Journal of Urban Economics 122 (2021) 103318

Contents lists available at ScienceDirect

Journal of Urban Economics

journal homepage: www.elsevier.com/locate/jue

The congestion costs of Uber and Lyft

Matthew Tarduno[a,b,1,*]

[a] Department of Agricultural and Resource Economics, University of California, Berkeley, United States
[b] Energy Institute at Haas, United States

A B S T R A C T

I study the impact of transportation network companies (TNC) on traffic delays using a natural experiment created by the abrupt departure of Uber and Lyft from Austin, Texas. Applying difference in differences and regression discontinuity specifications to high-frequency traffic data, I estimate that Uber and Lyft together decreased daytime traffic speeds in Austin by roughly 2.3%. Using Austin-specific measures of the value of travel time, I translate these slowdowns to estimates of citywide congestion costs that range from \$33 to \$52 million annually. Back of the envelope calculations imply that these costs are similar in magnitude to the consumer surplus provided by TNCs in Austin. Together these results suggest that while TNCs may impose modest travel time externalities, restricting or taxing TNC activity is unlikely to generate large net welfare gains through reduced congestion.

1. Introduction

Transportation network companies (TNC) like Uber and Lyft have grown rapidly over the past decade to become integral parts of urban transportation systems. A small but growing literature has attributed to these companies benefits that include billions in annual consumer surplus (Cohen et al., 2016), reductions in drunk driving (Greenwood and Wattal, 2015), and flexible work (Judd and Krueger, 2016; Angrist et al., 2017).

The costs of TNC expansion, however, have yet to receive commensurate treatment in the economics literature. Most notably, TNCs have been accused of contributing to traffic congestion (San Francisco Transit Authority, 2018; Schaller Consulting, 2018), but existing studies of the impact of TNCs on congestion are few, arrive at varied conclusions, and do not quantify the implied congestion costs (Li et al., 2019; Erhardt et al., 2019). Back of the envelope calculations suggest these costs could be substantial. A 2017 Inrix report, for example, placed the annual cost of congestion to US drivers at \$305 billion (Inrix, 2017)—roughly two orders of magnitude larger than estimates of national consumer surplus provided by Uber (Cohen et al., 2016). This suggests that if TNCs have even a modest impact on traffic congestion, the negative externalities associated with lengthening travel times could offset consumer surplus benefits. Understanding how and whether TNCs impact traffic congestion therefore plays a crucial role in determining appropriate policy response to the continued growth of these companies.

Two identification problems, however, make causal inference difficult when studying the relationship between TNC activity and traffic congestion. First, Uber and Lyft likely select entry locations based on trends in city-level characteristics unobservable to the econometrician. Comparisons that leverage differences in TNC entry dates across locations may therefore suffer from reverse causality. Second, within-city time series regressions may be biased by omitted variables (e.g., gentrification) which are serially correlated with TNC activity and also impact congestion.

In this paper I leverage a natural experiment in Austin, TX to circumvent these identification challenges: On May 9th, 2016, both Uber and Lyft unexpectedly exited Austin following a vote that upheld a city ordinance requiring driver background checks. I combine this variation in TNC activity with novel and granular Bluetooth traffic speed data, and setting-specific estimates of the value of travel time to answer two research questions. First, do transportation network companies impact traffic congestion? And if so, what are the travel-time related costs or benefits of TNC operation?

This setting informs two empirical strategies: a difference in differences comparing pre- versus post-May 9th traffic speeds in 2015 (where both companies operated year round) to 2016 (where both companies exit on May 9th), and a regression discontinuity in time. Across specifications, I find evidence of modest increases in traffic speeds following the exit of Uber and Lyft. Difference in differences results suggest that across all hours, traffic speeds increased roughly 1% following the exit of Uber and Lyft. 7 am. to 7 p.m. traffic speeds increased by 2.3%, with the largest TNC-related slowdowns occurring during the middle of the day (11 a.m. to 2 p.m.). Using setting-specific estimates of value of the travel time, I calculate that Austinites would be willing to pay roughly

* Correspondence to: Department of Agricultural and Resource Economics, University of California, Berkeley, United States.
E-mail address: tarduno@berkeley.edu

[1] I thank Michael Anderson, James Sallee, Meredith Fowlie, Aprajit Mahajan, Matthew Gibson, Alejandro Favela Nava, Jenya Kahn-Lang, the participants of the 2019 Giannini Foundation of Agricultural and Resource Economics Student Conference, as well as the Editor (Matthew Turner) and this paper's two anonymous reviewers for their valuable feedback. I also thank John Clary at the Austin Transportation Department for providing technical support during the data acquisition phase.

https://doi.org/10.1016/j.jue.2020.103318
Received 22 October 2019; Received in revised form 25 November 2020
Available online 21 January 2021
0094-1190/© 2020 Elsevier Inc. All rights reserved.

$33 to $52 million annually to avoid these slowdowns. Back of the envelope calculations suggest that these figures are a small fraction (4–6%) of total Austin-area congestion costs, and are roughly the size of estimates of the consumer surplus associated with TNC operation in Austin.

These findings improve on the existing literature in three ways. First, this is to my knowledge the only paper to use the exit of Uber and Lyft to study the impacts of TNCs on congestion. This translates to weaker identifying assumptions than those imposed in analyses leveraging the staggered expansion of these companies. Second, I extend existing analyses by mapping changes in travel speeds to changes in travel time costs, providing the first estimates of the congestion costs associated with TNC activity. And third, the spatial and temporal granularity in the Bluetooth data allows me to perform analyses that contribute to a more complete picture of the heterogeneous impacts of TNC activity on traffic congestion.

These findings also provide several important takeaways for policymakers. First, TNC activity can be viewed roughly as a transfer, as the consumer surplus enjoyed by TNC passengers is of similar size to the time loss incident on incumbent drivers. Second, it is difficult to rationalize TNC quantity restrictions purely on welfare grounds, as the lost consumer surplus may outweigh travel time gains. In other words, even if TNC regulation is more politically achievable than are price-based congestion controls, TNC regulation appears (at least in the Austin case) to be a poor tool to address congestion-related externalities. Relatedly, the relatively modest impacts of TNCs on traffic congestion in Austin suggest that congestion taxes targeted specifically at ridesharing companies are unlikely to result in large traffic-related welfare gains. Lastly, the fact that speeds slow in response to TNC activity suggests TNCs add vehicles miles traveled (VMT) to the transit system. In other words, the VMT avoided by sharing rides are outweighed by additional trips induced by the availability of TNCs.

The rest of this paper is organized as follows. Section 2 describes related literature and background. Section 3 details the events that precipitated the departure of Uber and Lyft from Austin. Section 4 outlines the data sources. I describe my empirical strategy and threats to identification in Section 5, and present results in Section 6. Section 7 concludes.

2. Background and related literature

Traffic congestion is a significant urban disamenity. It is costly (Inrix, 2018), it is associated with lower self-reported happiness (Anderson et al., 2016), and it comes with considerable co-costs in terms of noise and pollution (Currie and Walker, 2011). Although a tax is the canonical policy prescription for congestion (Vickrey, 1969), both theory and empirics suggest that because targeting individual contributions to congestion is difficult, realistic congestion pricing instruments (e.g., cordon charges) may fall well short of the welfare gains achievable by a hypothetical first best policy (Knittel and Sandler, 2018; Prud'Homme and Bocarejo, 2005). This, coupled with the potential political advantage of TNC regulation over comprehensive congestion taxation suggests that understanding the sign and magnitude of TNC related time costs or savings will be important for informing city-level policy. Indeed, several cities have already moved to regulate TNCs in the name of congestion. New York City, for example, cited congestion as a motivation for its 2018 ridesharing cap (New York Times, 2018). As of 2020, San Francisco, New York, and Chicago have all imposed "congestion fees," levied on TNC trips in the city center (New York Times, 2019). Outside of the US, cities like London and Vancouver have weighed congestion impacts as they deliberate over TNC policy (Reuters, 2019; Vancouver Sun, 2019).

As a number of other observers have noted, however, the impact of TNCs on traffic speeds is theoretically ambiguous. While survey data from Rayle et al. (2014) and Clewlow and Mishra (2017) suggest TNCs induce trips, and Mangrum and Molnar (2018) demonstrate that taxis—the closest analog to TNCs—increase congestion on the margin, Judd and Krueger (2016) show that in five of six US cities, Uber drivers spend a significantly higher fraction of their time with a passenger in their vehicle than do taxi drivers. This ride-sharing effect could attenuate or outweigh the effect of induced trips. There may also be complementarities between TNCs and public transit: Hall et al. (2018) use a difference in differences design on measures from the National Transit Database to conclude that Uber is indeed a complement to public transportation. It is unclear, though, whether complementarity between TNCs and public transit will result in more or fewer vehicle trips.

To date there exists little econometric work on whether TNCs cause traffic congestion, and existing results arrive at varied conclusions. Li et al. (2019), for example, use city-level congestion measures and differences in Uber's entry date to estimate the company's impact on congestion, concluding that Uber improves city-level congestion measures. Erhardt et al. (2019), on the other hand, use 2010 and 2016 Inrix traffic data and scraped measures of Uber activity to calibrate a traffic engineering model of San Francisco. They conclude that ridesharing companies were responsible for significant (30%) increases in vehicle hours traveled. In addition to the fact that these studies reach contradicting conclusions, the identification concerns outlined in the introduction suggest value in reassessing this question using a natural experiment.

3. Natural experiment

Austin, TX, is the 11th largest incorporated place in the United States and suffers from considerable congestion: According to Inrix, Austin ranked 14th nationally and 72nd globally in the number of average hours lost to congestion per driver. Cities with similar levels of per-driver congestion costs include San Diego, Berlin, and Manchester. Both Uber and Lyft began operating in Austin in 2014.

In December 2015, the Austin City Council passed Ordinance No. 20151217-075, which imposed a series of regulations on TNCs, including data requirements, restrictions on idling locations, and most controversially, fingerprinting requirements to facilitate driver background checks (The City Council of Austin, 2015). Proposition 1, sponsored by Uber and Lyft, attempted to overturn this ordinance. On May 7th, 2016, the Proposition was defeated in a citywide vote, with 56% of voters casting against (The Texas Tribune, 2016). In protest, Uber and Lyft exited the Austin market on May 9th (New York Times, 2016). 13 months later, Uber and Lyft re-entered Austin as Governor Greg Abbott signed into law HB 100, which overturned Austin's local ordinance (The 85th Texas Legislature, 2017). This variation in TNC activity provides the basis for my empirical identification.

During the yearlong absence of Uber and Lyft, Austin was not without ridesharing. A number of smaller TNCs entered the market or expanded their Austin presence following the defeat of Proposition 1. In date of their arrival in Austin, these companies are: GetMe (December 2015), Fare (Mid-May 2016), Fasten (June 1st, 2016), Tride Technologies (June 15th, 2015), and RideAustin (June 16th, 2016). Wingz, which provides rides to and from the airport, also started operating in Austin in May of 2016. A survey of Austin commuters conducted in November 2016 by Hampshire et al. (2017) offers a view of take up of these alternative rideshare companies. RideAustin held the largest market share (47.4%), followed by Fasten (34.5%), Fare (12.9%), GetMe (2.8%), Wingz (1.6%), and Tride (0.4%). Informed by the Hampshire et al. (2018) survey and the universe of RideAustin's 2016 trip-level data, I am able to infer the level of total TNC activity in Austin following the exit of Uber and Lyft. I can therefore identify a window following the Proposition 1 vote where alternative TNC activity is negligible (see Section 5.1).

4. Data

I use data collected from an array of Bluetooth sensors along major roadways (both highway and surface-level) operated by the Austin Department of Transportation. Located inside traffic signal cabinets,

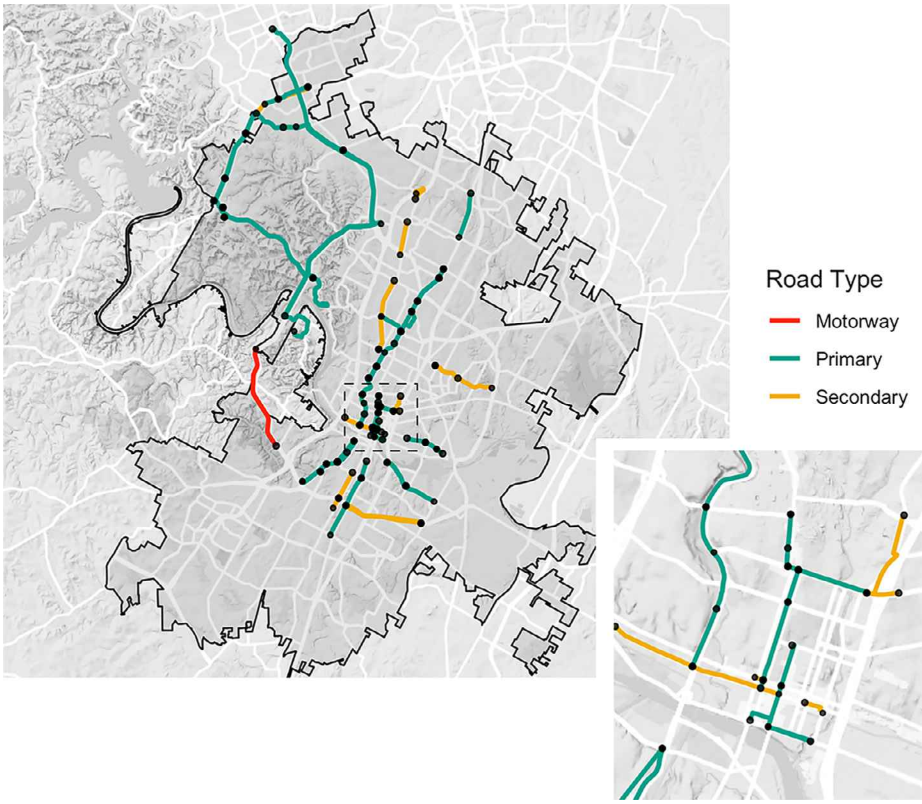

Fig. 1. Bluetooth Segment Locations in Austin, TX. *Notes:* Nodes represent terminal Bluetooth sensor locations for each of the 79 segments used in my analysis. Note that some sensors act as both origin and destination readers for different segments. Paths represent Google Maps recommended driving directions between endpoints of a given segment, colored by Open Street Map road type. Motorways are major divided highways, primary roads are large multi-lane roads that may or may not be divided. Secondary roads are typically two- to four-lane surface streets. The black line is the Austin city limit.

these sensors detect unpaired Bluetooth devices (e.g., smartphones, car systems) and estimate traffic speeds based on the movement of single devices (which are given unique anonymous identifiers) through the network of sensors.

I use an aggregated version of this dataset prepared by Post Oak Traffic Systems, which isolates device movements through specific road segments (henceforth *segments*), which are short sections along just one road. This company pre-processes the data in several ways. Data are aggregated at 15 min bins and represent the average speed across the segment for devices that appear at the origin reader first, and then the destination reader, and do not appear at any other sensors in the interim. These data are also filtered for outliers: only observations that fall within 75% of the IQR of the previous 15 observations are used in calculating speeds. This type of filtering is applied to combat bias from the movement of non-vehicle Bluetooth devices (like those carried by pedestrians) through the sensor network.

In addition to the data cleaning performed by Post Oak Traffic Systems, I further restrict my sample to consistently reporting sensors. Of the 430 total segments, I drop segments that report in fewer than 70% of days during each year (2015 and 2016) of the study period, leaving me with a panel of 79 segments. For robustness I also report results using a) all segments that report in more than 30% of study period days and b) only segments that report during 100% of study period days.

