Algorithmic Wage Discrimination: Evidence from Uber

Jyotsana Kala*

April 13, 2026

Preliminary: Comments welcome. Please do not cite or circulate.

Click here for the latest version.

Abstract

Digital platforms are increasingly relying on algorithms to personalize pay for gig workers. This paper provides the first empirical evidence of algorithmic wage discrimination in the labor market, using novel data on Uber drivers. Exploiting Uber's transition to algorithmic wage setting in 2022, I employ a staggered difference-in-differences research design to estimate its causal impact on labor market outcomes. Algorithmic wage setting significantly decreases drivers' average earnings per trip and increases personalization in pay leading to inequality among drivers. However, it improves allocative efficiency by reducing drivers' wait time between trips and passenger fare. The decline in earnings is systematically larger for drivers with lower rejection rates. These heterogeneous effects are consistent with Uber's algorithm learning about drivers' reservation wages from their engagement histories—specifically, their acceptance and rejection decisions—and using this information to tailor wage offers.

JEL Classification: J71, J31, R41, C73

Keywords: discrimination, algorithm, ridesharing, dynamic learning

*University of California, Irvine: jyotsank@uci.edu. I am grateful to Guillaume Rocheteau, David Neumark, Michael Choi, Damon Clark and seminar participants at the UCI Macro Brownbag, UCI Labor-Public Seminar, UCSB Macro Lunch Talk and CMC Seminar Series for their invaluable advice and feedback. I am also grateful to The Workers' Algorithm Observatory for generously providing their data, as well as for their time, patience, and valuable comments.

1 Introduction

Gig work housed on digital platforms has nearly tripled between 2019 and 2023, rising from about 1.9 million platform workers to approximately 5.8 million in just four years (Garin et al., 2025). A characteristic feature of this online labor market is that digital platforms, through their operation, have access to granular worker data. Gig workers in this sector, such as drivers for *Uber* and *Lyft*, or delivery personnel for *Instacart* and *DoorDash*, continuously interact with the platform by logging in, waiting for jobs, accepting or rejecting tasks, etc., generating data with each interaction. These data enable platforms to infer worker preferences and inform decisions regarding task allocation and compensation. Recent advances in information technology and artificial intelligence have further facilitated the collection and processing of large volumes of such granular data through machine learning algorithms.

The increasing availability of granular worker data opens the door to *algorithmic wage discrimination,* a practice where individual workers are paid different hourly wages for the same task based on personal preferences inferred by an algorithm (Dubal, 2023). Conceptually, this parallels first-degree price discrimination in product markets, where firms set personalized prices based on consumers' willingness to pay. Varian (1989) proves that first-degree price discrimination leads to a Pareto-efficient allocation, but transfers surplus from consumers to firms. By analogy, algorithmic wage discrimination could improve allocative efficiency in labor markets while redistributing surplus from workers to firms, raising concerns about equity. Consequently, it raises a central empirical question: do firms engage in this form of discrimination in practice? If yes, what is the impact on workers and consumers? In this paper, I use a novel dataset on Uber's labor market—assembled through an interdisciplinary collaboration—to document the first empirical evidence of algorithmic wage discrimination, and estimate its causal effect on workers' earnings. Additionally, I test its welfare implications on the efficiency of the market by estimating its causal impact on passenger fares and wait times. Finally, I develop a game-theoretic model of dynamic learning that captures *how* Uber's algorithm infers drivers' preferences from their actions over time, shedding light on the mechanism behind the empirical results while providing key insights for workers.

Empirically testing for algorithmic wage discrimination is challenging as there is no publicly available, granular dataset on platform workers' wages and job characteristics, and their individual attributes to infer their preferences. To fill this gap, I collaborate with the Workers' Algorithm Observatory, an academic research team that crowdsources driver data, to assemble a novel, granular dataset on Uber drivers. The dataset is a rich panel consisting of approximately 2 million trips completed by 647 Uber drivers across the U.S. over a period of four years, from 2021 to 2025. Each trip is accompanied by detailed trip-level information (such as distance, duration, pick-up and drop-off time and location, driver pay, passenger fare) and is matched to the corresponding driver with their individual characteristics related to productivity (captured by ratings) and preferences (inferred uniquely from rejection rate).

To examine algorithmic wage discrimination, I exploit a change in Uber's wage setting policy for its drivers in 2022, when Uber transitioned from a traditional rate card system (fixed pay per-mile and per-minute) to an algorithm-based pay structure. The policy was introduced across markets at different times between March and Sept 2022, while some states continue to rely on the old rate card system of setting driver pay. This staggered rollout provides quasi-experimental variation that I use to estimate the causal effects of algorithmic wage setting through a staggered difference-in-differences design. If Uber uses its algorithm to personalize wages, the effects of the policy should vary systematically by drivers' inferred preferences. Uber's algorithm conjecturally infers these preferences from drivers' participation history, particularly from the offers they reject (Dubal, 2023). Drivers who turn down more offers are generally inferred to have a higher reservation wage—they need to be offered more to take a trip. So, I test whether the effects on drivers' earnings differ systematically with their average rejection rates—the likelihood of rejecting a ride offer. This allows me to evaluate whether the adoption of algorithmic wage setting led to differential pay outcomes across drivers, consistent with evidence of algorithmic discrimination.

I begin the analysis by estimating a regression model linking drivers' earnings per trip (before any surge payment, bonus, or tips) with a comprehensive set of trip characteristics including location-specific time effects up to the hour to account for dynamic factors like variation in local demand throughout the day (effectively, traffic conditions[1]). Additionally, I estimate individual driver fixed effects. I find that drivers' base pay is associated with an increase of 94 cents with every additional mile, 47 cents with every additional minute, and 44 cents when driving under surge. Next, I exploit the shift in Uber's policy to algorithmic wage setting, to ask: Did the variation in pay due to individual driver differences change when Uber shifted to algorithmic wage setting? If yes, is this change systematically linked to drivers' inferred preferences rather than their skills or experience?

The key findings indicate that the shift to algorithmic wage setting reduced average drivers' base pay by $1 per trip, or 6.3 percent of their average trip earnings, after accounting for worker and trip characteristics. Furthermore, the impact varies systematically with their rejection rate. Drivers who are willing to accept lower pay are actually the ones who see the biggest drop in their earnings under this new system. For the typical driver in my study, a 10 percentage point decrease in rejection rate (implying the driver is becoming less picky), could mean losing an extra $14 in monthly income (roughly 1 percent) due to the change in policy. I also find that markets adopting the new system

[1]Including location-specific time effects removes between-location variation and controls for spatial heterogeneity in earning patterns, including location-specific economic shocks. This helps control for dynamic pricing wherein wages fluctuate over time based on real-time factors (such as demand, supply, competition, time of day, season, or external conditions) and not based on individual characteristics. Wages change for all workers the same way in response to market conditions at any given point in time.

exhibit wider dispersion in drivers' per-trip earnings (after accounting for trip-specific characteristics and traffic conditions) as compared to markets that retained the old rate-card system. Individual driver characteristics like a driver's skills or local knowledge play a role in determining their earnings regardless of the pay system. Crucially, however, the wider dispersion in pay under algorithmic wage setting is consistent with the hypothesis that Uber's algorithm is inferring each driver's reservation wage and customizing its offers accordingly.

Having established that algorithmic wage setting has distributional impact on workers, I move to a normative question: how does it affect social welfare and market efficiency? The workers already employed by the platforms are worse off as the platforms now offer them wages closer to their reservation wage. However, when platforms are able to infer the reservation wages of workers, they can hire some of the more "choosy" workers, which were earlier left out of the market due to their high (and unknown) reservation wages. As a result, the market expands. In the data, I observe this expansion through the entry of new drivers with higher rejection rates once the policy takes effect. With more coverage of workers, the matching between consumers and workers improves, potentially decreasing both waiting times and prices for consumers. The data confirm that the policy led to a \$2.88 reduction in passenger fare per trip and a 1.68 minute reduction in waiting time for an average driver. This highlights a trade-off between equity and efficiency; while the use of algorithm leads to inequality among workers, it has some efficiency gains which can improve social welfare.

Next, to understand the potential mechanisms driving the labor market results of algorithmic wage setting, I construct a game-theoretic model of wage setting between Uber and its drivers. The model serves two purposes. First, it provides micro-foundation for the main empirical finding, explaining how lower rejection rates can lead to wage declines. Second, it offers insights into how drivers can improve their earnings under algortihmic wage setting. The model captures, in a stylized form, how Uber's algorithm learns about drivers' preferences from their observed actions over time. Its main advantage over the empirical analysis is that it explicitly incorporates the strategic dimension of driver behavior, which is difficult to identify empirically given that rejected wage offers are not observable. Within the model, drivers can strategically act in response to Uber's learning process, thereby influencing the evolution of wage offers. This strategic interaction has important implications for understanding both driver behavior and potential policy interventions.

The model assumes that drivers are ex-ante heterogeneous in their reservation wages, with each driver's reservation wage being private information to them. In the first period, Uber receives an exogenous trip request (from a passenger) and makes a take-it-or-leave-it wage offer to the driver. The driver responds by either accepting or rejecting the offer, taking into account that their action reveals information about their reservation wage and may influence future offers (à la Hart & Tirole, 1988). In the second period, Uber makes a revised wage offer to the driver based on their action in the first period. The model predicts the existence of a separating equilibrium in which Uber benefits from learning drivers' reservation wages early by posting a screening wage in the first period. Once this information is obtained, Uber perfectly price discriminates in the second period, extracting all potential informational rents from drivers. As a result, wages decline—more so for drivers with lower reservation wages—and wage dispersion increases in subsequent periods. These predictions are consistent with the empirical patterns and offer a structural explanation for understanding how observed outcomes emerge from the interaction between Uber's algorithm and drivers' strategic behavior.

I conclude by using the model to draw implications for drivers to improve their earnings in environments characterized by dynamic learning and personalized wage setting by platforms. A key feature of the model is that the low reservation wage drivers (less selective) have an incentive to mimic those with high reservation wage (more selective) by rejecting some offers in the first period to secure a higher pay later on. This helps the less selective drivers to gain information rents. To isolate this mechanism, I extend the model to assume that drivers are non-strategic (myopic), focusing solely on current-period payoffs. The results show that average earnings in a separating equilibrium are lower when drivers are myopic, as it becomes easier and less costly for Uber to separate them across a broader set of parameter values. This implies that, once algorithmic wage setting is implemented, drivers who respond strategically over time can mitigate wage reductions, effectively pushing back against the algorithm.

Layout The remainder of the paper is structured as follows. Section 2 reviews the relevant literature. Section 3 describes the data and provides descriptive summary statistics. Section 4 discusses the empirical strategy. Section 5 presents the results and offers interpretation. Section 6 develops the theoretical model to understand the mechanisms and discusses policy recommendations, and Section 7 concludes.

2 Literature Review

This paper contributes to the emerging interdisciplinary literature on algorithmic personalized pricing in labor markets. It is the first to provide causal empirical evidence on algorithmic wage discrimination.

In legal scholarship, Teachout (2023) introduces the concept of algorithmic personalized wages and argues that such practices could exacerbate economic and racial inequality, proposing regulatory measures such as a ban on first-degree labor price discrimination. This normative concern is supported by Dubal's (2023) ethnographic evidence detailing the observed shift to algorithmic wage setting and its perceived impact on Uber drivers in the San Francisco Bay Area. I complement this qualitative work with a quantitative analysis of Uber's algorithmic wage setting policy and test whether it leads to personalization. My findings also call into question the proposed blanket bans on algorithmic wage setting, which—while appealing from an equity standpoint—may reduce overall market efficiency. I instead highlight policy approaches that balance equity and efficiency, such as rent-sharing mechanisms (e.g., lump-sum taxation of platform profits redistributed to workers) or regulations that facilitate entry and competition in the ride-hailing market.

