Pricing Efficiently in Designed Markets: Evidence from Ride-Sharing*

Jonathan V. Hall
Uber Technologies

John J. Horton
NYU Stern

Daniel T. Knoepfle
Uber Technologies

August 9, 2018

Abstract

In many designed markets, the platform eliminates price dispersion by setting the product market price, yet still allows supply-side free entry and exit. We explore the equilibrium of such markets, using data from Uber. Following price increases, drivers make more money per trip and—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, eventually bringing the hourly earnings rate back close to its previous level. The resulting higher fare/lower productivity equilibrium is generally inefficient.

JEL J01, J24, J3

*Thanks to Josh Angrist, Susan Athey, David Autor, 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, Ben Olken, 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. Please contact research@uber.com if you would like access to the data for replicating the results in this paper or for other research purposes. For disclosure statements, please see http://www.john-joseph-horton.com/papers/uber_disclosure.html.

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. 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—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. We use differences in the timing and size of city-specific, Uber-initiated changes to the fare faced by passengers to identify the effects of the fare on the market equilibrium. For our analysis, we abstract away from the specifics of Uber's taxi-like pricing of trips by constructing a "base fare" for each city week, which was the price for a typical trip during un-surged conditions.[1]

The price in the buyer side of the ride-sharing market is the base fare; the price in the seller side is the driver hourly earnings rate. 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.[2] However, the hourly earnings rate begins to decline shortly thereafter,

[1]During the period our data covers, trip prices faced by riders consisted of a per-trip charge, and per-mile and per-minute base rates; when demand so outstripped supply as to threaten the health of the dispatch system, these were "surged" by a multiplier greater than 1. Tolls, taxes, and other fees were added afterward as required. Trip payments for a driver were the rider fare less a commission, with relevant tolls, taxes, and fees reimbursed to the driver without Uber retaining a percentage.

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

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 the same basic pattern of initial pass-through of fare increases but a zero or even somewhat negative long-run effect.

The main reason the hourly earnings rate stays approximately the same is that the fraction of hours-worked that are spent with paying customers—which we call "utilization"—falls; a 10% fare increase lowers utilization by about 7%, using a "static" estimate of the effect (Borusyak and Jaravel, 2016). This single point estimate masks how the market actually evolves—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.[3]

The fall in utilization following a fare increase is due to a combination of fewer trips demanded and drivers working more hours, which has a business-stealing effect on other drivers. The increase in hours-worked appears to be primarily an intensive margin response, with drivers working more hours but not exiting or entering the market at different rates. For a 10% increase in the fare, drivers eventually work 6% more hours.

The pattern of the effects on hours-worked speaks to key feature of the market—namely that higher utilization is costly to workers, at least at the range of utilization levels we observe in the data. To see this, 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 costly to drivers, the platform faces a trade-

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

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 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 an 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.

Despite the large observed changes in utilization, we find that in general, changes in utilization alone were not sufficient to clear the market following fare increases. Higher fares were partially "undone" by less frequent and/or more moderate surge pricing (Chen and Sheldon, 2015; Hall et al., 2016). However, this surge effect is relatively small, in that for a 10% fare increase, our static estimate is that the average surge rate falls by about 2%. Furthermore, despite the goal of surge pricing to keep wait times more or less constant, we find that this goal was not met. 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).

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 recent times, Uber has decoupled rider and driver trip prices with upfront pricing, but drivers continue to earn trip pay determined by trip time and distance. Our conclusions should remain relevant to future scenarios where drivers are paid per trip, regardless of the exact structure.

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 market shares, we find no evidence that this is the case.

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.[4] 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 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.[5] These costs are relevant to a market-designing platform trying to set an efficient price.

Our results speak to the larger question of why some platforms take on price-setting, despite the well-known challenges of doing so. Reducing price dispersion is typically welcome, and simply enabling price comparison—say by using technology to reduce search costs—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 bi-lateral monopoly 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

[4]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.

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

free entry on the supply side.[6] Centralized price-setting can avoid the wasteful monopolistic competition equilibrium, which might explain in part why the hybrid price structure we explore is common in designed markets.

The rest of the paper is organized as follows. Section 2 describes the empirical context. Section 3 presents a model. Section 4 presents the empirical results. Section 5 concludes.

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 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 product. It is also available in more cities that UberBlack. See Hall and Krueger (Forthcoming) 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 are sent to the nearest available driver. At the end of the trip, the fare is automatically charged to the passenger's credit card. Uber handles all billing, customer support, and marketing.

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, as well as 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 wait times, with the intent of inducing supply and rationing demand in order to keep those wait times from increasing too

[6] 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 pricing.

much. During "un-surged" periods, the multiplier is 1.0.[7] 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. During the period covered by our data, 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.

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., 2013).

2.1 Measurement

A primary outcome of interest is the average hourly earnings rate of drivers. To construct this measure, 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.

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

For driver revenue, we omit reimbursements for known tolls and fees (such as airport fees), and deduct Uber's commission. 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, without costs subtracted. 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 hours-driven or rides-taken 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. Some promotional payments unrelated to driving, like those earned for referring another driver, are omitted from the numerator.

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 for dispatch. Merely having the app open without marking oneself available for dispatch does not contribute to our measure of hours-worked. Because of the computer-mediated nature of the market, hours-worked (as well as time on trip) is measured essentially without error (Varian, 2010).

Our definition of hours-worked does not perfectly capture what we might regard as working. For example, some drivers make themselves available for dispatch while "commuting" to where they normally seek passengers (such as the central business district in a city), increasing our hours-worked and decreasing our estimate of the hourly earnings rate. Drivers also report driving with multiple ride-sharing platforms simultaneously, turning one app off immediately after being dispatched by the other platform. This multi-homing strategy will also tend to inflate hours-worked and thus lower the implied hourly earnings rate.[8] Although these definitional ambiguities as to what con-

[8] Despite these possibilities, we find no evidence that the size of direct Uber competitors—and thus presumably the opportunity for multi-homing—affects our results.

stitutes "working" require us to be careful in interpretation, it is important to note that similar issues also arise in conventional work settings.[9] Furthermore, as our primary interest is in changes rather than levels, biases that do not differ systematically with the base fare are not a concern.

2.2 Data description

We construct a panel of 36 US cities over 105 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 UberBlack. To construct our panel, we started with the 50 largest US cities by total trip volume and then removed from the panel cities which had substantial changes to the areas of service availability or significant within-city geographical variation in pricing.[10] 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 in fact represent single markets.

The panel is slightly unbalanced in that we lack early data for a few cities that had relatively late introductions of UberX. The panel begins with the week in which driver earnings data is first reliably available; prior to 2014-06-02, historical driver earnings cannot be reconstructed with sufficient confidence.

2.3 Variation in the base fare index and identification

Uber has changed the base fare for UberX in every city in the panel, with nearly all cities experiencing multiple changes. Figure 1 highlights the weeks

[9]For example, there is a long legal debate on whether hours spend 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.

[10]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.

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 decision to change the fare in a city was made after consultation with the Uber employees responsible for the city in question. They were advised by Uber's internal "pricing team," which reportedly considered metrics like driver utilization and the fraction of trips taken under surge conditions. For our purposes, this creates an obvious selection concern in using changes in the fare for identification.

Despite the possibility of selection, there are several justifications for treating fare change variation as exogenous. First, we note that as every city in the panel had fare changes, it is not the case that latent differences exist between the kinds of cities that have fare changes and those that do not. Second, Figure 1 shows that many changes took place in numerous cities nearly simultaneously, making it clear that highly city-specific explanations were not driving many of the fare changes. While clearly coordinated, these fare changes are also not exactly simultaneous—there are several cases in the data where cities all have fare changes within a fairly small window, but the precise timing differs by only a few weeks. This feature is useful for our purposes, as it seems unlikely that the precise sequence of cuts reflects important latent differences.[11] Finally, leaks by various media sources indicate that a relatively rudimentary spreadsheet-based analysis was used to model city outcomes, making it doubtful that decision-makers were confidently conditioning on future potential outcomes.

To the extent cities were selected for fare changes on the basis of observable attributes, we know approximately what those attributes are, and we can look for pre-treatment trends in those outcomes. For each of our outcomes, we find

[11] However, fare changes also differed in magnitude and not just timing, and so there is a non-time related exogeneity concern. We have no insight into how the precise percentages were chosen for various cities.

Figure 1: UberX base fare index changes for US cities, by week

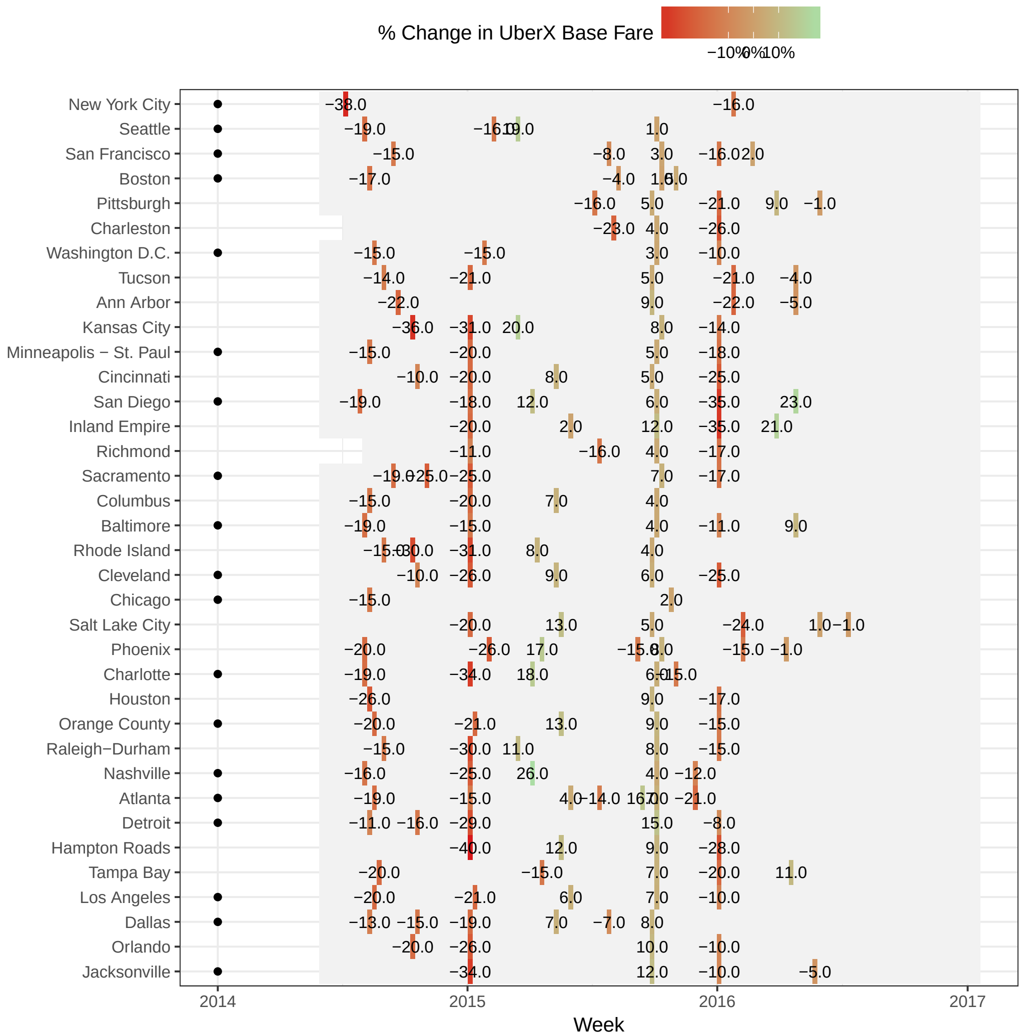

Notes: This figure indicates which cities in the panel had changes in the base trip price index, by week, and reports the size of that change, in percentage terms relative to the fare in the previous week. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The x-axis is in weeks. Whether the city had a viable UberBlack service is indicated by a dot next to the city name. Squares that are not shaded gray indicate that no data is available for that city week. Cities are listed in descending order of their average base fare over the entire panel. See Section 2.2 for a definition of the sample.

