Pricing Efficiently in Designed Markets: The Case of Ride-Sharing*

Jonathan V. Hall, Uber Technologies
John J. Horton, NYU Stern
Daniel T. Knoepfle, Uber Technologies

May 10, 2019

Abstract

In many designed markets, the platform eliminates price dispersion by setting the product market price, yet still allows free entry and exit. Using data from Uber, we show that the price profoundly affects the market equilibrium, but does not substantially change the hourly earnings rate of drivers in the long-run. Tracing out how the market re-equilibrates, we show that following fare increases, drivers make more money per trip—and initially more per hour-worked—and as a result, work more hours. However, this increase in hours-worked has a business stealing effect, with drivers spending a smaller fraction of hours-worked with paying customers. This reduced driver productivity, or "utilization" is the main reason the hourly earnings rate is unchanged. This price/utilization trade-off has implications for the efficiency of ride-sharing markets.

*Thanks to Josh Angrist, Susan Athey, David Autor, Erik Brynjolfsson, Peter Cohen, Keith Chen, Judith Chevalier, Will Dobbie, Arin Dube, Mark Duggan, Amy Finkelstein, Hank Farber, Andrey Fradkin, Ed Glaeser, Austan Goolsbee, Bob Hall, Steve Levitt, Alexandre Mas, Enrico Moretti, Haim Mendelson Ben Olken, Mandy Pallais, Paul Oyer, Jim Savage, Robin Yerkes Horton, and especially Jason Dowlatabadi for assistance, helpful comments and suggestions. 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 is an employee and shareholder of Uber Technologies and is John Horton's spouse.

1 Introduction

Market-designing platforms can use information technology to lower search costs, yet price dispersion has often proven stubbornly persistent in practice. Some platforms now take a more direct approach, eliminating price dispersion by simply setting the price faced by buyers on their platform. However, the platforms still typically use the market mechanism on the supply side of the market, allowing free entry of sellers, who in turn earn a fraction of receipts from the buyers they serve. With this hybrid structure, the platform's choice of a product market price affects both sides of the market, but in "opposite" directions—a higher price lowers demand and simultaneously increases supply, pushing the market out of equilibrium. How a new equilibrium emerges following a price change—and the implications this choice has for the functioning and efficiency of the market—is the focus of this paper.

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 base fare—or the cost of a ride under un-surged conditions. We use these changes to identify the effects of the fare on the market equilibrium. Our empirical design is essentially a difference-in-differences design with staggered treatments of varying intensity. As we will discuss at length, the assumptions required for causal inference in this panel setting are well-satisfied, as Uber's pricing decisions seem to be conditioned solely on market attributes that we observe.

In our analysis, a key outcome is the driver hourly earnings rate—essentially the price on the supply side of the market—and its determinants, particularly the fraction of driver hours-worked that are spent with paying customers. This fraction—which we call "utilization"—can be thought of as a measure of driver technical productivity. As we will show with a simple model, driver technical productivity is central to understanding the efficiency of a ride-sharing market.

For the driver hourly earnings rate, we find that when Uber raises the base fare in a city, the driver hourly earnings rate rises immediately as drivers make more money per trip.[1] However, the hourly earnings rate begins to decline shortly thereafter, and after about 8 weeks, there is no detectable difference in the driver average hourly earnings rate compared to before the fare increase. Measures of hourly earnings that include incentive payments from Uber and *direct* costs—though critically, not the cost of effort—show much less initial pass-through of fare increases and a somewhat negative long-run effect.

We can decompose the hourly earnings rate to understand, in an accounting sense, why there is little change. The main reason the hourly earnings rate stays approximately the same despite a higher per-trip payment is that the fraction of hours-worked that are spent with paying customers falls; a 10% fare increase lowers utilization by about 7%, using a "static" estimate of the effect. This single point estimate masks how the market actually evolves. Using a distributed lag model, we show that the decline in utilization is not immediate but rather emerges over the course of several weeks. The effect of a 10% fare increase is to eventually lower utilization by 10% in the long-run.[2] A second reason hourly earnings stayed roughly the same was that a price increase led to less frequent and/or more moderate surge pricing (Chen and Sheldon, 2015; Hall et al., 2016). 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. However, this effect on surge is relatively small, in that for a 10% fare increase, our static estimate is that the average surge rate falls by about 2%.

The fall in utilization following a fare increase is due to a combination of fewer trips demanded and drivers working more hours. The increase in hours-worked following a fare increase appears to be primarily an intensive margin response, with drivers working more hours but not exiting or entering the market at different rates.

In terms of magnitudes, for a 10% increase in the fare, drivers eventually work 6% more hours. Although it is tempting to view this response as a labor

[1] Although we frame our results as effects from a fare increase, we use both fare increases and decreases for identification.

[2] We use "long-run" here to mean the last by-week point estimate of the cumulative effect when using a distributed lags specification.

supply elasticity, it is important to note that we cannot separately identify the contribution of demand and supply to utilization changes, as our unit of analysis is the whole market. Furthermore, the price change is not a wage change, as the price change clearly affects both utilization and surge. This point that fare changes also alter utilization is critical, as we find that changes in utilization change observable driver costs (and likely change unobserved costs).

There are several reasons why higher utilization is costlier to drivers. Higher utilization leads to greater fuel expenditures, vehicle depreciation, and the opportunity cost of short breaks. It also simply requires greater driver effort compared to being without passengers—opening and closing doors; navigating busy streets accurately; dealing with luggage; ending personal phone calls; complying with passenger requests with respect to the radio, heat and AC, window position, and so on.[3]

To see the effects of these utilization costs in the data, we first note that when the hourly earnings rate increases immediately after a fare increase—but before utilization has fallen—drivers work more hours. This behavior is inconsistent with target-earning, at least in aggregate. As utilization begins to fall, pushing down the hourly earnings rate, hours-worked continues to *rise*. What appears to happen is that as utilization falls, driver costs fall as well, and so even though the driver hourly earnings rate is somewhat lower with a higher fare, this is offset by reduced costs to the driver.

Because higher utilization is costlier to drivers, the platform faces a trade-off in setting the product market price—the greater output from a higher utilization must be weighed against the cost. At the optimal product market price, a social planner would want a small increase in utilization to create

[3]That drivers have to work harder with a higher utilization has been a source of complaint by a recent NYC TLC policy change that led to mandated higher utilization. https://www.nytimes.com/2018/07/02/nyregion/uber-drivers-pay-nyc.html "The analysis also anticipates a net increase in pick-ups per hour, which would leave some possibly working harder for the same pay. But the analysis assumes car utilization will not intensify significantly, just a few minutes per hour, which would not offset the increase in earnings." https://web.archive.org/web/20181230200600/https://www.thenation.com/article/can-new-york-rein-uber/

output equal in value to the cost the additional output has to drivers. We show this formally, and demonstrate that the surplus-maximizing fare is also the lowest possible fare that can support a market equilibrium. If fares are lowered below this point, an equilibrium can only be restored through shifts in the demand curve, such as through degraded service quality.

In our model, passenger wait times are fixed. This allows us to focus our analysis on the economics of the price/utilization trade-off. Justifying this assumption to an extent is the existence of surge pricing, which raises prices to offset real-time imbalances of supply and demand that would otherwise cause the market to clear on wait-times (Hall et al., 2016). However, we find that wait times were not impervious to changes in the base price: for a 10% increase in the fare, median wait times fell by 6%. These changes have to be considered when assessing the efficiency implications of a fare change, though other evidence suggests passengers are relatively inelastic with respect to wait times compared to prices (Cohen et al., 2016).

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.[4]

[4]See: 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://blog.lyft.com/lyftline5?rq=price; 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.

The market adjustment process we uncover is fairly simple: when driving with Uber temporarily becomes a better deal, drivers work more hours and push down the hourly earnings rate through lowered utilization and somewhat less surge; the equilibration process runs in reverse when driving with Uber becomes a relatively worse deal. This adjustment process tends to push the driver hourly earnings to a fixed level, and so Uber faces a *de facto* horizontal labor supply curve with respect to the hourly earnings rate, at least within the range of fares and driver hourly earnings rates seen in our data. In short, the economic reason why the hourly earnings rate is insensitive to the product market fare is that the marginal driver has outside options that pin down the on-platform returns.

The market adjustment process we illustrate is similar to what is found in Hsieh and Moretti (2003), who show that real estate agent earnings are not, in the long-run (of 10 years) affected by house prices, despite agents being paid fixed, proportional commissions.[5] As in our paper, "product market" price changes leave the marginal product of labor unchanged because of a business-stealing decline in technical productivity (i.e., utilization). A key difference between our paper and Hsieh and Moretti (2003) is that higher technical productivity in our setting clearly imposes a cost on workers, both because of greater direct cost and the disutility of higher effort.[6] We also observe a clear trade-off in our data between prices and service quality as measured by wait times that was not present in their real estate agent setting. These cost and service quality changes are relevant to a market-designing platform trying to set an efficient price.

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).

[5] The 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. Equilibrium urban unemployment is, in a sense, a "between-worker" version of our within-worker utilization results.

[6] Hsieh and Moretti (2003) reasonably conclude agents are unlikely to face higher costs with additional hours-worked simply spent chasing leads and marketing.

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 quantity equilibrium, despite free entry on the supply side.[7] Centralized price-setting can avoid the wasteful monopolistic competition equilibrium.

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

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. UberBlack is the premium option, with newer, more luxurious cars and drivers that meet other conditions. UberX is the peer-to-peer option and is the largest and fastest-growing Uber rides product. It is also available in more cities than UberBlack. See Hall and Krueger (2018) for a discussion of the relative size of the two services. 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[8] At the end of

[7] 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.

[8] Uber'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

the trip, the fare is automatically charged to the passenger's credit card. Uber handles all billing, customer support, and marketing.

The price of a trip depends on a number of parameters set by Uber. There is a per-minute time multiplier and per-mile distance multiplier, as well as a fixed 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.[9] 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 (i.e., UberX or UberBlack), 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.

Near 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 end points. 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.

---

match remains strongly influenced by distance.

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

2.1 Measurement of hours-worked and driver hourly earnings rates

For our analysis, a primary outcome of interest is the average hourly earnings rate of drivers. To construct this measure 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 driver revenue, we omit reimbursements for known tolls and fees (such as airport fees), and deduct Uber's service fee. For our initial analysis, we do not calculate drivers' costs, and so it is important to regard our measure of the hourly earnings rate as a gross flow to both the driver's labor and capital. Later, we use a measure of net driver hourly earnings constructed by imputing costs. We use data on how changes in utilization affect miles-driven, and hence direct costs from fuel and wear and tear.

Drivers are eligible for promotional payments that typically depend on meeting various goals, such as a number of rides provided in a week. For our initial analysis, we do not include any driving-related promotional payments in our definition of the hourly earnings rate. When we do explore the effects of promotional payments, we allocate the payments as earnings in the week in which they were paid. Because of the non-linear nature of these promotional payments, measures of hourly earnings that include these measures do not reflect the true marginal earnings rate a driver would face. Some promotional payments unrelated to driving, like those earned for referring another driver, are omitted.

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 trip) is measured essentially without error, aside from rare technical glitches

(Varian, 2010). However, 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 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.[10] 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."[11]

3 Model

As the empirical analysis will make clear, there is a market-level trade-off between the price charged for a ride and the level of driver utilization: the higher the price, the lower the utilization. To explore this trade-off from a welfare standpoint, we develop a simple model.

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

[10] Such "multi-homing" behaviors lead to double-counting of some hours-worked (while the driver is waiting for dispatch) but undercounting 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 possibilities 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 §. For a theoretical analysis of the effects of multi-homing in ride-sharing markets, see Bryan and Gans (n.d.).

[11] For 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.

search process generates, such as non-existent/multiple equilibria or industry scale economies e.g., Arnott (1996) and Cairns and Liston-Heyes (1996). These models typically do not focus on driver utilization. There 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). In the "old" models of taxi markets, passenger demand depends on wait times, which adds substantial modeling complexity. We pursue a simpler modeling approach of a single demand curve, in part because of the presence of dynamic pricing.[12] However, so long as changes in wait times are assumed to be marginal, incorporating wait times into a welfare calculation is straightforward.

Our treatment of driver labor supply is also 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.[13]

There is a product market price, p, which generates demand $D(p)$ hours of the service, with $D'(p) < 0$. The service is produced by drivers, collectively working H hours. These hours of work are transformed into the service at a rate of x, where equilibrium x is the market level utilization, or, equivalently, driver technical productivity. The individual driver has utility $U(w, x) = w - c(x)$, where w is the equilibrium hourly earnings rate and $c(x)$

[12]See Hall et al. (2016) for evidence on the role of Uber's surge pricing in clearing the market when demand spikes. 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.

[13]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 hours measurement 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 present evidence that Uber driver's are responsive to hourly earnings in a neo-classical fashion and that there is little evidence of income targeting. Also using data from a Uber, Angrist et al. (2017) also find no evidence of income targeting.

is the disutility of utilization, with $c'(x) > 0$ and $c''(x) > 0$.[14]

We assume that the supply of hours-worked depends on driver utility, and so we write $H(U)$, with $H'(U) > 0$ and assume $H(0) = 0$. It is useful to think of $H(\cdot)$ as a uncompensated labor supply curve with $c(\cdot)$ being a dis-amenity, capturing both the direct and indirect costs of utilization.[15] We set aside any fee taken by the platform and assume that drivers claim all receipts, and so their gross hourly earnings rate is $w = px$. The total supply of the service when the product market price is p and the utilization is x is $S(p; x) = xH(U(px, x))$. Market clearing requires that

$$D(p^*) = S(p^*; x^*). \tag{1}$$

Figure 1 shows a market in equilibrium, labeled as point **A** in the figure, with a price of p^* and a quantity of transportation services Q_0.

3.1 Comparative statics of a fare change

Suppose the platform lowers the price by $\mathrm{d}p$. At this lower price, $\mathrm{d}p|D'(p)|$ more hours of the service are demanded. However, drivers now make less per hour-worked, with utility falling by $\frac{\partial U}{\partial p}\,\mathrm{d}p = x\,\mathrm{d}p$, which in turn lowers hours of the service provided by $xH'(U)\frac{\partial U}{\partial p}\,\mathrm{d}p$. The double-headed arrow in Figure 1 shows the gap in hours of the service that must be closed for the market to return to an equilibrium.

In the model, equilibrium could be restored through an increase in utilization, with a higher x pushing out the supply curve. In an actual ride-sharing market, there are several other ways the market could return to an equilibrium. One possibility is that surge pricing could increase, essentially "undoing" the fare decrease. Another possibility is that service quality could degrade, pulling in the demand curve. In the ride-sharing context, this service quality reduction would likely be an increase in passenger wait times, though drivers could also

[14] It is possible that at a sufficiently low x, drivers could be bored, making $c'(x) < 0$.

[15] Our assumption that $H(\cdot)$ is monotonically increasing rules out a backward bending labor supply curve, say due to income effects. Given that most drivers are part time (outside of NYC), ignoring income effects seems like a reasonable simplification.