Concurrently, in the computer science literature, Binns et al. (2025) empirically analyze Uber's algorithmic wage setting policy in the UK market, finding negative correlation between the adoption of algorithmic pay and driver earnings. Similarly, Pandey & Caliskan (2021) investigate the consumer side of the market, exploring whether AI pricing algorithms in the rideshare industry lead to disparate impact on consumers based on demographics. Relative to these studies, my paper studies the impact of Uber's algorithm in the U.S. market and estimates its causal effect on worker earnings and consumer fares. I also examine the mechanism explaining why such disparate impacts arise, and whether they are consistent with first-degree price discrimination in the labor market.

In economics, research on personalized pricing has primarily focused on product markets (Grennan, 2013; Shiller, 2013, 2016; Hannak et al., 2014; Srivastava, 2020; Dube & Misra, 2023; Brown & MacKay, 2023). This paper makes two contributions relative to this parallel work. First, I examine personalized pricing in the labor market, an empirically more challenging setting: unlike identical products sold at different prices, no two tasks assigned by digital platforms are exactly the same, complicating the distinction between dynamic (task-specific) and personalized (worker-specific) wage variation. Second, I use transaction-level data to infer workers' reservation wages, and employ a causal research design to credibly identify the effects of algorithmic wage setting. The underlying mechanism—personalization based on user engagement—parallels that in product markets.

This paper also contributes to a growing body of research using Uber data to study labor market outcomes in the rideshare industry (Cohen et al., 2016; Hall & Krueger, 2018; Hall et al., 2023; Chen et al., 2025; Castillo, 2025). Cook et al. (2021) examine the drivers of the gender earnings gap among one million Uber drivers, while Caldwell & Oehlsen (2023) use a field experiment to estimate labor supply elasticities by gender and outside options. These studies focus on group-level characteristics such as gender, whereas my analysis targets first-degree labor price discrimination, isolating worker-specific preferences as the mechanism of interest. Even after controlling for productivity-related factors specified in these papers, I find that systematic pay gaps persist among drivers. Chen et al. (2020) combine a field experiment with high-frequency panel data on driver wages and work decisions to estimate reservation wages and the value of job flexibility, showing theoretically how firms could use heterogeneity in preferences to improve efficiency. I test this mechanism empirically, using reservation wage proxies in my data to assess whether Uber's 2022 policy change leveraged heterogeneity in driver preferences to improve allocation efficiency.

A rapidly developing literature examines monopsony in online labor markets. Dube et al. (2020)

estimate labor supply elasticities on Amazon Mechanical Turk and document substantial monopsony power, while Fisher (2025) develops a two-sided model of digital platforms calibrated to Uber data, showing that Uber depresses drivers' earnings by roughly 15 percent. Compared to them, I find evidence of monopsony power in Uber's labor market implemented through personalized wage markdowns.

Finally, I contribute to the literature on signaling, learning, and repeated bargaining by modeling dynamic wage setting in algorithmically managed labor markets. Hart & Tirole (1988) model a long-term relationship between a seller and a buyer with private valuations. As trade progresses, buyers' valuations are partially revealed through their actions. They show that under non-commitment, the finite-horizon equilibrium converges to a no-discrimination outcome as the horizon tends to infinity. I adapt their framework into a simplified two-period model and characterize a separating equilibrium in finite time. Unlike Kennan (2001), I assume permanent rather than persistent valuation types for the informed agent. The mechanism in my model also closely parallels models of the ratchet effect found in industrial organization, regulation, and labor economics. In particular, Laffont & Tirole (1988) show that in a two-period principal-agent model, when uncertainty about agent ability is small, optimal contracts involve significant pooling. Similarly, in my setting, the degree of pooling versus separation depends on the underlying heterogeneity among drivers—when heterogeneity is low, pooling is high.

3 Data and Descriptive Statistics

The empirical analysis utilizes data from the Workers' Algorithm Observatory (WAO) on Uber's labor market. WAO is a research collaborative that crowdsources driver data to study the black-box algorithmic systems that determine pay, schedule, and other working conditions on gig platforms[2]. The sample consists of 1.91 million trips completed between Jan 2021-Jun 2025 by 647 Uber drivers across 46 states in the U.S.[3] Each observation is a trip linked to a specific driver. The data include detailed information on the trip, such as distance, duration, date, time, pickup and drop-off latitude and longitude, and pay information like driver earnings, tips, bonuses, Uber's commission, congestion surcharges, and passenger fares. Additionally, each trip specifies the ride category assigned by Uber, such as UberX, UberXL, Comfort, Uber Green or others. The data also identify the driver of each trip and include driver-specific characteristics like average ratings, rejection rate (the average probability that the driver rejects a ride), original hire date, and model and make of the driver's car. I make two sample restrictions to the data. I drop the bottom and the top one percentile of driver earnings, and drivers with less than 10 trips in the sample, both to avoid outliers influencing the empirical results. Table 1 reports the descriptive statistics for my sample, separately by trip characteristics and driver

[2]For more details on data collection by WAO, see Calacci et al. (2025).

[3]Alaska, North Dakota, South Dakota and Vermont are not represented in the sample.

Table 1: Descriptive Statistics

	Mean
Panel A. Trip characteristics	
Distance (miles)	7.75
Duration (min)	16.22
Base pay ($)	16.22
Passenger Fare ($)	25.74
Uber's take rate	0.23
Driver's take rate	0.64
Wait time (min)	17.41
Proportion of UberX	0.78
Proportion of Comfort	0.09
Proportion of UberXL	0.06
Proportion of Uber Green	0.05
Number of trips	1,909,272
Panel B. Driver characteristics	
Monthly rides	129.89
Monthly earnings	2102.83
Experience (months) since hire	79.03
Experience (months) in sample	21.07
Experience (rides) in sample	2956.33
Rating	4.96
Rejection Rate	0.26
Number of drivers	647

Notes: This table reports descriptive statistics for the sample of trips completed by 647 Uber drivers between Jan 2021 and June 2025. Data are from the Workers' Algorithm Observatory.

characteristics. Appendix A displays many of these descriptive statistics graphically.

3.1 Trip characteristics

The average trip distance in the sample is 7.7 miles and the average trip duration is 16.2 minutes. Around half of the trips are completed on the weekend, and a quarter between 10p.m.-2a.m. 78% of trips are classified as UberX rides. 75% of trips originate from Washington, Colorado, California, Florida and Oregon. Figure 1 shows the spatial distribution of trips in the sample.

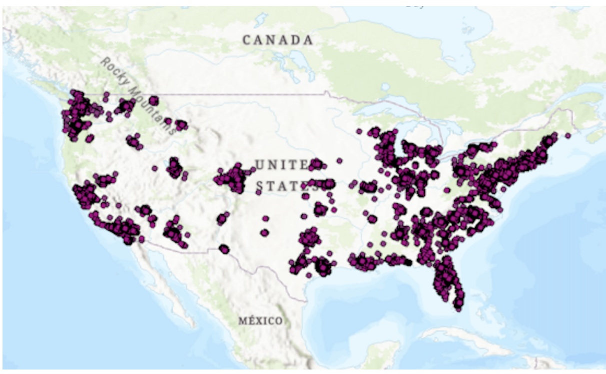

Figure 1: Spatial Distribution of Trips in the U.S.

Base pay per trip is defined as drivers' earnings before any surge payment, bonus, or tips. The average base pay per trip is \$16.2. Given the pick-up and drop-off time of each trip, I can observe the time between consecutive trips in the sample. When this time is less than an hour, I interpret it as the wait time, i.e., the time a driver spends waiting between two consecutive trips[4]. Drivers wait for about 17 minutes on average between trips. After factoring in both the time spent driving and waiting for the next trip, drivers earn about \$30 for each engaged hour[5]. Drivers' base pay is strongly positively correlated with both trip distance and duration. The average pass-through from passenger fare to driver base pay is 54%. Furthermore, Uber's take rate, that is, the percentage of passenger fare taken by Uber as service fee, increases as the per-minute fare increases. Consequently, the driver's take rate (share of the fare that goes to the driver) falls with the per-minute fare (see Figure 2). So, drivers are relatively better off taking several low-fare trips instead of a single high-fare trip.

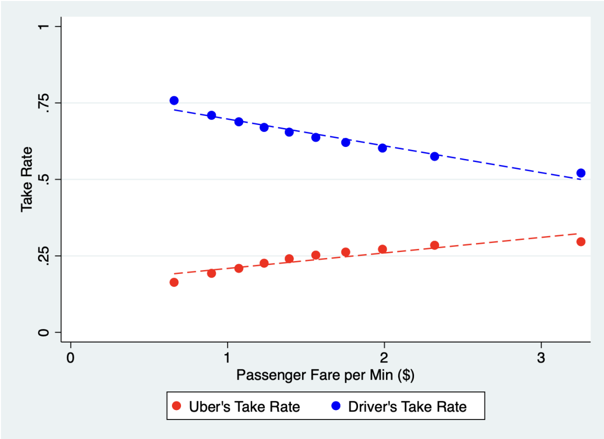

Figure 2: Take Rate of Uber Driver by Per-minute Passenger Fare

3.2 Driver characteristics

An average driver completes 130 trips per month and drives for 21 months in the sample. I use the original hire date and the last trip in sample to calculate a driver's total experience at Uber. An average driver has over six years of experience at Uber[6]. The productivity of drivers is measured through their average ratings. However, ratings are highly clustered in the narrow range of 4.8 to 5, with a sample mean of 4.96, limiting their informativeness. Tips, by contrast, provide a more meaningful measure of productivity as they capture passengers' revealed preferences about drivers' productivity. In my sample, tips amount to roughly 12% of the fare on average and exhibit a strong positive correlation with ratings (see Figure 3).

[4]Rides with wait times exceeding one hour are considered part of a different shift.

[5]This aligns with Uber's CEO Dara Khosrowshahi's earnings estimate of \$30 per engaged hour for drivers.

[6]This does not necessarily mean that the driver has been actively driving over the past six years. It simply indicates that drivers in the sample are not new to Uber, and they might appear in and out of the sample.

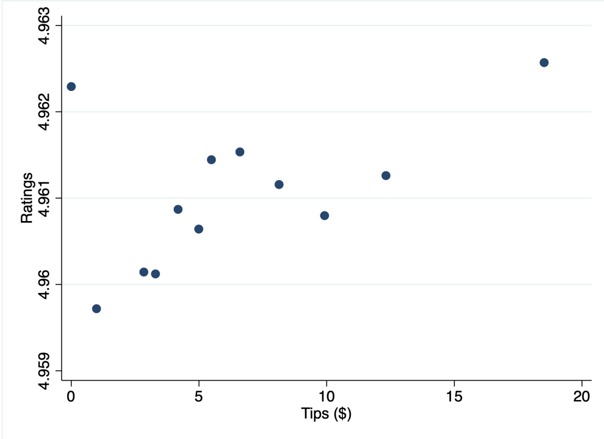

Figure 3: Tips and Ratings given by Passengers

Drivers' engagement with the platform is captured by their average rejection rate, defined as the probability of rejecting a ride offer based on their most recent three months of driving. The average rejection rate in my sample is 26 percent. Moreover, there are 25 drivers with 0 percent rejection rate, implying that they accepted every offer they were made over the last three months of driving. Figure 4 shows the distribution of drivers' rejection rates.

