Distributional and Productivity Implications of Regulating Casual Labor: Evidence from Ridesharing in Indonesia

Shotaro Nakamura*† Rizki Siregar‡

December 20, 2023

Download the latest version here

Abstract

Regulations intended to improve workers' earnings, such as minimum wage, may have a muted efficacy due to adjustment mechanisms and general equilibrium effects. Using comprehensive transaction data from one of Indonesia's dominant platforms, we study the market-wide implications of a minimum-fare policy on ridesharing apps. We estimate the causal effects of the policy by exploiting an exogenous variation in the policy's rollout. We find that, on average, the policy increases the trip price but does not significantly affect the overall transaction volume nor increase driver earnings or wages. These effects are driven by a higher excess labor supply, reducing the number of transactions per driver. The excess labor supply comes from lower-earning drivers but does not lead to their increased earnings. The policy also lowers driver productivity by increasing the share of less productive drivers in the workforce and reducing individual productivity due to crowding on the supply side.

JEL: J38, O18, R48
Keywords–Ridesharing, Minimum Wage, Casual Labor, Distributional Impact, Labor Productivity

*Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Federal Trade Commission, or its Commissioners. They also do not reflect the views of the firm from which we received data. The paper has gone through a check by the firms' employees to ensure confidentiality of their data and other proprietary information, but not on the empirical findings and views expressed in the paper. The authors report no conflict of interest.

†Federal Trade Commission shotaro.n.nakamura@gmail.com
‡Fakultas Ekonomi dan Bisnis, Universitas Indonesia rizki.siregar@gmail.com

1 Introduction

The informal sector, which we define as transactions that are not registered or regulated by an existing legal framework, is the dominant source of source in developing economies, accounting for approximately 70 to 80 percent of all employment in lower and lower-middle-income countries (World Bank 2018).[1] Informal employment is characterized by its lower pay, higher uncertainty, and lower productivity in developing economies (Ulyssea 2018; Kochar 1995; Kochar 1999; La Porta and Shleifer 2014). If the characteristics associated with informal and casual labor markets result from market frictions or failures, policy interventions that induce transfers to workers in the informal sector may improve their welfare and productivity. The price floor, such as minimum wage, is a standard policy instrument that intends to generate transfers to suppliers of labor.

It is unclear, however, if price floors would deliver intended outcomes on workers' earnings in informal and casual labor markets. In a neoclassical setting with no market failures or frictions, a binding price floor would reduce quantities demanded and induce excess supply. However, empirical findings deviate from this intuition; in the context of formal, salaried employment, the literature on minimum wage has found mixed evidence on employment that deviate from a neoclassical intuition (e.g., Card and Krueger 1994; Cengiz et al. 2019; Jardim et al. 2018). Price-floor policies may also have distributional and productivity implications, as an emerging body of work on minimum wage has shown positive impact (e.g., Engbom and Moser, forthcoming). Yet, the effects of regulations in informal and casual labor markets likely differ from those in formal and salaried settings based on differences in regulatory structures, the relevance of market frictions and failures, and the extent of

[1] For this statistic, informal employment is defined as "not [having] a contract, social security, and health insurance and is not a member of a labor union" World Bank 2018. Although various definitions of informal work exist, we follow the one set by ILO and emphasize that the vast majority of work in developing economies remains outside of the regulatory framework, including wage floors (International Conference of Labour Statisticians 1993).

employers' monopsony power.

In this paper, we focus on the market-wide impact of labor regulation directly imposed on a casual and informal labor market: two-wheel taxi rides on mobile ridesharing apps in Indonesia. In Indonesia and other developing economies, app-based platforms have largely absorbed the stock of ride-hailing and delivery drivers who had previously worked offline.[2] Digital platforms provide visibility into a type of informal labor market on which there had been limited, high-frequency data. The platforms also provide an opportunity for policymakers to intervene in a once-unregulatable labor market segment, as prices and other attributes of transactions can be controlled algorithmically.

We study the effects of a price-floor policy *per transaction* in the ridesharing market with a similar policy objective as in minimum wage: to raise workers' daily earnings. In 2019, the Indonesian government introduced a price floor on the amount that drivers received on a trip ("*driver fare*" henceforth) on all ridesharing platforms. By exploiting the city-level variation in the timing of policy implementation, we estimate its causal effects with difference-in-differences and synthetic control methods. The data-rich environment, in which we have access to the universe of transactions and worker-level productivity data from one of Indonesia's two major online platforms for ridesharing, allows us to identify the effect of policy on market-wide adjustments, earning and labor supply distributions, and productivity in detail.

Overall, we find that the transaction-level price-floor policy intended to increase workers' earnings in an informal and casual labor market may not achieve its objectives; the policy induces a large supply response that reduces the driver-to-customer match rate and productivity, canceling the effect of a higher piece rate on earnings and wages. We break down our empirical results into the following.

[2] It is difficult to get a reliable estimate of the share of ride-hailing transactions that are mediated by mobile apps. Based on publicly available industry reports, the share of app-based services in the taxi industry is around 40 to 60% in Indonesia and the ASEAN region (Statista 2022; Mordor Intelligence 2022).

First, we find that the policy increases the average driver fare through a binding price floor, but does not significantly affect the transaction volume or driver earnings. We find statistically significant 4.6% increase in average driver fare and statistically insignificant 0.2% increase in transaction volume. However, we also find that the policy only leads to a statistically insignificant 1.7% reduction in daily earnings from driver fare and a statistically insignificant 6.7% reduction in hourly wage, i.e., daily earnings divided by supply hour. The lack of statistically significant effects on driver earnings despite increased average driver fare suggests market adjustments.

Second, we find that increased excess supply explains the limited impact on average drivers' daily earnings despite the higher average driver fare. The labor supply measures collected by the platform allows us to construct a measure of excess labor supply as drivers' idle time and separate it from transactions hours. We find a statistically significant 24.3% increase in excess supply hours, i.e., the sum of all idle hours on the app from all drivers, yet no impact on transaction hours. As a result, we find statistically insignificant effects on the effective wage rate, i.e., earnings per hour available on the app. In the ride-hailing market, where drivers can easily adjust their own supply decisions—a setup that may generalize to informal and casual labor markets more broadly—the effects of a higher piece rate on driver wages and earnings are crowded out by excess supply.

Third, we investigate on what margins labor supply is adjusted, and the extent to which a subset of drivers may have benefited from the policy. We find that the effect on excess supply is driven by noisy adjustments on the extensive—the number of distinct individuals per day—as well as intensive—how much drivers work conditional on being present on a given day—margins of driver supply. We find that the policy increases the total labor supply of workers in the bottom three to four deciles of pre-policy earnings and labor supply by 15 to 40%. We also find, however, that the *average* earnings do not increase significantly for the low-income driver segments whose aggregate supply increases. We also do not find evidence that drivers can target higher-value trips by differentially accepting trips at the minimum price, possibly because of high matching frictions and the allocation algorithms that discourage such behaviors. These results suggest that a) marginal gains from higher driver fares are competed away by malleable supply, and b) the highly frictional nature of the market does not allow the drivers to target their labor supply to higher-value trips, leaving them no better off.

Fourth, we find that the minimum-fare policy and associated adjustments come at the cost of reduced driver productivity. The minimum fare policy results in 8 to 11% reduction in average driver productivity, as defined by a ratio of the quantity (distance, number of trips, or duration) of trips divided by supply hour. We break down the average effects to two channels; first, the policy increases supply hours from low-productivity workers, driving down the average productivity of the fleet via a compositional shift. Second, the policy reduces individual driver productivity across the pre-policy productivity distribution, even those with higher pre-policy productivity whose labor supply was not affected. These results suggest that increased driver availability reduces the productivity of inframarginal and more productive drivers by crowding out the supply side with less productive competitors. In summary, our results suggest that a regulation guaranteeing a minimum payment for a job could affect labor productivity by changing both who participates in the labor market and how they perform on the job.

Fifth, on the consumer side, we find heterogeneous incidence of the policy's cost, yet homogenous effects on the incidental benefit, i.e., shorter wait time from the availability of more drivers. When we examine the differential effects by customers' exposure to the policy, i.e., consumption of trips in the pre-policy period that would be regulated, we find that customers in the top deciles pay 20% more per trip and on daily expenditures. However, we also find a homogeneous, 5% reduction in a proxy of customers' wait time regardless of their policy exposure. The results explain why we saw limited demand-side responses, as customers whose transactions are targeted by the policy have inelastic demand. The results also imply that the incidence of the regulation falls most heavily on the same subset of consumers. Finally, uniform reduction in wait time suggests that the incidental benefit of the policy is suboptimally allocated relative to its cost.

To rationalize these findings, we introduce a static matching model of the ride-hailing market in which prices shift exogenously, workers respond by choosing labor supply, the total quantity of which affect average driver productivity and wage. We use this framework to select key outcomes and compare effect sizes with a similar study by Hall et al. (2021) that evaluates a platform policy by Uber in the United States. Both papers find increases in labor supply, but the magnitude in our analysis is more than five times larger than in Hall et al. (2021). The difference in the magnitude of the labor-supply effect explains the differences in effect sizes on downstream outcomes, i.e., driver productivity and wages. Due to several key contextual differences such as policy design and location, we are unable to infer what drives the large differences in the labor supply response. Yet, our findings are consistent with the idea that high levels of informality and large labor stocks in Indonesian cities trigger larger supply responses than in the United States.

This article makes several contributions to the literature on labor regulations, casual work in developing economies, and applications of standard labor-policy instruments in the gig economy. First, we expand on the insights from the minimum wage literature by addressing the same fundamental question—the disemployment effects of price floors—in a novel and increasingly relevant labor market: the gig economy. At the first glance, our results on transaction volume are in line with those suggesting limited effects of minimum wage on employment (e.g., Cengiz et al. 2019). Yet, upon further inspections, we find large labor-supply responses and crowding that suggest the importance of adjustment mechanisms that dampen the direct effects of higher fares. Our findings suggest that price-floor policies may result in very different outcomes in standard employments in the formal sector v.s. piece-rate or casual labor markets. In the former, monopsony power may limit the extent of labor supply responses, yet in the latter the supply is much more flexible by design.

Second, our platform-wide access to transaction-level data allows us to analyze the productivity effects of labor regulation and its mechanisms. The productivity measures are distinct from wages and available at the driver-day level for all drivers on one of Indonesia's two largest ridesharing platforms. Previous studies find positive effects of minimum wages on productivity and attribute them to the "efficiency wage" hypothesis, i.e., increases in wages are absorbed by increased worker efficiency (Ku 2022; Coviello et al. 2021; Dustmann et al. 2022). We provide an alternative perspective to this question by showing that productivity is *reduced* for two reasons. First, our data capture an entire labor-market segment of a dominant online platform rather than individual employers or chains, allowing us to detect market-wide implications and effects beyond direct treatment effects. Second, with productivity measures at the worker level rather than at the firm level, we observe not only what happens to individual labor productivity but also how workers of a given productivity level are reallocated. In other words, we capture a more comprehensive effect on productivity, including workers' intensive and extensive margins as well as spillovers.

Third, our work provides novel insights into informal labor in developing cities and the scope of regulation to increase workers' earnings. Our study offers unique insights into the implication of a standard labor policy instrument applied in casual and informal labor markets, where many workers with low socioeconomic status in developing cities find employment. Labor supply decisions of informal, casual-wage workers may be subject to liquidity constraints and reference-dependent preferences (e.g., Dupas et al. 2020). In such a context, one may expect that interventions that improve the terms of work may help increase informal laborers' earnings and productivity. The existing body of work provides insights on the efficacy of either minimum-wage policies in the formal sector or a quantity-side intervention in a form of employment guarantees for lower-wage earners.[3] Yet, to our knowledge, there is no existing work on the efficacy of a price floor policy applied directly in the informal labor market. Our work contributes to a gap in the literature with the following insight; a minimum wage-like policy on casual work may have limited efficacy in increasing the earnings of low-wage earners, and it comes at the cost of increased expenditure for price-inelastic consumers and reduced worker efficiency.

Fourth, we contribute to the understanding of labor-regulation instruments in gig-economy platforms. Our paper is, to our knowledge, one of the first to empirically evaluate a government-initiated price-control policy in an app-based ridesharing market. Our work is closely related to Hall et al. (2021), who assess driver responses and marketplace equilibrium after price shocks, and Horton (2017), who studies the productivity effect of a minimum-wage policy in an online labor market. Our empirical context is similar yet distinct as we focus on the impact of a price-floor policy rather than an average shift in fare. The regulatory structure makes it more likely to induce a distributional effect compared to a uniform price increase. Our context also differs in that it was a government policy that applied to all platforms rather than a platform-led pricing policy in markets served by a single provider (i.e., Uber). Therefore, our environment may be more reflective of the current competitive landscape for gig-economy platforms, and our findings may be more relevant from the regulatory perspective.

We proceed with a description of the empirical context of the fare regulation and the environment of the ridesharing market in Indonesia in Section 2. We then provide details on the data we use in Section 3. Section 4 presents our identification strategies. In Section

[3] In terms of minimum wage policies in developing economies, findings are less conclusive than in developed ones, and the studies necessarily focus on the formal sector (Neumark and Corella 2021). For a similar policy that regulates how much individuals are entitled to work rather than how much they get paid per work, there is evidence that guaranteed employment through public works schemes increases earnings and leads to a positive general equilibrium effect in the local economy (e.g., Imbert and Papp 2015; Beegle et al. 2017; Muralidharan et al. 2017).

5, we describe our findings. We discuss our results in Section 6 with a combination of a conceptual framework and comparisons of estimates with Horton (2017) for external validity. We conclude in Section 7.

2 Empirical context

2.1 Ridesharing and its regulations in Indonesia

Ridesharing platforms have served Indonesian consumers since the mid-2010s. The platforms provide services by automobile cars and motorcycles called *ojeks*, which are cheaper and more popular modes of ride-hailing in Indonesia. Currently, 60 to 70% of all ridesharing trips are conducted by *ojeks*. Ridesharing platforms also offer various other services by motorcycle drivers, such as food and package delivery. The motorcycle-based services, in particular, disrupted offline services that were the norm before the introduction of such platforms.

In 2019, the government of Indonesia enacted regulations on transactions on ridesharing platforms (Indonesian Ministry of Transportation 2019). It was motivated by growing concerns about the welfare of drivers on ridesharing platforms. Pressures from groups and associations of drivers demanded regulations that protect drivers in terms of, among others, safety and wage security. The government of Indonesia then enacted several laws starting in early 2019 for motorcycle taxi services. In this paper, we focus on one of the regulations, in particular, fare regulations of motorcycle taxi services. Hence, from here forward, we refer to services provided by motorcycle drivers in all of our analyses.

We exploit the variation by city in the timing and threshold of the minimum- and maximum-fare policy. From Indonesian-language news sources, we have identified the following timetable of policy announcements and implementation by the Ministry of Transportation, as shown in Table 1. From our reading of the news coverage, it seems that the May-1st rollout was only temporary, and the government suspended the regulation on May 15th. We have also identified that the policy went into effect for most cities where our ridesharing partner operates by the third phase of the regulation's implementation. Therefore, the main source of policy variation is the second phase (i.e., the July-1st implementation) versus the third (August-9th).

2.2 Structure of the fare regulation

The fare regulation consists of two components. First, drivers are guaranteed a payment per trip, regardless of the distance or duration of a trip. This component of the policy effectively works as a price floor of a transaction, and we refer to it as the minimum total fare. Second, the policy sets a limit on the range of per-kilometer price that drivers have to be paid, which we refer to as the per-kilometer rate. In reality, the pricing algorithm first estimates the fare based on distance, location, time, and other factors conditional on the constraints on the per-kilometer rate. The estimated fare is then compared to the minimum total fare, and the final driver fare is the maximum of those two values.

The per-km minimum and maximum fare, as well as minimum total fare, varies by groupings of cities, to which the Ministry of Transportation refers as "zones." Zone 1 covers major cities in the populous islands of Bali, Sumatra, and Java (except Jakarta Metropolitan area), Zone 2 consists of Jakarta Metropolitan area (often referred to as Jabodetabek), and Zone 3 includes the rest of the country. Table 2 shows the minimum and maximum per-km fares and minimum total fares imposed by the Ministry. Appendix Section A.1 lists the cities included in each rollout phase of the regulation's implementation. In addition, the fare regulation only targets taxi services. Other services catered by motorcycle drivers, such as delivery, are not subject to this policy.

2.3 Contextual specificity and external validity concerns

Our empirical context is uniquely suited to study the effects of price regulation in equilibrium and the mechanisms that lead to it. First, thickness on both sides of the ridesharing market means that individual "employers" (i.e., customers) likely do not have significant market power. Although the platform decides on transaction prices via its algorithm, it is doubtful that either of the two large platforms has the significant pricing power to deviate from their competitors' prices. The market share of the smaller of the two platforms is 40 to 50% as of 2021 (Measurable AI 2022).[4] The market condition helps us abstract away from monopsony power and instead focus on supply-side mechanisms and distributional effects. Second, the short average contract duration allows us to observe markets in a new equilibrium quickly after the introduction of the minimum-fare policy, unlike in standard wage-employment contracts. Third, unlike uniform price increases led by the platform operators like in Horton (2017), the policy in Indonesia was government-led and generated a price floor, which mimics a traditional minimum wage more closely and may be more likely to induce a distributional impact.

Yet, the data and policy contexts may limit the extent to which our findings may be generalized beyond the ridesharing context. First, the government policy regulated driver fares per transaction rather than driver wages. These differences in the regulated units may limit the relevance of our findings to the minimum wage literature. Second, the policy consisted of both minimum total fare and constraints on per-kilometer rates, making the direct comparison with the minimum-wage design difficult. Third, we cannot observe the long-term impact of labor regulation because of the limited duration of plausibly exogenous policy variation.

[4] Anecdotally, the two largest platforms have been engaged in fierce competition over price à la Bertrand in the last decade.