Figure 1: Equilibrium in a ride-sharing market after a decrease in the base fare

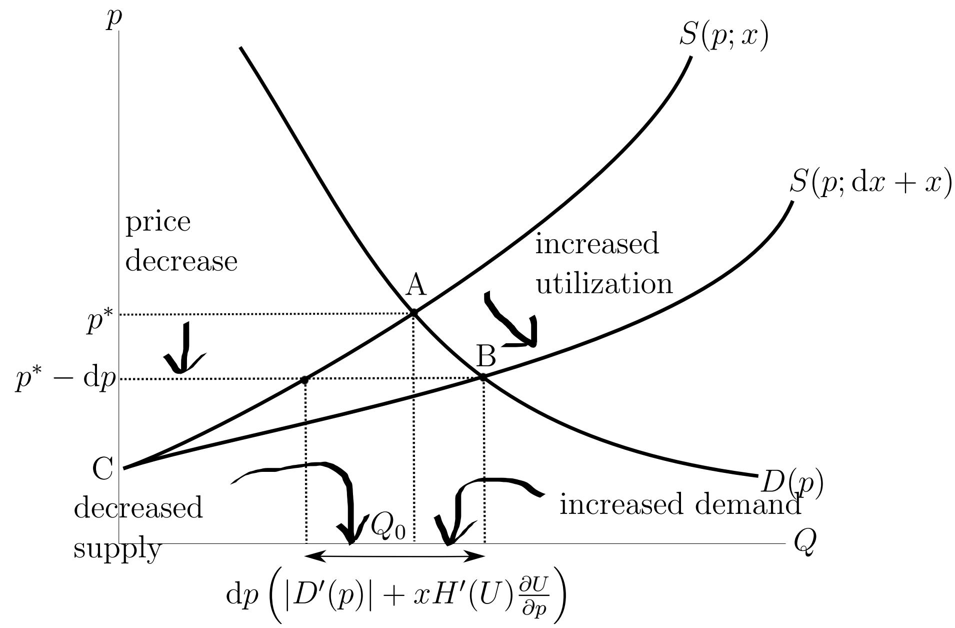

Notes: The $D(p)$ curve is the market demand for hours of transportation service, when the price, or fare, is p. The $S(p;x)$ curve is the hours of transportation supplied when the price is p and the driver utilization level is x. Q_0 is the market equilibrium quantity of hours of transportation service before a fare change. This figure illustrates the effects of a reduction in price from p to $p-\mathrm{d}p$.

provide less effort at customer service, or lower quality drivers could replace better quality drivers. We will discuss these other empirical possibilities in §5.

If the market is brought back to equilibrium purely through increases in utilization, then the supply curve shifts out to $S(p;\mathrm{d}x+x)$, where $\mathrm{d}x$ is the change in utilization. This brings the market to a new equilibrium indicated at **B** in Figure 1. Market clearing requires

$$\mathrm{d}p\left(|D'(p)|+xH'(U)\frac{\partial U}{\partial p}\right)=\mathrm{d}x\left(H+xH'(U)\frac{\partial U}{\partial x}\right),$$

and so the change in utilization with respect to the product market price is, in elasticity terms,

$$|\epsilon_p^x| = \frac{|\epsilon_p^D| + |\epsilon_U^H||\epsilon_p^U|}{1 - |\epsilon_U^H||\epsilon_x^U|}. \tag{2}$$

In this elasticity formula, the numerator quantifies the "gap" opened up by a fare change, while the denominator quantifies how it is closed through changes in utilization—the direct shifting of the supply curve (the "1" term) and the decrease in labor supply from the effects of the decreased utilization at a higher price (the $|\epsilon_U^H||\epsilon_x^U|$ term).

To visualize how changes in utilization clear the market, it is useful to have a figure where x is shown explicitly, rather than implicitly as in Figure 1. In Figure 2 we plot the gap between demand and supply versus the utilization, for a collection of product market prices. The curves, which are $\Delta(x;p) = D(p) - S(p;x)$, are drawn for product market prices of p_L, p_M and p_H, with $p_L < p_M < p_H$. Note that when $x = 0$, there is, mechanically, no supply and so the gap is simply the demand at that price. For all curves, as x begins to increase, the gap narrows, so long as $H + xH'(U)\frac{\partial U}{\partial x} > 0$.

For the $\Delta(x;p_H)$ curve, it first crosses the y-axis at **a**, corresponding to a market equilibrium. As x increases further, eventually the curve reaches an inflection point at **b**, where $H + xH'(U)\frac{\partial U}{\partial x} = 0$. This is the point where further increases in utilization actually lower the supply of the service because the direct effect of higher utilization "outside" the curve i.e., $\mathrm{d}xH$, is offset by the reduction in labor supply "inside" the curve. At **b**, there is excess supply, and so a further increase in x, which lowers output, is needed to clear the market. To the right of **b**, utilization increases are now *lowering* total output until we reach another equilibrium at **c**. As drawn, the $\Delta(x;p_H)$ curve—and all other curves—eventually turns horizontal, at the x where $px - c(x) = 0$ i.e, no drivers supply any hours because $U(px,x) \leq 0$.

Now consider the low price $\Delta(x;p_L)$ curve. Even at the maximum possible supply at **d**, demand exceeds supply, and so a price of p_L cannot support an equilibrium. At a higher price p_M, the $\Delta(x;p_M)$ curve, as drawn, has an inflection point that is just tangent to the y-axis, at **e**. This makes p_M the

Figure 2: Gap between demand and supply as a function of utilization, for different product market prices

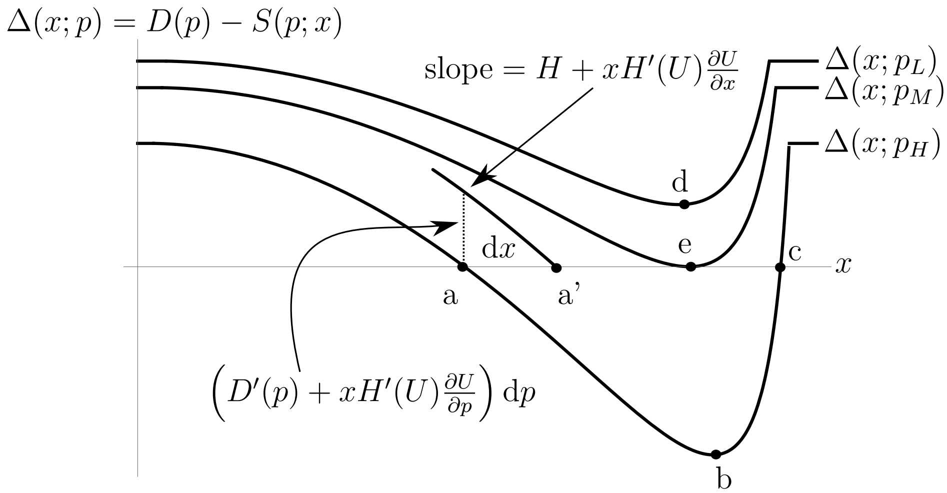

Notes: This figure illustrates the gap between demand and supply of the service at different levels of utilization, for a given product market price. Three prices are shown: p_L, p_M, and p_H. The change in the equilibrium from point **a** to **a'** corresponds to a decrease in the product market price from p_H to $p_H - \mathrm{d}p$.

lowest possible price for which the market clears, as any lower price shifts the $\Delta(x;p)$ curve up and so no equilibrium will exist. As one might intuit, the **e** equilibrium with price p_M has some attractive welfare properties—something we will show more directly below.

The interesting comparative statics of the equilibrium are when the market is at the **a** or **c** equilibrium, which are also the most likely equilibria in an existing market—the **e** equilibrium is unlikely, as it requires getting the price just right, and so there is no equilibrium associated with p_L. Assume that $p = p_H$ and that the market is at the **a** equilibrium. The interesting comparative statics are a price decrease from this level. A price decrease raises the $\Delta(x;p_H)$ curve by both the direct demand effect and the supply effect from the lower hourly earnings rate—it is the same double-arrow gap labeled in Figure 1. This brings the market to a new equilibrium, indicated by **a'**, which has a dx higher utilization.

Graphically, we can see that when the initial gap between supply and demand is larger because of a fare decrease—the vertical distance from **a** to the new curve—a bigger change in utilization is needed to restore equilibrium. This is why ϵ_p^D, the demand elasticity, "shows up" in the numerator of Equation 2. In contrast, the elasticity of supply cuts both ways—highly elastic supply causes the initial gap after a fare decrease to be larger, but the gap is more easily closed with smaller reductions in utilization because $\Delta(x;p)$ is steeper in x. This is why ϵ_U^H shows up in both the numerator and denominator of Equation 2.

Note that at the **c** equilibrium in Figure 2, lowering the price causes the utilization to *fall*. Even though a lower price increases demand—and a lower x exacerbates the problem—at this equilibrium, utilization is very costly to drivers and so the lower utilization induces a large labor supply response. Empirically, the sign of the change in utilization following a fare increase should indicate which equilibrium a market is in. To preview our results, we find that fare increases lower utilization on average, consistent with markets being in the **a** equilibrium (and not in a **c** equilibrium).

3.2 Efficient product market price

As the platform sets p, a natural question is what p is efficient in this model? In Figure 1, lowering p increases surplus by the triangle **ABC**, suggesting that the lowest possible p is efficient. From Figure 2, this lowest possible p was p_M, which allowed for an equilibrium at point **e**.

We can show that p_m is surplus-maximizing by computing the area of the **ABC** triangle from a marginal increase in utilization. At this optimal p_m, no further increase in surplus should be possible, meaning the area of **ABC** should vanish. Let $\triangle ABC(\hat{p}; \mathrm{d}x)$ be the area the triangle when the market price is $\hat{p}$ following a dp price decrease that causes a dx increase in utilization. As drawn, the pushing out of the supply curve creates a triangle wedge, though it is possible that the **C** could move up, causing $S(p; x + \mathrm{d}x)$ to cross $S(p; x)$ and so not be a triangle. The real quantity we care about is $\int_0^{\hat{p}} S(p; x + \mathrm{d}x) - S(p; x) dp$. If we differentiate this quantity by $\hat{p}$ and solve the first order condition, we have

$$\begin{aligned}\frac{\mathrm{d}\triangle ABC(\hat{p};x)}{\mathrm{d}\hat{p}} &= \frac{\mathrm{d}}{\mathrm{d}\hat{p}}\int_0^{\hat{p}}\left(H + xH'(U)\frac{\partial U}{\partial x}\right)\frac{\mathrm{d}x}{\mathrm{d}\hat{p}}\,\mathrm{d}p \\ &= \left(\frac{\mathrm{d}x}{\mathrm{d}\hat{p}} + \hat{p}\frac{\mathrm{d}^2 x}{\mathrm{d}\hat{p}^2}\right)\big(H + xH'(U)[\hat{p} - c'(x)]\big) = 0.\end{aligned} \tag{3}$$

We can see that surplus is maximized either when further changes in p cannot change the utilization, $\frac{\mathrm{d}x}{\mathrm{d}\hat{p}} + \hat{p}\frac{\mathrm{d}^2 x}{\mathrm{d}\hat{p}^2} = 0$, or when

$$p^* = c'(x) - \frac{H(U)}{xH'(U)}. \tag{4}$$

The $\frac{\mathrm{d}x}{\mathrm{d}\hat{p}} + \hat{p}\frac{\mathrm{d}^2 x}{\mathrm{d}\hat{p}^2} = 0$ solution reflects an equilibrium where utilization cannot be increased further, even if drivers would prefer it—utilization is capped at 1 and technical matching frictions might make the highest possible utilization lower than 1.

The interior solution described by Equation 4 is the price at which any further change in prices would change hours-worked by the same amount, in percentage terms. In words, at the optimal price, any further decrease in price cannot increase total output. The condition for the optimal price given in Equation 4 is the same that characterizes the minimum of the $\Delta(p;x)$ curves from Figure 2. To add the cost of increased wait times, we would need to add a $\int_{\hat{p}}^{\infty}\frac{\partial D(p,\underline{t})}{\partial \underline{t}}\frac{\mathrm{d}\underline{t}}{\mathrm{d}x}\frac{\mathrm{d}x}{\mathrm{d}\underline{t}}$ term to Equation 3, which would give a wait time aware welfare calculation.

3.3 Driver labor supply response

The driver labor response to changes in the base fare depends on how the new equilibrium changes their utility. As such, it is useful to indicate market equilibria relative to a collection of driver indifference curves. To do so, we ignore the fact that in any actual equilibrium, $w = px$ mechanically, and imagine that w and x could be set independently.

Figure 3 shows snippets of several driver indifference curves, in light gray, with w on the y-axis and x on the x-axis. In this (w, x) space, driver utility is increasing up and to the left, corresponding to a high hourly earnings rate and a low utilization. The snippets of indifference curves are drawn such that $u_0 > u_1 > \ldots u_n$. As higher utilization is costly, an increase in x requires a corresponding increase in w to keep drivers indifferent. Along an indifference curve, the slope is simply $c'(x)$, as utility is additively separable by assumption. As the cost is assumed to be convex in x, at higher values of x, an indifference curve gets steeper.

Figure 3: Optimal product market price and utilization

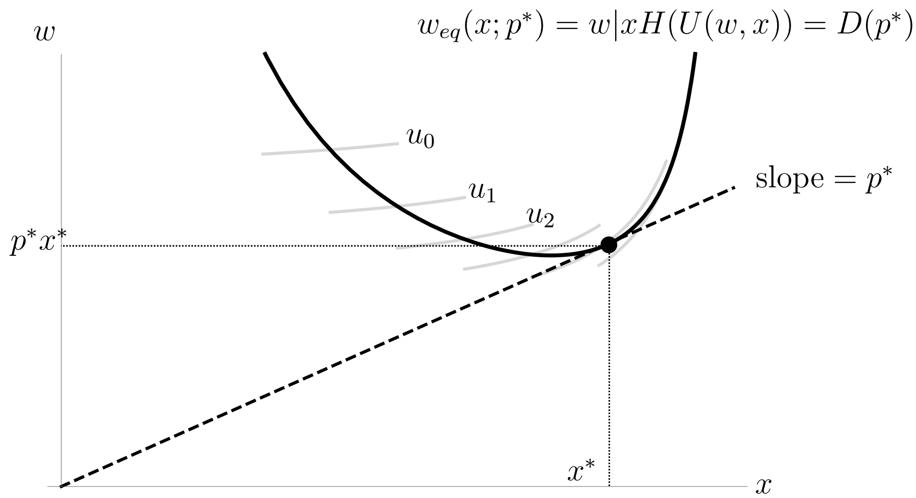

Notes: This figure illustrates shows several driver indifference curves with respect to the hourly earnings rate and the utilization. The heavy line, $w_{eq}(x; p^*)$, illustrates combinations of w and x that would clear the market for a given price p^*.

A given indifference curve with drivers at utility u corresponds to some fixed number of hours-worked, $H(u)$. For a given product market price, demand can be met with different numbers of drivers working at different levels of utilization. Let $w_{eq}(x; p)$ be the associated hourly earnings rate for an equilibrium with price p, at utilization x such that $D(p) = xH(U(w_{eq}(x; p), x))$. In Figure 3, this $w_{eq}(x; p)$ curve is plotted in a heavy line, for a price p^*. For a fixed amount of demand, with an increase in x, fewer hours of work are needed, and so we expect utility to decrease in x. Indeed, if u^* is the driver utility associated with a point on the $w_{eq}(x; p^*)$ curve,

$$\frac{\partial u^*}{\partial x} = \frac{-H(u^*)}{xH'(u^*)} < 0.$$

In contrast, $w_{eq}(x; p^*)$ is not monotonic in x, with

$$\frac{\mathrm{d}w_{eq}}{\mathrm{d}x} = c'(x) - \frac{H(U)}{xH'(U)}. \tag{5}$$

In this expression, because of the assumed convexity of $c(x)$, for low values of x, an increase in utilization raises costs very little, and so a relatively large reduction in w is needed to lower hours-worked. However, as x increases, eventually marginal costs rise enough so that $w_{eq}(x; p^*)$ starts to flatten out, and after the inflection point on $w_{eq}(x; p^*)$, a higher utilization actually requires a higher wage.

