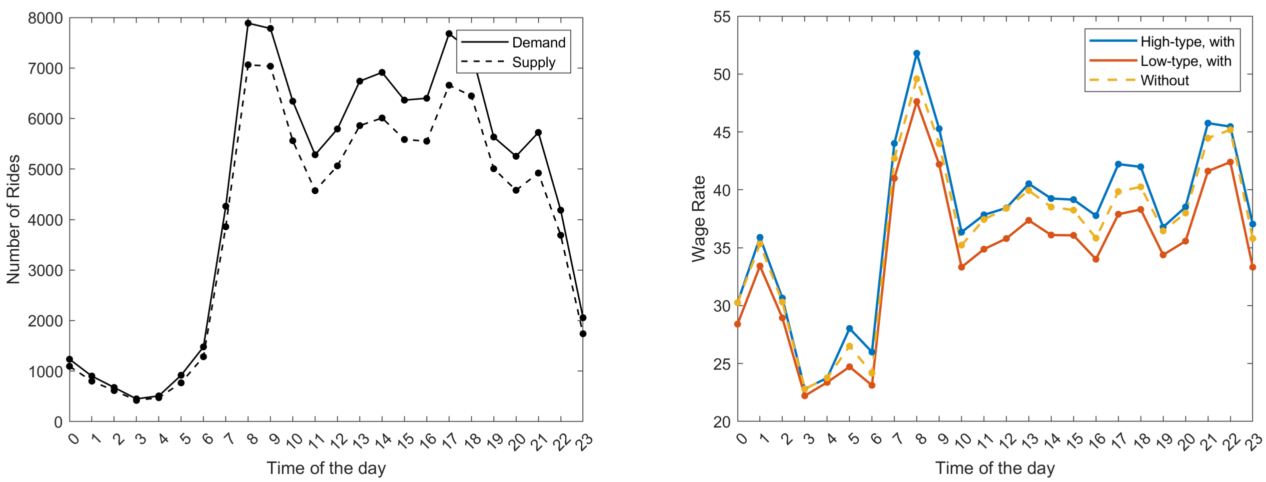

Figure K.1: Results from Eliminating the Wage Differential between W^H and W^L

Next, we study the counterfactual results of eliminating the wage differential between high- and low-performing drivers only in the treatment hour h. When we eliminate this wage differential in one particular hour, drivers will switch from being high performing because the benefit for doing so is now smaller. Because we fix the ride fares, and hence rider demand, there will be a labor shortage because of the lack of high-performing drivers. As a result, the equilibrium wage rate for low-performing drivers without a preferential algorithm will be higher than the wage rate when there is a preferential algorithm. Figure K.2 shows the elasticity of labor supply corresponding to the elimination of the wage differential in treatment hour h.