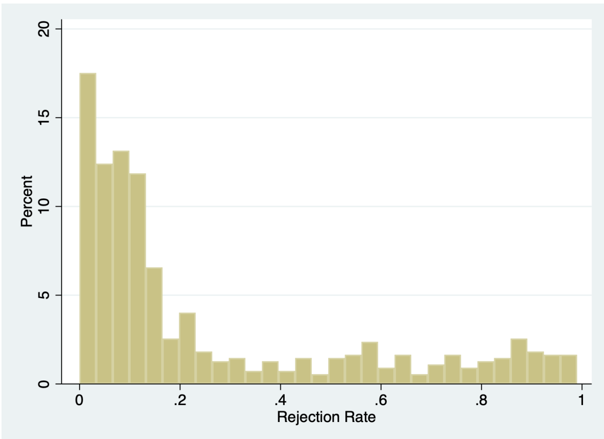

Figure 4: Distribution of Drivers' Rejection Rates

4 Empirical Strategy

The analysis begins with a trip-level baseline regression of driver's base pay on several trip characteristics and driver fixed effects to examine the determinants of drivers' earnings. This baseline specification is of the form:

$$Y_{i,x,z,t} = \alpha + \beta X_{x,z,t} + \lambda_z + \tau_t + \mu_{zt(h)} + \delta_i + \varepsilon_{i,x,z,t} \quad (1)$$

where $Y_{i,x,z,t}$ is the base pay of driver i for a trip x taken in zipcode z at time t. $X_{x,z,t}$ is a vector of trip characteristics, including distance and duration of the trip, a dummy variable for whether the

Table 2: Baseline regression

	(1) Pay	(2) Pay
Distance	0.94***	0.94***
	(0.03)	(0.03)
Duration	0.48***	0.47***
	(0.01)	(0.02)
Surge		0.44***
		(0.06)
Non-std ride		2.93***
		(0.10)
Time F.E.	Y	Y
Location F.E.	Y	Y
Time X Location F.E.	Y	Y
Individual F.E.	Y	Y
N	1,881,473	1,881,473
R^2	0.88	0.89

Notes: This table reports the results of the baseline regression estimated on the sample as described in the notes to Table 1. The dependent variable is driver pay per trip. Column 1 controls for distance and duration of the trip. Column 2 adds additional controls like whether the trip was taken under surge, and whether the trip was a non-standard ride (trips other than UberX). Results control for time, location, time-by-location, and individual driver fixed effects. Reported standard errors are clustered at the individual driver level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

trip is subject to surge pricing and a dummy variable for whether the trip constitutes a non-standard ride. A non-standard ride is UberXL, Comfort and Uber Green, all of which are priced higher for passengers than an UberX for their added features. λ_z includes two sets of location fixed effects characterized by pick-up and drop-off zipcodes of the trip. τ_t are a set of time fixed effects of the trip including the year, month, day of the week, and hour of the trip. I also interact the location and time fixed effects[7], given by μ_{zt}, to control for the hourly change in traffic conditions, or more broadly, supply and demand factors in each zipcode. δ_i are individual driver fixed effects. Standard errors are clustered at the individual driver level.

The results are reported in Table 2. As expected, base pay per trip is strongly positively correlated with trip distance and duration. Ceteris paribus, a one mile increase in the trip distance corresponds to an average increase of 94 cents in base pay. Additionally, each additional minute of duration is associated with an average increase of 48 cents in base pay. The results hold even when controlling for surge pricing and requesting a non-standard ride, both of which are positively associated with base pay.

[7]Hour of the trip is interacted with the zip code of the start location.

The baseline specification also estimates the individual driver fixed effects, δ_i, which capture how each driver's average earnings differ from the reference driver due to their individual characteristics after controlling for trip-related factors. The baseline model explains about 89% of variation in driver pay, 2% of which is attributable to individual driver fixed effects. Figure 5 shows the distribution of these individual driver fixed effects (in \$)[8].

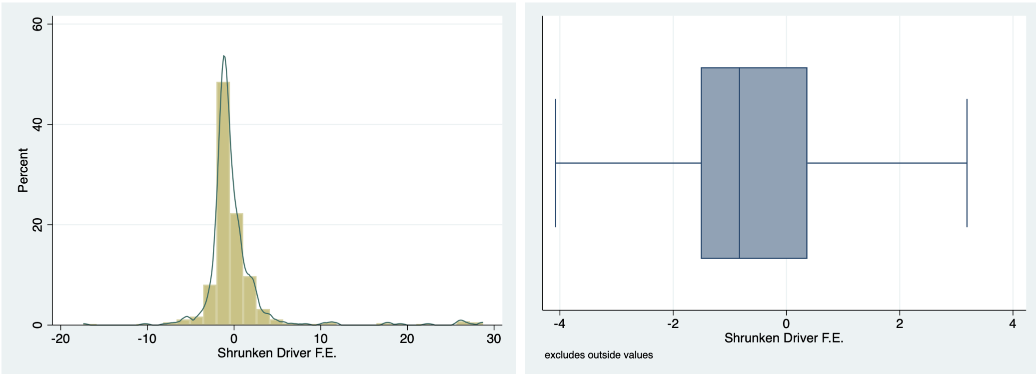

Figure 5: δ estimated from baseline regression

This confirms that even after controlling for trip-related factors such as distance, duration, and traffic conditions, some part of the variation in drivers' pay is correlated with their individual characteristics. Next, I investigate which specific individual characteristics account for this variation. For this, I augment the baseline regression with individual driver characteristics instead of individual driver fixed effects. The specification is as follows:

$$Y_{i,x,z,t} = \alpha + \beta X_{x,z,t} + \omega W_i + \lambda_z + \tau_t + \mu_{zt} + \varepsilon_{i,x,z,t} \tag{2}$$

where W_i is a vector of characteristics of driver i. Driver characteristics include their ratings, experience (as measured by the total number of trips completed in the sample and total number of months since hire), rejection rate, average wait time between trips and an indicator variable for whether they multi-app on Lyft (another ride sharing platform). Table 3 shows the results. Conditional on trip controls, drivers' base pay is positively associated with their ratings, average wait time between trips and rejection rate. This indicate that when drivers wait between trips, they tend to select themselves into accepting better paying trips. Drivers who are better rated are awarded with higher pay. For every 10 percentage point increase in ratings, base pay increases by 50 cents per trip. Drivers who multi-app on other rideshare platforms tend to accept lower pay, although not significantly so.

Conditional on trip-level controls, drivers' base pay is positively associated with their ratings, average wait time between trips, and rejection rate. This pattern suggests that when drivers wait

[8] The distribution of estimated driver fixed effects was first demeaned and then shrunk to account for measurement error. The box plot in the right panel of Figure 5 excludes outliers, defined as values beyond 1.5 times the inter-quartile range on either end of the distribution.

Table 3: Regression with Driver characteristics

	(1) Pay
Distance	0.97***
	(0.03)
Duration	0.48***
	(0.02)
Surge	0.61***
	(0.14)
Non-std ride	3.72***
	(0.27)
Rating	5.33**
	(2.38)
Experience (rides)	0.00003
	(0.00003)
Experience (months)	-0.003
	(0.003)
Avg Wait time	0.21**
	(0.090)
Multiapp	-0.39
	(0.25)
Rejection rate	1.32**
	(0.54)
Time F.E.	Y
Location F.E.	Y
Time X Location F.E.	Y
Individual F.E.	N
N	1,421,656
R^2	0.88

Notes: This table reports the results of the baseline regression augmented with driver characteristics—experience (number of rides completed and months spent driving), average wait time between trips, rejection rate, and whether the driver multiapps on Lyft. The sample includes trips by 551 Uber drivers from Jan 2021 to June 2025, restricted to those with complete characteristics. Results control for time, location, and time-by-location fixed effects. Reported standard errors are clustered at the individual driver level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

longer between trips, they tend to self-select into accepting higher-paying trips. Higher-rated drivers are also rewarded with higher pay: a 10-percentage-point increase in ratings is associated with a 50 cents increase in base pay per trip. In contrast, base pay does not appear to be correlated with drivers' experience. Drivers who multi-app across other rideshare platforms tend to earn less, although the difference is not statistically significant.

So far, I have demonstrated that pay varies among drivers completing the same trip based on individual driver characteristics. This variation can stem from multiple factors— 1) differences in driver skill levels (as reflected in their ratings), 2) differences in driver experiences (capturing learning effects as experienced drivers can identify more profitable areas/rides), or 3) differences in driver preferences, which impacts their acceptance of offers, i.e., Uber offering wages from a common distribution while drivers selecting wages based on their reservation wages. Consequently, the observed pay variation may simply reflect the underlying distribution of driver skills or/and preferences. Additionally, I conjecture that some of this variation arises from Uber using its algorithm to gradually learn about drivers' reservation wages and adjust offers accordingly.

Without comprehensive data on the complete history of a driver's engagement with the platform, i.e., data not only on the jobs they accept but also on those they reject, it is challenging to causally identify the role of Uber's learning on pay variation. To achieve this, I exploit an exogenous change in Uber's pricing policy for drivers in August 2022, when Uber transitioned from a per-mile rate card to an algorithm-based pay structure. Theoretically, the use of an algorithm gives the platform immense capability to process individual driver data and potentially discriminate by offering personalized wages. By leveraging this policy shift, I aim to isolate and quantify the portion of the temporal change in the distribution of drivers' fixed effects that can be attributed to Uber's use of algorithm to set pay.

4.1 Policy Background- Upfront Pricing

Uber's new pricing policy for drivers, called Upfront Pricing, was introduced amidst ongoing debates over whether Uber drivers should be classified as employees or independent contractors. To provide some background, Proposition 22, passed in California in November 2020, allowed app-based ride-hailing and delivery companies, such as Uber and Lyft, to classify their drivers as independent contractors rather than employees. This measure exempted these companies from Assembly Bill 5 (AB5), which mandated employee classification for many gig workers, thereby granting them benefits such as minimum wage, overtime pay, and unemployment insurance. Instead, Proposition 22 established an alternative set of benefits, including a minimum earnings guarantee, health care subsidies for eligible drivers, and occupational accident insurance. The measure was strongly supported by digital platforms, which collectively spent over $200 million on its campaign, making it one of the most expensive ballot initiatives in U.S. history.

In this context, driver unions frequently argued that gig workers in the platform economy, including Uber drivers, should not be classified as independent contractors, as they lacked the ability to set their own wages. Moreover, they were unable to view key details about a ride before accepting it. Under

Uber's previous rate-card pricing policy, drivers could see only the approximate distance to the pickup location but had no visibility into the ride's total distance, duration, destination, or, most importantly, the fare. Partly in response to these concerns, Uber's pricing policy was revised to provide drivers with complete ride information—particularly the fare—before they decided whether to accept or reject a ride[9]. This change aimed to enhance transparency and give drivers greater autonomy in managing their work. However, alongside this increased visibility, another fundamental shift occurred: fares were no longer determined by a fixed rate-card system but instead by an algorithm (the specifics of which Uber does not disclose to the drivers).

Several other changes were introduced as part of the new policy, including a destination filter that allowed drivers to find rides that took them closer to home at the end of their shifts, and the trip radar feature, which sent simultaneous trip offers to all drivers. Additionally, the Uber Pro program was launched, offering drivers discounts on fuel and select retailers. Another key change was rate rebalancing, where fares were updated in real-time if traffic conditions worsened during a trip, extending the time required to reach the destination. It was also announced that longer trips would pay less per mile than shorter trips, aiming to incentivize drivers to accept multiple short trips rather than waiting for a single long trip.

The new pricing policy was rolled out in two phases: Phase 1 in March 2022, covering 6 states, and Phase 2 in September 2022, expanding to 25 more states. The remaining 15 states in the sample continue to follow the previous rate-card system.[10] This staggered rollout provides a unique opportunity to exploit a difference-in-differences research design to estimate the causal effect of algortihmic wage setting on market outcomes.

4.2 Staggered Difference-in-Differences

To examine the effect of the policy on earnings, I begin by estimating a standard two-way fixed effects (TWFE) specification as follows:

$$Y_{i,x,z,s,t} = \alpha + \beta X_{x,z,s,t} + \gamma_1 Policy_{s,t} + \gamma_2 Policy_{s,t} * RR_i + \lambda_z + \tau_t + \mu_{zt(h)} + \delta_i + \varepsilon_{i,x,z,s,t} \tag{4}$$

where $Y_{i,x,z,s,t}$ is the base pay of driver i for a trip x in zipcode z of state s at time t, and RR_i is the average rejection rate which controls for drivers' preferences. $Policy_{s,t}$ is the value of the treatment for all drivers operating in state s at time t. Specifically, $Policy_{s,t}$ is 0 for all trips taken

[9]See Figure A9 for a comparison of information available to drivers in the ride offer before and after the pricing policy changed.