The 79 segments I use in my preferred specification are plotted in Fig. 1 and summarized in Table 1. The mean segment length is 0.72 miles, with minimum and maximum lengths of 0.06 and 3.8 miles, respectively. As shown in Fig. 1, my sample covers a range of road types. The smallest roads in my sample are two-lane roads, the largest are 7-lane roads, and the median segment is a 5-lane road. I observe 966,301 15-min speed reports during my study period. On average, a segment sees 4.77 devices move from origin to destination during each 15 min period, meaning that my data summarize roughly 4.6 million segment traverses. The average travel speed is 2.99 minutes per mile, which corresponds to 20.06 miles per hour. This figure is consistent with periods of significant congestion.

My variable of interest is minutes per mile, which has two advantages over miles per hour. First, a change of one mile per hour does not represent a constant damage over the domain of this variable: In terms of time lost, changing from 5 to 4 miles per hour is roughly 20 times as costly as changing from 20 to 19 miles per hour. Second, multiplying outcomes in minutes per mile by estimates of the value of time is a straightforward way to arrive at cost calculations from changes in traffic delays.

While novel and granular, the Bluetooth data bring challenges for estimation. First, in the raw data available on the Austin Open Data Portal, 61 of the 79 segments used in my analysis show the segment length changing over the course of the study period. While most of these adjustments are minor, and personal correspondence with Austin Transportation Department employees suggests that these adjustments likely reflect updated length measurements and not relocation of Bluetooth sensors, I nonetheless investigate the possibility that these segment length changes constitute a threat to identification in Appendix E. I use the updated length measurements for all speed calculations in all time periods. A second challenge is the possibility of Bluetooth sensors measuring the movement of pedestrians. If filtering does not eliminate all measurement error originating from Bluetooth devices used by Austinites walking or biking, and the use of these modes of transit is correlated with the period where Uber and Lyft exited Austin, the empirical strategies I describe below will arrive at biased estimates. I further investigate this in Section 5.4.

I compile several other datasets to augment my analysis. To control for weather-related shocks, I use precipitation and temperature data accessed through the National Oceanographic and Atmospheric Administration's National Centers for Environmental Information. To isolate a period of time where the impact of other TNCs is minimal, I use RideAustin's trip-level data. These data range from June 2nd, 2016 to April 13th, 2017, and are publicly available online (RideAustin, 2017).

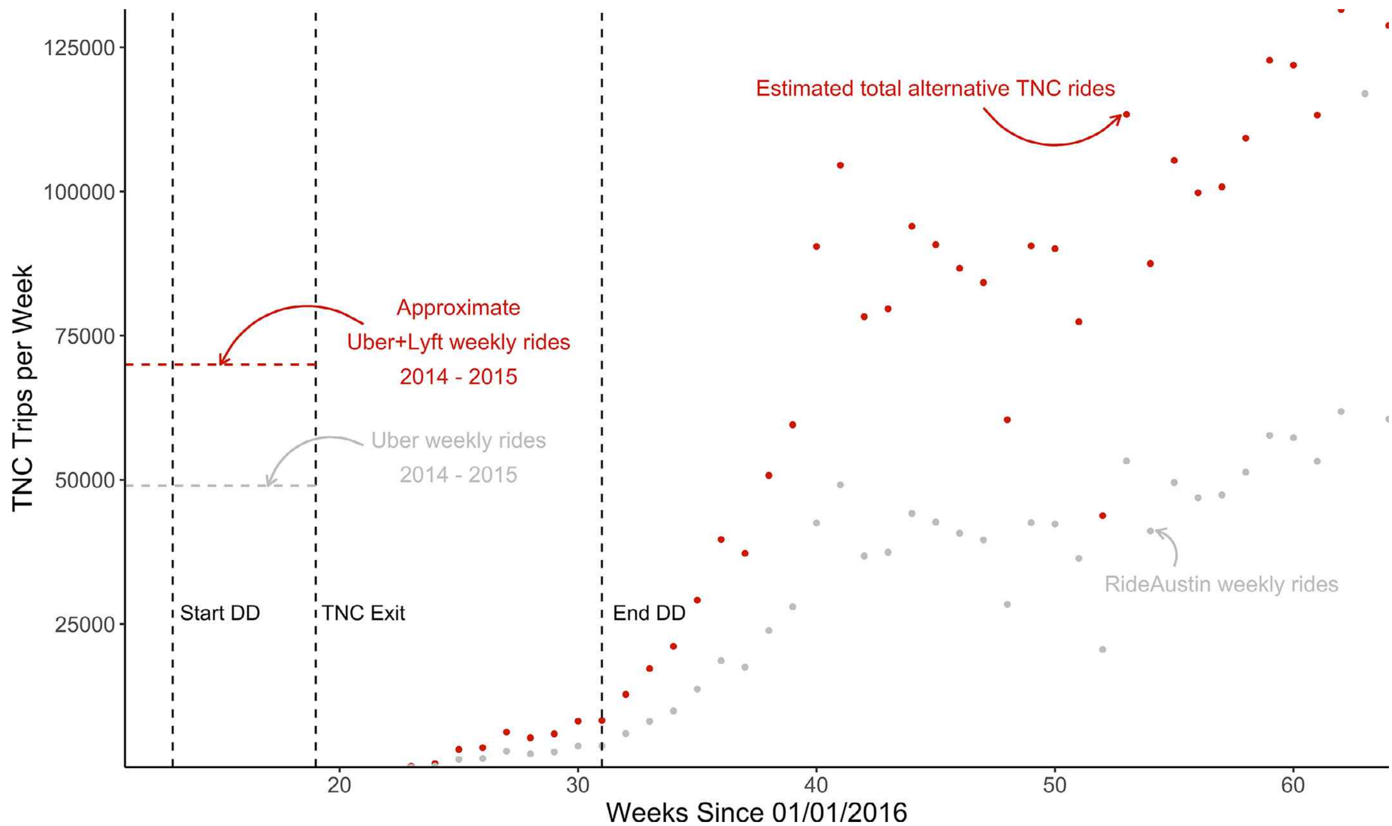

Fig. 2. Variation in TNC Activity. *Notes:* This figure displays the variation in ridesharing activity I use to identify the impact of TNCs on congestion. From left to right, the vertical lines represent the start of the 2016 difference in differences period (March 20th), the failure of Proposition 1 (May 9th), and the end of the 2016 difference in difference period (August 1st). The grey dotted line is the average number of Uber trips per week (as per an Uber Report on 2014–2015 operations). The red dotted line represents an estimate of total Uber and Lyft pre-exit activity, assuming a 30% Lyft market share. Note that because both Uber and Lyft entered Austin in 2014, the actual number of Uber trips in early May 2015 was likely much larger than 70,000 per week. The grey dots plot weekly RideAustin activity for the first 9 months of the company's operation. The red dots inflate the RideAustin data by the reciprocal of its November 2016 market share (47%) to provide an estimate for the total level of post Proposition 1 alternative TNC activity.

Lastly, I use two datasets to arrive at setting-specific value of time estimates. The first is the National Household Travel Survey (NHTS), which contains information on income and commuting habits. The second is a toll price and travel time dataset from the MoPac variable price freeway in Austin. These data were provided courtesy of the Central Texas Regional Mobility Authority and are further detailed in Appendix A.

5. Empirical strategy

5.1. Timeframe

I use Bluetooth traffic data from 2015 and 2016 to study the relationship between TNCs and congestion. I truncate this window to isolate periods where the variation in traffic speeds can be credibly attributed to the failure of Proposition 1. As described in Section 3, a number of TNCs entered the market following the exit of Uber and Lyft. Estimations using the entire yearlong suspension period as a comparison would therefore underestimate any changes relative to a TNC-free counterfactual. Informed by the universe of trips from RideAustin—the TNC with the largest market share during Uber and Lyft's absence—I truncate my estimation period on August 1st, 2016. Similarly, Austin hosts the South by Southwest Music Festival (SXSW) each March. I restrict my analysis to exclude the 2015 and 2016 festivals. This leaves me with data from March 20th to August 1st for both 2015 and 2016. The 2016 study period is plotted with TNC data in Fig. 2. Note that although the Austin Bluetooth data extend through 2019, significant portions of the spring are missing data from years 2017, 2018, and 2019, including Uber and Lyft's re-entry in May of 2017. While this rules out difference in differences specifications using later years, I am able to use data from different parts of 2017–2019 to perform placebo regression discontinuity estimates (see Appendix D).

5.2. Difference in differences

To study the effect of the exit of Uber and Lyft on travel times, I compare traffic speeds pre and post May 9th in 2016 (where Uber and Lyft exited) to 2015 (where both companies operated year-round). To capture heterogeneity in the congestion impacts across time of day, I perform this comparison within each hour of day, h (or equivalently, interacting each right-hand side term below with an hour of day dummy):

$$s_{i,y,t} = \alpha + \beta_h \delta_y \eta_t + \gamma_1 \delta_y + \gamma_2 \eta_t + \boldsymbol{\gamma}_3 \delta_y \boldsymbol{\theta}_i \cdot t + \boldsymbol{\gamma}_4 \boldsymbol{\theta}_i + \boldsymbol{\Gamma} \mathbf{X}_{y,t} + \epsilon_{i,y,t} \tag{1}$$

Where $s_{i,y,t}$ is the speed (in minutes per mile) measured over segment i on day t of year y. δ_y is a dummy that equals one for the year 2016, and η_t is a dummy that equals one for days (in any year) after May 9th. θ_i is a set of dummies for each road segment, and t is the signed number of days between a given date and May 9th of that year. $\mathbf{X}_t$ is a vector of controls that includes day of week fixed effects, holiday fixed effects, 10 10-degree daily temperature bins, and 10 daily precipitation level bins. The interacton betwen $\delta_y \eta_t$ is the treatment indicator, as it takes a value of 1 for observations after May 9th, 2016, and zero otherwise. $\delta_y \theta_i \cdot t$ are segment-year specific linear time trends.

The identifying assumption in the estimation of β_h—the effect of Uber and Lyft operation on travel speeds during a given hour of day h—is that conditional on seasonality and weather, the difference in travel speeds between 2016 and 2015 at hour h does not change after May 9th for reasons other than the operation of Uber and Lyft.

I calculate hour-specific congestion impacts with the goal of producing more accurate cost estimates. As I show in Appendix A, variable-toll data suggest that the value of travel time in Austin varies significantly from hour to hour. Similarly, the number of vehicles on the road peaks during rush hours. Together, this information suggests that the same change in traffic speeds could produce different aggregate congestion costs at different times of day. By matching hour-specific estimates of the impact of TNCs to hour-specific vehicle miles traveled (VMT) and hour-specific estimates of the value of travel time, my cost calculations account for temporal heterogeneity that pooled estimates may not reflect. To determine whether the convolution between hourly congestion impacts and hourly VOT is a first-order consideration, I also estimate a model pooling across hours of day. This estimator is Eq. 1, but run without interacting hour of day fixed effects with the right hand side variables. The rationale for this regression is to simulate what estimation and inference might look like using temporally aggregated data.

To investigate spatial heterogeneity, I estimate a model pooling over hours of day and allowing an idosyncratic treatment effect for each road segment. This model is equivalent to Eq. 1, but interacts the set of segment dummies with the treatment indicator, $\delta_y\eta_t$. $\boldsymbol{\beta}$ is now a 1x79 vector of segment-specific treatment effect estimates. Note that in this pooled Equation hour of day fixed effects are included in $\mathbf{X}_{y,t}$.

$$s_{i,y,t} = \alpha + \boldsymbol{\beta}\delta_y\eta_t\boldsymbol{\theta}_i + \gamma_1\delta_y + \gamma_2\eta_t + \boldsymbol{\gamma}_3\delta_y\boldsymbol{\theta}_i \cdot t + \boldsymbol{\gamma}_4\boldsymbol{\theta}_i + \boldsymbol{\Gamma}\mathbf{X}_{y,t} + \epsilon_{i,y,t} \tag{2}$$

5.3. Regression discontinuity

Lastly, I estimate a regression discontinuity model, again estimating hour-specific treatment effects (β_h) by interacting each term in the regression Equation with a set of hour of day fixed effects.

$$s_{i,t} = \alpha + \beta_h\eta_t + \boldsymbol{\gamma}_1\boldsymbol{\theta}_i + \boldsymbol{\gamma}_2\boldsymbol{\theta}_i \cdot t + \boldsymbol{\gamma}_2\boldsymbol{\theta}_1 \cdot t^2 + \boldsymbol{\gamma}_4\boldsymbol{\theta}_i + \boldsymbol{\Gamma}\mathbf{X}_t + \epsilon_{i,t} \tag{3}$$

The identifying assumption for β_h is that conditional on weather, potential outcomes (traffic speeds) in hour of day h are continuous about May 9th, 2016. While the identifying assumption for the RD is arguably weaker than that of the difference in differences estimator, the RD will produce estimates of the short-term response to the exit of Uber and Lyft. As such, I rely on the difference in difference estimator to produce my preferred annual congestion cost figures.

5.4. Threats to identification

Threat 1: Contemporaneous shocks. The identifying assumptions in both the RD and DID estimates rely on the absence of *year*post*-specific shocks to Austin area travel speeds. The end of the University of Texas, Austin (UT) school year, for example, presents a potential threat to identification if university-related traffic activity differed substantially between 2015 and 2016. In Appendix D, I use placebo exit dates to determine whether or not shocks that create regression discontinuity estimates on the order of my reported coefficients are empirically common. Fig. D.2 displays coefficient estimates using the actual exit date in relation to the distribution of coefficients from 134 regression discontinuities using placebo exit dates, 13 of which were chosen to line up with the beginning/end of a UT semester. 5 of the 134 placebo coefficients (4%) are more negative than the estimates using the actual TNC exit date, zero of which correspond to the start/end of a UT semester. This placebo test therefore suggests that shocks that produce RD estimates on the order of my estimates are empirically uncommon, and that my results are not likely a result of changing traffic patterns related to activity at UT.

Threat 2: Other modes of transportation. If the exit of Uber and Lyft led Austinites to substitute toward walking or biking and these trips were not dropped as outliers during data processing, β will not be identified. In other words, for other modes of transportation to bias my estimates, traffic speed must be mismeasured, and that mismeasurement must be correlated with the treatment.

Data on mode shares and mode speeds suggest that this type of bias cannot alone account for my results. Hampshire et al. (2017) suggest 1.8% of TNC users switched to bikes following Uber's exit. If TNCs made up 10% of Austin trips, and bikes constituted 1.53% (United States Census Bureau, 2015), this mode shifting represents an 11.8% increase in total bike trip volume. The average car in my sample took 2.99 min to traverse a mile—3.01 minutes per mile fewer than the 6 minutes per milemph) assumed by Google biking directions. These figures imply that for changes in bike shares to alone account for a change of 0.1 minutes per mileroughly the average treatment effect across daytime hours), bikes would need to constitute roughly 28% of observed Bluetooth samples *after* dropping extreme travel time outliers. This figure is inconsistent with the travel speeds implied by the movement of Bluetooth devices, which greatly exceed 10 miles per hour on average.

Nonetheless, I draw on a second traffic speed dataset to empirically examine this concern. In addition to Bluetooth sensors, the city of Austin also maintains pneumatic sensors that take periodic measurements of traffic speeds. While these measurements are not frequent enough to act as a replacement dependent variable, they do allow me to study the relationship between Bluetooth speed measurements and true traffic speeds by matching segments to pneumatic sensors.

While we should not expect pneumatic sensors to match segment speeds exactly (segments often include intersections), if there is significant switching to non-vehicular modes of transport that biases the Bluetooth speed measurements, this would be reflected in a change in the relationship between the two measurements. For example, say we have a segment-sensor pair, and prior to May 9th, 2016, when the pneumatic sensor reports a speed of 25 mph, the Bluetooth segment on average reports a speed of 20 mph. If there is bias from mode-switching, we would expect this relationship to change in the post period. Now, when the pneumatic sensor again registers 25 mph, the incresed number of non-filtered pedestrian datapoints biases the segment measurement downward, to, say, 18 mph.