Note that Equation 5, combined with Equation 4, which characterizes the efficient price, implies that

$$\frac{\mathrm{d}w_{eq}}{\mathrm{d}x} = p^*, \tag{6}$$

or that the $w_{eq}(x; p^*)$ curve is tangent to a line with a slope that is the product market price. Intuitively, this is the point where a small increase in utilization, $\mathrm{d}x$, creates additional output of value $p\,\mathrm{d}x$, but at a cost to the driver of $\mathrm{d}w$.

Note that for this tangency point in Figure 3 to actually be an equilibrium in which $w = p^*x^*$, the line with slope p also has to intersect the origin. It is clear from the figure that at the point of tangency, any lower value of p could not support an equilibrium, as a lower p shifts the w_{eq} curve up, while reducing the slope of the p curve tangent to $w_{eq}(x; p^*)$.

Returning to our question about the comparative static effects of a fare change on driver labor supply, consider again a market at equilibrium at **a** in Figure 2. For the sign of effects on the hourly earnings rate, all that matters is whether $|\epsilon_p^x| > 1$. If $|\epsilon_p^x| > 1$, then a fare increase lowers wages—a 10% increase in fares is offset by a more than 10% decline in utilization. For utility, it increases from a fare increase if $x\,\mathrm{d}p > \mathrm{d}x\,\big(p - c'(x)\big) > 1$. Because of the costs of higher utilization, a higher bar is required, with $\epsilon_p^x > p/(p - c'(x)$, which is always greater than 1. If $c'(x)$ is very large, then it is difficult for a fare increase to lower utility (and hence lower hours-worked).

4 Data

Our empirical goal is to estimate the causal effect of Uber's fare changes on various market outcomes. We first describe the data and present some graphical evidence on the effects of fare changes. We then present panel regression estimates followed by an assessment of the conditions for causal inference. We then switch to a distributed lag model for the remainder of our empirical analysis so we can explore market dynamics.

4.1 Description of the 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 a 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.[16] 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

[16] The 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 includes several markets that were not very mature in the start of the panel.

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.

4.2 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 4 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 4 shows several things. First, the hourly earnings rate time series shows no obvious trend, 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 then are more or less constant afterwards.[17]

The patterns shown in Figure 4 around the January cuts preview some of the main results from our regression analysis—namely little change in the hourly earnings rate despite large changes in the base fare index. Immediately after the January fare cuts, 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.

[17]Note that we use the median wait time in a city week as opposed to the mean, as wait times are trip level measures and subject to outliers.

Figure 4: Average UberX market attributes over time for the US city-week panel, as indices

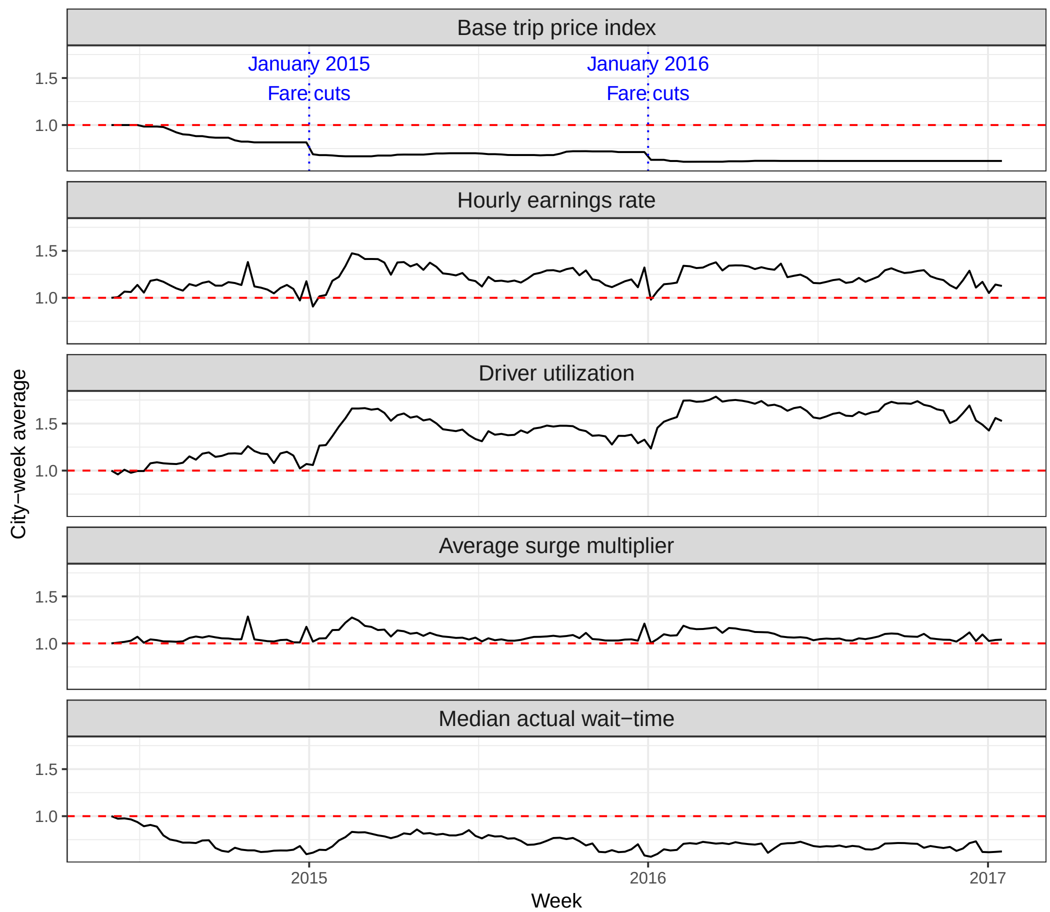

Notes: This figure plots the city-week panel weekly average for a collection of UberX market outcomes. All cities are weighted equally—see §4.1 for a definition of the sample. All series are turned into an index with a value of 1 in the first week. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The 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. Surge is the average value of the multiplier for all trips conducted in that city, during that week. Utilization is the total hours spent transporting passengers by drivers divided by the total number of hours-worked. The wait time for trip is the elapsed time from when a passenger requested a trip to when he or she was picked up.

4.3 Variation in base prices by city

As we saw in Figure 4, there were two large drops in January. However, Uber has changed the base fare for UberX in every city in the panel multiple times over the course of the panel. All these fare changes are shown in Figure 5a, with changes annotated in the week in which the UberX base fare index changed for the panel cities and reports the size of the change. 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.

Figure 5b plots the histogram of all fare changes in the panel. We can see that 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 local teams responsible for the various Uber markets. 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 deciding on fare changes—primarily driver utilization but also average surge and hourly earnings rates. 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.

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

Figure 5: UberX base fare index changes for US cities

(a) By-week changes

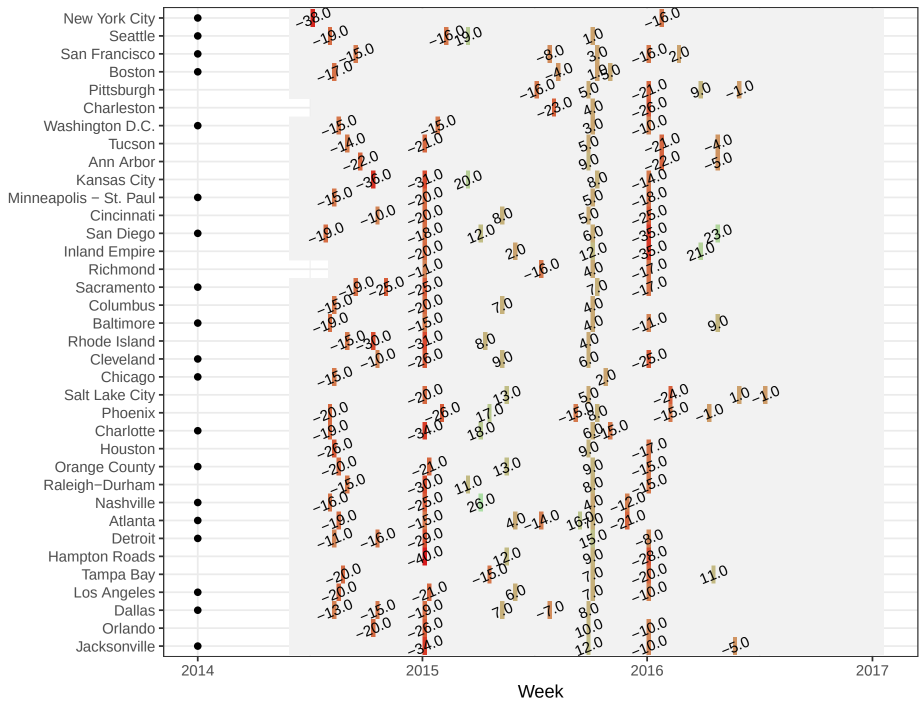

(b) Distribution of fare changes

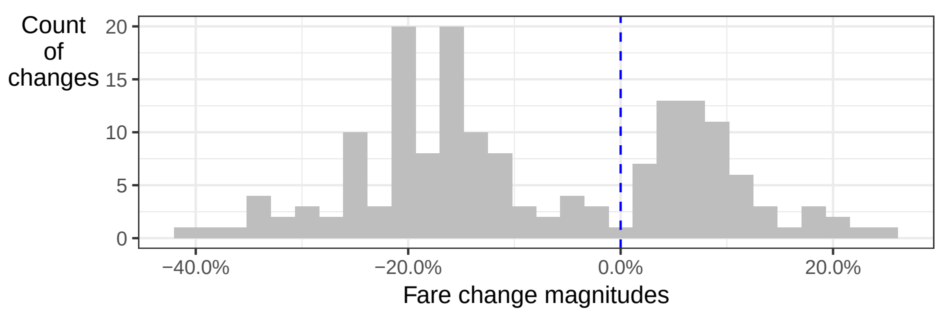

Notes: The top panel (a) indicates which cities in the panel had changes in the base trip price index, by week, and reports the size of that change, in percentage terms relative to the fare index in the previous week. Squares that are not shaded gray indicate that no data is available for that city week. See § 4.1 for a definition of the sample. The bottom panel (b) plots the histogram of all fare changes in the panel.

changes, it is not the case that latent differences exist between the kinds of cities that have fare changes and those that do not.[18] Second, Figure 5a 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 changes are nearly simultaneous, they are not perfectly simultaneous, and there is evidence of staggered roll-outs 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.

In addition to time-series variation in the timing of fare changes, there is also substantial variation in the magnitude of fare changes. As we will show, at least for the January cuts, much of this variation is explainable by the level of utilization in the city when the cuts were contemplated. With our regression specification, this kind of conditioning on observables is not a threat to identification. Furthermore, the more we can characterize what Uber was conditioning on, the lesser the worry that they were conditioning on something else we as researchers do not observe.

5 Results

In this section, we present regression estimates of the effects of changes in the base fare on market outcomes. Note that unlike in our model setup, 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. If the market can clear entirely though changes in utilization, then this distinction does not matter, as $m = 1$ and the demand curve remains fixed.

Before presenting our results, we first discuss the various threats to identification. Much of this discussion references appendices where we actually perform the analysis described—this section just provides the conclusion of

[18] 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). We will explore this point at length.

these various diagnostic checks.

For our actual analysis, we being by focusing on the driver hourly earnings rate—w in the model—and its factors, utilization and the surge multiplier, which we will regress on the base trip price index.[19] We then move to a distributed lag model which allows for a better representation of market dynamics. We then consider the role of driver costs, elucidating the role played by $c(x)$, and how the changes affect the driver labor supply response. We conclude by estimating the effect of the fare changes on market quantities, such as hours-worked and trips-taken.

Although we will later use a more complex specification that can account for dynamic effects, our baseline specification is

$$y_{it} = u_i + \beta_1 \log b_{it} + g_i t + d_t + \epsilon_{it}, \tag{7}$$

where y_{it} is some market-level outcome of interest in city i during week t, u_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.[20]

5.1 Threats to identification

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

[19] Note that this price index is not the price of an hour of transportation services, as in our model. Our index is instead the price for a standard trip, which we could convert to a time-based rate by dividing by the duration. However, this distinction is immaterial for our purposes, as we will be using the log of our price index as an independent variable.

[20] In §A.12, 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. Also in §A.13 we show residual plots for market quantity regressions. These diagnostic plots show that our preferred specification fits the data well for all cities.

likely to affect our results.

Our paramount identification concern is that cities receiving cuts are on different trends with respect to the outcome. We partially address this concern by including city-specific linear time trends in Equation 7. We can also diagnose whether there is a violation of parallel trends by plotting cumulative effects from a distributed lag model, prior to treatment, which we do in §5.3. We find no evidence of pre-trends in any of our outcomes. As a less model-dependent diagnostic approach, we can conduct an event study around the two January fare cuts to see, graphically, whether cities receiving cuts were on different trends, which we do in §A.5. Again, we find no evidence of a difference in pre-trends by "treatment" status.

Even with city-specific linear time trends, a concern in a long panel like 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 7. 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 §A.9. 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 7 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 §A.6, finding that we cannot reject the null of strict exogeneity.

Even if there are no pre-trends or un-captured 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 across the two January fare cuts to see if cold weather cities are exclusively getting cuts. We do this comparison in §A.5, finding that many warm-weather cities get cuts and many cold-weather cities do not. Furthermore, 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 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 §A.7, we find that the point estimates do not change with the inclusion of these city-calendar month effects.

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 §A.3 and find that including Uber's imputed ride-sharing share leaves all the point estimates unchanged. We also interact the base fare with this share measure to see whether the extents of ride-sharing competition affected the adjustment process. We find no evidence that they do, though it is important to note that, during the period covered by our panel, direct ride-sharing competition was in many cities nascent.

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 §A.8. We find no evidence that the inclusion of these controls affects the results.

Despite not finding any evidence that Uber was conditioning on future shocks, a lingering concern is that perhaps Uber had access to information that allowed it to make accurate predictions. We view this as unlikely. Using the January fare cuts, we show in §A.1 that magnitude of the cuts was clearly conditioned on city-specific measures of utilization.

We also have some limited insight into Uber's decision-making that supports this perspective. As Uber fare changes were controversial, they received media attention, which we summarize in §A.4. 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, all else equal, how many more trips a driver would have to complete to keep their earnings the same—with no consideration of how fare changes would change both supply and demand in reality. In short, if these spreadsheets were used, Uber's decision-making could be characterized as selecting on observables.

Even if β_1 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. We diagnose whether this is a problem in §A.10 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, though there is a loss of precision.

Another kind of heterogeneous effect comes from the nature of our panel, which is "thick" with changes, meaning that our implied controls are also affected by treatment effects that occurred earlier. Although we cannot do a fully stacked difference-in-difference, in §A.9 we estimate our model with sub-panels, again finding that the sub-panel estimates are similar to the full panel estimates. This also somewhat allays concerns that effects have changed over time as markets have matured.

5.2 Regression estimates of the effects of fare changes on the hourly earnings rate, utilization, and surge

Table 1 reports estimates of Equation 7 where the outcome variables are the log hourly earnings rate, log utilization, and log surge in Columns (1), (2) and (3), respectively. For each regression, standard errors are clustered at the level of the city.[21]

From Column (1), we can see that the effect of a base fare price change is positive but close to zero—a 10% fare increase would raise the hourly earnings

[21] We also conducted a block bootstrap at the city level to test for Bertrand et al. (2004) problems, but we found that the bootstrap standard errors were almost identical to the clustered standard errors, and so we only report clustered standard errors.

rate by just 0.7%. The confidence internal comfortably includes 0 and negative point estimates. From Column (2), we see part of the explanation for why hourly earnings do not increase—a higher fare reduces utilization, with a 10% increase in the base fare reducing utilization by about 7%. The rest of the 10% increase in fares is approximately undone by about a 2% decrease in the average surge multiplier, as we can see in Column (3). As $w = bmx$, the sum of the percentage change in surge, utilization and the base rate index should equal the percentage change in the hourly earnings rate.