[10]Policy timing is based on drivers' account of when they saw the new policy being implemented in their markets. Uber does not provide any information about the official timing of policy rollout. See Table A1 for policy timing by states.

by all drivers before March 2022 and switches to 1 for trips taken by drivers in treated states as the policy change takes effect. Other control variables are same as the baseline specification. Standard errors are clustered at the state level. Conditional on drivers' skills and preferences, γ_1 recovers the effect of the policy on drivers' average trip earnings. Additionally, I allow this effect to vary with drivers' preferences, i.e., their rejection rate. I take rejection rate as a proxy for the reservation wage of the driver. Higher the rejection rate, more picky is the driver as they need to be offered more to accept the ride. The policy allows Uber to infer a driver's reservation wage from their rejection rate and adjust their pay accordingly. As a result, drivers with lower rejection rates should see larger wage reductions than those with higher rejection rates. In addition, I also look at the effect of the policy on other outcomes like passenger fare and the average wait time of the driver.

Recent econometric work has highlighted potential biases in panel data estimates when adoption of a treatment is staggered and there are pre-trends or heterogeneous (dynamic) treatment effects (e.g., Callaway et al., 2024; Wooldridge, 2021). This could bias the causal effects estimated from specification (4). The methods that have been developed for addressing these biases are applicable to settings of dummy treatments, and treatments that do not turn on repeatedly (both of which are applicable to my context). Thus, I apply an alternative two-stage estimation framework proposed by Gardner et al. (2024) to identify the unbiased causal effect of the treatment. In this framework, group and period effects are identified in the first stage from the sample of untreated observations, and average treatment effects are identified in the second stage by comparing treated and untreated outcomes, after removing these group and period effects. The two-stage difference-in-differences (DD) specification is as follows:

$$Y_{i,x,z,s,t} = \alpha + \beta X_{x,z,s,t} + \gamma Policy_{s,t} + \lambda_z + \tau_t + \delta_i + \varepsilon_{i,x,z,s,t} \tag{5}$$

where γ recovers the causal effect of the policy on market outcomes.[11]

5 Results

In this section, I examine the effects of the algorithmic wage setting using the empirical strategy described above.

5.1 Earnings and Personalization

I first analyze the trends in earnings for the treated and control states. Figure 6 plots the average hourly driver earnings over time for two groups of markets: those that remained on the old rate card

[11] The "did2s" procedure itself primarily focuses on estimating the main treatment effect (ATT) robustly in a staggered DID setting which restricts my ability to estimate the interaction of the treatment with covariates.

system (shaded line) and those that switched to Upfront pricing after 2022 (bold line). Prior to the switch, earnings in both groups followed broadly parallel, upward trends. After the policy change, however, average hourly earnings in Upfront Pricing markets initially fell before stabilizing, while earnings in unaffected markets continued to rise.

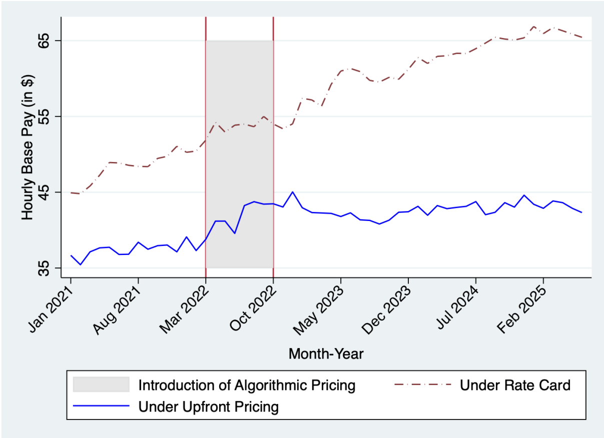

Figure 6: Average trip earnings by Treatment

Next, I causally test the impact of the policy change. If Uber's algorithm learns about drivers' reservation wages and systematically tailors offers in response, we should observe a widening dispersion in the distribution of driver fixed effects over time, accompanied by a decline in average earnings. As Uber refines its estimates of individual reservation wages, it can personalize wage offers more precisely. This refinement increases the dispersion in fixed effects—by aligning pay more closely with individual reservation wages—while reducing average earnings as Uber captures the informational rents that drivers previously retained through their private knowledge. To estimate the causal effect of the policy on average trip earnings, Table 4 presents the results from the two-way fixed effects specification described in equation 4.

The results show that the change in Uber's pricing policy to algorithmic wage setting leads to a reduction in drivers' average base pay by \$1.27 per trip, which is about 8% of their average trip earnings. The negative effect of the policy is robust to the exclusion of the mechanism from the specification, as seen in the first column of Table 4. Moreover, the reduction in pay is not uniform across all drivers; it varies systematically with their rejection rates. A 10 percentage point decrease in rejection rate leads to a further 11 cents reduction in pay due to the policy. Thus, drivers who are willing to accept lower pay (and hence, have lower rejection rates) are actually the ones who see the biggest drop in their earnings under this new system.

Further, I estimate the causal effect of the policy using the two-stage difference-in-differences approach (equation 5). I recover an estimate of -0.69 which is quite similar to the two-way fixed effects estimate, and implies that the policy led to a reduction in drivers' average base pay by \$0.69 per trip. Since treatment switches on at different times for different states, a graph in "event time"

Table 4: Two-Way Fixed Effects Regression

	Dependent Variable: Pay	
	(1) Baseline	(2) Interaction
Distance	0.97***	0.97***
	(0.19)	(0.19)
Duration	0.49***	0.49***
	(0.08)	(0.08)
Surge	0.50***	0.50***
	(0.12)	(0.13)
Ride Type	2.81***	2.81***
	(0.32)	(0.32)
Policy	-0.76	-1.27
	(0.63)	(0.68)
Policy * Rejection Rate		1.08***
		(0.59)
Time F.E.	Y	Y
Location F.E.	Y	Y
Time X Location F.E.	Y	Y
Individual F.E.	Y	Y
N	1,408,915	1,408,915
R^2	0.89	0.89

Notes: This table reports the two-way fixed effects regression results of the impact of algorithmic wage setting on driver pay per trip. All specifications include baseline controls, and time, location and time-by-location fixed effects. Column 1 estimates the average effect of the policy. Column 2 interacts rejection rate with the treatment. Reported standard errors are clustered at the state level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

is more informative. Figure 7 displays the event-study graph for average trip earnings using the two-stage difference-in-differences estimation. The causal effect of the policy on drivers' earnings is positive (although insignificant) in the year the treatment is introduced, but becomes negative in subsequent years. This provides strong evidence that the shift to algorithmic wage setting had a detrimental impact on drivers' average earnings.

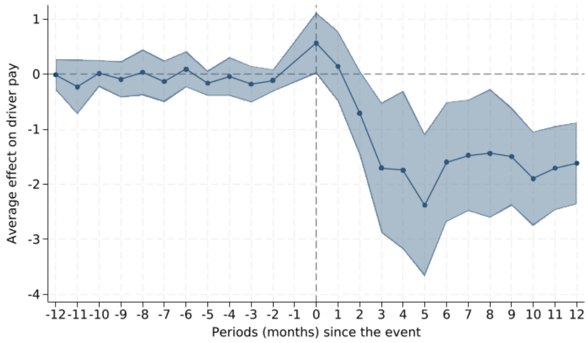

Figure 7: Event Study- Driver Pay

After estimating the impact on earnings, I examine the extent of personalization in pay. For this, I study the distribution of drivers' fixed effects by the status of treatment as shown in Figure 8. These fixed effects distributions are created by estimating equation (1) separately for four different groups-control group pre policy, control group post policy, treated group pre policy and treated group post policy. Before algorithmic wage setting was implemented, the distribution of individual driver fixed effects between the control and treated states (in red) looked pretty much similar. This shows that even with the old rate card system for determining driver pay, the earnings differed among drivers based on their individual characteristics. These characteristics could be related to drivers' productivity or differences in their tacit knowledge about which areas/times are better for driving. However, after the policy changed in 2022, there is a significant spread in the distribution of driver fixed effects in the treated states, while no such change in the distribution in control states (in blue). The only thing that changes post policy is Uber's ability to learn about drivers' behavior and use it to determine pay. Thus, this spread in distribution is consistent with the hypothesis that under algorithmic pricing policy, Uber learns about each drivers' reservation wage over time and tailors their wage offers closer to their reservation wage.

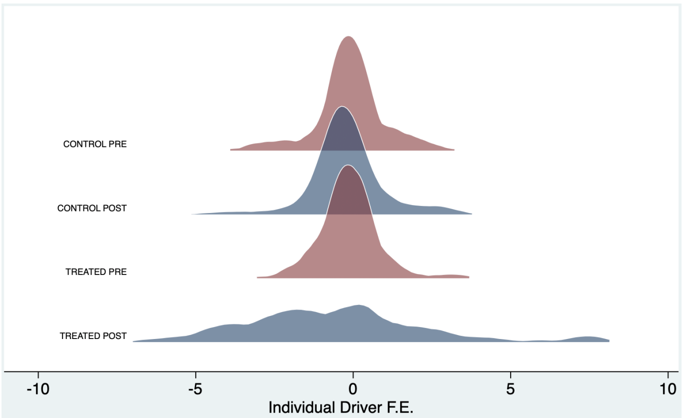

Figure 8: Distribution of Driver Fixed Effects by Treatment

5.2 Identification issues

A possible threat to identification is that the effect of policy might be correlated with other control variables. To alleviate this concern, I estimate equation (4) separately by treatment status. Table 5 shows the results. Upon estimating separate effects of every covariate on control and treated states, I find that the effect of rejection rate on pay is insignificant among control states while significantly positive among the treated states. This implies that the effect of the treatment varying by rejection rate is not merely an artefact of choosing a particular covariate for interaction; it is recovered even after a full interaction of all covariates by treatment.

Another key threat to identification is that the roll out of the policy could be endogenous, i.e., the policy was rolled out in markets where drivers were already less selective to begin with, and hence led

Table 5: Two-way Fixed Effects by Treatment

	(1) Control	(2) Treated
Distance	1.10***	0.55***
	(0.14)	(0.03)
Duration	0.53***	0.24***
	(0.05)	(0.03)
Surge	0.73***	0.64***
	(0.21)	(0.15)
Ride Type	3.41***	3.91***
	(0.36)	(0.63)
Rating	2.74***	3.35
	(0.59)	(3.93)
Rejection Rate	-0.27	2.48***
	(0.53)	(0.74)
Time F.E.	Y	Y
Location F.E.	Y	Y
Time X Location F.E.	Y	Y
Mean	17.39	11.51
N	1,125,098	288,840
R^2	0.91	0.83

Notes: This table reports the results of the regressions of driver pay per trip on baseline controls and driver characteristics capturing skills and preferences—such as rating and rejection rate—estimated separately for control and treated units. All specifications include time, location, and time-by-location fixed effects, and report means of the dependent variable. Reported standard errors are clustered at the individual driver level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

to lower earnings after the policy. To examine this, I plot the distribution of drivers' rejection rates in never-treated and ever-treated markets before treatment took effect. I also restrict to states which have both never-treated and ever-treated markets to control for any state-level policies that may affect rejection rates systematically in either of these markets. Figure 9 compares these distributions and indicate that there is no meaningful differences in driver preferences between the control and to-be-treated markets before the treatment switched on. This alleviates endogeneity concerns regarding the timing of policy rollout in different markets.

