Labor Market Equilibration: Evidence from Uber*

Jonathan V. Hall
Uber Technologies

John J. Horton
NYU Stern

Daniel T. Knoepfle
Uber Technologies

September 19, 2017

Abstract

Using a detailed city-week panel of US ride-sharing markets created by Uber, we study the effects of sudden, Uber-initiated price shocks on driver labor market outcomes, focusing on the short-run dynamics as the markets adjust and the long-run equilibrium. We find that driver hourly earnings rates—essentially market equilibrium wages—move immediately in the same direction as price changes, but that these effects are short-lived. Prevailing wages return to their pre-fare-change level within about 8 weeks. This return is achieved primarily through permanent changes in driver "utilization," or the fraction of hours-worked that are spent transporting passengers. Our results are consistent with drivers supplying labor highly elastically to Uber, most likely because drivers face no quantity restrictions on how many hours to to supply and new drivers face minimal barriers to entry.

JEL J01, J24, J3

1 Introduction

In the ride-sharing markets created by Uber, the fare schedule faced by passengers is set by Uber, but the hourly earnings of drivers is not—a driver's

*Thanks to Andrey Fradkin, Steve Levitt, Keith Chen, Judith Chevalier, Mark Duggan, John List, Ed Glaeser, Austan Goolsbee, Robin Yerkes Horton, Peter Cohen, 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.

hourly earnings depend, in part, on both the per-trip fare—which they receive a fixed percentage of after Uber's commission—but also on the driver's "utilization," or the fraction of his or her working hours that are spent driving passengers, earning money. In essence, a driver has an hourly earnings rate that is, mechanically, their marginal revenue product, and so by changing the product market price (by changing the fare schedule), Uber changes the marginal product of workers, at least until the market adjusts. The nature of the subsequent adjustment to this change depends on Uber's relationship to the broader labor market.

If Uber faces a horizontal labor supply curve—as would a small hiring firm in a large labor market—the market should adjust until drivers are back to earning their pre-change *market-wide* marginal revenue product. Free entry and exit by drivers, combined with low switching costs to alternative platforms and relatively low barriers to entry all would tend to flatten the labor supply curve. In contrast, if Uber faces an upward sloping labor supply curve—say, because of its unique flexible offering (Chen et al., 2017), regulatory barriers to entry, non-zero driver fixed costs, income targeting behavior, and so on—the driver hourly earnings rate could be permanently altered following a fare change. Even if Uber faces an upward-sloping labor supply curve, is not clear whether the fare change would have the same sign as the change in hourly earnings—if demand were sufficiently elastic, and supply sufficiently inelastic, a fare increase could lower driver utilization enough to lead to a lower hourly earnings rate, and for the same reason, a fare decrease could raise the driver hourly earnings rate.

In this paper, we explore how Uber-initiated, city-specific changes in the fare schedule faced by riders—the product market price—affect the driver hourly earning rate—the factor market price.[1] We do this by making comparisons between markets where Uber changed the fare schedule to other markets in other cities where the fare schedule did not change at that same time. We

[1] These changes are distinct from the use of dynamic or "surge" pricing that Uber uses to protect the market from short-run imbalances in supply and demand; the base fare schedule is similar to the per-mile and per-distance rates used in most conventional taxi markets.

use both fare increases and decreases, as well as variation in the size of fare changes. For this analysis, we only use data from UberX, Uber's peer-to-peer offering. As a second empirical strategy, we also exploit the fact that in some cities, Uber offers a luxury service, UberBlack, which generally has had constant fares over the period covered by our data. We can make this within-city comparison because UberBlack is not a perfect substitute for UberX on the product side, and there are enough UberBlack drivers that do not participate at all in the UberX market that we can treat the two markets as meaningfully distinct.

We view the two approaches as complementary, in that each can partially address the limitations of the other approach. The between-city comparison of UberX markets has the benefit that each market is geographically isolated. However, there is the lingering concern that cities were selected for fare changes on the basis of their characteristics, which is not a concern in the UberBlack within-city comparison. The downside of the within-city UberBlack approach is that the two products are substitutes for each other. We assess the credibility of our difference-in-differences research design, which we do in part by exploring whether cities–and services within cities—were selected for fare changes on the basis of pre-existing trends. We find little evidence of this kind of selection, which may seem surprising, but as we will show, the precise timing and targeting of fare changes is likely "ignorable" (Rosenbaum and Rubin, 1983).

Our findings from the two different approaches are reassuringly similar. When Uber raises the fare schedule in a city, driver hourly earnings jump up immediately, but this effect begins to decline shortly thereafter. After about 8 weeks, there is no detectable difference in average hourly earnings compared to the pre-fare change level. If anything, there is evidence that fare increases have lowered driver hourly earnings. As per trip prices are higher, how do drivers make the same amount per hour? Part of the explanation is that some of a fare increase gets "undone" by a reduction in how often surge pricing is in effect (Chen and Sheldon, 2015; Hall et al., 2016)—with higher fares, demand outstrips supply less frequently, meaning surge pricing occurs less often. However, the more important explanation for why hourly earnings do not change is that utilization falls—with higher fares, drivers spend a smaller fraction of their working hours on trips with paying customers. These changes in utilization are, by far, the more important of two margins: for a 10% fare increase, the average surge rate would fall by about 2%, but utilization would decrease by about 8%.

A limitation of our regression analysis is that we can say relatively little about the effects of fare changes on market quantities, such as total rides taken, hours-worked, drivers active and so on. The core econometric problem is that all of the markets we study are expanding rapidly, and some of these markets vary substantially in size. As such, outcomes that are not "scale-invariant" (such as rates like utilization and hourly earnings) tend to give imprecise point estimates in a regression framework. In some cases, there are large level differences in the pre-period, even though there is not necessarily evidence of a trend.

To address this empirical short-coming, we also conduct a between-city synthetic control analysis of the fare changes, constructing the synthetic control city from a donor pool of cities that are similar to the focal city. An advantage of the synthetic control approach is that we can remove from our sample cases in which the constructed synthetic control poorly approximates the focal city in the pre-period. This ability to throw away "bad" observations is a luxury of having many fare changes in our data. Using this approach, we find the same pattern of results from the panel approach, with fare increases leaving hourly earnings unchanged but permanently lowering both utilization and surge.

We also successfully use the synthetic control approach for "quantity" outcomes such as the number of trips taken. As expected, we find that fare increases reduce the quantity of trips taken. We find no evidence that the number of drivers active changes in response to a fare change. However, there is evidence that following a fare increase, total hours-worked declines. Fewer hours-worked, combined with lower driver utilization, is how the total number of supplied trips falls when the fare increases.

Although it might be tempting to use our analysis to estimate a demand curve for Uber rides, it is important to consider that following a price change, wait times could change, shifting the demand curve. Indeed, using our between-city panel approach, we find strong evidence that when the base fare increases, wait times fall, despite the goal of surge pricing to keep these wait times more or less constant. The sensitivity of wait times to the base fare is consistent with our finding that higher fares lower utilization—with lower utilization, the nearest car available for dispatch is likely to be closer to the requesting would-be passenger. Although these changes in wait times are potentially important for explaining how the market adjusts, they have little direct import for driver hourly earnings, which are only affected through the surge and utilization "channels."

There are several implications of our findings. Uber clearly has some control over the equilibrium in the product market, in the sense that by changing the base fare, they change the quantity of rides taken (without bringing the quantity to zero or infinity). This is perhaps unsurprising, given Uber offers a differentiated product. In contrast, Uber appears to have little control on the supply side of market, with driver hourly earnings returning to their pre-change level quickly. In short, drivers as a whole appear to be highly elastic in their labor supply with respect to Uber. While drivers clearly value the flexibility offered by the option of driving with Uber (Chen et al., 2017), it is not determinative for the marginal driver and/or the driver considering his or her marginal hour.

Our results tell a fairly simple economic story about how the supply-side of ride-sharing markets work. When driving with Uber temporarily becomes a better deal, drivers work more and push down hourly earnings through lowered utilization and somewhat less surge. The process runs in reverse when driving with Uber becomes a relatively worse deal, with drivers working fewer hours until there is higher utilization on the remaining hours-worked and somewhat more frequent surge. Interestingly, our finding that fare increases have reduced the number of hours-worked is consistent with the notion that, on average, fare *cuts* have made driving for Uber a better deal. This is perhaps less surprising that it might seem. We present a short theoretical argument in the discussion section of the paper that implies technological advances that cause a secular increase in utilization will cause Uber's un-adjusted fare to be too high, with the old fare being in the inelastic portion of the demand curve. When in this situation, Uber's profit-maximizing adjustment is a fare reduction, which will tend to raise driver hourly earnings (so long as the supply curve is not perfectly horizontal).

Although the Uber context is clearly special, a natural question is how many other tightly coupled product/labor markets clear through changes in utilization on the labor side rather than price. It is commonplace to hear workers with qualitatively similar degrees of flexibility to Uber drivers—small business owners, consultants, and freelancers—describe work as being "busy" or "slow"; rarely do they express the sentiment "exactly the same level of business as always, but at fluctuating prices." We are not the first paper to highlight this utilization-based market clearing: Hsieh and Moretti (2003) show that when housing prices increase, the earnings of real estate agents do not increase much if at all—agents simply sell fewer houses as more agents enter the industry. This is the Uber-equivalent of having fewer rides for some given number of hours-worked, i.e., a lower utilization. Because of the computer-mediated nature of the market (Varian, 2010), hours-worked and related measures are made essentially without error, which allows us to show directly how important this utilization margin is for market clearing.

The data granularity allows us to show how equilibration takes place over time. Even when the theoretical predictions in our setting are clear, there is little theoretical guidance on how quickly adjustments might take place.[2] With the rise of the so-called "gig economy" (Katz and Krueger, 2016), it seems like that there will be more markets structured in this way, with independent sellers and a mediating platform with some market design powers (Einav et al., 2016; Horton and Zeckhauser, 2016; Sundararajan, 2016).

A somewhat obvious implication of our findings is that ride-sharing mar-

[2]Some have argued that the speed of adjustment is critical for the monopsony versus competitive characterization of labor markets (Manning, 2003; Kuhn, 2004).

kets can exist in different equilibria, depending on the fare schedule faced by passengers. All else equal, the equilibria with a higher utilization are probably social preferable, given the negative externalities of driving (Parry et al., 2007; Edlin and Karaca-Mandic, 2006).[3] Furthermore, to the extent drivers have fixed costs of entry, the higher utilization equilibrium is preferable for the reasons presented in Mankiw and Whinston (1986).

The rest of the paper is organized as follows. Section 2 describes the empirical context. Section 3 presents a simple model and discusses some of the prior work on taxi markets. Section 4 presents the empirical results. Section 5 discuses the results and concludes with some thoughts on the implications of our results for Uber's pricing problem.

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 more than 60 countries. The core products of Uber are UberBlack and UberX (see Hall and Krueger (Forthcoming) for a discussion for the relative size of the two services). 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. Regardless of the product, riders 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 rider's credit card. Uber handles all billing, customer support, and marketing.

