

NBER WORKING PAPER SERIES

RIDE-SHARING MARKETS RE-EQUILIBRATE

Jonathan V. Hall
John J. Horton
Daniel T. Knoepfle

Working Paper 30883
<http://www.nber.org/papers/w30883>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2023

This manuscript was not subject to prior review by any party, as per the research contract signed at the outset of this project. The views expressed here are solely those of the authors. Both Hall and Knoepfle are employees and shareholders of Uber Technologies. Robin Yerkes Horton was an employee and shareholder of Uber Technologies and is John Horton's spouse. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w30883>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Ride-Sharing Markets Re-Equilibrate
Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle
NBER Working Paper No. 30883
February 2023
JEL No. J01,R4,R41

ABSTRACT

Following Uber-initiated fare increases, drivers make more money per trip and, initially, more per hour-worked. Drivers begin to work more hours. However, this increase in hours-worked—combined with a reduction in demand from a higher fare—has a business stealing effect, with drivers spending a smaller fraction of working hours transporting passengers. This market adjustment brings the hourly earnings rate back to about the rate that prevailed before the fare increase, in roughly two months. Passengers are partially compensated for higher prices by shorter wait times, but during the period covered by our data, fare increases likely reduced passenger welfare.

Jonathan V. Hall
Uber Technologies
jvh@uber.com

Daniel T. Knoepfle
Uber Technologies
knoepfle@gmail.com

John J. Horton
Massachusetts Institute of Technology
Sloan School of Management
100 Main St
Cambridge, MA 02142
and NBER
john.joseph.horton@gmail.com

Updated versions will be here: https://www.john-joseph-horton.com/papers/uber_price.pdf

Ride-Sharing Markets Re-Equilibrate*

Jonathan V. Hall
Uber Technologies

John J. Horton
MIT Sloan & NBER

Daniel T. Knoepfle
Uber Technologies

January 18, 2023

Abstract

Following Uber-initiated fare increases, drivers make more money per trip and, initially, more per hour-worked. Drivers begin to work more hours. However, this increase in hours-worked—combined with a reduction in demand from a higher fare—has a business stealing effect, with drivers spending a smaller fraction of working hours transporting passengers. This market adjustment brings the hourly earnings rate back to about the rate that prevailed before the fare increase, in roughly two months. Passengers are partially compensated for higher prices by shorter wait times, but during the period covered by our data, fare increases likely reduced passenger welfare.

1 Introduction

In many platform markets the price faced by buyers is ostensibly set by the platform. The platform still typically allows free entry of sellers, who in turn

*This manuscript was not subject to prior review by any party, as per the research contract signed at the outset of this project. The views expressed here are solely those of the authors. The code and data to replicate this study are or will be publicly available at http://www.john-joseph-horton.com/papers/uber_price.data. Both Hall and Knoepfle are employees and shareholders of Uber Technologies. Robin Yerkes Horton was an employee and shareholder of Uber Technologies and is John Horton's spouse.

earn a fraction of receipts from the buyers they serve. With this hybrid structure, the platform’s choice of a price affects both sides of the market, but—on the face of it—incompatible directions: A higher price lowers demand and simultaneously increases supply, pushing the market out of equilibrium. Our research question is a simple one—following platform price changes, how do these markets clear? And given how they clear, what are the implications for the functioning and efficiency of these markets?

Our empirical context is a collection of Uber-created ride-sharing marketplaces in the US. These marketplaces have experienced numerous city-specific, Uber-initiated changes to the time and distance “multipliers” that determine the price of a trip under typical conditions. The actual price passengers faced in the market also depends on Uber’s use of dynamic or “surge” pricing (Chen and Sheldon, 2015; Hall et al., 2016). We abstract away from the specifics of Uber’s taxi-like pricing of trips by constructing a price index, which is the price for a typical trip during un-surged conditions for each city-week.

We use changes in this price index to identify the effects of the fare on the market equilibrium. As we will discuss at length, the assumptions required for causal inference in this setting are well-satisfied, as Uber’s base fare pricing decisions seem to be conditioned on market attributes that we observe. Furthermore, we can account for other potential factors, such as local economic conditions, weather, and even the market share of competing ride-sharing services.

We find that when Uber raises the base fare, passengers face higher prices. This is not mechanical, as changes in surge pricing could fully “undo” the fare changes. However, surge does play a buffering role: with a higher base price, demand outstrips supply less often, and so the platform does not need to use surge pricing quite as much to clear the market. A 10% fare increase causes the average surge rate to fall by about 2%.

On the driver side, with a higher base fare, the driver’s hourly earnings rate rises immediately as drivers make more money per trip. However, the hourly earnings rate begins to decline shortly thereafter. After about 8 weeks, there is no clear difference in the driver’s gross average hourly earnings rate

compared to before the fare increase.

The main reason for this short-run/long-run difference is that the driver utilization—or the fraction of an hour-worked spent serving passengers—falls substantially with a higher fare: a 10% fare increase lowers utilization by about 7%. Combined with the decline in the surge multiplier, this fall in utilization is offsetting from the driver’s perspective with respect to the hourly earnings rate: a 10% increase in fare raises driver hourly earnings by 0.7%, with a 95% CI that includes 0. However, consistent with hourly earnings increasing by some amount, we do find that with a higher fare, driver hours-worked increases, both on the extensive and intensive margins.

On the passenger side, we find that with higher prices, wait times fell considerably. This reflects the fact that with lower utilization, all else equal, the nearest car available for dispatch is closer. A 10% increase in the base fare reduced wait times by about 6%. However, this quality improvement was not enough to offset reductions in demand from higher prices: With a higher base fare, the overall hours of transportation fell, as did the number of completed trips.

As fare changes have real and persistent effects on the market, it is clear, empirically, that ride-sharing markets have multiple equilibria, and that the platform’s pricing choices are consequential. To explore the welfare implications of these movements to different equilibria, we develop a simple model of a ride-sharing market. In the model, there is an equilibrium trade-off between passenger prices and driver utilization. Drivers generally prefer higher-price/lower-utilization equilibria; passengers prefer the reverse. But from an equilibrium with a sufficiently high fare, both drivers and passengers want fare cuts; and at sufficiently low fare equilibrium, both drivers and passengers want fare increases. However, despite the possibility of aligned interests with respect to fare changes, our empirical results indicate fare changes were not made in the Pareto improving range, and that fare increases made drivers better off and passengers worse off.

The main contribution of this paper is in offering a high-level description of how ride-sharing markets function, and the role of platform pricing in deter-

mining the equilibrium. Although our context is ride-sharing, other markets likely have similar economics, even if there is no centralized platform setting prices, so long as the supply side is endogenously “busy,” there is more or less free entry, and some market-clearing happens through non-price margins.

Our results are qualitatively similar to [Hsieh and Moretti \(2003\)](#), who show that real estate agent earnings are not affected by house prices, despite agents being paid fixed, proportional commissions (in the long run of 10 years). As in our paper, the reasons for little pass-through of product market price to wages are low entry barriers and business stealing.¹ However, our paper highlights the market usefulness of un-utilized time ([Hall, 1983](#)), as this time allows for higher service quality in the form of shorter wait times. The productivity limit in this market—100% utilized and hence fully “productive” drivers—is actually a marketplace disaster ([Castillo et al., 2017](#)). This price/quality trade-off has not been emphasized in the literature, but we find it is practically important to a market-designing platform.

The re-equilibration process we illustrate is fairly straightforward. When driving with Uber suddenly becomes a better deal, drivers work more hours and so more drivers are chasing fewer potential trips. This lowers driver utilization and hence the driver hourly earnings rate. The equilibrium decline in passenger demand from higher prices is, however, offset somewhat by the improved wait times enabled by lower utilization. The process runs in reverse when driving with Uber becomes a temporary worse deal through fare cuts.

The re-equilibration process can be directly observed because of the computer-mediated nature of the market. This adjustment process highlights a point often emphasized by economists but rarely seen so clearly—namely that the immediate and direct effect of some policy change or shock can be quite different from the ultimate effect following market adjustment. A fare increase initially has the anticipated direct effects—drivers initially make more money and passengers pay more for essentially the same service. But over time, as

¹The market equilibrium that arises bears similarities to [Harris and Todaro \(1970\)](#) who argue that rural-to-urban migration in developing countries tends to equalize the expected urban income and the expected rural income, despite higher urban wages.

drivers and passengers make different decisions in response to the new incentives, a new equilibrium was reached—and that equilibrium was quite different from what prevailed just weeks earlier.

The rest of the paper is organized as follows. We describe the empirical context in §2. Our panel data and the variation in prices are described in §3. We then develop a model of a ride-sharing market in §4. We then estimate the effects of fare changes on various market outcomes in §5. The threats to valid causal inference are discussed in §6. We conclude in §7.

2 Empirical context

Uber connects passengers with drivers-for-hire in real-time, creating a collection of city-specific, geographically-isolated markets. It currently operates in more than 340 cities, in over 60 countries. The core rides products of Uber are UberBlack and UberX. See [Hall and Krueger \(2018\)](#) for a discussion of the relative size of the two services. We focus exclusively on UberX in this paper.

Regardless of the product, passengers use the Uber app to set their location and request a ride. These trip requests were originally sent to the nearest available driver.² At the end of the trip, the fare is automatically charged to the passenger’s credit card. Uber handles all billing, customer support, and marketing.

2.1 The price of a trip

The price of a trip depends on a number of parameters set by Uber.³ There is a per-minute time multiplier and per-mile distance multiplier, as well as a fixed

²Uber’s matching has since evolved away from myopic trip-by-trip matching to better optimize for overall network efficiency; however, the quality of a ride request-available car match remains strongly influenced by distance.

³Near the end of our panel, Uber began using “up-front” pricing in which passengers are quoted a fare at the start of the trip, based on the expected values for the distance and duration, given the user-provided trip start and endpoints. The identifying variation in the base fare comes before up-front pricing was widely implemented. Furthermore, early versions of this pricing simply replaced expected values with realized values, hence not appreciably changing the price level.

initial charge, and service fees in some markets. To calculate the actual fare paid by the passenger, the parameters are multiplied by the realized time and distance of a trip, which is then multiplied by the surge multiplier that was in effect when the trip was taken. The surge multiplier is set algorithmically in response to supply and demand imbalances. During “un-surged” periods, the multiplier is 1.0.⁴ There is a minimum charge that applies if the calculated fare is below that minimum.

As we will see, Uber has changed the time and distance multipliers for UberX in every city in our data. When Uber has made a change in a given city, it has typically changed the time and distance multipliers by the same percentage. To avoid the complexity of tracking different fare components separately, we construct price indices. For a given service, city and week, the index is the total fare for an un-surged 6-mile, 16-minute trip. This trip is approximately the median trip time and distance for the US.

2.2 Measurement of hours-worked

We define driver hours-worked as the total time a driver spent “online” with the Uber platform, which includes all time on-trip, en-route to pick up a passenger, or simply being available to receive dispatch requests. Merely having the app open without marking oneself available to receive dispatch requests does not count in our measure of hours-worked. Because of the computer-mediated nature of the market, this hours-worked quantity (as well as time on the trip) is measured essentially without error, aside from rare technical glitches (Varian, 2010).

Our definition of hours-worked does not perfectly capture what one might regard as working, or what is commonly reported as work hours in government statistics. Being available to receive ride requests does not, in itself, exclude other time uses. For example, drivers may mark themselves available for dispatch while performing personal errands or “commuting” to where they normally seek passengers (such as the airport or central business district

⁴Cohen et al. (2016) uses variation in surge pricing to estimate the elasticity of demand for UberX at several points along the demand curve.

in a city), inflating our hours-worked measure. Similarly, drivers may mark themselves available for ride requests while performing other flexible work.

Drivers also report driving with multiple ride-sharing platforms simultaneously, going offline on the Uber app only after being dispatched by another platform, or in some cases keeping both apps on and simply turning down dispatches as needed.⁵

Although these definitional ambiguities require us to be careful in interpretation, they would mainly create complications if our interest was in “levels” rather than in “changes” and our interest is primarily in changes to the market equilibrium.⁶

2.3 Measurement of driver hourly earnings rates

To construct a measure of the driver hourly earnings rate for a city in a given week, we divide the total weekly driver revenue by the total hours-worked. This method is equivalent to averaging all driver-specific estimates of the hourly earnings rate and weighting by individual hours-worked. For the driver gross hourly earnings rate, we omit reimbursements for known tolls and fees (such as airport fees) and deduct Uber’s service fee.

Drivers are eligible for promotional payments that typically depend on meeting various goals, such as the number of rides provided in a week. When we explore the effects of promotional payments, we allocate the payments as earnings in the week in which they were paid. Some promotional payments

⁵Such “multi-homing” behaviors lead to double-counting of some hours-worked (while the driver is waiting for dispatch) but under-counting others due to time spent on trips for the other platform. However, under reasonable assumptions (e.g., Poisson arrivals of trips), the measured utilization of a multi-homing driver on one platform is the same as the utilization as a non-multi-homing driver. Thanks to Jason Dowlatabadi for helping us see this point. Despite the possibility that competitor platforms matter, we find no evidence that the share of direct Uber competitors—and thus, presumably, the opportunity for multi-homing—affects our results. See §B.4. For a theoretical analysis of the effects of multi-homing in ride-sharing markets, see [Bryan and Gans \(2019\)](#).

⁶For example, there is a legal debate on whether hours spent preparing to work—such as commuting and putting on work clothes—are compensable. See “Fact Sheet 22: Hours Worked Under the Fair Labor Standards Act“ <https://www.dol.gov/whd/regs/compliance/whdfs22.pdf> by the US Wage and Hour Division of the Department of Labor.

unrelated to driving, like those earned for referring another driver, are omitted.

2.4 Accounting for driver costs

The gross hourly earnings rate measure does not include the costs to drivers, such as fuel, wear and tear on the vehicle, and other consumables. As many of these costs depend on the number of miles driven, average costs likely change with the utilization, as unutilized drivers waiting for dispatch can reduce expenses by driving more slowly or even stopping completely. Although we lack speed data for the full panel, we do have city-specific average speeds of drivers for July 2017, conditioned on whether or not the driver was with a passenger.

As expected, the average speed is substantially lower when the driver is without passengers. The average speed difference is about 5.4 MPH, or a 30% difference from the baseline speed. We do not know on a driver-by-driver basis how much this speed reduction lowers costs, as the reduction depends on the driver’s vehicle. However, we can make some assumptions to construct an estimate of average costs based on average utilization.