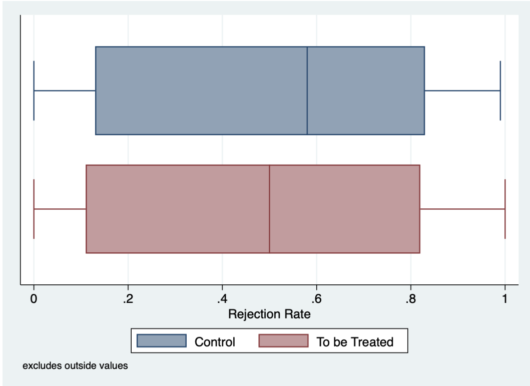

Figure 9: Rejection Rate over Treatment

A potential identification concern arises if the policy itself affects drivers' rejection rates, rendering rejection rate a "bad control" in the spirit of Angrist and Pischke (2009). Since the policy alters driver compensation, it may in turn influence the propensity to accept or decline ride requests. Ideally, one would exploit within-driver variation in rejection rates over time, using only pre-policy values to characterize driver selectivity. Unfortunately, the data do not record rejection rates longitudinally; rather, this variable is computed as an average over a driver's recent activity (last three-months of driving).

To recover a plausible pre-policy rejection rate, I adopt the following imputation strategy. I estimate a driver-level regression of rejection rate on a vector of driver characteristics — including average wait time, rating, tenure (measured in both months and completed rides), and a multi-apping indicator — along with the driver's estimated probability of treatment. Formally, I estimate:

$$RR_i = \alpha + \beta X_i + \gamma Pr(\text{Treat})_i + \varepsilon_i \tag{6}$$

where X_i comprises the driver characteristics enumerated above. I then impute a counterfactual, pre-policy rejection rate for each driver by setting $Pr(Treat)_i = 0$, thereby netting out the component of observed rejection behavior attributable to treatment exposure. This yields a driver-specific imputed rejection rate that approximates what selectivity would have looked like absent the policy.

Table 6 presents estimates of specification (4) using both actual and imputed rejection rates alongside the baseline. The results are stable across columns: the overall effect of the policy on driver earnings is negative, and the interaction with rejection rate is positive and statistically significant when using actual rejection rates, indicating that more selective drivers fare relatively better under the policy. The attenuation of the interaction coefficient under imputed rejection rates is consistent with classical measurement error, and the qualitative conclusions are unchanged. Taken together, these robustness checks allay concerns that endogenous sorting in rejection behavior is driving the main findings.

Table 6: TWFE Regression with Imputed Rejection Rate

	(1) Baseline	(2) Actual Rejection Rate	(3) Imputed Rejection Rate
Ride Controls[a]	Yes	Yes	Yes
Policy (γ_1)	-0.76	-1.27*	-1.85*
	(0.63)	(0.68)	(1.06)
Policy × RR (γ_2)		**1.08***	**6.63**
		(0.59)	(3.97)
Fixed Effects	Yes	Yes	Yes
Observations	1,408,915	1,408,915	1,408,915
R^2	0.89	0.89	0.89

Notes: [a] Includes Distance, Duration, Surge, and Ride Type.
Actual Rejection Rate is observed; Imputed Rejection Rate is estimated from driver history.
Standard errors clustered at the driver level. * p<0.1, ** p<0.05, *** p<0.01.

5.3 Fares and Wait Time

So far, I document the adverse effect of algorithmic wage setting on drivers' earnings and how it increases the extent of personalization in their pay. However, from a social planner's point of view, algorithmic wage setting gets rid of an important market friction—information friction due to workers holding private information about their preferences to work. When the algorithm can figure out these preferences, it can better match drivers with their preferred rides, hence reducing waiting time for both drivers and customers. I observe this in the data by regressing average wait time of the driver, as per the two-way fixed effects specification in Table (4). Additionally, I allow the effect of the treatment to vary with the rejection rate. Results are shown in Table 7.

I find that the policy reduced drivers' wait time between trips by 1.61 minutes. Moreover, a 10 percentage point decrease in the rejection rate led to an additional decrease in the average waiting time by 4 seconds due to the policy (though not significantly so). The effect on wait time varies by rejection rate presumably because once a passenger requests a ride, Uber systematically offers wages to drivers starting from the low reservation type and making its way sequentially to the high reservation wage type. This highlights a crucial tradeoff for drivers: those with lower rejection rates, on one hand, experience a steeper drop in earnings per trip but on the other hand, wait lesser in between trips after the policy.

The use of an algorithm to set pay leads to efficient matching through one additional channel—via an expansion of the market. When drivers hold information about their reservation wages private to them, some drivers with high (and unknown) reservation wages are left out of the market. These workers can possibly be integrated into the market under algorithmic wage setting. This leads to an increase in the supply of drivers which reduces the fares for passengers. This is observed in the data

Table 7: Two-way Fixed Effects with Wait Time between Trips

	(1) Wait time	(2) Avg Wait Time
Distance	0.06***	0.002
	(0.02)	(0.003)
Duration	-0.02***	0.002
	(0.01)	(0.004)
Surge	-1.19***	-0.02
	(0.35)	(0.06)
Ride Type	0.10	0.10
	(0.26)	(0.28)
Policy	-1.59***	-1.61**
	(0.70)	(0.72)
Rejection Rate * Policy	-0.41	0.65
	(1.11)	(0.95)
Time F.E.	Y	Y
Location F.E.	Y	Y
Time and Location F.E.	Y	Y
Individual F.E.	Y	N
N	912,034	1,408,915
R^2	0.17	0.70

Notes: This table reports two-way fixed-effects estimates of the effect of algorithmic wage setting on drivers' waiting time between trips. Column 1 measures waiting time at the trip level (actual wait time per trip), while Column 2 measures it at the driver level (average wait time across all trips). All specifications include baseline controls, driver characteristics (rating and rejection rate), the treatment variable, and its interaction with rejection rate, as well as time, location, and time-by-location fixed effects. Reported standard errors are clustered at the state level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

by regressing passenger fare on trip controls and the treatment, again following the specification in equation (4). As shown in Table 8, passenger fares decrease by up to \$4.85 per trip due to the policy. Moreover, the reduction in fare is more for passengers who are matched with less selective drivers. In fact, with a 10 percentage point decrease in the driver's rejection rate, passenger fare decreases by an additional 40 cents due to the policy. This implies that passengers matched with more selective drivers are charged more so that it can be translated into relatively higher earnings for these drivers post the policy. This lends even more credibility to my hypothesis that Uber uses its algorithm to offer personalized wages to drivers based on their reservation wages.

Table 8: Two-way Fixed Effects with Passenger Fare

	(1)	(2)	(3)
Distance	1.75***	1.80***	1.80***
	(0.30)	(0.31)	(0.31)
Duration	0.57***	0.58***	0.58***
	(0.06)	(0.05)	(0.05)
Surge	7.15***	7.55***	7.52***
	(0.90)	(1.12)	(1.11)
Ride Type	7.47***	7.44***	7.42***
	(0.73)	(0.72)	(0.71)
Rating		8.21***	8.11***
		(1.32)	(1.15)
Rejection Rate		2.12**	0.11
		(0.86)	(0.62)
Policy	-2.60**	-2.88**	-4.85***
	(1.19)	(1.09)	(1.35)
Rejection Rate * Policy			4.10***
			(0.84)
Time F.E.	Y	Y	Y
Location F.E.	Y	Y	Y
Time and Location F.E.	Y	Y	Y
Individual F.E.	Y	N	N
N	1,877,536	1,418,875	1,418,875
R^2	0.83	0.84	0.84

Notes: This table reports the two-way fixed effects regression results of the impact of algorithmic wage setting on passenger fare per trip. All specifications include baseline controls, and time, location and time-by-location fixed effects. Column 1 includes driver fixed effects. Column 2 includes driver characteristics like rating and rejection rate. Column 3 interacts rejection rate with the treatment. The sample size declines relative to Column 1 due to restricting the sample to drivers with complete characteristics. Reported standard errors are clustered at the state level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

6 Model- Dynamic Learning

Having demonstrated empirical evidence of algorithmic wage discrimination and estimated its causal effect, I now develop a partial equilibrium model of wage setting between Uber and a representative driver following Hart & Tirole (1988). The model serves to formally analyze the mechanisms through which learning over time leads to declining driver compensation.

6.1 Environment

Time is discrete, denoted by t. In the benchmark model, there are two time periods- $t = \{1, 2\}$ and two agents- Uber, U and driver, D. Both agents are risk-neutral and discount the future at the rate

δ. In every period, Uber receives an exogenous ride request, P_t, which denotes the fare consumers are charged for this ride. Upon receiving the request, Uber decides the wage, w_t, to offer the driver for this ride. The driver has an ex-ante preference for driving, represented by a disutility parameter, z. This parameter quantifies the disutility the driver experiences from driving and is drawn from a two-point distribution at the beginning of the first period. This distribution assigns the driver, one of two possible values- z_ℓ or z_h with probabilities γ_ℓ and γ_h, respectively. Driver's preference is not observed by Uber and is private information to the driver. At the beginning of each period, Uber makes a take-it-or-leave-it wage offer to the driver. The driver, upon receiving this offer, decides to either accept it or reject it and wait for the next offer. The game has the following extensive form:

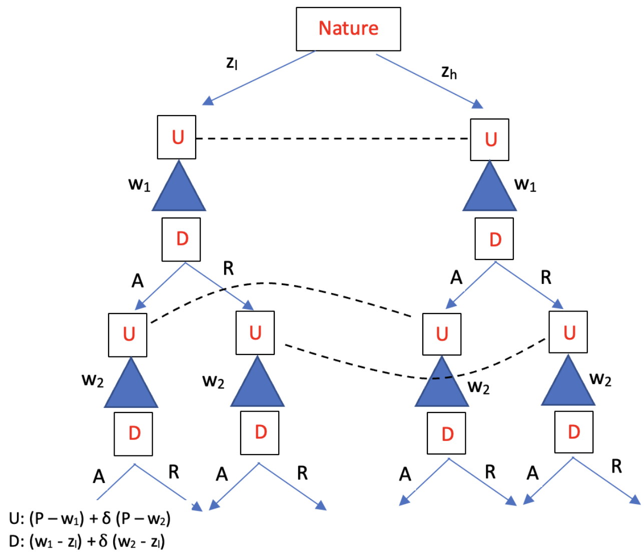

Figure 10: Bayesian Extensive Game

Actions in any period t are denoted by the vector $a_t = (z, w_t, q_t)$ where $q_t = \{0, 1\}$ is the driver's decision to accept or reject the offered wage. First, nature moves and assigns a preference z to the driver, then Uber makes a wage offer w_t to the driver, and finally, the driver decides whether to accept the offer ($q_t = 1$). Uber's payoff in period t is $q_t(P_t - w_t)$, i.e., the difference between the consumer fare and driver pay, conditional on the driver accepting the ride. Driver's payoff is the expected surplus from a ride, $q_t(w_t - z)$, which represents the difference between the wage, net of disutility, conditional on driving. For example, following the first branch of the game tree depicted in Figure 10, Uber's payoff at the end of the game will be $(P - w_1) + \delta(P - w_2)$ since both the first and the second period offers have been accepted by the driver with low-disutility, z_ℓ (henceforth referred to as L-type driver). Consequently, the driver's payoff at the end of the game will be $(w_1 - z_\ell) + \delta(w_2 - z_\ell)$.