To operationalize this anecdote, I match segments to pneumatic sensors, and run a regression of segment speeds on sensor speeds, allowing for a differential slope term interacted with a post May 9th2016 dummy. If I find a statistically (and economically) significant difference in slopes, I treat this as evidence of mode choice related bias. This exercise is detailed in Appendix B. I match 39 Bluetooth segments to pneumatic road sensors. In a simple regression with month of year and road segment fixed effects, I find little evidence to support pedestrian-induced bias in my estimates. As shown in Table B.1, the coefficient on the interaction between the post dummy and the pneumatic segment speed is not statistically different from zero, nor is it of meaningful magnitude.

Threat 3: TNC driving speeds. If TNC vehicles drive significantly slower or faster than the average non-TNC vehicle in a way that remains after filtering, the above estimates of β_h will be biased. During congested conditions it is unlikely that this should occur: if congestion slows all drivers, then travel time measurements from any subset of vehicles should be representative of average speeds. At free-flow traffic speeds, however, it is possible that TNCs drive faster (due to profit motive) or slower (idling to find riders) than non-TNC vehicles.

To test these concerns, I use public trip-level data from the startup RideAustin, which entered the market following the departure of Uber and Lyft. Following Mangrum and Molnar (2018), who construct "taxi races" to test whether different types of taxi travel at different speeds, I match RideAustin trips to Bluetooth segments, allowing me to test the null hypothesis that TNC vehicles drive at the same speeds as the average mix of vehicles.

This exercise is detailed in Appendix C. Over 221 trip-segment matches, I find that on average RideAustin vehicles traveled 0.03 minutes per mile slower while traversing a given segment than did the average device during the same time period. This difference is not statistically significant, nor should it meaningfully bias my results. Assuming

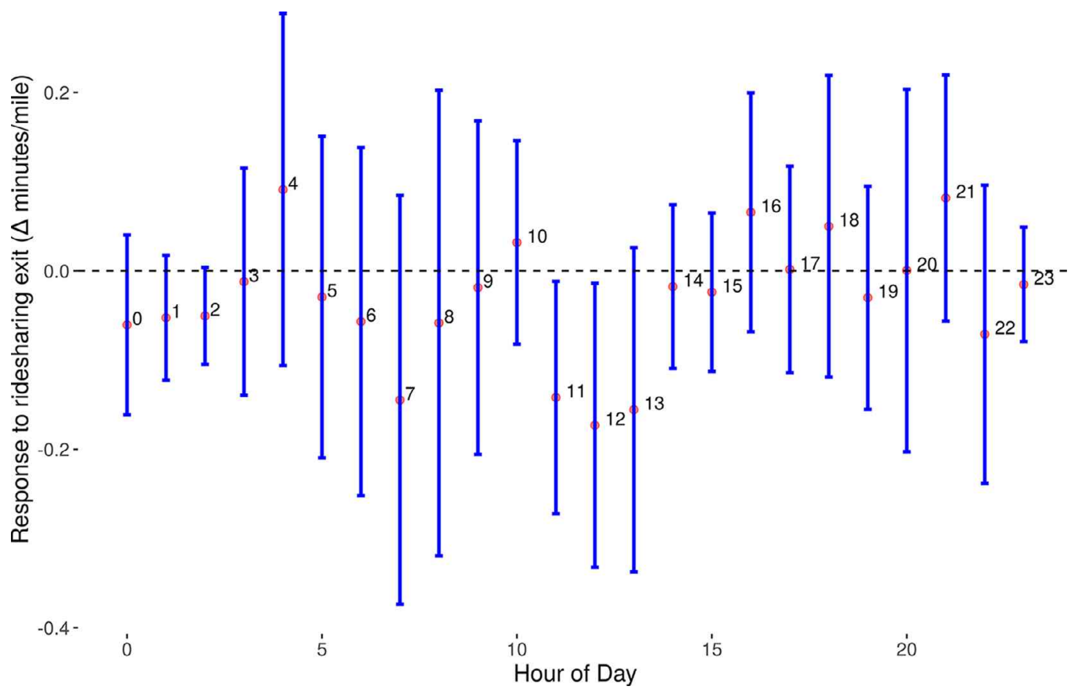

Fig. 3. Difference in Differences Results. *Notes:* Results from Eq. 1, a difference in differences comparing pre vs. post May 9th traffic speeds in 2015 (where both Uber and Lyft operated in Austin) to pre vs. post May 9th traffic speeds in 2016 (where both TNCs exited Austin). Points represent the estimated effect of TNC departure on traffic speeds (in minutes per mile) by hour of day. Controls include day of week, holiday, and segment fixed effects, segment-specific linear trends in days since May 9th, and flexible controls for temperature and precipitation. Bars reflect 95% confidence intervals from two-way standard errors clustered by segment-week. Traffic speed data were accessed through the City of Austin's Open Data Portal.

TNCs account for 10% of vehicle trips, for example, this difference in speeds implies a bias on the order of 0.003 minutes per mile—one to two orders of magnitude smaller than my estimates of the impact of TNCs on traffic speeds. To the extent that speed differences do generate bias, they will lead me to overstate improvements in traffic speeds resulting from a TNC ban.

6. Results and discussion

6.1. Traffic speeds

Across multiple specifications, I find evidence of modest increases in traffic speeds following the exit of Uber and Lyft. Results from my preferred specification (Eq. 1) are displayed in Table 2 and Fig. 3. Point estimates of changes in minutes per mile are largely negative, suggesting reduced congestion after the exit of Uber and Lyft. While the 95% confidence intervals for hour-specific estimates of changes in travel times generally include zero, an F-test rejects the null hypothesis of $\beta_h = 0\ \forall h$ ($p < 0.0001$). Although TNCs appear to negatively impact morning rush hour conditions, I estimate little change in evening rush hour speeds. The largest improvements in travel times following TNC exit come, surprisingly, between 11 a.m. and 2 p.m. Point estimates for off-peak hours (8 p.m. to 6 a.m.) are small and straddle zero. This pattern could be a result of TNCs comprising a higher share of vehicles during the middle of the day than during peak hours. Additionally, evening rush hour effects could be muted if TNC users are more likely to share cars during the evening than they are during the morning and early afternoon.

Fig. 6 and Table 3 display results from Eq. 3, a regression discontinuity by hour of day. These figures are qualitatively similar to, but larger in magnitude than the difference in differences results, suggesting that the short-term impacts of TNC exit may be more pronounced than the medium-term impacts. Fig. 7 plots residuals from a pooled regression discontinuity performed on daytime traffic speeds in 2016 (when Uber and Lyft exited) and 2015 (where both companies operated year-round).

Table 1
Road segment summary statistics.

	mean	sd	min	max
Average Speed (mph)	24.74	9.61	2.17	95.04
Minutes Per Mile	2.99	1.79	0.63	27.69
Segment Length	0.72	0.57	0.06	3.80
Number samples	4.77	3.77	1.00	45.00
Number of Lanes	4.70	0.91	2.00	7.00

Summary statistics for traffic data along 79 road segments in Austin, TX. Speed data reflect the average travel time for Bluetooth devices that move from origin sensor to destination sensor during a given 15-min interval. As described in Section 4, data are also filtered for outliers. Traffic speed data were accessed through the City of Austin's OpenData Portal.

The 2015 regresion discontinuity estimates a null effect, offering evidence that the 2016 regression discontinuity results are not driven by seasonal changes in traffic patterns.

Table 4 shows the results from running versions of Eqs. 1 and 3, pooling across hours. The pooled difference in differences results suggest that on average, speeds increase by 0.026 minutes per mile ($p = 0.15$), or roughly 0.9% following TNC exit. Consistent with the hour-specific estimates, restricting the pooled DID analysis to daytime hours (7 a.m. to 7 p.m.) generates larger estimates of speed increases following TNC exit ($\beta = -0.068$, $p = 0.2$). This coefficient translates to a 2.3% increase in daytime traffic speeds. Fig. 4 displays the raw speed data for daytime traffic by week of year for my study window, and provides evidence of the absence of pre-trends. Fig. 5 plots an event study version of Eq. 1, where separate treatment effects are estimated for each week. Consistent with the growth of RideAustin and other Austin-area TNC alternatives through the second half of 2016, the event study shows the treatment effect decaying over time: 2016 traffic speeds are significantly lower than those in 2015 for 10 weeks following the exit of Uber and Lyft, but

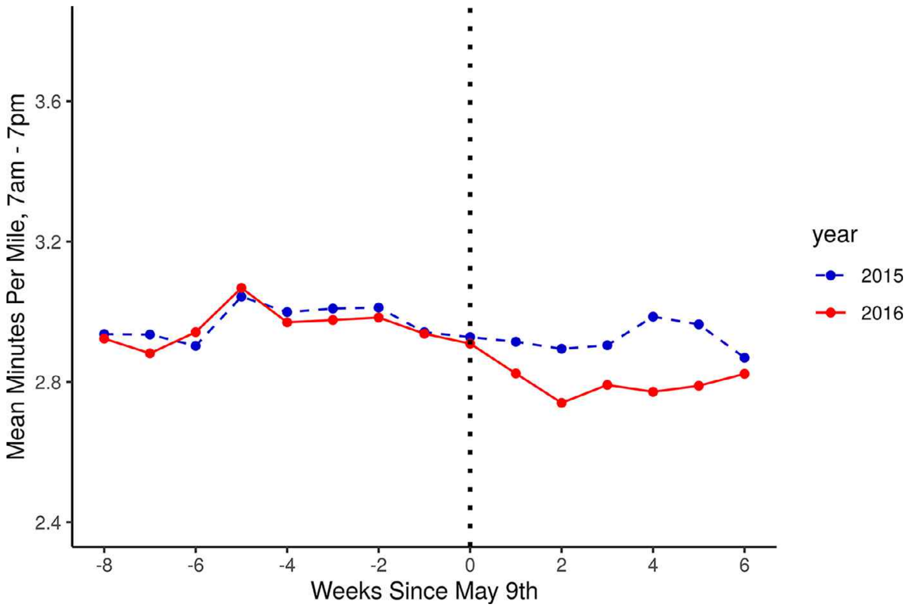

Fig. 4. Parallel Trends. *Notes:* This figure shows raw average speed (minutes per mile) between 7 a.m. and 7 p.m. over 79 road segments in Austin, TX, plotted by week of year for 2015 and 2016. Data were accessed through the City of Austin's Open Data Portal. The dotted line represents the week of May 9th, where Uber and Lyft ceased operation in Austin in 2016. Note that week zero is partially treated, as May 9th, 2016 was a Monday.

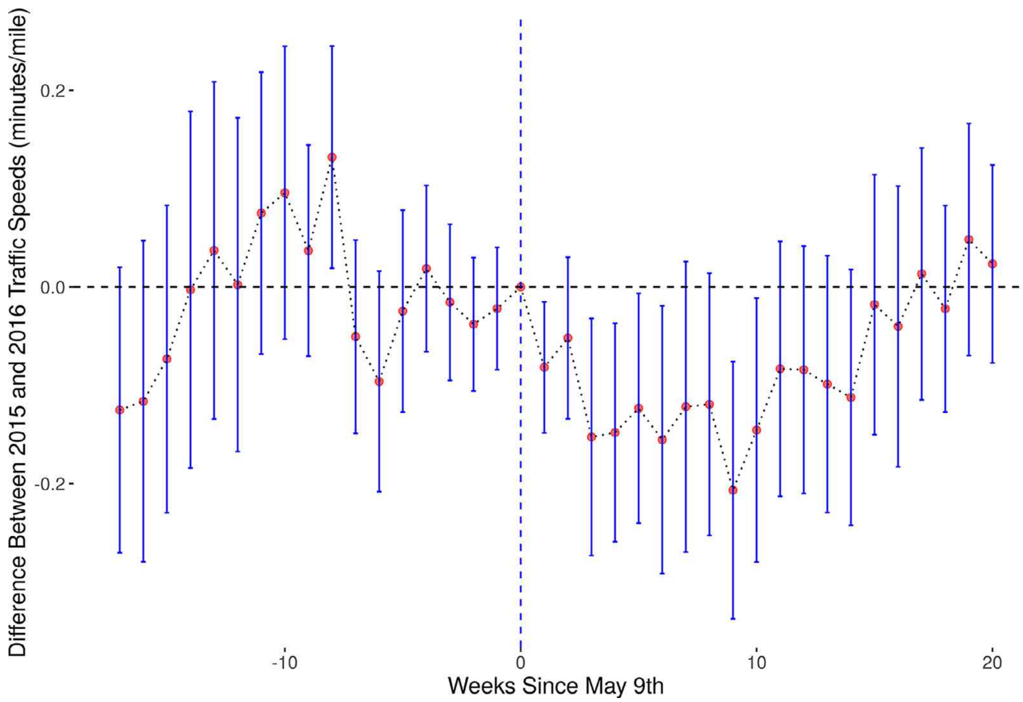

Fig. 5. Event Study. *Notes:* This figure plots results from a difference in differences regression with separate coefficients for 17 pre and 20 post-exit weeks. Points represent the estimated difference between 2015 and 2016 traffic speeds (in minutes per mile) relative to the difference in the week leading up to May 9th. Controls include day of week, holiday, SXSW, and segment fixed effects, as well as flexible controls for temperature and precipitation. Bars reflect 95% confidence intervals from standard errors clustered by segment.

by week 17, point estimates suggest that traffic speeds had returned to the baseline 2015–2016 difference.

As in the hour-specific estimates, the pooled RD estimates are larger in magnitude than are the DID results: Travel times decreased by 0.102 minutes per mile ($p = 0.01$) across all hours and by 0.134 ($p = 0.003$) minutes per mile for daytime hours. These coefficients correspond to travel time reductions of 3.4% and 4.5%, respectively.