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

	Dependent variable:		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.070	-0.715^{***}	-0.208^{***}
	(0.066)	(0.068)	(0.036)
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	0.784	0.842	0.472
Adjusted R^2	0.775	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 7. 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 §4.1 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: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

5.3 Distributed lags and leads and parallel trends

To explore how the market adjusts over time, we need a more richly specified regression model. We switch to using a finite distributed lags model, estimating

$$y_{it} = u_i + \sum_{\tau=\textsc{NumPre}}^{\textsc{NumPost}} \beta_\tau \log b_{it-\tau} + d_t + g_i t + \epsilon_{it}, \tag{8}$$

where u_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. For this analysis, the same sample is used as in Table 1, though it is shortened by $\{\textsc{NumPre} + \textsc{NumPost}\}$ weeks to account for the leads and lags. 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 $\sum_{\tau=\textsc{NumPre}}^{0} \hat{\beta}_\tau = 0$ 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.[22]

The implied weekly effects from Equation 8 are plotted in Figure 6 for the log hourly earnings rate, log utilization, and log surge. There 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.[23] For each outcome, the long-run effect from Table 1 is

[22]In §A.12 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.

[23]In §A.12, we report our preferred regression specification but vary the post-period band-

plotted at the 0 week, which is the week that the actual fare was changed.

In the top panel of Figure 6, the outcome is the log hourly earnings rate. For this regression, and all distributed lag regressions, standard errors are clustered at the level of the city. Examining the pre-period, there is no obvious trend. Following a fare increase, the driver hourly earnings rate increases immediately, though there is considerably less than full pass-through; the elasticity point estimate is only about 0.5. 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 from Table 1, the long-run estimate near the end of the post-period is negative, albeit with a 95% CI that (barely) includes zero, with a point estimate of -0.3.

In the middle panel, the outcome is the log 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 -0.8, which is close to the estimate of the static effect estimate from Equation 7. By the end of the post-period, the effect is -1.

In the bottom panel, 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 the end of the post-period, the effect is close to -0.3.

5.4 Accounting for driver costs

The gross hourly earnings rate measure does not include the costs to drivers. These costs likely change with the utilization, as un-utilized drivers waiting for dispatch can conserve fuel and reduce wear and tear by driving more slowly

---

width. 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.

Figure 6: Effects of a base fare increase on the driver hourly earnings rate and its components

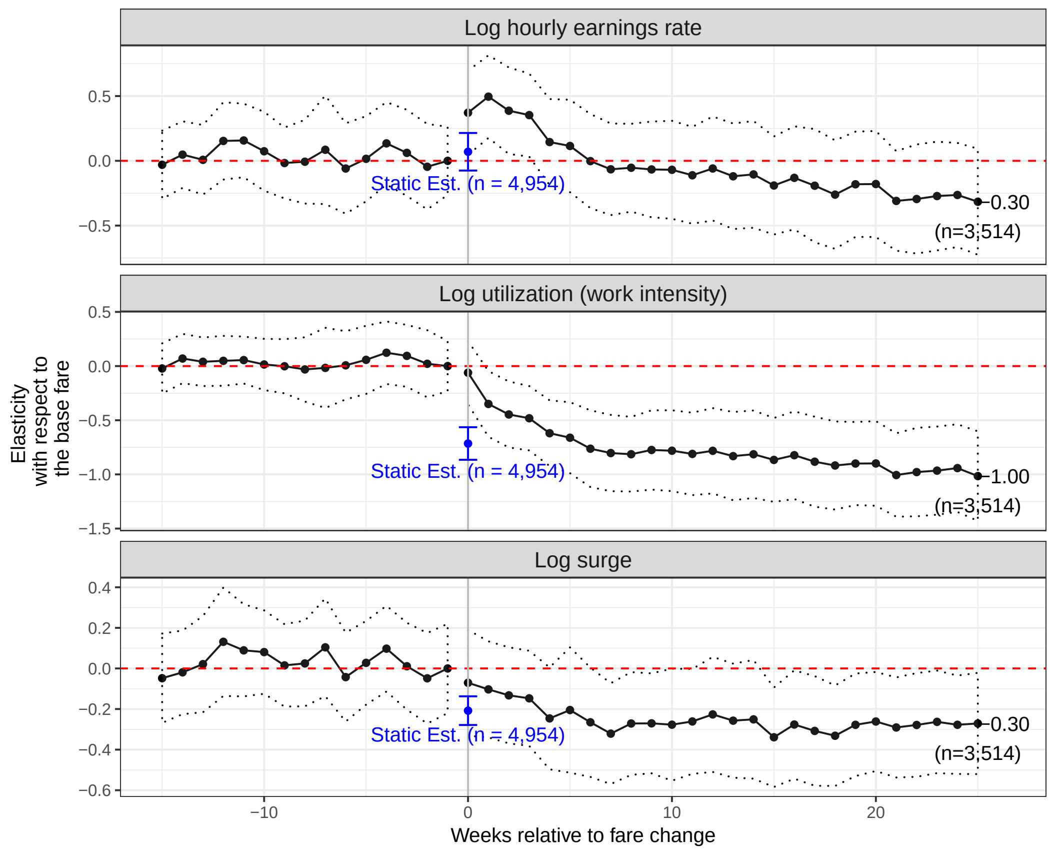

Notes: This figure plots the by-week cumulative effects of changes in the UberX base fare on the log hourly earnings rate, the log average utilization, and the log surge, along with 95% confidence intervals. These effects are from an estimation of Equation 8. The sample is a panel of US cities—see § 4.1 for a description. The x-axis are weeks relative to a fare change. 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 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. The confidence interval shows at week zero indicates the static effect corresponding to Equation 7.

or even stopping completely. As such, a lower utilization equilibrium is less costly to drivers, even aside from the reduced effort costs.

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, 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 lowers costs, as the reduction depends on the driver's vehicle. However, we can make some assumptions to construct a measure of net earnings more closely approximates a net measure of earnings i.e., $w - c(x)$ from our model.

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 on average $xs^{\text{Pax}} + (1-x)s^{\text{NoPax}}$ miles. The cost-per-hour in city i is then

$$\hat{c}(x) \approx \left(x_{it} s_i^{\text{Pax}} + (1 - x_{it}) s_i^{\text{NoPax}} \right) \bar{C} \tag{9}$$

where $\bar{C}$ is a cost parameter that is linear in miles-driven. Because we are modeling costs as linear in miles-traveled in Equation 9, the elasticity of costs with respect to utilization does not depend on the level of costs but rather the level of utilization and the percentage increase in speed when active, or Δs. For the elasticity of costs with respect to utilization, we do not actually need to know $\bar{C}$, as the elasticity is

$$\frac{\partial \log \hat{c}(x)}{\partial \log x} = \frac{\Delta s_i x_{it}}{1 + x_{it} \Delta s_i}, \tag{10}$$

and if we use the panel mean values for utilization, 0.58, and the panel mean percentage increase in speed carrying passengers, we get $\epsilon_x^{\hat{c}(x)} \approx 0.15$, and hence $\epsilon_b^{\hat{c}(x)} \approx -0.106$.

The effect of fare changes on net hourly earnings depends on $\bar{C}$. We use the rate of \$0.32/per mile for $\bar{C}$, which is used in Cook et al. (2018). This multiplier is intended to capture the full direct costs of the marginal mile driven, but not the costs of effort.[24] Using this rate and the city-specific speed

[24] https://web.archive.org/web/20181230201119/https://www.irs.gov/newsroom/standard-mileage-rates-for-2018-up-from-rates-for-2017

data, we 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).

The effects of fare changes on costs are shows in Figure 7. In the top panel, log costs are the outcome. We can see, as expected, that a fare increase lowers costs. However, the effects are modest—the end of period elasticity is only -0.09, which is fairly close to our ball-park estimate of -0.106 calculated simply from the panel-wide average speed differences and utilization. Note that these effects are not just mechanically the utilization effects scaled, as we make use of city-specific measures of changes in driving speed with and without passengers.

In middle panel, the outcome is the log hourly earnings rate, net of direct costs, with costs inferred using the Equation 9 method. With the inclusion of driving costs, the denominator is smaller, and so the range of effect sizes (which are percentages) is larger than when costs are not included. The same basic pattern of results holds—a fare increase raises the hourly earnings rate at first, but then it declines, eventually turning negative.

The Equation 9 method likely over-states the per-mile costs drivers face when they are not with passengers. For one, fuel efficiency is better when a vehicle is occupied only by the driver than by the driver and one or more passengers. Not having passengers lowers cleaning costs and expenses for consumables. Drivers can also use miles driven without passengers to head in a direction they are going anyway (such as to commute at the end of a shift or to run errands). For these reasons, in the bottom panel we just include our estimate of direct costs, essentially ignoring any costs associated with an hour-worked that is not with a passenger. We can see that this long-run estimate is quite close to zero.

5.5 Driver labor response

Although we can measure the effects of fare changes on the driver hourly earnings rate and we have precise measures of driver hours-worked, we cannot

Figure 7: Effects of a base fare increase on driver direct costs and net hourly earnings

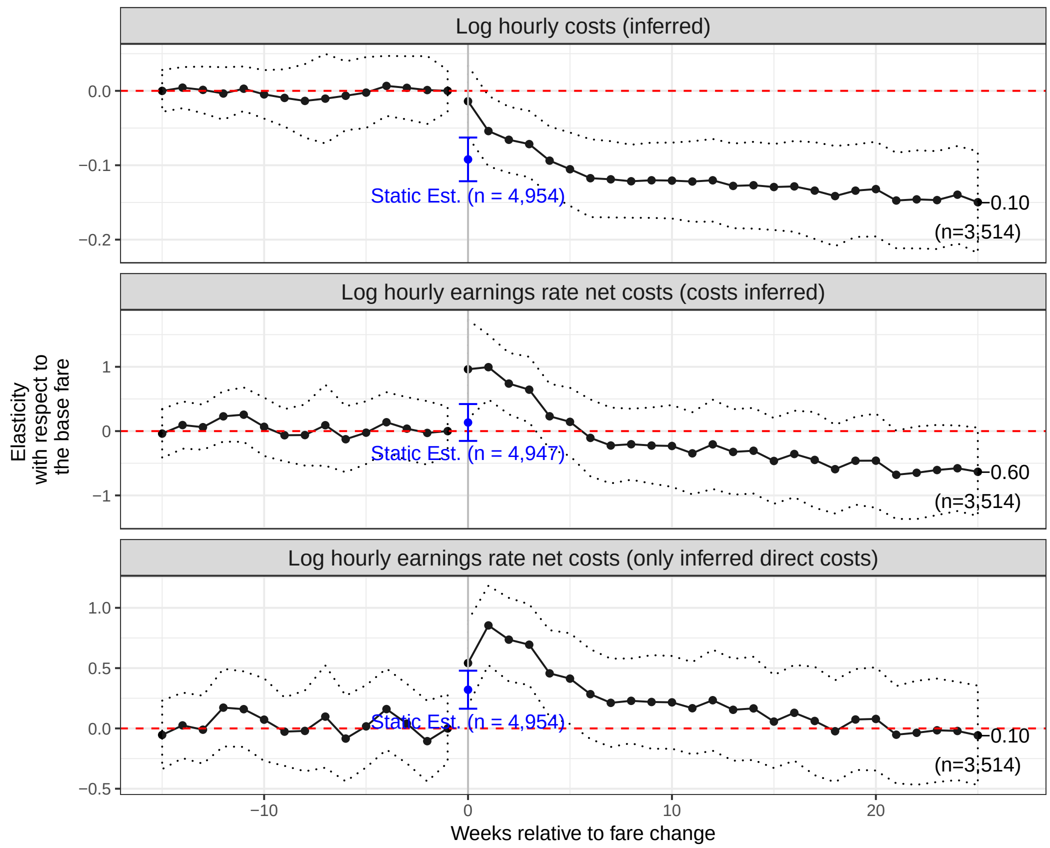

Notes: This figure plots the by-week effects of changes in the UberX base fare on drivers costs, from an estimation of Equation 8. The sample is a panel of US cities—see §4.1 for a description. The x-axis are weeks relative to a fare change. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The confidence interval shows at week zero indicates the static effect corresponding to Equation 7.

measure a labor supply elasticity. Given the change in hourly earnings and costs, driver utility could change, and hence hours-worked could change, as $H'(u) > 0$. Our measures of the driver labor response include hours-worked per driver, the number of drivers working at least some number of hours, and a measure of "churn," which is the fraction of drivers active in the previous week that are active the next week. The effects of fare increases on these outcomes are shown in Figure 8.

In the top panel, the outcome is the log hours-worked per driver. This measure is calculated by taking the total hours-worked that week for all drivers, divided by the total number of active drivers. We can see that following a fare increase, hours-worked per driver jumps immediately and then continues to rise, eventually flattening out but never obviously declining. The long-run effect at the end of the period is 0.6.

In the middle panel, the outcome is the log number of active drivers. Unlike with the hours/driver measure, we see no evidence of a sharp initial jump. However, as the weeks pass, there is some evidence of more drivers, which then declines somewhat, eventually turning negative by the end of the period. However, all of these estimates are fairly imprecise, with a 95% CI always including zero. In the bottom panel, the outcome is the percentage of drives churned. We can see it falls initially after the fare increase—consistent with the slight increase in drivers active—but then the effect on the churn measure returns back to close to 0.

The relationship between the change in the hourly earnings rate and the change in hours-worked per driver sheds light on the decision-making of drivers. If we compare Figure 8 and Figure 6, we can see that as weeks go by following a fare increase, hours-worked rise and stay higher while the hourly earnings rate keeps declining. In fact, by week 8, the hourly earnings rate—without or without costs included—is lower than before the fare increase, but hours-worked is higher. In short, the effect on hours-worked has the "wrong" sign if we naively interpret gross or net hourly earnings rates as wages—wages are lower, but drivers are working more.

A parsimonious explanation is that greater utilization is costly to drivers beyond what we tried to capture with $\hat{c}(x)$ in §5.4. Because of this cost, following a fare increase, even if the hourly earnings rate is somewhat lower, the reduced utilization at the higher product market price equilibrium is enough to generate an equilibrium increase in hours-worked.