Suppose drivers have an average speed of s^{PAX} when active and s^{NoPAX} when inactive i.e., without passengers. If the utilization in a city is x , in a given hour of work, a driver drives $xs^{\text{PAX}} + (1 - x)s^{\text{NoPAX}}$ miles. The cost-per-hour in city i is then just this average miles per hour times the imputed cost per mile. We use the rate of \$0.30/per mile, the median estimate from [Zoepf et al. \(2018\)](#). This multiplier is intended to capture the full direct costs of a mile driven, but not the costs of effort. Using this rate and the city-specific speed data, we can calculate measures of *net* hourly earnings. For the inactive speed, we apply the 30% adjustment to the active speed that week (assuming the July 2017 measured difference applies to all periods).

3 Data

Our panel consists of 36 US cities over 138 weeks, beginning with the week of 2014-06-02 and ending with the week of 2017-01-16. All cities in the panel

have an UberX service, though only some have an UberBlack service. To construct our panel, we started with the 50 largest US cities by total trip volume at the end of the panel. From this panel, we then removed cities that had substantial changes to the areas of service availability or significant within-city geographical variation in pricing.⁷ These cities include Boulder, Denver, Indianapolis, Las Vegas, Philadelphia, Austin, Portland, Palm Springs, San Antonio, Ventura, New Orleans, and Miami, and the “cities” of Connecticut, New Jersey, and Greater Maryland, which were managed as cities in Uber’s system but did not functionally represent single markets.

The panel is slightly unbalanced in that we lack early data for Charleston and Richmond (20 total missing weeks) which had relatively late introductions of UberX. The panel begins with the week in which driver earnings data is first reliably available; prior to 2014-06-02, historical driver earnings cannot be reconstructed with sufficient confidence for our purposes.

3.1 Panel-wide averages over time

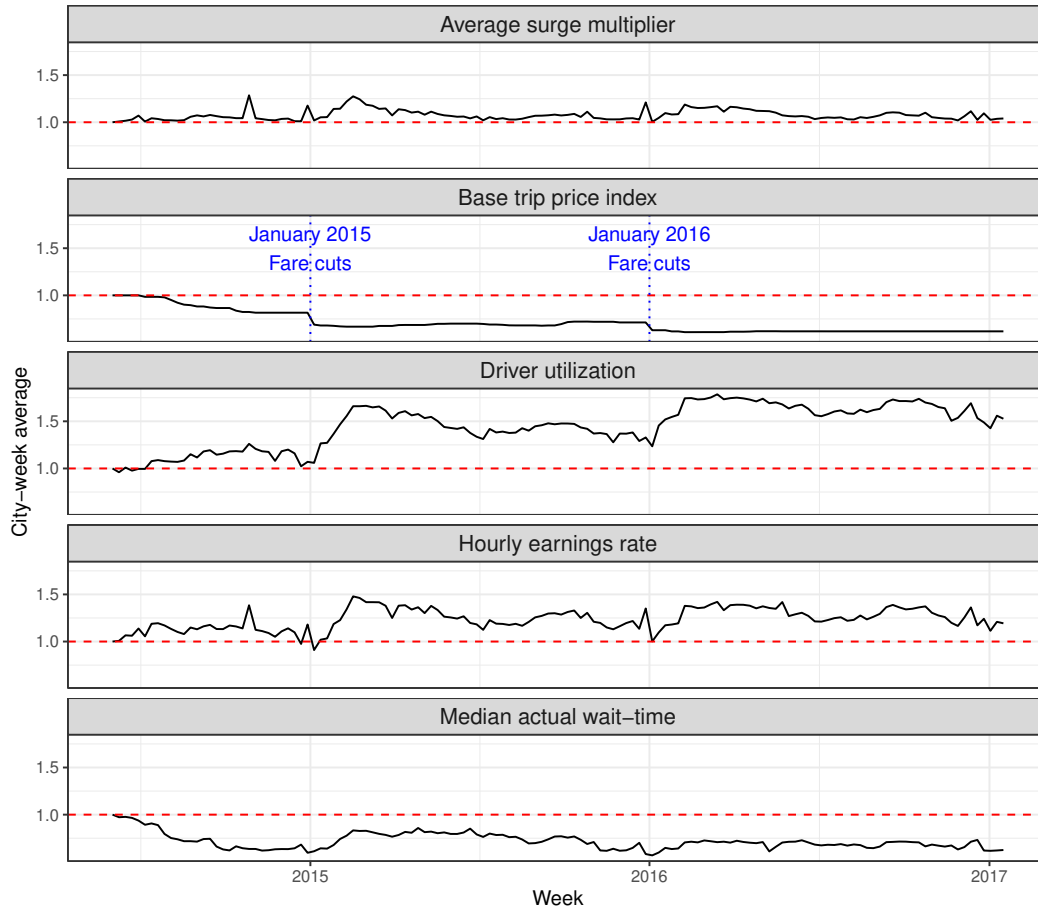
We first simply plot the weekly averages for the base fare index and our main outcome measures, pooled over all cities in the panel. Figure 1 shows, from top to bottom, the mean base price index, hourly earnings rate, utilization, average surge, and median wait time. All series are normalized to have a value of 1 in the first period of the panel.

In the top panel, we can see that there has been a long-run decline in the price index, though it has not been strictly monotonic. There are two clear sharp drops in the price index at the start of both 2015 and 2016 when Uber cut fares substantially in many cities. We will refer to these as the January fare cuts.

For the other market outcomes, Figure 1 shows several things. Perhaps most notably, the hourly earnings rate time series shows no obvious trend,

⁷The 50-city starting point is, of course, somewhat arbitrary, but this cutoff ensures a long panel of cities with substantial markets. As it is, not all cities in our panel are complete because even the top 50 include several markets that were not very mature at the start of the panel.

Figure 1: Average UberX market attributes over time for the US city-week panel, as indices



Notes: This figure plots the city-week panel weekly average for a collection of UberX market outcomes. All cities are weighted equally—see §3 for a definition of the sample. All series are turned into an index with a value of 1 in the first week.

though it does fluctuate. In contrast, driver utilization has increased substantially. There is little systematic change in average surge levels. Wait times were high early in the panel, but fell substantially by early 2015 and are more or less constant afterward.

3.2 The January fare cuts as event studies

The patterns shown in Figure 1 around the January cuts preview some of the main results from our regression analysis—namely little persistent change in the hourly earnings rate despite large changes in the base fare index. Immediately after the January fare cuts, the average surge increases, as does utilization and wait times. However, only utilization seems to show a persistent change in levels. Of course, we can present more credible evidence in a regression framework, but this “event study” approach offers visual confirmation of the direction of some of the effects.

3.3 Variation in base fare prices by city

As we saw in Figure 1, there were two large price reductions in January. However, Uber has changed the base fare for UberX in every city in the panel multiple times over the period covered by the panel. All these fare changes are shown in Figure 2a, with changes annotated in the week in which the UberX base fare index changed for the panel cities. The size of the change (as a percentage from the week before) is listed. A gray tile indicates that no change occurred that week.

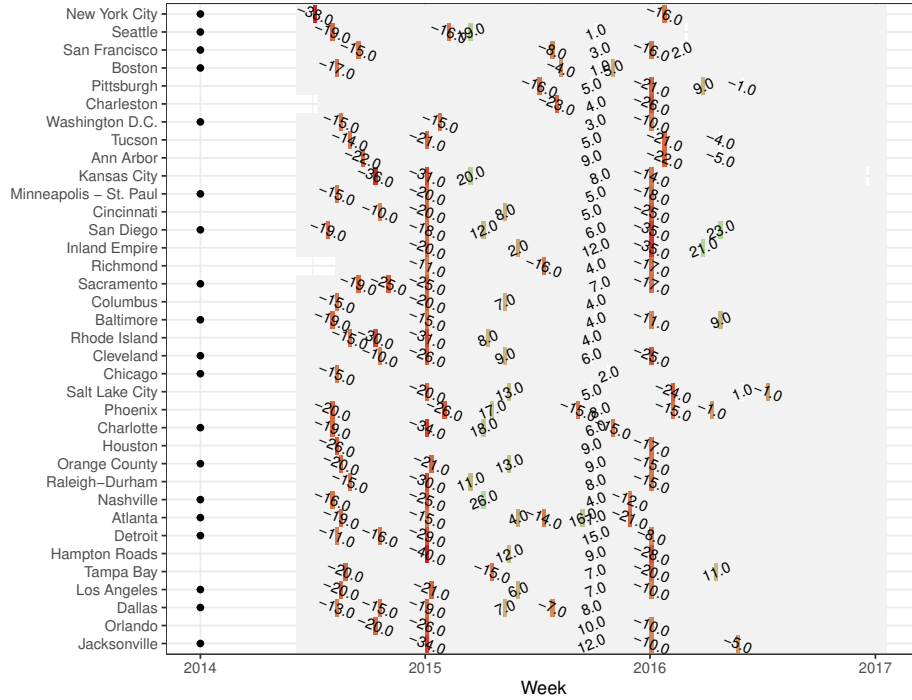
Cities are listed in descending order of their average base fare over the period. A black dot next to the city’s name indicates that that city had an UberBlack service. The large January fare cuts are clearly visible as a vertical band.

Figure 2b plots the histogram of all fare changes in the panel. We can see fare reductions are more common than fare increases, though fare increases are not rare. Fare reductions are also larger in magnitude than fare increases, on average. When discussing effects, we will describe the effects as if all changes were increases, though we use both increases and decreases for identification.

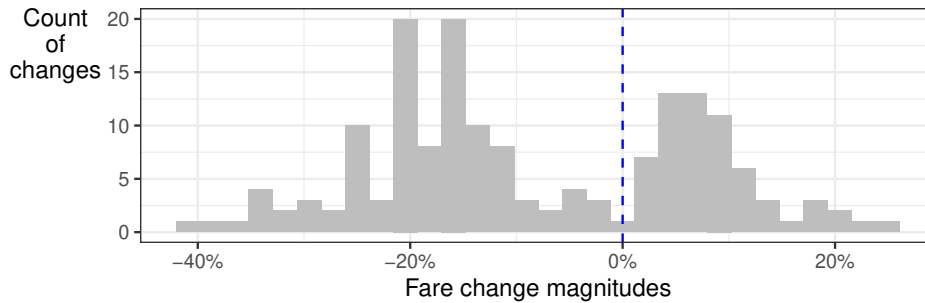
The decision to change fares in a particular city was made centrally by Uber’s internal pricing team, but in consultation with the responsible local teams. None of the authors of this paper were involved in this decision-making, but our understanding is that the pricing team considered market metrics when

Figure 2: UberX base fare index changes for US cities

(a) By-week changes



(b) Distribution of fare changes



Notes: The top panel (a) indicates which cities in the panel had changes in the base trip price index, by week, and reports the size of that change, in percentage terms relative to the fare index in the previous week. Squares that are not shaded gray indicate that no data is available for that city week. See §3 for a definition of the sample. The bottom panel (b) plots the histogram of all fare changes in the panel.

deciding on fare changes—primarily driver utilization but also average surge and hourly earnings rates. In short, prices appear to have been selected in response to observables.

Our analysis of the attributes of cities selected for large January fare cuts supports this view that utilization differences largely explained selection, at least with respect to the magnitude of fare cuts. Uber also claimed to be considering future changes in demand—namely reductions in demand due to impending winter weather—though this heuristic was apparently imperfectly followed, as we will show. Regardless, this kind of conditioning can readily be handled with our empirical approach, as we will discuss.

Although we will discuss identification issues at length in §5, it is useful to note that several gross features of the variation in price changes suggest a credible panel analysis is possible.

First, as every city in the panel had fare changes, it is not the case that latent differences exist between the kinds of cities that have fare changes and those that do not. A counter-point, however, is that we do not have true controls that never experienced any fare changes, which can create other complications (Goodman-Bacon, 2018). However, we are able to extensively break up our long panel into sub-panels to look for the kinds of heterogeneity in effects that would indicate a problem.

Second, Figure 2a shows that many changes took place in numerous cities nearly simultaneously, ruling out highly city-specific explanations for the existence of fare changes.

Third, although many changes are nearly simultaneous, they are not perfectly simultaneous. There is evidence of a staggered roll-out of some changes, with timing differing by a few weeks. It seems unlikely that the precise sequence of cuts reflects important latent differences between cities.

4 Model of a ride-sharing market

To explore the welfare implications of price changes and guide the empirics, we develop a simple model of a ride-sharing market. Although there are extant

models of taxi markets, they tend to focus on the micro details of search and matching, and the unique market properties this search process generates, such as non-existent/multiple equilibria or industry-level scale economies due to the nature of cruising e.g., [Arnott \(1996\)](#); [Frechette et al. \(2015\)](#) and [Cairns and Liston-Heyes \(1996\)](#).⁸ Our view is that we can usefully abstract away from the micro-details with market-level supply and demand curves that can be thought of as capturing population preferences over some substantial amount of time.

There is a price, p , for an hour of transportation that is set by the platform. This includes the base fare and the average surge multiplier.⁹ Drivers provide hours of labor that are turned into hours of transportation at a rate x , which is the endogenous market-level utilization. From an hour of work, drivers get $u_{drv} = (p - c)x$ where c is the added flow expense of being on a trip because of increased wear and tear, greater fuel expenditure, greater effort, and so on, versus being off-trip. Drivers collectively supply $H((p - c)x)$ hours, with $H'(\cdot) > 0$.¹⁰

⁸There is some newer work that builds on these insights ([Castillo et al., 2017](#)), some of which directly estimate models of search ([Frechette et al., 2015](#); [Buchholz, 2015](#)).

⁹See [Hall et al. \(2016\)](#) for evidence on the role of Uber’s surge pricing in clearing the market when demand spikes. [Banerjee et al. \(2016\)](#) discusses how dynamic pricing can add robustness to a system in the face of uncertainty in model parameters. See [Castillo et al. \(2017\)](#) for a discussion of the importance of surge pricing to prevent nearly discontinuous changes in wait times when demand outstrips supply.

¹⁰Our treatment of driver labor supply is simple, ignoring behavioral considerations, such as income targeting ([Camerer et al., 1997](#); [Thakral and Tô, 2017](#)) and even whether labor supply changes are due to extensive or intensive margin adjustments. Rather, we assume that labor supply can be captured with a single supply curve of total hours-worked. There is some evidence that behavioral labor supply considerations are relatively unimportant. [Farber \(2005, 2008\)](#) argues that income targeting findings are mostly due to division bias, and that driver behavior is mostly consistent with the neoclassical labor supply model. Errors in the measurement of hours-worked tend to attenuate an estimate of the labor supply since the measurement of hours is also used to calculate the wage. A key advantage of our empirical setting is that we can measure hours-worked essentially without error. [Farber \(2015\)](#) shows that there is substantial heterogeneity in individual labor supply elasticities and that drivers that do not learn to work more when wages are temporarily high are not long for the taxi driving profession. Using data from Uber, [Chen and Sheldon \(2015\)](#) also presents evidence that Uber drivers are responsive to hourly earnings in a neoclassical fashion and that there is little evidence of income targeting. Also using data from Uber, [Angrist et al. \(2017\)](#) also find no evidence of income targeting.