Formally, the game tree represents a Bayesian extensive-form game consisting of the following:

1. **Players**: $N = \{Nature, Uber(U), Driver(D)\}$

2. **Histories**: $H = \{(\phi), (z_\ell), (z_\ell, w_1), (z_\ell, w_1, A), (z_\ell, w_1, A, w_2), (z_\ell, w_1, A, w_2, A), ...\}$

3. **Player Function**: $P(\phi) = Nature, P(z_\ell) = U, P(z_\ell, w_1) = D, ...$

4. **Nature/Types**: $P(z_\ell|\phi) = \gamma_\ell, P(z_h|\phi) = \gamma_h$

5. **Information Set**: $I_{Nature} = \{\phi\}$,

   $I_U^1 = \{(z_\ell), (z_h)\}$,

   $I_U^2 = \{(z_\ell, w_1, A), (z_h, w_1, A)\}$,

   $I_U^3 = \{(z_\ell, w_1, R), (z_h, w_1, R)\}$,

   $I_D^1 = \{(z_\ell, w_1)\}, ...$

6. **Payoffs**: $u_U(z_\ell, w_1, A, w_2, A) = (P - w_1) + \delta(P - w_2)$, $u_D(z_\ell, w_1, A, w_2, A) = (w_1 - z_\ell) + \delta(w_2 - z_\ell)$,...

The strategy σ^i for $i \in N$ is a function that assigns an action $A(h)$ to each history where $P(h) = i$ with $A(h) = A(h^{'})$ whenever $h, h^{'} \in I_i$. The strategies of Uber and the driver in each period are as follows:

- $\sigma_1^U = w_1$ (decides w_1)
- $\sigma_1^d = Pr(q_1 = 1|z, w_1) : \{z_l, z_h\}\ X\ R \rightarrow \{0, 1\}$ (decides q_1)
- $\sigma_2^U = w_2|w_1, q_1 : R\ X\ \{0, 1\} \rightarrow R$ (decides w_2)
- $\sigma_2^d = Pr(q_2 = 1|z, w_2) : \{z_l, z_h\}\ X\ R\ X\ \{0, 1\}\ X\ R \rightarrow \{0, 1\}$ (decides q_2)

Strategy profile is a complete set of strategies of each player, $\sigma = \{\sigma^U, \sigma^D, \sigma^{Nature}\}$.

Beliefs, μ_i for $i \in N$ are probability distributions over each possible history within a particular I_i, at each I_i. The beliefs of Uber at the end of each period are as follows:

- $\mu_1(z_\ell) = Pr(z_\ell)$: probability distribution over $z; \mu_1(z_\ell) = \gamma_\ell, \mu_1(z_h) = \gamma_h$
- $\mu_2(z_\ell, w_1, q_1) = Pr(z_\ell|w_1, q_1) : R\ X\ \{0, 1\} \rightarrow R$

The set of beliefs is given by, $\mu = (\mu_i)\forall i \in N$. Strategies and beliefs form a Perfect Bayesian equilibrium (PBE).

6.2 Equilibrium

A (weak) PBE of a Bayesian extensive game is a pair (σ^*, μ^*) that satisfies:

- **Sequential rationality-** At every information set I_i player i's strategy maximizes her payoff, given the actions of all the other players, and given her beliefs, i.e.,

$$\sigma_U^* \in argmax\ u_U(\sigma_d^*, \mu),\ \sigma_d^* \in argmax\ u_d(z, \sigma_U^*)$$

- **Bayesian updating-** At information sets reached with positive probability when σ^* is played, beliefs are formed according to σ^* and Bayes' rule when necessary. For example, the posterior beliefs in t = 2 follow Bayes' rule (on the equilibrium path):

$$\mu_2 = Pr(z_\ell | q_1 = 1) = \frac{Pr(z_\ell)Pr(q_1 = 1|z_\ell)}{Pr(q_1 = 1)} = \frac{\gamma_\ell \sigma^*_{d,1,z_\ell}}{\gamma_\ell \sigma^*_{d,1,z_\ell} + \gamma_h \sigma^*_{d,1,z_h}}$$

- **Off-equilibrium beliefs-** At information sets that are reached with probability zero when σ^* is played, beliefs may be arbitrary but must be formed according to Bayes' rule when possible.

I further assume that all ties for the driver are resolved in the favor of accepting to keep the number of equilibria tractable.

Let $\bar{\mu}$ be defined as $(P - z_h)/(P - z_\ell)$. The set of equilibria may then be characterized as follows:

1. When $\gamma_\ell < \bar{\mu}$, there exists a unique equilibrium in which Uber offers a high wage in both the periods and is unable to screen drivers. This pooling equilibrium is depicted in Figure 11 (the path in red represents the equilibrium play).

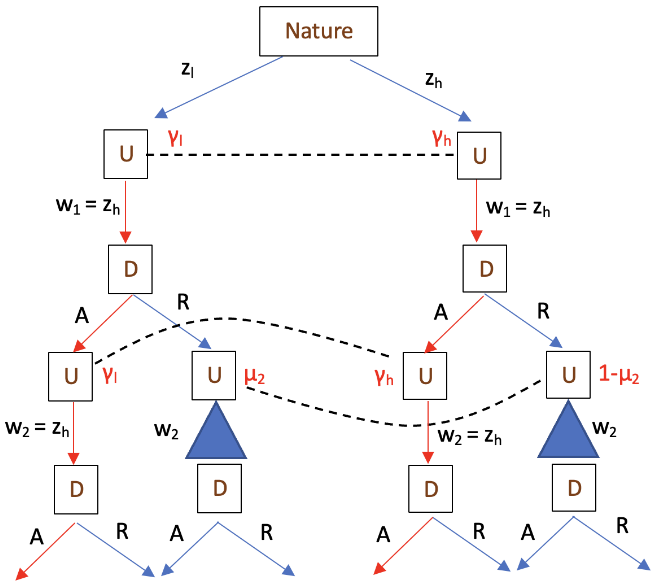

Figure 11: Pooling Equilibrium

2. When $\gamma_\ell > \bar{\mu}$, there exists a unique equilibrium in which Uber offers a screening wage in the first period and then perfectly price discriminates in the second period. This separating equilibrium is depicted in Figure 12.

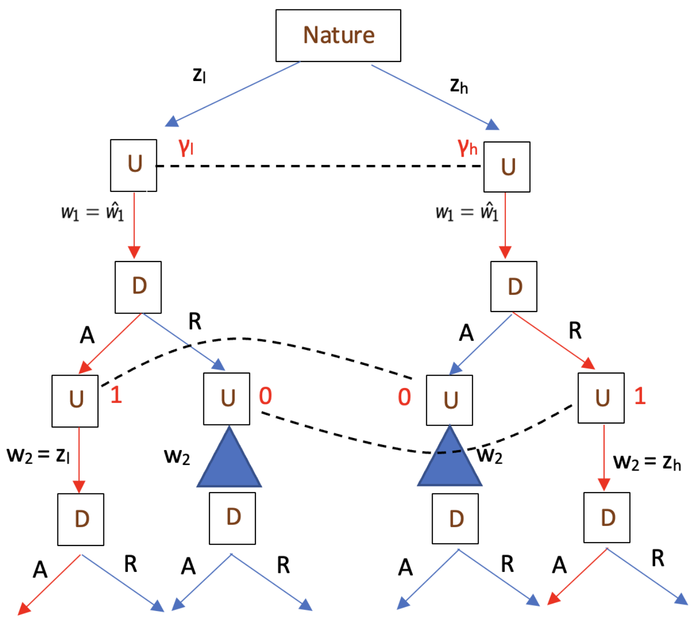

Figure 12: Separating Equilibrium

3. When $\gamma_\ell = \bar{\mu}$, there exist multiple equilibria: the separating equilibrium discussed above, and a pooling equilibrium in which Uber offers a high wage in the first period and randomizes between a high wage and a low wage in the second period.

Intuitively, if the proportion of L-type drivers in the population is high enough, Uber finds it profitable to screen drivers in the first period. In other words, when the extent of driver differentiation is high, Uber benefits from learning drivers' reservation wages early on. Then, in the second period, it can make customized offers for the same ride based on that information.

The model predicts that, in a separating equilibrium, the average wage of workers declines between the two periods. This decline occurs solely because Uber can now extract the information rents that workers previously received due to private knowledge of their type. Moreover, the decline is larger for the low-type as compared to the high-type worker, which also leads to dispersion in pay over time. These results align with the empirical evidence, which shows that Uber's learning over time results in a reduction in drivers' wages and dispersion in earnings.

6.3 Extension- Non-strategic drivers

A key feature of the model is that low-type drivers have incentive to mimic high-type drivers by rejecting current wage offers to secure a higher pay later on. This is because in a separating equilibrium, a rejection is perceived as the driver having a high reservation wage. To isolate the effect of this particular mechanism, I consider an extension of the baseline model where drivers are "naive" or non-strategic. These drivers focus solely on their immediate payoff and do not account for future consequences. As a result, they have no incentive to reject an offer that yields positive utility in the current period merely to signal a different type and secure a better offer in the next period.

Let $\tilde{\mu}$ be defined as $(P - z_h)/(P - z_\ell + (z_h - z_\ell))$. The set of equilibria with non-strategic drivers are characterized as follows:

1. When $\gamma_\ell < \tilde{\mu}$, there exists a unique equilibrium in which Uber offers a high wage in both the periods and is unable to screen drivers.

2. When $\gamma_\ell > \tilde{\mu}$, there exists a unique equilibrium in which Uber offers a screening wage in the first period and then perfectly price discriminates in the second period.

3. When $\gamma_\ell = \tilde{\mu}$, the separating equilibrium and the pooling equilibrium discussed above can co-exist.

4. When $\gamma_\ell = \bar{\mu}$, there exist multiple equilibria: the separating equilibrium discussed above, and a pooling equilibrium in which Uber offers a high wage in the first period and randomizes between a high wage and a low wage in the second period.

I find that average earnings are lower when drivers are myopic as compared to when they are strategic. Strategic drivers may reject first-period offers to mimic high-type drivers, which forces the equilibrium wage in the first period to be set above the low-type's reservation wage, z_ℓ. In contrast, naive drivers accept any offer that meets or exceeds their reservation wage, leading the first-period equilibrium wage to fall to that level. Also, note that $\tilde{\mu} < \bar{\mu}$, which implies that Uber can screen workers more frequently—i.e., over a wider range of parameter values—when they are naive. This is because screening naive workers is less costly: low-type workers only need to be paid their reservation wage to differentiate themselves from high-type workers. This implies that workers can essentially fight back against the algorithm and raise their earnings by being strategic and rejecting some good offers early on.

7 Conclusion

Recent advances in information technology and artificial intelligence have enabled digital platforms to better monitor workers. Advanced technological systems now track and quantify the entire history of workers' engagement with the platform, as well as their personal characteristics. While debates often center on privacy infringements and the need for transparent, consent-based technological governance, a more nuanced issue often goes unexamined: how firms leverage this amassed data to further extract rents from workers. The insights gleaned from these vast datasets, processed by machine learning, inform important decisions like setting individual compensation. In this paper, I provide the first empirical evidence of algorithmic wage discrimination in the labor market using a novel dataset on Uber drivers from 2021 to 2025. The sample includes detailed trip-level data for approximately 2 million Uber trips, matched to drivers with information on their individual characteristics related to productivity and preferences.

The analysis reveals several key patterns. First, even after controlling for trip-level characteristics and traffic conditions, 2% of the variation in pay is attributable to persistent differences across individual drivers, suggesting that drivers are not uniformly compensated for similar work. Individual driver differences may reflect variations in skills (partly captured by ratings) or tacit knowledge from experience, but these do not imply that Uber systematically sets pay based on reservation wages. To isolate the causal effect of Uber's learning on base pay variation, I employ a staggered difference-in-differences design exploiting the shift in Uber's pricing policy for drivers in 2022 from a per-mile rate card to an algorithmic pay structure. Building on the observation that pay varies across drivers for similar trips, I examine whether this transition altered pay variation and whether such changes are linked to drivers' personal preferences (reservation wages) rather than their skills or experience. The results show that algorithmic wage setting increased personalization in pay and reduced drivers' average earnings per trip by \$2.41. Moreover, the decline is steeper for those drivers who reject less. This is accompanied by a reduction in passenger fare by \$2.88 per trip and average wait time for drivers by 1.68 minutes.