As noted in Section 3, a number of ridesharing firms entered the market after the exit of Uber and Lyft. If alternative TNC activity was substantial during the study period, my results will be attenuated relative to the counterfactual of a TNC-free Austin. Although RideAustin's data provide some insight into the level of alternative TNC activity in Austin, it is unclear whether the growth of RideAustin in 2016 is representative of the growth of all alternative TNCs, or whether RideAustin grew by cutting into the market share of firms like Fasten and Fare, which arrived earlier. In Table 5, I report estimates of the impact of TNCs on traffic speeds in Austin under each of these two possible trajectories of TNC activity in 2016: In rows 1 and 3 (*RideAustin Data*), I assume 10,200 TNC trips per day during the pre-period (see Uber (2015)), and use RideAustin's time-series data—inflated by the reciprocal of its market share—to produce a time-varying measure of TNC activity following the failure of Proposition 1. This time series is plotted in red in

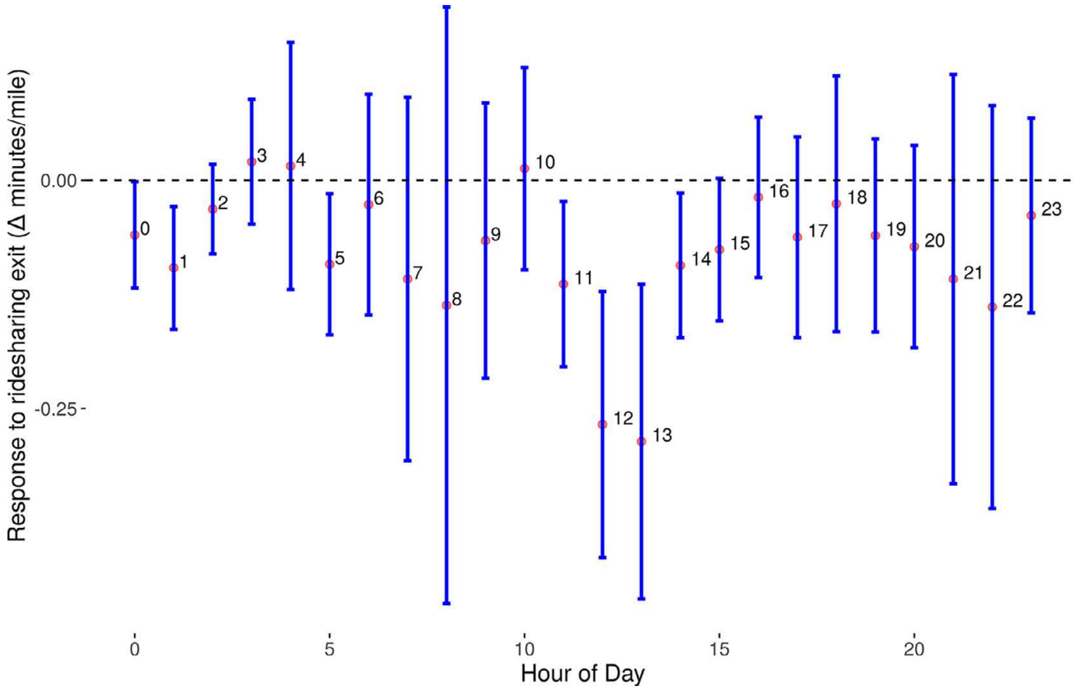

Fig. 6. Regression Discontinuity Results. *Notes:* Results from Eq. 3, a regression discontinuity performed on traffic speeds across 79 road segments in Austin, TX. The bandwidth is March 20th - August 1st of 2016, which (asymmetrically) spans the May 9th departure of Uber and Lyft. Points represent the estimated effect of TNC exit on traffic speeds by hour of day. A negative point indicates an estimated increase in traffic speed. Controls include day of week, holiday, and segment fixed effects, segment-specific second degree polynomials in days since May 9th, and flexible controls for temperature and precipitation. Bars reflect 95% confidence intervals from two-way standard errors clustered by segment-week. Traffic speed data were accessed through the City of Austin's Open Data Portal.

Table 2
Difference in differences results.

Hour of Day	β_h	se	p
0	−0.0606	0.0515	0.2588
1	−0.0525	0.0358	0.1640
2	**−0.0505**	0.0278	0.0908
3	−0.0121	0.0649	0.8547
4	0.0911	0.1006	0.3805
5	−0.0293	0.0920	0.7547
6	−0.0568	0.0996	0.5778
7	−0.1446	0.1169	0.2365
8	−0.0583	0.1331	0.6680
9	−0.0188	0.0954	0.8469
10	0.0318	0.0581	0.5921
11	**−0.1418**	0.0664	0.0508
12	**−0.1730**	0.0812	0.0512
13	−0.1555	0.0927	0.1157
14	−0.0176	0.0469	0.7125
15	−0.0238	0.0453	0.6083
16	0.0657	0.0683	0.3524
17	0.0016	0.0591	0.9790
18	0.0500	0.0862	0.5715
19	−0.0300	0.0638	0.6451
20	0.0004	0.1036	0.9966
21	0.0817	0.0703	0.2646
22	−0.0709	0.0852	0.4191
23	−0.0152	0.0328	0.6497
F-test			0.0000
N			966,301

Notes: Results from Eq. 1, a difference in differences comparing pre vs. post May 9th traffic speeds in 2015 to pre vs. post May 9th traffic speeds in 2016 (where both Uber and Lyft exited Austin). Controls include segment-specific linear in day trends, a precipitation dummy, day of week fixed effects, and year and post May 9th dummies. Standard errors are clustered by segment-week. β_h represent the estimated effect of TNC departure on traffic speeds (in minutes per mile) by hour of day. Bold coefficients are significant at the 10% level. The final row reports the p-value from a joint hypothesis test of $\beta_h = 0\ \forall h$.

Table 3
Regression discontinuity results.

Hour of Day	β_h	se	p
0	**−0.0598**	0.0297	0.0637
1	**−0.0959**	0.0343	0.0142
2	−0.0316	0.0251	0.2287
3	0.0202	0.0348	0.5711
4	0.0156	0.0690	0.8245
5	**−0.0920**	0.0395	0.0352
6	−0.0267	0.0618	0.6727
7	−0.1081	0.1016	0.3050
8	−0.1370	0.1666	0.4248
9	−0.0661	0.0769	0.4047
10	0.0128	0.0566	0.8249
11	**−0.1137**	0.0462	0.0275
12	**−0.2675**	0.0744	0.0029
13	**−0.2861**	0.0878	0.0057
14	**−0.0932**	0.0404	0.0368
15	**−0.0759**	0.0398	0.0775
16	−0.0187	0.0449	0.6833
17	−0.0623	0.0561	0.2856
18	−0.0257	0.0715	0.7246
19	−0.0605	0.0539	0.2808
20	−0.0727	0.0565	0.2188
21	−0.1082	0.1144	0.3601
22	−0.1388	0.1126	0.2379
23	-0.0385	0.0544	0.4910
F-test			0.0000
N			501,010

Notes: Results from Eq. 3, a regression discontinuity performed on traffic speeds across 79 road segments in Austin, TX. The bandwidth is March 20th - August 1st of 2016, which (asymmetrically) spans the May 9th departure of Uber and Lyft. Controls include hour of day, day of week, holiday, and segment fixed effects, segment-specific second degree polynomials in days since May 9th, and flexible controls for temperature and precipitation. Standard errors are clustered by segment-week. β_h represent the estimated effect of TNC departure on traffic speeds (in minutes per mile) by hour of day. Bold coefficients are significant at the 10% level. The final row reports the p-value from a joint hypothesis test of $\beta_h = 0\ \forall h$.

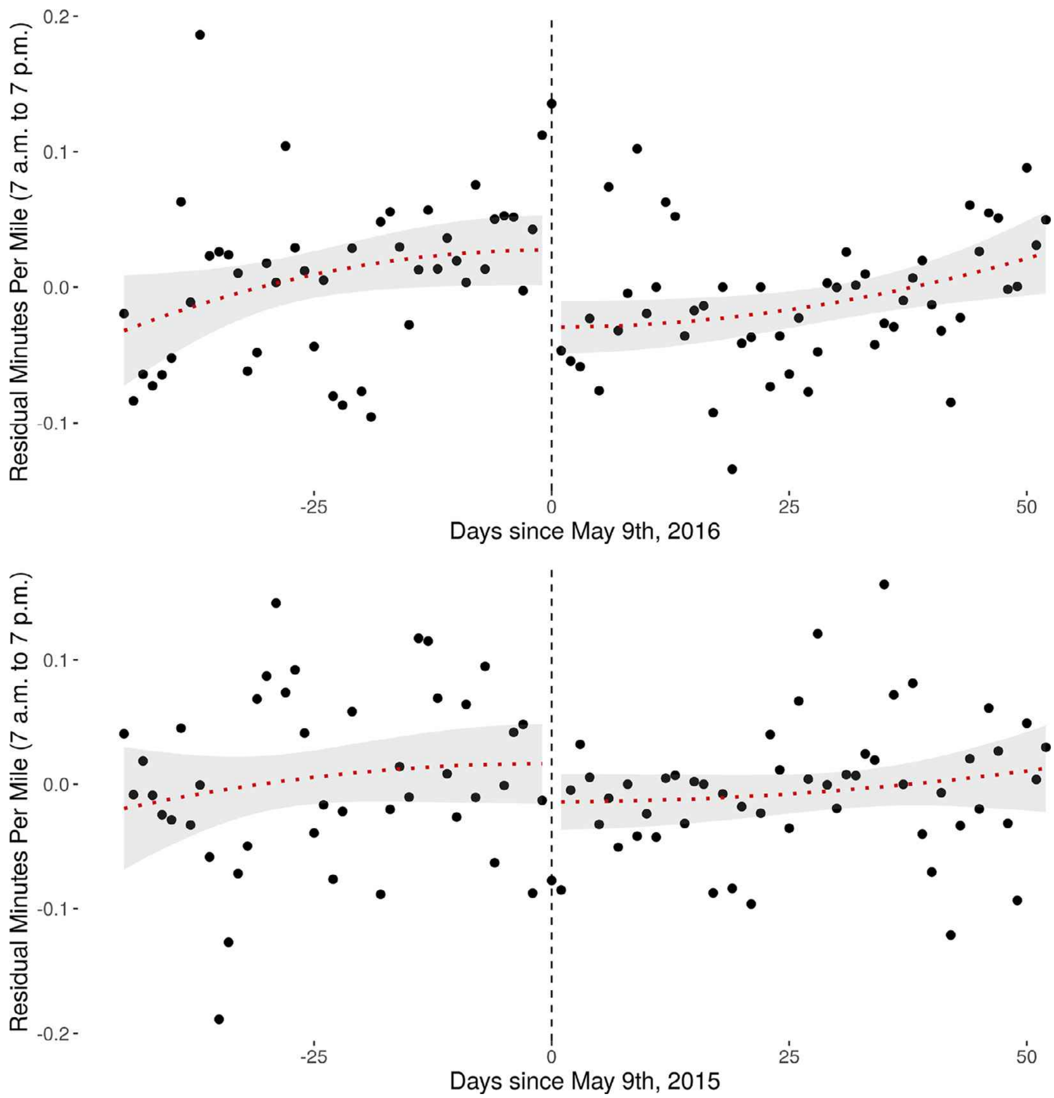

Fig. 7. Regression Discontinuity Residual Plots. *Notes:* This figure plots the daily mean residuals from a pooled version of Eq. 3, omitting the $\delta_y \eta_t$ (year*post) indicator. The dependent variable is minutes per mile, measured over 79 road segments in Austin, TX. The bandwidth is March 20th - August 1st of 2016 (or 2015), which (asymmetrically) spans the May 9th departure of Uber and Lyft. Controls include hour of day, day of week, holiday, and segment fixed effects, segment-specific second degree polynomials in days since May 9th, and flexible controls for precipitation and temperature. The dotted line represents a second degree polynomial in days since May 9th; the shaded region is the 95% confidence interval. Traffic speed data were accessed through the City of Austin's Open Data Portal.

Fig. 2. Under this assumption, neither the fuzzy DID nor the fuzzy RD specification differs significantly from the results in Table 4. In rows 2 and 4 (*Hampshire Data*), I again use 10,200 TNC trips per day for the pre-exit figure, but then assume that a constant 4180 (41% of 10,200) TNC trips per day during are completed during the treatment period. This figure reflects results from a November 2016 survey conducted by Hampshire et al. (2017), where 41% of survey respondents reported that they completed an Uber or Lyft reference trip using another TNC following the failure of Proposition 1. This trajectory assumes that RideAsutin's growth is not representative of the alternative TNC market, and instead entirely reflects RideAustin winning over customers from other already-established Uber and Lyft alternatives.

Intuitively, the results using the *Hampshire Data* assumption are roughly 1.7 times larger than the estimates from Table 4. This offers a useful bound for this exercise investigating attenuation. If the RideAustin data is even partially representative of the growth of alternative ridesharing companies in Austin, then TNC activity in May–August of 2016 was lower than the 41% replacement reported by Hampshire et al. (2018) in November. I therefore view row 4—which implies a 7.6% reduction in travel speeds when moving from zero TNC activity to full TNC activity—as an upper bound for the congestion impacts of TNCs in Austin.

My estimates of changes in traffic speeds together with data on the number of total TNC and non-TNC vehicle trips in Austin allow me to estimate the implied elasticity of congestion with respect to TNC volumes. According to the 2017 NHTS, Austin-area households take an average of 3.6 vehicle trips a day. Austin's 37,000 households, then, generate roughly 1.35 million vehicle trips per day. The available data on Uber and Lyft suggest that the two services together completed roughly 10,200 trips per day prior to their 2016 exit from Austin (see Fig. 2). Multiplying this figure by a factor of two to reflect the capacity factor estimated by Judd and Krueger (2016) suggests that Uber and Lyft together accounted for 1.5% of Austin-area vehicle trips prior to the failure of Proposition 1. My pooled estimates of the impact of TNCs on Austin-area congestion therefore imply congestion elasticities with respect to TNC volume of between 0.6 and 2.3. My preferred specification, which suggests a 2.3% increase in daytime traffic speeds following the exit of Uber and Lyft, corresponds to an elasticity 1.5. These estimates lie on

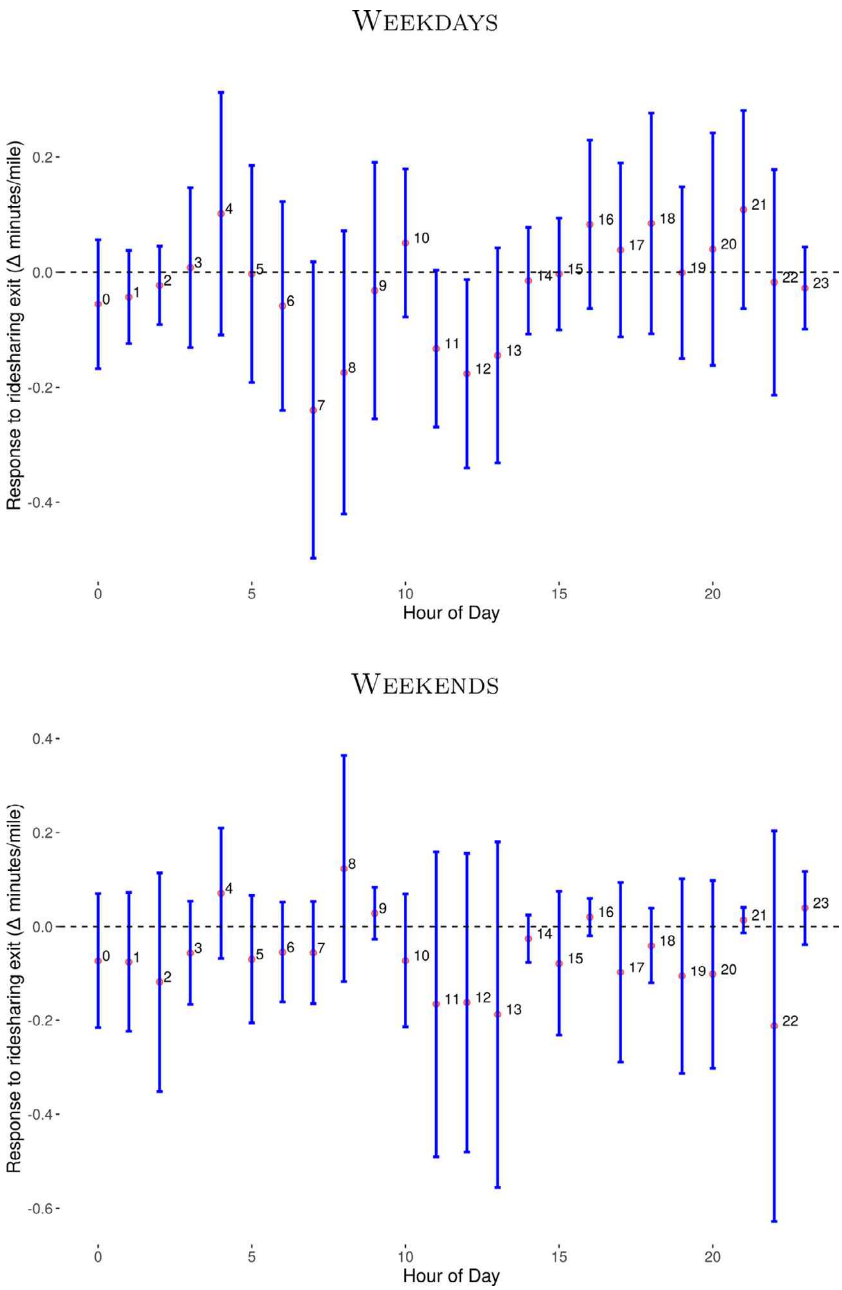

Fig. 8. Weekday vs. Weekend Effects. *Notes:* This figure plots difference in differences estimates (Eq. 1) of the impact of Uber and Lyft's exit on travel speeds separately for weekdays and weekends. Points represent the estimated effect of TNC departure on traffic speeds (in minutes per mile) by hour of day. Controls include day of week, holiday, and segment fixed effects, segment-specific linear trends in days since May 9th, and flexible controls for temperature and precipitation. Bars reflect 95% confidence intervals from two-way standard errors clustered by segment-week. Traffic speed data were accessed through the City of Austin's Open Data Portal.

the lower end of the range of congestion elasticities from existing studies. Anderson et al. (2016), for example, estimate a congestion elasticity of 2.7 in Beijing; findings from Leape (2006) imply an elasticity of 2.5 in London, and results from Eliasson (2009) imply an elasticity of 1.5 in Stockholm. The relatively small congestion elasticity implied by my estimates may reflect the lower levels of congestion in Austin relative to cities like London and Beijing, or the offsetting effect of trips saved by the 'ridesharing effect' of TNCs.