Passengers are sensitive to the price of an hour of transportation, but also the quality. The primary quality dimension of interest, at least in the short-run, is the wait time, or the time elapsed from requesting a ride until actually starting a trip (see [Buchholz et al. \(2020\)](#) on the value would-be passengers place on waiting time).¹¹ We assume that the wait time is determined by the market tightness, which is the ratio of passengers demanding trips and the drivers actively working. In equilibrium, this market tightness is the same as utilization, as we will show.¹² Intuitively, the reason for this relationship between utilization and wait time is that, all else equal, a higher utilization means the nearest available driver is farther away.

The cost of the wait time to passengers is $\phi(x)$, with $\phi'(x) > 0$, and $\phi''(x) > 0$ and so passenger utility from an hour of transportation is $u_{pax} = -(p + \phi(x))$. There is a market-level demand for hours of transportation, $D(p + \phi(x))$, with $D'(\cdot) < 0$. Market-clearing requires that

$$D(\underbrace{p + \phi(x)}_{-u_{pax}}) \equiv xH(\underbrace{(p - c)x}_{u_{drv}}). \quad (1)$$

Note that supply hours are scaled by x , the utilization, and that equilibrium utilization is also the equilibrium market tightness, i.e., $x^* = D^*/H^*$.

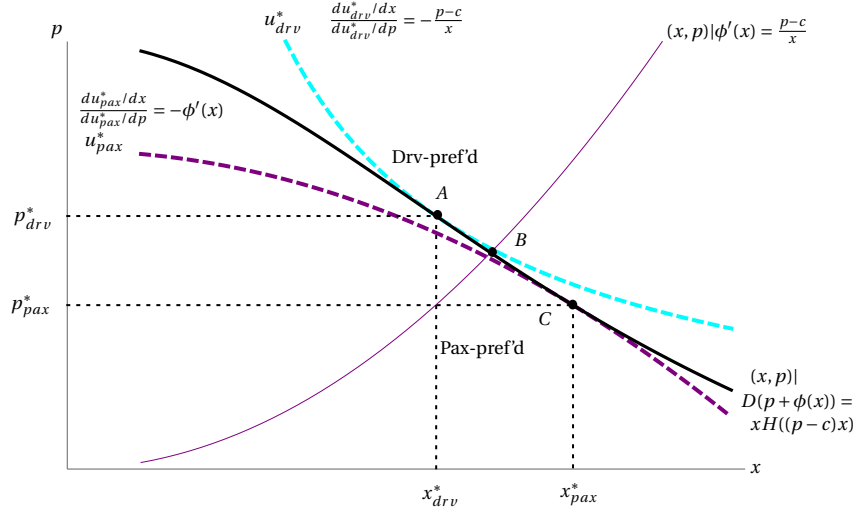
In [Figure 3](#), we indicate possible equilibria with a heavy, downward-sloping line. These are the prices and utilization levels that satisfy [Equation 1](#). The x-axis is the level of driver utilization, x and the y-axis is the price, p . The origin is the passenger’s bliss point—a very low price and low utilization, which means short wait times. The driver’s bliss point is up and to the right—a very high price and high utilization, which gives a high hourly earnings rate.

[Proposition 1](#) shows that the market equilibrium utilization, x , is decreas-

¹¹Although there could be a decline in service quality not measured by wait times, the lack of composition effects on the driver extensive margin (as we will see) makes this somewhat unlikely, though drivers could presumably be ruder or less helpful, keeping a less clean car, and so on.

¹²Although we do not explicitly model the micro-details of matching, the specific wait time/utilization assumption we make is easy to motivate with a reasonable matching function between passengers and empty cars—See [Appendix A.4](#).

Figure 3: Possible equilibria and driver and passenger indifference curves



ing in the price, p : there is an equilibrium trade-off between price and utilization.

Proposition 1. *The market-level utilization is declining in the price: $\frac{dx_{eq}}{dp_{eq}} \leq 0$. See Appendix A for proof.*

There is an equilibrium for each choice of p , but there are three equilibria that are of special interest, labeled A , B , and C in Figure 3: A , the driver-preferred equilibrium, which maximizes $(p - c)x$ and hence hours of work (as $H'(\cdot) > 0$); B , the Pareto equilibrium (where driver and passenger indifference curves are tangent to each other); and C , the passenger-preferred equilibrium, which minimizes $p + \phi(x)$ and hence maximizes hours of transportation (as $D'(\cdot) < 0$). Proposition 2 characterizes the three equilibria in terms of driver and passenger indifference curves and a set of possible equilibria. Let p_{drv}^* and p_{pax}^* be the prices at A and C , respectively. An important implication of Proposition 2 is that $p_{drv}^* > p_{pax}^*$ and, following from Proposition 1, $x_{drv}^* < x_{pax}^*$.

The reason the driver's indifference curve is tangent at A is that from the driver's preferred equilibrium, any change in equilibrium price (i.e., movement along the heavy black curve of possible equilibria) would make a driver worse off. A similar logic applies at C for the passenger's indifference curve.

At the point, B , the Pareto equilibrium, the driver and passenger indifference curves are tangent to each other. However, they are *not* tangent to the possible equilibria curve—from B , drivers would still prefer a higher price and passengers would prefer a lower price.

Although the label of B as the Pareto equilibrium would seemingly imply some special desirability, it is just that at B , both passengers and drivers would be indifferent to the same small increase in price that came with a specific utilization increase/wait-time decrease. However, this trade-off is not even possible at B so long as A and C are separate equilibria.

Proposition 2. *The three “interesting” equilibria are:*

$$\begin{aligned}
 A) \text{ Driver-preferred } (p_{drv}^*, x_{drv}^*) : & \quad \frac{p_{drv}^* - c}{x_{drv}^*} = \left| \frac{dp_{eq}}{dx} \right| \\
 B) \text{ Pareto } (p_P^*, x_P^*) : & \quad \phi'(x_P^*) = \frac{p_P^* - c}{x_P^*} \\
 C) \text{ Passenger-preferred } (p_{pax}^*, x_{pax}^*) : & \quad \phi'(x_{pax}^*) = \left| \frac{dp_{eq}}{dx} \right|.
 \end{aligned}$$

with $p_{drv}^* \geq p_P^* \geq p_{pax}^*$ and $x_{drv}^* \leq x_P^* \leq x_{pax}^*$. See Appendix A for proof.

In considering the effects of fare changes on the market equilibrium, there is one unambiguous prediction: a higher fare will lower utilization (Proposition 1). For other market outcomes, the predicted effects depend on the equilibrium that prevailed when the change was made. When the price is above what drivers would prefer (higher than A), a fare decrease can raise both hours-worked and hours of transportation and is Pareto-improving. Below this price but above the passenger’s preferred point (in between A and C), fare cuts are desired by passengers but not by drivers. At the passenger’s preferred price and below (prices less than C), further price cuts make both sides worse off and so price increases would be desired from this price level.

In the model, because hours-worked is monotonic in driver utility and hours of transportation is monotonic in passenger utility, changes in these quantities following fare changes are sufficient statistics for changes in welfare for each side of the market. The signs of these changes in quantities also partially

identify the location of the equilibrium from which a change was made, namely whether the equilibrium price that prevailed was 1) above the price at A , 2) below the price at A but above the price at C or 3) below the price at C . The market-level comparative statics are summarized in Proposition 3.

Proposition 3. *The effects of a fare increase from p on market quantities are:*

	x	D	H	
$p > p_{drv}^*$	↓	↓	↓	<i>Fare cut Pareto-improving</i>
$p_{drv}^* > p > p_{pax}^*$	↓	↓	↑	<i>Ambiguous; Pax want cuts; Drivers want increases</i>
$p < p_{pax}^*$	↓	↑	↑	<i>Fare increase Pareto-improving.</i>

See Appendix A for proof.

It might be tempting to use this model to derive platform incentives, but given that these platforms likely do enjoy scale economies (Arnott, 1996), a static perspective on profit-maximization would likely be highly misleading as a description of their actual behavior.

5 Results

In this section, we present estimates of the effects of changes in the base fare on market outcomes. We will first present the static, long-run effects of fare changes, and then report results from dynamic specifications that allow us to trace out the market adjustment. It is important to emphasize that we are *not* estimating supply or demand elasticities, but rather characterizing the movement from one market equilibrium to another.

Unlike in the model, the platform does not choose p directly. Instead, the platform changes a base fare index, b , with the actual p faced by passengers, on average, being $p = mb$, where m is the average surge multiplier, which is also a choice made by the platform.

Regarding connecting the model to empirics, fare increases should lower utilization and decrease wait times. However, the effect on hours of transportation and hours of work depends on “where” the market is on the possible equilibria curve. While we might think of the direction of effects partially identifying the equilibrium that prevailed in all markets, different markets could

be in different equilibria when fares were changed. This difference could potentially lead to us making a kind of ecological fallacy if we presumed that the average effect applied to each unit. This composition possibility requires us to be cautious in interpretation, but to the extent we think Uber was pricing similarly in all markets, the sign of effects on quantities should indicate “where” these markets existed according to Proposition 3. Empirically, there is no evidence for concern about composition.¹³ When we divide the panel into shorter sub-panels, we find similar results in each panel, suggesting that despite large changes in average fares over time (generally downward), the same directional effects prevailed and that all markets were operating in the A to C “range.”

5.1 Static estimates of effects on market-level outcomes

Our baseline specification is

$$y_{it} = \alpha_i + \beta_1 \log b_{it} + g_i t + d_t + \epsilon_{it}, \quad (2)$$

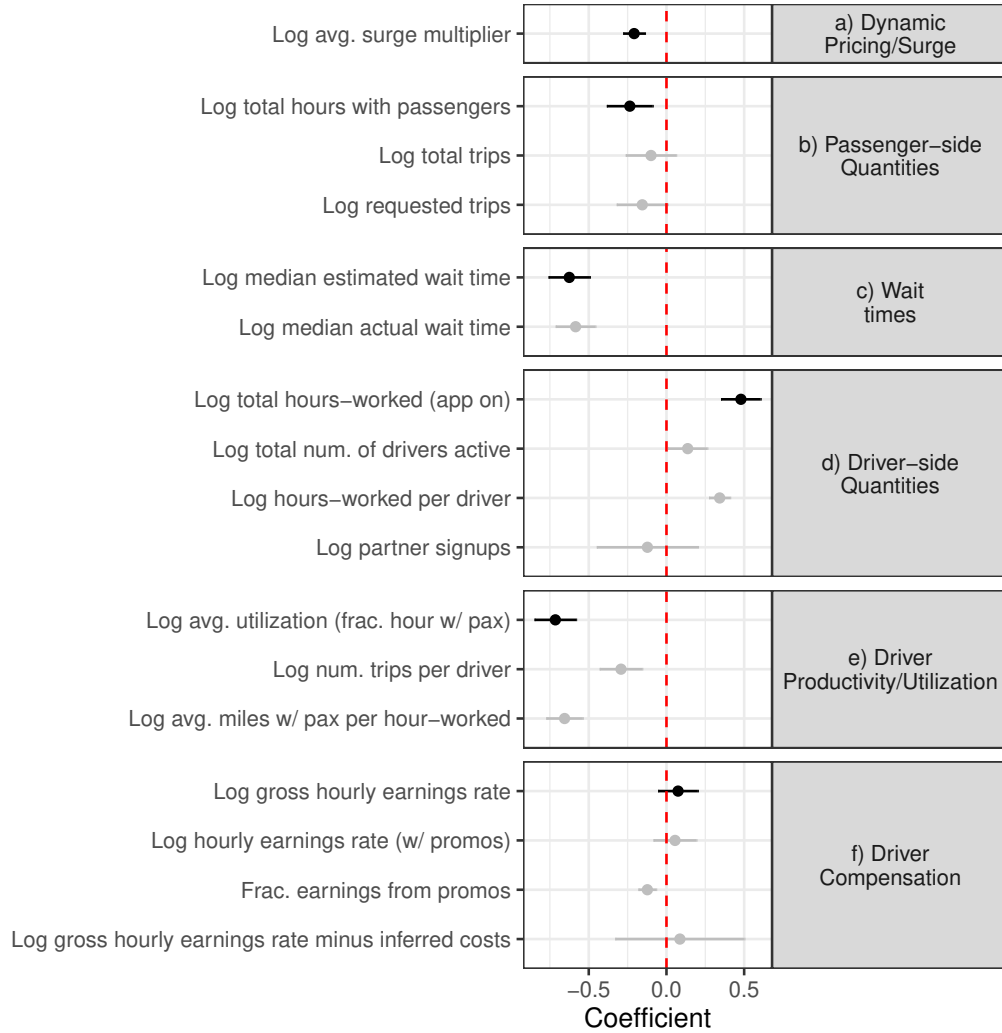
where y_{it} is some market-level outcome of interest in city i during week t , α_i is a city-specific fixed effect, b_{it} is the base trip price index, g_i is a city-specific linear time trend, d_t is a week-specific fixed effect and ϵ_{it} is an error term.¹⁴

Figure 4 plots estimates of β_1 from Equation 2 for a collection of market outcomes, along with 95% CIs. Each facet of the figure reports estimates of different measures of the same “kind” of market outcome. For example, the facet labeled “Passenger-side Quantities” reports the effect of a fare increase on the log total hours of transportation (which is the quantity described in the model), but also the number of requested trips and completed trips. In each facet, our primary outcome estimated is plotted in black and shown first; related but secondary outcomes are shown below and plotted in gray.

¹³This analysis is in §B.6.

¹⁴In §B.9, we report all regressions reported here, but without city-specific time trends. Generally, these trends improve the precision of the estimates (particularly for market quantities), as forcing all cities to only differ by a level over the entire panel leads to systematic residuals for some cities.

Figure 4: All point estimates of effects of Uber fare changes on market outcomes



Notes: All estimates are of β_1 from Equation 2. These regressions in table form are in §C.

5.1.1 Dynamic pricing/surge

In the facet of Figure 4 labeled “Dynamic Pricing/Surge” the only outcome is the average surge multiplier. With a higher base fare, the average surge multiplier declines. However, the effect is not large enough to undo a base fare change (it would need to have a point estimate of -1 , not the ≈ -0.2 ob-

served). We can conclude that in equilibrium, passengers faced higher average prices after a fare increase.