To explore the mechanisms behind the labor market implications of algorithmic wage setting, the paper develops a two-period sequential learning model between Uber and its drivers. In this model, Uber learns about drivers' reservation wages from their rejection decision in the first period and accordingly personalizes the wage offers in the second period. The model predicts that average driver earnings decline, more so for those who reject less, as Uber learns about drivers' preferences over time. The model also implies that less selective drivers can potentially improve their earnings by behaving strategically, i.e., by mimicking the behavior of more selective drivers.

Platform work has become a vital source of income for many low-wage workers in the U.S.; however, the growing role of algorithmic monitoring and wage setting raises important concerns about their earnings and income inequality. At the same time, it has economic implications for market efficiency. As digital labor continues to expand, a key policy challenge is to design regulatory frameworks that harness the efficiency gains of algorithmic optimization while ensuring equitable outcomes for workers. While this paper documents the existence and consequences of algorithmic wage discrimination and explores the mechanisms behind it, many questions remain—such as the general equilibrium effects of firm entry on total welfare in an algorithmically managed economy—representing open avenues for future research.

References

[1] Ajunwa, I., Crawford, K., & Schultz, J. (2017). Limitless worker surveillance. Calif. L. Rev., 105, 735.

[2] Binns, R., Stein, J., Datta, S., Van Kleek, M., & Shadbolt, N. (2025, June). Not Even Nice Work If You Can Get It; A Longitudinal Study of Uber's Algorithmic Pay and Pricing. In Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (pp. 1484-1497).

[3] Brown, Z. Y., & MacKay, A. (2023). Competition in pricing algorithms. American Economic Journal: Microeconomics, 15(2), 109-156.

[4] Calacci, D., Rao, V. N., Dalal, S., Di, C., Pua, K. W., Schwartz, A., ... & Monroy-Hernández, A. (2025). FairFare: A Tool for Crowdsourcing Rideshare Data to Empower Labor Organizers. arXiv preprint arXiv:2502.11273.

[5] Caldwell, S., & Oehlsen, E. (2023). Gender, outside options, and labor supply: Experimental evidence from the gig economy. University of California, Berkeley Working Paper.

[6] Castillo, J. C. (2025). Who benefits from surge pricing?. Econometrica, 93(5), 1811-1854.

[7] Chen, M. K., Rossi, P. E., Chevalier, J. A., & Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. Journal of political economy, 127(6), 2735-2794.

[8] Chen, K., Feinerman, K., & Haggag, K. (2025). Flexible pay and labor supply: Evidence from uber's instant pay. Management Science.

[9] Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of uber (No. w22627). National Bureau of Economic Research.

[10] Cook, C., Diamond, R., Hall, J. V., List, J. A., & Oyer, P. (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. The Review of Economic Studies, 88(5), 2210-2238.

[11] Dubal, V. (2023). On algorithmic wage discrimination. Columbia Law Review, 123(7), 1929-1992.

[12] Dube, A., Jacobs, J., Naidu, S., & Suri, S. (2020). Monopsony in online labor markets. American Economic Review: Insights, 2(1), 33-46.

[13] Dubé, J. P., & Misra, S. (2023). Personalized pricing and consumer welfare. Journal of Political Economy, 131(1), 131-189.

[14] Faig, M., & Jerez, B. (2005). A theory of commerce. Journal of Economic Theory, 122(1), 60-99.

[15] Fisher, Jack. (2024). Monopsony Power in the Gig Economy.

[16] Fisher, Jack. (2024). Worker Welfare in the Gig Economy.

[17] Fudenberg, D., & Tirole, J. (1983). Sequential bargaining with incomplete information. The Review of Economic Studies, 50(2), 221-247.

[18] Garin, A., Jackson, E., Koustas, D., & Miller, A. (2024). The Evolution of Platform Gig Work, 2012-2023. SOI Working Paper.

[19] Gardner, John, Neil Thakral, Linh Tô, & Luther Yap. (2024). Two-stage differences in differences. Mimeo.

[20] Grennan, Matthew. "Price discrimination and bargaining: Empirical evidence from medical devices." American Economic Review 103.1 (2013): 145-177.

[21] Guda, H., & Subramanian, U. (2019). Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. Management Science, 65(5), 1995-2014.

[22] Hall, J. V., & Krueger, A. B. (2018). An analysis of the labor market for Uber's driver-partners in the United States. Ilr Review, 71(3), 705-732.

[23] Hall, J. V., Horton, J. J., & Knoepfle, D. T. (2023). Ride-sharing markets re-equilibrate (No. w30883). National Bureau of Economic Research.

[24] Hannak, A., Soeller, G., Lazer, D., Mislove, A., & Wilson, C. (2014, November). Measuring price discrimination and steering on e-commerce web sites. In Proceedings of the 2014 conference on internet measurement conference (pp. 305-318).

[25] Hart, O. D., & Tirole, J. (1988). Contract renegotiation and Coasian dynamics. The Review of Economic Studies, 55(4), 509-540.

[26] Kennan, J. (2001). Repeated bargaining with persistent private information. The Review of Economic Studies, 68(4), 719-755.

[27] Laffont, J. J., & Tirole, J. (1988). The dynamics of incentive contracts. Econometrica: Journal of the Econometric Society, 1153-1175.

[28] Pandey, A., & Caliskan, A. (2021, July). Disparate impact of artificial intelligence bias in ridehailing economy's price discrimination algorithms. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (pp. 822-833).

[29] Rhodes, A., & Zhou, J. (2024). Personalized pricing and competition. American Economic Review, 114(7), 2141-2170.

[30] Shiller, B. R. (2013). First degree price discrimination using big data (p. 32). Brandeis Univ., Department of Economics.

[31] Shiller, B. (2016). Personalized price discrimination using big data (No. 108).

[32] Srivastava, R. (2020). First-Degree Price Discrimination: Evidence from Informal Markets in India.

[33] Teachout, Z. (2023). Algorithmic Personalized Wages. Politics & Society, 51(3), 436-458.

Appendix A. Descriptive Statistics

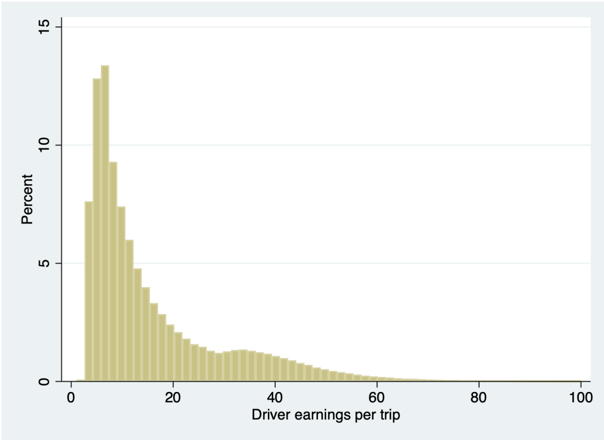

Figure A1: Driver Base Pay

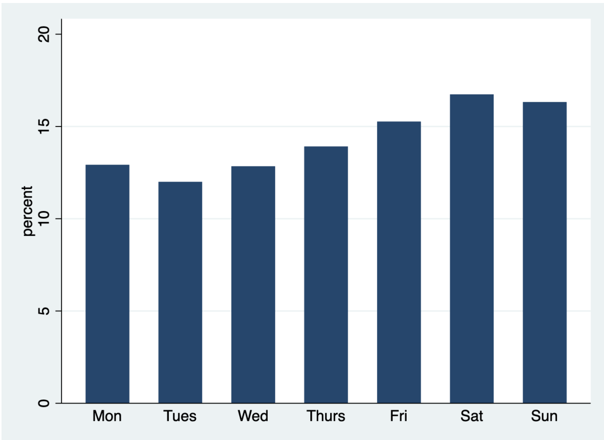

Figure A2: Day of the Trip

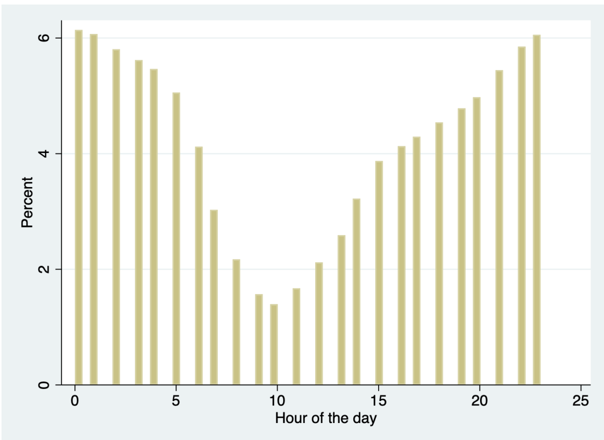

Figure A3: Hour of the Trip

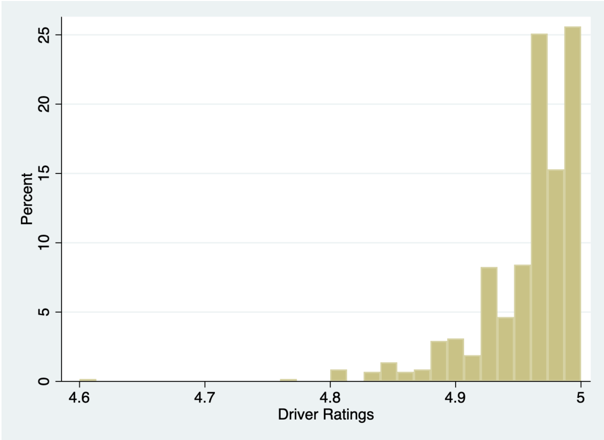

Figure A4: Driver Ratings

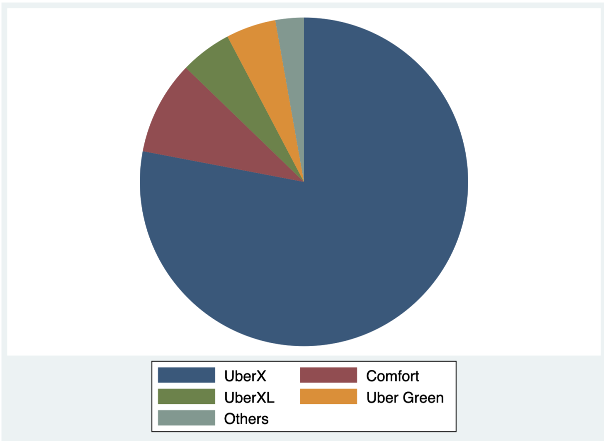

Figure A5: Type of Trip Requests

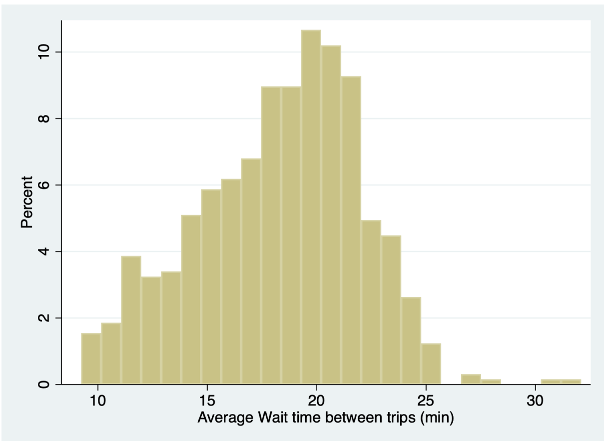

Figure A6: Average wait time between trips

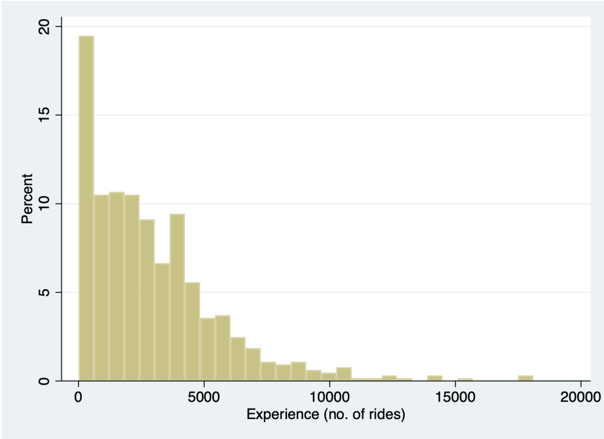

Figure A7: Drivers' Experience (total trips completed in the sample)