In Appendix D, I investigate the robustness of the results presented in this section. Fig. D.1 plots both the difference in differences and regression discontinuity results for bandwidths ranging from 20 to 70 days around May 9th. The conclusion that daytime traffic speeds increase following the exit of Uber and Lyft holds across bandwidth choices. To test the likelihood that the regression discontinuity estimates presented above are the result of a contemporaneous shock to Austin-area traffic speeds, I compare my estimates to coefficients from 134 regression discontinuities using placebo exit dates. The results of this exercise are shown in Fig. D.2. 5 of the 134 placebo coefficients (4%) fall below the estimate using the true exit date, suggesting that it is empirically unlikely that my RD estimates are the result of an unobserved Austin-area transit shock.

6.2. Heterogeneity and external validity

Results from Eq. 2, which allows for segment-specific congestion responses, are plotted in Figs. 9 and 10. Two themes emerge. First, there is no clear spatial pattern in congestion impacts: I estimate negative and positive travel time impacts both for segments in the city center and for outlying roads. Second, Fig. 9 shows that segments that experienced exceptionally large changes in traffic speed were characterized by exceptionally high levels of pre-period traffic congestion, suggesting construction or other segment-specific shocks may explain these estimates. Absent these outliers, the segment-specific effects exhibit relatively low variance. An important avenue for future work would be to investigate

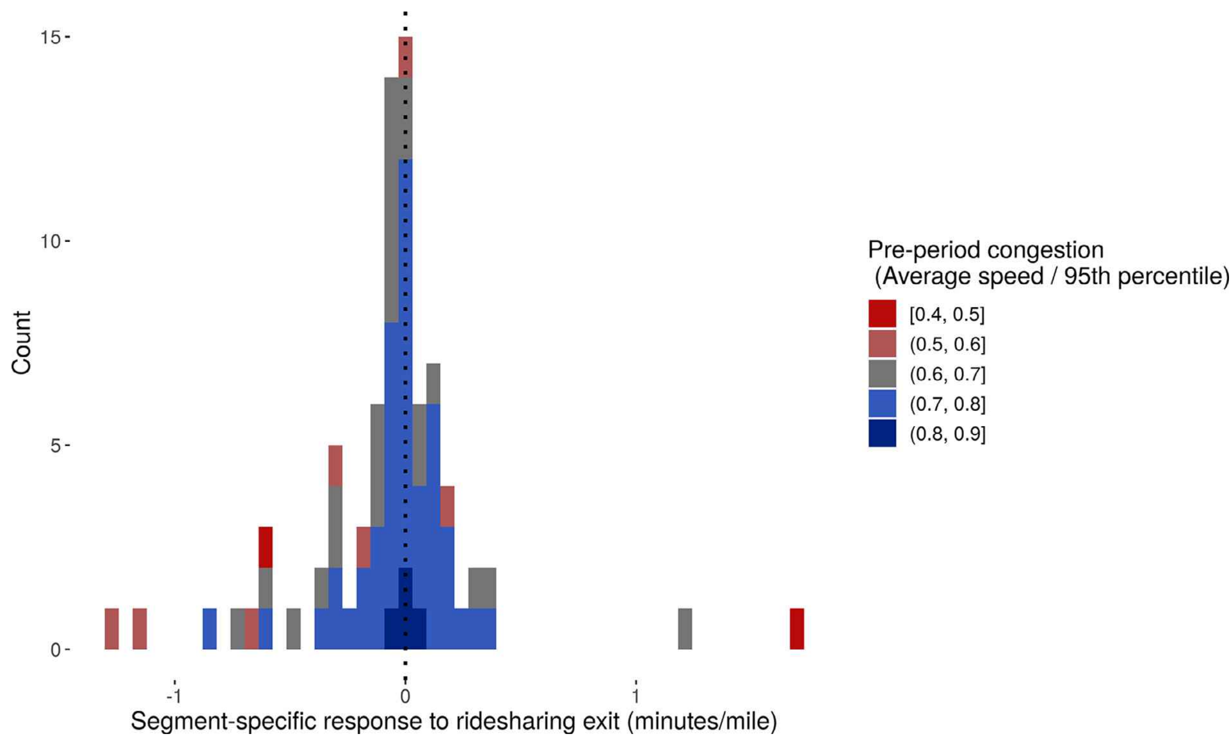

Fig. 9. Distribution of Segment-Specific Responses. *Notes:* Results from Eq. 2, a difference in differences comparing pre vs. post May 9th traffic speeds in 2015 (where both Uber and Lyft operated in Austin) to pre vs. post May 9th traffic speeds in 2016 (where both Uber and Lyft exited Austin), allowing for segment-specific congestion responses. Bars represent the number of segments with idiosyncratic changes in traffic speeds falling withing a given bin. Cells are colored by the pre May 9th 2016 congestion level, as measured by the ratio of average speed to the 95th percentile of speed. Traffic speed data were accessed through the City of Austin's Open Data Portal.

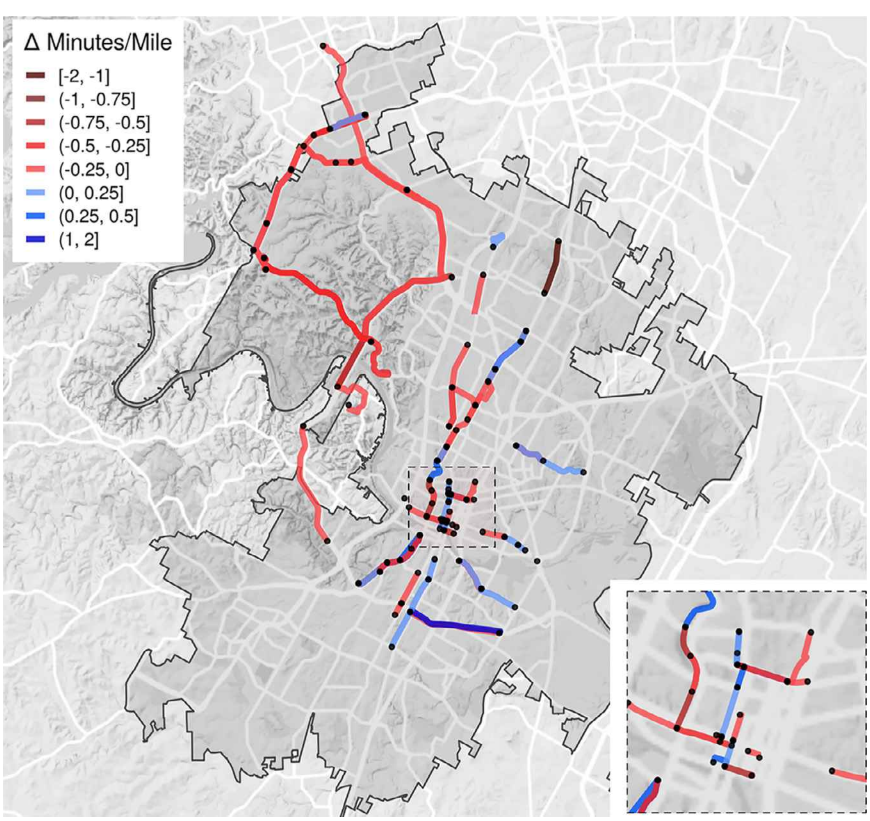

Fig. 10. Segment-Specific Responses. *Notes:* Results from Eq. 2, a difference in differences comparing pre vs. post May 9th traffic speeds in 2015 (where both Uber and Lyft operated in Austin) to pre vs. post May 9th traffic speeds in 2016 (where both Uber and Lyft exited Austin), allowing for segment-specific congestion responses. Paths represent Google Maps recommended driving directions between endpoints of a given segment, colored by the sign and magnitude of the estimated segment-specific change in traffic speed. The black line is the Austin city limit. Traffic speed data were accessed through the City of Austin's OpenData Portal.

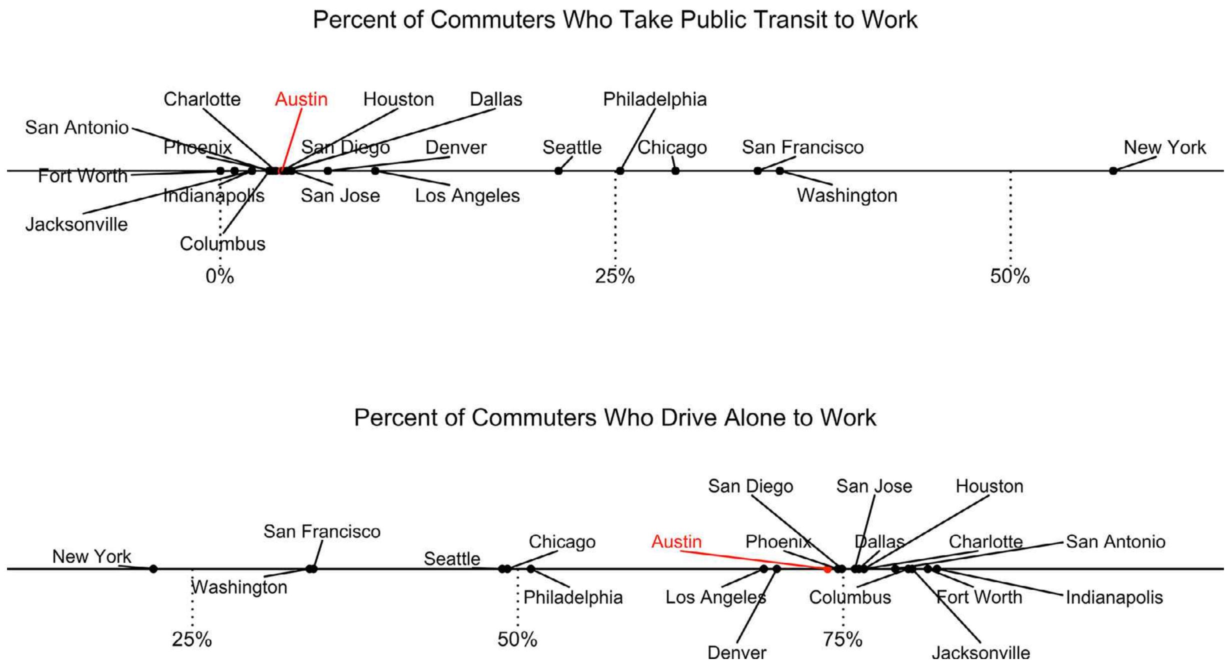

Fig. 11. External Validity. *Notes:* This figure depicts vehicle and public transit use across the 20 largest US cities for the year 2017. Data for both subfigures come from the U.S. Census Bureaus 2017 American Community Survey. Similarities in commuting behavior between Austin and other "Sun Belt" cities suggests that the findings in this paper may be most applicable to this group of metros.

Table 4
Pooled estimates.

	β (Δ minutes/mile)	*se*	*p*	Implied annual cost ($)
Difference in Differences (All hours)	−0.0261	0.0170	0.1479	−33,096,514
Difference in Differences (7 a.m. - 7 p.m.)	−0.0684	0.0529	0.2004	−63,985,000
Regression Discontinuity (All Hours)	−0.1015	0.0353	0.0052	−129,010,337
Regression Discontinuity (7 a.m. - 7 p.m.)	−0.1335	0.0433	0.0028	−124,930,327

Notes: The first two rows display results from a variation of Eq. 1, a difference in differences specification that estimates the pooled impact of TNC exit on traffic speeds across hours of day. β represent the estimated effect of TNC departure on traffic speeds, measured in minutes per mile. Controls include segment-specific linear in day trends, controls for precipitation, day of week fixed effects, hour of day fixed effects, and year and post-May 9th dummies. Standard errors are clustered by segment-week. Row 1 shows the results of this regression using speed data on all hours, and column 2 shows results restricted to 7 a.m. to 7 p.m. The final column displays annual costs implied by multiplying β by annual Ausin-area VMT, and then by $15.40, which is 50% of the average per-worker wage rate for Austin households, according to the 2017 NHTS. Traffic data were accessed through the City of Austin's OpenData Portal. Rows 3 and 4 repeat this exercise for Eq. 3.

whether these outliers indeed represent extreme congestion reductions from TNC operation in select locations, or whether they can be explained by data absent from this setting.

The external validity of the results presented in this paper hinges on whether Austin is representative of other metropolitan areas in terms of commuter preferences and the substitutability of transit options. To determine which cities have transit systems that resemble Austin's, Fig. 11 depicts how public transit use and personal vehicle travel vary across the 20 largest metro areas in the US. Commuters in the majority of large American cities (especially those located in the 'Sun Belt') exhibit mode choices similar to those in Austin, where commuters heavily favor solo commutes in personal vehicles. In cities with extensive public transit systems (e.g. New York, Washington, San Francisco), however, commuter choices are quite different than they are in Austin. This suggests caution when applying the results described in this paper to address policy questions in these metro areas.

6.3. Equilibrium response

It is worth discussing the extent to which the brief disruption of TNC activity in Austin provides insights into the equilibrium differences in congestion levels between a city with and a city without TNCs. In addition to the first-order changes in traffic flow caused by the absence of TNCs, the full equilibrium response to the exit of TNCs would reflect a combination of short and long-run adjustments made by road users. More specifically, road users will spatially re-optimize in response to differential speed changes, and city residents may change long run by vehicle purchase or sorting decisions.

Because spatial re-optimization over route choices likely occurs in the short-run, and I use a large sample of segments covering different types of roadways, my estimates likely reflect this spatial substitution. My results do not, however, reflect long-term adjustments: The above analysis compares traffic speeds in Austin with and without TNCs, holding fixed decisions on car purchasing behavior and locational sorting. Some of the reduction in congestion that I measure likely comes from individuals who, prior to 2016, chose not to purchase a vehicle because they had access to ridesharing. In the long run, the actions of these marginal car owners would erode the traffic improvements that resulted from the exit of Uber and Lyft.

6.4. Congestion costs

Armed with estimates of hour-specific changes in travel times, I calculate the external congestion cost associated with TNC operation as

Table 5
Fuzzy RD and Fuzzy DD.