5.1.2 Passenger-side quantities

In the facet of Figure 4, labeled “Passenger-side Quantities,” the first outcome is the log total hours with passengers. This outcome corresponds to D^* in the model. We can see that a higher base fare reduces the equilibrium number of hours of transportation. The effect size is about a 2.5% reduction in total transportation hours from a 10% increase in the base fare. Other quantity outcomes in this facet—the log total completed trips, as well as the log total requested trips (regardless of whether they were completed)—show broadly similar, if somewhat smaller, effect sizes.

These quantity point estimates—a 10% fare decreasing trips by about 2.5%—are smaller than the demand elasticities estimated by Cohen et al. (2016), which uses step variation in surge pricing. The difference in the estimates highlights the difference in the nature of the quantity we are estimating, namely the change from one market equilibrium to another, rather than the slope of a demand curve. As we will see, with a higher fare, wait times fell considerably and this surely matters to would-be passengers (Buchholz et al., 2020), and so a fare change should not be thought of as the market moving along a demand curve.

An implication of the reduction in hours of transportation is that fare increases were being made from an equilibrium where $p > p_{pax}^*$, i.e., the price was not so low that passengers welcomed the fare increase (which follows from Proposition 3).

5.1.3 Wait times

In the facet of Figure 4 labeled “Wait times,” the two outcomes are the median estimated and actual wait times. Both decreased substantially with a higher fare. Presumably, this would have some offsetting effect on demand (recall $\phi'(x) > 0$), though to reiterate we cannot measure the effect, as is confounded

with the effects of the fare change.

5.1.4 Driver-side quantities

In the facet of Figure 4 labeled “Driver-side Quantities,” the first outcome is the log total number of hours worked by drivers. This outcome corresponds to H^* in the model. We can see that with a higher base fare, drivers work more hours in total. The increase in hours-worked is consistent with the fare increase happening from a $p < p_{drv}^*$, i.e., drivers welcomed the fare increase. Combined the passenger-side analysis that $p > p_{pass}^*$, we have evidence that fare changes were happening from an equilibrium in the range of equilibria between the driver and passenger most-preferred equilibria i.e., between A and C in Figure 3.

We also report the effects of a fare increase on the number of active drivers and the number of hours-worked per active driver. Both increase substantially with a higher fare, though the intensive margin effect is larger. We also report the effects on the log number of driver sign-ups, though this outcome is so imprecisely estimated that little can be concluded.

5.1.5 Driver productivity/utilization

In the facet of Figure 4 labeled “Driver Productivity/Utilization,” the first outcome is driver utilization or x^* in the model. Consistent with Proposition 1, a higher fare causes a lower utilization. The effect is substantial: a 10% increase in the base fare reduces equilibrium utilization by about 7%. We also report the log number of trips per hour-worked, which also declines, as does the log number of passenger miles per hour-worked. These outcomes point to a lower driver technical productivity with higher fares. However, given the increase in the fare, productivity has not necessarily changed, as drivers are performing fewer but more valued trips on the margin.

5.1.6 Driver compensation

In the facet of Figure 4 labeled “Driver Compensation,” the first outcome is the log gross hourly earnings rate. This would correspond to $\approx px$ in the model, as costs are not included. The effect of a fare increase is slightly positive, but the confidence interval includes zero: a 10% increase in fare raises driver hourly earnings by 0.7%. If driver hourly earnings did increase, it is consistent with fare increases being made from within the A to C equilibrium range.

Unlike in the model, the platform can pay drivers beyond what they earn from trips through “promotional payments.” The next outcome in the facet is the log hourly earnings rate including promotional payments. The effect is slightly positive and similar to the gross hourly earnings measure. While the confidence intervals clearly overlap and relative comparisons are hard to make, we would expect the measure with promotional payments included to show an effect closer to zero, as fewer drivers would fall below a floor level of earnings (which could trigger promotional payments).

We can look directly for evidence of promotional payments making up for lowered cumulative earnings. The next outcome down in the facet is the fraction of a driver’s earnings coming from promotional payments. Note that this measure is not in logs, as there are numerous zeros, as relatively few cities in the panel actually use promotional payments at all. The promotional payment fraction declines slightly, consistent with the effect on the hourly earnings rate including promotional payments being slightly closer to zero.

The last measure is the log gross hourly earnings rate minus inferred costs, which is a measure that imputes miles driven per hour and uses a \$0.30/mile rate. The effects of a fare increase are also slightly positive for this measure, but far less precise than the other measures.¹⁵

Although there is no strong evidence of an increase in driver hourly earnings, given the increases in hours-worked observed in the driver-quantities analysis, it seems likely that the net effect on driver utility from a fare increase was positive, consistent with markets being at an equilibrium “between” A and C

¹⁵Likely due to the inferred costs shifting the outcome in levels to the left, and thus inflating the variance in a logged outcome.

when fares were changed.

5.2 Dynamics of market adjustment

There was little to no long-run pass-through of the fare changes into the driver hourly earnings rate, but large changes in wait times and utilization. However, presumably, this effect was not realized immediately, but rather was the result of a market adjustment that took place over time. To explore how the markets adjusted over time, we estimate a finite distributed lags model,

$$y_{it} = \alpha_i + \sum_{\tau=\text{NUMPRE}}^{\text{NUMPOST}} \beta_{\tau} \log b_{it-\tau} + d_t + g_i t + \epsilon_{it}, \quad (3)$$

where α_i is a city-specific fixed effect, b_{it} is the fare index in city i at time t , τ the number of weeks from the focal week, g_i is a city-specific linear time trend and d_t is a week-specific fixed effect. The number of pre-period week indicators is NUMPRE and the number of post-period weeks indicators is NUMPOST. Note that with this specification, multiple fare changes can be included in the estimate and a “focal” week does not need to be specified.

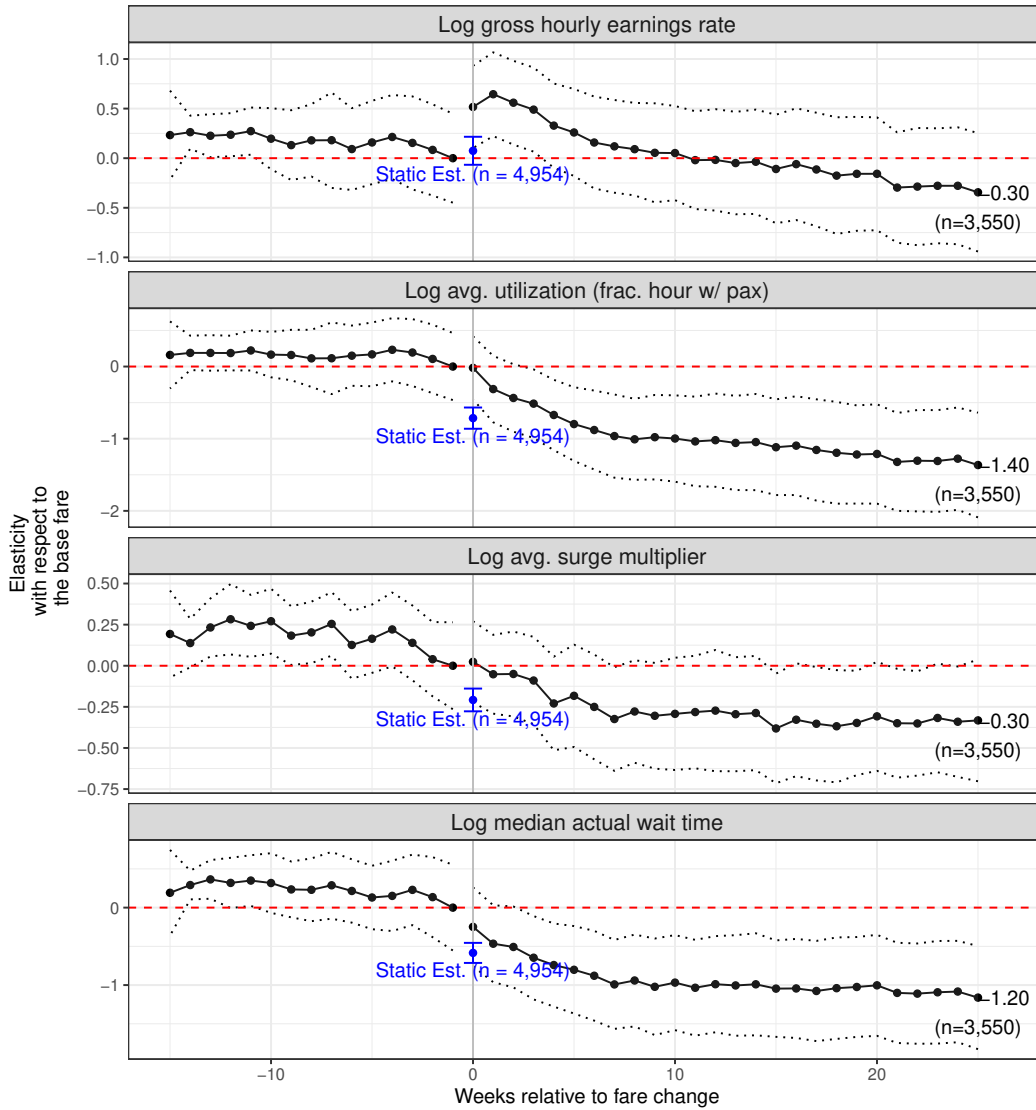
We impose the restriction when estimating the model that $\sum_{\tau=\text{NUMPRE}}^0 \hat{\beta}_{\tau} = 0$ i.e., that the cumulative effect in the week prior to the fare change is 0. This allows for cities having fare changes to differ from those not having changes in a given week by a level amount, but the inclusion of multiple per-period windows still allows us to detect whether those cities were on different trajectories with respect to the outcome.¹⁶

The implied weekly effects from Equation 3 are plotted in Figure 5 for the log hourly earnings rate, log utilization, and log surge. For comparison, the “static” Equation 2 effect for each outcome is also plotted at the “0” time. For each regression, standard errors are clustered at the level of the city.¹⁷ There

¹⁶In §B.9 we report the same distributed lag models as in the main body, but without imposing the zero effect at week -1 assumption. For some outcomes, not imposing this restriction leads to pre-period effects that are systematically higher or lower, but as expected, we see no evidence of trends. Further, the pre-period level differences are generally fairly modest in magnitude.

¹⁷We also conducted a block bootstrap at the city level to test for Bertrand et al. (2004)

Figure 5: Effects of a base fare increase on the driver hourly earnings rate and its components



Notes: This figure plots the by-week cumulative effects of changes in the UberX base fare on market outcomes. These effects are from an estimation of Equation 3; at $t = 0$, the static estimates from Equation 2 are shown. The sample is a panel of US cities—see §3 for a description. The x-axis is weeks relative to a fare change.

problems, but we found that the bootstrap standard errors were almost identical to the clustered standard errors, so we only report clustered standard errors.

are 15 pre-periods and 25 post-periods. These leads and lags were selected by visually inspecting various combinations and seeing where the cumulative effect “flattens out” in the post period and then doing a sensitivity analysis around the window length choice.¹⁸

In the top facet of Figure 5, the outcome is the log gross hourly earnings rate. Prior to the fare change, we see no evidence of a worrisome trend. If we look at effects over time, following a fare increase, the driver hourly earnings rate increases immediately, though there is considerably less than a full pass-through; the elasticity point estimate is only about 0.6 in the first week after. In the weeks that follow, this increase in the hourly earnings rate declines, with the point estimate at week 8 being 0.1. Unlike the static estimate, the long-run estimate near the end of the post-period is negative, albeit with a 95% CI that includes zero.

In the second facet from the top of Figure 5, the outcome is the driver utilization. In the pre-period, there is no evidence of a trend. Utilization falls following a fare increase, though the effect is not immediate—in the 0 week, the effect is almost precisely 0, whereas we observed that the driver hourly earnings rate jumped immediately. However, by week 8, the elasticity point estimate is -1, which is close to the estimate of the static effect from Equation 2.

In the second facet from the bottom of Figure 5, the outcome is the log average surge. There is no obvious trend in the pre-period and the pre-period weekly point estimates are all close to zero. The average multiplier gradually declines following a fare increase. By week 8, the effect is pretty close to the static estimate that a 10% increase in fares reduces average surge by 2%.

In the bottom facet of Figure 5, the outcome is the log median wait time. There is no evidence of a pre-trend, and post-fare cut, the path is similar

¹⁸In §B.9, we report our preferred regression specification but vary the post-period bandwidth. The various plots illustrate that results are not sensitive to somewhat larger and smaller lead/lag windows. Because of the structure of our data, larger pre-period windows do cause a loss of usable data. Given that the state of the literature on lead/lag selection seems more art than science, we felt a visual approach checked for robustness with different periods was preferable to something more model-driven.

to the path of driver utilization. With a higher fare, the wait time declines substantially.

6 Identification issues

There are a number of threats to identifying the effects of fare changes on market outcomes. For some of these putative threats, our regression specification alone is sufficient. For other threats, our approach is diagnostic, in that we assess the existence or extent of the problem, often by using other sources of data or conducting various robustness checks. For threats for which we have no diagnostic test, we rely on secondary sources to assess whether the issue is likely to affect our results. In this section, we discuss the issue and make references to appendices where the actual analysis is reported.

6.1 What was Uber doing?

The variation in fares was created by Uber, but we do not fully know their decision-making process. Perhaps Uber—but not us—had access to information that allowed it to make accurate predictions about city-specific trajectories, and was conditioning their decisions accordingly. We view this as unlikely.

As Uber fare changes were controversial, they received media attention, which we summarize in §B.1. Contemporaneous media reports of Uber’s decision-making—including reviews of leaked spreadsheets that were used as decision support tools—strongly imply that Uber “forecasting” models were in fact simply accounting exercises. The spreadsheets calculated how many more trips a driver would have to complete to keep their earnings the same given what was observed in the past—with no consideration of how fare changes would change both supply and demand in reality. In short, if these spreadsheets were used as described, Uber’s decision-making could be characterized as selecting on observables that we observe.

As a direct piece of evidence for this selection on observables claim, we can use the January cuts as a kind of case study. Using the January fare cuts,

we show in §B.1 that the magnitudes of the cuts were clearly conditioned on city-specific measures of utilization that prevailed before the cut, but not on trends.

Uber’s pricing decisions could be selecting on trends rather than levels. We partially address this concern by including city-specific linear time trends in Equation 2. We can also diagnose whether there is a violation of parallel trends by plotting cumulative effects from a distributed lag model, prior to treatment. We find no evidence of pre-trends in any of our outcomes—this can be seen for some outcomes in Figure 5.