no evidence of a violation of the parallel trends assumption. To the extent cities were being selected in differed in trends, all of our regressions include city-specific trends (though for most outcomes, they appear unnecessary).[12]

3 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 search process generates, such as non-existent/multiple equilibria or industry scale economies e.g., Arnott (1996); Cairns and Liston-Heyes (1996). 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 the taxi industry discussed above, passenger demand depends on wait times, which adds substantial modeling complexity. We can pursue a simpler modeling approach—namely using a single demand curve—because dynamic pricing is intended to prevent the market from clearing on wait times.[13]

The unique features of ride-sharing markets—such as dynamic pricing and centralized algorithmic dispatch—make many of these search considerations less important. However, setting aside search and wait times is clearly a simplification of reality. As we will see, for the market conditions we observe, price changes have unambiguous—albeit relatively small—effects on wait times. From other evidence, it is clear that riders demand more rides when wait times are lower: Cohen et al. (2016) estimate that riders are only about a quarter as elastic to wait times as they are to prices. Thus, the wait time response produces a small countervailing effect to any price change that we do not model, though the fundamental dynamics of how the market adjusts are likely unaffected.

[12] We report regressions without these trends in Appendix B.

[13] 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.

We assume a fixed demand curve that depends only on price. Our treatment of driver labor supply is also simple, ignoring behavioral considerations, such as income targeting (Camerer et al., 1997; Thakral and Linh, 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.[14]

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 wage and $c(x)$ is the disutility of utilization, with $c'(x) > 0$ and $c''(x) > 0$.[15]

We assume that the supply of hours-worked depends on the 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.[16] We set aside any tax taken by the platform and assume that workers 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 $S(p; x) = xH(U(px, x))$.

[14]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.

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

[16]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, ignoring income effects seems like a reasonable simplification.

Market clearing requires that

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

Figure 2 shows a market in equilibrium, with a price of p^*, labeled as point A in the figure.

Figure 2: Equilibrium in a ride-sharing market after a decrease in the base fare

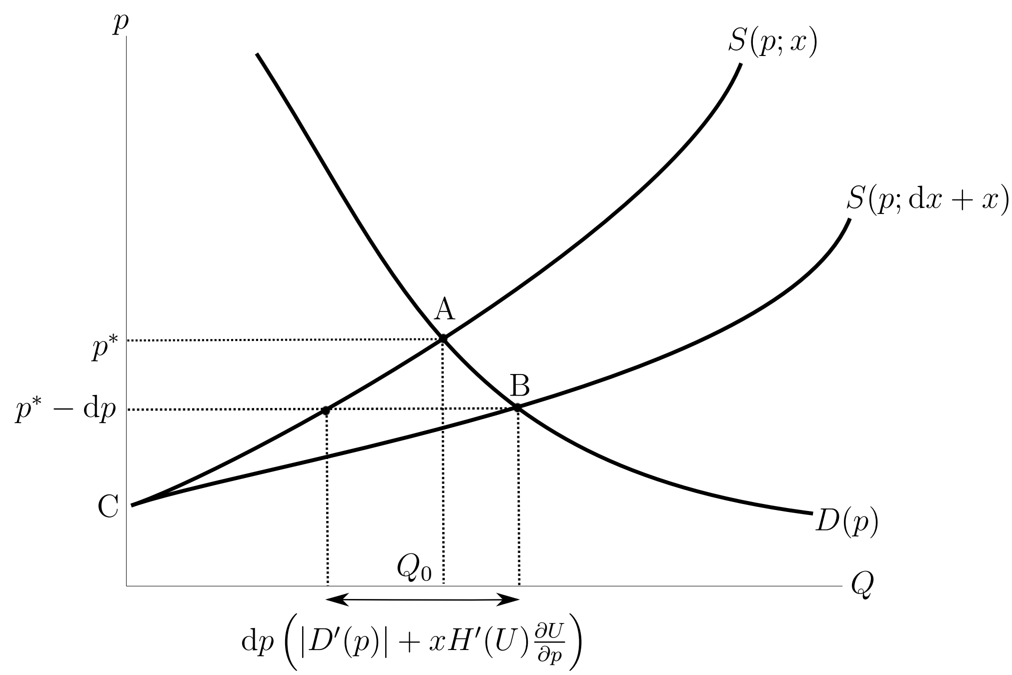

Notes: The $D(p)$ curve is the market demand for hours of transportation service, when the price, or fare, is p. The $S(p; x)$ curve is the hours of transportation supplied when the price is p and the driver utilization level is x. Q_0 is the market equilibrium quantity of hours of transportation service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

3.1 Effects of a fare decrease

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 $x^2 H'(U)\,\mathrm{d}p$. The double-headed arrow in Figure 2 shows the gap in hours of the service that must be closed for the market to return to an equilibrium.

There are several 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. A third possibility is that utilization could increase, with a higher x pushing out the supply curve.

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 2. 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,

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

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 2. In Figure 3 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, or $D(p)$. 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

Figure 3: Gap between demand and supply as a function of utilization, for different product market prices

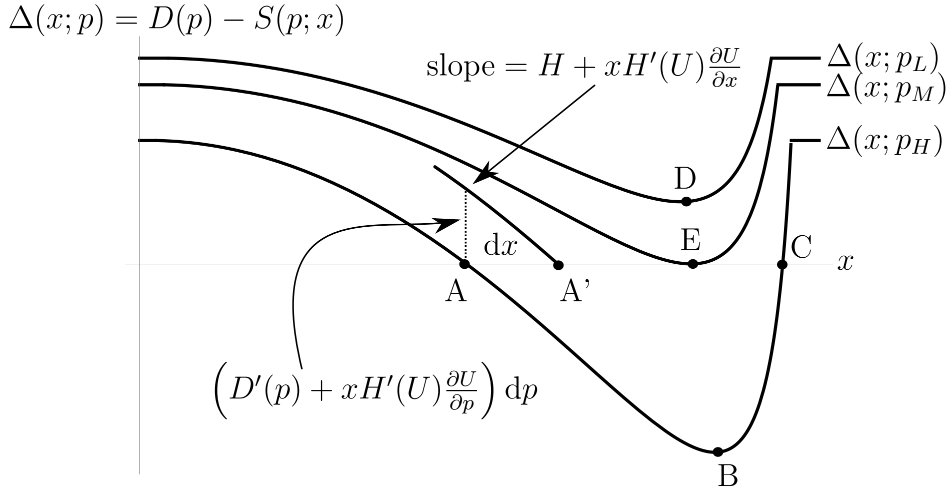

Notes: This figure illustrates the gap between demand and supply of the service at different levels of utilization, for a given product market price. Three prices are shown: p_L, p_M, and p_H. The change in the equilibrium from point A to A' corresponds to a decrease in the product market price from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

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 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 point where $px - c(x) = 0$ i.e, no workers 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 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.

In an already-clearing existing marketplace, it is likely that the market would be at an equilibrium like the A or C equilibrium, as p_M requires getting the price just right and a price like p_L cannot support an equilibrium. Assume that $p = p_H$. 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 2. Returning to Figure 3, assume the market is in equilibrium, at A. Following the fare decrease, the new equilibrium occurs at A', which has a dx higher utilization.

Graphically, we can see that when the initial gap between supply and demand is larger, say because of highly elastic demand, 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 3, 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.

3.2 Efficient product market price

As the platform sets p, a natural question is what price is efficient. In Figure 2, lowering p increases surplus by the triangle ABC, suggesting that the lowest possible p is efficient. From Figure 3, this lowest possible p was p_M, which allowed for an equilibrium at point E. We can show that this point is surplus-maximizing directly by computing the area of the ABC curve. Let $TS(\hat{p})$ be the total surplus at market price $\hat{p}$. If we differentiate $TS(\hat{p})$ by $\hat{p}$ and solve the first order condition, we have

$$\begin{aligned}\frac{\partial TS(\hat{p})}{\partial \hat{p}} &= \frac{\partial}{\partial \hat{p}} \int_0^{\hat{p}} \left(H + xH'(U)\frac{\partial U}{\partial x} \right) \frac{\mathrm{d}x}{\mathrm{d}\hat{p}} \,\mathrm{d}p \\ &= \frac{\mathrm{d}x}{\mathrm{d}\hat{p}} \big(H + H'(U)(p - c'(x)) \big) = 0.\end{aligned}$$

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

$$p^* = c'(x) - \frac{H(U)}{xH'(U)}. \tag{3}$$

The $\frac{\mathrm{d}x}{\mathrm{d}\hat{p}} = 0$ 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 is described by Equation 3, which states, in elasticity terms, that at the optimum, $|\epsilon_U^H \epsilon_p^U| = 1$. In words, the optimal price is where any further decrease in price cannot increase total output. The condition for the optimal price given in Equation 3 is the same that characterizes the minimum of the $\Delta(p; x)$ curves from Figure 3.

3.3 Driver labor supply response

The driver labor response to changes in the base fare depend 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, by assumption $w = px$, and de-couple the hourly earnings rate and the utilization, with drivers simply paid w regardless of the utilization or product market price.

Figure 4 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 4: Optimal product market price and utilization

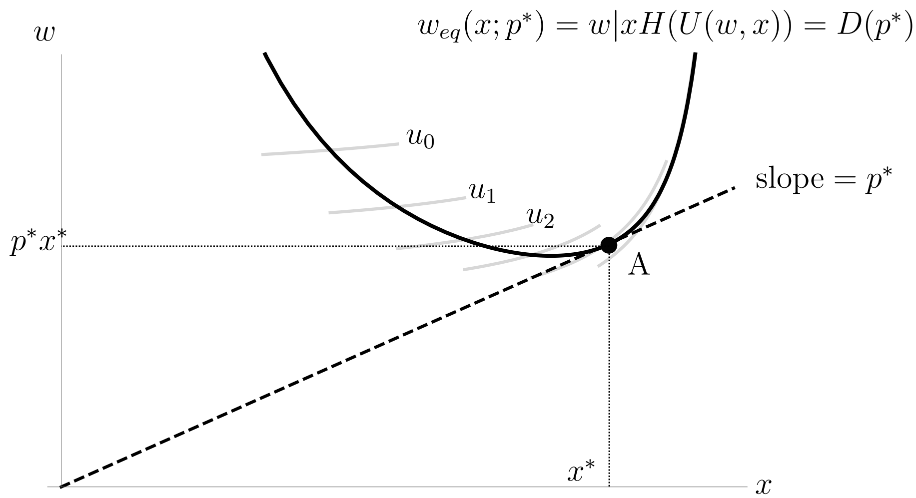

Notes: This figure illustrates shows several driver indifference curves with respect to the hourly earnings rate and the utilization. The heavy line, $w_{eq}(x;p^*)$, illustrates combinations of w and x that would clear the market for a given price p^*.