	β (Δ minutes/mile)	*se*	*p*
Fuzzy DID, RideAustin Data	0.0721	0.0573	0.2084
Fuzzy DID, Hampshire Data	0.0978	0.0237	0.0010
Fuzzy RD, RideAustin Data	0.1324	0.0133	0.0000
Fuzzy RD, Hampshire Data	0.2263	0.0227	0.0000

Notes: The first two rows display results from a fuzzy difference in differences specification that estimates the impact of TNC activity on Austin traffic speeds between 7 a.m. and 7 p.m. The regression coefficient represents the change in travel times (in minutes per mile) resulting from a change from full TNC operation [1] to no TNC operation [0]. The first row (*RideAustin Data*) uses RideAustin's trip-level data together with estimates of RideAustin's market share to construct a measure of daily TNC activity. The second row assumes that the level of TNC activity during the period following the exit of Uber and Lyft was a constant 41% of pre-exit levels. This assumtion is based on a November 2016 survey conducted by Hampshire et al. (2018). Rows three and four report results from fuzzy regression discontinuity designs that estimate the impact of TNC activity on Austin traffic speeds between 7 a.m. and 7 p.m., using the same TNC activity assumptions.

follows:

$$\Delta\text{congestion cost} = \sum_h \Delta \text{ minutes per mile}_h * \text{miles driven}_h * \text{value of time}_h \tag{4}$$

Δ *minutes per mile* are the coefficients, by hour of day, h, estimated above. To estimate *miles driven*$_h$, I use periodic traffic counts to estimate the share of VMT by hour of day in Austin, and multiply these shares by estimates of daily VMT provided by the Texas Department of Transportation. Note that this operation assumes that my estimates represent an average effect for all VMT within Austin City limits. Table 6 provides evidence in support of this assumption: According to data maintained by the Texas Department of Transportation, roads included in my preferred specification resemble those not included in terms of congestion, VMT, and speed. Finally, I calculate Austin-specific hourly value of travel time (VOT) estimates (*value of time*$_h$, above) using data from Austin's MoPac freeway (see Appendix A). I also present results using a VOT heuristic from Small and Verhoef (2007): 50% of the wage rate. I estimate the wages of Austin-area drivers by calculating the per-worker income for car-commuting Austin households in the 2017 National Household Travel Survey (NHTS). Standard errors for all congestion cost estimates follow formulas developed by Goodman (1960).

I report the results of this exercise in Table 7. Both rows reflect changes in travel times estimated in my preferred specification (Eq. 1), a DID across years. Using time-varying (MoPac) and uniform (NHTS) VOT estimates, I calculate the daily congestion costs associated with Uber and Lyft activity at $92,071 and $127,983, respectively. These estimates correspond to annual costs of $33 million ($p$=0.181) and $46 million ($p$=0.049). Disaggregating this sum by weekend and weekday effects (Table 9) produces slightly larger figures annual cost figures: $39 million ($p$=0.041) and $52 million ($p$=0.003). In Appendix D, I report results from this exercise using a regression where observations are weighted by the number of Bluetooth devices recorded in each 15-min window. The rationale for this specification is to investigate whether the above results are biased when segments with differing traffic flows are implicitly given the same weight in determining regression coefficients. The cost estimates from this weighted regression are similar to the non-weighted results and imply an annual congestion cost associated with Austin-area TNC activity of $54 million. Note that each of these aggregate cost measures relies on the assumption that VOT is uniform across the city. While misattributing VOT estimates by location may bias these estimates, there are two reasons why this bias is likely small: First, in a recent investigation of the heterogeneity of urban VOT, Buchholz et al. (2020) find that the majority of the variation in VOT is across individuals rather than across locations within a city. Second, the lack of a spatial gradient in the segment-level congestion estimates (see Fig. 10) means that a cost calculation using a modest VOT gradient between the city and the suburbs (as in Fig. 8 of Buchholz et al. (2020)) would produce similar aggregate welfare measures as those reported above in Tables 4 and 7, so long as the VOT estimates I use reflect the rough spatial average of the Austin metro area.

Table 7
Congestion cost estimates.

	Daily cost ($)	*se*	*p*	annual cost ($)
Time-varying VOT	−92,071	97,547	0.1806	−33,605,827
Uniform $15.40 VOT	−127,983	72,156	0.0489	−46,713,725

Notes: Estimates of the travel-time congestion costs of TNC operation in Austin, TX. The first row displays the result of the exercise described in Eq. 4, which matches hour-specific changes in travel time to hour-specific willingness to pay estimates (detailed in Appendix A) and hour-specific traffic volume measurements. The second row uses a VOT of $15.40, which is 50% of the average per-worker wage rate for car-commuting Austin households, according to the 2017 NHTS. Standard errors for cost figures are calculated following Goodman (1960). See Table 8 for hour-by-hour cost estimates.

Several outside studies provide valuable context when interpreting these congestion cost estimates. First, according to the Inrix Global Scorecard, the aggregate 2017 travel time cost in Austin was $2.8 billion, $810 million of which was attributed to traffic slowdowns (Inrix, 2017). Back of the envelope calculations using my estimates therefore suggest that Uber and Lyft together accounted for 1–2% of all travel time costs in Austin, and 4–6% of congestion costs. Second, estimates of consumer surplus associated with TNCs offer a useful benchmark for policymakers weighing the benefits of TNC operation against the costs. The results from Cohen et al. (2016) allow me to produce two

Table 6
Austin metro validity.

	Not Sampled	Sampled	p
Annual Delay per Mile (person-hours)	118,372.98	101,368.11	0.58
Texas Congestion Index	1.36	1.39	0.61
Peak Period Average Speed	33.22	29.52	0.16
Freeflow Speed	41.20	37.97	0.28
Average Daily VMT	204,488.84	166,805.50	0.46
Peak Period Annual Hours of Delay (person-hours)	377,230.86	241,909.86	0.34

Notes: This table uses data maintained by the Texas Department of Transportation (TXDoT) to compare observable characteristics of Austin-area roads that do not appear in my sample (column 1) to those that do (column 2). Column (3) reports p-value of the corresponding *t*-test for a difference in means. The data cover 86 road segments in Austin, 35 of which overlap with the 79 Bluetooth segments used in the above analysis. Note that roads sections are coded as sampled even if the Bluetooth segment does not cover the entire corresponding TXDoT road segment.

Table 8
Summary statistics for components of cost calculation.

Hour of Day	β	*se*	VOT	sd	VMT	sd	Cost/hour (MoPac VOT)	Cost/hour (NHTS VOT)
0	−0.06	0.04	8.66	3.05	201,112	1667	−1759	−1759
1	−0.05	0.04	8.4	2.96	125,170	825	−921	−921
2	−0.05	0.04	8.29	2.91	95,873	566	−669	−669
3	−0.01	0.05	8.3	2.92	64,993	265	−109	−109
4	0.09	0.05	8.32	3	45,746	73	578	578
5	−0.03	0.06	9.69	3.57	79,183	89	−375	−375
6	−0.06	0.07	12.67	8.28	216,472	1181	−2595	−2595
7	−0.14	0.08	25.64	30.81	510,299	7646	−31,539	−31,539
8	−0.06	0.12	28.78	37.99	680,662	11,764	−19,034	−19,034
9	−0.02	0.08	14.67	17.84	729,092	8322	−3344	−3344
10	0.03	0.05	10.16	4.08	726,981	3464	3919	3919
11	−0.14	0.04	10.1	3.75	786,248	4234	−18,771	−18,771
12	−0.17	0.04	10.05	4.83	910,484	4754	−26,387	−26,387
13	−0.16	0.04	10.41	6.06	943,347	5733	−25,444	−25,444
14	−0.02	0.04	11.75	9.67	919,202	6131	−3176	−3176
15	−0.02	0.04	20.79	19.93	937,669	5146	−7716	−7716
16	0.07	0.04	29	26.84	950,548	4552	30,189	30,189
17	0	0.04	24.5	22.83	953,432	4493	618	618
18	0.05	0.04	23.27	24.29	940,423	4676	18,220	18,220
19	−0.03	0.04	11.55	9.64	834,215	3861	−4822	−4822
20	0	0.04	9.76	3.68	644,943	3076	47	47
21	0.08	0.04	8.42	3.4	529,941	2619	6082	6082
22	−0.07	0.04	8.67	3.36	429,036	2491	−4397	−4397
23	−0.02	0.04	8.57	3.06	306,256	2127	−665	−665
Daily Cost							**−92,071**	**−127,983**

Notes: This table displays the components of the aggregate cost calculations in Table 7. The first two columns reproduce the difference in differences results from Table 2. Columns 3 and 4 show the mean and standard deviation of the value of travel time (VOT) estimates from the MoPac Freeway (see Appendix A). Columns 5 and 6 report the mean and standard deviation of hourly VMT. These estimates are generated by multiplying the hourly share of total Austin Traffic (as per Austin traffic count data) by estimates of aggregate Austin-area VMT (as per the Texas Department of Transportation). The penultimate column reports the hourly costs estimates using the VOT from column 3; the final column reports hourly cost estimates applying a constant $15.40 per hour.

Table 9
Congestion costs disaggregating weekdays and weekends.

	Annual cost ($)	*se*	*p*
Time-varying VOT	−39,989,479	21,348,072	0.0410
Uniform $15.40 VOT	−52,084,544	15,784,030	0.0026

Notes: Estimates of the travel-time congestion costs of TNC operation in Austin, TX, using separate estimates for weekdays and weekends. The first row follows Eq. 4, and uses separate hourly Value of Travel Time (VOT) schedules for weekdays and weekends derived from MoPac Data. The second row applies a uniform VOT time across all hours. Following the heuristic provided by Small and Verhoef (2007), this VOT is 50% of the per-worker wage rate for Austin-area car commuting households, as per the 2017 NHTS. Standard errors for cost figures are calculated as described in Section 6.

back of the envelope estimates of consumer surplus in Austin. First, Cohen et al. (2016) conclude that in the four cities they examine, $1.57 of consumer surplus are generated for every dollar spent on TNCs. Uber reported that in 2015, its drivers grossed $27 million in the Austin area (Uber, 2015). According to Uber's S1 filing, 83% of the payments to the Uber app went to drivers in 2016, suggesting that Austin consumers pay roughly $32.5 million to Uber annually. Inflating this figure by the multiplier from Cohen et al. (2016), and assuming an equal expenditure-consumer surplus ratio between Lyft and Uber implies a total TNC-related consumer surplus for the city of Austin of $72.9 million annually. As a second way of estimating consumer surplus in Austin, I rescale the Cohen et al. (2016) estimate of national consumer surplus from Uber ($6.8 billion). Multiplying this figure by Austin's share of the US urban population (0.49%) yields an estimate of the 2015 Austin-area consumer surplus of $33 million annually. Again inflating this figure to account for Lyft's market share suggests a TNC-related consumer surplus for the city of Austin of $47 million. These exercises therefore suggest that the congestion costs resulting from Uber and Lyft activity in Austin are similar in magnitude to the consumer surplus benefits provided by these companies.

7. Conclusion

Using a natural experiment in Austin, TX, I study whether transportation network companies impact traffic speeds. I estimate that TNCs are responsible for a 2.3% increase in Austin-area travel times between 7 a.m. and 7 p.m. This figure masks important heterogeneity, with the largest TNC-related slowdowns occurring between 11 a.m. and 2 p.m. By matching setting-specific changes in traffic speeds to hour-specific estimates of the value of travel time, I find that Austinites would be willing to pay $33 to $52 million annually to avoid the slowdowns induced by TNC activity. Back of the envelope calculations using estimates of TNC consumer surplus from Cohen et al. (2016) suggest that the cost of TNC-related congestion in Austin is similar in magnitude to the consumer surplus generated by these companies.

These results have important implications for urban transportation policy. While a comprehensive congestion tax is the most efficient response to congestion externalities, charging all road users the social marginal cost of their actions is both technologically and politically challenging (King et al., 2007). Given these difficulties, a natural question is whether policies that target a related good (TNCs) would improve welfare, and if so, how such a policy would perform relative to the first-best.

My results suggest that quantity restrictions—like those imposed by New York in 2016—are unlikely to produce substantial welfare benefits, as the congestion benefits from restricting TNC activity are roughly offset by lost consumer surplus. Estimates of the congestion costs associated with TNC activity also speak to the efficacy of TNC taxation as a means of addressing traffic externalities. Models of imperfect Pigouvian taxation show that the welfare gains of a second-best tax relative to a first-best tax are a function of how well second best taxes target heterogeneous externalities, as well as the correlation between demand

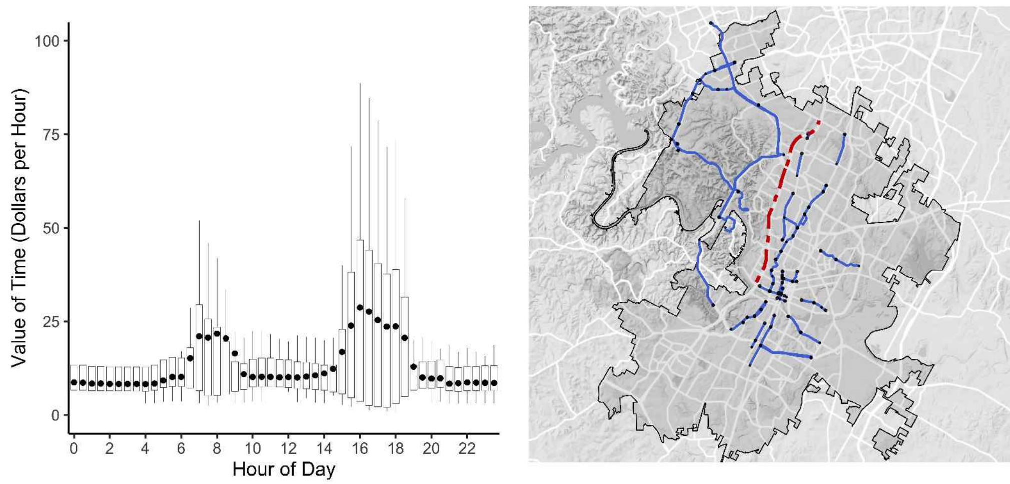

Fig. A1. Revealed Preference Value of Travel Time Estimates. *Left Pane:* Willingness to pay for travel time reductions in Austin, TX. Constructed using 2017 data from the MoPac variable-toll lane (Highway Loop 1), provided by the Central Texas Regional Mobility Authority. Dots represent means of observed equilibrium prices divided by expected time savings by hour of day.
Right Pane: MoPac expressway (red dashed) plotted with the 79 road segments used in estimation of the changes in travel speeds (blue). Nodes represent Bluetooth sensors, and the black line is the Austin City limit.

elasticities and idiosyncratic externalities (Knittel and Sandler (2018) and Diamond (1973)). Taxing ridesharing, then, will perform well as a second-best congestion pricing scheme if a) ridesharing activity is well-correlated with congestion, and b) if demand for ridesharing is elastic relative to the demand for other congesting trips. The results from this natural experiment suggest that TNC activity is not highly correlated with urban congestion externalities: back of the envelope calculations suggest TNCs are responsible for a small fraction of the congestion costs in Austin, and TNC-related externalities often occur at times where the value of travel time is relatively low. Taken together with results from Cohen et al. (2016), who characterize TNC trips as inelastic, uniform TNC taxation is unlikely to produce welfare gains that approach those realized under a first-best congestion price. For policymakers politically constrained to target congestion taxes at TNCs, the significant temporal and spatial variation in congestion impacts estimated in this paper suggests that there may be substantial improvements in policy efficiency from including spatial and temporal variation in pricing. TNC congestion fees levied in New York and San Francisco, for example, may be made more efficient by adding variation to the tax rate by time of day or origin/destination zone.

These results also pose several questions that may inform future research. First, identifying the drivers of the spatial and temporal heterogeneity in the congestion impacts of TNCs would be useful to city policymakers attempting to better target TNC-based congestion policies. Second, replicating this type of analysis in other settings with similar identification opportunities would provide a valuable test of external validity, especially in cities with markedly different commuting and public transit landscapes. Finally, the fact that speeds slow in response to TNC activity suggests TNCs add vehicles to the road. In other words, my results suggest that the ride-induction effect dominates the ride-sharing effect. This conclusion will be important to test in other cities, as the impact of TNCs on VMT is an important uncertainty in the prediction of transportation sector emissions.