The price of an UberX and UberBlack trip depends on a number of parameters set by Uber. There is a per-minute time multiplier and per-mile distance multiplier, as well a fixed initial charge. To calculate the actual fare paid by the rider, the parameters are multiplied by the realized time and distance of a

[3]Though this claim must be tempered by the higher cost drivers presumably have at higher levels of utilization.

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 much. During "un-surged" periods, the multiplier is 1.0.[4] There is a minimum charge, and in some markets there are various service fees. Recently, Uber has begun using "up-front" pricing in which riders 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 this switch to up-front pricing was widely implemented.

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.

We explore how changes in the price index affect a variety of market outcomes. Our most important outcome is the average hourly earnings rate of drivers. This rate requires a precise measure of hours-worked. We define hours-worked as the total time a driver spent "online" with the Uber platform; this includes all time on-trip, en-route to pick up a rider, or simply available for dispatch. Merely having the app open without marking oneself available for dispatch does not contribute to hours-worked.

Our definition of hours-worked is imperfect. 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 the hourly earnings rate.

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

Drivers also report driving with multiple ride-sharing platforms simultaneously, turning one off immediately after being dispatched by the other service. This multi-homing strategy will also tend to inflate hours-worked and thus lower the implied earnings rate.

We calculate the hourly earnings rate as total weekly net driver revenue divided by the total hours-worked. This is method is equivalent to averaging all driver-specific estimates and weighting by hours-worked. For revenue, we omit reimbursements for known tolls and fees (such as airport fees), and deduct Uber's commission. 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. As first, we *do not* include any driving-related promotional payments in our estimate of the hourly earnings rate. However, we do explore the role of promotional payments after first presenting the effects of fare changes on the "organic" measure of the hourly earnings rate.

Drivers are eligible for promotional payments that typically depend on meeting various goals, such as hours-driven or rides-taken in some week. We allocate promotional 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.

2.1 Data description

We construct a panel of 43 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 UberX service, though only some have UberBlack. 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.[5] These cities include Boulder,

[5]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 of even the top 50 includes several markets that were not very mature in the start of the panel.

Denver, Indianapolis, Las Vegas, Philadelphia, 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 Portland, Charleston, New Orleans, and Richmond because these cities 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 earnings cannot be reconstructed with sufficient confidence.

2.2 Variation in the base fare index

Uber has changed the base fare for UberX in every city in the panel, with some cities experiencing multiple changes. Figure 1 highlights the weeks in which the UberX base fare index changed for the panel cities and reports the size of the change. A grey tile indicates no change occurred that week. Cities are listed in descending order of their average base fare over the period. A black dot next the city's name indicates that city had an UberBlack option available.

According to Uber, the decision to change the fare in a city was made after consultation with the 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. This creates an obvious selection concern, though we will argue that there are several justifications for treating fare changes as exogenous for our between-city comparison purposes, and that furthermore, our fine-grained time scale makes the parallel trends assumptions readily assessable.

First, as every city in the panel had fare changes, it is not the case that latent differences exist between the kinds of cities that have fare changes and those that do not. 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. Note that in both 2015 and 2016, many cities experienced large cuts in the base fare right after

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

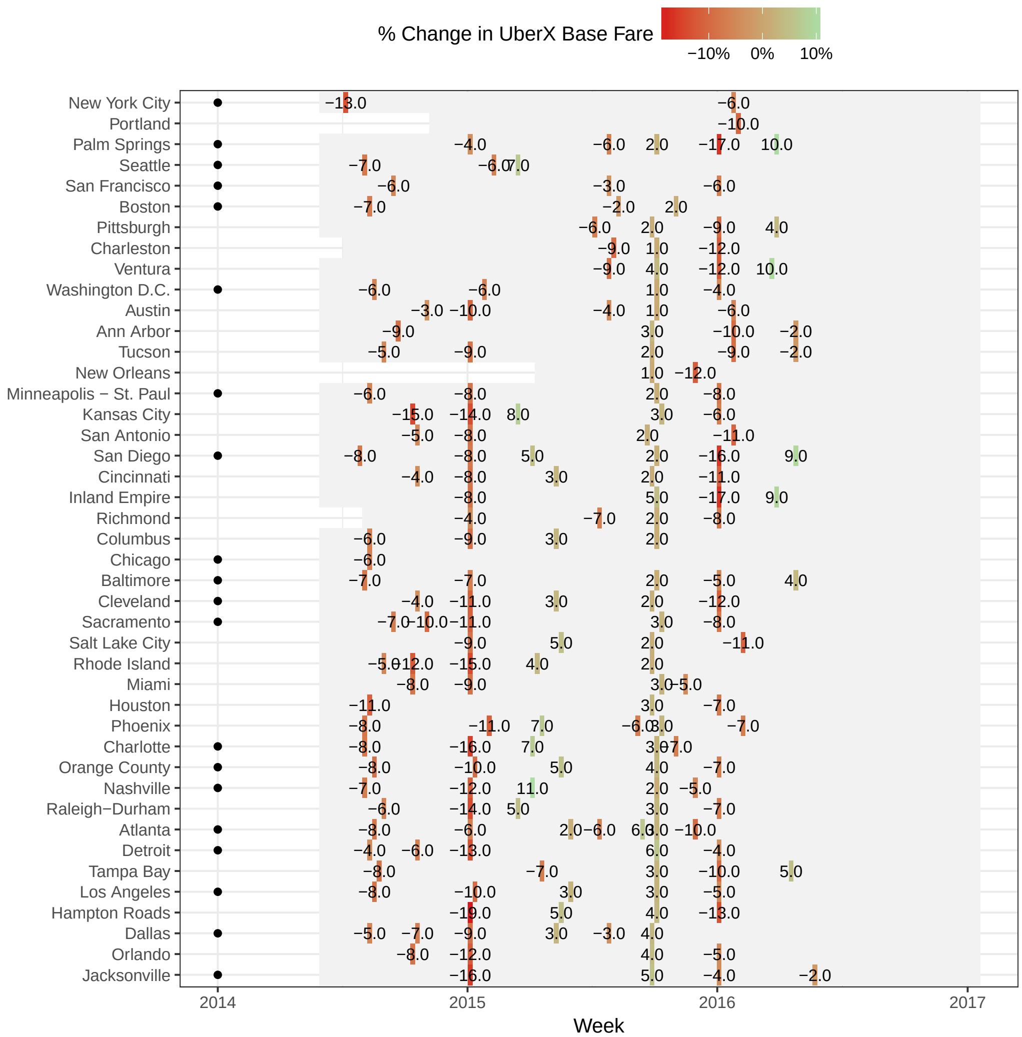

Notes: This figure indicates which cities in the panel had changes in the base trip price index, by week, and plots the size of that change, in percentage terms relative to the fare in the previous week. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, 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. Cities are listed in descending order of their average base fare over the entire panel. See Section 2.1 for a definition of the sample.

the start of the new year. 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, and the precise sequence in which changes occur is more likely to be ignorable.[6] Third, leaks by various media sources indicate that a 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, which we can do because of the highly granular nature of our data. In addition to checking for evidence of selection, our UberBlack within-city research design is not subject to the same concern that latent city-specific factors were driving selection, as is the case with the between-city design.

3 Conceptual framework

In this section, we present a simple model of the ride-sharing market in a city. The main purpose of this model is to illustrate the potential margins of adjustment in response to a change in the base fare. There are extant models of taxi markets, but 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. The features of ride-sharing markets make many of these search considerations less important, and so we develop our own simple model. However, we do briefly review this literature before presenting our model.

Arnott (1996) focuses on the potential for scale economies in the taxi industry, which in turn might justify subsidization. Castillo et al. (2017) build on this insight to discuss the role of dynamic pricing in avoiding highly inefficient equilibria. Cairns and Liston-Heyes (1996) present a model in which both sides of the market search for each other, showing that competitive equi-

[6] 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.

librium does not necessarily exist.[7] Frechette et al. (2015) present a calibrated equilibrium model of the NYC taxi market, with a focus on matching frictions and their role in determining the market equilibrium.

In the models of the taxi industry discussed above, demand depends on the wait times of riders, which adds substantial complexity. We can pursue a simpler modeling approach because dynamic pricing prevents the market from clearing on wait times for riders, at least in intent.[8] Surge pricing is designed to keep wait times below a certain threshold, making it reasonable to model the market with a single demand curve that depends only on price.

Our treatment of driver labor supply is also simple, ignoring "micro" or behavioral considerations, such as income targeting (Camerer et al., 1997) 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.[9]

The model we present should be thought of as describing a market over the course of weeks or months. Market clearing is not due to drivers deciding to extend shifts or riders forgoing a surged trip, but rather about drivers deciding how much to drive with Uber over the course of weeks (if at all) and riders deciding how many hours of transportation service to buy over the course of

[7]Both supply and demand are determined by the wait time (either for the cab on the passenger side, or for the next fare, on the supply side) and the fare and that the market could clear any number of different equilibria.

[8]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.

[9]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.

weeks (if any).

3.1 Model

Riders demand hours of transportation services. The price per hour of transportation services is p, which is the base fare, b, times the average surge multiplier, or $p = bm$. There is demand for hours of transportation, $D(p)$, with $D'(p) < 0$. The quality of transportation services is constant and unrelated to any other features of the market. Homogeneous drivers have an hourly earnings rate of w and collectively supply $S(w)$ hours of *driving*, with $S'(w) > 0$. Hours of driving are turned into hours of transportation at a rate x, where x is the fraction of hours-worked drivers spend on trips carrying passengers. Market clearing requires that

$$D(p) = xS(w).$$

The average hourly earning rate of drivers is simply the price times the utilization: $w = px$.

Figure 2 shows a market initially in equilibrium at some base fare p and Q hours of transportation services. Uber does not change p directly, but rather changes b, which in turn can change m in equilibrium. Suppose Uber lowers the fare by $\mathrm{d}b$, causing the fare faced by riders to change by $\mathrm{d}p = m\,\mathrm{d}b$. At this new base fare $p - \mathrm{d}p$, there are $\mathrm{d}pD'(p)$ more hours of transportation demanded. Before x or m adjusts, hourly earnings fall by $\mathrm{d}px$, reducing hours driven by $S'(w)\,\mathrm{d}w$ and hence hours of transportation by $xS'(w)\,\mathrm{d}w$. This gap between what is demanded and what is supplied is $\mathrm{d}p\left(D'(p) - xS'(w)\right)$, which is indicated in Figure 2 with a double-headed arrow.

Our first question of interest is how the market adjusts when b is exogenously changed—specifically the effect the change has on the driver hourly earnings rate. We assume, for now, that the surge multiplier is fixed and that adjustment can only occur through changes in utilization. Proposition 1 gives an answer to this adjustment question in terms of the labor supply and transportation demand elasticities.