3 Data

We construct our data sets for analysis using the collaborating platform's database, consisting of the universe of transactions and driver-daily measures. We restrict our sample to all completed motorcycle trips that had non-zero payments to associated drivers. Our data set contains trips from 64 Indonesian cities, 55 of which are in the data for our analysis. The data covers the period of January 1st to August 8th, 2019.

The transaction-level data set contains information about the price and service type of the trip. Each trip is associated with one of the service types, the most popular of which are taxis, food delivery, and non-food delivery. We use data from all service types for the analysis unless specified as being restricted to "regulated" service types, in which case we restrict the data to taxis. For each trip, we can identify payments made to the driver and the price charged to the customer. We define these measures as the driver- and customer-fare, respectively, and use them for analysis. We note that the amount customers are charged is not necessarily what they end up paying, as they may receive discounts through credits or bundled offers. Unfortunately, we are not able to identify the actual amount paid by the customers as we are not able to distinguish discounts and credit usage from non-cash payments in the database to which we have access.

The transaction data also contain information about the trips' characteristics and associated driver and customer IDs. We use booking and completion time stamps to identify the date and duration of the trip. The data also contain route distance estimates used by the pricing algorithm, which we use to compute per-km fare rates. We also use driver and customer IDs to construct driver- and customer-daily level measures.

Other than the transaction-level data, we use the driver-daily level database to measure driver availability. This data source captures a daily measure of driver availability by tracking how long a driver has the platform app open. We take this measure as a proxy of daily supply hours, defined as the amount of time a driver was available for hire on a given day. We also use this measure as the denominator of our productivity measures.

3.1 Constructing aggregated analysis data sets

We aggregate the transaction- and driver daily-level data for analysis because we are interested in identifying market-wide effects and distributional impact. We construct city-day level panel data to assess the city-wide average effects on average prices, frequencies, and other outcomes. We describe the empirical strategy that uses this city-day level panel data in Section 4.1.1.

We also aggregate the data at the city-day-"bin" level to identify effects on the distribution of outcomes and to estimate conditional average treatment effects by pre-policy characteristics of drivers and customers. For effects on the distribution, whose results are presented in Section 5.1.1, we aggregate the transaction data at the level of city, day, and driver-fare bin of 1,000 Rupiah. The outcome variable for each row of this panel data is the number of transactions for which the driver fare was in a given fare bin for a given city on a given day. For conditional average treatment effects by pre-policy characteristics, such as results presented in Section 5.3, we construct panel data of averages and counts at the city-day-bin level, where the bin is defined as the decile of the driver- or customer-characteristics in the month prior to the start of our policy variation.

4 Identification strategies

We consider two methods for our identification strategy: difference-in-differences and synthetic controls. One key aspect is the extent to which the parallel trends hold. As discussed in Section 2, the government rolled out the new minimum and maximum wage policies non-randomly, and larger cities tended to fall under the new regulation earlier. Violating the parallel-trends condition may mean that the treatment effect estimates are biased.

Figure 1 shows trends of average fees by the rollout phases, with policy rollout timing marked by vertical dash lines of the corresponding color. It shows that cities in Phase 1 have somewhat different trends from those in Phases 2 and 3, which may be because cities in Phase 1 included many of the largest cities in Indonesia. On the other hand, Phases 2 and 3 have remarkably similar time trends on the average fare.

Figure 1 also helps visualize the effect of fare regulation on average prices, which increase sharply right at policy introduction. For confidentiality reasons, we standardize to the average driver fare in cities included in Phase 2 from January 1st all the way up to the day before the policy implementation, and we call this value X. For cities in Phase 2, the average fare from the implementation date to the end of the year increased to $1.21X$. For cities in Phase 3, the pre-policy average (i.e., up to August 8th) was $0.82X$, but increased to $1.12X$ post-policy implementation. As such, we will deploy difference-in-differences and synthetic control methods to paint a holistic picture of the policy effect. We present results from our checks on parallel trends and other robustness checks in Appendix Section 5.2.2.

4.1 Difference-in-differences

Our preferred empirical approach is the difference-in-differences (DiD) method, where we take advantage of variation in the timing of policy rollout. Due to concerns about the violation of parallel trends, we will restrict our sample to cities in rollout phases 2 and 3, where we are confident of exact policy implementation timing and parallel trends hold, as shown in Figure 1. We restrict the data to dates between January 1, 2019, to August 8, 2019, the last day of policy variation between phases 2 and 3. We conduct difference-in-differences analysis on city-day averages, as well as on fare (and other) bins, as follows.

4.1.1 DiD over average outcomes

We estimate the following equations on the city-day level outcomes. Subscript c denotes city, while t denotes day as the unit of time.

$$Y_{c,t} = \beta_0 + \beta_1 * I_{c,t}(c \in Treat, t > 0) + \gamma_c + \delta_w + \rho_d + \epsilon_{c,t} \tag{1}$$

We study the impact of the fare regulation on various city-day outcomes, $Y_{c,t}$. The granularity of the data allows us to construct four groups of our primary city-day outcomes. First, we analyze the impact on transactions, such as average fare, average distance, and average per-km rate. The second group of outcomes is quantities, such as total supply hours, total driver earnings, total distance, number of drivers, and number of bookings. Third, we also estimate the impact of the regulation on driver-daily outcomes, such as drivers' daily supply hours, daily earnings, and wages. We differentiate earnings and wage, with wage defined as daily earnings divided by supply hour. Lastly, we study the impact of regulation on various measures of productivity, such as the ratios of the number of rides, distance worked, and time spent on providing rides over the total hours spent on the platform.

We define the first day of treatment, July 1st, 2019, as t equals 1. Then, the indicator for treatment, $I_{c,t}$, is one if city c is included in phase 2 of the fare regulation's rollout for all days starting on the first day of treatment. This indicator variable is zero otherwise. We include city fixed effects, γ_c, as well as two sets of time fixed effects: calendar-week fixed effects, δ_w, and day-of-the-week fixed effects, ρ_d. Lastly, $\epsilon_{c,t}$ is an idiosyncratic error term.

Our coefficient of interest is β_1 which represents the impact of the fare regulation on city-day outcomes. This coefficient captures the within-city changes in the outcome variables. It provides the average treatment effect of the fare regulation. We also cluster the standard errors at the city level.

4.1.2 DiD over frequency by fare bins

We use the city-level difference-in-differences (DiD) over various bins to uncover the heterogeneity in localized effects of fare regulation. Again, we exploit the timing of the policy implementation in cities included in phases 2 and 3.

We estimate the effect of minimum-fare policies on transactions and distributions of wages. The basic specification is the following:

$$E_{j,c,t} = \sum_{j=0}^{31} \beta_j I_{j,c,t}(c \in Treat, t > 0, j) + \mu_{j,c} + \rho_{j,t} + \epsilon_{j,c,t} \tag{2}$$

$E_{j,c,t}$ is the number of transactions in j*1000-Rupiah per-km fee bin in city c on day t. $I_{c,j,t}$ is a treatment indicator term, equaling 1 if the minimum fare policy is implemented in city c on day t and the fare falls between $j * 1000$ and $(j+1) * 1000$ Rupiah. These fee bins are at the 1000-Rupiah increment, except for fees bigger than 30,000 Rupiah, denoted as $j = 31$. Table 2 shows that minimum fare varied by zones (and was, in fact, given as a range rather than a particular value), and cities in different zones are spread over phases 2 and 3 of the policy rollout. In practice, however, the new minimum fare in most cities was 8,000 Rupiah. In order not to introduce any other noise by adjusting these fare bins relative to possibly misidentified minimum thresholds of a given city, we use fare bins in levels.

We also include wage-bin-by-city and wage-bin-by-time fixed effects, $\mu_{j,c}$ and $\rho_{j,t}$ respectively. These terms allow us to control for city-specific fee distributions as well as trends in the nationwide fee distributions. $\epsilon_{c,j,t}$ denotes the error term. We cluster our standard errors by the city level.

Using the estimated $\hat{\beta}_j$s, we can calculate changes in transaction volume over sections of the fee distribution in response to the minimum-fare policy. One of our primary questions is whether the policy reduces transactions below the new minimum-fare threshold as well as increases transactions just above. We define transaction losses as the change in the number of transactions below the new minimum fare over the policy period.

As we can capture the estimated $\hat{\beta}_j$s over the whole distributions of transactions, we can also compute and illustrate the cumulative effect: $\sum_{j=0}^{31} \hat{\beta}_j$. This cumulative effect can be compared to the average treatment effects. Nevertheless, our main aim is to uncover the heterogeneity in impacts to understand better the mechanisms driving the average treatment effects.

4.2 Synthetic control-based method

To address potential concerns arising from the violation of the parallel-trends condition, we also estimate the treatment effect using a version of the synthetic controls estimator first proposed by Abadie and Gardeazabal (2003) and Abadie et al. (2010). At its core, the synthetic control methods allow us to construct a counterfactual using a weighted average of untreated units that best resembles treated units' pre-treatment trends. We deviate from the classic implementation of the synthetic control method, where there would be one treated unit and a set of donors in the control group. Rather, we have multiple treated units and multiple donors in the control group. We approach estimation and inference as follows.

First, we construct a synthetic control for each city in the treated group. In our case, these are the cities in the rollout Phase 2, which we denote with subscript $i \in I$ Cities in the donor pool are those in the rollout Phase 3, denoted with subscript $j \in J$. We restrict the data to April 11 to May 31, and June 11 to August 8, 2019 for the following reasons. First, we match to the pretrends closer to the point of policy variation (July 1st) by dropping data from earlier dates, i.e., those before April 11. Second, we find, as we discuss in further detail in Section 5.2.2, that there are significant, one-time drops in the number of transactions and drivers during the week of *Eid al Fitr* in early June. We found in our preliminary analysis that matching on data from this period generates significant noise, so they are dropped when constructing synthetic controls. Third, we trim the data beyond August 8, the last day before the Phase 3 implementation, to ensure that the donor pool cities are not under the policy regime throughout our sample period.

For an outcome Y at time t each treated city i in Phase 2, we construct $\hat{Y}_i t^N$, i.e., a potential outcome for i at t if it were $\{N\}ot$ treated:

$$\hat{Y}_{i,t}^N = \sum_{j=I+1}^{I+J} \hat{w}_{j,i} Y_{j,t} \tag{3}$$

The weights $\hat{\mathbf{W}} = (\hat{w}_{I+1,i}, ..., \hat{w}_{I+J,i})'$ minimize the distance between i and the donors on the predictors used to determine the relative importance of each donor. Because we do not have relevant covariates that we can gather outside of the data received from the online platform, we use means of the primary outcome variables from the pre-treatment period as predictors with equal weight.[5] They are:

- log(average driver fare)
- log(number of trips)
- log(average daily customer expenditure)
- log(average daily income)
- log(average number of rides per supply hour), i.e., our main measure for productivity

Next, We construct "placebo" synthetic control for each of the donor-pool cities, using the rest of the donor-pool cities as its donor. In other words, for city j in t

$$\hat{Y}_{j,t}^N = \sum_{k=\in J-\{j\}} \hat{w}_{k,j} Y_{k,t} \tag{4}$$

[5]Because our pre-treatment periods are disjoint due to the omission of data from the *Eid* period, we use two means from April 11 to May 31 and June 11 to June 30.

After having constructed synthetic controls for both treated and control cities, we estimate the treatment effect as a weighted average of the difference between actual and synthetic control outcomes over the post-treatment period. Since we have multiple treated cities, we use the log pre-treatment transaction volume as weights for the average treatment effect and denote it as γ_i. In other words, the treatment estimate is:

$$\hat{\hat{\tau}}_e = \sum_{i=1}^{I} \gamma_i \hat{\tau}_i = \sum_{i=1}^{I} \gamma_i (Y_{i,t} - \hat{Y}_{i,t}^N) \tag{5}$$

For inference, we measure the extent to which $\hat{\hat{\tau}}$ is an outlier relative to the standard error of the estimator we propose. In order to assess this, we observe the distribution of "placebo" estimate, defined exactly as in Equation 5, except the i subscript is swapped by j, and $\hat{Y}_{i,t}^N$ with $\hat{Y}_{j,t}^N$. We denote this placebo estimate $\hat{\hat{\tau}}_p$. We construct 999 such estimates via sampling of donor cities with replacement.

For each estimate, $\hat{\hat{\tau}}_e$ or $\hat{\hat{\tau}}_p$, we take the ratio of mean squared root errors (RMSRE) of the synthetic control from the actual outcome between the post- and pre-treatment periods. Here, for the pre-treatment period we use data from April 11 to June 31, including the *Eid* period. We rank the estimates based on the RMSRE. The ranking of how much larger the post-treatment MSRE is relative to the pre-treatment one, divided by 1,000, is the p-value.

We present estimates from difference-in-differences in the main result sections. We find that these synthetic control-based inference procedure estimates are not qualitatively different. The results are included in Appendix Section B.6.

5 Results

The results section is organized as follows. First, in Section 5.1, we present our estimates of effects on average price and quantity. We then show results on adjustment mechanisms in price distribution in Section 5.1.1. We then discuss demand- and supply-side responses in Section 5.2. We proceed with discussions on the distributional impact of this policy on both the demand and supply sides in Section 5.2. We conclude with results showing the implications on labor productivity in Section 5.4. Additional results are included in Appendix Section B, to which we will refer in the remainder of this section when relevant.

5.1 Price and quantity

The first-order empirical question is if and in which direction the policy shifts the average driver fare and quantity demanded, i.e., the number of transactions.

Table 3 summarizes the treatment effects on price and quantity. It shows a 4.6% increase in average driver fare, statistically significant at the 5% level. If we restrict our sample to the type of transactions regulated by the policy, i.e., taxi services, we see an approximately 13-percent increase in the average price that is statistically significant at the 1% level. We further identify which aspects of the price regulation seem to have bound in Section 5.1.1. On transaction volume, we find a minimal point estimate at 0.2 percent, although it is statistically insignificant, and the standard error is much larger than that of the average price. In the subset restricted to the regulated segment, we see a noisy yet large reduction of 10% of transactions. When taken together, the combination of a statistically significant impact on price and a noisy effect on the number of transactions lead to a statistically insignificant yet positive impact on total driver revenue in columns 5 and 6.

We also find shifts in transactions between regulated and unregulated labor subsegments. As discussed in Section 2, the government's policy only regulated taxi trips but not other services carried out by the two-wheel drivers, such as food and delivery. We present results on these possible spillovers in Appendix Section B.3.

A higher average driver fare and a small change to the overall transaction volume should, all else equal, increase daily earnings per driver. Yet, we find that the policy does not significantly increase driver-daily earnings or wages. Table 4 shows that the policy insignificantly reduces drivers' average daily earnings by 1.7% and insignificantly increases it by 2% from the regulated segment. We also find a noisy yet large 6.7% reduction in the imputed average wage, defined as total daily earnings from trips divided by supply hours.

Results from Tables 3 and 4 suggest the presence of some mechanisms through which an increase in drivers' average fare and a noisy yet small effect on overall transaction volume somehow leads to a relatively small and insignificant effect on driver daily earnings. Though noisily estimated, a large reduction in average wage seem to indicate the roles of demand- and supply-side adjustment mechanisms in bringing down drivers' daily earnings to the pre-policy level. We explore this mechanism in Section 5.2.

5.1.1 Binding aspects of the price regulation

A statistically significant increase in the average driver price indicates that at least some aspect of the price regulation, consisting of minimum fare and a floor and ceiling for the per-kilometer rates, is binding. To narrow in on the binding aspects of the regulation, we estimate changes to the distributions of driver fares and conditional average effects on per-kilometer rates by the fare bin, as described in Equation 2.

First, we find that the policy shifts the distribution of transaction prices toward the new minimum threshold from below. We adopt this method to identify changes in employment over wage bins from Cengiz et al. (2019). Figure 2 shows the result by 1,000-rupiah bins of driver fare, with red dash lines showing the cumulative effect from the lowest bin. The figure shows noisy yet large reductions in transactions below the new threshold of 8,000 rupiah and a significant increase at this new threshold. The cumulative effect remains relatively close to zero throughout the rest of the distribution, though it seems to converge further to zero. The changes in the distribution on the right-hand tail depend somewhat on the choice of fixed effects, but the shift from the left to the new threshold is robust across various specifications. We, therefore, find that the minimum total fare aspect of the regulation is binding.

To identify if the constraints on per-kilometer rates is also binding, we estimate changes to the mean and median of the total driver fare divided by distance driven. We do so because we cannot separate fixed (i.e., non distance-dependent) fare from distance-based fare. We estimate the treatment effects by the trip-distance bins at a 500-meter increment. If the fixed fare increases but the per-kilometer fare does not, then we should see positive effects for short distances and the effects fade for longer distances. If, on the other hand, the increases on per-kilometer rates matter, then we should see positive effects even on long distances.

The results are shown in Figure 3. We find positive effects on the "per-km fare," here defined as the total driver fare divided by distance driven, for shorter trips, but statistically significant effects beyond 3 kilometers. Effects on other quantiles, shown in Appendix Figure D.6, however, suggest that the 75th and 95th percentile values of the per-kilometer rate is statistically significantly higher for trips longer than 5 kilometers. We, therefore, find suggestive evidence that the per-kilometer rate is also binding for a subset of transactions.

5.2 Demand and supply-side adjustments

We conjecture that the overall lack of a positive effect on drivers' earnings is driven by increased driver availability, leading to crowding on the supply side. Such supply responses likely cancel out gains from higher fares on earnings in the new equilibrium. To demonstrate this, we show the policy's effect on driver allocation to customers. First, we identify the effect on the overall supply, as shown in Table 5. We find an 8.7% increase in the total supply hours, significant at the 10% level. On the other hand, the policy insignificantly increases the total transaction hours, i.e., hours of all trips demanded, by 1.7%. We have also found in Table 3 that the number of transactions is not significantly affected. These results indicate that the fare policy increases excess supply.