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 intensities. 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 4, 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

$$\frac{\mathrm{d}w_{eq}}{\mathrm{d}x} = c'(x) - \frac{H(U)}{xH'(U)}. \tag{4}$$

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 4, combined with Equation 3, which characterizes the efficient price, implies that

$$\frac{\mathrm{d}w_{eq}}{\mathrm{d}x} = p^*, \tag{5}$$

or that the $w_{eq}(x; p^*)$ curve is tangent to a line with a slope that is the product market price. In Figure 4, we indicate this tangency point at A. 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 at A in Figure 4 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 3. 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 Results

In this section, we present estimates of the effects of changes in the base fare on market outcomes, using a between-city analysis of the UberX market. Note that unlike in our model setup, the platform does not choose p directly, as this is affected by surge. The platform changes a base fare index, BaseFare, with the actual p faced by passengers, on average, being $p = m \times \textsc{BaseFare}$, where m is the average surge multiplier. If the market can clear though changes in utilization, then this distinction does not matter, as $m = 1$ and the demand curve remains fixed.

To start, we focus 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.[17] 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.

4.1 Panel-wide averages over time

Before presenting the panel regression results, 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 5 shows, from top to bottom, the mean base price index, hourly earnings rate, utilization, average surge, and median wait

[17] 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.

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. In the panel below, we see that during that same period, the hourly earnings rate time series exhibits no similar trend. In contrast, driver utilization has increased substantially over the same period. There is little systematic change in average surge levels. In the bottom panel, we see that wait times were high early in the panel, but fell substantially by early 2015 and then are more or less constant afterwards.[18]

The patterns shown in Figure 5 preview some—but not all—of the main results from our regression analysis—namely little change in the hourly earnings rate despite large changes in the base fare index. Note that the two large drops in the base fare index occur at the start of 2015 and 2016. Immediately after, average surge increases, as does utilization, though only utilization seems to show a persistent change in levels. The driver hourly earnings rate series appears more or less stable, as does the median wait time series (following a decrease in early 2014).

4.2 Effects of fare changes on the hourly earnings rate, utilization, and surge

To begin our analysis of the effects of fare changes, we set aside the possibility that market adjustment takes time and simply report "static" estimates of the effects of fare changes. Our preferred specification is

$$y_{it} = \alpha_i + \beta_1 \log \textsc{BaseFare}_{it} + \gamma_i t + \delta_t + \epsilon_{it}, \tag{6}$$

where y_{it} is some market-level outcome of interest in city i during week t, α_i is a city-specific fixed effect, $\textsc{BaseFare}_{it}$ is the base trip price index, γ_i is a

[18]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 5: Average UberX market attributes over time for the US city-week panel, as indices

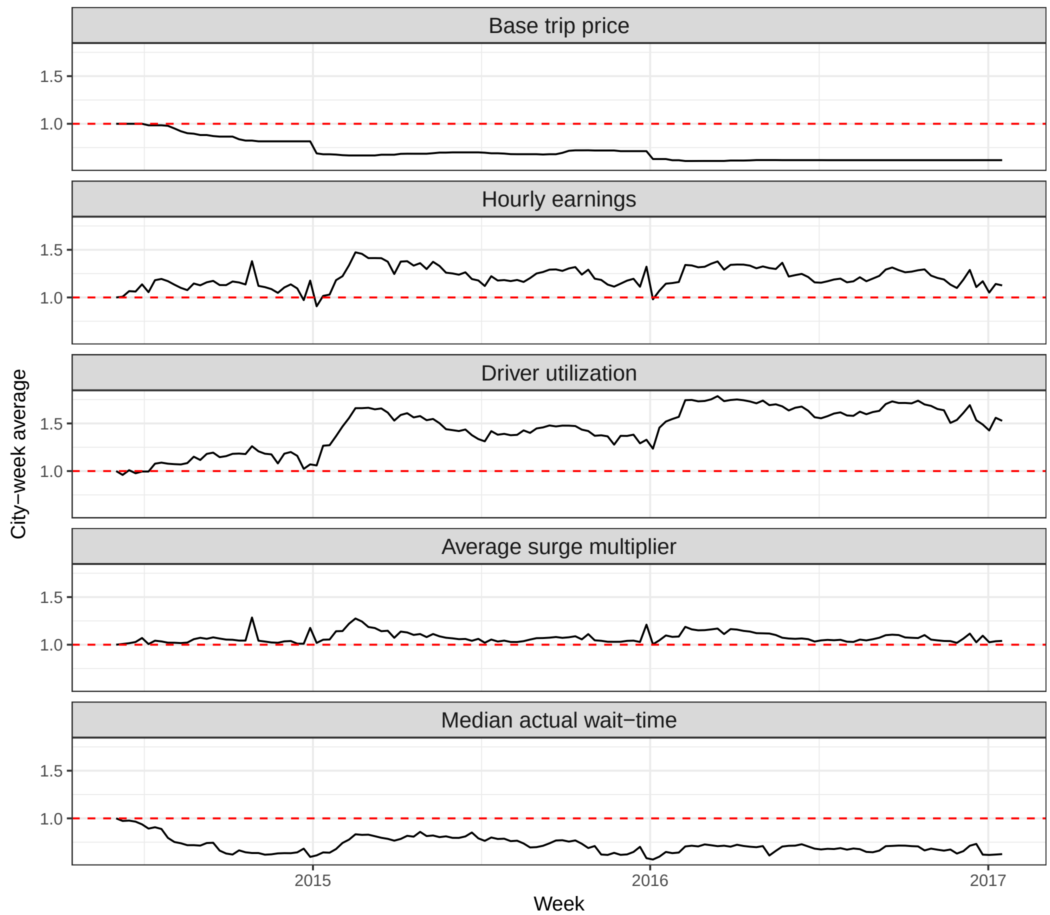

Notes: This figure plots the city-week panel weekly average for a collection of UberX market outcomes. All cities are weighted equally—see Section 2.2 for a definition of the sample. All series are turned into an index with a value of 1 in the first week. The base fare index is the price to passengers 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.

city-specific linear time trend and δ_t is a week-specific fixed effect.[19]

[19] In the Appendix B, we report all regressions reported here, but without city-specific time trends. Generally, these trends improve the precision of the estimates (particularly for

Table 1 reports estimates of Equation 6 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.[20]

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 rate by just 0.7%. 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).

A natural question is how the presence of alternative ride-sharing platforms might affect the results. Alternative platforms would presumably make both drivers and riders more elastic. Despite this possibility, we have no evidence this is the case—interacting Uber's imputed at-the-moment ride-sharing market share with the price index has no detectable effect on the point estimates. See Appendix A.2 for this analysis.

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 some floor amount. We explore the effects of fare changes on measures of driver hourly earnings that include these promotional payments in Appendix A.1. 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.

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,

---

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 Appendix B we show residual plots for market quantity regressions. These diagnostic plots show that our preferred specification fits the data well for all cities.

[20]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.

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 6. 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 Section 2.2 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 $p \leq .001 : ***$.

estimating

$$y_{it} = \alpha_i + \sum_{\tau=\text{NumPre}}^{\text{NumPost}} \beta_\tau \log \text{BaseFare}_{it-\tau} + \delta_t + \gamma_i t + \epsilon_{it}, \tag{7}$$

where α_i is a city-specific fixed effect, BaseFare_{it} is the fare index in city i at time t, τ the number of weeks from the focal week, γ_i is a city-specific linear time trend and δ_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. We impose the restriction when estimating the model that $\sum_{\tau=\text{NumPre}}^{0} \hat{\beta}_\tau = 0$ i.e., that the cumulative effect in the week prior to the fare change is 0. This allows for cities having fare changes to differ from those not having changes 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.[21]

The implied weekly effects from Equation 7 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.[22] For each outcome, the long-run effect from Table 1 is plotted at the 0 week, which is the week that actual fare was changed.

In the top panel of Figure 6, the outcome is the log hourly earnings rate. Examining the pre-period, there is no obvious trend. Following a fare increase,

[21] In Appendix B 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.

[22] In Appendix B, we report our preferred regression specification but vary the post-period bandwidth. The various plots illustrate the 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. As such, we do maintain our current pre-period window length even in the direction comparisons. Given 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.

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 6. 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.

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

Although we lack 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 reduces costs, as it depends on the driver's vehicle. However, we can make some assumptions to construct a measure of

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

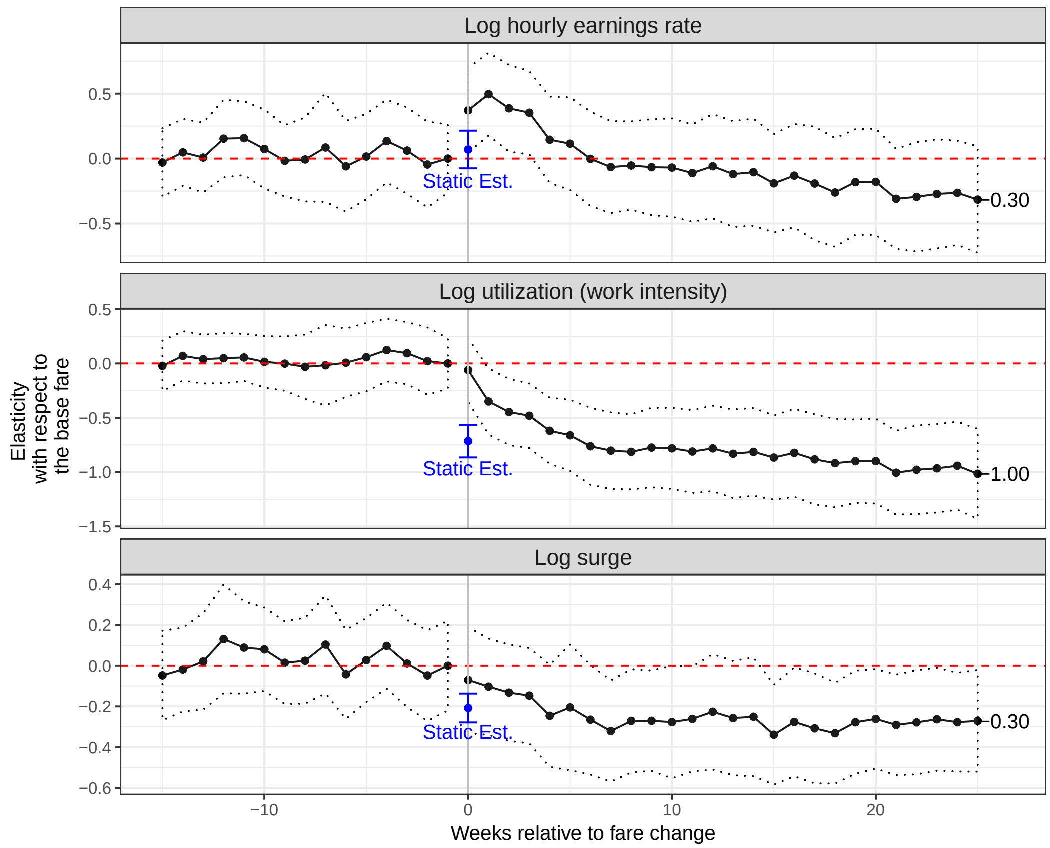

Notes: This figure plots the by-week cumulative effects of changes in the UberX base fare on the log hourly earnings rate, the log average utilization, and the log surge, along with 95% confidence intervals. These effects are from an estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 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 6.