6.2 Other potential factors

Even if there are no unmodeled city-specific dynamics prior to a cut, a worry is that cities selected for fare changes were *about* to have some change in conditions that motivated the change, thus violating the strict exogeneity assumption. Three likely candidates are (1) weather-related demand shocks, (2) the action of competitors, and (3) local economic conditions.

For weather, we can compare the two January fare cuts to see if cold weather cities are exclusively getting cuts. We do this comparison in §B.2, finding that in each set of January fare cuts, many warm-weather cities get cuts and many cold-weather cities do not. This analysis also supports the view that cities receiving cuts were not on different trends, as we can see graphically, that cities receiving January cuts were not on different trajectories relative to “control” cities.

Further supporting the argument that weather was not important for selecting cities for January fare cuts, across the two January cuts, there is substantial variation in who is “treated” and who is not. We can exploit the fact that not all cities had the same “treatment assignment” around the two January cuts to perform a placebo test. Our long time period also allows us to include city-calendar month-specific fixed effects to capture city-specific changes due to climate differences. Doing this analysis in §B.3, we find that our key point estimates do not change substantively with the inclusion of these

city-calendar month effects.

Uber is not the only ride-sharing company, and in the period covered by our data, Uber faced direct competition from other ride-sharing platforms in some cities. Despite the possibility that the market adjustment process could be affected by the presence of competitors—namely by making both sides more elastic—using data on city-specific ride-sharing platform shares, we find no evidence that this is the case. However, it is important to note that during most of the period covered by our panel, competition from alternative ride-sharing platforms was nascent. Furthermore, to the extent competitor ride-sharing platforms followed the pricing decisions of Uber, a fare change could have been, in a sense, market-wide. We have some limited evidence that Uber’s competitors matched Uber’s fare changes but lack the comprehensive by-week pricing data we have for Uber.¹⁹

For direct competitors, we calculate Uber’s share of the ride-sharing category in each city-week, using monthly data and then see whether controlling for share changes the results. We do this analysis in §B.4 and find that including Uber’s imputed ride-sharing share leaves the key point estimates more or less unchanged. We also interact the base fare with this share measure to see whether the extent of ride-sharing competition affected the adjustment process. We find no evidence that they do, though it is important to re-emphasize, during the period covered by our panel, direct ride-sharing competition was in many cities non-existent.

For local economic conditions, we can also control for city-specific local economic conditions, as measured by the MSA unemployment rate, which we do in §B.5. We find no evidence that the inclusion of these controls affects the results.

Even with city-specific linear time trends, a concern in a long panel like

¹⁹See: <https://www.geekwire.com/2014/lyft-slashes-prices-nationwide-20/>;
<https://techcrunch.com/2014/04/07/lyft-spring-pricing/>;
[https://www.forbes.com/sites/briansolomon/2016/01/25/
is-uber-trying-to-kill-lyft-with-a-price-war/#1d24e9f56573](https://www.forbes.com/sites/briansolomon/2016/01/25/is-uber-trying-to-kill-lyft-with-a-price-war/#1d24e9f56573); <https://blog.lyft.com/lyftline5?rq=price>;
[https://www.digitaltrends.com/mobile/
lyft-uber-price-cuts/](https://www.digitaltrends.com/mobile/lyft-uber-price-cuts/); <http://fortune.com/2016/01/15/lyft-price-cuts/>;
<https://blog.lyft.com/posts/start-off-2016-with-lower-prices?rq=price>.

ours is that even a city-specific linear trend is not sufficient to meet the strict exogeneity assumption. For example, cities with lower-than-expected utilization, even conditional on the included controls, might be targeted for fare changes, creating a correlation between ϵ_{it} (as specified) and b_{it} in Equation 2. A related problem in a long panel like ours (in which units are being treated multiple times) is the possibility of treatment effect heterogeneity creating “bad” controls, leading to undesirably weighted estimates (Abraham and Sun, 2018; Athey and Imbens, 2018; Goodman-Bacon, 2018).

One diagnostic approach to assess this possibility is to divide the sample into shorter panels and compare sub-panel estimates to the overall estimates, which we do in §B.6. We find that both the “short T” point estimates are quite similar to the full panel estimates for all of our outcomes. This lack of difference suggests our Equation 2 specification is sufficient, as well as undercuts the notion that treatment effects might be differing over time. Another diagnostic approach is to conduct the statistical test for strict exogeneity proposed by Wooldridge, which we do in §B.7, finding that we cannot reject the null of strict exogeneity.

6.3 Cuts versus increases

Even if β_1 from Equation 2 is identified, this single parameter estimate could mask substantial heterogeneity in effects. For power reasons, we cannot explore every possible interaction effect, but there are some that could be particularly consequential. One worry is that the effects of fare cuts could be different from the effects of fare increases, implying “kinked” demand and supply curves. We diagnose whether this is a problem in §B.8 by estimating our model with sub-panels in which the variation in the base fare is all of one “kind” i.e., all increases or all decreases. We find that the point estimates for all our outcomes are the same sign and of similar magnitudes. In short, cuts and increases seem to “work” the same way.

7 Discussion and conclusion

The key finding of the paper is that following a fare change, ride-sharing markets adjust primarily through changes in driver utilization and changes in wait times. This occurs because drivers respond to temporarily higher “wages” by working more hours, which has a business stealing effect. In the long run, a fare increase seems to leave driver hourly earnings nearly unchanged, or perhaps slightly higher. The lack of price effects *on average* seems to apply even to the introduction of Uber into US cities—[Berger et al. \(2017\)](#) presents evidence that the introduction of Uber lowered the average hourly earnings of professional drivers, but as [Angrist et al. \(2017\)](#) points out, the increase in earnings from self-employed drivers left the average unchanged.

By showing the connection between the product market price and market efficiency, our results speak to the larger question of why some platforms take on price-setting, despite the well-known challenges of doing so ([Hayek, 1945](#)). Short of setting a price, a platform could simply make price comparison easier, which appears to be sufficient in some cases ([Jensen, 2007](#)), but not all cases ([Dinerstein et al., Forthcoming](#)). In our setting, price comparison would be relatively costly to buyers given the “perishable” nature of the service and the large differences in match quality created by the spatial component of for-hire transportation ([Castillo et al., 2013](#)). As such, it seems probable that without centralized price setting, the logic of [Diamond \(1971\)](#) could lead to an inefficient high price/low utilization equilibrium, despite free entry on the supply side.²⁰

Uber has more recently decoupled rider and driver trip prices with upfront pricing, but drivers generally continue to earn trip pay determined by trip time and distance. Even with the move towards upfront pricing, Uber still has to decide on an approximate price level, or what the “average” trip will cost. As such, more sophisticated pricing does not sidestep the issue of choosing a price level. The platform could switch to an auction model, with drivers

²⁰[Filippas et al. \(2018\)](#) reports the results of an experiment conducted in a computer-mediated marketplace, showing that the platform substantially raised utilization when it centralized (and lowered) pricing.

submitting bids, though given the “perishable,” time-sensitive nature of the service, this would likely be surplus-dissipating, especially given the trend in online markets away from auctions for primarily taste-based reasons (Einav et al., 2018). Our conclusions should remain relevant to future scenarios where drivers are paid per trip, regardless of the exact price structure.

With a higher driver utilization, each hour of work is more productive, allowing Uber to meet the same amount of passenger demand with aggregate hours of work—though subject to the caveat that wait times matter. Although utilization is, as we show, highly sensitive to the fare, it also is presumably affected by technological considerations. Many of Uber’s platform improvements can be interpreted as attempts to raise utilization through technological means, such as “back to back trips” (matching drivers before their current trip is finished based on predicted drop-off time and location) and having passengers re-locate slightly before pick-up.

This paper has focused on market-level attributes and outcomes. A natural direction for future work would be to take an individual driver’s perspective. In particular, it would be interesting to consider driver micro labor supply decisions, focusing on the role of individual differences in costs. Although we have modeled drivers as only caring about utilization to the extent it affects their hourly earnings, it seems probable that drivers vary in their preferences over the different utilization equilibria, both because of their personal preferences about being “busy” as well as their capital, with drivers with less fuel-efficient vehicles preferring the low utilization equilibrium.

References

- Abraham, Sarah and Liyang Sun**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” 2018.
- Allegretto, Sylvia A, Arindrajit Dube, and Michael Reich**, “Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data,” *Industrial Relations: A Journal of Economy and Society*, 2011, *50* (2), 205–240.
- Angrist, Joshua D., Sydney Caldwell, and Jonathan V. Hall**, “Uber vs. Taxi: A driver’s eye view,” *Working Paper*, 2017.
- Arnott, Richard**, “Taxi travel should be subsidized,” *Journal of Urban Economics*, 11 1996, *40* (3), 316–333.
- Athey, Susan and Guido W Imbens**, “Design-based analysis in difference-in-differences settings with staggered adoption,” Technical Report, National Bureau of Economic Research 2018.
- Banerjee, Siddhartha, Ramesh Johari, and Carlos Riquelme**, “Dynamic Pricing in Ridesharing Platforms,” *SIGecom Exch.*, sep 2016, *15* (1), 65–70.
- Berger, Thor, Chinchih Chen, and Carl Benedikt Frey**, “Drivers of disruption? Estimating the Uber effect,” *Working Paper*, 2017.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How much should we trust differences-in-differences estimates?,” *The Quarterly Journal of Economics*, 2004, *119* (1), 249–275.
- Bewley, Truman F**, *Why wages don’t fall during a recession*, Harvard university press, 2009.
- Bryan, Kevin A. and Joshua S. Gans**, “A theory of multihoming in rideshare competition,” *Journal of Economics & Management Strategy*, 2019, *28* (1), 89–96.

- Buchholz, Nicholas**, “Spatial equilibrium, search frictions and efficient regulation in the taxi industry,” *Working Paper*, 2015.
- , **Laura Doval, Jakub Kastl, Filip Matějka, and Tobias Salz**, “The Value of Time: Evidence From Auctioned Cab Rides,” Technical Report, National Bureau of Economic Research 2020.
- Cairns, Robert D. and Catherine Liston-Heyes**, “Competition and regulation in the taxi industry,” *Journal of Public Economics*, 01 1996, 59 (1), 1–15.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler**, “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics*, 1997, 112 (2), 407–441.
- Castillo, Juan Camilo, Daniel T. Knoepfle, and E. Glen Weyl**, “Surge pricing solves the wild goose chase,” *Working Paper*, 2017.
- Castillo, Marco, Ragan Petrie, Maximo Torero, and Lise Vesterlund**, “Gender differences in bargaining outcomes: A field experiment on discrimination,” *Journal of Public Economics*, 2013, 99, 35–48.
- Chen, M. Keith and Michael Sheldon**, “Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform,” *Working paper*, 2015.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe**, “Using big data to estimate consumer surplus: The case of Uber,” *Working Paper*, 2016.
- Cullen, Zoë and Chiara Farronato**, “Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms,” *Working Paper*, 2018.
- Diamond, Peter A.**, “A model of price adjustment,” *Journal of Economic Theory*, 1971, 3 (2), 156–168.

- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan**, “Consumer price search and platform design in Internet commerce,” *American Economic Review*, Forthcoming.
- Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan**, “Auctions versus Posted Prices in Online Markets,” *The Journal of Political Economy*, 2018, *126* (1).
- Farber, Henry S.**, “Is tomorrow another day? The labor supply of New York City cabdrivers,” *The Journal of Political Economy*, 2005, *113* (1), 46–82.
- , “Reference-dependent preferences and labor supply: The case of New York City taxi drivers,” *American Economic Review*, 06 2008, *98* (3), 1069–1082.
- , “Why you can’t find a taxi in the rain and other labor supply lessons from cab drivers,” *The Quarterly Journal of Economics*, 11 2015, *130* (4), 1975.
- Filippas, Apostolos, Srikanth Jagabathula, and Arun Sundararajan**, “Exit, voice, and availability: A randomized field experiment on sharing economy marketplace design,” *Working Paper*, 2018.
- Frechette, Guillaume, Alessandro Lizzeri, and Tobias Salz**, “Frictions in a competitive, regulated market: Evidence from taxis,” *Working Paper*, 2015.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” Technical Report, National Bureau of Economic Research 2018.
- Hall, Jonathan V and Alan B Krueger**, “An analysis of the labor market for Uber’s driver-partners in the United States,” *Industrial and Labor Relations Review*, 05 2018, *71*, 705–732.
- , **Cory Kendrick, and Chris Nosko**, “The effects of Uber’s surge pricing: A case study,” *Working Paper*, 2016.

- Hall, Robert E**, “Is Unemployment a Macroeconomic Problem?” *The American Economic Review*, 1983, 73 (2), 219–222.
- Harris, John R and Michael P Todaro**, “Migration, unemployment and development: a two-sector analysis,” *The American economic review*, 1970, 60 (1), 126–142.
- Hayek, Friedrich August**, “The use of knowledge in society,” 1945.
- Hsieh, Chang-Tai and Enrico Moretti**, “Can free entry be inefficient? Fixed commissions and social waste in the real estate industry,” *Journal of Political Economy*, October 2003, 111 (5), 1076–1122.
- Jensen, Robert**, “The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector,” *The Quarterly Journal of Economics*, 2007, 122 (3), 879–924.
- Thakral, Neil and Linh Tô**, “Daily labor supply and adaptive reference points,” *Working paper*, 2017.
- Varian, Hal R**, “Computer mediated transactions,” *The American Economic Review*, 2010, 100 (2), 1–10.
- Wooldridge, Jeffrey M**, *Econometric analysis of cross section and panel data*, MIT press, 2010.
- Zoepf, Stephen M, Stella Chen, Paa Adu, and Gonzalo Pozo**, “The economics of ride-hailing: Driver revenue, expenses and taxes,” *Working paper*, 2018.

A Proofs

Proposition 4 was not discussed in the main body, but it is useful for other proofs below.

Proposition 4. *The Pareto set is defined by $\phi'(x) = (p - c)/x$ and the Pareto set price is increasing in utilization, x .*

Proof. At points in the Pareto set,

$$\begin{aligned}\frac{du_{pax}/dx}{du_{pax}/dp} &= \frac{du_{drv}/dx}{du_{drv}/dp} \\ \phi'(x) &= \frac{p - c}{x}.\end{aligned}$$

The Pareto set price is increasing in x , as $p'(x) = (\phi''(x) + p/x^2) > 0$, as $\phi''(x) > 0$. □

A.1 Proof of Proposition 1

Proof. If we differentiate the market clearing condition by p , which is exogenous, we have

$$\frac{dp_{eq}}{dx_{eq}} = \frac{H + H'p - D'\phi'(x)}{D' - H'x} \leq 0, \tag{4}$$

as $\phi'(x) > 0$ and $D'(\cdot) < 0$. □

A.2 Proof of Proposition 2

Proof. The equilibrium that maximizes driver utility is

$$\max_{p,x} (p - c)x \quad \text{s.t.} \quad D(p + \phi(x)) = xH((p - c)x).$$

which gives us

$$\begin{aligned}\frac{dU_{drv}/dx}{ddU_{drv}/dp} &= \frac{dp_{eq}}{dx_{eq}} \\ -(p-c)/x &= \frac{dp_{eq}}{dx_{eq}}\end{aligned}$$

The drivers prefer an equilibrium where an increase in prices leaves their hourly earnings unchanged.

