

UBER AND TRAFFIC FATALITIES

Michael L. Anderson and Lucas W. Davis*

Abstract—Previous studies of the effect of ridesharing on traffic fatalities have yielded inconsistent conclusions. We revisit this question using proprietary data from Uber measuring monthly rideshare activity at the Census tract level. We find a consistent negative effect of ridesharing on traffic fatalities, with impacts concentrated during nights and weekends. Our results imply that ridesharing has decreased U.S. traffic fatalities by 5.2% in areas where it operates. The annual life-saving benefits are \$6.8 billion. Back-of-the-envelope calculations suggest that these benefits are of similar magnitude to producer surplus captured by Uber shareholders or consumer surplus captured by Uber riders.

I. Background

THERE is a long history in economics of empirical analyses of traffic fatalities (Peltzman, 1975; Levitt & Porter, 2001; Cohen & Einav, 2003; Ashenfelter & Greenstone, 2004; White, 2004; Anderson, 2008a; Abouk & Adams, 2013; Jacobsen, 2013; Anderson & Auffhammer, 2014; DeAngelo & Hansen, 2014). One of the most significant recent changes in road transportation has been the introduction of ridesharing. The existing literature has hypothesized two primary mechanisms by which ridesharing could impact traffic fatalities. On the one hand, ridesharing affects the *composition* of drivers on the road, for example, substituting inebriated drivers with sober drivers, potentially reducing total traffic fatalities. On the other hand, ridesharing also affects the overall *scale* of driving, with the convenience and low cost of ridesharing possibly leading to an increase in total travel in passenger vehicles.

A number of recent studies empirically estimate the effects of ridesharing on traffic fatalities.¹ This existing work focuses on the timing of Uber entry into markets and yields inconsistent, often contradictory, conclusions. Brazil and Kirk (2016) and Zhou (2020) exploit the timing of Uber rollout across U.S. counties and find no associations with traffic fatalities or drunk driving respectively. Dills and Mulholland (2018) finds that the relationship between Uber entry and traffic fatalities can be negative or positive, depending on

the specification.² Greenwood and Wattal (2017) and Peck (2017) focus on the timing of Uber rollout within California and New York City (NYC) respectively; both studies find reductions in alcohol-related fatalities after the introduction of Uber. In contrast, Barrios et al. (2023) exploits city-level timing of Uber and Lyft rollout across the United States and concludes that ridesharing *increases* traffic fatalities, while Cairncross et al. (2025) applies synthetic-control methods and finds a statistically insignificant relationship between ridesharing and traffic fatalities in Vancouver.

In summary, the existing literature studies the effects of market entry by Uber and finds that it may cause traffic fatalities to decrease, increase, or remain unchanged.³ Our study represents, to the best of our knowledge, the first work using proprietary Uber ridership data to estimate the effects of ridesharing on traffic fatalities. With these more detailed data we find a robust negative impact of ridesharing on traffic fatalities. Impacts are negative and statistically significant across a range of alternative specifications and concentrate during evenings and nights, as expected.

The paper performs several back-of-the-envelope calculations to put our results in context. Scaled up to reflect current ridership levels, our results imply that ridesharing reduces total U.S. traffic fatalities by 5.2% in areas where it operates. Based on conventional estimates of the value of statistical life (VSL), the annual life-savings benefits are \$6.8 billion. We compare these impacts to the market capitalization of Uber and estimates in the literature for the total consumer surplus from ridesharing.

II. Data

We combine data from two sources: the National Highway Traffic Safety Administration (NHTSA) Fatality Analysis Reporting System (FARS) and rideshare activity from Uber. FARS data represent a census of fatal crashes in the U.S. FARS contains detailed information on each crash, including geographic coordinates, roadway type, date and time, and suspected involvement of alcohol. We include FARS data from 2012 to 2016.

Appendix table A1 reports summary statistics from the FARS data. In the U.S. as a whole there are an average of 34,077 fatalities annually over the sample period. In areas that see rideshare activity by 2016—our analytic sample—there are an average of 15,672 fatalities annually. In these

²In the baseline results (table 3, panel A), Dills and Mulholland (2018) reports eight estimates; five negative, three positive, and none statistically significant.

³One exception is Morrison et al. (2017), which studies the disruption and resumption of Uber service in Las Vegas, Reno, San Antonio, and Portland. It finds mixed evidence—in Portland, the resumption of Uber service correlated with a drop in alcohol-involved crashes, but in the other cities, it did not.

Received for publication March 7, 2022. Revision accepted for publication September 5, 2023. Editor: Tavneet Suri.

*Anderson and Davis: University of California.

The authors have not received any financial compensation for this project nor do they have any financial relationships that relate to this research. They thank Luna Yue Huang for excellent research assistance, Jonathan D. Hall, Yizhen Gu, Matt Tarduno, the editor and three anonymous reviewers, and seminar participants at UC Berkeley, LSE, Montana State University, and the National Bureau of Economic Research for helpful comments, and Jonathan Wang, Santosh Rao Danda, and Cory Kendrick for assistance in accessing Uber data.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01385.