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

$$\hat{c}(x) \approx \left(x_{it}s_i^{\text{PAX}} + (1-x_{it})s_i^{\text{NOPAX}}\right)\bar{C} \tag{8}$$

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 8, 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

$$\frac{\partial \log \hat{c}(x)}{\partial \log x} = \frac{\Delta s_i x_{it}}{1 + x_{it}\Delta s_i}, \tag{9}$$

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

The effect of fare changes on net hourly earnings depend 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.[23] Using this rate and the city-specific speed data, we calculate measures of *net* hourly earnings. For the inactive speed, we apply the 30% adjustment to the active speed that week.

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

[23]https://www.irs.gov/newsroom/standard-mileage-rates-for-2018-up-from-rates-for-2017

make use of city-specific measures of changes in driving speed with respect to utilization.

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

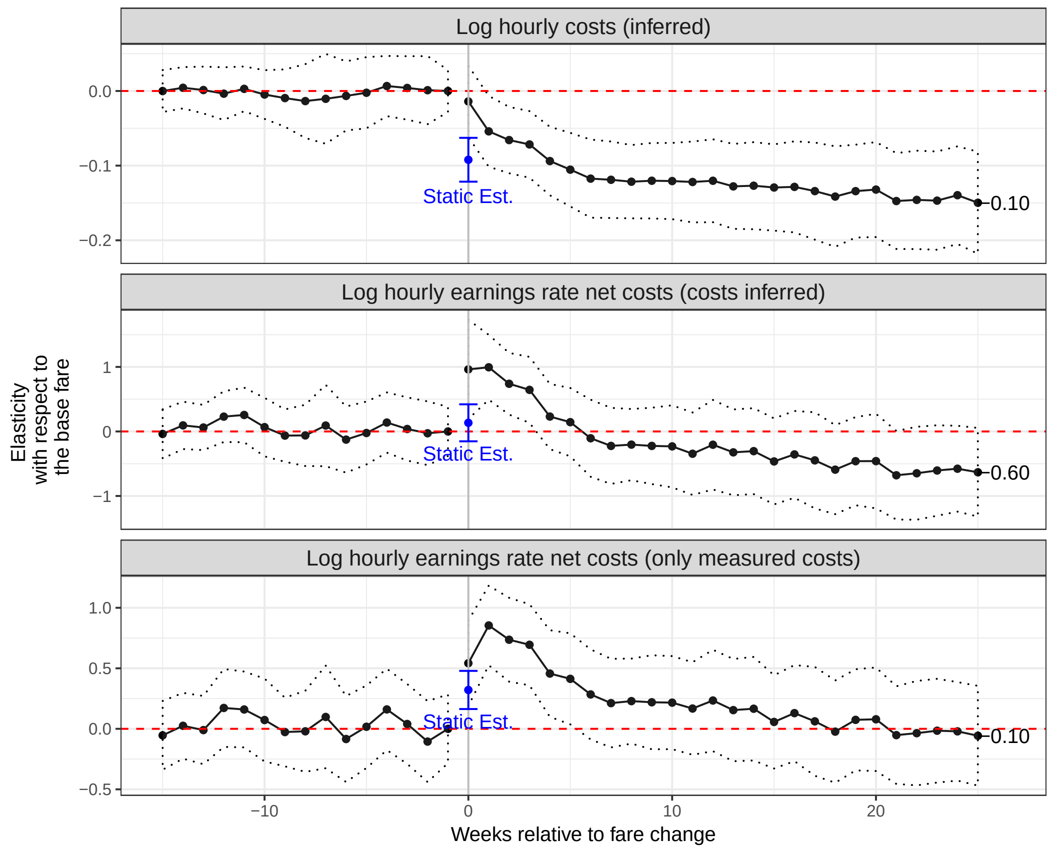

Notes: This figure plots the by-week effects of changes in the UberX base fare on drivers costs, from an estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 for a description. The x-axis are weeks relative to a fare change. The base fare index is the price to passengers of an 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 6.

In middle panel, the outcome is the log hourly earnings rate, net of direct costs, with costs inferred using the Equation 8 method. With the inclusion of costs, the denominator is smaller, and so the range of effect sizes 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. However, there is now stronger evidence that hourly earnings actually goes below its previous level.

The Equation 8 method likely over-states the per-mile costs drivers face when they are not with passengers. For one, a lighter load increases fuel efficiency. 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 end a shift or 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.

4.4 Driver labor response

Given the change in hourly earnings and costs, we expect driver utility to change and hence hours-worked to 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 intensive margin, 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 churn measure returns back to close to 0.

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

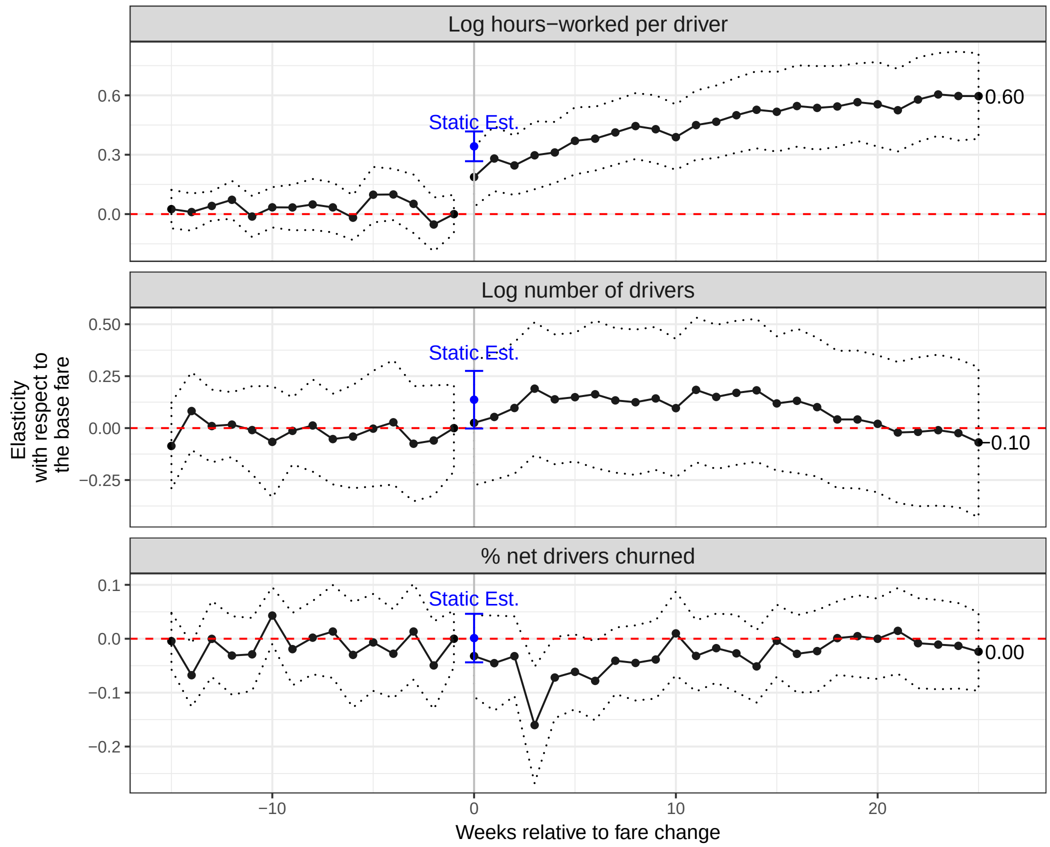

Notes: This figure plots the effects of changes in the UberX base fare on hours-worked per driver, log number of active drivers and the percentage of drivers churning. These effects are from an OLS estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 for a description. The x-axis 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 6.

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 Section 4.3. 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.

4.5 Passenger wait-times

As we discussed when presenting the model in Section 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. One demand curve shifter could be passenger wait time, as shorter wait times are preferred to longer wait times by all would-be passengers. We explore this possibility in Figure 9, where the outcome is the median wait time. We use median wait times rather than averages as reported wait-times are prone to outliers.

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 as the number of drivers does not decrease). These changes in wait times imply that markets could not clearly entirely through changes in utilization, and that furthermore, surge pricing did not, as implemented during the period covered by the experiment, hold product attributes exactly fixed.

4.6 Overall 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. How-

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

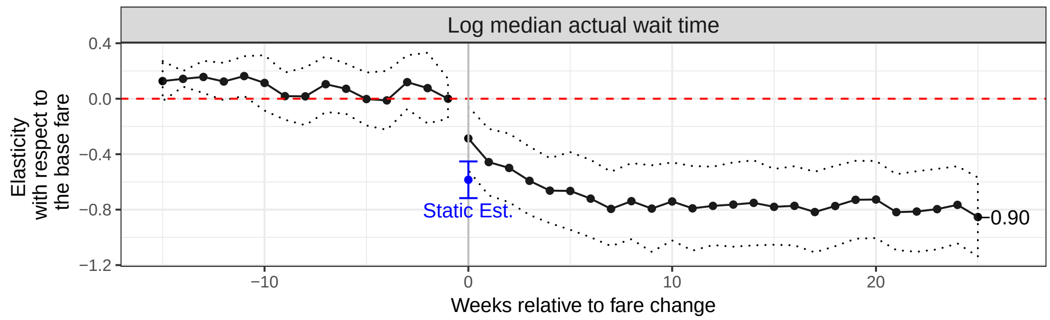

Notes: The outcome in this figure is the log median wait time (in seconds). These effects are from an OLS estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 for a description. The x-axis are weeks relative to a change in the base fare index. The base fare index is the price to passengers of an un-surged, 6 mile, 16 minute trip in that city, in that week. The horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to an estimation of Equation 6. 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.

ever, when outcomes are logged and a city-specific linear trend is included, a regression model can fit the data reasonably well.[24] In contrast to our "rate" 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 incor-

[24]See Appendix B.2 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 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.

rect.[25]