Proposition 4 already describes the Pareto equilibrium. The passenger preferred equilibrium maximizes the hours of transportation. It is equivalent to minimizing the part inside the demand curve, subject to the market clearing constraint, or

$$\min_{p,x} u_{pax} = p + \phi(x) \quad \text{s.t. } D(p + \phi(x)) \equiv xH((p-c)x).$$

which is satisfied when

$$\begin{aligned}\frac{du_{pax}/dx}{du_{pax}/dp} &= \frac{dp_{eq}}{dx} \\ -\phi'(x) &= \frac{dp_{eq}}{dx}.\end{aligned}$$

At the passenger's preferred equilibrium, a small increase in price has a disutility equal to the utility of marginally shorter wait time.

For the relationship between the possible points, what matters is the relative slopes of the Pareto curve and the possible equilibria curve. At the Pareto equilibrium, (p_P^*, x_P^*) , the slope of the possible equilibria curve is steeper than the passenger indifference curve:

$$|p'_{eq}(x)| = \left| \frac{Dx}{D' - xH'} + \phi'(x) \right| > |\phi'(x)|. \quad (5)$$

As such, there are lower price and higher utilization equilibria interior to the passenger's indifference curve at the Pareto equilibrium, and all higher price lower utilization equilibria are less preferred. Thus, the passenger's bliss point

has a lower price and a higher utilization relative to the Pareto point: $x_{pax}^* \geq x_P^*$ and $p_{pax}^* \geq p_P^*$. As the driver and passenger indifference curves are tangent at the Pareto equilibrium, from the driver’s perspective, there are higher price and lower utilization equilibria interior to their indifference curve at the Pareto equilibrium, and all lower price higher utilization equilibria are less preferred. Thus, $x_{drv}^* \leq x_P^*$ and $p_{drv}^* \leq p_P^*$. \square

A.3 Proof of Proposition 3

Proof. From Proposition 1, we know that regardless of the price, a fare cut increases utilization. From a p that is higher than the driver’s preferred price, p_{drv}^* , a reduction in fares moves drivers along the possible equilibria curve closer to their preferred equilibrium, raising their utility. This increases hours-worked, H . It also moves passengers closer to their preferred point, and so D increases. When prices are this high, a fare reduction is Pareto improving and there is no trade-off.

For $p < p_{drv}^*$ but $p > p_{pax}^*$ —prices between the driver and passenger preferred points—we are moving away from the driver’s preferred equilibrium, and so hours-worked decline, whereas we are moving close to the passengers preferred equilibrium, and so hours of transportation increase. In this range, fare changes are not Pareto improving. Finally, for $p < p_{pax}^*$, we are moving away from the preferred equilibrium of both types, and so both hours-worked and trips decline. In this range, a fare increase would be Pareto improving. \square

A.4 Wait times

Assume a constant returns to scale matching function, $m(\text{Cars}, \text{Pax's})$ between empty cars and passengers. The instantaneous probability that a passenger is matched is $m((1-x)H, D)/D$, and so the expected time until matching with a Poisson process in equilibrium is

$$\mathbb{E}[\Delta t] = m((1-x)/x, 1)^{-1}, \quad (6)$$

which depends solely on x . As such, $\frac{d\Delta t}{dx} > 0$, or wait times are increasing in utilization. Perhaps the assumption of constant returns to scale is not a good one—see [Arnott \(1996\)](#) for the notion that taxi services should be subsidized to reap scale economies—but other examples of online matching markets—even those with a substantial geographic component—seem to show constant returns to scale e.g., ([Cullen and Farronato, 2018](#)). [Frechette et al. \(2015\)](#) offers evidence of increasing returns to scale when NYC Taxi markets are not very busy, but essentially constant returns to scale during high usage times.

B Identification

B.1 Evidence for selection on observables

Despite the plausibility of city-specific forecasts driving decision-making, we view this as unlikely. Instead, the evidence is most consistent with Uber conditioning on observable attributes of a city available to us as researchers—namely the current level of utilization in a city.

Part of the evidence on Uber’s decision making comes from media reports.²¹ BuzzFeed News independently examined the spreadsheets Uber used to explore pricing and reported that these spreadsheets were not forecasting models at all, in that they “...don’t predict the true effects of price cuts” but rather, according to Uber, simply “simulate various scenarios that could happen.”

As best we can tell, the spreadsheets were intended to look at how a change in the price parameters would mechanically affect what a driver would have earned had they completed the same trips as before—and to calculate how many more trips a driver would have to provide to keep earnings per hour the same:

the spreadsheets seem to estimate how many more rides price cuts would *have to* (our emphasis) generate in order to keep gross driver earnings stable. But that increase in rider demand is not guaranteed.

There was apparently no forecast made in the spreadsheet models about what the effects of the price changes would have on demand or supply. Instead, this “forecasting” was actually Uber considering current conditions in the city—a data generating process that a suitably specified fixed effect panel model can accommodate.

Note that if Uber were conditioning on anticipated treatment effects—such as choosing cities with low utilization for fare cuts precisely because these cities

²¹<https://web.archive.org/web/20181227210723/https://www.buzzfeednews.com/article/carolineodonovan/uber-documents-suggest-price-cuts-dont-always-raise-driver-w> Uber did claim that these were not the only tools used to explore pricing, but we know of no other forecasting tools being used.

would have large treatment effects—strict exogeneity still allows us to identify the average effect, but the interpretation would be different than a case with homogeneous treatment effects. The estimated treatment effect in this case would be a weighted average of effects.²²

B.1.1 Evidence for selection on utilization for the January cuts

In the bottom row of Figure 6b, the size of each city’s fare cut is plotted versus the city’s utilization 3 weeks before the fare cut. This point is indicated with a black dot, which is also the base of an arrow. The tip of that arrow is that same city’s utilization 3 weeks after the fare cut. A loess smoothed line is fit through the initial utilization points. The 2015 January cut is on the left and the 2016 January cut is on the right. For both January periods, we can see that cities with lower utilization were clearly targeted for larger fare reductions. The lower the utilization, the larger the cut, on average.

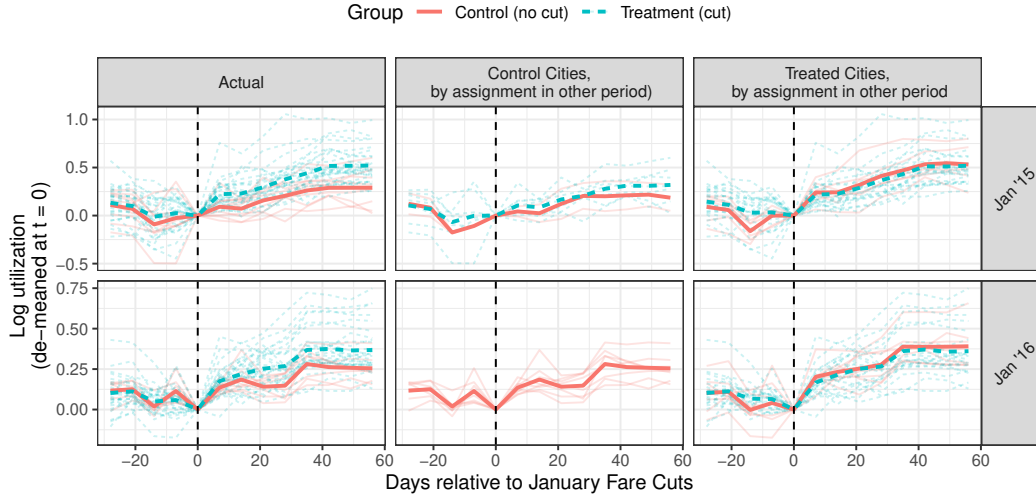
In the top panel Figure 6b, cities that did not receive January cuts are arranged on a number line by their pre-period utilization level. All have a fare “cut” of 0%, but to prevent over-plotting, they are arranged vertically by their utilization. We can see that the utilization of the cities not receiving cuts had substantial overlap with those cities having cuts, suggesting that selecting cities for cuts was not based strongly on utilization, even if the magnitude of the cuts was based on utilization, conditional upon receiving a cut.

By examining the length and direction of the arrows, we can see how utilization changed after the fare cut. We can see across both periods, cities with fare reductions almost always have increases in utilization. In contrast, cities without fare reductions, in the top panel, both increases and decreases are common (and are all fairly small). Among those cities that have fare reductions, we can see that cities with the largest fare reductions tend to have long

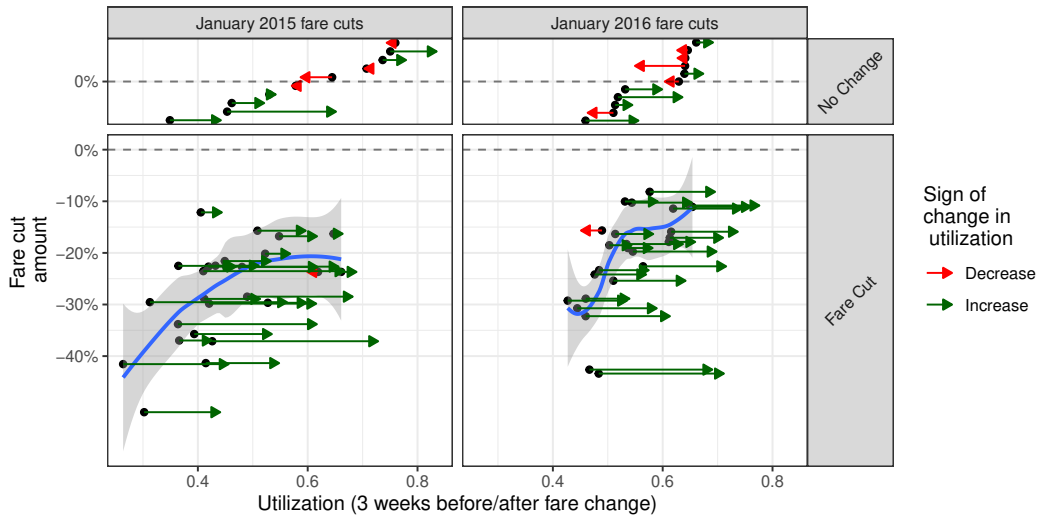
²²Our intuition is that estimates would be weighted averages of each city’s probability of being selected for a fare cut of that magnitude, though we have not explored this rigorously. At the binary treatment case, things are more straightforward. Consider two periods where the treatment is turned on for some units in the second period, and $y_{it} = \alpha_i + W_{it}\tau(\alpha_i)$, where $\tau(\alpha_i)$ captures the notion that treatment effects depend on the city-specific effect. If $Pr(W = 1|\alpha_i)$ is the probability of treatment, then $\mathbf{E}[\Delta y] = \mathbf{E}[\tau(\alpha_i)Pr(W = 1|\alpha_i)]$.

Figure 6: By-week city demeaned utilization rates around the two January fare cuts

(a) By-week city demeaned utilization rates around the two January fare cut periods



(b) Pre-fare change utilization and January cut magnitudes



Notes: In the left column in the top panel shows the actual demeaned utilization around the January fare cut week, with the solid red line indicating non-cut cities (the control) and the blue dashed line showing the treatment. Data are from the 2015 and 2016 January fare cuts, in the top and bottom rows, respectively. The bottom panel shows the January fare cut magnitude versus the utilization three weeks before the fare change (base of the arrow) and the utilization 3 weeks after (tip of the arrow) for both January fare cuts. This is a loess smoothed line in the scatter plot, based on the pre-change utilization. Those cities without a fare change are shown in the top panel.

arrows i.e., experience the largest increases in utilization.

Uber was clearly conditioning on utilization in determining the magnitude of cuts (and said as much publicly). But there is still clearly residual variation. The most parsimonious explanation for price variation in the data is that Uber was simply learning to price through experimentation. The company claimed as much when announcing fare cuts: “[w]e’ve learned over the years that we do best when we test new things. With each new test—small or large—we learn more about the choices riders make, and how those choices impact earnings for drivers.”²³

B.2 Parallel trends, with evidence from the January cuts

In our setting, where the treatment unit is whole markets, the main identification concerns mirror the concerns of the minimum wage literature. In the empirical minimum wage literature, an ever-present concern is that states that raise the minimum wage are experiencing rising economic fortunes, creating a spurious correlation between minimum wage levels and employment. In our empirical context, perhaps cities with already-increasing utilization were selected for fare cuts. In minimum wage studies, the typical empirical approach is to include unit-specific time-trends, but also to use a distributed lag model and then look for evidence of pre-trends (Allegretto et al., 2011). We can do both, and as we will show, there is no evidence of pre-trends for any of our outcomes.

A more graphical approach for assessing parallel trends is to simply plot city-specific means for the outcome around an “event.” In our setting we do not have a single treatment event—price variation is spread out over the entire panel. We can—and do—plot cumulative effects from a distributed lag model, but we can also use the large January cuts to explore, graphically, the parallel trends assumption in a model-free way. Despite many cities receiving fare

²³<https://web.archive.org/web/20181224143345/https://www.uber.com/newsroom/beating-the-winter-slump-price-cuts-for-riders-and-guaranteed-earnings-for-drivers/>

cuts, which we call “treated,” there are numerous cities that did not have fare changes, leaving some cities to serve as a control.

In Figure 6a, we plot the by-week utilization rates for all our cities, demeaned to 0 on the day of the cut, for both of the January periods. The two “rows” of the figure show data for the January 2015 and 2016 cuts in the top and bottom rows, respectively.

In the leftmost column, the actual by-week values are plotted for each city, by whether the city received a cut and is in the treatment (dashed) or did not and is in the control (solid). The averages for these two groups are also plotted in a heavy line. In the column labeled “Actual” we can see for both sets of cuts, (1) no evidence that, on average, treated and untreated cities were on different trajectories before the cut and (2) clear evidence of an increase in utilization for treated cities after the fare cut. This evidence of a treatment effect is consistent with what we observed in Figure 1 in our quasi-event study.