Next, we demonstrate that the policy adjusted the extensive-margin labor supply relative to demand, i.e., the number of distinct drivers on the market in a given day, such that there are fewer customers per available driver. We then show that this adjustment affects the intensive margin, i.e., the number of transactions per driver; we find that the per-person number of trips for both drivers and customers is reduced.

Table 6 shows the extensive margin effect, i.e. changes to numbers of distinct drivers and customers. We find noisy effects on driver and customer headcounts, although the number of distinct drivers seems to have increased by 6.5% for all services, and the number of customers decreased by 6.3% for the regulated segment. When looking at a ratio of distinct customers to drivers, however, we find a significant reduction of 4.5% for all services. When restricted to the regulated segment, the effect sizes are larger at -7.4% and significant at the 5% level. We also note that these effects are not signs of a large and significant influx or exodus of drivers into or from the platform; in Appendix Section B.3.1 we show that the policy does not significantly increase the rate of new drivers' entry.

Table 7 shows our findings on demand- and supply-side responses on the intensive margin. We find that both drivers and customers reduce the number of trips per person; when looking at all service types, drivers reduce their daily number of trips by 6.3% (significant at 10% level), and customers by 1.7% (insignificant at conventional levels). If we restrict to the taxi service type, we identify statistically significant effects of -10.9% and -3.5% for drivers and customers, respectively. These intensive-margin results indicate that the per-customer demand for rides is negatively affected by the policy, though its magnitude is relatively small. Consequently, because of the increased labor supply, there is a larger negative impact on the number of trips allocated per driver.

In summary, the minimum-fare policy increases excess supply by inducing demand and supply adjustments on a combination of intensive and extensive margins, such that fewer trips are allocated to a given driver. Lower allocation seems to have canceled out the effect of a higher per-trip fare, resulting in a statistically insignificant effect on daily earnings with a relatively small coefficient.

5.2.1 Demand-side response and compensation via wait time

Tables 6 and 7 suggest some reduction in the quantity demanded by the customers as a result of policy, though the magnitude is relatively small and most of the effects may be canceled out by the shift into non-regulated service types. In this section, we explore the extent of the policy's incidence on customers and potential compensations in wait time. As such, we look at the average treatment effects on per-driver daily estimated expenditure and a proxy of wait time. There are, however, some limitations to these outcome measures, and results should be taken with caution; first, the customer fare, a daily customer-level sum of which is used as their daily expenditure, is inclusive of any credits and discounts they may receive. Second, our proxy of the wait time is "wait distance," i.e., the linear distance between where the driver accepted the order and where it is picked up.

Table 8 shows the results on driver-level daily expenditures and their wait time per trip. We find that the policy increased average daily expenditure on trips by customers by 8.0% and 19.6% if restricted to the regulated segment. The increased customer expenditure, however, is coupled with a significant reduction in wait time; we find a significant 14.9% reduction in wait distance for all services and a 7.6% reduction in the regulated segment. These results suggest that customers may have incurred a higher cost of trips, though we do not know if or the extent to which this effect is canceled out by credits and discounts. We also find that customers are, on average, compensated for seemingly higher fares by shorter wait time. However, we find in Section 5.3 that this compensation is not differentially higher for trips whose fares are increased due to the policy.

We also estimate demand and supply elasticities from the policy variation, though results on this should also be taken with caution due to the measurement concern on the customer fare. We use the variation in average fare from the policy. Consumer- and producer-facing prices differ because the platform charges transaction fees and may implement other policies that may dissociate driver payments from customer fees. We show that the increased driver fare per government policy is passed through to the consumer in Appendix Section B.4, and that noisily estimated demand and supply elasticities suggest that the average elasticity is low for demand and high for supply in Appendix Section B.5.

5.2.2 Identification: Parallel trends

As discussed in Section 4, the causal interpretation of our coefficient estimate hinges on the assumption that, in the absence of the minimum-fare policy, trends of the outcome variables for the phase-2 rollout cities are identical to those of the phase-3 cities. Although the counterfactual itself is unobservable, we test for evidence of parallel pre-trends in the four main outcomes: average price, number of transactions, driver-average daily earnings, and driver-average number of trips per supply hour (i.e. a productivity measure).

We use a distributed lag model to estimate the presence of pre-trends. We identify weekly treatment effects 6 weeks into the policy variation period, as well as 7 weeks prior. We use the following equation for the city-day level outcomes where subscript c denotes city, and t and w denote day and calendar week as units of time respectively.

$$Y_{c,w,t} = \beta_0 + \Sigma_{k=-6}^{6} \beta_1^k * I_{c,t}(c \in Treat, w = k) + \gamma_c + \rho_w + \epsilon_{c,t} \tag{6}$$

$\hat{\beta_1^k}$ is the DiD estimate for week k, where $k = 0$ corresponds to the week right before the start of policy variation. Figure 4 shows the results. Subfigure (a) shows that the weekly pre-policy coefficients ($k \leq 0$) on the log-average driver fare outcome are small statistically indistinguishable from 0. We then observe a significant positive jump at $k = 1$, and remains at the same level into week 6. We observe similar trends, except with a negative impact post-policy, on the productivity outcome in Subfigure (e). On the log-number of transactions and drivers, however, we find evidence of differential pre-treatment patterns, especially in weeks -4 to -2, as shown in Subfigures (b) and (c). This significant pre-trend in the transaction outcome may also contribute to the noise in the average treatment effect. In Subfigure (d), we also find differential pre-trends in the driver-average daily revenue outcome, though the magnitude of it is significantly smaller.

Upon closer inspection of the data, we find that the points of significantly different pre-treatment patterns in Subfigures (b) and (c) coincide with Eid al-Fitr, a major Islamic holiday marking the end of Ramadan. We find that the transaction volume and the number of market participants drop substantially in most cities. We find that this drop-off happened to be significantly larger for the Phase-2 (i.e., "treated") cities than in Phase-3 ones. We also find that this differential drop-off in transaction is particularly pronounced in lower-earning and productivity deciles, as Appendix Section B.1 shows.

We address the deviations from parallel trends condition in the following ways. First, we estimate the average treatment effects on data excluding the Eid al-Fitr period on outcomes for which the parallel trend is shown not to hold in Figure 4. The results are presented in Appendix Tables C.9 and C.10, showing that the estimates and the statistical significance of the results are qualitatively unchanged. Second, we present the results from analysis using the synthetic control-based inference procedure in Appendix Section B.6, and find qualitatively similar results as ones from difference-in-differences. Third, for analysis using panel data at the city-bin-day level, we include a bin-date fixed effects on top of city-bin ones to control for bin-specific trends to the best extent possible.

5.3 Distributional impact

In Section 5.2, we identify average demand- and supply-side responses to the regulation. However, price-floor policies often aim to protect and transfer welfare to the most vulnerable or lowest-earning workers in the labor market. In this subsection, we assess if and how the policy achieves its potential goal of redistribution. We focus on two aspects of driver and customer heterogeneity: earnings/expenditure and potential exposure of their supply/consumption to the policy.

We run the DiD model by pre-policy transaction-volume quantiles as described in Section 4.1.2. Aside from the main specification, which includes city-bin and bin-date fixed effects, we present an alternative version that includes city-bin and date fixed effects. We select four measures from which to construct deciles of individuals' transaction volume and policy exposure, all from the month prior to policy variation (June 1 to 30, 2019). We restrict the sample to existing drivers and customers. The measures are:

- Drivers' total earnings from trips
- Drivers' share of total earnings that would have qualified for the minimum fare
- Customers' total expenditures for trips
- Customers' share of total expenditures that would have qualified for the minimum fare

5.3.1 Distributional impact on the supply side

Figures 5 and D.7, as well as their Appendix analogues, show the conditional average treatment effects by drivers' pre-policy earnings and policy exposure deciles, respectively. The interpretations of subfigures are as follows. In subfigures (a), we show the conditional average treatment effects on driver fare by decile. We then show the effect on aggregate supply by the decile in subfigure (b). To identify mechanisms on the aggregate-supply effect, we break down the effects into extensive and intensive margins, i.e. the number of distinct drivers and supply hours per driver, in subfigures (c) and (d). We then identify the effects on driver earnings and wages on subfigures (e) and (f), respectively.

Figure 5 shows that while low-earning drivers increase their supply, they do not experience increased earnings. Subfigure (a) shows that the policy increases the per-trip driver fare fairly uniformly across earning deciles, though the effect seems to be slightly larger for higher income deciles. The larger effect on the highest earnings deciles does not seem to be robust to the choice of fixed effects and the span of the pre-policy period over which to calculate the pre-policy decile. Subfigure (b) shows that the policy increased the labor supply of workers in the 4 lowest pre-policy earning deciles, by approximately 15 to 40%. We then break down these effects on aggregate supply in subfigures (c) and (d) into the extensive and intensive margins, respectively. Subfigure (c) shows that the policy increases the number of distinct drivers participating from the bottom 3 to 4 deciles by 10 to 20%, depending on the decile and the set of fixed effects. Subfigure (D) shows that, conditional on driver's availability, the policy increases the driver supply hours by 5 to 20%, again depending on the decile and fixed effects. The consequences of the supply responses, and its intensive- and extensive-margin effects, on earnings and wages are shown in subfigures (e) and (f). We find that the estimates on the effects on earnings for the low earners are smaller, at around 5 to 10%, and statistically indistinguishable from zero. In sum, we find evidence consistent with the overall supply-side responses highlighted in Section 5.2; large increases in excess supply comes from lower earners, but their effects on post-policy earnings are canceled out by the supply-side crowding.

Appendix Figure D.7 also provides suggestive evidence that confirms the supply side mechanisms in response to higher fare. The figures show that the policy increases the average fare for drivers most exposed to the policy, but it did not increase their driver earnings. In sum, we find that the price policy induces supply responses from lower earners and those most exposed to the policy, but the impact on earnings and wages is limited.

5.3.2 Distributional impact on the demand side

Figures 6 and 7 show the distributional impact by the differential impact on customers' outcomes by their expenditure and policy exposure, respectively. In these figures, we identify the distributional impact on the quantity and price of customers' orders, as well as their daily expenditure and the compensation mechanism in shorter wait time, in subfigures (a), (b), (c), and (d), respectively.

Figure 6 shows a relative lack of differential impact by customers' pre-policy expenditure levels. We find a uniform increase in customer fare and their expenditure across the pre-policy expenditure bins, accompanied by a uniform reduction in the wait time. Figure 7, on the other hand, shows a more distinct pattern by the policy exposure decile. We find that those most exposed to the policy reduce their number of trips slightly and face significantly higher per-trip fares. This results in higher expenditure for the most-exposed customers. Interestingly, higher fare and expenditure are not coupled with a greater reduction in the wait distance, as shown in subfigure (d).

In summary, we find that the incidence of the policy falls significantly on customers whose preferred types of trips were the target of the minimum fare. They are relatively demand-inelastic and absorb the higher cost of a trip induced by the policy. It is possible that the minimum-fare policy functions as a transfer from customers with high, inelastic demand to workers with low earnings, as may be an intention of a price-floor policy. The evidence for such a mechanism is relatively weak and should not be taken for granted. We also find that the most exposed customers are not compensated for their higher incidence with shorter wait time any more than less frequent ones. This result indicates that shorter wait time is a byproduct of supply-side adjustments in the market at large rather than a mechanism to offset the impact of certain customers' higher fares.

5.4 Effects on labor productivity

A potential implication of increased labor supply that competes over a fixed amount of transactions, along with other distortions and reallocations that the pricing policy might have introduced, is that drivers are less effective at finding and executing their trips. In this section, we investigate the impact on driver productivity. We identify productivity effects on two channels; reallocation of work to less productive workers and reduction in individual productivity. We find suggestive evidence that the policy reduces average driver productivity, and this is due to both increased participation of less productive drivers, as well as a reduction in individual productivity regardless of their pre-policy productivity levels.

We define three measures of daily labor productivity as follows: the number of paid rides, distance traveled on those paid rides, and time engaged in them, all over supply hours. The last measure–time engaged in paid rides per supply hour—is also referred to as utilization rate. We note that, given a positive effect on total labor supply in Table 4 and a small and statistically insignificant effect on an overall number of transactions in Table 3, there should be a negative mechanical effect on average productivity measures. In fact, we find in Table 9 that using these three measures, there is approximately an 8 to 10% reduction in labor productivity, some of which is statistically significant at the 10% level.

5.4.1 Do less productive workers participate more?

One possible mechanism behind the reduced average driver productivity is that the policy induces less productive drivers to participate more frequently, thereby increasing the share of inherently less productive drivers in the fleet. This seems plausible, as we find in Section 5.3 that lower-earning drivers increase their labor supply. If they are also less productive according to our measures, then the resulting stock drivers on the market would be less productive.

We find evidence for this mechanism in Figure 8. We identify effects on the log daily number of distinct drivers by the pre-policy productivity decile, with the city-bin fixed effects, as well as either date or bin-date fixed effects. We use distance traveled per supply hour as the pre-policy productivity measure. We find that drivers in the two to three lowest pre-policy productivity deciles increase daily participation significantly, depending on the set of fixed effects. On the other hand, drivers in the highest deciles of pre-policy productivity may reduce daily participation, though the results do not hold with more robust fixed effects. These results suggest that the policy induces a compositional shift in the fleet toward less productive drivers.

5.4.2 Are drivers less productive on average?

Another possible mechanism is that, as a result of the crowded labor market or changes to the on-the-job incentives, drivers are less productive individually than they were before the introduction of the policy. To assess this possibility, we again look at treatment effects by pre-policy productivity deciles. The outcome measure this time, however, is the productivity measure itself. Results in Figure 9, indicate a noisy reduction in individual productivity of the similar magnitude as in Table 9, depending on the decile.The results, therefore, imply that crowding on the supply side, in turn, reduced individual productivity.

6 Discussions

Our empirical results show that a price-floor policy in a ridesharing market increases labor supply from the lower-earning and less productive drivers. We also find that excess supply crowds the supply side, canceling out the effects of higher per-trip fares on driver wages. Increased labor supply and inelastic demand also reduce individual productivity by lowering the match rate.

In this section, we introduce a conceptual framework to highlight the adjustment mechanisms of the minimum-fare policy through the supply-side response and its impact on the match rate. We illustrate that an exogenous trip-price increase leads to an increase in labor supply and a reduction in match rate. We find a positive effect on wages, but the relative magnitude to the labor supply response depends on parameter values, such as the convexity of the labor cost function.

We then compare the effect sizes from our analysis with another study with the most similar empirical setting—a study by Hall et al. (2021) on the effect of Uber's base-fare changes on drivers' responses, transaction outcomes, and productivity. Based on the conceptual framework and its prediction, we compare the results with Hall et al. (2021) and hypothesize on the reasons behind the similarities and differences of our results.

6.1 Conceptual framework: setup

The conceptual framework focuses on market-wide effects and abstracts away from details on the search and matching processes. Many papers studying pricing policies on ridesharing platforms, such as Frechette et al. (2019) and Castillo (2020), focus on dynamic and spatial aspects of search and match and benefits of centralized allocation system and surge pricing, respectively. Instead, we take an approach similar to that of Hall et al. (2021) by focusing on market-wide effects on labor supply and earnings.

We imagine a marketplace in equilibrium where platform operators derive piece-rate price p for a representative trip in the market via an algorithm.[6] In other words, we treat the price p as being set exogenously by the platform rather than being endogenously determined to clear the market. Driver availability and demand are then endogenously determined in response to the exogenously set price. Trip allocations are made to drivers via an algorithm

[6]In reality, two platforms compete against each other for customers while maintaining an efficient fleet of drivers. We abstract away from aspects of platform competition because, anecdotally, these platforms engage in price competition and offer relatively homogeneous prices and services on a given trip attribute. We focus on market-wide responses to government regulation on pricing, assuming that the two platforms respond similarly.

rather than to the lowest bidders, leading to some number of idle drivers (i.e., excess supply) in equilibrium. The assumptions we make are similar to those of Hall et al. (2021).

Drivers are matched with a trip at a rate x, which is endogenous to demand and supply. A driver i decides how many hours l_i to make themselves available on the app. Not all units of l_i translates into earnings, as some are spent idling. Drivers determine l_i based on the exogenously determined price of a ride p, match rate x, and their cost function in terms of labor supplied, $c()$, where $c'(.) > 0$ and $c''(.) > 0$. A driver's utility function is $U_i^d = pxl_i - c(l_i)$. We include individual subscripts to account for driver heterogeneity in a latter subsection. Market-wide supply is $L = \sum_{i \in \mathbb{I}} l_i$.

On the demand side, a customer's quantity demanded is a function of price and wait time w, which itself is a function of match rate x. Customer's utility is $U^c = f(p, w(x))$, and the market demand $D(f(p, w(x))$, where $D'(.) < 0$. Because we do not find significant demand-side responses empirically in Section 5.2, we simplify the demand-side equation to be of homogeneous agents responding deterministically to the price and wait time.

In equilibrium, drivers maximize their utility by selecting l_i subject to the market price p and match rate x, demand D, and the match rate x. In other words:

$$px - c^{'}(l_i^*) = 0 \tag{7}$$

and,

$$x^* = \frac{D}{L} \tag{8}$$

6.2 Conceptual framework: Comparative statics

Next, we identify the effect of exogenous price increases on driver supply, match rate, and wages. For now, we treat drivers as homogeneous agents of a unit mass. Implicitly differentiating the driver first-order condition in Equation 7 with respect to p, we get:

$$\frac{dl_i}{dp} = \frac{x + p\frac{dx}{dp}}{c''} \tag{9}$$

Differentiating Equation 8 with respect to p gives us:

$$\frac{dx}{dp} = \frac{1}{L}\frac{dD}{dp} - \frac{x}{L}\frac{dL}{dp} = -\frac{x}{L}\frac{dL}{dp} = -\frac{x}{L}\frac{dl_i}{dp} \tag{10}$$

The last equality is true if $\frac{dD}{dp} = 0$, a scenario which seems to fit with empirical evidence and assertion we make to simplify our comparative static results. Plugging 10 into 9 gives:

$$\frac{dl_i}{dp} = \frac{xL}{c''L + xp} > 0. \tag{11}$$

Proposition 1 *With homogeneous drivers in the market, an exogenous price shock increases driver supply by* $\frac{xL}{c''L+xp} > 0$.