Appendix A. Revealed preference VOT estimates

The MoPac (Texas State Highway 1) is a north-south route in Austin. Starting in November 2017, the Central Texas Regional Mobility Authority opened an 11-mile variable-price express lane on the MoPac (see Fig. A.1). The price of using this lane adjusts in order to keep the express lane moving at free-flow speeds: tolls increase when the express lane is busy and decrease when it is underused. Toll rates are posted at the northbound and southbound entrances.

Using 30-min resolution data provided by the Central Texas Regional Mobility Authority on MoPac prices and average travel times on the tolled and non-tolled lanes, I recover time-varying estimates of the value of travel time. Commuters entering the MoPac see the toll price, but not the difference in travel times between the tolled and non-tolled lanes. I therefore produce value of travel time estimates by dividing the observed toll price on a given date and time by the *expected* travel time savings for that time of day. To calculate the expected time savings, I take the average time difference between tolled and non-tolled lanes by half hour of day. For example, if the average difference between tolled and non-tolled lanes is 4 min between 9:00 a.m. and 9:30 p.m., and I observe an average price between 9:00 and 9:30 on a given day of $1, then the implied value of travel time for that half-hour block on that day is $15 per hour. I then aggregate these observations across days to recover an average value of travel time for each half hour of day. Note that because the tolled lane is always faster than the non-tolled lane *in expectation*, this method rationalizes driver's use of the toll lane when ex-post travel times are equal between tolled and non-tolled lanes.[2]

These estimates are summarized in Fig. A.1, and are broadly consistent with value of travel time estimates from related settings (Small and Verhoef, 2007). Importantly, however, VOT peaks during morning and evening rush hour periods, possibly reflecting different commuters, or heterogeneity in the value of time based on trip purpose. Note that the motivation for using MoPac data is to study the convolution between the value of travel time and estimated congestion impacts, not to produce novel estimates of the value of travel time.

Appendix B. Threats to identification from other modes of transportation

In addition to Bluetooth sensors, the city of Austin maintains pneumatic sensors which take periodic measurements of traffic speeds. Pneumatic sensors are stretched across traffic lanes, and therefore will not be influenced by pedestrian activities. Additionally, pneumatic sensors classify observations by axel length, meaning activity from bicycles

[2] Note that the MoPac has a price floor of $0.25; 55% of the observations (largely in off-peak hours) are at the price floor. VOT estimates excluding these observations produce similar results.

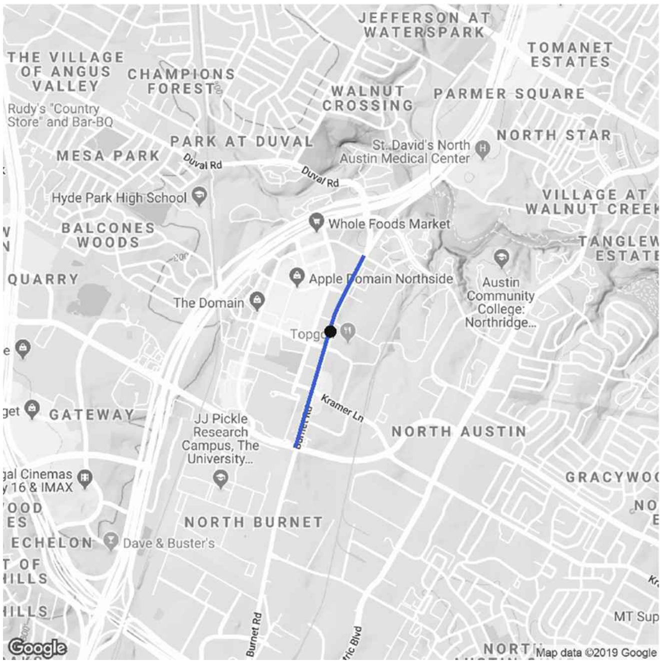

Fig. B1. Segment-sensor matching. *Notes:* An example of a Bluetooth segment (blue path) matched to a pneumatic sensor (black dot). Segments were matched to sensors by both location and direction of travel.

Table B1
Tests for Bias.

	Point estimate	*se*	*p*
β_1	0.0366	0.0221	0.0979
β_2	–0.0094	0.0222	0.6727

Notes: Results from Eq. 5, which tests whether the relationship between pneumatic speed measurements ($s_{i,t}$) and Bluetooth speed measurements ($y_{i,t}$) changes after the exit of TNCs in Austin. In addition to the variables listed, this regression includes segment and month of year fixed effects.

will not be reflected in vehicle speed measurements. I identify 39 instances where Bluetooth road segments overlie pneumatic sensors (see Fig. B.1), and test whether the relationship between the two speed measures changes significantly after the Proposition 1 vote.

For other modes of transportation to bias my estimates, speed must be mismeasured on Bluetooth segments, and that mismeasurement must be correlated with the treatment. To test for this bias, I perform the following regression:

$$y_{i,t} = \alpha + \beta_1 s_{i,t} + \beta_2 D_t \cdot s_{i,t} + \gamma_1 \boldsymbol{\phi}_t + \gamma_2 \boldsymbol{\theta}_i + \epsilon_{i,t} \tag{5}$$

Where $y_{i,t}$ is the Bluetooth speed measurement on segment i and time t, and $s_{i,t}$ is the pneumatic road segment speed measurement on segment i and time t. D_t is a treatment dummy, which equals one for days after May 9$^{\text{th}}$, 2016. $\boldsymbol{\phi}_t$ is a set of month of year fixed effects, and $\boldsymbol{\theta}_i$ is a set of segment fixed effects. I test the null hypothesis $\beta_2 = 0$, and interpret a significant result as evidence of bias in the Bluetooth speed masurements. Results from Eq. 5 are displayed in Table B.1.

Appendix C. Threats to identification from TNC driving speeds

If TNC vehicles drive significantly slower or faster than the average non-TNC vehicle in a way that remains after filtering, the above empirical strategy will be biased. To test for this potential bias, I match RideAustin trips to segments, allowing me to test the null hypothesis that TNC vehicles drive at the same speeds as the average mix of vehicles.

Table C1
TNC vehicle speeds.

Δ Minutes per Mile	*se*	*p*
–0.0324	0.1643	0.8438

Notes: A comparison of means between Bluetooth-recorded travel speeds and TNC vehicle speeds. The coefficient Δ *Minutes per Mile* represents the difference in means between travel speeds (in minutes per mile) for these TNC trips and the average travel times recorded by the corresponding segment over the 15-min period where the TNC trip occurred.

I match TNC trips to segments based on the following criteria: for a given segment, the sum of the distance between segment and TNC trip termini must be less than 500 m, and the distance traveled by the TNC vehicle must be within 10% of the segment length. I identify 1901 such matches, with several segments recording multiple TNC trips that fit this criteria. I then replicate the type of data filtering applied by Post Oak Traffic Systems. Recall that in the data I use for my analysis, only observations that fall within 75% of the interquartile range (IQR) of the 15 most recent observations are used to calculate average speeds. While I do not have access to the IQR data, I do have the standard deviation of speed measurements for any given 15-min interval. I use this to estimate the IQR, and drop RideAustin trips that fall outside of IQR estimate for the corresponding segment, time, and date, as these trips likely be dropped from the Bluetooth speed measurements.

After applying these filters, I am left with 221 trip matches, which are summarized in Table C.1. The regression coefficient reflects a difference in means between TNC trip speeds and the segment speeds recorded at corresponding times. On average, TNC vehicles traversed segments 0.03 minutes per mile slower than did the average recorded vehicle. This difference is not statistically significant, and under reasonable assumptions about TNCs as a share of total vehicles (5–15%), should not generate meaningful bias in the results reported above.

A shortcoming of this test is that I only observe drivers while they have a passenger in the car. This test will not identify differences caused by passengerless TNCs driving systematically above or below the flow of traffic. Note that the most extreme cases of this activity (e.g., idling while waiting for passengers) will be dropped by the IQR filter.

Appendix D. Robustness

In this section, I investigate the robustness of my results to alternative specifications and alternate road segment groups. Table D.1 shows the results of difference in difference regressions using alternate linear trend and weighting specifications. Table D.2 shows the results of running Eq. 1 on different sub and supersets of the Bluetooth data used in the main analysis. My results are stable over these specifications: The F-test rejects the null that the hour-specific effects are jointly zero. Additionally, the estimated annual congestion cost of TNC activity is of a similar magnitude to my preferred specification in each of these alternative specifications.

Fig. D.1 plots estimated regression discontinuity and difference in differences coefficients pooling over hours of day for symmetric bandwidths ranging from 20 to 70 days about May 9$^{\text{th}}$. The regression discontinuity results are stable over this range. The difference in differences results show positive, but not statistically significant point estimates for a small minority of bandwidths. Across all bandwidths, for both RD and DID specifications, point estimates of congestion impacts between 7 a.m. and 7 p.m. are negative. Taken together, the results presented in Fig. D.1 are consistent with the conclusions in the body of this paper: daytime travel speeds in Austin likely increased following the exit of Uber and Lyft.

Fig. D.2 displays regression discontinuity estimates of the impact of TNC departure on traffic speeds using the actual TNC exit date in relation to the distribution of coefficients from 134 regression discontinu-

Table D1
Alternate specifications.

Hour of Day	Model 1		Model 2		Model 3		Model 4	
	β_h	p	β_h	p	β_h	p	β_h	p
0	−0.0606	0.26	−0.0632	0.24	−0.0789	0.21	−0.1059	0.14
1	−0.0525	0.16	−0.0550	0.13	−0.0850	0.13	**−0.1405**	0.09
2	**−0.0505**	0.09	**−0.0508**	0.09	−0.1007	0.12	−0.0195	0.74
3	−0.0121	0.85	−0.0158	0.81	−0.1051	0.14	0.1506	0.27
4	0.0911	0.38	0.0841	0.41	−0.0795	0.38	0.1073	0.57
5	−0.0293	0.75	−0.0365	0.71	**−0.1850**	0.10	0.0886	0.66
6	−0.0568	0.58	−0.0765	0.43	**−0.2831**	0.03	−0.0599	0.75
7	−0.1446	0.24	−0.1647	0.16	**−0.3566**	0.07	0.0425	0.85
8	−0.0583	0.67	−0.1072	0.52	−0.1727	0.26	−0.3543	0.23
9	−0.0188	0.85	−0.0589	0.58	−0.1561	0.25	−0.0924	0.55
10	0.0318	0.59	0.0133	0.84	−0.1190	0.25	0.0001	1.00
11	**−0.1418**	0.05	**−0.1445**	0.04	−0.0621	0.29	−0.1327	0.14
12	**−0.1730**	0.05	**−0.1771**	0.04	**−0.1745**	0.05	**−0.1885**	0.10
13	−0.1555	0.12	**−0.1615**	0.09	−0.1082	0.12	**−0.2408**	0.08
14	−0.0176	0.71	−0.0224	0.63	−0.0379	0.42	−0.0103	0.84
15	−0.0238	0.61	−0.0304	0.49	−0.0257	0.60	−0.0026	0.97
16	0.0657	0.35	0.0578	0.37	0.0515	0.33	0.0715	0.32
17	0.0016	0.98	−0.0028	0.96	0.0248	0.60	0.0360	0.49
18	0.0500	0.57	0.0459	0.59	−0.0102	0.85	0.0615	0.33
19	−0.0300	0.65	−0.0285	0.65	−0.0254	0.63	−0.0327	0.57
20	0.0004	1.00	−0.0065	0.95	−0.0192	0.76	0.0324	0.73
21	0.0817	0.26	0.0775	0.27	0.0090	0.90	**0.1261**	0.08
22	−0.0709	0.42	−0.0763	0.37	−0.0491	0.51	−0.0502	0.57
23	−0.0152	0.65	−0.0235	0.53	**−0.1450**	0.07	**−0.1187**	0.04
F-test		0.00		0.00		0.00		0.00
Annual Cost		−33,605,827		−49,764,126		−91,753,234		−54,047,223

Notes: Results from Eq. 1. Model 1 reproduces the results from my preferred specification, with year-segment specific linear trends. Model 2 includes only year-specific linear trends (i.e, pools across segments). Model 3 includes only segment-specific linear trends (i.e., pools across years). Model 4 uses the same specification as column 1, but weights observations (which are 15-min average travel times) by the number of vehicles observed. Bold coefficients are significant at the 10% level. The penultimate row reports the p-value from a joint hypothesis test of $\beta_h = 0\ \forall h$; the final row reports annual costs using hourly VOT estimates from Austin's MoPac freeway (see Appendix A).

Table D2
Alternate segment groups.

Hour of Day	Group 1		Group 2		Group 3	
	β_h	p	β_h	p	β_h	p
0	−0.0606	0.26	−0.0379	0.44	−0.0399	0.78
1	−0.0525	0.16	**−0.0482**	0.08	−0.1551	0.12
2	**−0.0505**	0.09	−0.0412	0.19	−0.0854	0.42
3	−0.0121	0.85	−0.0125	0.74	0.0498	0.74
4	0.0911	0.38	0.0547	0.39	0.0785	0.64
5	−0.0293	0.75	−0.0124	0.87	−0.0177	0.88
6	−0.0568	0.58	−0.0147	0.84	0.2108	0.36
7	−0.1446	0.24	−0.0899	0.35	−0.1519	0.51
8	−0.0583	0.67	−0.1080	0.35	−0.1177	0.82
9	−0.0188	0.85	−0.0448	0.51	−0.2941	0.27
10	0.0318	0.59	−0.0162	0.73	−0.1159	0.42
11	**−0.1418**	0.05	**−0.1416**	0.05	**−0.2689**	0.07
12	**−0.1730**	0.05	−0.1246	0.19	**−0.3233**	0.03
13	−0.1555	0.12	−0.0956	0.28	−0.3944	0.13
14	−0.0176	0.71	0.0029	0.95	−0.1260	0.43
15	−0.0238	0.61	−0.0217	0.56	−0.0168	0.91
16	0.0657	0.35	0.0435	0.47	0.0059	0.96
17	0.0016	0.98	−0.0307	0.62	0.0587	0.49
18	0.0500	0.57	−0.0222	0.76	−0.0479	0.77
19	−0.0300	0.65	−0.0520	0.30	−0.1182	0.41
20	0.0004	1.00	0.0036	0.97	0.0546	0.70
21	0.0817	0.26	0.1056	0.10	0.0967	0.58
22	−0.0709	0.42	−0.0468	0.60	−0.1515	0.36
23	−0.0152	0.65	−0.0243	0.48	−0.1332	0.46
F-test		0.00		0.00		0.00

Notes: Results from Eq. 1, applied different groups of road segments. Group 1 is my preferred specification, which uses all traffic segments which report in more than 70 percent of days in each year. Group 2 relaxes this level to segments that report in 30 percent of days. Group 3 uses only segments that report in every day of the study period. Bold coefficients are significant at the 10% level. The final row reports the p-value from a joint hypothesis test of $\beta_h = 0\ \forall h$.