An obvious objection to this approach—despite no evidence of a violation in parallel trends—is that perhaps the cities that have cuts are selected, most likely on the basis of impending negative demand shocks due to weather (e.g., conditioning on ϵ_{it}). However, if we examine cities in Figure 2a that experienced January fare reductions, there is little evidence that cuts are universally weather-related. For example, in the first week of January 2015, we see no fare reductions in New York City, Boston, and Pittsburgh—not locations known for balmy winters—but large reductions in, among others, Tucson, Dallas, Houston, and Orange County. If we move forward one year to January 2016, New York City and Pittsburgh do get a fare cut and Boston does not; Dallas does not get a cut, but Houston and Orange County do get cuts.

Going beyond a qualitative assessment of how likely impending weather explains “treatment assignment,” we can also explore the paths of cities that differed in their assignment across the two January fare cuts, allowing for a kind of placebo test. The idea is that if changes in utilization were caused purely by selection related to weather, we should find spurious effects even when a city was not treated.

In the middle column of Figure 6a, the sample consists of only those cities

that were in the control in the focal year i.e., did not receive a cut. As before, the top rows show data from the January 2015 cuts and the bottom row shows the January 2016 cuts. In this column, the dashed line is the average for cities that were treated or will be treated in the *other* January fare cut. If those cities “naturally” were going to have a rise in utilization, we should see the dashed line rise as in the “Actual” column. In 2015, we see some slight evidence of an increase in utilization for the pseudo-treated in the post-period, but it is negligible and far less than observed in “Actual.” In 2016, there are no pseudo-treated cities.

In the right column, the sample is only those that were treated (i.e., had a cut) in the focal January, with solid line being the average for those cities that were in the control in the other January period. Again, we see no evidence of a “treatment effect” in either January period. There still could be some change that Uber could foresee unrelated to weather that caused or did not cause a fare change and that was related to future utilization, but there is no evidence that weather was a culprit.

There are other potential factors Uber could have been conditioning on. In the subsequent sections, we consider the potential role played by weather, competitors, and local economic conditions. In each case, we compare the point estimates to the baseline estimates from the main body. For convenience, the effects on the components of the driver gross hourly earnings rate are reported in Table 1.

B.3 Controlling for city-specific weather by season

If Uber was lowering fares in some cities expected to have a demand reduction due to weather, it would create a correlation between the base fare and demand (and hence many of our market level outcomes).

One approach to dealing with this concern is to include city-specific calendar-month fixed effects. E.g., there would be a fixed effect for New York City and January, which would be 1 for the New York City weeks during January 2014, January 2016 and January 2017. These fixed effects are intended to pick up

Table 1: Effects of fare changes on market outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.075 (0.064)	-0.715*** (0.067)	-0.208*** (0.035)
City FE	Y	Y	Y
City-specific linear trend	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R ²	0.794	0.842	0.472
Adjusted R ²	0.785	0.835	0.448

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The sample for each regression is the same, and is a city-week panel of UberX markets. See §3 for a description of the sample. The hourly earnings rate is the total earnings by drivers (excluding costs and Uber’s commission but not including any promotional payments) divided by total hours-worked, for drivers in that city, that week. Utilization is the total hours spent transporting passengers by drivers divided by the total number of hours-worked. Surge is the average value of the multiplier for all trips conducted in that city, during that week. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

city-specific differences in demand or supply due to different weather patterns.

Table 2 reports regressions for our main outcomes using city-calendar month fixed effects. All of the point estimates are very close to the original point estimates from Table 1. This casts doubt on the notion that Uber was conditioning on weather when deciding prices in a manner that simply led to a spurious correlation and the failure of the strict exogeneity assumption.

Table 2: Effects of fare changes on market outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.071 (0.068)	-0.715*** (0.067)	-0.214*** (0.036)
City FE	Y	Y	Y
City-specific linear trend	Y	Y	Y
City-Calendar Month FE	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R ²	0.801	0.850	0.492
Adjusted R ²	0.791	0.842	0.466

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2, but with the inclusion of calendar-month and city interactions. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The sample for each regression is the same, and is a city-week panel of UberX markets. See §3 for a description of the sample. The hourly earnings rate is the total earnings by drivers (excluding costs and Uber’s commission but not including any promotional payments) divided by total hours-worked, for drivers in that city, that week. Utilization is the total hours spent transporting passengers by drivers divided by the total number of hours-worked. Surge is the average value of the multiplier for all trips conducted in that city, during that week. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

B.4 Effects of ride-sharing competitors on main results

Another potential threat to identification is the actions of competitors. It is beyond the scope of this analysis to try to model the competition between ride-sharing platforms and the larger for-hire industry, or the broader transportation market. However, we can at least assess whether our panel results are sensitive to the presence of a substantial ride-sharing competitor. To do this, we interact our base price index with Uber’s share of the ride-sharing category.

Our estimates of Uber’s share come from the market research company “Second Measure,” which in turn uses credit card data. The reported measures are from each July, from 2014 to 2017. From these measures, we impute weekly measures matching our panel with a linear model. For cities in which no competitor was operating that week, we impute Uber’s share as 1.

In Table 3 we report our long-run regressions, mirroring our analysis in Table 1, though we leave out a price-specific trend to reduce variance in exchange for some (small) amount of bias. However, we first use the imputed Uber share as an outcome variable in Column (1). The coefficient is positive, large in magnitude but insignificant.

In the next columns, we report estimates for the hourly earnings rate, utilization and average surge. The base trip price index is interacted with the imputed Uber share of ride-sharing in that city that week. For all outcomes, the level of Uber’s category share has no detectable effect on the point estimate. The results suggest that the degree of direct rivalry in the market had no discernible effect on how Uber’s marketplace adjusted following fare changes.

Despite the possibility that direct competitors would matter, we have no evidence this is the case—interacting Uber’s imputed at-the-moment ride-sharing share with the price index has no detectable effect on the point estimates. This may simply reflect the fact that during the period covered by our analysis, Uber was the sole ride-sharing platform in many cities and held a dominant position in others, with Uber’s share of ride-sharing in the US being around 85% as

Table 3: Effects of fare changes on market outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>		
	Log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	-0.119 (0.102)	-0.892*** (0.104)	-0.224*** (0.029)
Uber share	0.017 (0.014)	0.010 (0.011)	0.003 (0.002)
Uber share \times Log base fare index	-0.005 (0.005)	-0.002 (0.004)	-0.001 (0.001)
City FE	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R ²	0.735	0.774	0.440
Adjusted R ²	0.725	0.765	0.419

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2, but with the base price interacted with Uber’s imputed market share. The base fare index is the price to passengers of an un-surfed, 6 mile, 16 minute trip in that city, in that week. The sample for each regression is the same, and is a city-week panel of UberX markets. See §3 for a description of the sample. The hourly earnings rate is the total earnings by drivers (excluding costs and Uber’s commission but not including any promotional payments) divided by total hours-worked, for drivers in that city, that week. Utilization is the total hours spent transporting passengers by drivers divided by the total number of hours-worked. Surge is the average value of the multiplier for all trips conducted in that city, during that week. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

late as January 2016.²⁴

B.5 Controlling for the city-specific unemployment rate

We might also be concerned that perhaps other economic shocks would matter, such as the local unemployment rate. Again, when we control for the city-specific monthly employment rate by MSA, we get no important change in our point estimates. In Table 4 we include each city’s MSA monthly unemployment rate as a regressor. All of the point estimates are very close to the original point estimates from Table 1. This casts doubt on the notion that Uber was conditioning on weather when deciding prices in a manner that simply led to a spurious correlation and the failure of the strict exogeneity assumption.

B.6 Static estimates with “short T” sub-panels

A concern with our empirical specification is that a single fixed effect and city-specific linear time trend is not sufficient to meet the strict exogeneity assumption. For example, cities with lower-than-expected utilization, given the fixed effect and linear trend, might be targeted for fare changes. One approach to explore this hypothesis is to divide the sample into shorter-T periods, but still include city-specific fixed effects. With this approach, it is more likely that strict exogeneity is met in each of the sub-panels.

We do this in Figure 7, reporting estimates for two smaller periods, $[0, T/2]$ and $(T/2, T]$. With these shorter panels, the city-specific linear time trend becomes harder to estimate and so we eliminate it. However, we keep the city-specific time trend for the full panel analysis, labeled “Full.”

Related to the concern about effect heterogeneity, there is a growing empirical interest in what can be identified by dynamic panel data models (Abraham and Sun, 2018; Athey and Imbens, 2018; Goodman-Bacon, 2018). A common thread in this literature is that problems arise when treatment effects vary

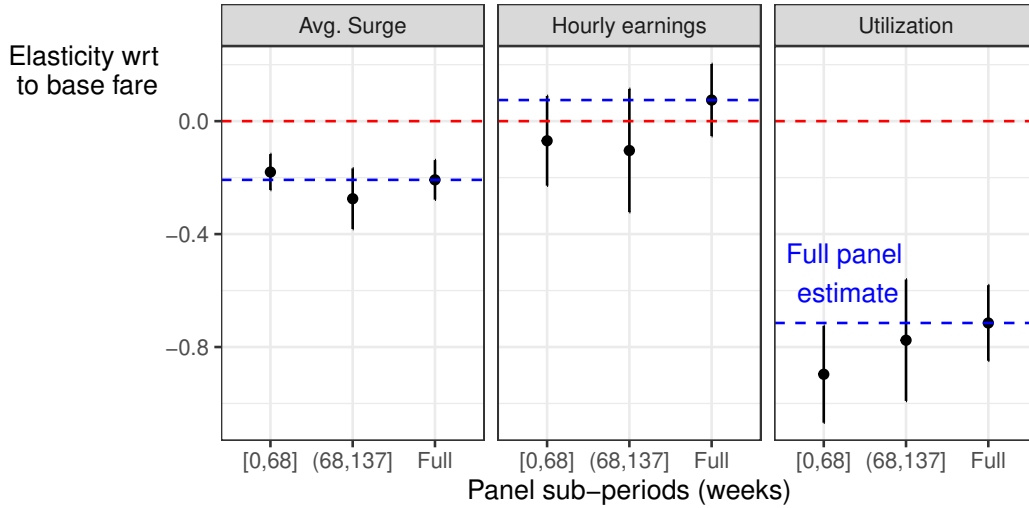
²⁴<https://web.archive.org/web/20181231153043/https://www.thestreet.com/markets/ipo/do-uber-recent-numbers-support-its-ipo-valuation-14783564>

Table 4: Effects of fare changes on market outcomes from a city-week panel of UberX markets, controlling for MSA monthly unemployment rate

	<i>Dependent variable:</i>		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.086 (0.068)	-0.697*** (0.069)	-0.209*** (0.036)
Unemployment rate	0.001 (0.014)	-0.003 (0.016)	0.004 (0.007)
City FE	Y	Y	Y
City-specific linear trend	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,816	4,816	4,816
R ²	0.787	0.841	0.475
Adjusted R ²	0.777	0.834	0.451

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2 but also includes as a regressor the MSA-level unemployment rate. The base fare index is the price to passengers of an unsurged, 6 mile, 16 minute trip in that city, in that week. The sample for each regression is the same, and is a city-week panel of UberX markets. See §3 for a description of the sample. The hourly earnings rate is the total earnings by drivers (excluding costs and Uber’s commission but not including any promotional payments) divided by total hours-worked, for drivers in that city, that week. Utilization is the total hours spent transporting passengers by drivers divided by the total number of hours-worked. Surge is the average value of the multiplier for all trips conducted in that city, during that week. Fixed effects are included for the city and for the week. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Figure 7: Point estimates for the effects of base fare changes with different panel lengths



Notes: This figure reports estimates similar to Table 1, but with the panel cut in half on the time dimension and the city-specific linear time trend removed. For comparison purposes, the point estimates from Table 1, which includes the time trend are shown. These estimates are labeled “Full” and are indicated with a horizontal blue dashed line.

by time and across units, creating estimates that are weighted averages of different effects.

For simplicity, this work typically considers a binary treatment that occurs at some point in an experimental unit and then stays “on” for the remainder of the panel, such as state-wide policy change. Our empirical setting is more complex, in that the independent variable—the base fare—is continuous and changes multiple times over the course of the panel. Despite our setting being different, we can at least assess some of the flavor of the concern about weighted combinations of effects by estimating effects using different sub-panels and seeing how the point estimates differ.

B.7 Testing the strict exogeneity assumption

One statistical test of the strict exogeneity assumption is to include a “lead” of our explanatory variable in our regression (Wooldridge, 2010),

$$y_{it} = \alpha_i + \beta_1 \log b_{it} + \gamma \log b_{it+1} + g_i t + d_t + \epsilon_{it}. \quad (7)$$

Under the null hypothesis of strict exogeneity, $\gamma = 0$. When we run this regression, using utilization as the outcome, $\hat{\gamma}$ is -0.087, with an SE of 0.065, giving a t-stat of -1.336 and a p-value of 0.182.

B.8 Heterogeneous effects by direction of fare change

As we saw in Figure 2a, our identifying variation includes both price increases and decreases. We might suspect that fare increases and fare decreases have different effects on a ride-sharing market. For example, price decreases are more likely to be heavily promoted by Uber than price increases, perhaps hastening their effects. On the driver side, to the extent we analogize the hourly earnings rate to a wage paid by a firm, there are good reasons to think fare decreases might elicit a different behavioral reaction (Bewley, 2009).

Despite reasons to suspect heterogeneous effects, there are counter-arguments. Kinked demand curves are typically hard to justify theoretically and perhaps even less so in our empirical context—it is not the case that passengers had years of constant prices around which to form reference points. Base fare changes are not uncommon in our data, as our short-run price changes due to surge.

Empirically, given the structure of our data, it is hard to use fare increases and decreases separately, at least over the entire panel. However, we can compare the effects of January cuts—which are, of course, only cuts—to the overall panel estimates. There is also a period in 2015 when nearly all the variation in prices was price increases.

We use these sub-panels with variation of all the same type and estimate our baseline panel model. Figure 8 reports estimates of Equation 2. Point

estimates are shown for our main outcomes of interest. For the sub-samples, we remove the city-specific linear time trend.

The samples are: (1) the full panel estimate, (2) windows around the January 2015 cuts (from 2014-12-08 to 2015-03-09), (3) windows around the January 2016 cuts (from 2015-12-07 to 2016-03-07), and (4) the “interior” of 2015 when all of the variation in prices were price increases. For (2) and (3), post-periods of different lengths are used, with the number of weeks of post-cut data included indicated above the error bars. For (4), we use from 2015-03-02 to 2015-06-29.

For the two January cuts, we use different post-period lengths. These lengths are shown above the top of the confidence interval, in weeks. They are in length order and a line connects them.