Uber has, in some markets, paid promotional payments to drivers. Many of these payments are various forms of earnings guarantees. Typically, if drivers drive some minimum number of hours, they are guaranteed to make at least

Figure 8: Effects of fare increase on driver labor output in equilibrium

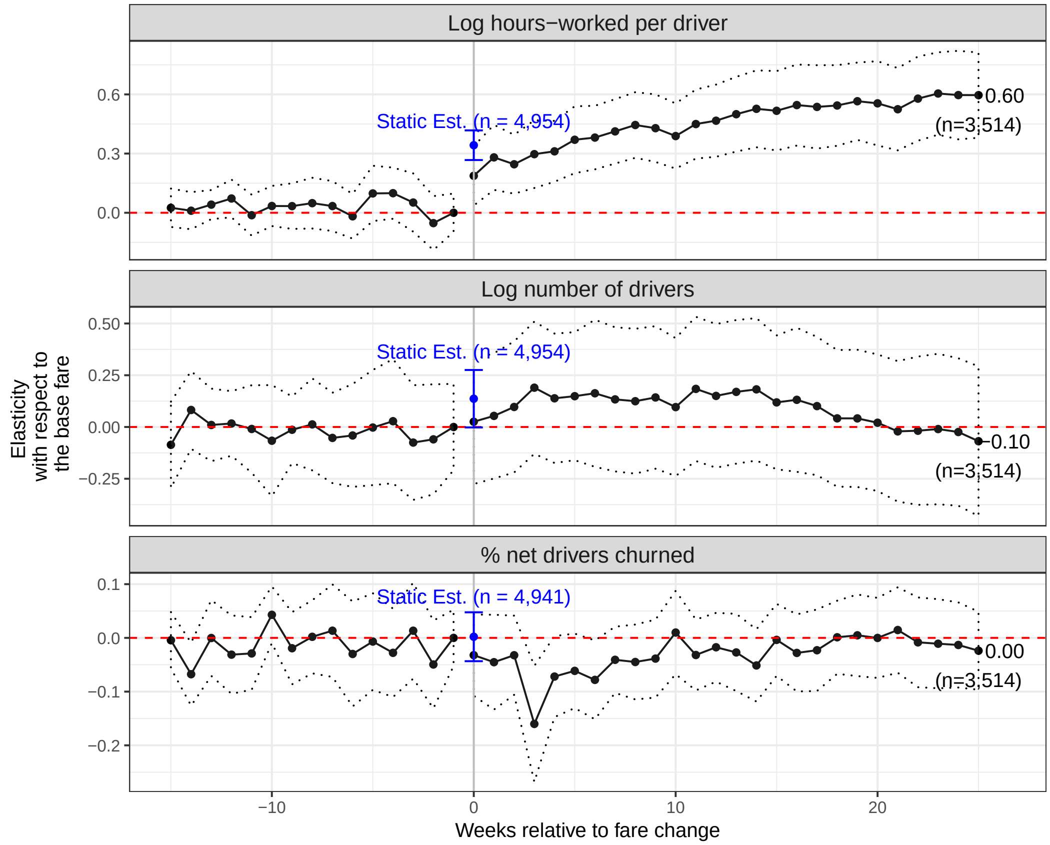

Notes: This figure plots the effects of changes in the UberX base fare on hours-worked per driver, log number of active drivers and the percentage of drivers churning. These effects are from an OLS estimation of Equation 8. The sample is a panel of US cities—see §4.1 for a description. The x-axis are weeks relative to a fare change. The independent variable is the base fare index. 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 confidence interval shows at week zero indicates the static effect corresponding to Equation 7.

some floor amount. Including these payments to construct a "wage" is misleading in the sense that these promotional payments are non-linear and averages using these measures do not reflect the marginal hourly earnings rate a driver faced.

Despite the limitations of this measure, We explore the effects of fare changes on measures of driver hourly earnings that include these promotional payments in §A.11. The same basic pattern of results appear, though there is some evidence that promotional payments "buffer" the effects of base fare changes in the short-run, leading to much less pass-through.[25]

5.6 Passenger wait times and service quality

As we discussed when presenting the model in §3, one way the market could clear following a base fare change was a *shift* in the demand curve, rather than movement along the demand curve. In our context, the likely demand curve shifter is passenger wait time, as shorter wait times are preferred to longer wait times by all would-be passengers. If a fare increase reduces utilization, all else equal, the nearest empty car that can be dispatched to a passenger is closer. As a second channel, if a fare increase increases hours-worked, then this would also tend to lower wait times.[26]

We explore this possibility in Figure 9, where the outcome is the median wait time. We can see that a 10% increase in fares reduces wait times at the end of the period by about 6%. The reason for this change is presumably that with less demand and/or lower driver utilization, for a given would-be passenger requesting a ride, the nearest empty car is likely to be closer (so long

[25] There has been some debate about whether these guarantees are empirically important. "But drivers BuzzFeed News spoke with say the guarantee comes with so many complicated stipulations that it's almost worthless. The guaranteed rates are available 24 hours a day, but at different rates depending on certain time slots (for example, drive on Friday from 6 p.m. to 9 p.m., earn \$20). But, crucially, those rates are attainable only by drivers who accept and complete a certain percentage of rides (for example, 1.5 rides per hour completed with a 90% acceptance rate). If drivers fail to meet any of these criteria, the guarantee doesnt apply. Uber declined to tell BuzzFeed News what percentage of drivers received guarantees."

[26] There is likely to be a mechanical trade-off between utilization and wait-times if hours-worked stays the same. Consider a would-be passenger requesting a ride. Let $F(r)$ be the probability that a given car is r or fewer units of distance away. If we normalize the "maximum" distance to 1 and assume that probability of a car in a particular region is constant (i.e., a Poisson point process), then the probability that the closest care is r distance away or further is $D(r) = (1 - F(r))^{(1-x)N}$ where N is the number of drivers working at that moment. We can use this to get the pdf and compute the expected distance, which is then $\bar{r} = (2 + 3n + n^2)^{-1}$, where $n = (1 - x)N$, which in turn is proportional to the wait time. A higher utilization thus increases wait time unless there is an offsetting increase in N.

as the number of drivers does not decrease). These changes in wait times imply that markets could not clear entirely through changes in utilization, and that furthermore, surge pricing did not, as implemented during the period covered by the panel, hold product attributes exactly fixed. Although there could be a decline in service quality not measured by wait times, the lack of composition effects on the extensive margin (as we will see) make this somewhat unlikely.

Figure 9: Effects of fare changes on log actual median wait times

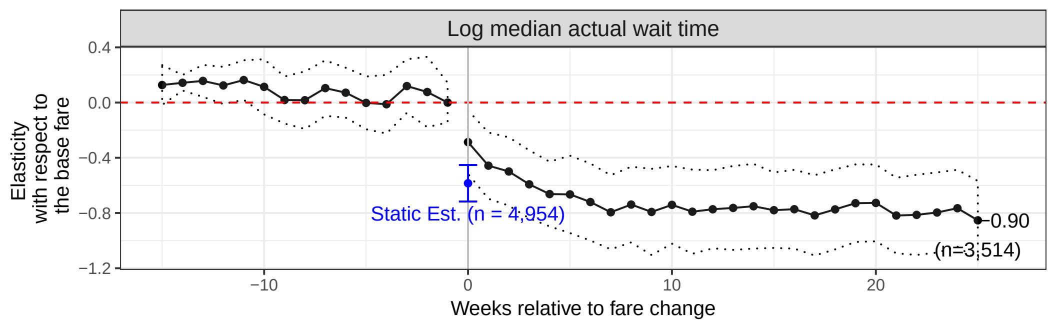

Notes: The outcome in this figure is the log median wait time (in seconds). These effects are from an OLS estimation of Equation 8. The sample is a panel of US cities—see §4.1 for a description. The x-axis are weeks relative to a change in the base fare index. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The vertical dashed blue line and confidence interval at $t = 0$ indicates the static effect corresponding to an estimation of Equation 7. Fixed effects are included for the city and for the week. 95% CIs are shown for each point estimate and standard errors are clustered at the city level.

5.7 Market quantities

So far in our analyses, the outcomes have been rates rather than absolute quantities. Because of differences in city sizes and growth trajectories, fitting a single regression model to all of the data is potentially challenging. However, when outcomes are logged and a city-specific linear trend is included, a regression model can fit the data reasonably well.[27] In contrast to our "rate"

[27]See §A.13 for by-week, by-city residual plots. Many of the largest residuals can be accounted for by one-off events (e.g., the Super Bowel in Phoenix). Other cities show residuals

regressions, city-specific linear trends are essential for our "quantity" regressions.

The effect of the base fare index on market quantities is shown in Figure 10. In the top panel, the outcome is the log number of trips taken, while in the middle panel, the outcome is the log number of hours with passengers. For both outcomes, the effects at the end of the period are negative, with values of -0.32 and -0.49, respectively. As with all of our outcomes, it is important to remember that the unit of analysis is the whole market and although it is tempting to treat these point estimates as demand elasticities, this is incorrect.[28]

In the bottom panel, the outcome is log hours-worked. It rises immediately after a fare increase and continues to climb before flattening out with a long-run value of 0.53 by the end of the period (this outcome is just the product of the extensive and intensive margins presented in Figure 8). The comparison of hours with passengers and hours-worked clearly illustrates the utilization effects: with higher base fares, drivers work more hours but actually provide fewer hours of transportation services.

6 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. 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 unchanged or even slightly lower. The lack of

with a clear periodicity (e.g., Tuscon and Ann Arbor) corresponding to the academic year, as these are both large college towns. We make no attempt to include additional controls for these features of our data, as doing so would be somewhat *ad hoc* and we prefer to let this kind of variance be accounted for by the error term.

[28]Because of our finding of changed services quality and the change in the level of surge pricing—both factors would tend to bias the estimate towards zero relative to the true demand elasticity (i.e., less surge pricing undoes some of the fare increase and service quality improves). That caveat aside, the long-run point estimates are close to the micro demand estimates found in Cohen et al. (2016).

Figure 10: Effects of fare changes on market quantities

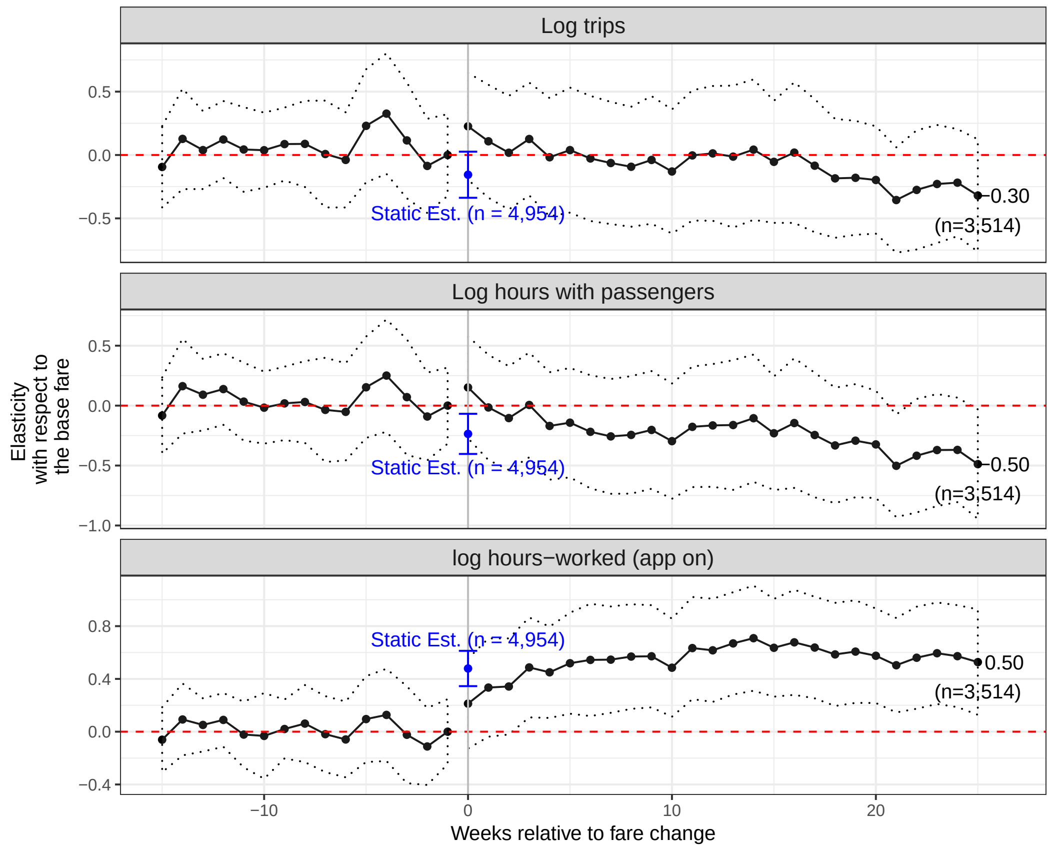

Notes: The outcomes in this figure are the log number of trips, the log number of hours with passengers, and the log number of hours-worked. These effects are from an OLS estimation of Equation 8. The sample is a panel of US cities—see §4.1 for a description. The x-axis are weeks relative to a change in the base fare index. 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 vertical dashed blue line and confidence interval at $t = 0$ indicates the static effect corresponding to an estimation of Equation 7. Fixed effects are included for the city and for the week. 95% CIs are shown for each point estimate and standard errors are clustered at the city level.

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) point out, the increase in earnings from self-employed drivers left the average unchanged.

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. The platform could switch to an auction model, with drivers 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.

One way to characterize the ride-sharing platform's pricing problem is as choosing a preferred fare/utilization equilibrium. Our theoretical model suggests this price is likely to be as low as possible, or equivalently, utilization as high as possible. Both prices and utilization have been moving in this direction. However, our evidence that lowered fares increased wait times and led to more surge suggests that some fare cuts led to degraded reliability, leaving the market unable to clear solely through changes in utilization.

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. 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 interpreted as attempts to raise utilization though 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 perspective. In particular, it would be interesting to consider driver micro labor supply decisions, focusing on the role of the individual differences in costs. 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., Sydnee 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.

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.*

Buchholz, Nicholas, "Spatial equilibrium, search frictions and efficient regulation in the taxi industry," *Working Paper*, 2015.

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.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer, “The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers,” *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.

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 provide: 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.

A Additional analyses and robustness checks

A.1 Evidence for selection on observables

Despite no evidence that Uber was selecting cities on the basis of predicted changes in demand or supply, we have strong evidence that they were conditioning on utilization when deciding on the size of price cuts, at least for the January cuts.

In the bottom row of Figure 11b, 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 11b, cities that did not receive January cuts 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 (consistent with Table 1). 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 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]eve 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."[29]

A.2 Uber conditioning on future shocks

Parallel trends in the pre-period are not sufficient to identify the effects of fare changes. An important threat to identification would be if any value of ϵ_{it} was correlated with b_{it}. The most likely way this could occur is if Uber was conditioning accurately on future demand or supply shocks when choosing a price. To see why, suppose demand was going to fall sharply in January for certain cities, but not for others. In the absence of any fare change, and if drivers made no adjustments, cities about to experience a demand reduction would see their utilization fall. By lowering prices in cities about to experience a decline, Uber would stimulate demand and counter-act this decline. Note that with the inclusion of week-specific fixed effects, a general change in demand or supply is not a concern.

Empirically, we can control for city-specific weather by including indicators for the calendar month (i.e., 12 dummies for January, February, and so on) interacted with the city. We do this in §A.7 and show that the inclusion of these effects barely changes any of the point estimates.

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. This analysis is in §A.8.

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

A.3 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 2 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 2: Effects of fare changes on market outcomes from a city-week panel of UberX markets

	Dependent variable:		
	Log hourly earnings	utilization	surge
	(1)	(2)	(3)
Log base fare index	−0.125	-0.892^{***}	-0.224^{***}
	(0.102)	(0.105)	(0.030)
Uber share	0.017	0.010	0.003
	(0.013)	(0.011)	(0.002)
Uber share × Log base fare index	−0.004	−0.002	−0.001
	(0.005)	(0.004)	(0.001)
City FE	Y	Y	Y
Week FE	Y	Y	Y
Observations	4,954	4,954	4,954
R^2	0.724	0.774	0.440
Adjusted R^2	0.714	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 7, but with the base price interacted with Uber's imputed market share. 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 §4.1 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: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

late as January 2016.[30]

A.4 Was Uber forecasting demand?

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.[31] 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.

---

[30]https://web.archive.org/web/20181231153043/https://www.thestreet.com/markets/ipos/do-uber-recent-numbers-support-its-ipo-valuation–14783564

[31]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.

Note that if Uber were conditioning on anticipated treatment effects—such as choosing cities with low utilization for fare cuts precisely because these cities 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.[32]

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 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. We do this by estimating "early" and "late" panels (§A.9) and by isolating sources of variation that are all of the same "type" (§A.10), finding no evidence of strongly differing effects.

A.5 Parallel trends, with evidence from January cuts

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.

[32] Our 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 $\mathbf{E}[\Delta y] = \mathbf{E}[\tau(\alpha_i)Pr(W = 1|\alpha_i)]$.