Corollary 1.1 *With homogeneous drivers in the market, an exogenous price shock reduces match rate* x *by* $-\frac{x^2}{c''L+xp} < 0$.

We define wage as xp, i.e., the expected earnings per unit of labor supply l_i. Using equations above and totally differentiating xp, we derive the response in wage to a change in price p:

$$\frac{dxp}{dp} = x + p\frac{dx}{dp} = \frac{xc''L}{c''L + xp} > 0. \tag{12}$$

Proposition 2 *With homogeneous drivers in the market, an exogenous price shock increases driver wage by* $\frac{xc''L}{c''L+x+xp} > 0$.

Propositions 1 and 2 show that with an exogenous price increase, driver supply and wages increase. We also see that the effect on wages is a combination of those on prices and match rate. The effect on wage coming from the changes in match rate is negative, as shown in

Corollary 1.1. The effect of exogenous price increases on wages is, therefore, tampered by a reduction in match rate.

There are some areas in which our empirical results do not necessarily align with the findings from our conceptual framework; we do not find evidence for increased average wage. This could be because of parameter values; for instance, our theoretical framework would predict a small effect on the wage if workers' cost function is not very convex (i.e., low c''). Second, the matching rate may decrease exponentially to crowded supply side in reality, which our conceptual framework does not capture. Third, agents may be heterogeneous, and so are their labor-supply responses and effects on their match rate, a key parameter in driver productivity.

6.3 Comparisons with other empirical estimates of pricing policy

We compare our empirical estimates to other work that studies the effect of pricing in ridesharing markets on labor supply, match rate, and wages. Our main point of comparison is Hall et al. (2021) (HHK for the remainder of this subsection), which studies the effect of Uber's base-fare changes on drivers' responses, transaction outcomes, and productivity. Following the setup and results of the conceptual framework, we focus on the following outcomes:

- Number of trips
- Aggregate supply hours
- Productivity (kilometers driven on trip/supply hour)
- Wage

The rationale behind these choices are as follows; first we compare the effect on the number of transactions, as we consider it to be determined by the demand-side responses.

We confirm that the aggregate demand effect between our context and HHK's are similar, in that we both find limited responses. We then identify the relative effect sizes on the aggregate supply hours, i.e., the choice variable of the conceptual framework. We then analyze how the choices on supply influences productivity (and its analogue, the match rate between drivers and customers) and driver wage.

To compare effect sizes, we standardize our estimates into elasticities to driver fare, i.e., the percentage change in the outcome variable in response to a percentage change in the average driver fare. We use the two-stage least squares (2SLS) regression, where we instrument for the driver fare with the policy variation variable $I_{c,t}(c \in Treat, t > 0)$, the regressor of the main specification in Equation 1. We provide further detail of the 2SLS approach we use in Appendix Section B.5.

We find that our estimates are in line with those of HHK, albeit with differences in magnitude. Table 10 shows the elasticity estimates for some of our key outcome variables. The table also reports corresponding estimates from HHK. First, Column 1 shows that the effect on the number of transactions is small and indistinguishable from 0 in both our estimates and HHK, making it easier to compare the supply-side effects. Second, Column 2 shows that our supply hour elasticity estimate of 1.76 is, though statistically insignificant, more than five times greater than that of HHK (0.34). The drastic difference in the supply response may drive the difference in the magnitude of productivity elasticity, defined as kilometers driven on the trip per supply hour, in Column 3. We find the productivity elasticity of -1.87, nearly three times as much as in HKK (-0.66). Lastly, Column 4 shows that the elasticity on wage is negative in our estimation at -1.37, as opposed to small and negative in HHK of 0.075, though neither are statistically significant.

We speculate that the larger labor-supply response in our context than those in HHK drives the differences in the effects on productivity and wages. We first note the differences in contexts, such as the locale and the structure of the price policy—a price floor versus a uniform increase in the base fare. These contextual differences make it difficult to attribute the cause of a larger labor-supply response in our findings. However, conditional on this effect on labor supply, our results, combined with HHK's, seem to confirm that supply-side crowding reduces productivity and wages. We provide a conceptual framework to elucidate this mechanism in the next section.

7 Conclusion

Regulation to improve workers' earnings via a price floor on a unit of labor may not achieve its intended policy targets when considering its effects in market equilibrium. In this article, we study the impact of the introduction of a minimum fare policy for ridesharing app workers in Indonesia. We find that the exogenously shifted price of labor would have a limited impact on workers' average earnings in equilibrium. We identify demand- and supply-side mechanisms through which such muted impact would result. We also find that the average effect on driver earnings masks the heterogeneous impact, where lower earners supply more labor but may not necessarily earn more. We also find suggestive evidence that adjustments from the policy led to reduced worker productivity via changes in the composition of workers and individual reduction in productivity. These findings have several important implications for the economics literature and policy-making.

First, the findings add to our understanding of regulating informal work in developing countries. Previous work such as Muralidharan et al. (2017) studies policies that guaranteed both the quantity (i.e., number of days employed) and price (daily earnings) of work to casual laborers, leading to positive earnings and structural transformation. In our context, however, only the price on the piece rate is guaranteed, and the market-equilibrium effect cancels out higher payments per transaction for an average worker. These differences in the policy design—and likely the scale and the fiscal commitment from the government—may have contributed to different outcomes between our analysis and one by Muralidharan et al. (2017). Our findings suggest that a simple minimum wage-like policy would not trigger increased earnings for all workers, even when enforcement of such policy is made feasible by the increased presence of online platform-mediated marketplaces. Instead, the policy may induce more labor supply from lower earners and induce competition over a given amount of demand. We are, however, unable to offer definitive insight as to whether the policy would result in a meaningful redistribution toward lower earners.

Second, our findings also provide novel insights into market-wide implications of and adjustment mechanisms triggered by the implementation of minimum wage, despite some contextual differences. We find noisy yet small effect sizes on the overall transaction volume and shifts in the price distribution, similar to findings of Cengiz et al. (2019). Our results suggest that a small overall effect on disemployment can occur without monopsony power, although the statistical precision on that statement is limited. Our results are also in contrast with Jardim et al. (2018), who finds lower net earnings via reduced hours worked. We find a relatively tight estimate of around zero on average earnings. Furthermore, we find that through crowding on the supply side, the policy *reduces* productivity in equilibrium. This result stands in contrast to previous work that found that an increased minimum wage would attract more efficient workers when offered by more localized, specific employers rather than across the market (e.g., Ku 2022; Coviello et al. 2021). The fact that we find a reduction in productivity implies that regulation on labor price may have negative allocative consequences; with the policy, it now takes more working hours to deliver the same amount of transactions, while the excess labor supply could have been more productive elsewhere. Hence, our findings suggest that policymakers may want to weigh the allocative costs against any potential benefit of the policy.

Lastly, our findings also provide insights into the efficacy of regulating labor markets on online platforms. Our results are in line with findings from previous work on pricing policy like Horton (2017), who show increased labor supply and crowding to cancel out the effects of higher piece-rate on earnings. Our additional contributions are the distributional impact that may be driven by the policy design; minimum fares rather than uniform increases in fare may affect some drivers positively and others negatively. Our findings, therefore, suggest not only limited efficacy of minimum fare on earnings at the cost of lower labor productivity but also potentially uneven effects on drivers based on their pre-policy productivity and the extent of reliance on ridesharing as a source of income.

References

Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132, 2003.

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.

Kathleen Beegle, Emanuela Galasso, and Jessica Goldberg. Direct and indirect effects of malawi's public works program on food security. *Journal of Development Economics*, 128: 1–23, 2017.

David Card and Alan B Krueger. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *The American Economic Review*, 84 (4):772, 1994.

Juan Camilo Castillo. Who benefits from surge pricing? *Available at SSRN 3245533*, 2020.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454, 2019.

Decio Coviello, Erika Deserranno, and Nicola Persico. Minimum wage and individual worker productivity: Evidence from a large us retailer. 2021.

Pascaline Dupas, Jonathan Robinson, and Santiago Saavedra. The daily grind: Cash needs and labor supply. *Journal of Economic Behavior & Organization*, 177:399–414, 2020.

Christian Dustmann, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge. Reallocation effects of the minimum wage. *The Quarterly Journal of Economics*, 137(1):267–328, 2022.

Niklas Engbom and Christian Moser. Earnings inequality and the minimum wage: Evidence from brazil. *American Economic Review*, forthcoming.

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

Jonathan V Hall, John J Horton, and Daniel T Knoepfle. Pricing in designed markets: The case of ride-sharing. Technical report, Working paper, Massachusetts Institute of Technology, 2021.

John J Horton. Price floors and employer preferences: Evidence from a minimum wage experiment. *Available at SSRN 2898827*, 2017.

Clement Imbert and John Papp. Labor market effects of social programs: Evidence from india's employment guarantee. *American Economic Journal: Applied Economics*, 7(2): 233–63, 2015.

Indonesian Ministry of Transportation. Decree of the minister of transportation of the republic of indonesia number 348, 2019.

International Conference of Labour Statisticians. Resolution concerning statistics of employment in the informal sector, adopted by the fifteenth international conference of labour statisticians. In *The Fifteenth International Conference of Labour Statisticians*, 1993.

Ekaterina Jardim, Mark C Long, Robert Plotnick, Emma Van Inwegen, Jacob Vigdor, and Hilary Wething. Minimum wage increases and individual employment trajectories. Technical report, National Bureau of Economic Research, 2018.

Anjini Kochar. Explaining household vulnerability to idiosyncratic income shocks. *The American Economic Review*, 85(2):159–164, 1995.

Anjini Kochar. Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural india. *Review of Economics and Statistics*, 81(1): 50–61, 1999.

Hyejin Ku. Does minimum wage increase labor productivity? evidence from piece rate workers. *Journal of Labor Economics*, 40(2):325–359, 2022.

Rafael La Porta and Andrei Shleifer. Informality and development. *Journal of economic perspectives*, 28(3):109–26, 2014.

Measurable AI. Ride-hailing race in indonesia: Gojek versus grab. https://blog.measurable.ai/2022/01/18/ride-hailing-marketshare-in-southeastasia-indonesia-gojek-versus-grab/, 2022. Accessed: 2022-10-27.

Mordor Intelligence. Asean taxi market - growth, trends, covid-19 impact, and forecasts (2022 - 2027). https://www.mordorintelligence.com/industry-reports/asean-taxi-market/, 2022. Accessed: 2022-10-27.

Karthik Muralidharan, Paul Niehaus, and Sandip Sukhtankar. General equilibrium effects of (improving) public employment programs: Experimental evidence from india. Technical report, National Bureau of Economic Research, 2017.

David Neumark and Luis Felipe Munguia Corella. Do minimum wages reduce employment in developing countries? a survey and exploration of conflicting evidence. *World Development*, 137:105165, 2021.

Statista. Ride-hailing and taxi - indonesia. https://www.statista.com/outlook/mmo/shared-mobility/shared-rides/ride-hailing-taxi/indonesia/, 2022. Accessed: 2022-10-27.

Gabriel Ulyssea. Firms, informality, and development: Theory and evidence from brazil. *American Economic Review*, 108(8):2015–47, 2018.

World Bank. *World development report 2019: The changing nature of work*. The World Bank, 2018.

8 Tables

Table 1: Timeline of fare regulation

Time	Description
11 March 2019	Safety regulation released. The Minister of Transportation stated that the Ministry is still working on the minimum fare regulation. He hinted it will be around IDR 2,400/km.
25 March 2019	The Ministry of Transportation released the fare regulation.
1 May 2019	Start of the implementation of the regulation.
1 May 2019	First phase of the fare regulation's implementation: Jakarta, Bandung, Yogyakarta, Surabaya, dan Makassar.
1 July 2019	Second phase of the fare regulation's implementationn: 41 cities.
9 August 2019	Third phase of the fare regulation's implementation: 123 cities.
2 September 2019	Fare regulation is implemented in all cities where ride-share platforms operate.

Source: Minister of Transportation Decree No. 348 Year 2019 and various news sources. List of news sources are available upon request.

Table 2: Structure of the fare regulation

Zone	per-km min-max fares	minimum total fare
Zone 1	Rp 1,850-2,300	Rp 7,000 - 10,000
Zone 2	Rp 2,000-2,500	Rp 8,000 - 10,000
Zone 3	Rp 2,100-2,600	Rp 7,000 - 10,000

Source: Minister of Transportation Decree No. 348 Year 2019.

Table 3: Average treatment effects on driver fare and number of transactions

	log(Avg driver fare)		log(N trips)		log(Sum driver fare)	
	All services	Regulated	All services	Regulated	All services	Regulated
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.0461**	0.1286***	0.0021	-0.0976	0.0483	0.0310
	(0.0177)	(0.0322)	(0.0829)	(0.0914)	(0.0751)	(0.0813)
Observations	12,760	12,760	12,760	12,760	12,760	12,760
R^2	0.93870	0.91247	0.98193	0.98331	0.98381	0.98426
Within R^2	0.03673	0.14282	3.5×10^{-6}	0.00605	0.00189	0.00066
Day fixed effects	✓	✓	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓	✓	✓

Notes: All dependent variables are in log. "Avg driver fare": The city-day average of the price drivers receive for a ride. "N trips": Number of trips per city per day. "Sum driver fare": City-day level aggregate of the driver fare. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 4: Average treatment effects on driver daily earnings, and wage

	log(Avg earnings/day)		log(Avg wage)
	All services	Regulated	
	(1)	(2)	(3)
Treat	-0.0167	0.0199	-0.0674
	(0.0248)	(0.0347)	(0.0503)
Observations	12,760	12,760	10,962
R^2	0.92638	0.94523	0.82260
Within R^2	0.00132	0.00134	0.00902
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. "Avg earnings/day": City-day average of drivers' daily earnings from fares. "Avg wage": City-day average of drivers daily earnings, divided by their daily total available hours on the app. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 5: Average treatment effects on driver supply hours, and trip duration

	log(Sum supply hrs)	log(Sum transaction hrs)	log(Sum idle hrs)
	(1)	(2)	(3)
Treat	0.0865*	0.0169	0.2430***
	(0.0500)	(0.0738)	(0.0886)
Observations	10,962	12,760	10,912
R^2	0.98475	0.98198	0.92888
Within R^2	0.00996	0.00024	0.01580
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. "Sum supply": City-day aggregate of drivers' daily total available hours on the app. "Sum transaction hrs": City-day aggregate of the total duration of all trips conducted. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 6: Average treatment effects on the extensive margin market participation

	log(N drivers)		log(N customers)		log(N customer/driver)	
	All services	Regulated	All services	Regulated	All services	Regulated
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.0649	0.0111	0.0195	-0.0630	-0.0454*	-0.0741**
	(0.0605)	(0.0599)	(0.0749)	(0.0826)	(0.0262)	(0.0327)
Observations	12,760	12,760	12,760	12,760	12,760	12,760
R^2	0.98460	0.98515	0.98345	0.98470	0.88919	0.95239
Within R^2	0.00494	0.00014	0.00035	0.00295	0.01159	0.02297
Day fixed effects	✓	✓	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓	✓	✓

Notes: All dependent variables are in log. "N drivers": Number of distinct drivers at the city-day level. "N customers": Number of distinct customers at the city-day level. "N customer/driver": Number of distinct customers divided by the number of distinct drivers. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 7: Average treatment effects on the intensive margin market participation

	log(Avg n. bookings/driver)		log(Avg n. bookings/customer)	
	All services	Regulated	All services	Regulated
	(1)	(2)	(3)	(4)
Treat	-0.0628*	-0.1087**	-0.0173	-0.0346***
	(0.0322)	(0.0409)	(0.0115)	(0.0116)
Observations	12,760	12,760	12,760	12,760
R^2	0.91089	0.95232	0.84742	0.83562
Within R^2	0.01750	0.03698	0.00751	0.03185
Day fixed effects	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓

Notes: All dependent variables are in log. "Avg n. bookings/driver": Average number of daily trips per driver. "Avg n. bookings/customer: Average number of daily trips per customer. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 8: Average treatment effects on customer expenditure and wait time

	log(Avg expenditure/day)		log(Avg wait distance)	
	All services	Regulated	All services	Regulated
	(1)	(2)	(3)	(4)
Treat	0.0801*	0.1957***	-0.1488***	-0.0759*
	(0.0408)	(0.0391)	(0.0472)	(0.0387)
Observations	12,760	12,760	7,620	7,620
R^2	0.82291	0.89509	0.85736	0.90391
Within R^2	0.02451	0.17482	0.03352	0.01501
Day fixed effects	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓

Notes: All dependent variables are in log. "Avg expenditure/day": Average of total customer fare per customer per day, at the city-day level. "Avg wait distance": Average of the linear distance between the accepted driver and the pickup point per trip, at the city-day level. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 9: Average treatment effects on productivity

	log(Avg util. rate) (1)	log(Avg km/supp hr) (2)	log(Avg rides/supp hr) (3)
Treat	-0.0800	-0.0919*	-0.1050*
	(0.0516)	(0.0521)	(0.0556)
Observations	10,962	10,962	10,962
R^2	0.75871	0.87826	0.83735
Within R^2	0.00852	0.01600	0.02264
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. "Avg util. rate": average utilization rate, i.e., the number of hours spent on paying trips divided by active hours spent on app, i.e., "supply hours." "Avg km/supp hr": average cumulative distance traveled on paying trips per supply hour. "Avg rides/supp hr": average daily number of rides per supply hour. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table 10: Elasticities to driver fare and comparisons from Hall et al. (2021)

	log(N trips) (1)	log(Sum supply hrs) (2)	log(Avg km/supp hr) (3)	log(Avg wage) (4)
log(Avg customer fare)	0.0219			
	(0.8594)			
log(Avg driver fare)		1.756	-1.867**	-1.370
		(1.242)	(0.7973)	(0.8560)
Observations	12,760	10,962	10,962	10,962
R^2	0.98178	0.98004	0.84658	0.80637
Within R^2	-0.00832	-0.29534	-0.24009	-0.08163
Day fixed effects	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓
HHK estimates	-0.099	0.342***	-0.655***	0.075
HHK SEs	(0.081)	(0.034)	(0.059)	(0.064)

Notes: All dependent variables are in log. Point estimates are elasticity 2SLS estimates based on Equation 15. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

9 Figures

Figure 1: Average driver fare, by policy rollout phase

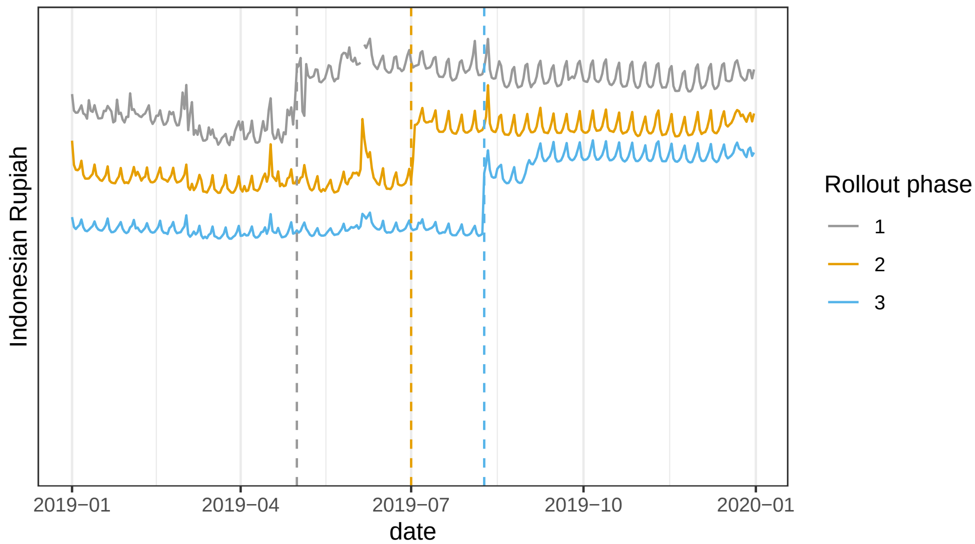

Notes: We plot the daily averages of per-trip fares received by drivers for each rollout phases. The vertical scale does not have value labels for confidentiality reasons, but the range of the axis starts from 0. Vertical dashed lines are the timing of policy rollout for the phases in the corresponding color. The first rollout phase was on 1 May 2019. The second rollout phase was on 1 July 2019. The third rollout phase was on 9 August 2019.