Figure 10: Effects of fare changes on market quantities

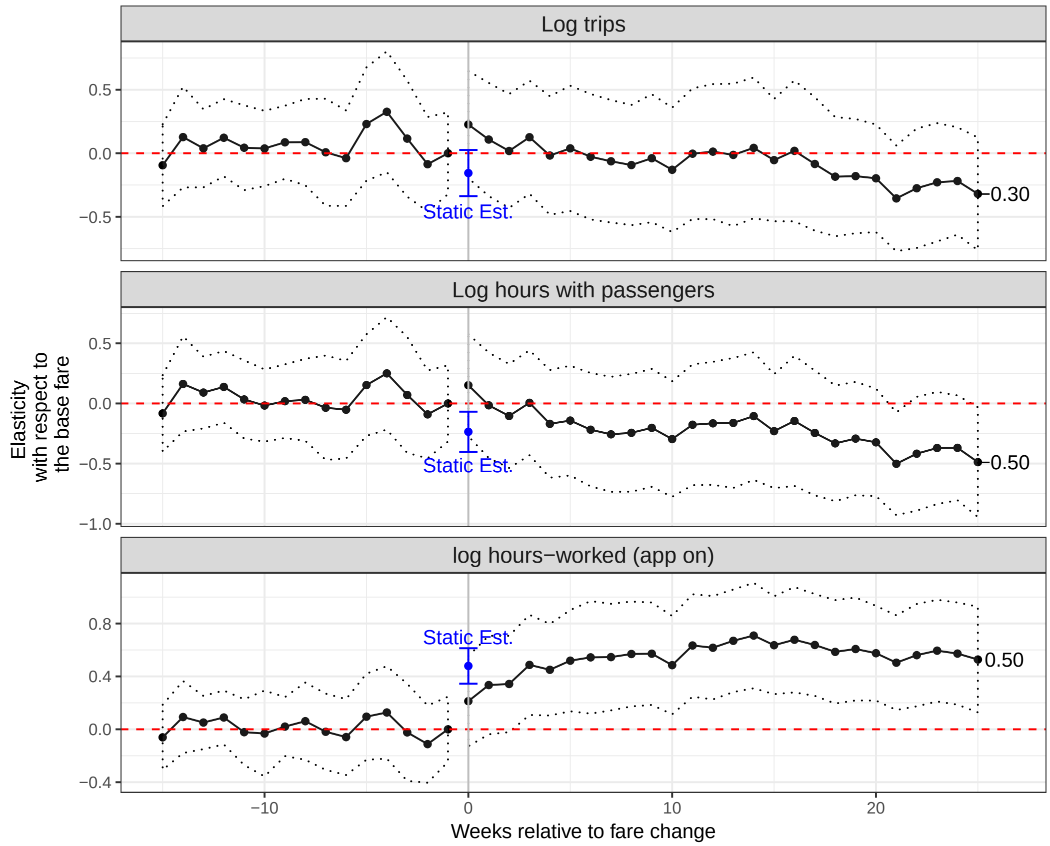

Notes: The outcomes in this figure are the log number of trips, the log number of hours with passengers, and the log number of hours-worked. These effects are from an OLS estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 for a description. The x-axis are weeks relative to a 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 horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to an estimation of Equation 6. 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.

In the bottom panel, the outcome is log hours-worked. It rises immediately

[25]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).

after a fare increase and continues to climb before flattening out with a long-run 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.

5 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 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.

In our empirical analysis, we assumed that Uber was changing fares without conditioning on market-specific attributes. When ride-sharing markets were less mature and there was less practical experience in managing these markets, assuming Uber would change fares as-if at random is plausible. With hindsight about how the market seems to adjust, Uber might make different choices—namely by picking 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. However, our evidence that lowered fares increased wait times and lead to more surge suggests that some fare cuts likely went “too far,” 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 fewer drivers.

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 "forward dispatch" (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. With the move to up-front pricing, Uber could potentially charge more for utilization-reducing trips, such as those to areas where a return trip with a paying passenger is relatively unlikely.

To the extent we think of Uber's fare selection problem as applying a markup rule to its unit costs, technological improvements in utilization would lower those unit costs, implying that the optimal fare adjustment has typically been a fare reduction. This line of argument provides an intriguing as-if explanation for why Uber has continually reduced the base fare over time. It would also explain why, if anything, the hourly earnings rate rose following fare cuts—if Uber was cutting fares because a higher utilization was obtainable, the pre-cut fare was unprofitably in the elastic region of the demand curve. However, the increase in the hourly earnings rate is also consistent with higher utilization being costlier for drivers, making this line of reasoning fairly speculative.

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

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.

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.

Borusyak, Kirill and Xavier Jaravel, "Revisiting event study designs," *Working Paper*, 2016.

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 D Levin, and Neel Sundaresan, "Sales mechanisms in online markets: What happened to Internet auctions?," Technical Report, National Bureau of Economic Research 2013.

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.

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*, Forthcoming.

_ **, 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.

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 Tô Linh, "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.

A Additional analyses

A.1 The role of promotional payments

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 some floor amount. Given their structure, we might expect these guarantees to act as automatic stabilizers, counteracting the immediate effects of a fare change on "organic" earnings. It is important to note that this with-promotions included measure is problematic, in that it is not the actual marginal payment drivers face.

To explore the role of these promotional payments on market adjustment, in Figure 11, 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 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.

A.2 Effects of ride-sharing competitors on main results

It is beyond the scope of this analysis to try to model the competition between ride-sharing platforms and the larger for-hire industry. 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 market. Our estimates of Uber's market 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. For cities in which no competitor was operating that week, we impute Uber's share as 1.

Figure 11: Effects of fare changes on hourly earnings with and without promotional payments included

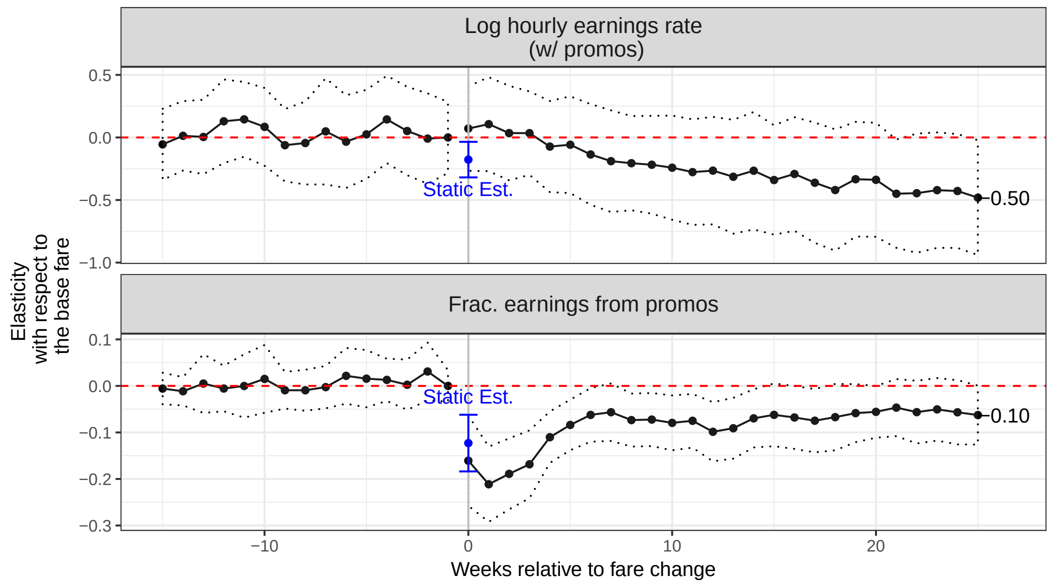

Notes: This figure plots the effects of changes in the UberX base fare index on three outcomes (from left to right): (1) the log hourly earnings rate not including promotional payments, (2) the log total hourly earnings rate includes promotional payments, and (3) the fraction of hourly earnings coming from promotional payments. These effects are from an OLS estimation of Equation 7. The sample is a panel of US cities—see Section 2.2 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 horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to an estimation of Equation 6. 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.

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 market share in that city that week. For all outcomes, the level of Uber's market share has no detectable effect on the point estimate. The

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 6. 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 Section 2.2 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 $p \leq .001 : ***$.

results suggest that the degree of rivalary in the market had no discerible effect on how Uber's marketplace adjusted following fare changes.

B Online Appendix

B.1 Alternative specifications

Figure 12: Alternative specifications for Figure 6

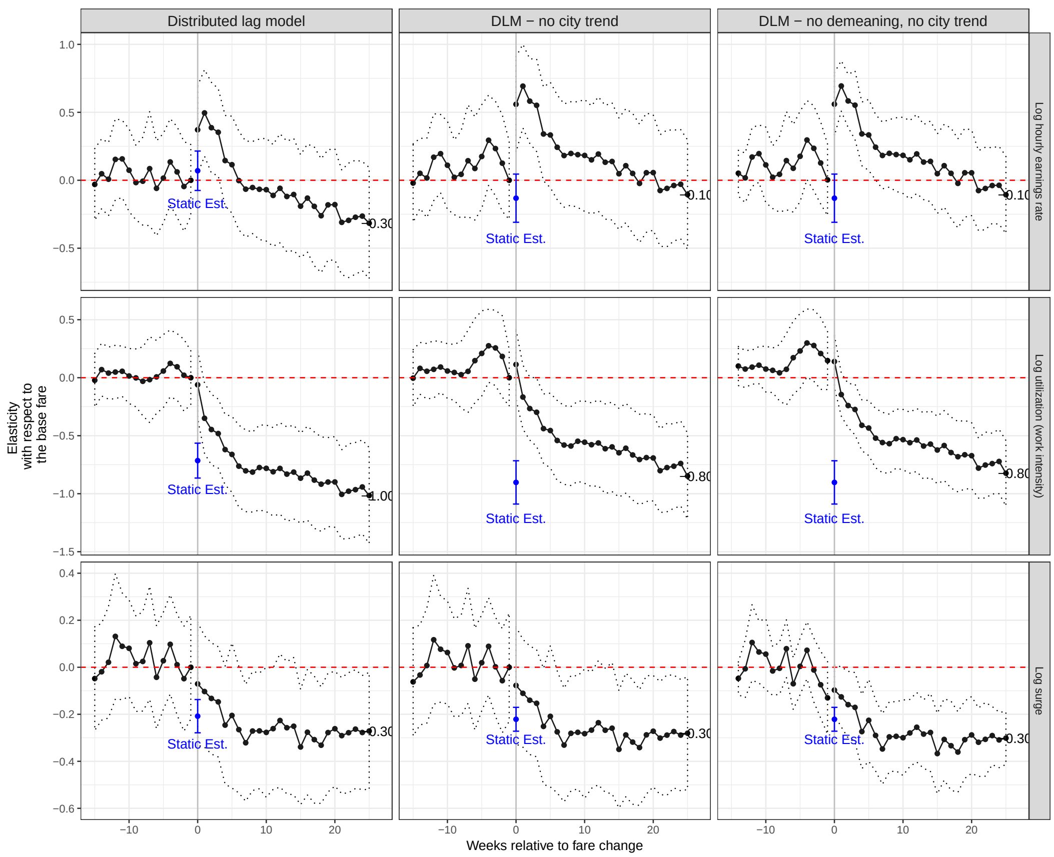

Notes: Alternative specifications.