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 in §5.3, 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 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 11a, 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 4 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 5a 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

Figure 11: By-week city demeaned utilization rates around the two January fare cuts

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

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

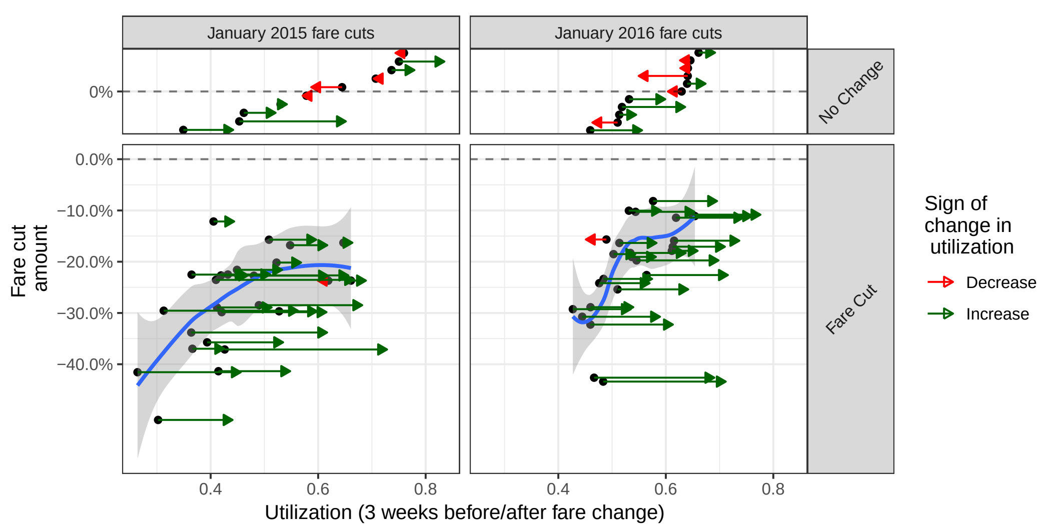

Notes: In the left column in the top panel shows the actual demeaned utilization around the January fare cut week, with the solid red line indicating non-cut cities (the control) and the blue dashed line showing the treatment. Data are from the 2015 and 2016 January fare cuts, in the top and bottom rows, 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.

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, allowing for a kind of placebo test.

In the middle column of Figure 11a, 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.

A.6 Testing 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),

$$y_{it} = u_i + \beta_1 \log b_{it} + \gamma \log b_{it+1} + g_i t + d_t + \epsilon_{it}. \tag{11}$$

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 0.067, giving a t-stat of -1.308 and a p-value of 0.191.

A.7 Controlling for city-specific weather by season

In §A.2, we discussed how Uber could be conditioning on expected changes in demand, such as due to weather—Uber claimed as much in certain press releases. 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 city-specific differences in demand or supply due to different weather patterns.

Table 3 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.

A.8 Controlling for the city-specific unemployment rate

In §A.2, we discussed how Uber could be conditioning on expected changes in demand or supply, such as due to local economic conditions. In Table 4 we include each city's MSA monthly unemployment rate as a regressor. All of

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

	Dependent variable:		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.065	-0.715^{***}	-0.214^{***}
	(0.070)	(0.069)	(0.037)
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	0.792	0.850	0.492
Adjusted R^2	0.781	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 7, 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 §4.1 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: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : * * *$.

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

	Dependent variable:		
	log hourly earnings rate	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	0.081	-0.697^{***}	-0.209^{***}
	(0.070)	(0.071)	(0.037)
Unemployment rate	0.001	-0.003	0.004
	(0.014)	(0.016)	(0.008)
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	0.777	0.841	0.475
Adjusted R^2	0.766	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 7 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 §4.1 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: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

A.9 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 12, 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."

Figure 12: Point estimates for the effects of base fare changes with different panel lengths

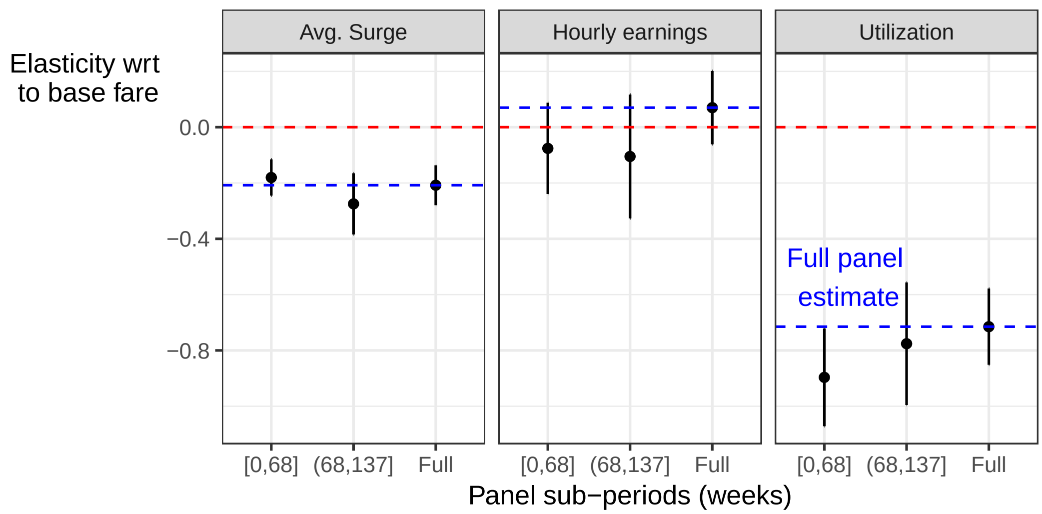

Notes: This figure reports estimates similar to Table 1, but with the panel cut in half on the time dimension and the city-specific linear time trend removed. For comparison purposes, the point estimates from Table 1, which includes the time trend are shown. These estimates are labeled "Full" and are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

A.10 Heterogeneous effects by direction of fare change

As we saw in Figure 5a, 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 13 reports estimates of Equation 7. 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,

Figure 13: Effects of a base fare increase on the driver hourly earnings rate and its components using different samples

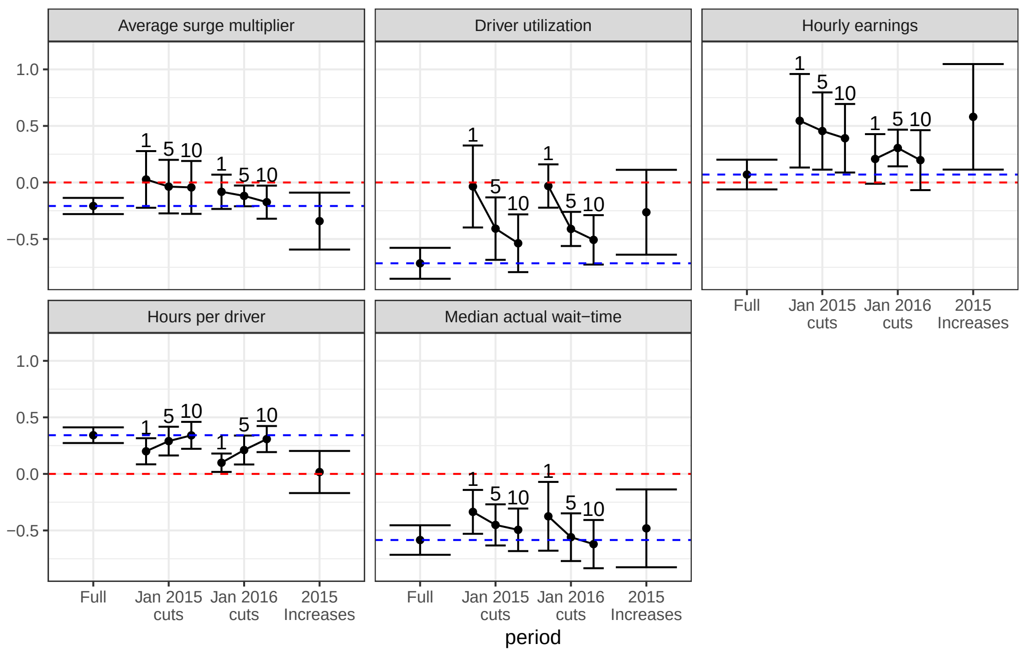

Notes: This figure reports estimates of Equation 7 using different sample definitions. The samples are: (1) the full panel estimate, (2) windows around the January 2015 cuts, (3) windows around the January 2016 cuts, and (4) the "interior" of 2015 when all of the variation in prices were price increases. 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.

and so it is unsurprising that point estimates change with longer panels.

A.11 The role of promotional payments

To explore the role of these promotional payments on market adjustment, in Figure 14, we plot the effects of a fare increase on the hourly earnings rate with promotional payments included. We also report the effects of the fare change on the fraction of hourly earnings that are due to promotional payments in the bottom panel. When promotional payments are included, there is essentially no detectable pass-through of the fare increase.

In the bottom panel, we can see that the fraction of earnings from promotional payments declines immediately and substantially. This would be expected even if there was no change in promotional payment program eligibility; with the change in fares, we would expect, at least in the short-run, for amount earned from these programs to vary.

Figure 14: Effects of fare changes on hourly earnings with and without promotional payments included

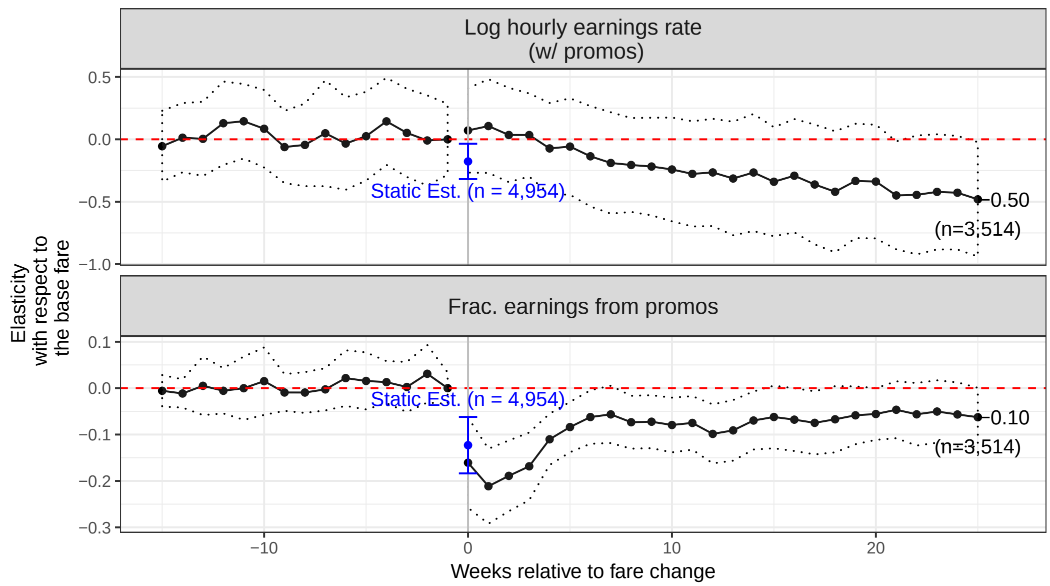

Notes: This figure plots the effects of changes in the UberX base fare index on three outcomes (from left to right): (1) the log hourly earnings rate not including promotional payments, (2) the log total hourly earnings rate includes promotional payments, and (3) the fraction of hourly earnings coming from promotional payments. These effects are from an OLS estimation of Equation 8. The sample is a panel of US cities—see § 4.1 for a description. The x-axis are weeks relative to a change in the base fare index. 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 vertical dashed blue line and confidence interval at $t = 0$ indicates the static effect corresponding to an estimation of Equation 7. Fixed effects are included for the city and for the week. 95% CIs are shown for each point estimate and standard errors are clustered at the city level.

A.12 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 15 illustrates the pattern we use for all outcomes-. In the left column, the cumulative effects from Equation 8 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 $\sum_{\tau=\textsc{NumPre}}^{0} \hat{\beta}_\tau = 0$. In Figure 16, we report cumulative effects 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 15: Alternative specifications for Figure 6

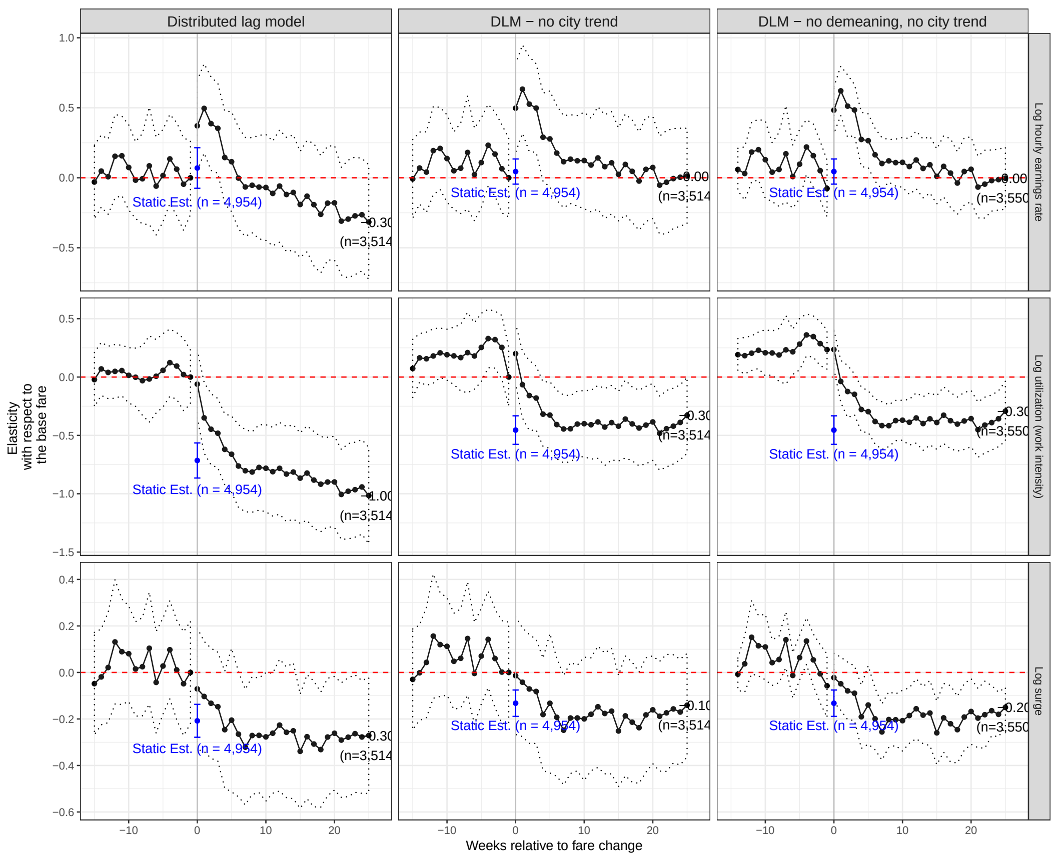

Notes: Alternative specifications.