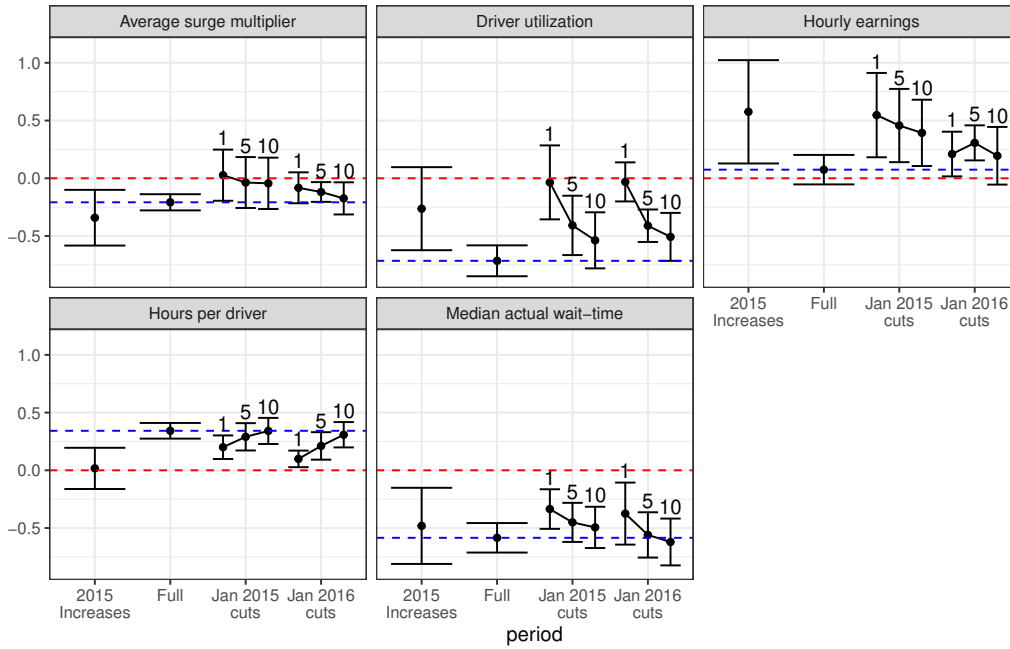
With these shorter panels, there is a clear loss of precision. However, we can see that all point estimate effects are directionally the same as those found in the full panel. We know there are dynamics to many of these outcomes, and so it is unsurprising that point estimates change with longer panels.

B.9 Alternative regression specifications

In the main body of the paper, we reported our preferred specifications for the various city outcomes. However, there was some freedom in this choice with respect to (1) whether city-specific linear time trends were included, (2) whether the pre-period cumulative effect was constrained to be zero and (3) the number of post-period lags to include. In this section, we report estimates of our effects using difference choices.

Figure 9 illustrates the pattern we use for all outcomes-. In the left column, the cumulative effects from Equation 3 are plotted. These just recapitulate the results from the main body of the paper. In the middle column, we report the same distributed lag model (DLM) results but remove the city-specific linear trends. In the right column, we report the same DLM, but without the city trend and without demeaning in the -1 period i.e., we do not impose the restriction that $\sum_{\tau=\text{NUMPRE}}^0 \hat{\beta}_{\tau} = 0$. In Figure 10, we report cumulative effects

Figure 8: Effects of a base fare increase on the driver hourly earnings rate and its components using different samples

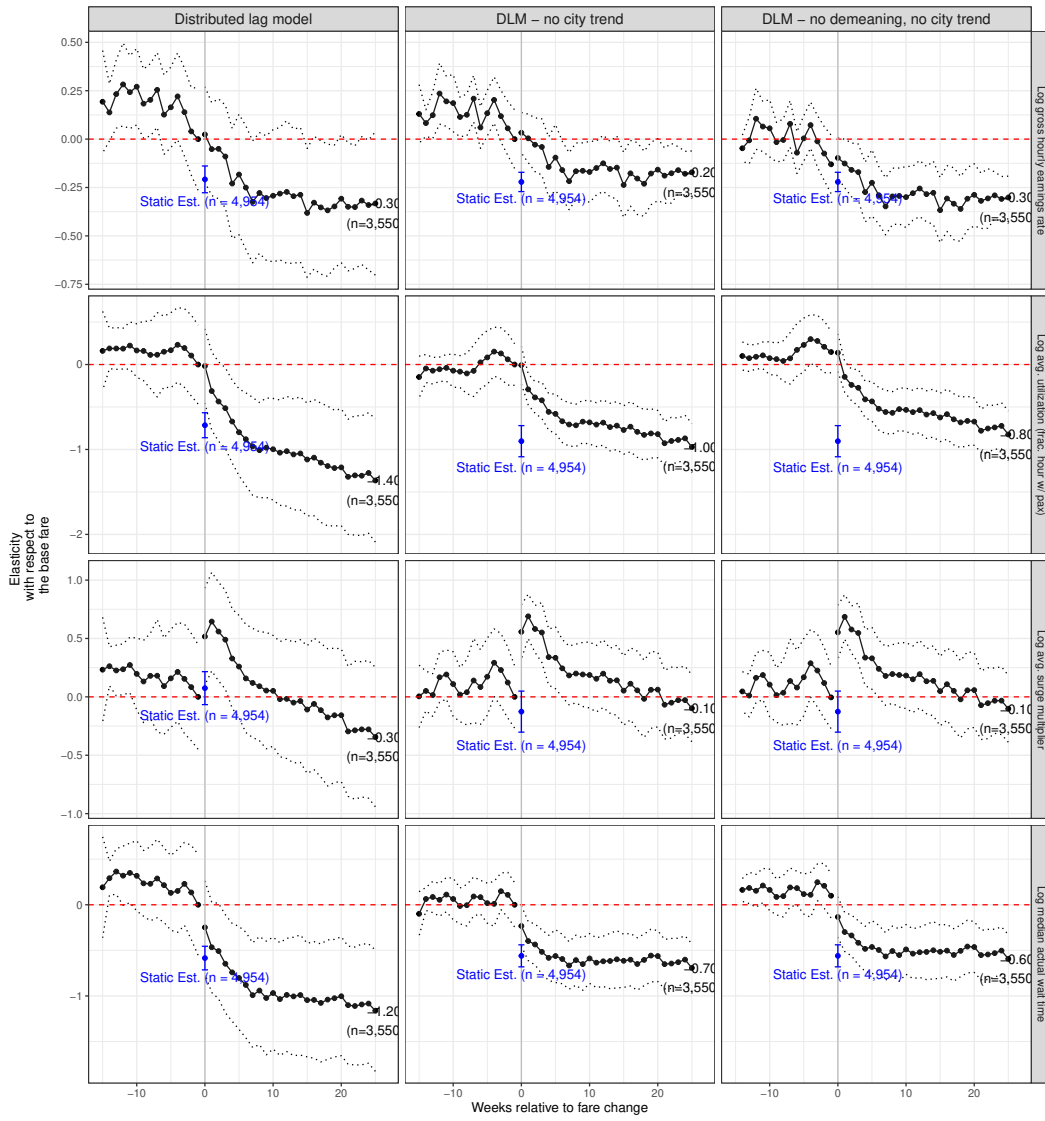


Notes: This figure reports estimates of Equation 2 using different sample definitions. The samples are: (1) the full panel estimate, (2) windows around the January 2015 cuts, (3) windows around the January 2016 cuts, and (4) the “interior” of 2015 when all of the variation in prices were price increases. For (2) and (3), post-periods of different lengths are used, with the number of weeks of post-cut data included indicated above the error bars. For (4), we use from 2015-03-02 to 2015-06-29.

with our preferred specification, but for a collection of post-period bandwidths.

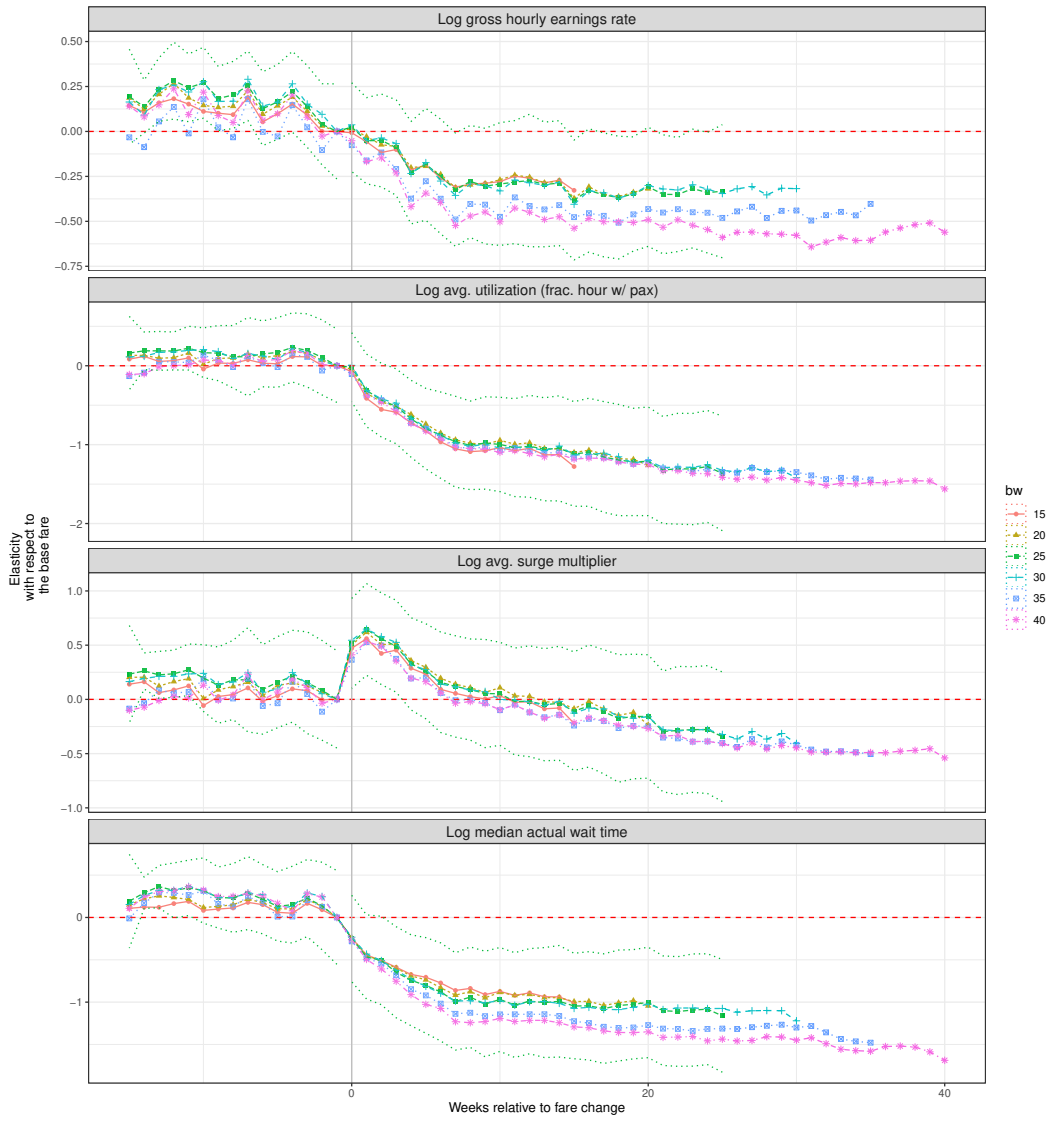
We present these same alternative specification/alternative bandwidth plots for all of our other main outcomes in the figures below.

Figure 9: Alternative specifications for Figure 5



Notes: Alternative specifications.

Figure 10: Alternative post-period bandwidths for Figure 5



Notes: Alternative post-period bandwidths.

C Tables

Table 5: Effects of fare changes on Dynamic Pricing/Surge outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>	
	Log avg. surge ...	
Log base fare index	-0.208*** (0.035)	
City FE	Y	
City-specific linear trend	Y	
Week FE	Y	
Observations	4,954	
R ²	0.472	
Adjusted R ²	0.448	

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Table 6: Effects of fare changes on Passenger-side Quantities outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>		
	Log total trips...	Log requested t...	Log total hours...
	(1)	(2)	(3)
Log base fare index	-0.099 (0.081)	-0.156* (0.082)	-0.236*** (0.074)
City FE	Y	Y	Y
City-specific linear trend	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R ²	0.989	0.988	0.990
Adjusted R ²	0.989	0.988	0.989

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Table 7: Effects of fare changes on Wait times outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>	
	Log median esti...	Log median actu...
	(1)	(2)
Log base fare index	-0.625*** (0.067)	-0.585*** (0.064)
City FE	Y	Y
City-specific linear trend	Y	Y
Week FE	Y	Y
Observations	4,954	4,954
R ²	0.854	0.889
Adjusted R ²	0.848	0.884

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Table 8: Effects of fare changes on Driver-side Quantities outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>			
	Log hours-worke...	Log total num. ...	Log total hours...	Log partner sig...
	(1)	(2)	(3)	(4)
Log base fare index	0.342*** (0.034)	0.137** (0.063)	0.479*** (0.064)	-0.122 (0.163)
City FE	Y	Y	Y	Y
City-specific linear trend	Y	Y	Y	Y
Week FE	Y	Y	Y	Y
Observations	4,954	4,954	4,954	4,954
R ²	0.913	0.993	0.992	0.953
Adjusted R ²	0.909	0.993	0.991	0.951

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Table 9: Effects of fare changes on Driver Productivity/Utilization outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>		
	Log num. trips ...	Log avg. utiliz...	Log avg. miles ...
	(1)	(2)	(3)
Log base fare index	-0.292*** (0.068)	-0.715*** (0.067)	-0.655*** (0.059)
City FE	Y	Y	Y
City-specific linear trend	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R ²	0.860	0.842	0.919
Adjusted R ²	0.853	0.835	0.915

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Table 10: Effects of fare changes on Driver Compensation outcomes from a city-week panel of UberX markets

	<i>Dependent variable:</i>			
	Frac. earnings ... (1)	Log gross hourl... (2)	Log hourly earn... (3)	Log gross hourl... (4)
Log base fare index	-0.123*** (0.029)	0.075 (0.064)	0.055 (0.069)	0.086 (0.208)
City FE	Y	Y	Y	Y
City-specific linear trend	Y	Y	Y	Y
Week FE	Y	Y	Y	Y
Observations	4,954	4,954	4,954	4,927
R ²	0.605	0.794	0.774	0.747
Adjusted R ²	0.588	0.785	0.764	0.735

Notes: This table reports OLS regressions of city-week outcomes on the log base fare index. The estimating equation is Equation 2. Standard errors are clustered at the level of the city. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.