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

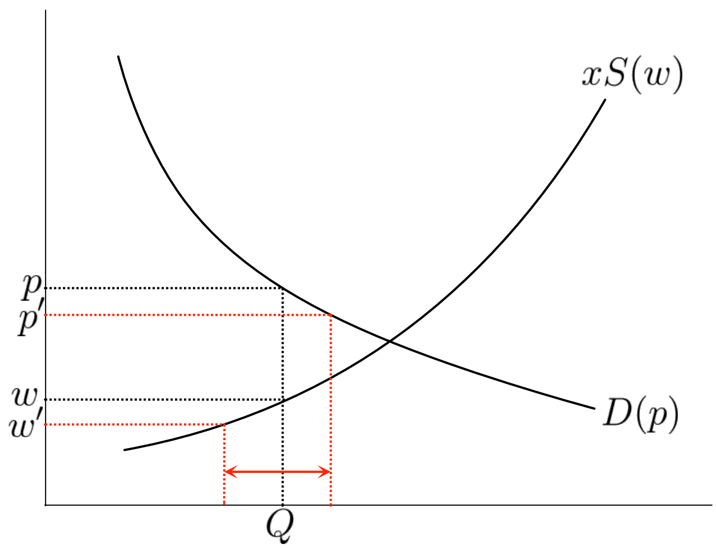

Notes: The $D(p)$ curve is the market demand for hours of transportation service, when the price is p. The $S(w)$ curve is the hours of driving supplied and $xS(w)$ is the hours of transportation services provided, where w is the driver hourly earnings rate and x is the driver market utilization. Q is the equilibrium quantity of hours of transportation service before a fare change. This figure illustrates the effects of a reduction in the base fare, from p to p' before any equilibrium adjustment. The double-headed arrow indicates the gap between hours-supplied and hours-demanded following the fare change.

Proposition 1. *If the market only clears on driver utilization, the elasticity of the driver's hourly earnings, w, with respect to the base fare, p, is*

$$\epsilon_p^w = \frac{1+\epsilon_p^D}{1+\epsilon_w^S}.$$

Proof. The elasticity of the driver's hourly earnings with respect to p is $\epsilon_p^w = 1 + \epsilon_p^x$. The elasticity of utilization with respect to p can be derived from the market clearing requirement. Differentiating $D(p) \equiv xS(w)$ by p, we have $\epsilon_p^D = \epsilon_p^x + \epsilon_w^S \left(1 + \epsilon_p^x\right)$, and so

$$\epsilon_p^x = \frac{\epsilon_p^D - \epsilon_w^S}{1 + \epsilon_w^S}.$$

□

Interestingly, increasing utilization works both "inside and outside" the supply curve, by making each hour supplied more productive, but also by causing movement along the curve as drivers respond to higher hourly earnings. It immediately follows from Proposition 1 that when drivers are infinitely elastic, hourly earnings cannot fall, as $\epsilon_w^S = \infty$ implies that $\epsilon_p^w = 0$. If drivers have finite labor supply elasticities, than the direction of the effect of a fare change on the hourly earnings rate depends on the demand elasticity. Conditional upon a finite labor supply elasticity, if $|\epsilon_p^D| > 1$, an increase in the fare will lower driver hourly earnings, and vice versa when $|\epsilon_p^D| < 1$. The intuition is that with elastic demand, changes in the fare cause large shifts in the quantity of rides demanded, which in turn affects driver utilization enough to more than offset the direct effect of the fare change.

Proposition 1 only considers adjustments in x as the means to clear the market. However, as Figure 2 illustrates, there are other ways for the market to clear that are not covered by Proposition 1. First, the average surge multiplier can increase, essentially "undoing" the fare change. Second, the supply curve could be shifted out, say due to more aggressive driver recruitment. Third, the demand curve could be shifted in, say because wait-times increase. We will investigate all three of these hypotheses empirically.

Note that our original definition of driver hourly earnings, $w = px$, assumes that there is only a fixed cost to driving an hour, and that the value of x does not enter separately into driver decision-making. However, drivers would face higher per-hour costs from a higher x, such as through higher gas expenditure (both because they have a passenger in the car and because they cannot park and wait for dispatch), greater dis-utility of effort from driving a passenger, increased vehicle wear-and-tear, increased cleaning expenses and so on. If we assume that costs are linear in x, i.e., $w = (p - c)x$, with $c > 0$, Proposition 1 would still hold, but to the extent there is individual driver heterogeneity in c, those drivers that that find a higher utilization equilibrium less costly would be inframarginal.

4 Results

In this section, we present both the between-city (UberX to UberX) and within-city (UberX to UberBlack) estimates of the effects of base fare index changes on market outcomes using a panel model. Our outcomes of interest are driver hourly earnings—w in the model—and the endogenous components, utilization and the surge multiplier, which are x and m in the model, respectively. Recall the $w = bmx$, where b is the base price.[10] After reporting the panel results, we present the synthetic control estimates for both the same outcomes used in the panel analysis—hourly earnings, utilization, and surge—but also for market quantities, namely the number of trips taken, the total hours-worked and the number of active drivers. We then return to the between-city panel to explore the effects of fare changes on passenger wait times.

4.1 Panel-wide averages over time

Before presenting the panel regression results, we first simply plot the weekly averages for various measures, pooled over all cities in the panel. Figure 3 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. First, 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. During that same period, hourly earnings have shown no similar trend. Driving utilization has generally increased substantially.

[10] Our constructed price index is not exactly b, which in the model is the cost of an hour of transportation services. 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.

Actual wait times were high early in the period, but then fell substantially by early 2015.

This figure previews some of the main results. Note that the two large drops in the fare index occur at the start of 2015 and 2016. Immediately after, average surge increases, as does utilization. Hourly earnings increase following the fare decreases, but exhibit much of an overall trend. There seems to be some evidence, at least in 2015, that after the fare decrease, wait times increased, but in later time periods, there is much less change in wait times.

4.2 Between-city approach

To begin, we set aside the possibility that market adjustment takes time and simply report "long-run" estimate of the effects of fare changes, assuming that the full adjustment occurs immediately.[11] Our specification is

$$\log y_{it} = \beta \log b_{it} + \gamma_i + \delta_t + \epsilon_{it}, \tag{1}$$

where y_{it} is some market-level outcome of interest in city i during week t, γ_i is a city-specific fixed effect, δ_t is a week-specific fixed effects, and b_{it} is the base trip price index. Table 1 reports estimates of Equation 1 where the outcome variables are the log hourly earnings, 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.[12]

From Column (1), we can see that an increase in the base fare decreases hourly earnings, though the effect is close to zero, with a confidence interval that comfortably includes zero. The point estimate implies a 10% increase in the base fare would lower driver hourly earnings by a little less than 1%. 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

[11] This specification is sometimes called the "static" specification—see Borusyak and Jaravel (2016) and references therein.

[12] 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.

Figure 3: 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 2.1 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 riders 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 is average elapsed time from when a passenger requested a trip to when they were picked up.

fare reducing utilization by about 8%. Finally, in Column (3) we can see that the "rest" of a 10% increase in fares is counteracted by about a 2% decrease in the average surge multiplier.

The panel model of Equation 1 cannot account for the possibility that the markets adjust over time following a fare change. To account for ad-

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

	Dependent variable:		
	log hourly earnings	log utilization	log surge
	(1)	(2)	(3)
Log base fare index	-0.092	-0.827^{***}	-0.218^{***}
	(0.083)	(0.088)	(0.025)
City FE	Y	Y	Y
Week FE	Y	Y	Y
Observations	5,852	5,852	5,852
R^2	0.663	0.732	0.432
Adjusted R^2	0.652	0.724	0.414

Notes: This table reports OLS regressions of several city-week outcomes on the log base fare index. The estimating equation is Equation 1. The base fare index is the price to riders 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 2.1 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 $p \leq .001 : ***$.

justment over time, we need a more richly specified regression model. Let $\textsc{FareChange}_{it}^{\tau}$ be an indicator for whether at time t, the base fare in city i is different than it was τ weeks earlier or later. To "cherry pick" the week when a fare change occurred/will occur exactly τ weeks away from the current week, we define an indicator

$$z_{it}^{\tau} = \textsc{FareChange}_{it}^{\tau} \prod_{s=0}^{\tau} (1 - \textsc{FareChange}_{it}^{s}).$$

We can then estimate

$$\log y_{it} = \beta_{\text{LR}} \log b_{it} + \sum_{\tau=A}^{B} \beta_{\tau} z_{it}^{\tau} \Delta \log b_{it}^{\tau} + \gamma_i + \delta_t + \epsilon_{it}, \tag{2}$$

where b_{it} is the fare index in city i at time t, $\Delta \log b_{it}^{\tau}$ is the difference in (logs) between the current fare index and the fare index τ weeks prior or later (depending on the sign of τ), δ_t is a week-specific fixed effect and γ_i is a city-specific effect. The number of pre-period week indicators is A and the number of post-period weeks indicators is B. The estimated effect for a fare change that occurred τ weeks ago is $\hat{\beta}_{\text{LR}} + \hat{\beta}_{\tau}$ if $0 < \tau \leq B$, otherwise it is just $\hat{\beta}_{LR}$. The coefficients on the pre-period indicators can be used to assess whether the cities selected for fare changes were systematically different, or on different trajectories, with respect to the outcome.

The implied weekly effects from Equation 2 are plotted in Figure 4 for each of the same outcomes as used in Table 1—the outcomes are the log hourly earnings, log utilization, and log surge, using $A = B = 15$. For each outcome, the long-run effect from Table 1 is plotted as a dashed horizontal line in the post-period.

From the left-most panel of Figure 4, we can see that following a fare increase, driver hourly earnings increase immediately, though the pass-through is considerably less than 1.0; the point estimate is only about 0.3. In the weeks that follow, this increase in hourly earnings declines, with the point estimate turning negative by week 8, at which point it is close to the long-run estimate from Table 1. Turning to the pre-period, there is no obvious trend in the collection of coefficients. However, in most periods, the point estimate is higher, though for nearly all point estimates, the 95% CI comfortably includes zero.

Figure 4: Between-city estimates of the market adjustments to a change in the base fare index

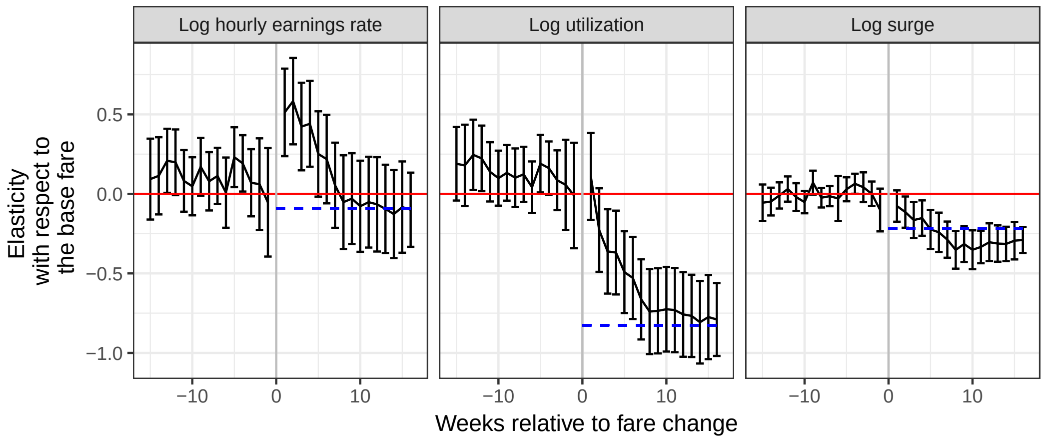

Notes: This figure plots the effects of changes in the UberX base fare on several market outcomes. These effects are from an OLS estimation of Equation 2. The sample is a panel of US cities—see Section 2.1 for a description. The x-axis are weeks relative to a fare change. The independent variable is the base fare index. The base fare index is the price to riders 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 horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to Equation 1. 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 next panel to the right, the outcome is the log utilization. Following a fare increase, driver utilization gradually falls. By week 8, the effect is close to the long-run effect. In the pre-period, there is no evidence of a trend, though in all cases, the utilization point estimate is somewhat higher, which would be consistent with Uber targeting cities for fare increases that had higher utilization, or equivalently given our empirical specification, targeted cities for fare cuts when utilization was low.