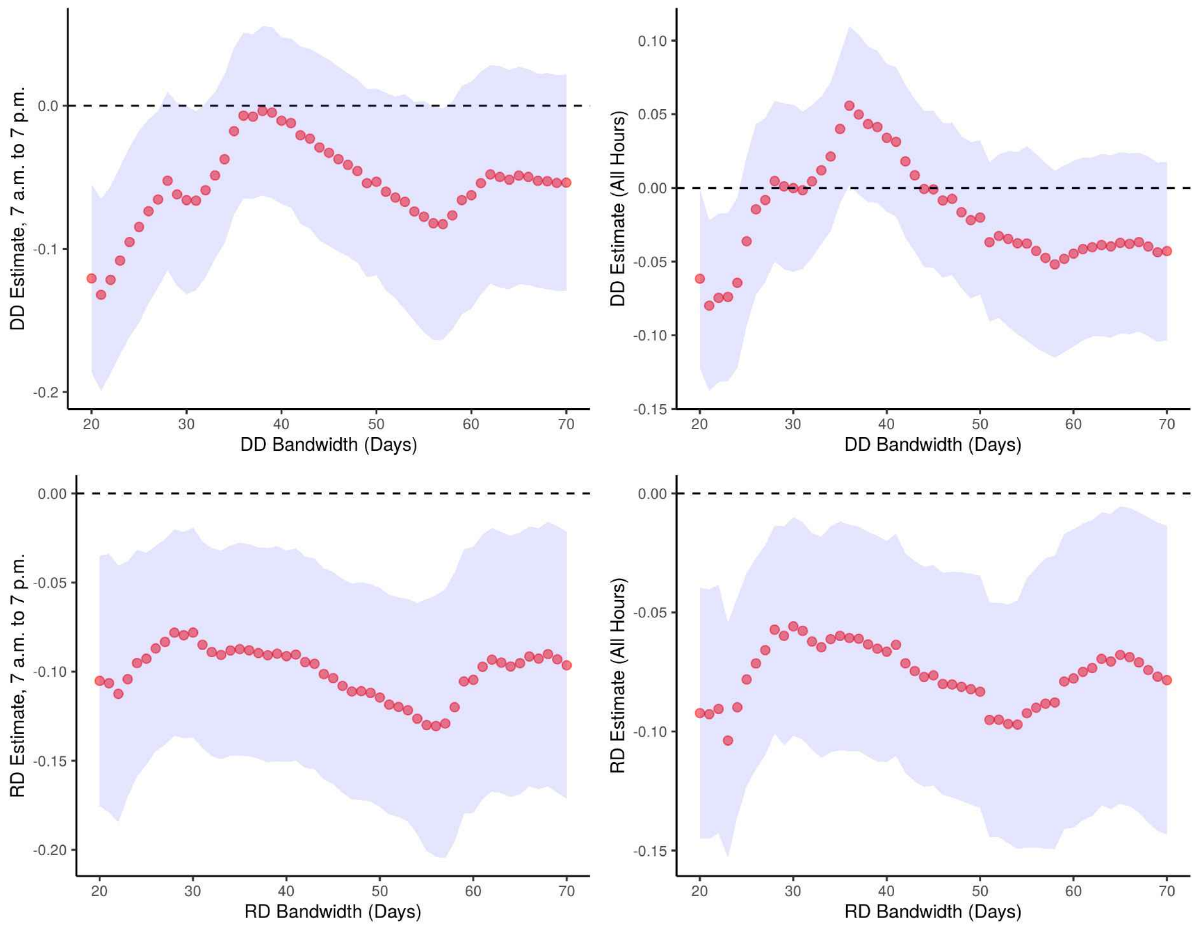

Fig. D1. Bandwidth Sensitivity. *Notes:* This figure plots estimated difference in differences and regression discontinuity coefficients pooling over hours of day (a variation of Eq. 3) for symmetric bandwidths ranging from 20 to 70 days about May 9th, 2016. Bars represent the 95% confidence interval using standard errors clustered by segment-week. Specifications with a bandwidth of over 45 days include a dummy for the SXSW festival. For reference, my preferred specification uses an asymmetric bandwidth of 45 pre-period days and 52 post-period days.

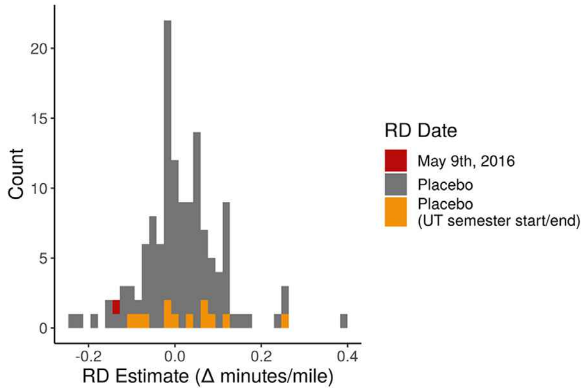

Fig. D2. Regression Discontinuity Placebo Tests. *Notes:* This figure displays regression discontinuity estimates of the impact of TNC departure on traffic speeds in Austin, TX, using the actual TNC exit date (in red) in relation to the distribution of coefficients from 134 regression discontinuities using placebo exit dates (in grey and yellow). The grey cells are results from 121 placebo regression discontinuity dates that run every 20 days from 2015 to 2019 with a 30 day bandwidth. The yellow cells are results from using the start/end of the UT semester as placebo regression discontinuity dates. I omit placebo regressions with significant amounts of missing data (missing more than 75% of days within the RD bandwidth). Controls in each regression are day of week, holiday, and segment fixed effects, segment-specific second-degree polynomials in days since May 9th, and flexible controls for temperature and precipitation. Specifications with March dates include a dummy for the SXSW festival.

ities using placebo exit dates. The results of this exercise suggest that a) it is unlikely for shocks to the Austin-area traffic system to create changes in travel speeds on the order of the estimates reported in Table 4, and b) it is unlikely that the RD results reflect contemporaneous changes in traffic resulting from the end of the University of Texas, Austin semester.

Appendix E. Segment length revisions

The raw Bluetooth traffic data available on the Austin Open Data Portal show the lengths of traffic segments varying over time. Of the 79 segments that I use in my analysis, 61 segment lengths changed during the study period, with the majority of these adjustments occurring on March 24th, 2016. On average, these adjustments were small: only four segment lengths were adjusted by more than 4%. According to the data providers, these adjustments most likely reflect updated distance measurements, not the repositioning Bluetooth readers. As such, I use updated segment lengths to calculate average speeds in all time periods.

To verify that this data quality issue does not constitute a threat to identification, I run a regression discontinuity about each segment length change. Of the 61 segments where lengths were adjusted, 47 had adequate data to run a regression discontinuity about the date where the segment length was changed. 8 of these 47 regression discontinuities registered statistically significant changes in traffic speeds. These results of these RDs are plotted in Fig. E.1.

Table E.1 shows that omitting these 8 segments from my study pool does not substantively change the results presented in the body of this paper. Columns 1 and 2 reproduce estimates from Table 4, and columns

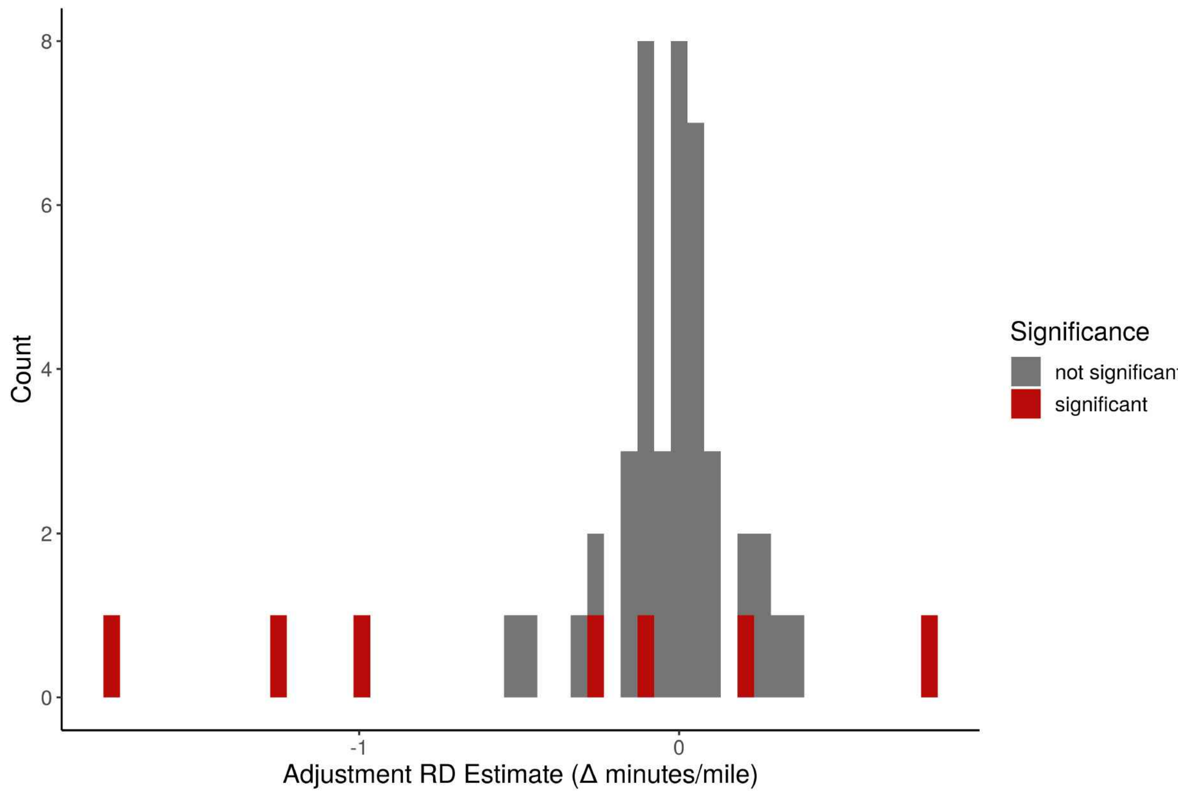

Fig. E1. Speed Around Segment Length Revisions. *Notes:* This figure plots coefficients from 47 regression discontinuity specifications estimated where the raw Bluetooth data showed a change in a segment length. Each cell corresponds to the estimated change in traffic speed (in minutes per mile) estimated about the date of a single length revision. Controls in each regression include day of week, holidays, SXSW, and hour fixed effects, segment-specific linear trends in days since adjustment, and flexible temperature and precipitation functions. Significant results are potted in red. Standard errors in all regressions are Newey-West.

Table E1
Sensitivity to segments with anomalous length revision data.

	β (Table 4)	*se*	β (Restricted sample)	*se*
Difference in Differences (All hours)	−0.02605	0.01701	−0.02972	0.02462
Difference in Differences (7 a.m. - 7 p.m.)	−0.06838	0.05295	−0.06137	0.03211
Regression Discontinuity (All Hours)	−0.10155	0.03535	−0.06498	0.01899
Regression Discontinuity (7 a.m. - 7 p.m.)	−0.13351	0.04331	−0.08567	0.02529

Notes: The first two columns reproduce the results from Table 4, which reports difference in differences and regression discontinuity results pooled across the hours of the day (Eqs. 1 and 3, respectively). Columns 3 and 4 show results from the same regressions run on a sample that excludes the 8 segments that registered a statistically significant change in traffic speeds about the adjustment date.

3 and 4 re-run these specifications excluding the 8 segments that registered a statistically significant change about their readjustment date. Across all four specifications, analyses using a restricted pool of segments suggest a modest increase in traffic speeds following the exit of Uber and Lyft from Austin. Note that the magnitude of the regression discontinuity results are sensitive to the exclusion of these 8 segments.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jue.2020.103318.

CRediT authorship contribution statement

Matthew Tarduno: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Visualization.

References

Anderson, M.L., Lu, F., Zhang, Y., Yang, J., Qin, P., 2016. Superstitions, street traffic, and subjective well-being. J. Public Econ. 142, 1–10.
Angrist, J., Caldwell, S., Hall, J., 2017. Uber vs. taxi: A Driver's eye view. NBER Working Paper 23891.
Buchholz, N., Doval, L., Kastl, J., Matějka, F., Salz, T., 2020. The Value of Time: Evidence From Auctioned Cab Rides. Technical Report. National Bureau of Economic Research.
Clewlow, R., Mishra, G., 2017. Disruptive Transportation: the Adoption, Utilization, and Impacts of Ride-Hailing in the United States. Institute of Transportation Studies, University of California, Davis Research Report UCD-ITS-RR-17-07.
Cohen, P., Hahn, R., Hall, J., Levitt, S., Metcalfe, R., 2016. Using Big Data to Estimate Consumer Surplus: The Case of Uber. Technical Report. National Bureau of Economic Research.
Cramer, J., Krueger, A.B., 2016. Disruptive change in the taxi business: the case of uber. Am. Econ. Rev. 106 (5). 177–82
Currie, J., Walker, R., 2011. Traffic congestion and infant health: evidence from E-ZPass. Am. Econ. Rev. 3, 65–90.
Diamond, P.A., 1973. Consumption externalities and imperfect corrective pricing. Bell J. Econ. Manag. Sci. 526–538.
Eliasson, J., 2009. A cost–benefit analysis of the stockholm congestion charging system. Transp. Res. Part A: Policy Pract. 43 (4), 468–480.
Erhardt, G.D., Roy, S., Cooper, D., Sana, B., Chen, M., Castiglione, J., 2019. Do transportation network companies decrease or increase congestion? Sci. Adv. 5 (5), eaau2670.
Goodman, L.A., 1960. On the exact variance of products. J. Am. Stat. Assoc. 55 (292), 708–713.

Greenwood, B.N., Wattal, S., 2015. Show me the way to go home: an empirical investigation of ride sharing and alcohol related motor vehicle homicide. Fox Sch. Bus. Res. Pap. (15–054).
Hall, J.D., Palsson, C., Price, J., 2018. Is uber a substitute or complement for public transit? J. Urban Econ. 108, 36–50.
Hampshire, R., Simek, C., Fabusuyi, T., Di, X., Chen, X., 2017. Measuring the impact of an unanticipated disruption of uber/lyft in austin, tx. Lyft in Austin, TX (May 31, 2017).
Inrix, 2017. INRIX Global Traffic Scorecard.
Inrix, 2018. INRIX Global Traffic Scorecard.
King, D., Manville, M., Shoup, D., 2007. The political calculus of congestion pricing. Transp. Policy (Oxf) 14 (2), 111–123.
Knittel, C.R., Sandler, R., 2018. The welfare impact of second-best uniform-Pigouvian taxation: evidence from transportation. Am. Econ. J.: Econ. Policy 10 (4). 211–42
Leape, J., 2006. The london congestion charge. J. Econ. Perspect. 20 (4), 157–176.
Li, Z., Hong, Y., Zhang, Z., 2019. Do ride-sharing services affect traffic congestion? an empirical study of uber entry. Working Paper.
Mangrum, D., Molnar, A., 2018. The marginal congestion of a taxi in new york city. Working Paper.
New York Times, 2016. Uber and Lyft End Rides in Austin to Protest Fingerprint Background Checks.
New York Times, 2018. Uber Hit With Cap as New York City Takes Lead in Crackdown.
New York Times, 2019. Your Taxi or Uber Ride in Manhattan Will Soon Cost More.
Prud'Homme, R., Bocarejo, J.P., 2005. The london congestion charge: a tentative economic appraisal. Transp. Policy (Oxf) 12 (3), 279–287.
Rayle, L., Shaheen, S., Chan, N., Dai, D., Cervero, R., 2014. App-Based, on-Demand ride services: comparing taxi and ridesourcing trips and user characteristics in san francisco. University of California Transportation Center.
Reuters, 2019. Uber and other taxi firms to pay London congestion charge.
RideAustin, 2017. Comprehensive Ride Data.
San Francisco Transit Authority, 2018. TNCs and Congestion.
Schaller Consulting, 2018. The New Automobility: Lyft, Uber and the Future of American Cities.
Small, K., Verhoef, E., 2007. The economics of urban transportation. Routledge.
The 85th Texas Legislature, 2017. Texas House Bill 100.
The City Council of Austin, 2015. Ordinance 20151217-075.
The Texas Tribune, 2016. Austin's Proposition 1 Defeated.
Uber, 2015. Case Study Shows Our Impact in Austin.
United States Census Bureau, 2015. American Commnity Survey.
Vancouver Sun, 2019. Dan Fumano: Vancouver wants to charge Uber and Lyft users a congestion fee.
Vickrey, W.S., 1969. Congestion theory and transport investment. Am. Econ. Rev. 59 (2), 251–260.