Figure 16: Alternative post-period bandwidths for Figure 6

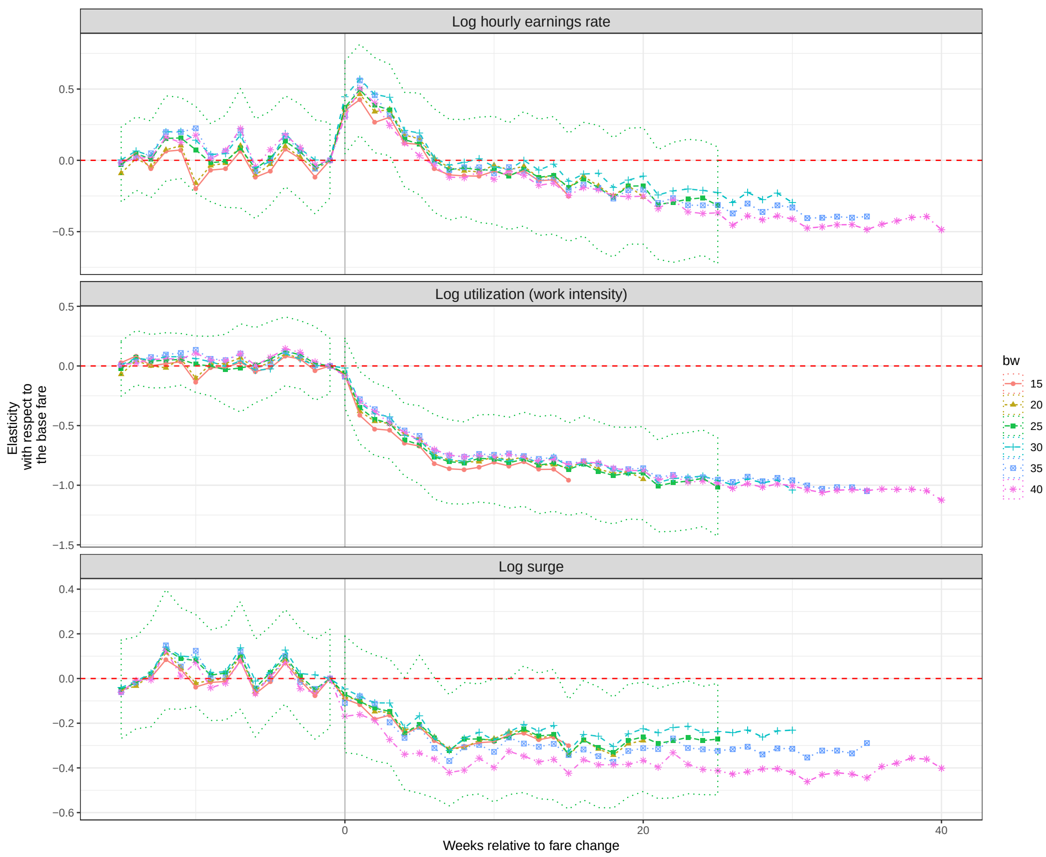

Notes: Alternative post-period bandwidths.

Figure 17: Alternative specifications for Figure 7

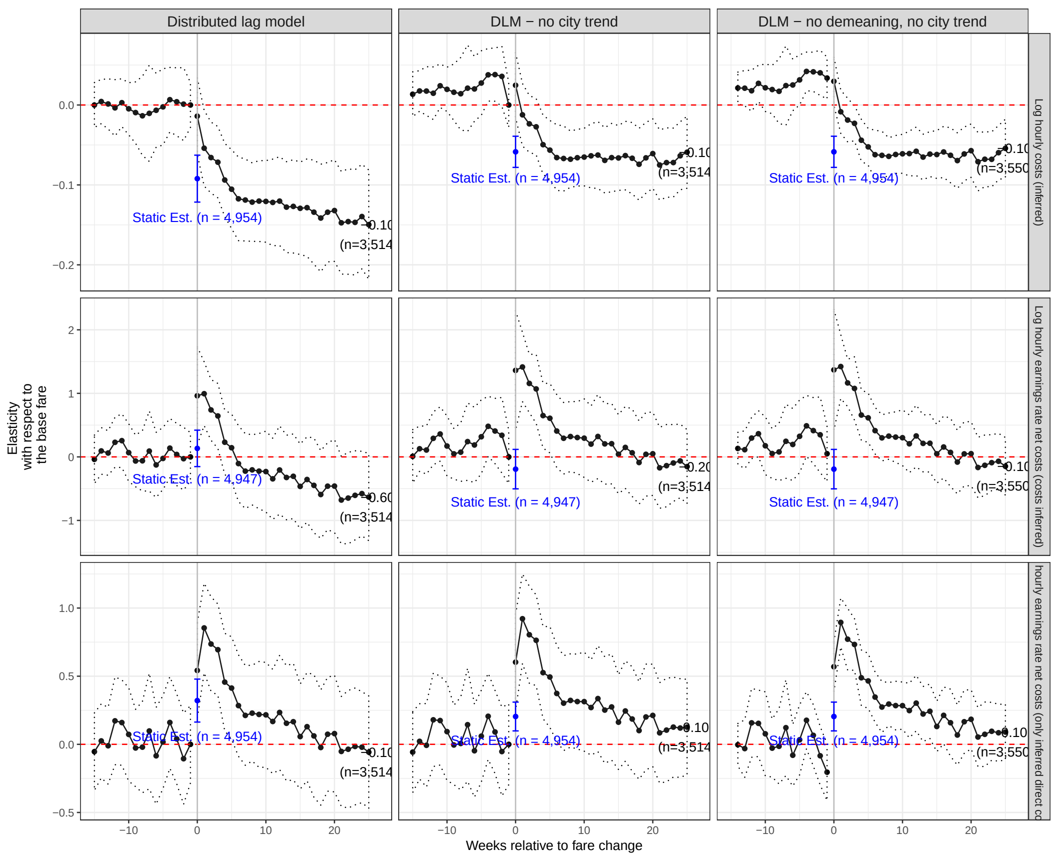

Notes: Alternative specifications.

Figure 18: Alternative bandwidths for Figure 7

Notes: Alternative post-period bandwidths.

Figure 19: Alternative specifications for Figure 8

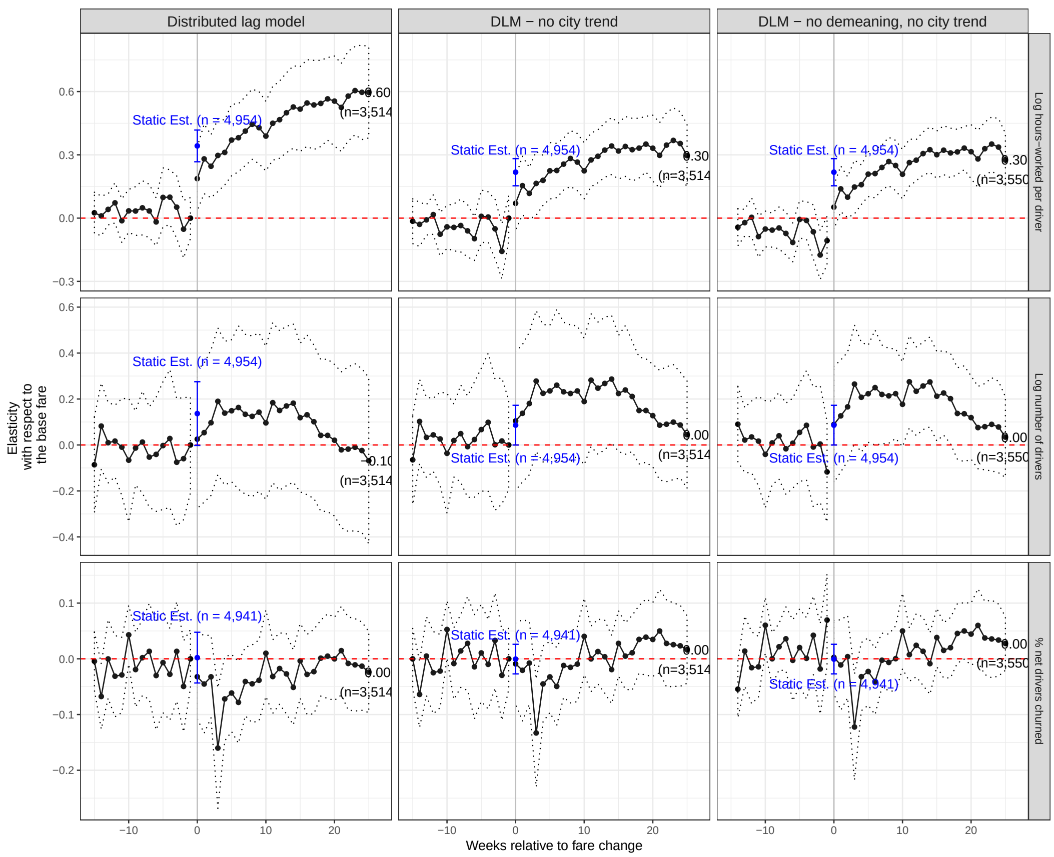

Notes: Alternative specifications.

Figure 20: Alternative bandwidths for Figure 8

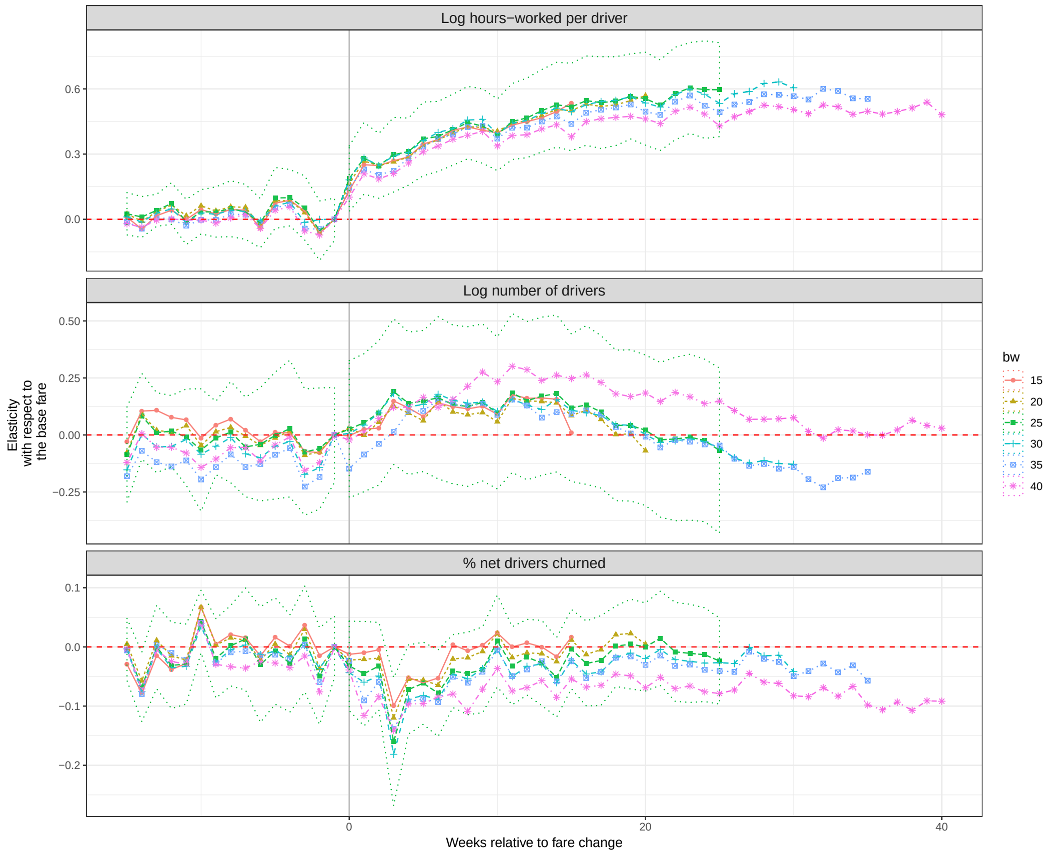

Notes: Alternative bandwidths.

Figure 21: Alternative specifications for Figure 9

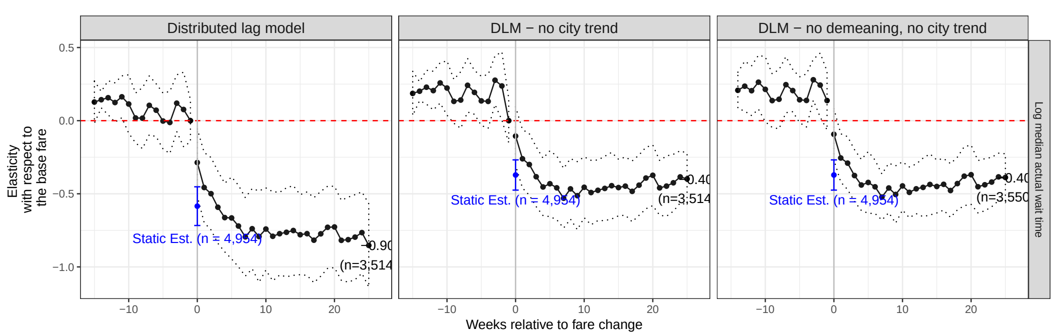

Notes: Alternative specifications.

Figure 22: Alternative bandwidths for Figure 9

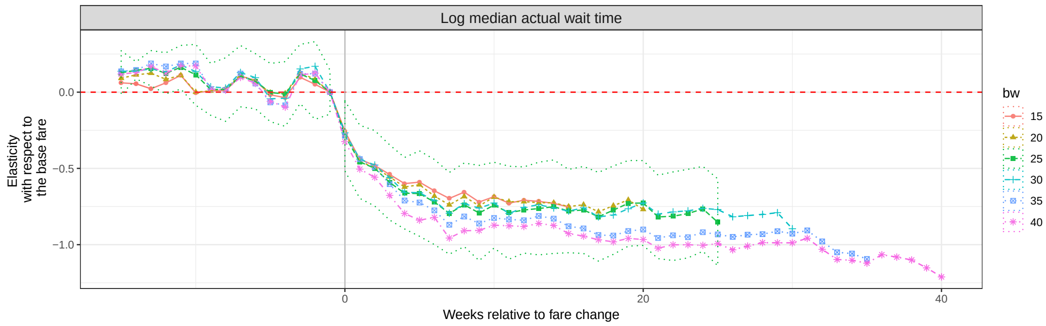

Notes: Alternative bandwidths.

Figure 23: Alternative specifications for Figure 10

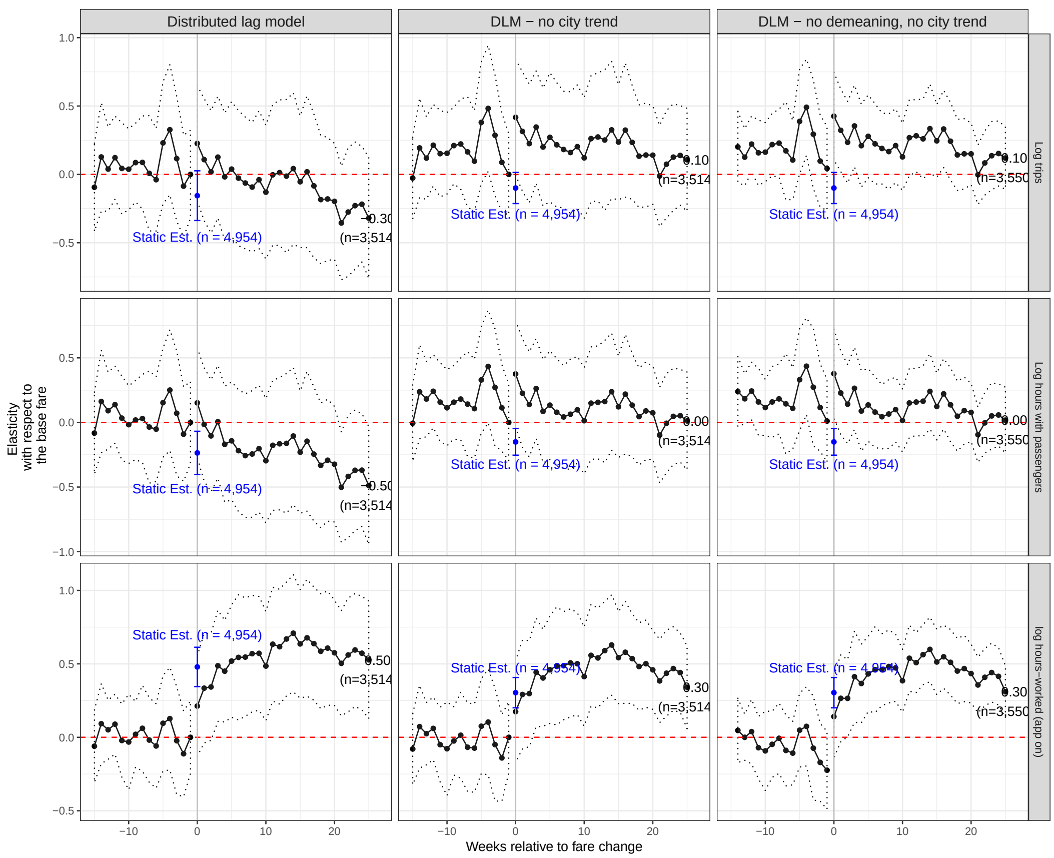

Notes: Alternative specifications.