In the next panel to the right the outcome is the log average surge. The average multiplier declines following a fare increase. However, the long-run estimate is still a bit closer to zero than the 15 week point estimate, suggesting the market has not fully adjusted by week 15. There is no obvious trend in the pre-period and the point estimates are all close to zero.

4.3 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—if drivers drive some number of hours, they are guaranteed to make at least some floor. Given their structure, we might expect some of them to act as automatic stabilizers, counter-acting the immediate effects of whatever the change.

To explore the role of these promotional payments, in Figure 5, we plot the effects of a fare increase on the hourly earnings rate with and without promotional payments included. The panels are, in order: (1) the driver's hourly earnings (not including promotions), which is the same outcome as in Figure 4, (2) the driver hourly earnings including promotions, and (3) the fraction of hourly earnings that are due to promotional payments.

Comparing the two leftmost panels, we see that promotional payments likely "soften" the effects of are changes. For example, the pass-through immediately after the fare increase is about 0.5, whereas when promotional payments are included, this pass through is closer to 0.3. In the rightmost panel, we can see that the the fraction of earnings from promotional payments drops immediately after a fare increase. This is why the hourly earnings with promotional payment "peak" is below the measure that is purely from earnings due to fares. Despite this difference, we can see that the longer-run effect of promotional payments is close to zero. It seems that the promotional payments "channel" likely prevents extreme short-term hourly earnings fluctuations, but has negligible long-run effect, with the long-run being only about a month.

Figure 5: Effects of fare changes on hourly earnings with and without promotional payments included

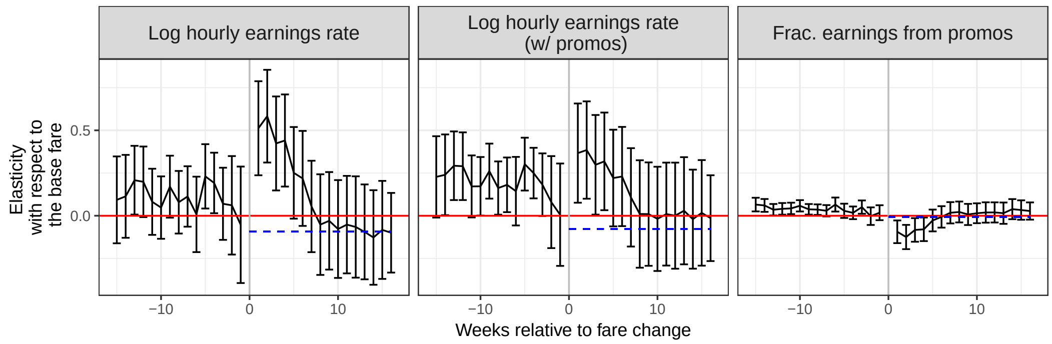

Notes: This figure plots the effects of changes in the UberX base fare on three outcomes (from left to right): (1) hourly earnings not including promotional payments, (2) total hourly earnings, which includes promotional payments, (3) the fraction of hourly earnings coming from promotional payments. These effects are from an OLS estimation of Equation 2. The sample is a panel of US cities—see Section 2.1 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 riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. As all outcomes are in logs, the point estimates can be interpreted as elasticities, which are estimated using both fare increases and decreases. The horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to Equation 1. 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.

4.4 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 for Uber trips. One demand curve shifter could be passenger wait time. In Figure 6, we report panel estimates of a fair change on the log median expected wait times in the left panel, and the log median actual wait times in the right panel. The expected wait time is that Uber reports to would-be riders.

The figure makes it clear that with a fare increase, wait times declined. The effects are substantial, with a 10% increase in fares reducing wait times by about 5%. The pricing effect on wait times is understandable—with less demand and/or lower driver utilization, for a given rider requesting a ride,

Figure 6: Effects of fare changes on actual and estimated log median arrival times

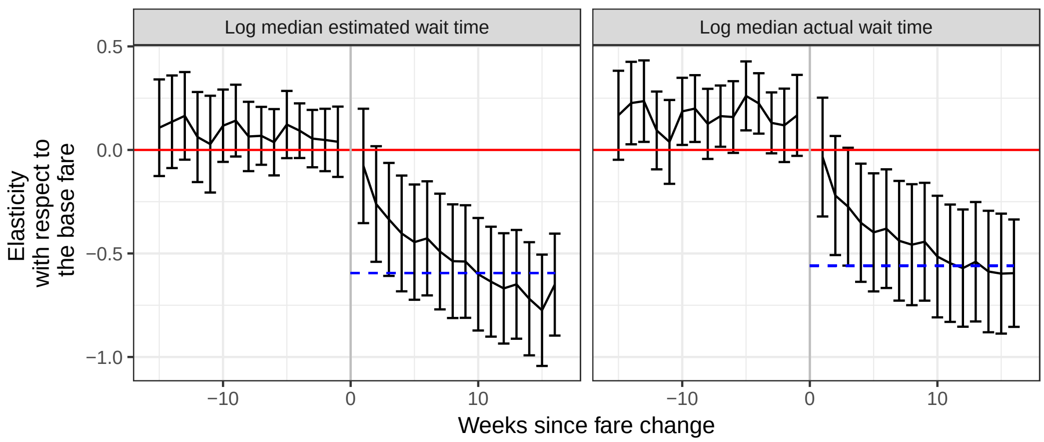

Notes: The outcomes in this panel are the log wait times (in seconds), both estimated by Uber (in the left panel) and the realized wait times (in the right panel).

the nearest empty car is likely to be closer (so long as the number of drivers does not decrease). The synthetic control results suggest that there was no detectable change in the number of drivers in response to fare changes, so it seems likely that the change in wait times is explained by changes in utilization. These wait time results have little import for driver earnings, though they do suggest that shifting the demand curve is partially a margin on which the market adjusts, and that surge pricing does not, as implemented during the period covered by the experiment, hold product attributes fixed. As such, at least some of the increase in demand from a fare increase was rationed by increased wait times (Lindsay and Feigenbaum, 1984).

4.5 Within-city approach using UberBlack

We now consider the effects of a change in the base fare index, but use UberBlack within the same city as the comparison. The sample for this analysis is restricted to cities in the panel with an UberBlack service, which are indicated in Figure 1 with a black dot next to the city name. The UberBlack panel was constructed in the same way as the UberX panel, with the exception that we only consider drivers that drove for UberBlack exclusively and were not "cross-dispatched" to UberX trips. Figure 7 plots the base fare indices for both UberX and UberBlack for the UberBlack panel, with the UberBlack price index indicated by a solid line and UberX price index by a dashed line. In all cities, the UberBlack base fare index has been constant since late 2014. During that same period, as we saw in Figure 1, there has been substantial variation in UberX fares.

For each city-week, we compute $\Delta \log b_{it} = \log b_{it}^{\text{UBERX}} - \log b_{it}^{\text{UBERBLACK}}$ and the analogous differences in the various log outcome measures, hourly earnings, utilization and surge. We then estimate the regression

$$\Delta \log y_{it} = \beta \Delta \log b_{it} + \gamma_i + \epsilon_{it}, \tag{3}$$

where γ_i is a city-specific fixed effect. This specification is analogous to Equation 1 which we used in the between-city analysis, with the exception that a week fixed effect is not included, as this would absorb the treatment effect. An attractive feature of the within-city design is that any city-specific time effect is removed by differencing, improving precision, whereas with the between-city approach, only a panel-wide time effect can be included.

Table 2 reports estimates of Equation 3. In contrast to the between-city analysis, the outcomes are now the difference, in logs, between the UberX market outcome for that week and the UberBlack outcome for that week. However, the coefficients on the independent variable, which is the difference in the log base price index, has approximately the same interpretation as before—it is an elasticity of that UberX outcome with respect to the differences in the UberX base fare index. However, we also use the change in UberBlack fares (see Figure 7) as identifying variation, and so a more precise definition of the effect is a percentage change in the difference between the two services—though most of the variation comes solely from changes in the UberX fare, giving the point estimates a similar interpretation to the between-city estimates.[13]

[13]If we exclude pre-2015 data from the panel and thus only rely on UberX variation, the

Figure 7: UberX and UberBlack base fare indices for the UberBlack panel, by week

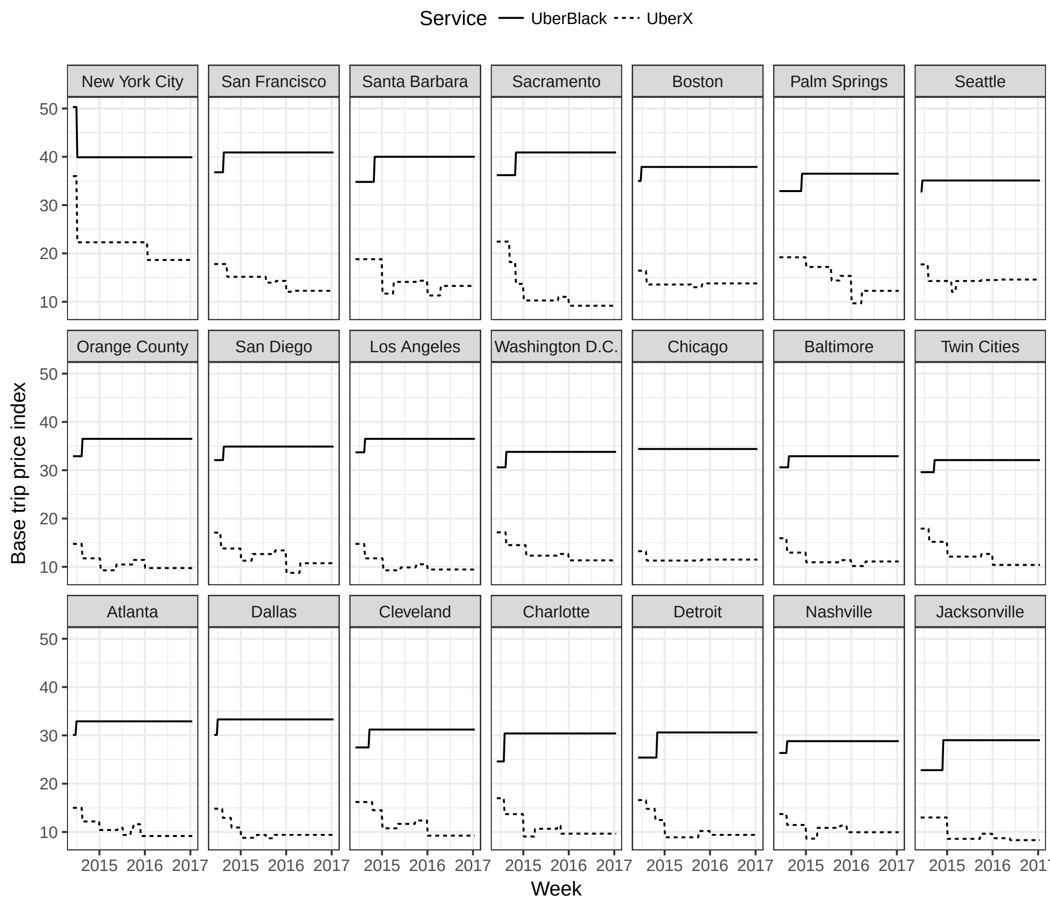

Notes: This figure plots the weekly base fare indices for UberX and UberBlack services for US cities with both services. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week.