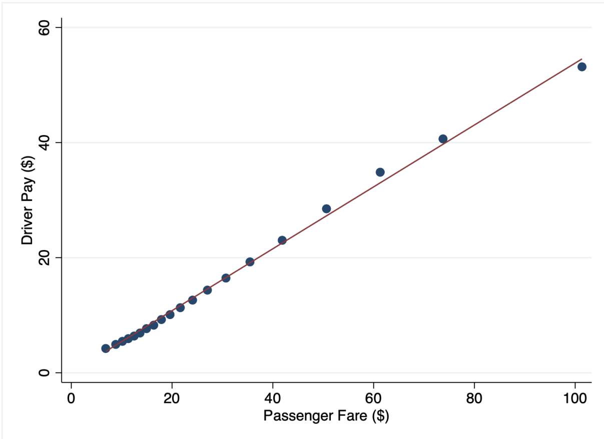

Figure A8: Passenger Fare vs Driver Pay

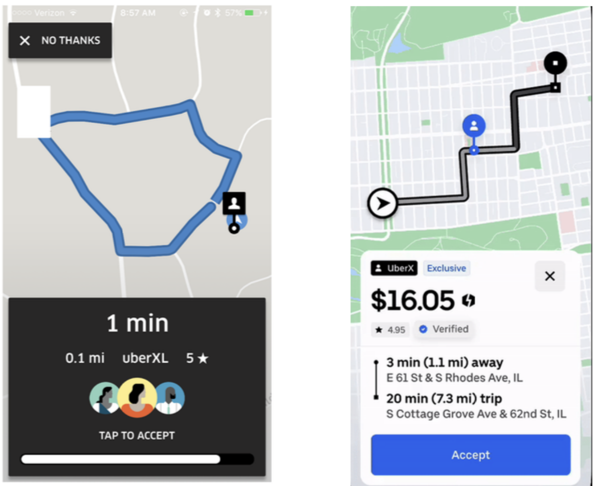

Figure A9: A ride offer before and after the policy change

Table A1: States by Treatment Timing

State	March 2022	September 2022	Never Treated
Alabama		✓	
Arizona		✓	
Arkansas			✓
California		✓	
Colorado		✓	
Connecticut		✓	
Delaware			✓
Florida	✓		
Georgia	✓		
Hawaii			✓
Idaho		✓	
Illinois		✓	
Indiana		✓	
Iowa			✓
Kansas			✓
Kentucky		✓	
Louisiana	✓		
Maine			✓
Maryland		✓	
Massachusetts		✓	
Michigan		✓	
Minnesota		✓	
Mississippi			✓
Missouri	✓		
Montana			✓
Nebraska		✓	
Nevada		✓	
New Hampshire			✓
New Jersey			✓
New Mexico		✓	
New York		✓	
North Carolina		✓	

Continued on next page

Table A1 – *Continued*

State	**March 2022**	**September 2022**	**Never Treated**
Ohio	✓		
Oklahoma		✓	
Oregon			✓
Pennsylvania		✓	
Rhode Island		✓	
South Carolina		✓	
Tennessee		✓	
Texas	✓		
Utah		✓	
Virginia			✓
Washington			✓
West Virginia			✓
Wisconsin		✓	
Wyoming			✓

Appendix B. Additional Regressions

Here, I explore some additional regression specifications. After estimating the individual driver fixed effects from equation (1), I examine which observable driver characteristics explain this variation. For this, I regress the estimated driver fixed effects on the driver characteristics available in the data, both related to productivity and preferences. The specification is as follows:

$$\delta_i = \mu + \omega W_i + \epsilon_i \tag{B1}$$

where δ_i is the estimated fixed effect for driver i and W_i is a vector of characteristics of driver i. Driver characteristics are the same as in specification (2). Table A2 shows the results.

Table A2: Regression with Driver characteristics

	(1) Driver F.E.
Rating	9.05***
	(0.13)
Experience (rides)	-0.00005***
	(0.000002)
Experience (months)	0.0002*
	(0.0001)
Avg Wait time	-0.01***
	(0.002)
Multiapp	-0.69***
	(0.01)
Rejection rate	2.50***
	(0.02)
N	1,425,757
R^2	0.03

Notes: This table reports regression results using trips by 551 Uber drivers from Jan 2021 to June 2025 with complete driver characteristics. The dependent variable is the driver fixed effect estimated in Table 2, Column 2. Controls include driver characteristics—rating, experience (rides completed and months driving), average wait time between trips, rejection rate, and whether the driver multi-apps on Lyft. Reported standard errors are clustered at the individual driver level. The coefficients are statistically significant at the ***1%, **5%, or *10% level.

The model explains only 3% of the variation in estimated driver fixed effects, suggesting that many relevant individual driver characteristics are missing from the data. In reality, Uber likely has access to a richer set of driver information—including their age, sex, race, home location (which can reveal routes that help drivers return home), cancellation rates, and, importantly, offer history—that is not available to the econometrician analyzing this sample[12]. Caveat to the omitted variable bias, the results indicate a positive correlation between ratings and drivers' individual contribution towards pay. Higher rejection rates are associated with higher fixed effects. Drivers who multi-app on other rideshare platforms are more likely to accept offers face a reduction in pay due to their individual characteristics. These results are qualitatively similar to those in Table 3.

[12]See Uber's privacy notice for drivers for a complete list of driver data that Uber collects.

Appendix C. NYC TLC Data

In this section, I use data from the New York City Taxi & Limousine Commission (NYC TLC) for April 2022 to cross-validate some of the claims made using the WAO data with an alternative dataset. NYC TLC, created in 1971, is the agency responsible for licensing and regulating New York City's yellow taxis, green taxis, for-hire vehicles (FHVs), commuter vans, and paratransit vehicles. NYC TLC requires for-hire vehicle companies like Uber and Lyft to share trip-level data of the universe of their drivers every month. This granular data consists of trip variables like distance, duration, date, time and location of pick-up and drop-off. It also contains detailed compensation variables like base passenger fare, tolls, tips, taxes, congestion surcharge, airport fee and driver pay. The data consists of roughly 15 million trips completed each month by all Uber drivers in New York City.

The monthly average base pay of Uber drivers is \$19.56 per trip, the average trip distance is 5.1 miles and the average duration of the trip is 19.4 minutes. The data indicates that drivers' base pay is strongly correlated with the total miles and duration of the trip. Moreover, there is a high pass-through from passenger fare to driver base pay. Next, I observe the wait time of the passenger in the data, i.e., the time between when the trip was requested by the passenger and when the trip started. The wait time per trip ranges between 0 and 60 minutes with an average of 6 minutes (Figure B1).

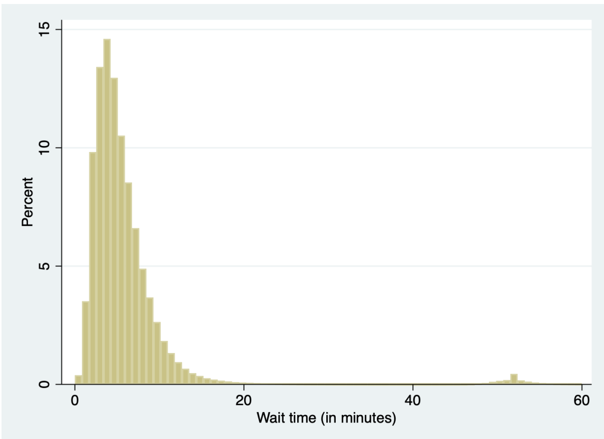

Figure B1: Distribution of waiting time

I take a longer wait time of the trip as a proxy for a lower reservation wage of the driver. This is assuming that a trip request that takes a long time to be accepted must be an unfavorable one for the drivers. Hence, those drivers who end up accepting them must be the ones with lower reservation wage. I observe that these trips are usually the ones with long distances and long duration, such as airport commutes in many cities (Figure B2).

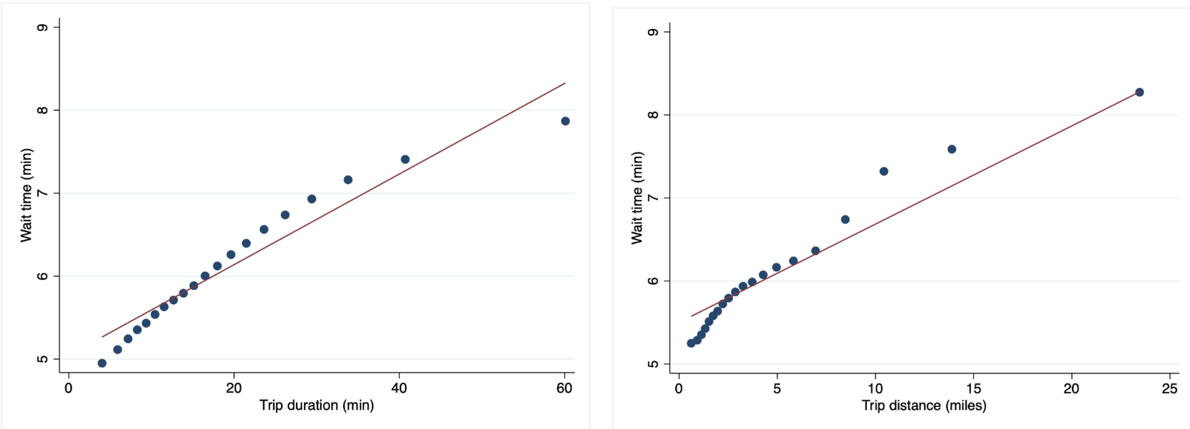

Figure B2: Determinants of waiting time

I also observe that conditional on trip duration and distance, wait times are negatively correlated with driver base pay. So a long wait time (interpreted as lower reservation wages) is associated with a lower driver pay (Figure B3).

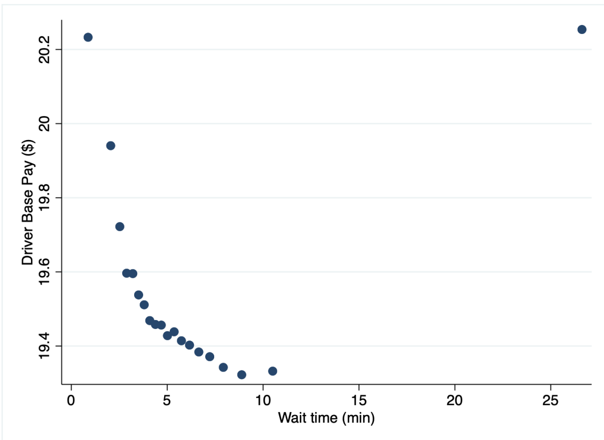

Figure B3: Driver Base Pay vs Wait time of the trip

Wait time of the trip is a noisy measure of driver's reservation wage. A possible limitation of this measure is that a shorter wait time need not necessarily translate into a higher reservation wage. For example, let us say we observe a \$10 trip in the data that gets accepted fairly quickly. This could be accepted by a driver who is willing to accept any ride and hence, has a lower reservation wage (say \$2). Alternatively, it could be accepted by a driver who has been rejecting many rides (of less than \$10) and hence, has a higher reservation wage (say \$9). Since the data do not identify drivers across trips, accurately estimating their reservation wages remains challenging. Nonetheless, longer wait times can still be taken as a good measure of a lower reservation wage as over a long period of time, it is reasonable to assume that the ride was rejected by many drivers. The results with this alternative dataset are consistent with the main findings, providing support for their external validity.