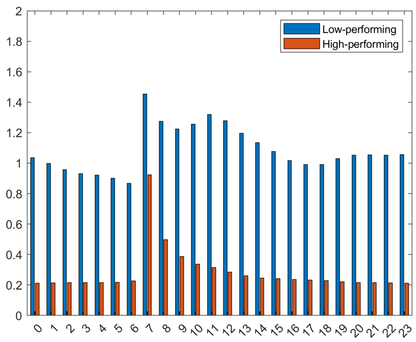

(a) 7 AM

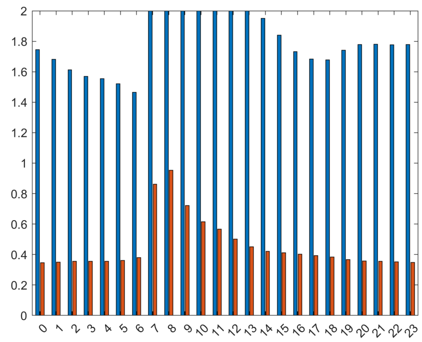

(b) 8 AM

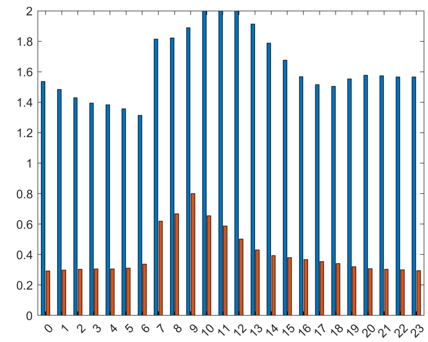

(c) 9AM

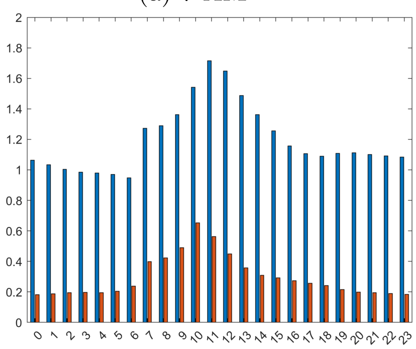

(a) 10 AM

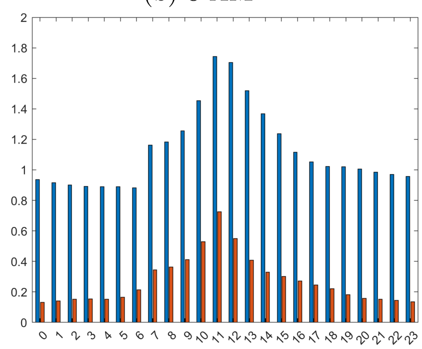

(b) 11 AM

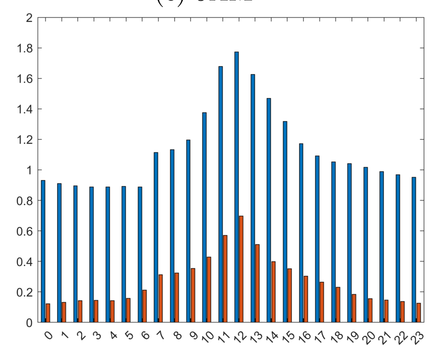

(c) 12 PM

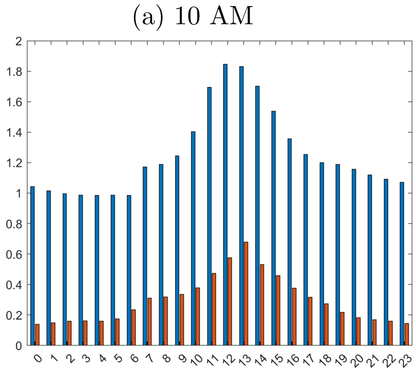

(a) 1 PM

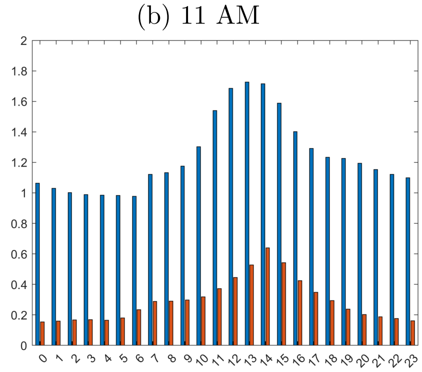

(b) 2 PM

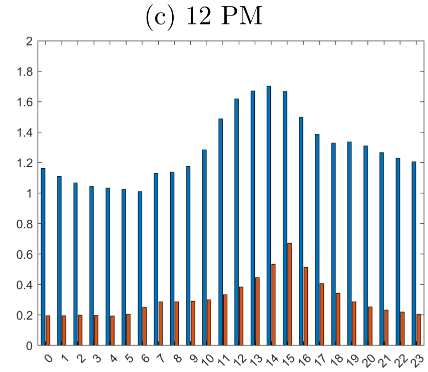

(c) 3 PM

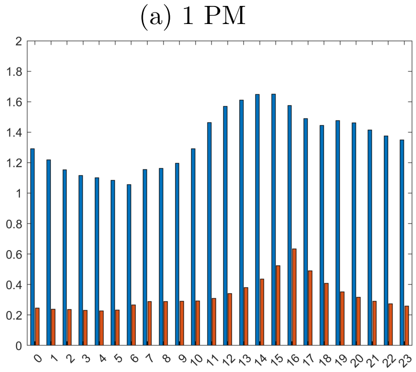

(a) 4 PM

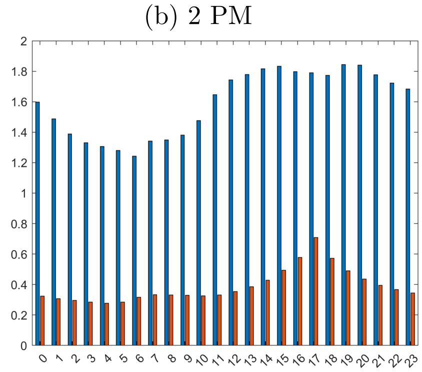

(b) 5 PM

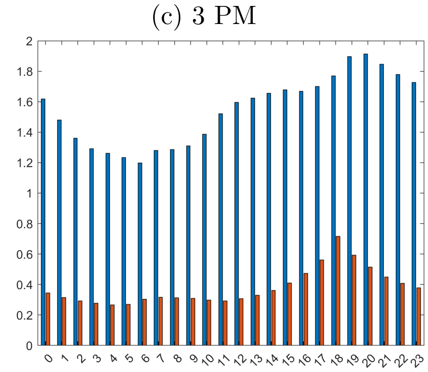

(c) 6 PM

Figure K.2: Absolute Elasticity of Low-Performing and High-performing Drivers