Starting with hourly earnings, Column (1) of Table 2, the effect of a fare increase is to lower hourly earnings, with a 10% increase leading to a 2% decrease, which is close to the between-city estimate. Unlike in the between-city comparison from Table 1, this reduction is conventionally significant. As before, we can decompose the effect. In Column (2), we can see that a 10%

pattern of results is the same, though estimates are less precise, as a non-trivial fraction of UberX fare variation occurred in this pre-period.

base fare increase lowers utilization by about 11%, which is larger than the between-city estimate. In Column (3), a 10% fare increase lowers average surge by about 1%, which is slightly below the between-city estimate.

Table 2: Effects of fare changes on market outcomes (in logs) using UberX versus UberBlack fare variation

	Dependent variable:		
	log hourly earnings	log utilization	log surge
	(1)	(2)	(3)
Difference in log base fare indices	-0.207^{***}	-1.137^{***}	-0.111^{***}
	(0.069)	(0.078)	(0.015)
City FE	Y	Y	Y
Observations	2,898	2,898	2,898
R^2	0.334	0.516	0.174
Adjusted R^2	0.330	0.513	0.168

Notes: This table reports OLS regressions of the within-city difference in UberX and UberBlack by week outcomes on the difference in the log base fare index for that week. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. The estimating equation is Equation 3. The sample for each regression is the same, and is a city-week panel for a collection of US cities with both UberX and UberBlack. See Section 2.1 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. 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 $p \leq .001 : ***$.

To account for market adjustment that takes time, we use the same pre/post-period indicator approach based on Equation 2. Figure 8 plots the corresponding point estimates for the within-city analysis.

Starting in the leftmost panel, the outcome is the log hourly earnings. Hourly earnings increase almost immediately, with almost full pass through. However, this increase begins to fall. At the end of post-period, the point estimate is not quite as low as the long-run estimate (indicated with the dashed blue line) and is almost exactly at zero. There is no evidence of a pre-period trend, but perhaps some evidence of slightly higher levels of earnings in cities

Figure 8: Within-city estimates of the effects of fare changes on market outcomes, over time

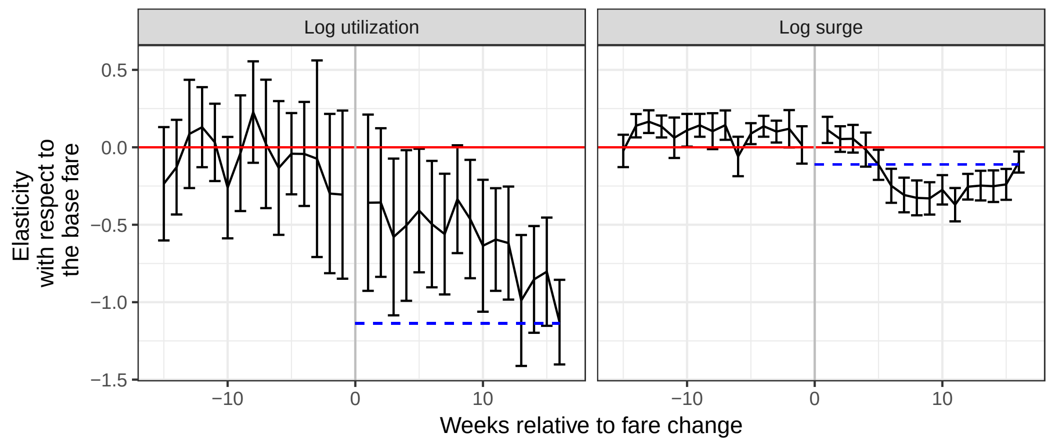

Notes: This figure plots the effects of changes in the difference in base fare between UberX and UberBlack on the difference in market outcomes for UberX and UberBlack, using several market outcomes. The estimating equation is based on Equation 2, but using UberBlack within the same city as the comparison. The x-axis are weeks relative to a fare change. The independent variable is the difference in the base fare index. The base fare index is the price to riders 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 horizontal dashed blue line indicates the effect corresponding to Equation 3, which does not include any dynamic adjustment. 95% CIs are shown for each point estimate and standard errors are clustered at the city level.

targeted for a fare increase.

In the middle panel of Figure 8 the outcome is the log utilization. Utilization declines following a fare increase. By week 15, the point estimate is quite close to the long run estimate, which is close to -1. There is no evidence of a pre-period trend or a difference in the per-period levels.

In the rightmost panel, the outcome is the log surge multiplier. We can see that surge declines following a fare increase, but that the post-period shows some evidence of a u-shaped pattern, with average surge eventually increasing, leaving a long-run estimate close to zero. There is no strong evidence of a per-

period trend, though there is perhaps some evidence of higher levels.

4.6 Synthetic control approach

Because of the growth of Uber, the panel regression approach works poorly for some outcomes of interest, such as market quantities. The synthetic control approach allows us to evaluate how well a counter-factual control approximates a city experiencing a fare cut. To perform a synthetic control analysis, we first found, for each city, broadly similar cities. We did this by ranking all cities relative to the focal city by taking (a) the average difference in log driver hourly earnings and (b) their correlation in this measure over time, normalizing both measures and then weighting the two measures equally. We selected the top 20 cities for each focal city. We did this rather than include all cities in the panel as potential donors, as we found including all cities led to over-fitting.

Next, for each fare change, we took the restricted list of donor cities and excluded cities that had fare changes within the full 15 week "window," forward and backward. If the resulting donor pool is empty, we exclude that fare change from the sample. Using the selected donor pool, we then ran the synthetic control algorithm described in Abadie et al. (2011), using the driver hourly earnings as the primary outcome and then the other components (surge and utilization) as other covariates, all in logs.

Using the synthetic control, we then compute the time course of the differences between the actual focal city and the synthetic control. For each outcome, we use the same synthetic control. Treatment effects are scaled by the percentage change in the base fare to make all point estimates comparable to the panel estimates, i.e.,

$$\frac{\left(y_{it}^{T} - y_{it}^{C}\right)}{\Delta b_i / b_i} \tag{4}$$

in order to create elasticities comparable to our panel results (this is similar to the approach used by Dube and Zipperer (2015) in the context of combining synthetic control estimates from minimum wage case studies). For each outcome and each fare change, we have a synthetic control estimate. We then combine these estimates—we will discuss aggregation when we present the results.

4.7 Synthetic control estimates of the effects on fare changes on driver outcomes

To measure the goodness of a particular synthetic control estimate, we measure the difference between the treated and control for a given outcome, and then find the maximum over the entire pre-period. We then restrict our attention to those synthetic control estimates were the maximum absolute difference (MAD) is below a threshold. We then just average the per-period estimates that remain.

Figure 9 presents the synthetic control estimates by averaging all the weekly Equation 4 estimates for the hourly earnings, utilization and surge outcomes, using different MAD thresholds—namely 1, 2, 4 and "Inf" (i.e., the whole sample). The figure shows how important it is to select only synthetic control estimates that perform reasonably well in the pre-period. For example, in the hourly earnings and utilization panels, the "Inf" estimate (i.e., averaging all synthetic control estimates) performs very poorly in the pre-period and then gives implausible post-period estimates. In contrast, trimming out bad matches gives per-period weekly averages that are quite close to zero (as expected) and gives post-period results that are pattern-wise, similar to the panel regression results.

The downside of trimming, of course, is that it reduces the sample size. Each estimate is labeled with the fraction of all fare cuts that are included with that particular cutoff. For example, for hourly earnings, the 2 MAD cutoff means that only a slightly more than 16% of fare changes can be included. In contrast, that same cutoff for utilization allows for nearly a quarter of observations to be included and about 37% of observations for the surge outcome.

To select our preferred estimate, we select a MAD to remove the worse

Figure 9: Synthetic control estimates of the effects of fare changes on market quantities, by pre-period match criteria

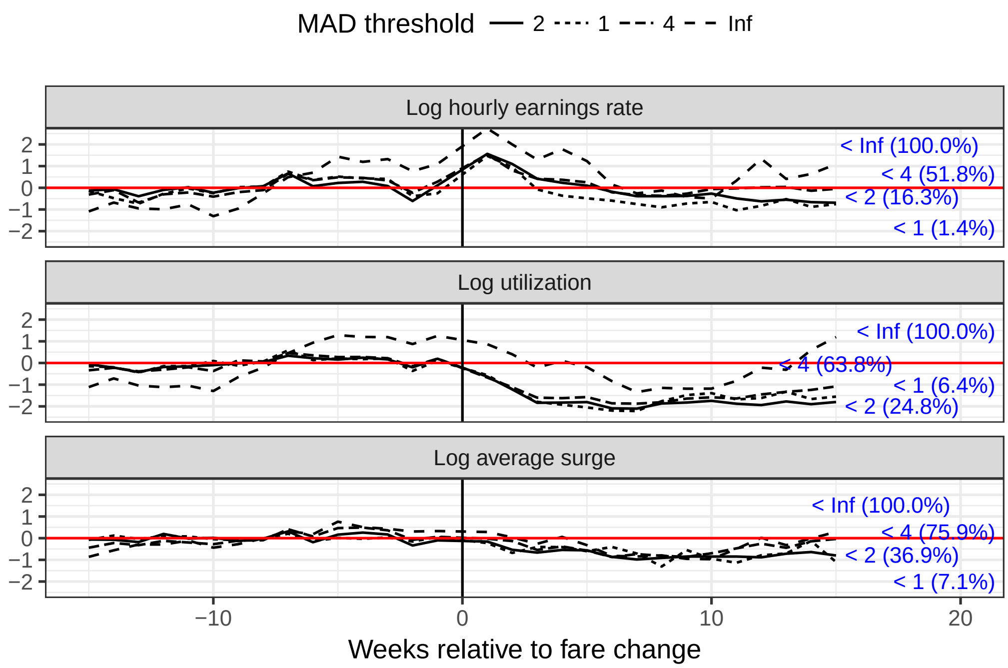

Notes: This figure plots the the average difference between cities experiencing a fare change and a synthetic control city comprising similar cities that contemporaneously did not have a fare change. All synthetic control effets are scaled by the change in the base fare index to make the results comparable to the panel estimates. The base fare index is the price to riders 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 different lines represent different selection criteria employed for whether a synthetic control estimate is included in the average: for each line, the cut-off value is shown (which is either Inf, 4, 2 or 1) and next it, the fraction of all fare changes this restriction will include.

estimates and then weight the remaining estimates by the inverse of the pre-period mean squared error. The first step assumes that some estimates are so bad as to only add noise; the second step is motivated by the notion that the most efficient estimator of a mean weights the observations by their inverse variance.