Figure 2: Treatment effect on frequency by 1,000-rupiah bins

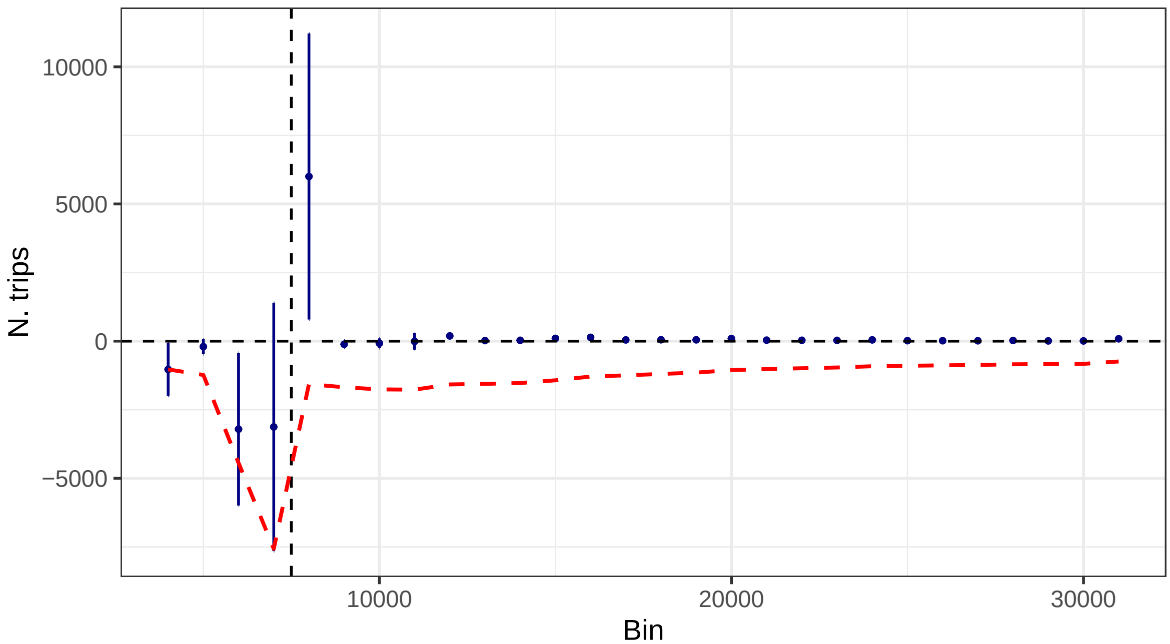

Notes: The dependent variable is the number of transactions in 1000-Rupiah bins. Point estimates in blue are the estimated impact on each bin, $\hat{\beta}^j$, as in equation 2. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The dashed red line shows the cumulative effect, i.e., the sum of the point estimates, from the left to right. Regression is run on a bins-city by day panel data. The model includes fare-bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the city-bin level.

Figure 3: Treatment effects on average and median driver fare divided by distance driven, by the 500-meter trip distance bin

(a) Average per-km fare

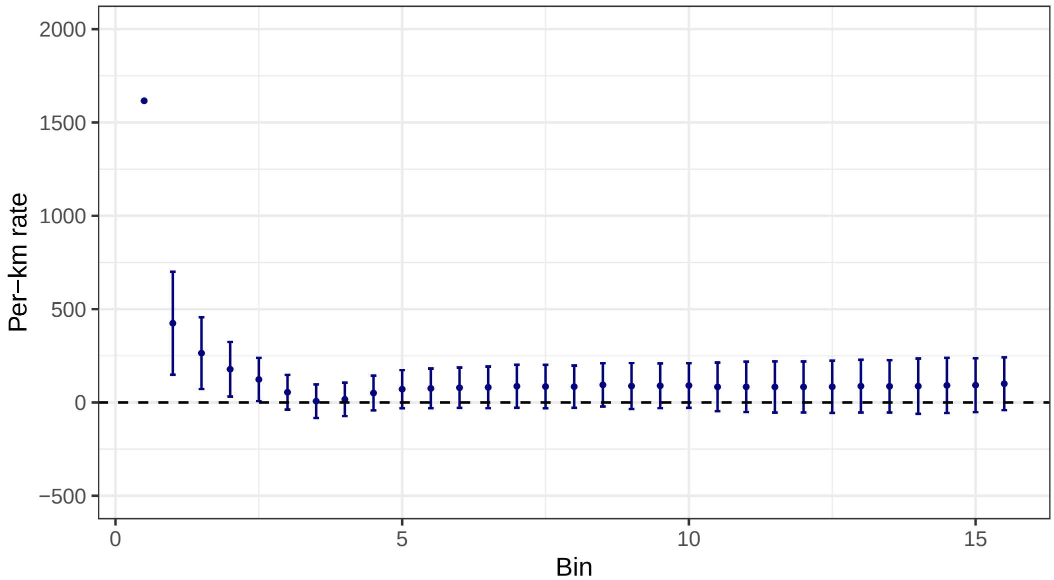

(b) Median per-km fare

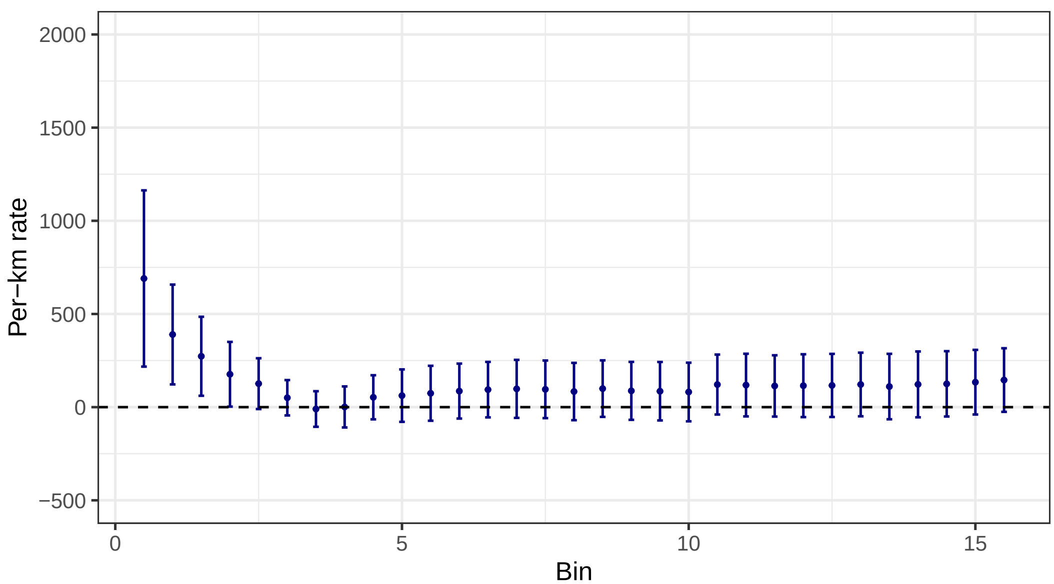

Notes: The dependent variables are the conditional average and median measures of total driver fare divided distance driven for subfigures (a) and (b), respectively. Point estimates in blue are the estimated impact for each 500-meter trip distance bin. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the city-bin level..

Figure 4: Weekly treatment effects on average driver fare and number of bookings

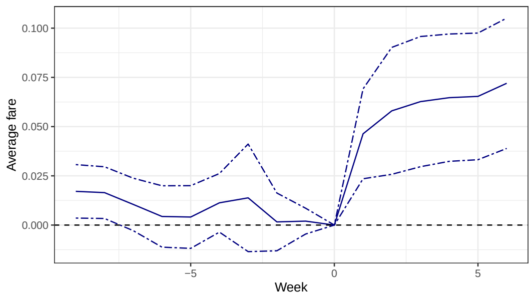

(a) Log(average driver fare)

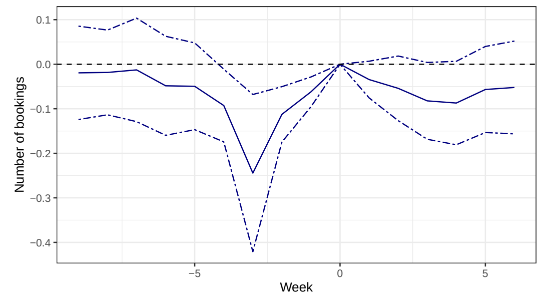

(b) Log(number of transactions)

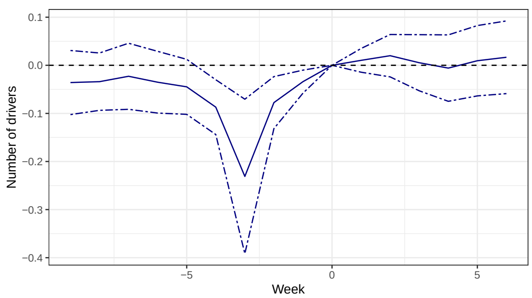

(c) Log(number of drivers)

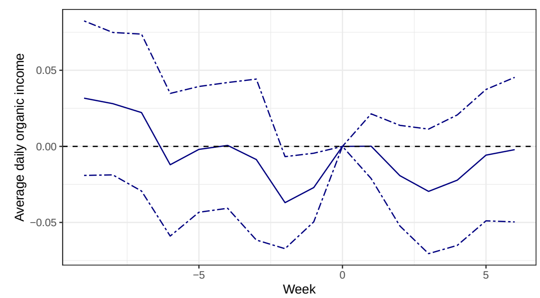

(d) log(driver average daily revenue)

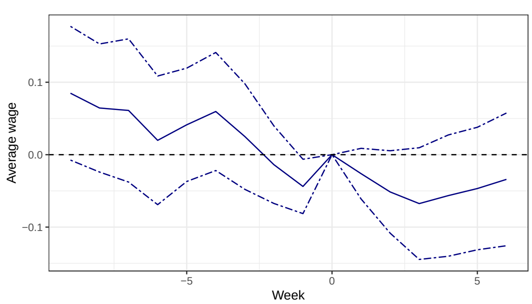

(e) log(average wage)

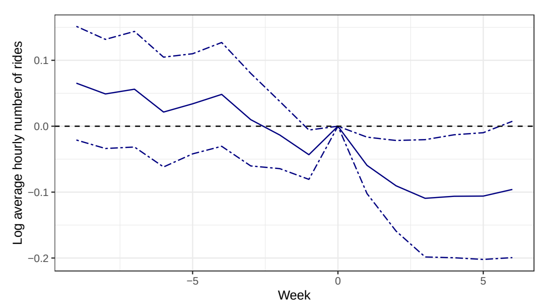

(f) log(average num. trips/supply hour)

Notes: All dependent variables are in log. Point estimates, in solid lines, are the estimated weekly average treatment effects $\hat{\beta}_1^k$, where k corresponds to the weeks since measure on the horizontal axis, as in equation 6. The 95% confidence intervals for coefficients are shown as range bound by the dashed lines. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and the errors clustered at the city level.

Figure 5: Conditional average treatment effects by drivers' total earning deciles

(a) Driver fare/trip

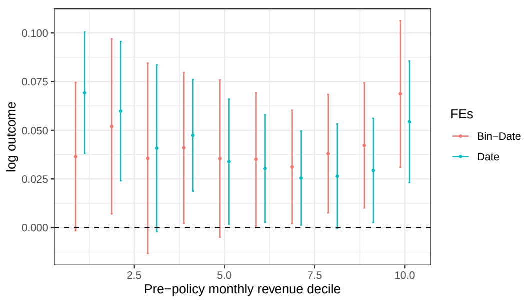

(b) Supply hours/decile

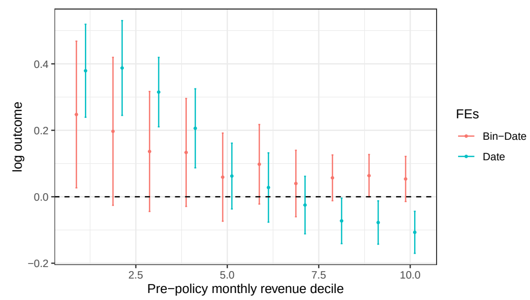

(c) N. distinct drivers/bin

(d) Supply hours/driver

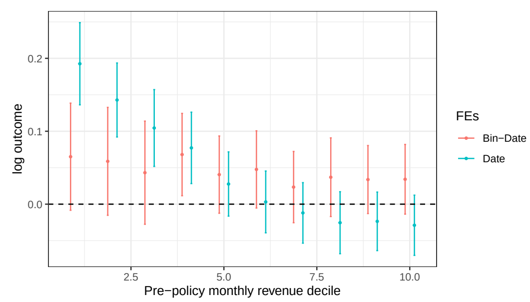

(e) Driver earnings/driver

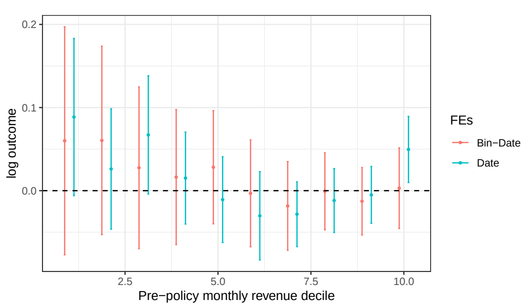

(f) Driver wage

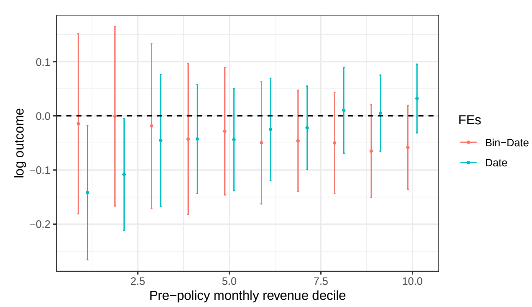

Notes: The dependent variables are the log-transformed outcomes listed in subfigure captions. The deciles are defined by the pre-policy transaction volume of drivers, defined as the sum of all driver fares over the pre-policy period of June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates in blue are the estimated impact on each bin of pre-policy average daily transaction. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and either fare-bins-by-day fixed effects or day fixed effects. Standard errors are clustered at the city-bin level. The sample is restricted to drivers who joined the platform by June 30, 2019.

Figure 6: Treatment effects by customers' pre-policy total expenditure deciles

(a) Number of trips/customer

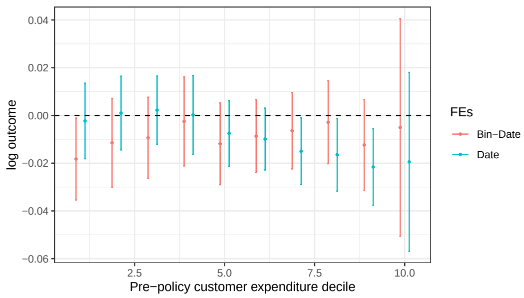

(b) Customer fare/trip

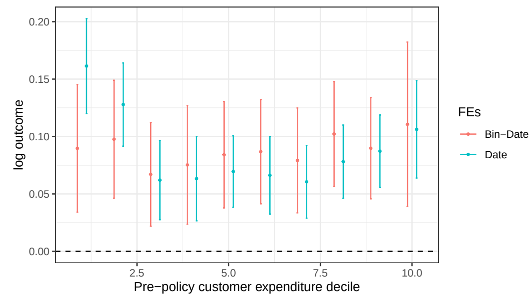

(c) Daily expenditure/customer

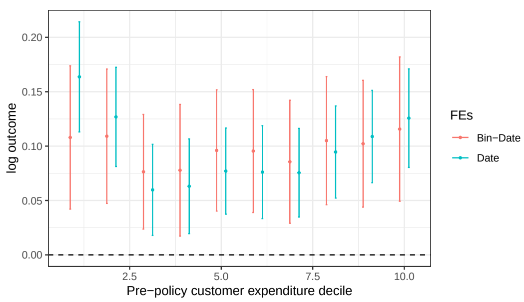

(d) Wait distance/trip

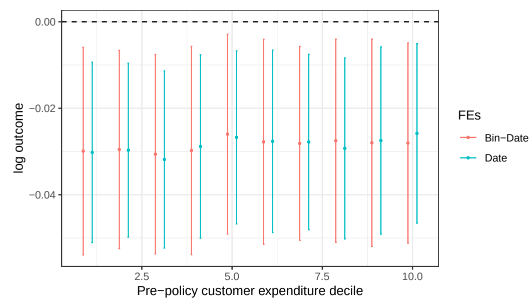

Notes: The dependent variables are the log-transformed outcomes listed in subfigure captions. The deciles are defined by the pre-policy customer expenditure, defined as the sum of all customer fare over the pre-policy period of June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates in blue are the estimated impact on each bin of pre-policy average daily transaction. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the city-bin level. The sample is restricted to customers who had made at least one order by June 30, 2019.