B.2 Residual plots

Figure 13: Alternative post-period bandwidths for Figure 6

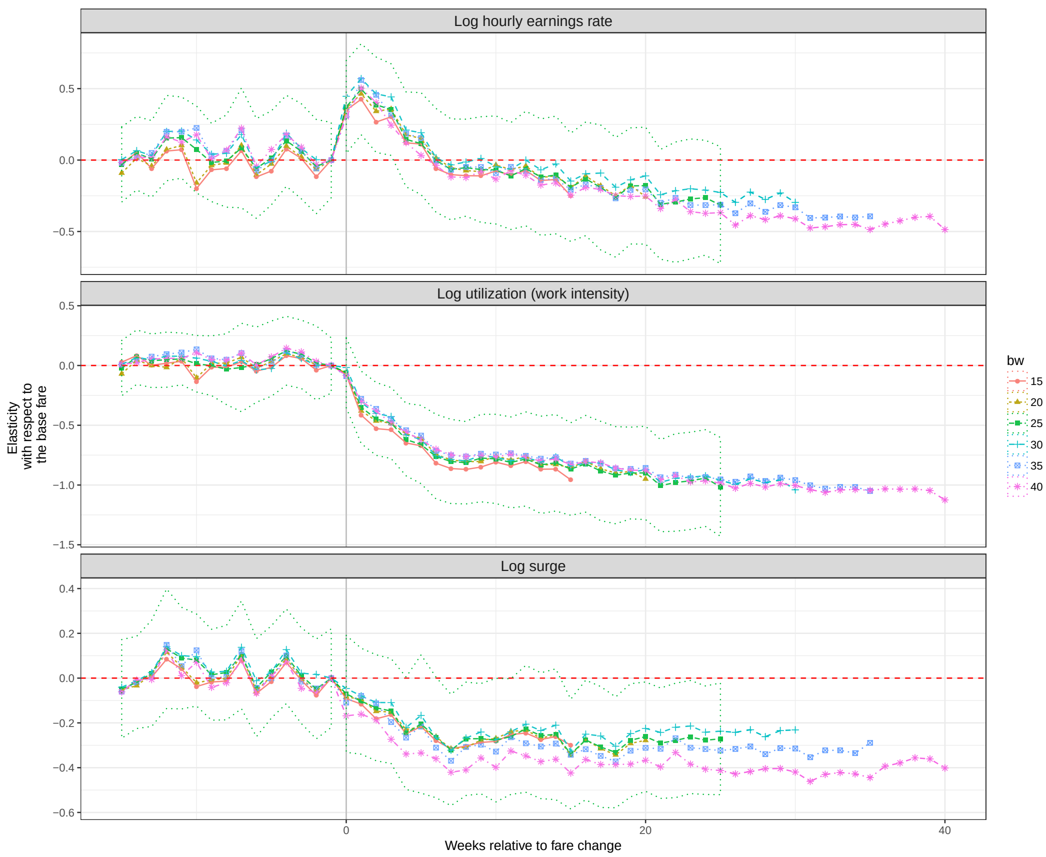

Notes: Alternative post-period bandwidths.

Figure 14: Alternative specifications for Figure 7

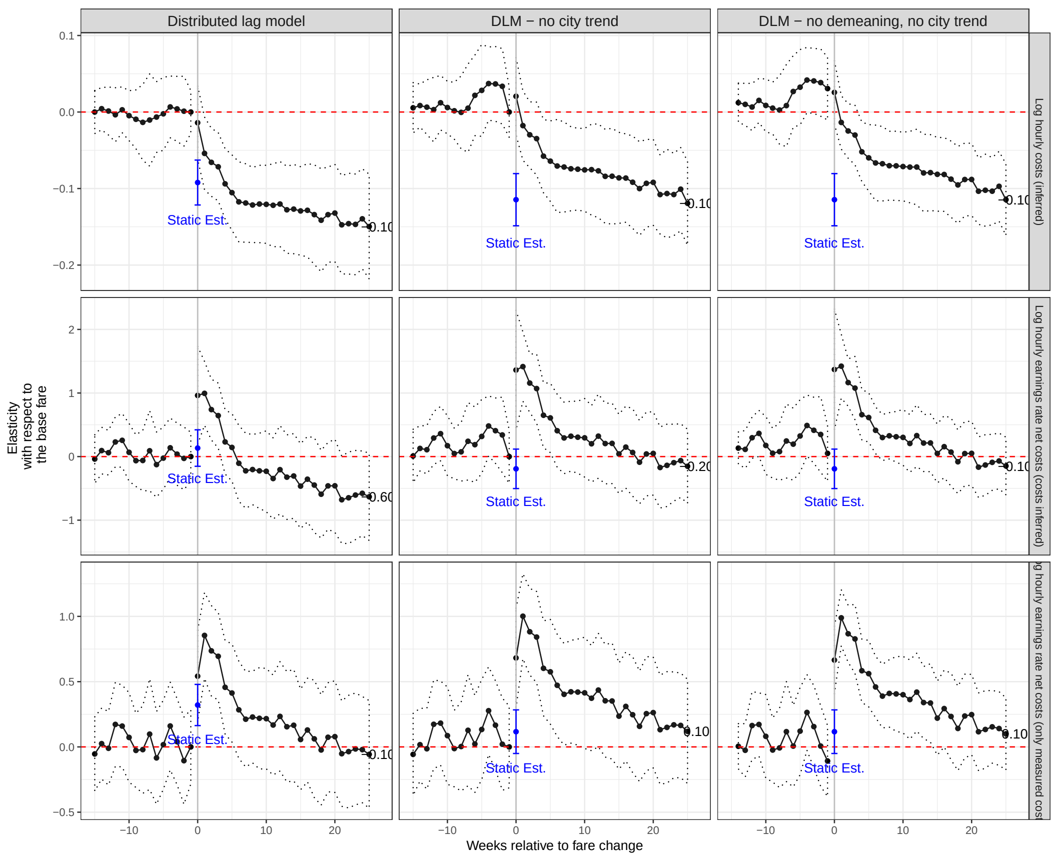

Notes: Alternative specifications.

Figure 15: Alternative bandwidths for Figure 7

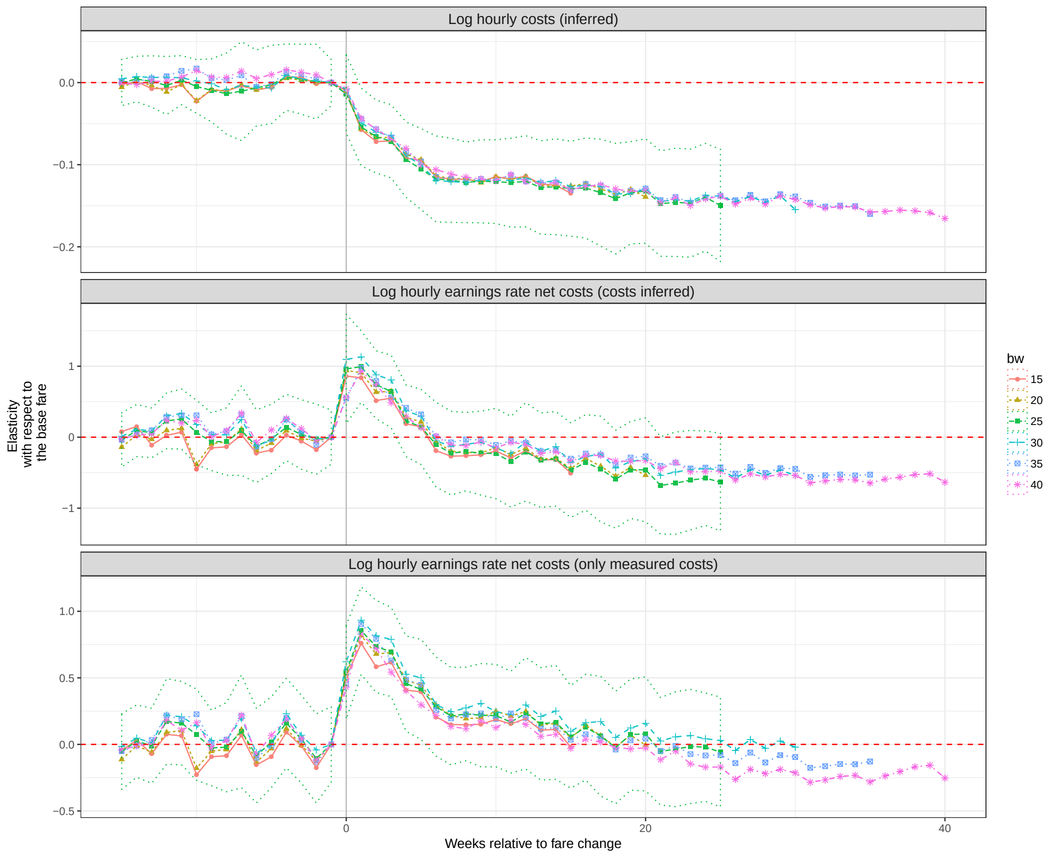

Notes: Alternative post-period bandwidths.

Figure 16: Alternative specifications for Figure 8

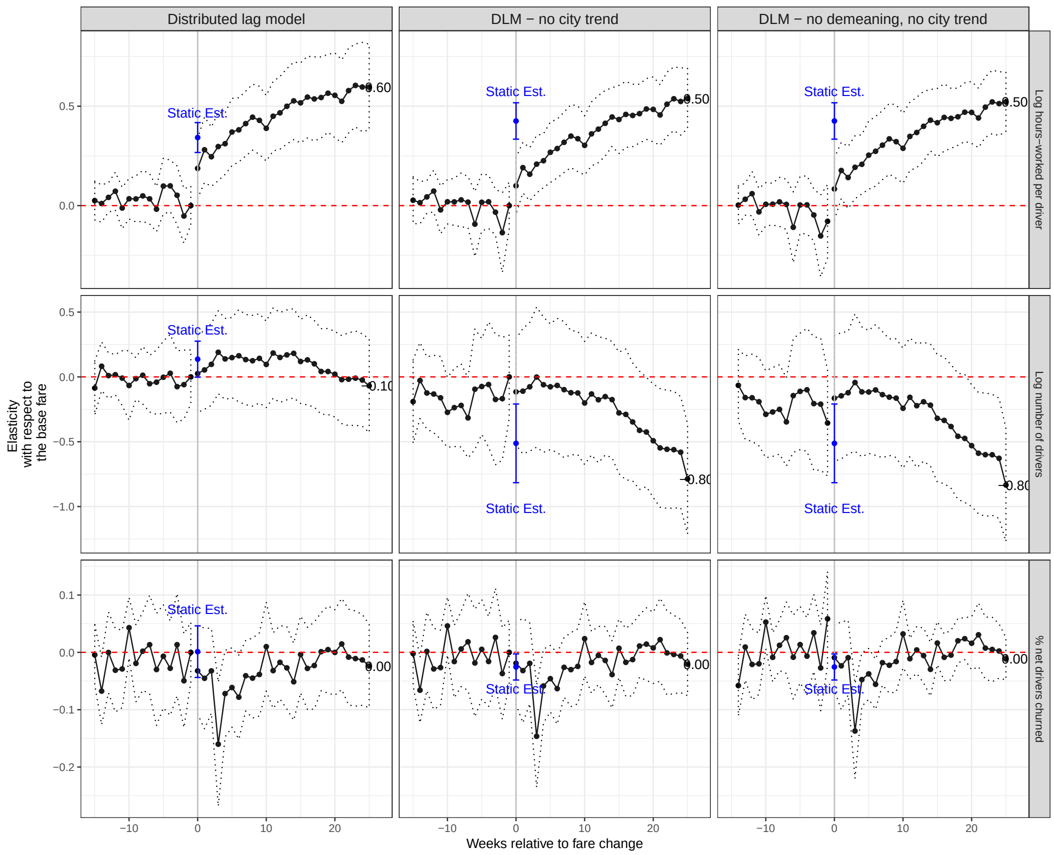

Notes: Alternative specifications.