Figure 10 plots the the by-week estimated treatment effect for each of the fare changes for which a reasonable synthetic control could be constructed, using the hourly earnings, utilization, and surge as outcomes. Weighted standard errors are reported. The cut-off used is 2.5 and the per estimate weighting is the inverse of the MSE. The individual synthetic control estimates meeting the 2.5 MAD cut-off that make up the sample are plotted in light gray.

In the top panel of Figure 10, we can see a large initial increase in hourly earnings following a fare increase, followed by a decline, with the effect eventually turning negative. This is the same pattern found in both the within and between panel analyses. In the next panel down, the outcome is the driver's log utilization. Utilization strongly declines following an increase in the fare, again matching the pattern found in the panel. In the bottom panel, the outcome is the log average surge. We can see a clear decline following a fare increase, as in the panel analyses, and perhaps some evidence of a u-shaped pattern.

Although the pattern of results from the synthetic control approach are similar to the panel results, the light gray individual estimates in Figure 10 make it clear that there is substantial heterogeneity in the course of outcomes for different cities, with many paths far away from the average of the estimates. From the perspective of any individual driver in any particular city, the central tendency in the data might not be experienced.

4.8 Synthetic control estimates of the effects of fare changes on market quantities

We now turn to the effects of fare changes on market quantities, such as total driver hours-worked, the number of active drivers, and the number of trips. Figure 11 shows the average synthetic control effect estimate using various cut-offs for pre-period match. The importance of restricting the sample is even clearer for this set of outcomes–the unrestricted estimates for all three outcomes show an average far away from constant near-zero in the pre-period.

The three restrictions all given patterns of results that are qualitatively

Figure 10: Synthetic control estimates of the effects of fare changes on driver outcomes

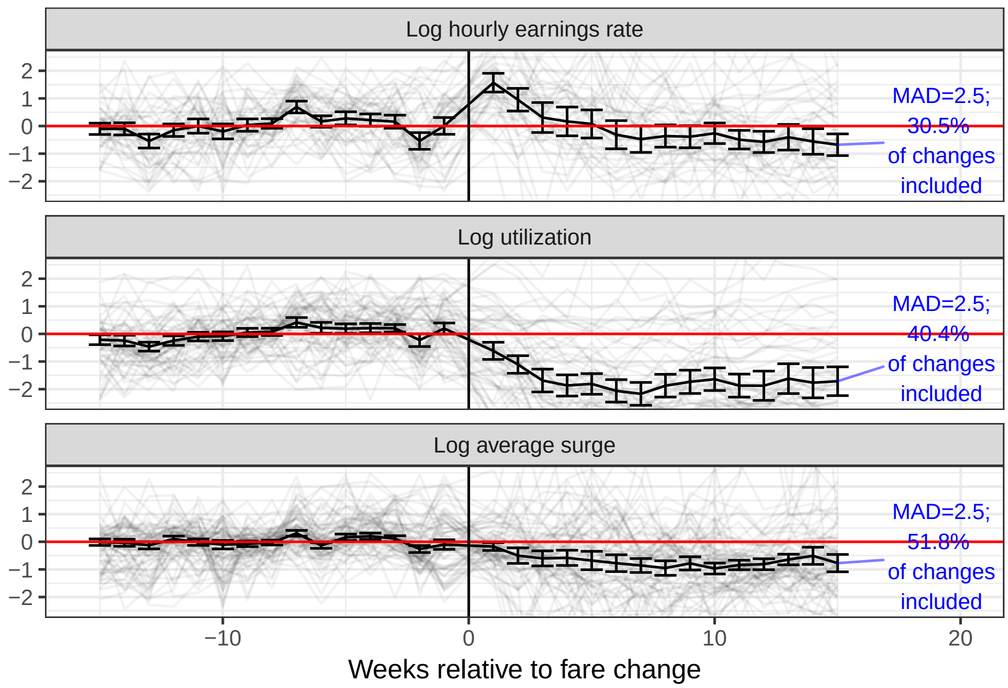

Notes: This figure plots the the difference between cities experiencing a fare change and a synthetic control city comprising similar cities that contemporaneously did not have a fare change. Only synthetic control estimates where the maximum absolute pre-period difference between the focal city and synthetic control was less than 2.5 are included. The sold dark line indicates the average value, weighted by the pre-period MSE, with error bars indicating a 95% CI. All synthetic control effets are scaled by the change in the base fare index to make the results comparable to the panel estimates. The base fare index is the price to riders 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.

similar, though there are large differences in the point estimates. For example, for the log trips outcome, the highly restricted sample (< 1) has a final period point estimate close to -2, but only includes 0.7% of the sample, which is actually just a single pair! Consistent with the increased variance that arises from using a single sample, this series shows a jump in trips at period 2 that is almost certainly a statistical artefact, as it disappears with the larger samples afforded by relaxing the restriction to 2. With 2, the sample jumps to about 12%, and with 4, it jumps to a bit more than 40% of the sample.

Figure 11: Synthetic control estimates of the effects of fare changes on market quantities, by pre-period match criteria

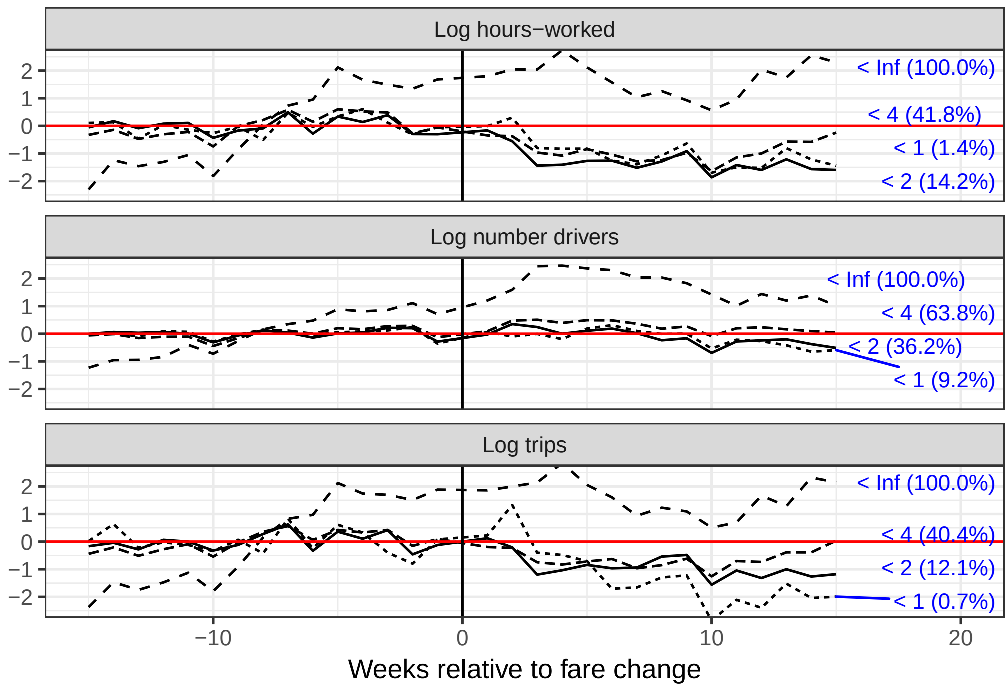

Notes: This figure plots the average difference between cities experiencing a fare change and a synthetic control city comprised of similar cities that contemporaneously did not have a fare change. The different lines represent different selection criteria employed for whether a synthetic control estimate is included in the average: for each line, the cut-off value is shown (which is either Inf, 4, 2 or 1) and next it, the fraction of all fare changes this restriction will include. Hours-worked is the total number of hours worked (i.e., had the Uber app on and were available for dispatch or were driving passengers, or enroute to pick up passengers) by drivers in that that city, in that week. The number of drivers is the total number of drivers in that city that worked at least some number of hours. The number of trips is the number of trips completed.

Figure 12 plots synthetic control estimates of the effects of a fare change on the log number of hours-worked, the log number of active drivers, and the log number of trips taken. We maintain our use of 2.5 as our pre-period selection criterion.

Figure 12: Synthetic control estimates of the effects of fare changes on market quantities

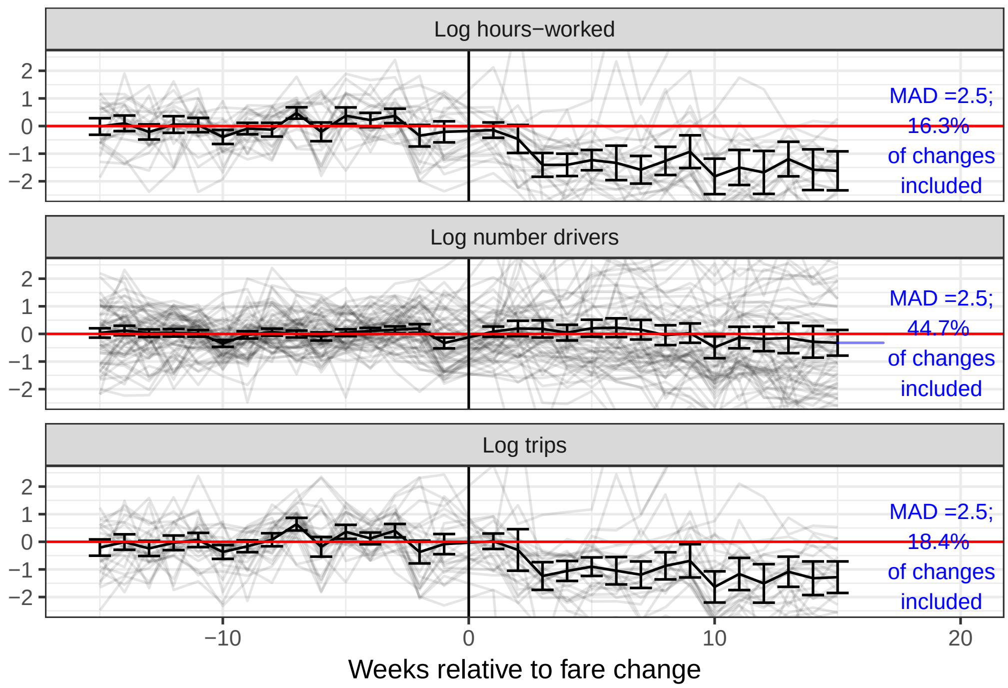

Notes: This figure plots the difference between cities experiencing a fare change and a synthetic control city comprised of similar cities that contemporaneously did not have a fare change. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. Hours-worked is the total number of hours worked (i.e., had the Uber app on and were available for dispatch or were driving passengers, or enroute to pick up passengers) by drivers in that that city, in that week. The number of drivers is the total number of drivers in that city that worked at least some number of hours. The number of trips is the number of trips completed. Only synthetic control estimates where the maximum absolute pre-period difference between the focal city and synthetic control was less than 2.5 are included. The sold dark line indicates the average value, weighted by the pre-period MSE, with error bars indicating a 95% CI.