Figure 7: Conditional average treatment effects by customers' pre-policy exposure

(a) Number of trips/customer

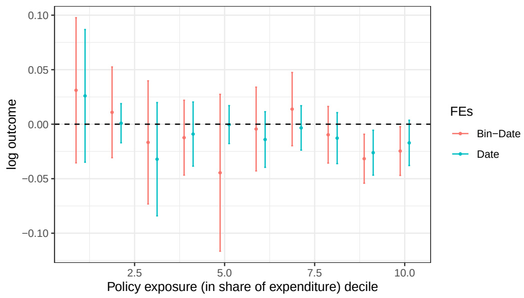

(b) Customer fare/trip

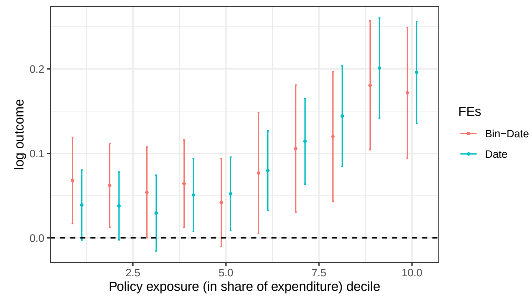

(c) Daily expenditure/customer

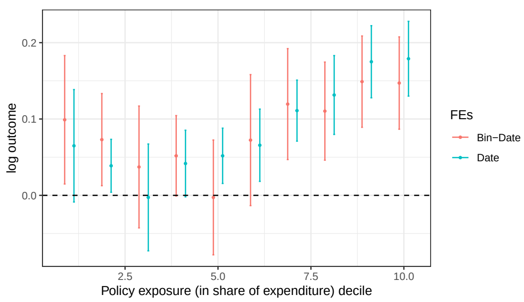

(d) Customer-average wait distance

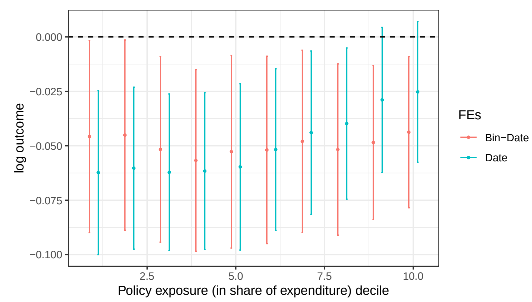

Notes: The dependent variables are the log-transformed outcomes listed in subfigure captions. The deciles are defined by the potential policy exposure of customers during the pre-policy period, defined as the share of the customer's monthly expenditure that would qualify for the policy. The pre-policy period is defined as June 1 to June 30, 2019. Point estimates in blue are the estimated impact on each bin of pre-policy average daily transaction. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the city-bin level. The sample is restricted to customers who had made at least one order by June 30, 2019.

Figure 8: Conditional average treatment effects on the number of distinct drivers, by drivers' pre-policy productivity

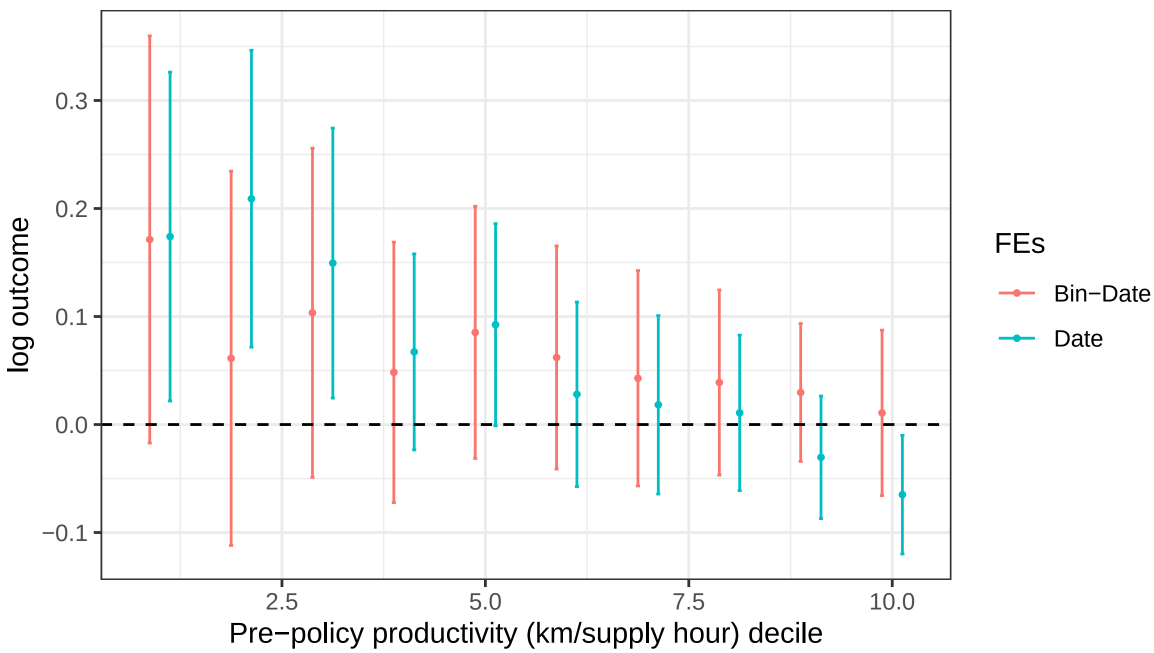

Notes: The dependent variable is the log-transformed number of distinct drivers per city per day. The deciles are defined by the pre-policy productivity level of drivers, defined as distance driven on trips per supply hour. Pre-policy period is June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates and the 95% confidence intervals are shown as dots and whiskers. The models include city-bin fixed effects, as well as either bin-date or date fixed effects. Standard errors are clustered at the city-bin level.

Figure 9: Conditional average treatment effects on productivity, by drivers' pre-policy productivity

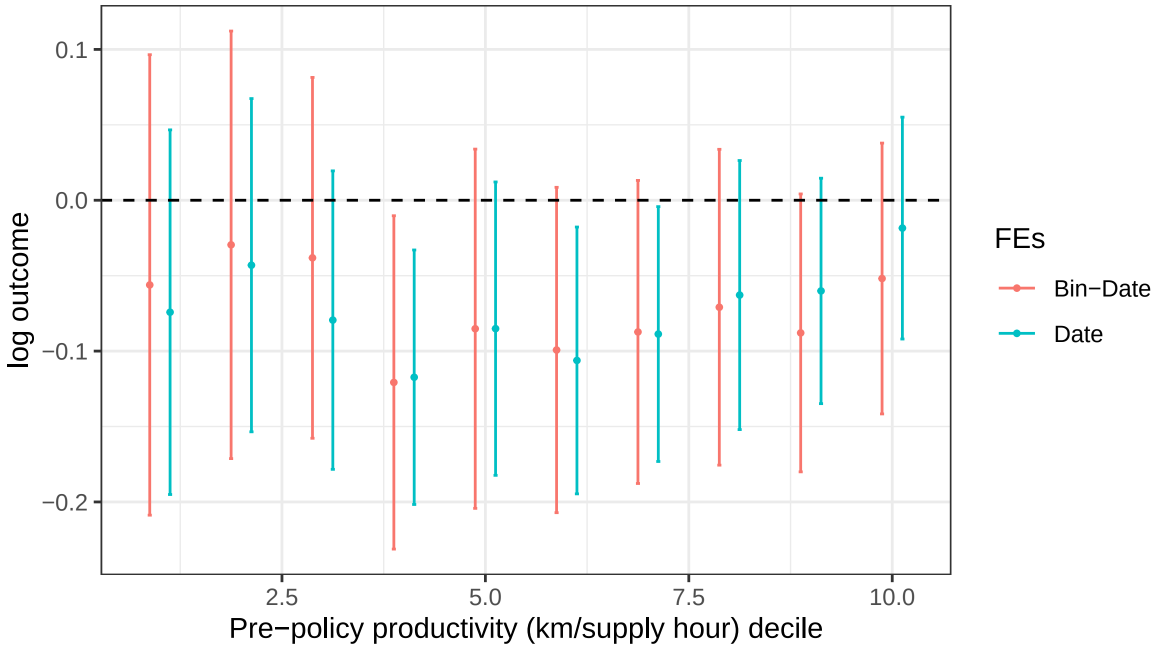

Notes: The dependent variable is the log-transformed average distance driven on trips per supply hour per city-bin-day. The deciles are defined by the pre-policy productivity level of drivers, defined as distance driven on trips per supply hour. Pre-policy period is June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates and the 95% confidence intervals are shown as dots and whiskers. The models include city-bin fixed effects, as well as either bin-date or date fixed effects. Standard errors are clustered at the city-bin level.

A Policy context

A.1 Fare regulation's implementation by city

Phase 1 (5 cities): Jakarta, Bandung, Yogyakarta, Surabaya, Makassar

Phase 2 (41 cities):

- Zone 1: Banda Aceh, Kota Medan, Kota Batam, Kota Pekanbaru, Kota Palembang, Kota Bandar Lampung, Kota Metro, Kota Belitung, Kota Bandung, Kota Semarang, Kota, Solo, Kota Yogyakarta, Kota Surabaya, Kota Denpasar, Kab. Probolinggo, Kab. Pasuruan, Kab. Kudus, Madura.

- Zone 2: Jakarta, Kab. Bogor, Kota Bogor, Kota Depok, Kab. Tangerang, Kota Tangerang, Kota Tangerang Selatan, Kab. Bekasi, Kota Bekasi.

- Zone 3: Kota Pontianak, Kota Palangkaraya, Kota Samarinda, Kota Balikpapan, Kota Banjarmasin, Kota Mataram, Kota Kupang, Kota Manado, Kota Gorontalo, Kota Palu, Kota Makassar, Kota Kendari, Kota Ambon, Kota Jayapura.

Phase 3 (88 cities/districts):

- Zone 1: Kota Sabang, Kota Bukittinggi, Kabupaten Agam, Kabupaten Lima Puluh Kota, Kabupaten Tanah Datar, Kota Padang Panjang, Kota Payakumbuh, Kota Duri, Kabupaten Bengkalis, Kota Tanjung Pinang, Kota Jambi, Kabupaten Muaro Jambi, Kabupaten Kisaran, Kabupaten Asahan, Kabupaten Karo, Kabupaten Toba Samosir, Kota Tanjung Balai, Kota Padangsidempuan, Kota Padang Lawas Utara, Kabupaten Tapanuli Selatan, Kabupaten Serdang Bedagai, Kota Pematangsiantar, Kabupaten Simalungun, Kota Tebing Tinggi, Kota Rantau Prapat, Kabupaten Labuhan Batu, Kabupaten Batang, Kabupaten Cilacap, Kabupaten Kebumen, Kabupaten Banyumas, Kabupaten Brebes, Kabupaten Purworejo, Kota Pekalongan, Kabupaten Pekalongan, Kabupaten Pemalang, Kabupaten Banjarnegara, Kabupaten Purbalingga, Kota Salatiga, Kabupaten Banyuwangi, Kabupaten Bojonegoro, Kabupaten Jember, Kabupaten Bondowoso, Kabupaten Jombang, Kabupaten Kediri, Kota Kediri, Kabupaten Nganjuk, Kota Madiun, Kabupaten Magetan, Kabupaten Ngawi, Kabupaten Ponorogo, Kota Mojokerto, Kabupaten Mojokerto, Kota Serang, Kabupaten Lebak, Kota Cirebon, Kabupaten Cirebon, Kabupaten Garut, Kabupaten Indramayu, Kabupaten Kuningan, Kabupaten Majalengka, Kota Tasikmalaya, Kabupaten Tasikmalaya, Kabupaten Subang, Kota Sukabumi, Kabupaten Sukabumi, Kabupaten Cianjur, Kabupaten Purwakarta, Kabupaten Sumedang, Kabupaten Ciamis, Kabupaten Pangandaran, Kota Banjar, Kota Malang, Kabupaten Malang, Kota Batu, Kota Tegal, Kabupaten Tegal, Kabupaten Demak, Kabupaten Kendal, Kabupaten Pati, Kabupaten Jepara

- Zone 3: Kota Bitung, Kota Tomohon, Kota Palopo, Kota Tarakan, Kota Ternate, Kota Sorong, Kabupaten Merauke, Kota Pare-Pare

Phase 4: all cities and districts where ride-share platforms operate.

B Additional results

B.1 Effects over time

One concern is that the exogenous price shocks may have differential short-term vs is the long-term. If it is so and, for instance, the causal effect of the policy diminishes over time, the estimates we present may not be relevant for the effect in equilibrium. Appendix Section 5.2.2 addresses this concern, by showing that the causal effects of policy are significant and persistent for the 6 weeks in which we have policy variation for the price and productivity outcomes. Furthermore, we argue in Section 1 that the estimate we identify *are* the equilibrium effects because of short contract duration and low cost of supplying labor on the platform once registered. Yet we cannot rule out that drivers and customers adjust their responses over time.

This question may be particularly relevant regarding the policy effect, which we show to be driven by both entry of lower productivity driver and reduction in individual productivity in Section 5.4. It is possible that less productive drivers, who now work more, may learn to improve their productivity as they work more, therefore diminishing the negative impact over time. To identify if such dynamic effects exist, we estimate Equation 6 from Appendix Section 5.2.2 separately for each decile of the pre-policy driver productivity.

Figure B.1 shows the result. We find that the negative productivity effect is noisily estimated and statistically indistinguishable from 0 for the first and second deciles. For higher deciles, however, we observe negative and persistent effects on productivity. If the aforementioned mechanism of productivity adjustment were to take place, then we would expect that estimates for lower deciles be positive in the later weeks post policy introduction. We do not observe such a trend.

Figure B.1: Weekly treatment effects on average driver productivity, by pre-policy productivity (km driven per supply hour) decile

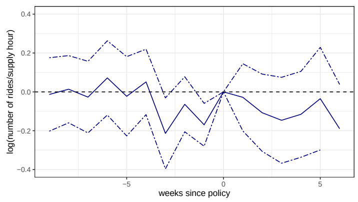

(a) 1st decile

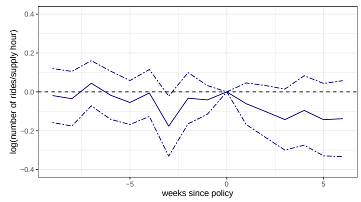

(b) 2nd decile

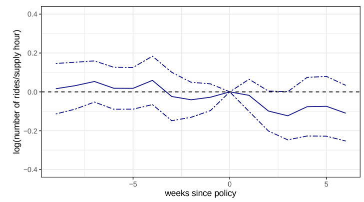

(c) 3rd decile

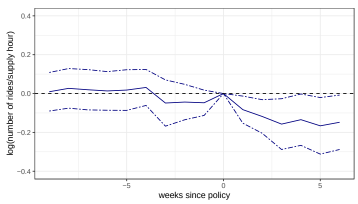

(d) 4th decile

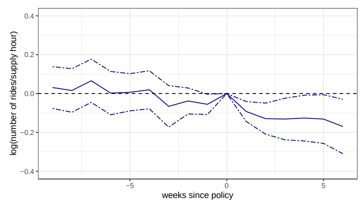

(e) 5th decile

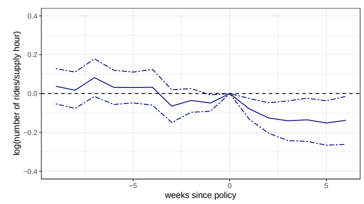

(f) 6th decile

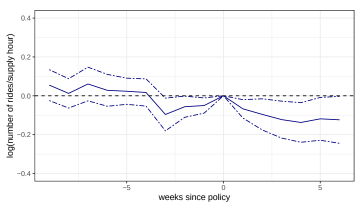

(g) 7th decile

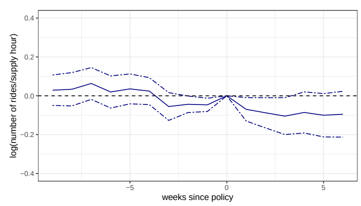

(h) 8th decile

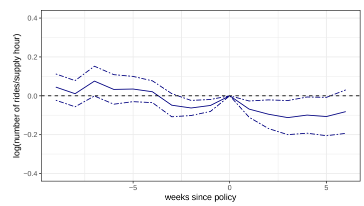

(i) 9th decile

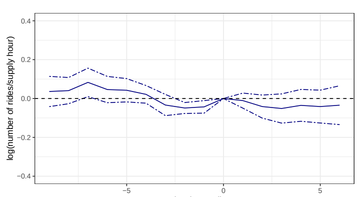

(j) 10th decile

Notes: All dependent variables are in log. Point estimates, in solid lines, are the estimated weekly average treatment effects $\hat{\beta}_1^k$, where k corresponds to the weeks since measure on the horizontal axis, as in equation 6. The 95% confidence intervals for coefficients are shown as range bound by the dashed lines. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and the errors clustered at the city level.