Figure 17: Alternative bandwidths for Figure 8

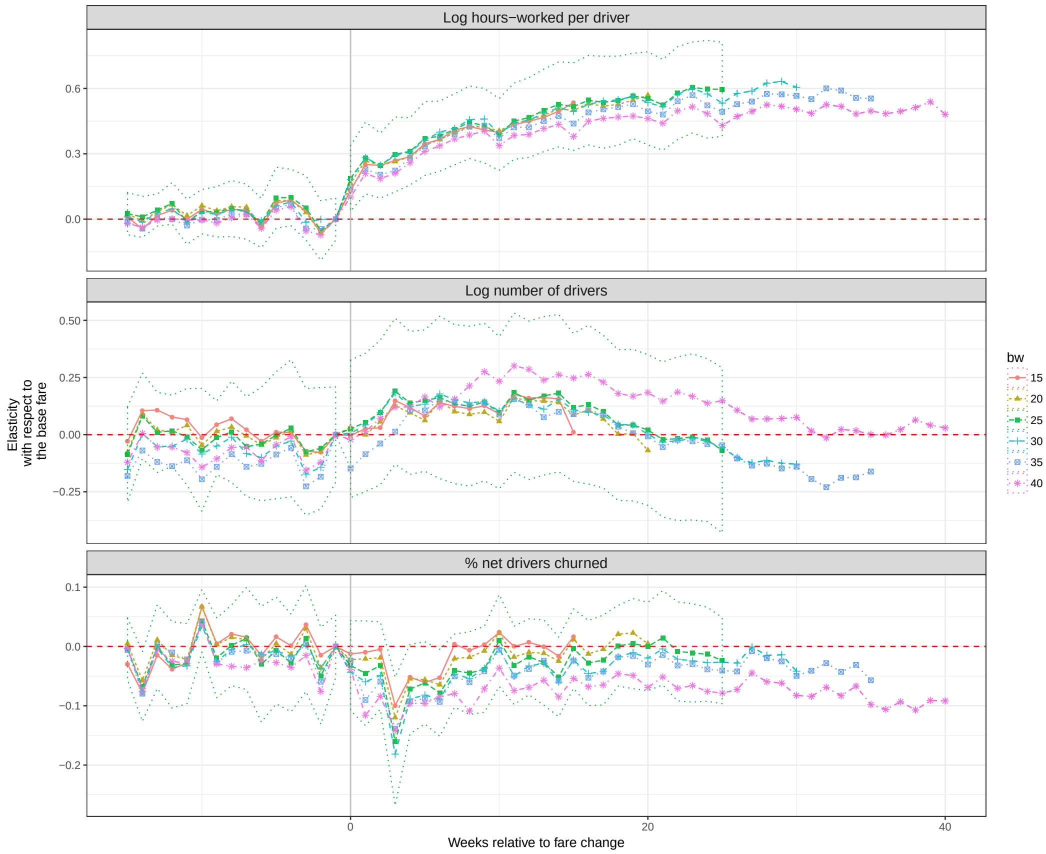

Notes: Alternative bandwidths.

Figure 18: Alternative specifications for Figure 9

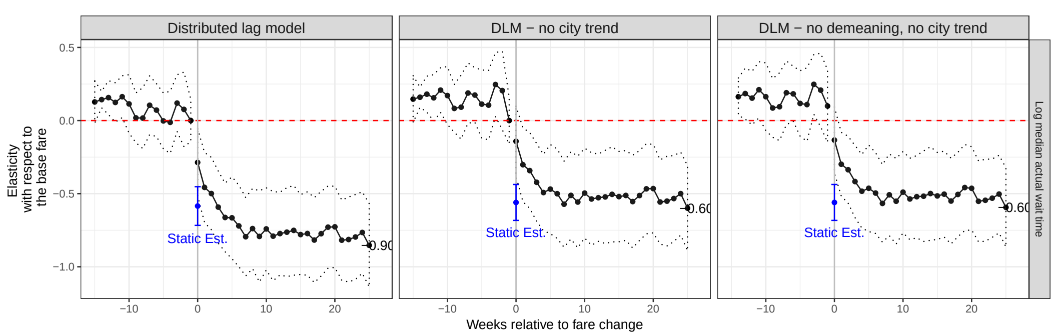

Notes: Alternative specifications.

Figure 19: Alternative bandwidths for Figure 9

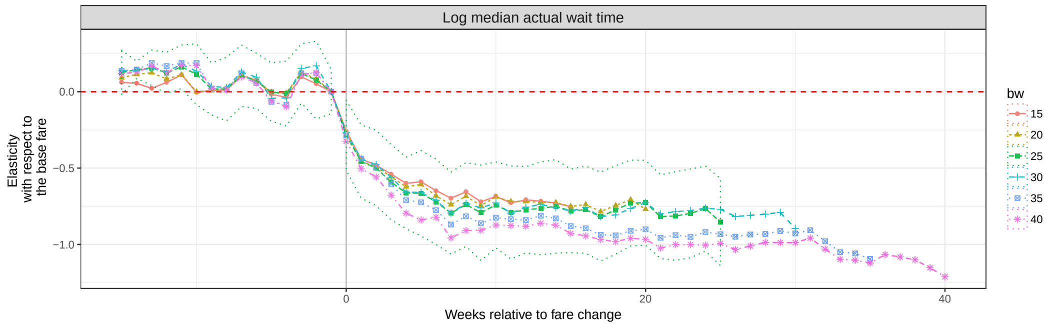

Notes: Alternative bandwidths.

Figure 20: Alternative specifications for Figure 10

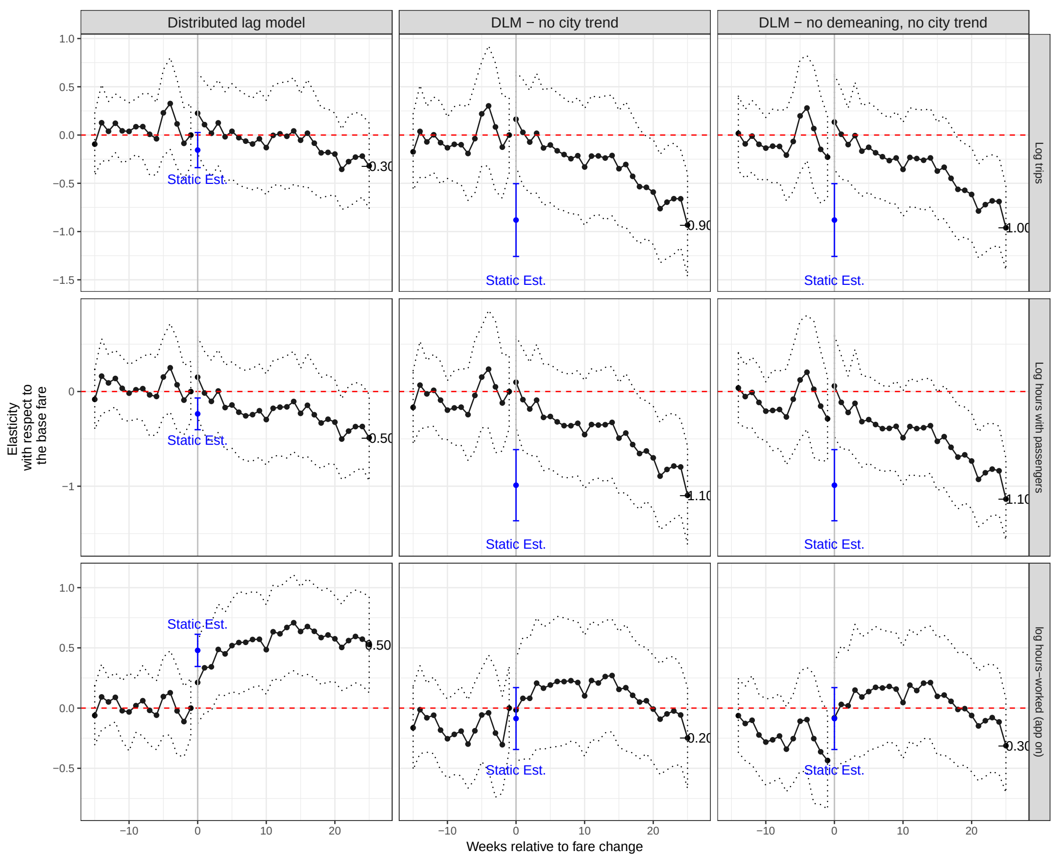

Notes: Alternative specifications.

Figure 21: Alternative bandwidths for Figure 10

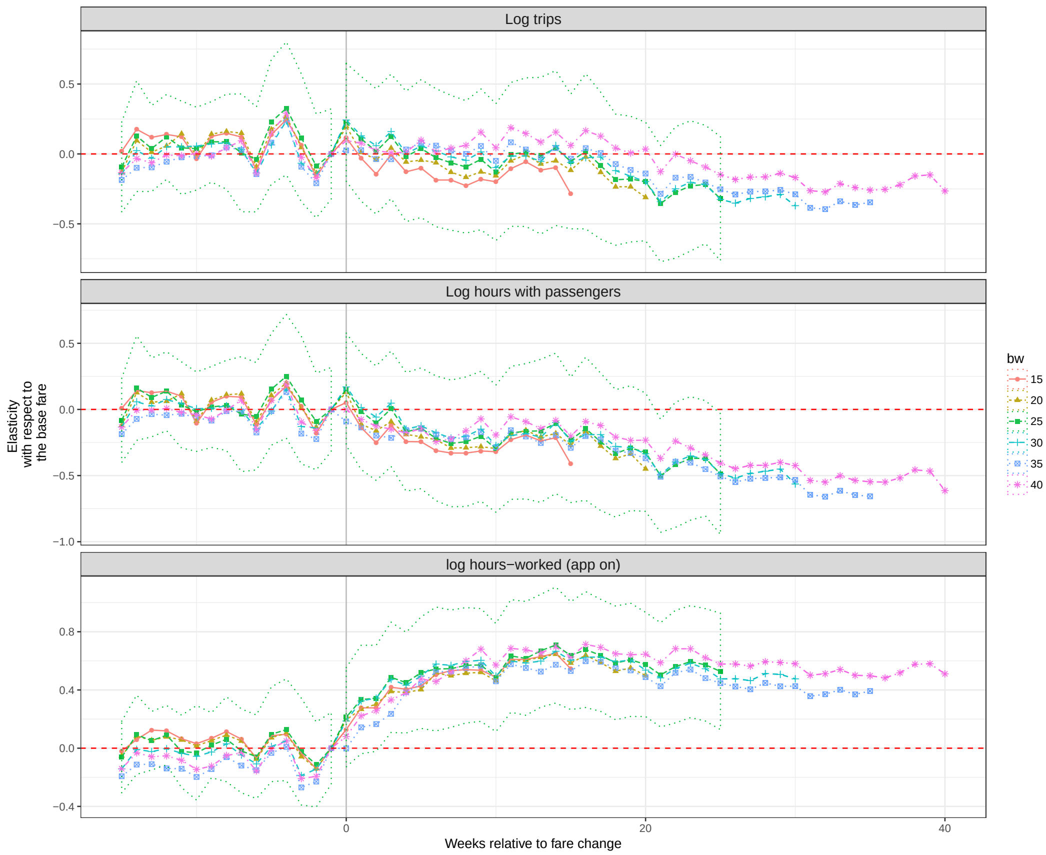

Notes: Alternative bandwidths.