In the top panel, we can see that the number of trips taken declines after a fare increase. Although we had indirect evidence that this would be the case, this is the first time we have looked at the effects on trips-taken directly. The synthetic control plot indicates why the regression performed poorly in terms of obtaining pre-period balance—there are substantially fewer paths plotted, as far fewer synthetic control estimates achieve good balance in the pre-period.

In the middle panel, the outcome is the log number of drivers active on the platform. There is little evidence of a change, though the estimates are quite imprecise. Note that the lack of a pre-period increase in the number of drivers suggests that one adjustment margin discussed earlier—namely a pushing out of the supply curve—was not a factor in explaining how the market adjusted. Although drivers do enter and exit freely, Uber could have increased efforts to acquire or "on-board" drivers to push out the supply curve.

In the bottom panel, the outcome is the log number of hours-worked. The total number of hours worked does decline, at least after about week 3, post fare increase, though there is substantial heterogeneity. Recall that by this week, from Figure 10, hourly earnings are still slightly higher. The large change in quantity of hours supplied—despite only small changes in the hourly earnings rate—is consistent with the characterization of drivers as supplying labor highly elastically to Uber. The change in hours-worked, combined with the change in utilization, together approximately explain the reduction in trips taken.[14]

4.9 Comparison of estimates

We have presented three different ways of obtaining estimates of the effects of fare changes. We plot all three approaches in Figure 13. The between-city results are fit using only those cities that comprise the within-city sample (i.e., that had UberBlack). The synthetic control estimates are presented using the pre-period cutoff of 2.5 (as in Figure 10). The general pattern of results is similar across methods. Following a fare increase, hourly earnings increase at first, then decline over time. In all cases, the utilization rate falls, as does the average surge level.

Despite the qualitatively similar pattern, there are substantial differences in some of the point estimates. For example, the synthetic control estimate for hourly earnings shows a much larger initial pass-through of the fare change but then a larger reduction in hourly earnings by the last period. There is

[14] There can be changes in trip composition that make this not a pure accounting identity.

some evidence that there is greater initial pass-through of a fare cut in the within-city analysis, though this could simply be sampling variation, as the confidence intervals generally overlap.

Figure 13: Comparison of of the effects of fare changes using different empirical approaches

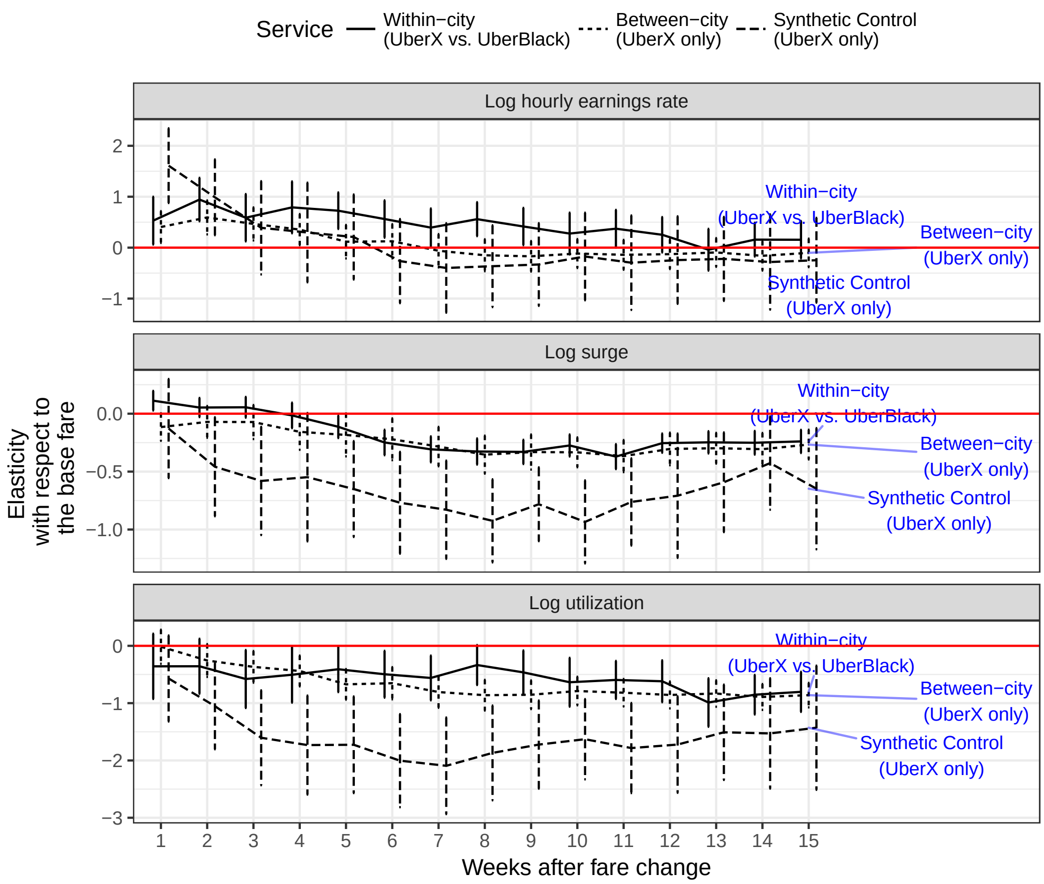

Notes: This figure shows the by-week estimates of the effects of a fare change on various market outcomes. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. The hourly earnings rate is the total earnings by drivers (excluding costs and Uber's 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. Three different estimates are shown: the between-city panel estimates, the within-city panel estimates, and the synthetic control estimates, using the 2.5 pre-period cut-off and weighting by the inverse pre-period MSE.

As our interest is primarily in the long-term effects of fare changes, in Table 3, we report the point estimates (and standard errors) for each of our three measures related to driver hourly earnings, at week 15 post fare change. Starting in the leftmost column, the effects of changes in the base fare on "long-run" driver hourly earnings is minimal, varying not only in magnitude, but also in sign. The most negative (and least precise) is from the synthetic control, which suggests a 10% increase would lower driver hourly earnings by about 2.5%, whereas the within-city estimates, suggest it would raise earnings by about 1.6%. None of these estimates are difference from zero at conventional significance levels.

For utilization, the within- and between-city point estimates are quite close, with effect around -8% for a 10% fare increase. In contrast, the (imprecise) synthetic control estimate finding a much larger effect. For surge, again, the within- and between-city point estimates are quite close, with a 10% fare increase reducing average surge by about 2.5%. The synthetic control results are both larger and much less precise.

Table 3: Comparison of Week 15 elasticities by different estimation method

Estimation	Log hourly earnings	Log utilization	Log surge
Between-city	-0.102 [0.14]	-0.86 [0.11]	-0.267 [0.06]
Within-city	0.155 [0.18]	-0.803 [0.17]	-0.239 [0.05]
Synthetic Control	-0.247 [0.42]	-1.429 [0.54]	-0.648 [0.26]

Notes: This table reports the elasticities for various outcomes with respect to the base fare index, using three different methods, along with the associated standard errors. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. Three different estimates are shown: the between-city panel estimates (see Section 4.2), the within-city panel estimates (see Section 4.5), and the synthetic control estimates (see Section 4.6), using the 2.5 pre-period cut-off and weighting by the inverse pre-period MSE. The reported elasticity is the for the 15th week following a fare change. 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. See Section 2.1 for a description of the sample.

4.10 Labor supply elasticities

We have characterized the labor supply to Uber as not only being highly elastic, but actually indistinguishable from infinity: we observe large changes in hours-worked, despite essentially no detectable change in hourly earnings rates. In Figure 14, the leftmost panel plots the log total hourly earnings. In the next panel over, the outcome is log hours-worked per driver. In the first period after the fare increase, the hourly earnings rate increases about 40%. That same period, hourly-worked increases by about 20%, implying a 0.5 labor supply elasticity. However, this is clearly not an elasticity that applies in later weeks: note that by week 15, there is no difference in observed hourly earnings, and yet there is about a 20% reduction in hours-worked, implying an infinite labor supply elasticity.

Figure 14: Effects of fare changes on driver hourly earnings rate, hours per driver, and driver "churn" measures

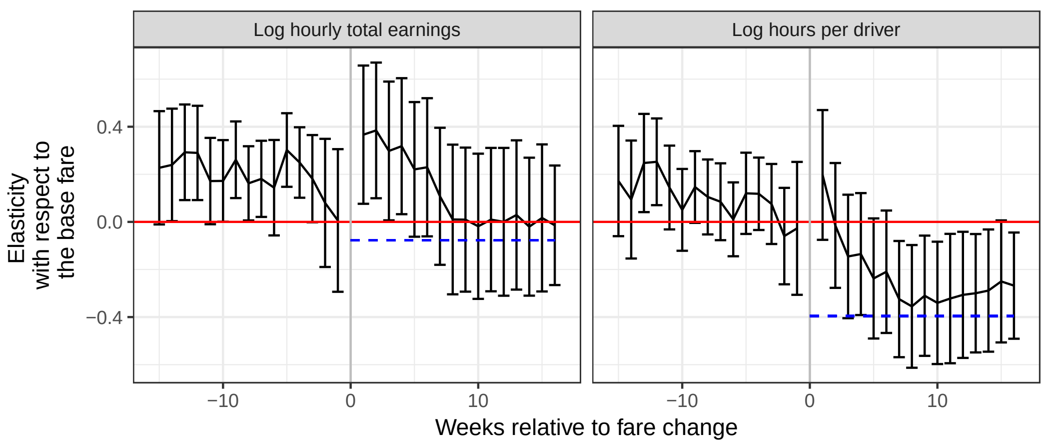

Notes: This figure plots the effects of changes in the UberX base fare on hourly total earnings (which includes promotional payments) and (2) hours-worked per driver. These effects are from an OLS estimation of Equation 2. The sample is a panel of US cities—see Section 2.1 for a description. The x-axis are weeks relative to a fare change. The independent variable is the base fare index. The base fare index is the price to riders of an un-surged, 6 mile, 16 minute trip in that city, in that week. As all outcomes are in logs, the point estimates can be interpreted as elasticities, which are estimated using both fare increases and decreases. The horizontal dashed blue line in the post period indicates the "long-run" effect corresponding to Equation 1. 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.

How do we reconcile findings with highly credible (and finite) micro estimates from this exact same empirical context—namely Angrist et al. (2017) and Chen et al. (2017)? Part of the explanation is realizing that market labor supply elasticities can, of course, look quite different from the average elasticities of drivers on the platform, as the market labor supply depends the marginal drivers. As a simple but extreme example, imagine that drivers have idiosyncratic reservation wages, but no intensive margin elasticity—they either work of they do not. If there are sufficient numbers of drivers with reservation wages near the market wage, then we might find the market labor supply is highly elastic, even though each inframarginal driver has a labor supply elasticity of zero at the market wage. In short, it is clear that a group of highly inelastic workers at the going market wage can create a highly elastic market labor supply.

In the Uber context, is seems quite likely that the region around the reservation wage is "thick" with workers that highly elastic because their reservation wage is what they can get on the competitor platform. As such, both platforms would have lots of bunching of reservation wages right near the market wage.