Figure B.2: Weekly treatment effects on the number of distinct drivers, by pre-policy productivity (km driven per supply hour) decile

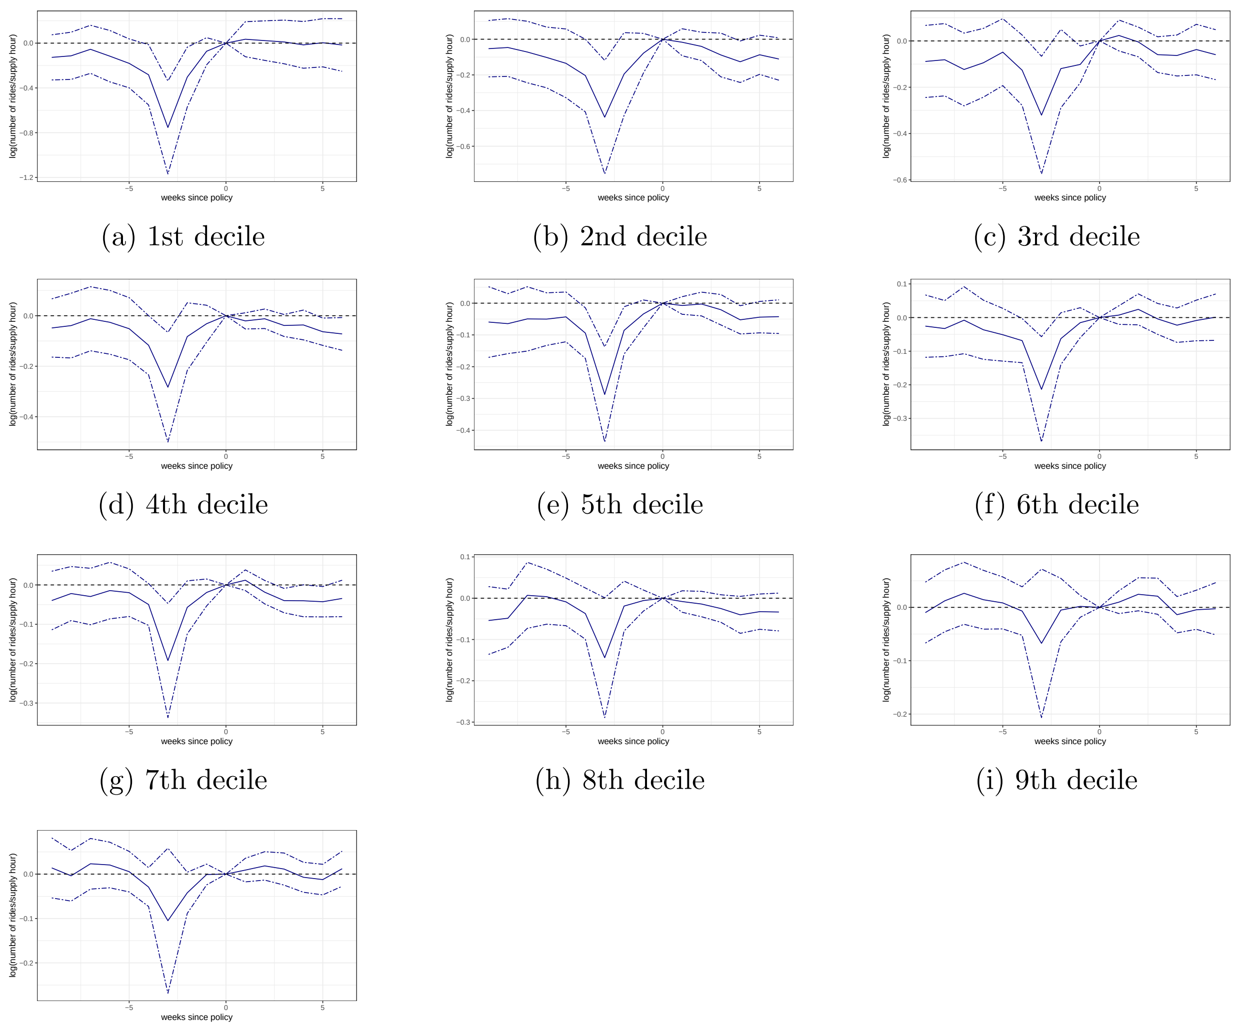

(a) 1st decile

(b) 2nd decile

(c) 3rd decile

(d) 4th decile

(e) 5th decile

(f) 6th decile

(g) 7th decile

(h) 8th decile

(i) 9th decile

(j) 10th decile

Notes: All dependent variables are in log. Point estimates, in solid lines, are the estimated weekly average treatment effects $\hat{\beta_1^k}$, where k corresponds to the weeks since measure on the horizontal axis, as in equation 6. The 95% confidence intervals for coefficients are shown as range bound by the dashed lines. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and the errors clustered at the city level.

B.2 Average treatment effects on trip attributes

Table B.1: Average treatment effects on driver fare and number of transactions

	log(Avg driver fare) (1)	log(Avg distance) (2)	log(Avg driver fare/km) (3)
Treat	0.0461**	0.0142	0.0393
	(0.0177)	(0.0165)	(0.0298)
Observations	12,760	12,760	12,760
R^2	0.93870	0.89941	0.56430
Within R^2	0.03673	0.00252	0.00124
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

B.3 Adjustment mechanism: Policy spillovers to unregulated segment

In Section 5.1.1, we show that the policy shifted the distribution of driver fare from below-threshold levels and higher fare toward the new threshold. This distributional shift underlines the effect on the total number of transactions, on which we estimated small yet noisy treatment effect in Section 5.1.

We find suggestive evidence that some of the null effects are driven by spillovers to unregulated, non-taxi service, including food and other delivery services. First, Tables B.2 and B.3 show that reduction in the number of taxi bookings is coincided with an increase in non-taxi counterparts, which are not regulated by the minimum fare policy. However, the estimates for th enon-taxi segments are noisily estimated and cannot be distinguished from zero.

Second, Figure B.3 shows that the changes in distribution of fares had different patterns for taxi and non-taxi segments. For non-taxi segments, the minimum fare policy does not reduce transaction quantities on lower price ranges, yet only increases it at the policy threshold that should only apply to taxis. Although the effect is attenuated by reduction in transactions at higher fares, the figure shows an increase in transaction for the non-taxi segment without any significant losses in price ranges below the minimum-fare threshold. This pattern is indicative of positive spillover effects in the unregulated segment, perhaps as a shift in transactions from the regulated counterpart. Although effects of the overall city-level effect on the number of transactions by service type are still imprecisely estimated, these point estimates and changes to the distribution of fares by service type suggest that certain types of taxi trips were replaced by comparable non-taxi ones.

The shift in transactions from taxi to non-taxi service types is corroborated by Figure B.4, which shows changes in distributions of transactions across trip distance, separated by service type. The figure shows that a reduction in a mass of short-distance taxi trips is more than made up by an increase in non-taxi trips of similar distance. This would be in line with the idea that taxi trips of certain attributes were switched by non-taxi of the same attributes, such as switching to food delivery from a restaurant instead of visiting it.

Table B.2: Average treatment effects on average price by service type

	log(Avg driver fare)		
	(1)	(2)	(3)
Service type	Taxi	Non-taxi	All
Treat	0.1286***	-0.0107	0.0461**
	(0.0322)	(0.0095)	(0.0177)
Observations	12,760	12,760	12,760
R^2	0.91247	0.93400	0.93870
Within R^2	0.14282	0.00177	0.03673
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table B.3: Average treatment effects on average price by service type

	log(N trips)		
	(1)	(2)	(3)
Service type	Taxi	Non-taxi	All
Treat	-0.0976	0.0516	0.0021
	(0.0914)	(0.0747)	(0.0829)
Observations	12,760	12,760	12,760
R^2	0.98331	0.97862	0.98193
Within R^2	0.00605	0.00213	3.5×10^{-6}
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Figure B.3: Treatment effect on frequency by 1,000-rupiah bins, by service type

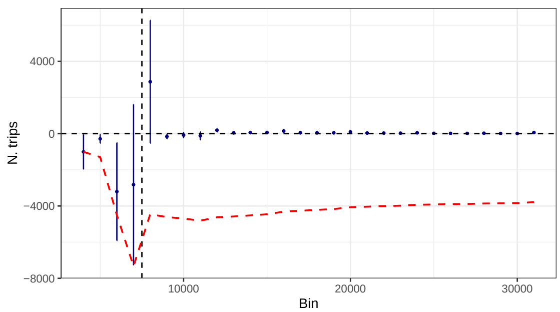

(a) Taxi

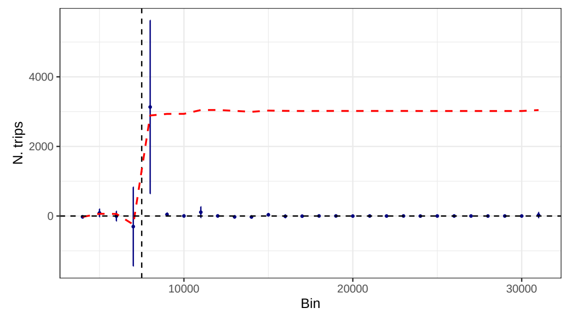

(b) Non-taxi

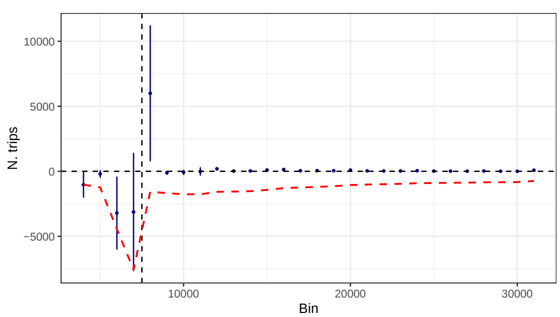

(c) All job types

Notes: The data is restricted to taxi trips and non-taxi trips for subpanels a and b, respectively. The result on subpanel C is for all job types, and is identical to Figure 2. The dependent variable is the number of transaction in 1000-Rupiah bins. Point estimates in blue are the estimated impact on each bin, $\hat{\beta}^j$, as in equation 2. The 95% confidence intervals for coefficients are shown by the range plots in blue. The dashed red line illustrates the cumulative effects, i.e., the sum of the point estimates. Regression is run on a panel of bins-city over day. The model includes fare-bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the product of zones and policy rollouts.

Figure B.4: Treatment effect on frequency by 500-meter distance bins, by service type

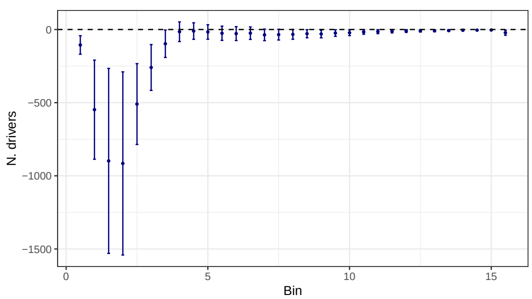

(a) Taxi

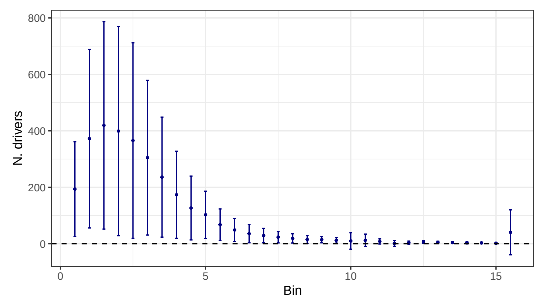

(b) Non-taxi

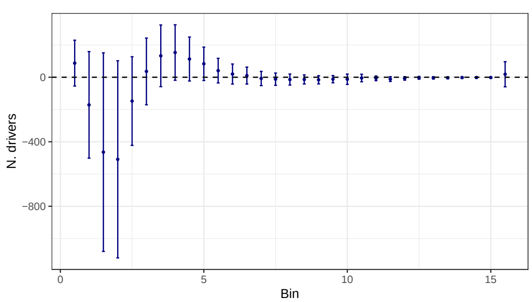

(c) All job types

B.3.1 New driver entries and permanent exits

In Section 5.2, we find a noisy yet positive impact of the policy on the number of workers per day. This effect may be driven by new driver entries to the platform, increased days worked of the existing drivers, or both. The extent of new worker entries also matters for the analysis on distributional impact and productivity in Sections 5.3 and 5.4, where we condition the data to drivers who worked in June, 2019.

We assess the magnitude of new worker entries and permanent exits by constructing a city-month level panel data of driver entries and exits. Driver entry and exit dates are defined as the minimum and maximum dates of the timestamps on their transactions. We define the entry month as the first month in which drivers had their first trip, and the exit month as the month in which they conducted their last trip, plus 1. We construct a data set covering February (the first month in which we can observe entry or exit) to July, 2019. We convert the entry and exit variables to the share of the number of drivers of a given city in June, 2019.

We find statistically insignificant effects on new driver entry or permanent exit of that the minimum-fare policy. Table B.4 shows the results. The effect sizes on new entry are 1.7 percentage points relative to the pre-policy fleet size (i.e., the number of distinct drivers in operation). The effect size on the log outcome indicates that this effect corresponds to a 22.1% increase, i.e., a large percentage increase off a small value. We refrain from interpretating the effect sizes of the exiters, as they have inconsistent signs between the share- and log-share outcomes.

B.4 Passthrough of driver-cut increases on customer fare

We investigate the extent to which the effects on transactions and spillovers are driven by the demand side response. First, we note that the exogenous increase in drivers' cut of the fare

Table B.4: Average treatment effects on driver entry and exit

	Share(entrants) (1)	Share(exiters) (2)	log(Share(entrants)) (3)	log(Share(exiters)) (4)
Treat	0.0172	0.0033	0.2213	-0.0871
	(0.0186)	(0.0070)	(0.1924)	(0.1117)
Observations	348	295	348	295
R^2	0.70367	0.67418	0.68840	0.83072
Within R^2	0.00463	0.00023	0.00777	0.00248
Month fixed effects	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓

Notes: Point estimates are of average effects. Standard errors are reported in parentheses. Regressions are run on a panel of cities over month as unit of time. All regressions include city fixed effects and month fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

leads to an increase in the amount billed to the consumer. This is not a given in ridesharing apps and other two-sided markets, where the platform operators may opt to absorb price shocks in one side (i.e. drivers) of the market so as not to affect the other (ii.e. customers). Table B.5 shows that the the policy increased both the drivers' cut as well as fare faced by the customers, and that these increases only occur in the taxi segment. In this appendix section we show that the pass-through of driver-cut increases on consumer price is greater than 100%, likely indicating complete pass-through plus administrative costs charged by the platform.

Table B.5: Average treatment effects on customer- and driver-fare

	log(Avg driver fare) (1)	log(Avg customer fare) (2)
Treat	0.0461**	0.0974**
	(0.0176)	(0.0447)
Observations	12,760	12,760
R^2	0.93644	0.82979
Within R^2	0.03547	0.03517
Week fixed effects	✓	✓
DoW fixed effects	✓	✓
City fixed effects	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

We estimate the passthrough rate using the two-stage least squares (2SLS) regressions. Using exogenous variation in the driver fare, we estimates the passthrough rate of changes in driver cut to customer fare, by instrumenting the driver fare ($ln(P_{c,t}^{driver})$) with the exogenous policy variation ($I_{c,t}(c \in Treat, t > 0)$) in the following equation:

$$ln(P_{c,t}^{customer}) = \beta_0 + \beta_1 * ln(P_{c,t}^{driver}) + \gamma_c + \rho_t + \epsilon_{c,t} \tag{13}$$

Table B.6 shows that the for taxi segment, a 10% increase in driver cut amount is associated with 18% increase in customer fare. The greater than 100% passthrough likely accounts for platform fees and/or additional discounts (in credits and free rides) that are given to customers. The passthrough rate is negative but imprecisely estimated for non-taxi segments.

Table B.6: **Elasticity of customer fare to driver cut**

	log(Avg customer fare)	
	All services	Regulated
	(1)	(2)
log(Avg driver fare)	2.111***	1.791***
	(0.7822)	(0.3512)
Observations	12,760	12,760
R^2	0.88062	0.90666
Within R^2	0.29631	0.28041
Day fixed effects	✓	✓
City fixed effects	✓	✓

Notes:

B.5 Elasticity estimates

We estimate the elasticities of key outcome variables to price via two-stage least squares (2SLS) regressions. We use the policy variation variable to instrument for the price. We use the relevant price measures depending on the outcome; for outcomes that we consider to be "demand side," we use the prices customers face. The "demand-side" outcome measure for our analysis is the number of trips. For the "supply-side" measures, i.e., driver supply hours, productivity, an average wage, we use the driver fare as price.

For the demand-side elasticity measure, we estimate the relationship between the number of trips and the price consumers face by instrumenting log average customer fare ($ln(P_{c,t}^{customer})$) with the exogenous policy variation ($I_{c,t}(c \in Treat, t > 0)$). The second-stage equation for the two-stage least-squares estimator is as follows, in which we regress the log of transaction volumes, $ln(Q_{c,t})$ on the instrumented average customer fare in log:

$$ln(Q_{c,t}) = \beta_0 + \beta_1 * ln(P_{c,t}^{customer}) + \gamma_c + \rho_t + \epsilon_{c,t} \tag{14}$$

Similarly, we estimate supply-side elasticity measures to the price they face, i.e., driver fare. We instrument log average driver fare ($ln(P_{c,t}^{driver})$) with the exogenous policy variation ($I_{c,t}(c \in Treat, t > 0)$) for the following equation, with the log of average supply hours, $ln(SupplyHr_{c,t})$, as the outcome variable:

$$ln(SupplyHr_{c,t}) = \beta_0 + \beta_1 * ln(P_{c,t}^{driver}) + \gamma_c + \rho_t + \epsilon_{c,t}. \tag{15}$$

B.6 Results from synthetic control-based inference

We find that the synthetic control-based inference procedure described in Section 4.2 yields similar results as the difference-in-differences counterpart. Table B.7 shows the results on our six main logged outcomes: average driver fare, number of trips, number of drivers, total supply hour, average driver daily earnings, and productivity (average distance driven on trips/supply hour). The dataset is the city-day panel, and the estimates are comparable to Tables 3, 4, 5, 6, and 9. The second and third columns show the pre- and post-policy differences between the weighted treatment units and their synthetic controls. The fourth column shows the p-values based on RSMRE rankings. The differences between the values on the third and second columns can be interpreted as the difference-in-differences estimate between the treatment and placebos, and the fourth column has its p-value.