Figure 22: Alternative specifications for Figure 11

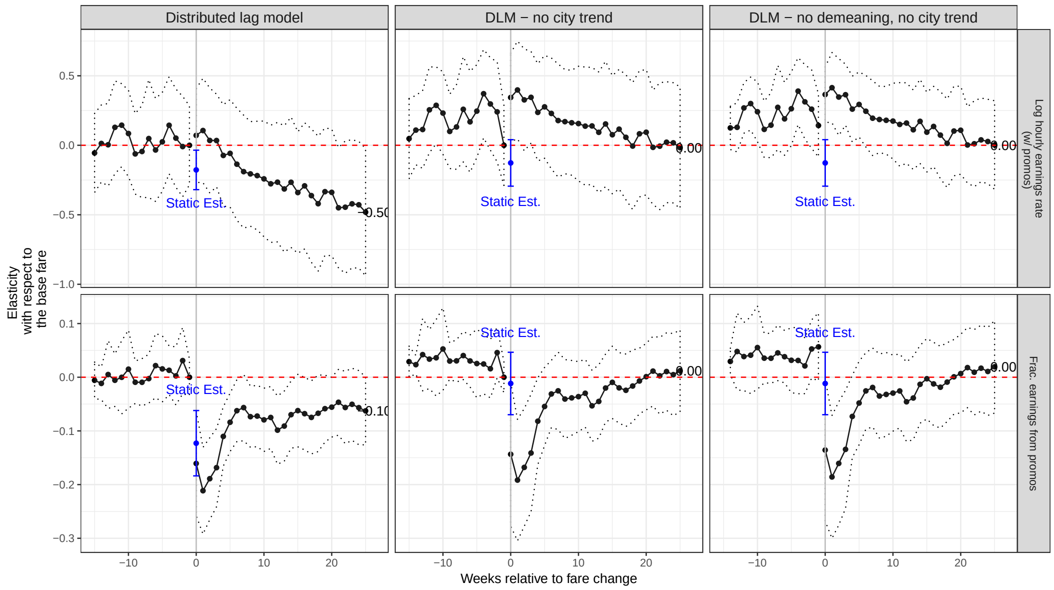

Notes: Alternative specifications.

Figure 23: Alternative bandwiths for Figure 11

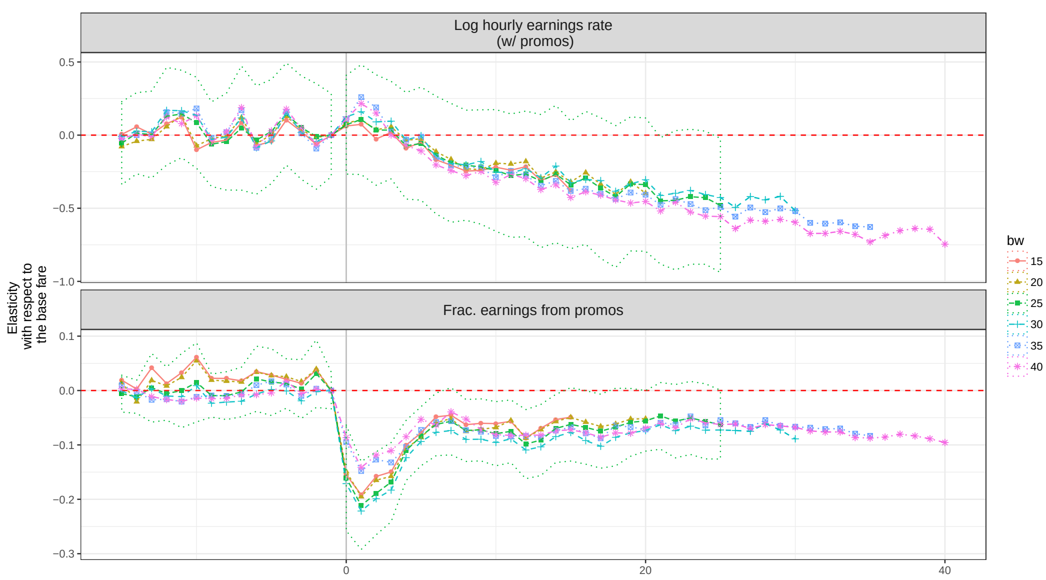

Notes: Alternative bandwidths.

Figure 24: Residual plots for log number of trips, by city and week

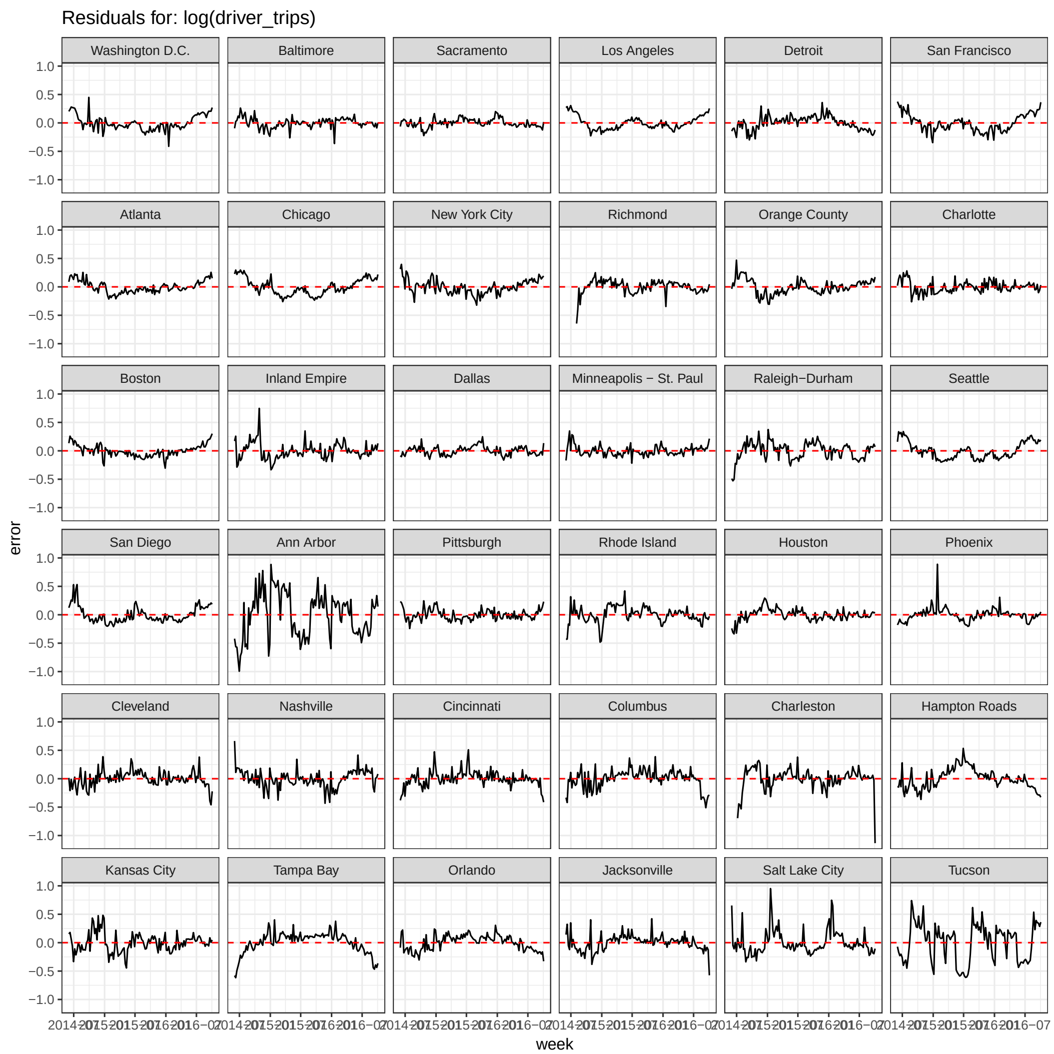

Notes: This figure plots the by-city week residuals when the outcome is the number of driver trips.

Figure 25: Residual plots for log number of hours-worked, by city and week

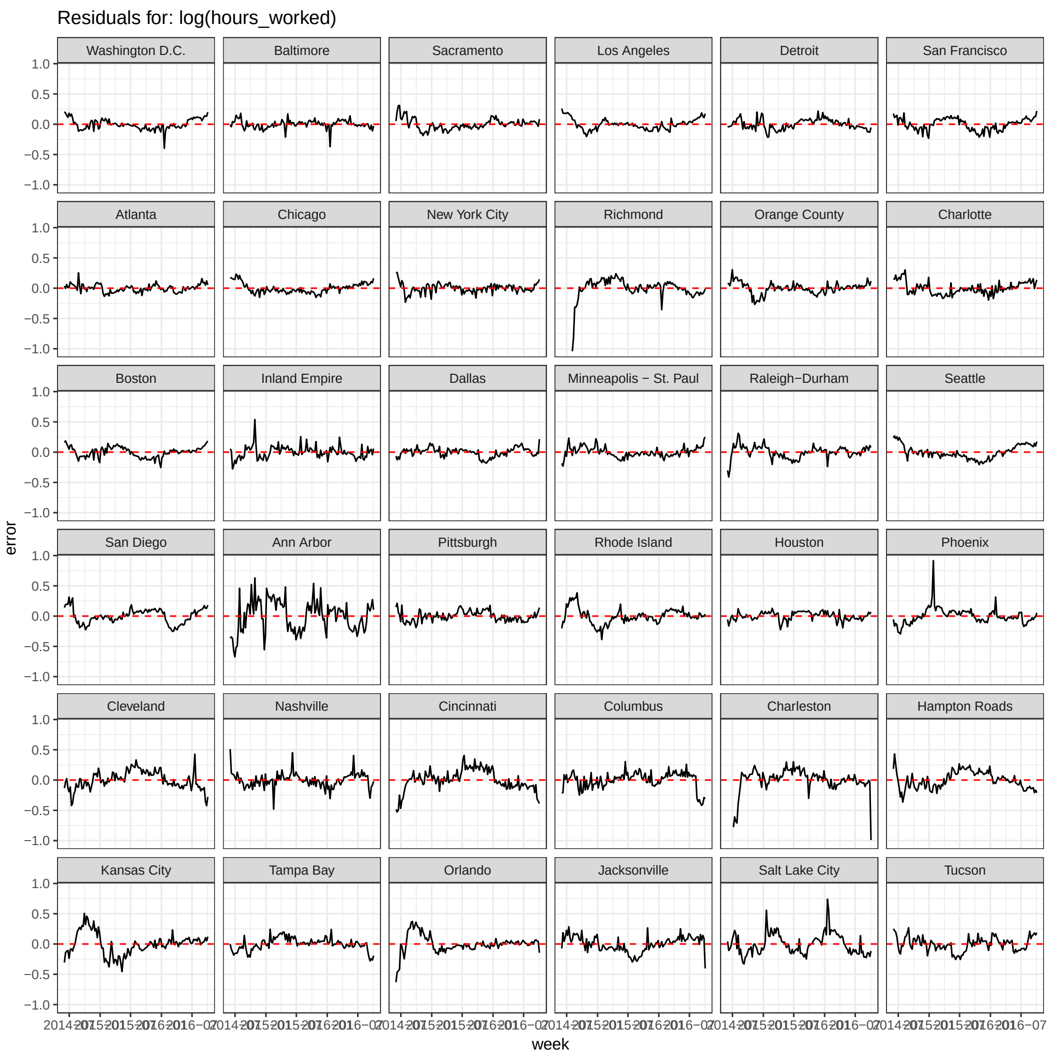

Notes: This figure plots the by-city week residuals when the outcome is the number of driver trips.