Figure 24: Alternative bandwidths for Figure 10

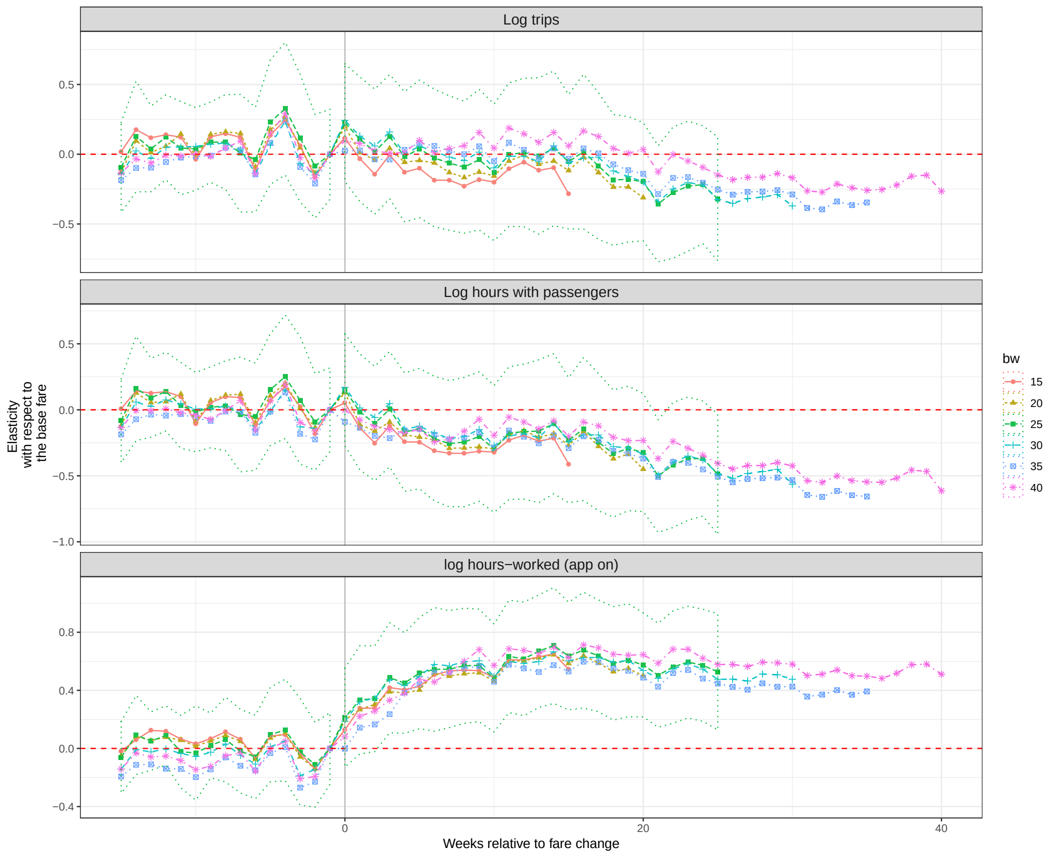

Notes: Alternative bandwidths.

Figure 25: Alternative specifications for Figure 14

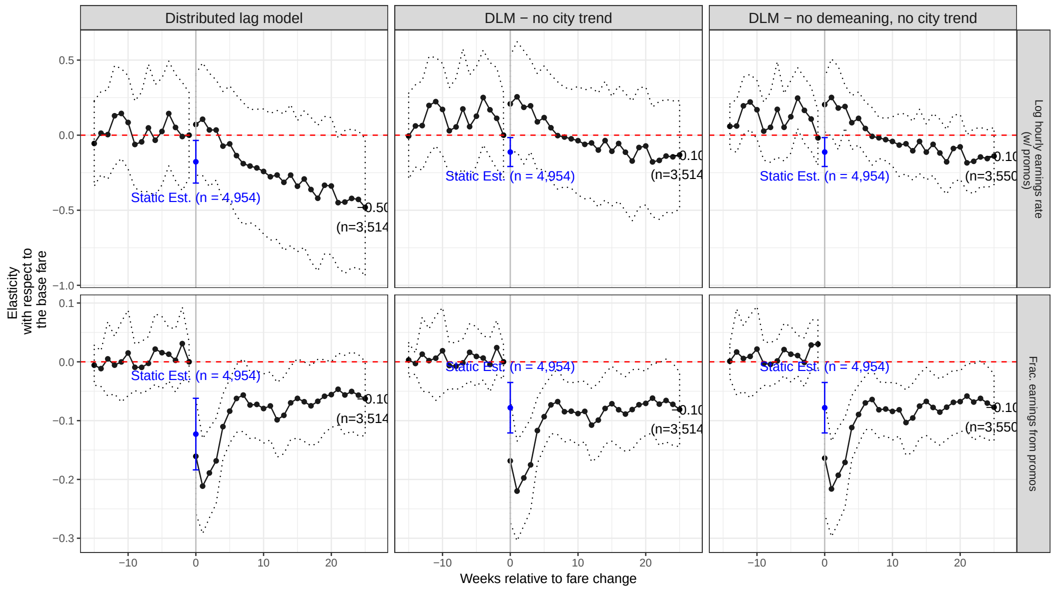

Notes: Alternative specifications.

Figure 26: Alternative bandwidths for Figure 14

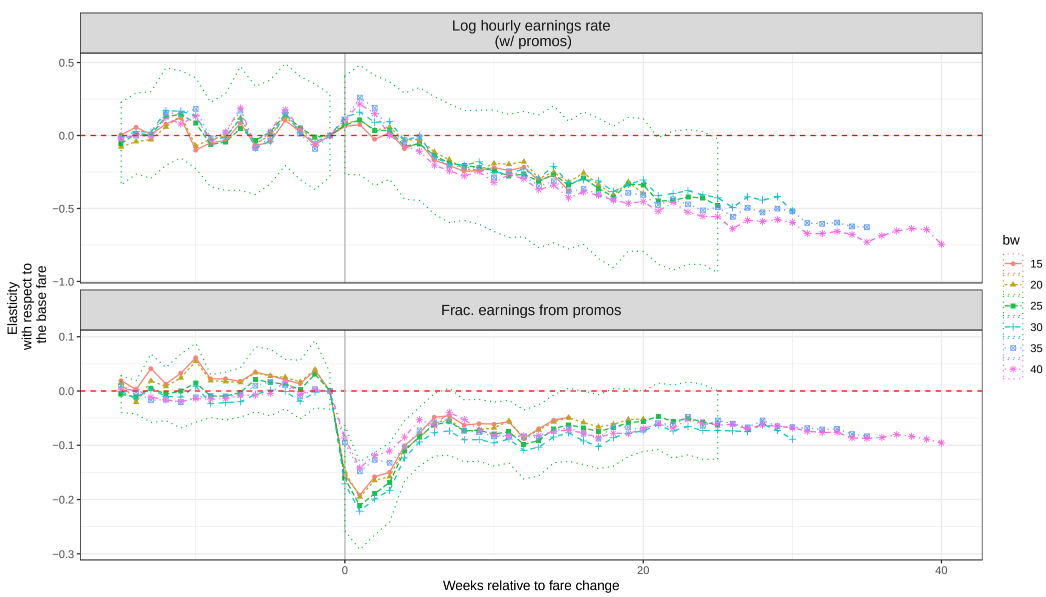

Notes: Alternative bandwidths.

A.13 Residual plots

Figure 27 shows residual plots for the log number of trips in a city, using Equation 7. Figure 28 shows residual plots for the log number of hours-worked in a city, using Equation 7.

Figure 27: Residual plots for log number of trips, by city and week

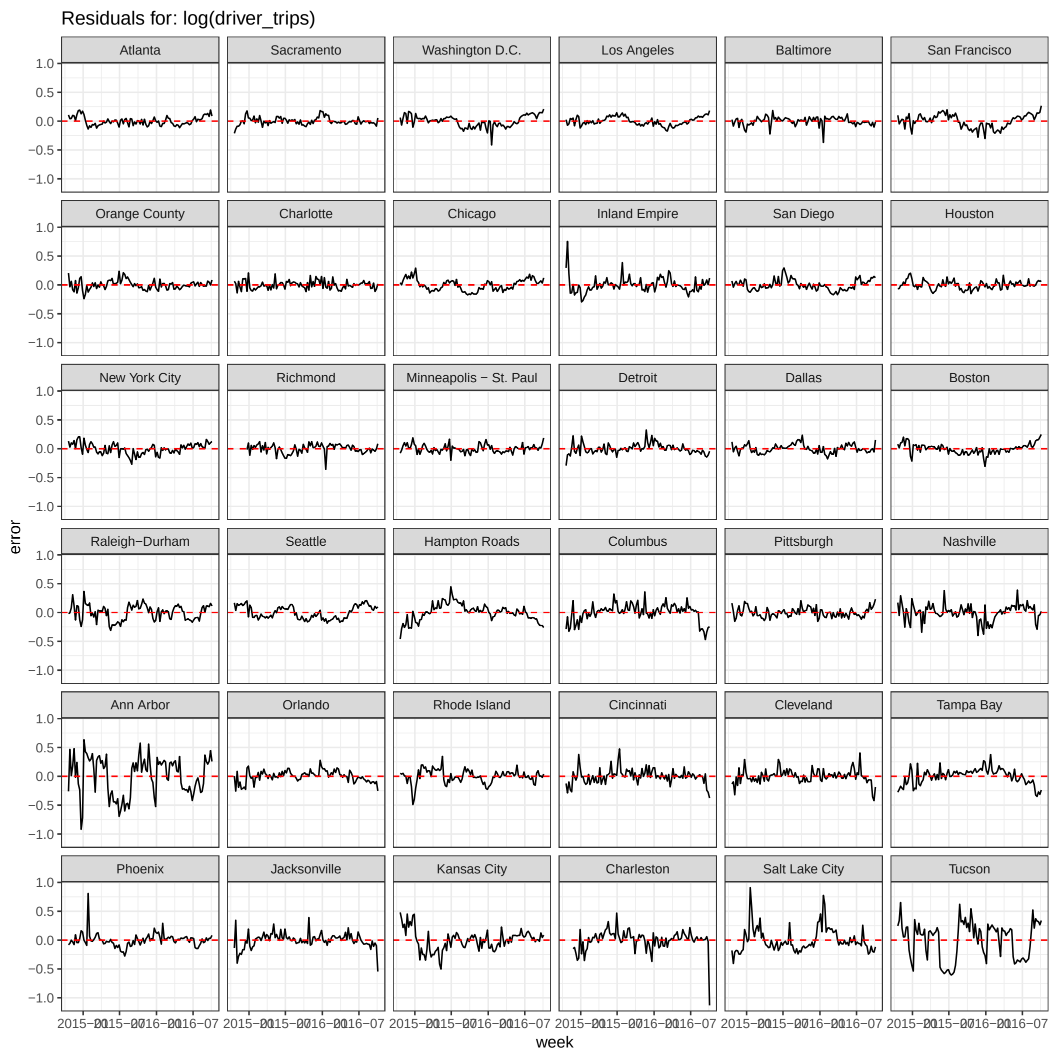

Notes: This figure plots the by-city week residuals when the outcome is the number of driver trips.

Figure 28: Residual plots for log number of hours-worked, by city and week

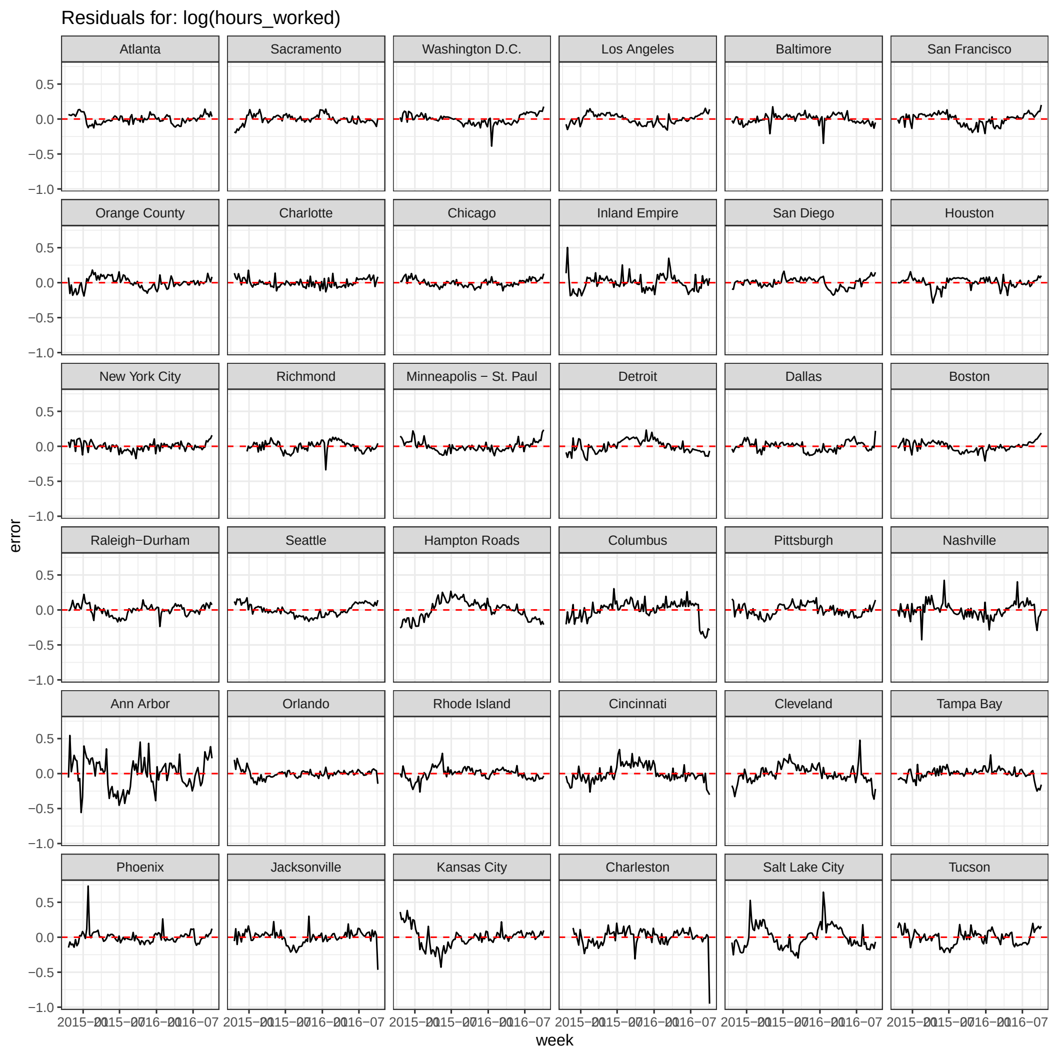

Notes: This figure plots the by-city week residuals when the outcome is the number of driver trips.