Table B.7 shows results that are consistent with our difference-in-difference estimates. We find a 7.3% increase in average driver fare but do not find statistically significant effects on the number of trips, drivers, and their driver earnings. We find an 11.3% reduction in driver productivity. One notable difference, however, is that the effects on total driver supply hours is insignificant for the synthetic control-based analysis, with a treatment effect of around 1.0%.

Figure B.5 shows the time-series plots of weighted treatment effects and their placebos from Day 101 (April 11, 2019) to Day 216 (August 8, 2019). We find that, as we see in Table B.7, the differences between weighted treated outcomes and their synthetic controls (i.e., the pre-policy values of the red line) are greater than zero. This may be due to the weekly cyclicality or other daily variations in the outcomes data, where the optimally weighted synthetic control was biased. As such, it seems from the figure that the causal estimate should be the differences in deviations between weighted treated units and the placebos post- and pre-policy.

Figure B.5 also confirms the instantaneous and persistent effects on average driver fare, as shown in the raw data in Figure 1 and in distributed lag DiD model in Figure 4. We also find that for outcomes for which the treatment estimates are less statistically significant, the differences between pre- and post-treatment levels are less stark than for statistically significant estimates.

Table B.7: **Average treatment effects on driver fare and number of transaction**

	Pre-policy difference	Post-policy difference	P-value (RMSRE-based)
log(Avg driver fare)	0.022	0.095	0.016
log(N trips)	0.149	0.043	0.820
log(N drivers)	0.118	0.079	0.559
log(Total supply hours)	0.132	0.142	0.245
log(Avg daily income)	0.047	0.019	0.918
log(Avg km/supply hr)	0.018	-0.095	0.040

Notes: All dependent variables are in log. Outcome variables are listed on the first column, with corresponding estimates and p-values on the second to fourth columns. The second and third columns report the pre- and post-policy differences between the weighted treatment effects and their synthetic controls. The fourth column shows the p-values based on RSMRE rankings of the treatment effect, relative to their "placebo" counterparts.

Figure B.5: **Synthetic control-based estimates and placebo simulations**

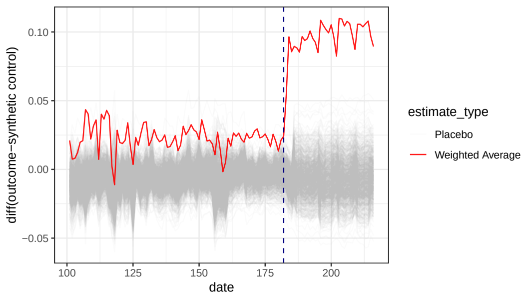

(a) log(Avg. driver fare)

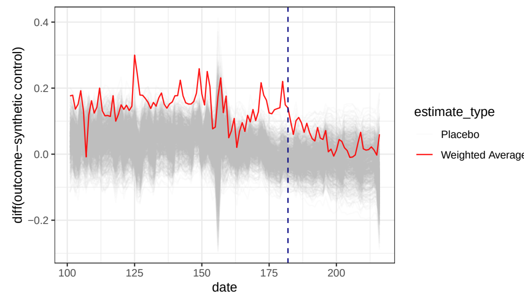

(b) log(N. trips)

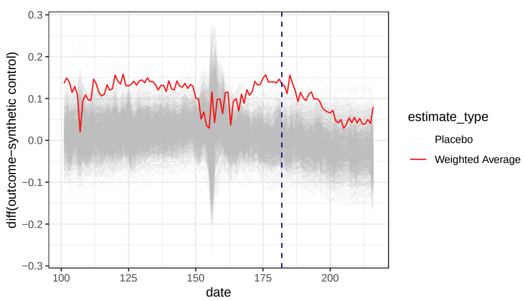

(c) log(N. drivers)

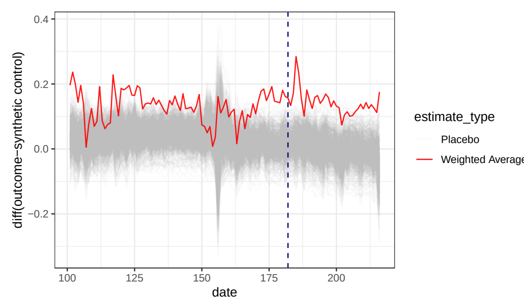

(d) log(Total supply hours)

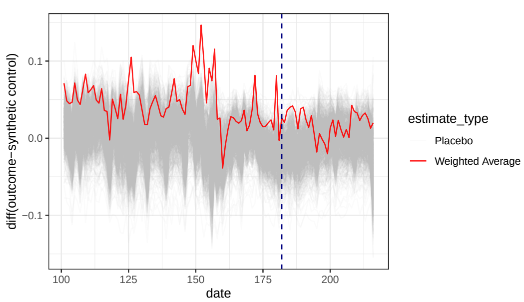

(e) log(Avg. daily income)

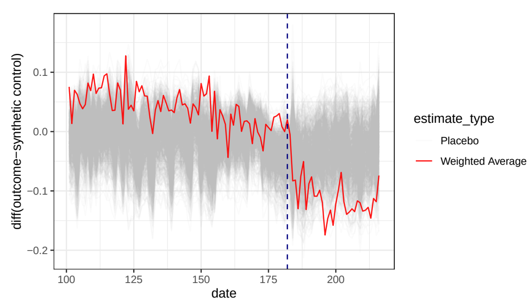

(f) log(Avg. n rides/supply hr)

C Appendix Tables

Table C.8: Average treatment effects on customer fare and total spending

	log(Avg customer fare)		log(Sum customer fare)	
	All services	Regulated	All services	Regulated
	(1)	(2)	(3)	(4)
Treat	0.0974**	0.2303***	0.0995*	0.1327**
	(0.0450)	(0.0476)	(0.0544)	(0.0585)
Observations	12,760	12,760	12,760	12,760
R^2	0.83656	0.89504	0.98499	0.98572
Within R^2	0.03658	0.19079	0.00926	0.01437
Day fixed effects	✓	✓	✓	✓
City fixed effects	✓	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table C.9: Robustness check: Average treatment effects on the number of transactions, excluding the Eid al-Fitr period

	log(N trips)		
	Full sample	1 week	2 weeks
	(1)	(2)	(3)
Treat	0.0021	-0.0044	-0.0079
	(0.0829)	(0.0847)	(0.0867)
Observations	12,760	12,354	11,948
R^2	0.98193	0.98332	0.98361
Within R^2	3.5×10^{-6}	1.66×10^{-5}	5.53×10^{-5}
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Labels above the column number refer to the sample restriction. "Full sample" uses all data from January 1 to August 8, 2019. Column with the "1 week" label excludes data from June 2 (one day before the start of Eid al-Ftr) to June 8, 2019. Column with the "2 weeks" label excludes data from June 2 to June 15, 2019. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table C.10: Robustness check: Average treatment effects on the number of drivers, excluding the Eid al-Fitr period

	log(N drivers)		
	Full sample	1 week	2 weeks
	(1)	(2)	(3)
Treat	0.0649	0.0582	0.0562
	(0.0605)	(0.0615)	(0.0628)
Observations	12,760	12,354	11,948
R^2	0.98460	0.98603	0.98627
Within R^2	0.00494	0.00454	0.00441
Day fixed effects	✓	✓	✓
City fixed effects	✓	✓	✓

Notes: All dependent variables are in log. Point estimates are of average treatment effects, $\hat{\beta}_1$, as in equation 1. Labels above the column number refer to the sample restriction. "Full sample" uses all data from January 1 to August 8, 2019. Column with the "1 week" label excludes data from June 2 (one day before the start of Eid al-Ftr) to June 8, 2019. Column with the "2 weeks" label excludes data from June 2 to June 15, 2019. Standard errors are reported in parentheses. Regressions are run on a panel of cities over day as unit of time. All regressions include city fixed effects and time fixed effects, and standard errors clustered at the city level. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

D Appendix Figures

Figure D.6: Treatment effects on percentiles of the driver fare/trip distance by trip distance bin

(a) 5th percentile

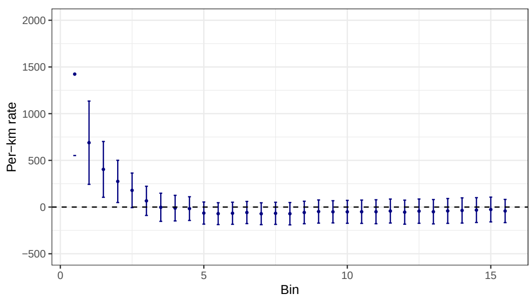

(b) 25th percentile

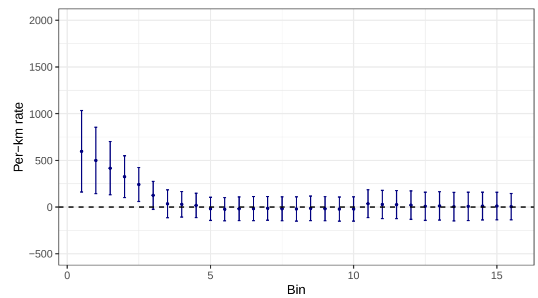

(c) 75th percentile

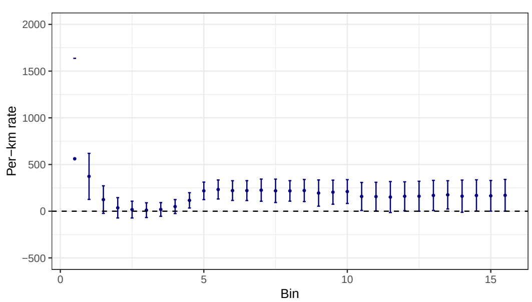

(d) 95th percentile

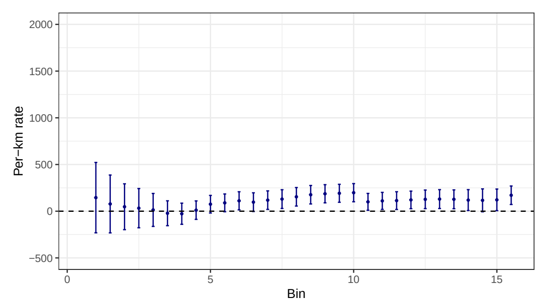

Notes: The dependent variables are the percentile measures of total driver fare divided by the distance driven in kilometers. Point estimates in blue are the estimated impact for each 500-meter trip distance bin. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and fare-bins-by-day fixed effects. Standard errors are clustered at the city-bin level..

Figure D.7: Conditional average treatment effects by drivers' pre-policy exposure

(a) Driver fare/trip

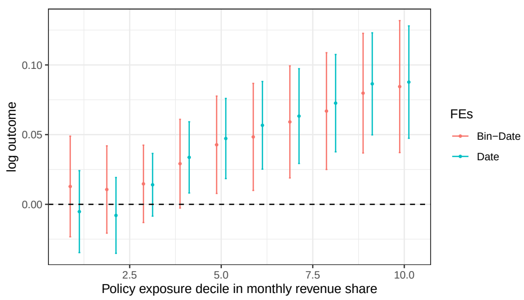

(b) Supply hours/decile

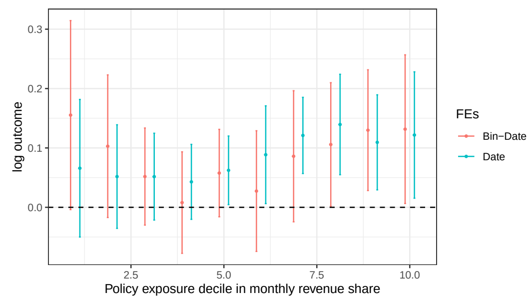

(c) N. distinct drivers/decile

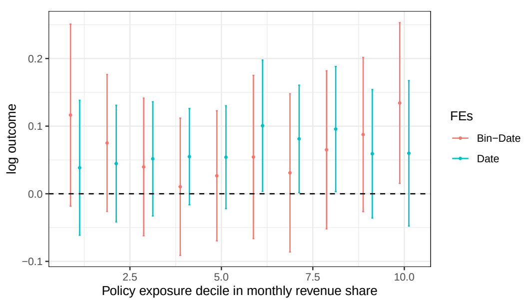

(d) Supply hours/driver

(e) Driver earnings/driver

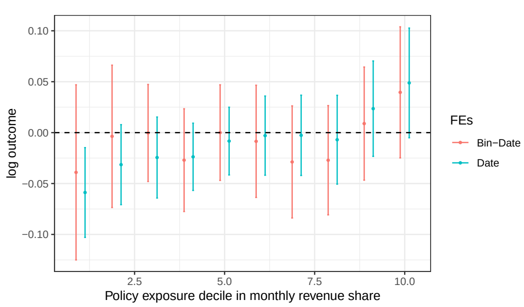

(f) Driver wage

Notes: The dependent variables are the log-transformed outcomes listed in subfigure captions. The deciles are defined by the potential policy exposure of drivers during the pre-policy period, defined as the share of the driver's monthly earnings that would qualify for the policy. The pre-policy period is defined as June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates in blue are the estimated impact on each bin of pre-policy average daily transaction. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and either fare-bins-by-day fixed effects or day fixed effects. Standard errors are clustered at the city-bin level. The sample is restricted to drivers who joined the platform by June 30, 2019.

Figure D.8: Conditional average treatment effects by drivers' pre-policy wage

(a) Driver fare/trip

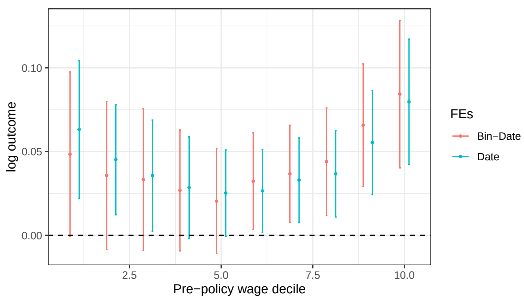

(b) Supply hours/decile

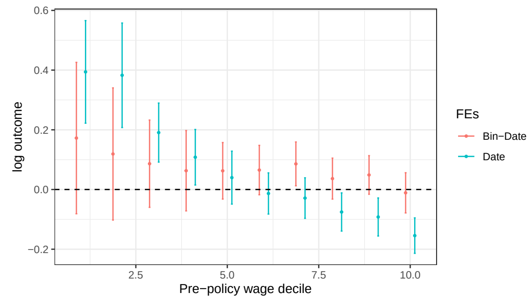

(c) N. distinct drivers/decile

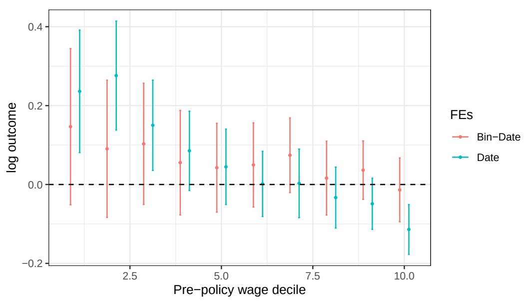

(d) Supply hours/driver

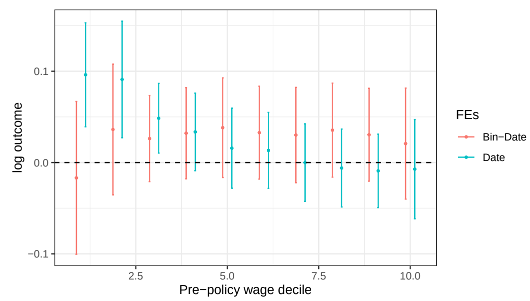

(e) Driver earnings/driver

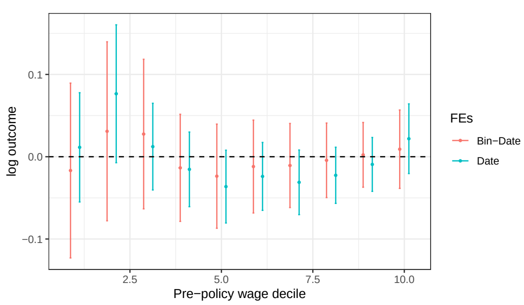

(f) Driver wage

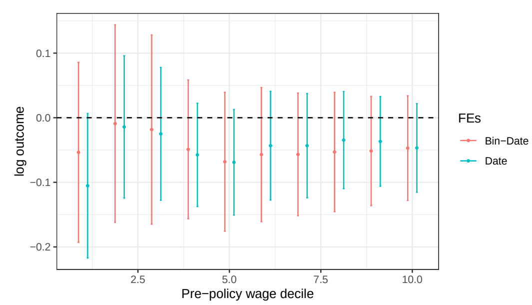

Notes: The dependent variables are the log-transformed outcomes listed in subfigure captions. The deciles are defined by the potential policy exposure of drivers during the pre-policy period, defined as the share of the driver's monthly earnings that would qualify for the policy. The pre-policy period is defined as June 1 to June 30, 2019. Data are restricted to transactions from drivers who were present in the data as of June 30, 2019. Point estimates in blue are the estimated impact on each bin of pre-policy average daily transaction. The 95% confidence intervals for coefficients are shown as the whisker bar in blue. The model includes bin-by-city fixed effects and either fare-bins-by-day fixed effects or day fixed effects. Standard errors are clustered at the city-bin level. The sample is restricted to drivers who joined the platform by June 30, 2019.