In terms of the Chen et al. (2017) results, we can see how drivers could be highly elastic with respect to the platform, but still evince "micro" behaviors consistent with a finite elasticity. Suppose there exists a competitor ride-sharing platform that offers exactly the same hourly earnings rate and the same pattern of wages. Would-be drivers could be indifferent between the two platforms—suppose there is an ϵ of switching cost. Drivers would still pick and choose when to work, weighing off changes in their outside option to what the current platform wage was, obtaining substantial surplus but also exhibiting a "reasonable" Frisch labor supply elasticity. Now, suppose that the alternative platform could offer hourly earnings rates $w + \Delta w$, with $\Delta w > \epsilon$, no matter how many drivers switched. All drivers would switch, and we would observe an effectively infinite labor supply elasticity *to the platform* even though on the new platform, they go back to working in a manner consistent with a finite labor supply elasticity.

5 Discussion and conclusion

The key findings of the paper are that following a fare change, the markets adjust primarily through changes in driver utilization, and there is little long-run effect on driver hourly earnings (with a "long-run" being about 8 weeks), consistent with drivers supplying labor highly elastically to Uber. Although our panel specifications show how this adjustment occurred, the main elasticity result can more or less be seen by simply comparing the time series over the base fare index and hourly earnings—despite large reductions in fares, there has been no detectable trend in average hourly earnings. Interestingly, this lack of price effects *on average* seems be 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.

Our estimates come from fare changes made over a certain range—presumably much larger changes would have different effects. A key assumption in our empirical analysis is that fare changes are exogenous. When markets were less mature and there was less practical experience, assuming Uber would change fares as-if at random is a plausible assumption. With hindsight about how the market seems to adjust, Uber might make different choices. Building on the model presented in Section 3, we model Uber's incentives with respect to fare changes and then consider our predictions versus our empirical results.

In our original presentation of the model in Section 3, we did not consider Uber's commission, or the fraction of receipts it collects. Suppose that Uber collects a fraction τ. Uber sets both p (or sets b, with knowledge that the change in m will lead to their desired p) and a commission $\tau \in [0, 1]$. For a given p, Uber picks the τ that would just allow the market to clear. As such, Uber has an incentive to drive up utilization as high as possible, as this lets them provide a given quantity of rides with the lowest total payment to drivers, minimizing costs.

Let $\bar{x}$ be the highest possible utilization for the market, given the technical limitations of the platform. It could be the case that $\bar{x}$ depends on the scale of the market—as in Arnott (1996)—but we will set aside this consideration, as many of the approaches to increasing utilization are technical rather than economic, such as "forward dispatch" (matching drivers before their current fare is finished based on predicted drop-off time), having riders re-locate slightly before pick-up, using up-front pricing to charge more of utilization-reducing trips, and so on.

Uber's profit maximization problem, given a fixed utilization, $\bar{x}$, is illustrated Figure 15. There is a single demand curve, but two supply curves that differ in their units—one for driver hours, $S(w)$, and one for hours of transportation, $\bar{x}S(w)$. At a given p, market clearing requires that $D(p) = \bar{x}S(w)$. Let the quantity of trips taken at the market clearing price be Q, which means drivers provide Q hours of transportation, which requires $Q/\bar{x}$ hours of work. The portion shaded red in the diagram is $wQ/\bar{x}$ which is the total amount paid to drivers. The lightly shaded area is total revenue minus the wage bill or

$$\begin{aligned} p\bar{x}\frac{Q}{\bar{x}} - w\frac{Q}{\bar{x}} &= (p - w/\bar{x})Q \\ &= \tau p Q, \end{aligned}$$

which is Uber's profits, assuming zero marginal intermediation costs.[15] Uber makes the the blue rectangle, $\tau p Q$, as large as possible, subject to market clearing.

To actually solve this problem, consider that for a given p, Uber chooses a τ so that exactly $D(p)$ hours of transportation will be provided. The optimal commission for a fixed p and $\bar{x}$ is thus

$$\tau^* = 1 - w^*(D(p)/\bar{x})/p\bar{x},$$

where $w^*(\cdot)$ is the inverse labor supply curve. It will be useful to define hourly

[15]The value $p\bar{x}$ can be constructed graphically by taking the diagonal line from the origin to $(Q/\bar{x}, p)$, which has a slope of $p\bar{x}/Q$. This diagonal line evaluated at Q then has a height of $p\bar{x}$.

Figure 15: Uber's profits at the optimal commission

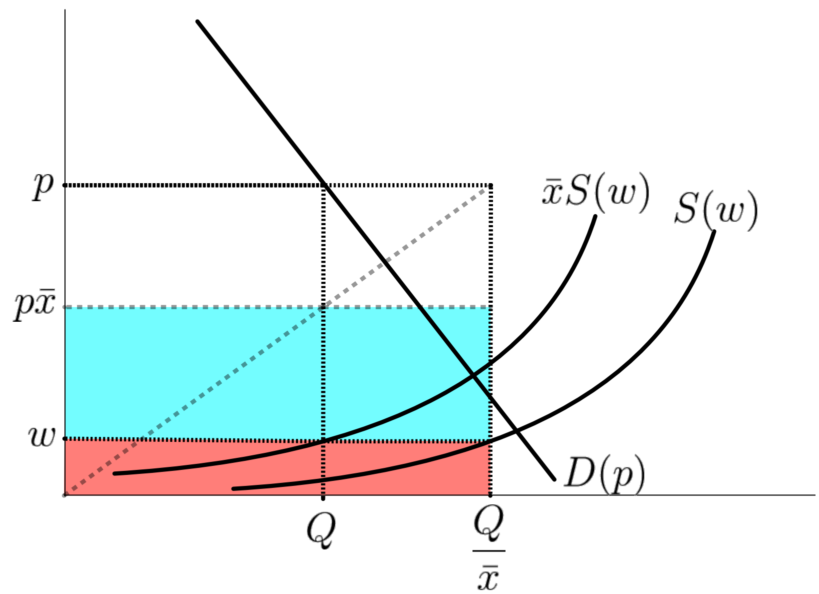

earnings as a function of the price, rather than quantity, or $w(p) = w^*(D(p)/\bar{x})$. With the optimal commission, Uber's profits are

$$\pi(p) = D(p)\left(p - w(p)/\bar{x}\right). \tag{5}$$

Note that $w(p)/\bar{x}$ is the per-ride hour cost Uber faces when $D(p)$ hours of transportation are demanded. If this cost did not depend on p, then Uber's problem would simply be the monopolist's pricing problem. Instead, Uber must consider the effect their choice of p will have on their unit labor costs. With a higher p, it earns more from each ride, both because the price is higher, but also because it can meet that demand with fewer drivers, which means they earn a lower hourly rate for the monopsony reason.

To the extent Uber faces an upward sloping labor supply curve, it has an incentive to restrict the quantity of rides more than a monopolist facing a competitive labor market. Proposition 2 formalizes this argument, giving what is essentially a marginal cost markup rule (Equation 6) for a double monopolist/monopsonist.

Proposition 2. *The optimal base fare is a markup of the costs of providing a unit of transportation ($w/\bar{x}$) such that*

$$p^* = \frac{w}{\bar{x}} \left(\frac{1}{1 + \frac{1}{\epsilon_p^D}} \right) \left(1 + \frac{1}{\epsilon_w^S} \right). \tag{6}$$

Proof. The first order condition from maximizing Equation 5 is

$$D'(p) \left[p - w(p)/\bar{x} \right] + D(p) \left[1 - w'(p)/\bar{x} \right] = 0 \tag{7}$$

which we can write as $\epsilon_p^D \tau = w'(p)/\bar{x} - 1$. As the market needs to clear, $D(p) \equiv \bar{x} S(w(p))$ and so

$$w'(p) = \frac{\epsilon_p^D}{\epsilon_w^S} \bar{x} \left(1 - \tau \right).$$

We can use this to re-write Equation 7 as

$$\tau^* = \frac{\epsilon_p^D - \epsilon_w^S}{\epsilon_p^D \left(1 + \epsilon_w^S \right)}$$

and finally, using $\tau^* = 1 - w/p\bar{x}$, we get the markup rule above. □

There are several implications of Proposition 2. First, if our empirical finding that drivers supply labor highly elastically to Uber holds, then Uber applies a standard markup based solely on the elasticity of demand, ignoring the monopsonist component of the markup rule, as the ϵ_w^S term in Equation 6 goes to 1. Second, as $\bar{x}$ increases, such as through technological advances or scale economies, the profit-maximizing base fare declines and/or the optimal τ increases. Technological advances provide an intriguing as-if explanation for why Uber has continually reduced the base fare, on average. Furthermore, our finding that if anything hourly earnings rose following a fare cut is consistent with Uber being unprofitably in the too-elastic region of the demand curve when the fare was cut.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synth: An R package for synthetic control methods in comparative case studies," *Journal of Statistical Software*, 2011, *42* (i13).

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," Technical Report, 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.

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.

Chen, M. Keith and Michael Sheldon, "Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform," *Working paper*, 2015.

— , **Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen**, "The value of flexible work: Evidence from Uber drivers," *Working Paper*, 2017.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, "Using big data to estimate consumer surplus: The case of uber," Technical Report 2016.

Dube, Arindrajit and Ben Zipperer, "Pooling multiple case studies using synthetic controls: An application to minimum wage policies," *Working Paper*, 2015.

Edlin, Aaron S and Pinar Karaca-Mandic, "The accident externality from driving," *Journal of Political Economy*, 2006, *114* (5), 931–955.

Einav, Liran, Chiara Farronato, and Jonathan Levin, "Peer-to-peer markets," *Annual Review of Economics*, 2016, *8*, 615–635.

Farber, Henry S., "Is tomorrow another day? The labor supply of New York City cabdrivers," *The Journal of Political Economy*, 02 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.

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.

Horton, John J and Richard J Zeckhauser, "Owning, using and renting: Some simple economics of the "sharing economy"," Technical Report, National Bureau of Economic Research 2016.

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.

Katz, Lawrence F and Alan B Krueger, "The rise and nature of alternative work arrangements in the United States, 1995-2015," *Working Paper*, 2016.

Kuhn, Peter, "Is monopsony the right way to model labor markets? a review of Alan Manning's 'Monopsony in Motion'," *International Journal of the Economics of Business*, 2004, *11* (3), 369–378.

Lindsay, Cotton M. and Bernard Feigenbaum, "Rationing by waiting lists," *The American Economic Review*, 1984, *74* (3), 404–417.

Mankiw, N Gregory and Michael D Whinston, "Free entry and social inefficiency," *The RAND Journal of Economics*, 1986, *17* (1), 48–58.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2003.

Parry, Ian WH, Margaret Walls, and Winston Harrington, "Automobile externalities and policies," *Journal of Economic Literature*, 2007, *45* (2), 373–399.

Rosenbaum, Paul R and Donald B Rubin, "The central role of the propensity score in observational studies for causal effects," *Biometrika*, 1983, *70* (1), 41–55.

Sundararajan, Arun, *The sharing economy: The end of employment and the rise of crowd-based capitalism*, MIT Press, 2016.

Varian, Hal R, "Computer mediated transactions," *The American Economic Review*, 2010, *100* (2), 1–10.