¹Additional research has focused on other potential externalities of ridesharing. For example, ridesharing affects congestion (Hall et al., 2018; Tarduno, 2021), labor markets (Berger et al., 2018; Chen et al., 2019), and alcohol consumption (Teltser et al., 2021).

areas, approximately 33% of fatalities involve multivehicle collisions, 40% involve single-vehicle collisions, and 26% involve pedestrians or bicyclists.

Alcohol involvement is reported in approximately 30% of fatal crashes, and we expect deterring drunk-driving to be a primary mechanism through which ridesharing affects fatalities. FARS contains several data elements pertaining to alcohol involvement. We use a variable that codes the number of drinking drivers involved in a crash. A driver qualifies as drinking if he has a positive blood alcohol concentration (BAC) or if police report alcohol involvement.⁴ We classify the crash as involving alcohol if at least one driver was drinking. This undercounts the true number of alcohol-involved crashes because police may not always detect or report alcohol involvement, and alcohol data are “often missing”, resulting in an “undercount [of] the actual number of drunk drivers” (National Highway Traffic Safety Administration, 2016, p. 72). Some fatal crashes with no reported alcohol involvement therefore involve alcohol. Thus our analysis focuses on models that specify any fatal crash as the dependent variable, but we also explore models that specify any alcohol-involved fatal crash as the dependent variable in section IIIC.

Rideshare activity data come from Uber. These data report trip counts aggregated to the Census tract-by-month level for all Uber trips originating in a given Census tract, excluding those for which the origin or destination is an airport.⁵ The data cover 1 July 2012 through 1 January 2017 for all Census tracts in the United States, excluding those in Seattle and NYC.⁶ The start date coincides with the launch of UberX—the service that people generically refer to as “Uber”—and the end date represents the latest data Uber were willing to share with us. Trip counts are normalized to the level of a specific Census tract in San Francisco during May 2015, and Census tracts with less than 0.05% of this level of trip activity in a given month are rounded down to zero. While the authors had input into the data extraction parameters, decisions on what data to release ultimately rested with Uber staff.

To construct the analytic data set, we merged FARS data and Uber trip data at the tract-by-month level. Specifically, we assigned each crash in FARS to a Census tract and month. We then totaled different outcomes—for example, total fatalities, total alcohol-involved fatalities, total fatalities by time-of-day—at the tract-by-month level. Our largest estimation sample contains 2.7 million tract-by-month observations.

Uber trips are recorded based on Census tract of origin. The average Census tract, however, contains approximately

4,000 inhabitants, and a large county contains dozens of Census tracts. Many trips thus traverse more than one Census tract—the median Uber trip during our sample period was approximately five miles in length, and the 90th percentile trip was approximately ten miles in length.⁷ To account for the multitract nature of most trips, we constructed the treatment of interest to be an inverse-distance-weighted average of Uber trip activity in nearby Census tracts. Specifically, for tract i we took a weighted average of trip activity in all tracts whose centroids are within 10 miles of tract i 's centroid, with weights proportional to $distance_{ij}^{-1}$, where $distance_{ij}$ measures the distance in miles between tract i and tract j .⁸ We confirm that our results are not sensitive to the exact specification of the weight, but given the spatially correlated nature of the treatment our analysis is best conceptualized as a city-level analysis.

One limitation of our data is that we do not measure activity for other ridesharing companies. During our sample period, however, Uber captured most of the U.S. ridesharing market. For example, Lyft—Uber's largest U.S. competitor by far—had approximately 6% and 14% of Uber's 2015 and 2016 U.S. bookings respectively (the last two years of our sample; see appendix A1.1 for details). Thus, our data represent a good proxy for overall ridesharing activity during our sample period.

III. Empirical Strategy and Results

A. Main Results

To determine the effect of Uber rideshare activity on traffic fatalities, we estimate a series of regressions of the form:

$$y_{it} = \beta \cdot \text{Rideshare}_{it} + \delta_t + \gamma_i + \varepsilon_{it}. \quad (1)$$

The dependent variable y_{it} is an indicator for the presence of any fatal crash in tract i in month t .⁹ When reporting

⁷Source: Communications with Uber staff.

⁸Tract i 's weight is normalized to one, and a tract whose centroid is one mile from tract i 's centroid receives a weight of 0.2. In a perfect grid each tract would be surrounded by four other tracts (suggesting a weight of 0.25, which yields virtually identical results), but we chose 0.2 because in an extremely small number of cases the distance between two neighboring tracts can be as little as 0.25 miles (the average tract size in a moderately-dense city is on the order of one square mile; for example, San José, CA is 181 square miles and contains approximately 200 Census tracts). In a simple gravity model in which trips radiate uniformly in all directions and continue indefinitely, the share of trips originating in j that cross tract i is proportional to the inverse of the distance between the two tracts. However, there are two complicating factors in our context. First, trips do not continue indefinitely; they die off with distance. Second, trips do not radiate uniformly, but tend to concentrate within travel corridors. The first factor suggests that the weight should decline more strongly with distance, while the second factor suggests that the weight should decline less strongly with distance. We assume that the two factors roughly offset each other, but we check that our results are robust to weights that decline somewhat more or less steeply than $distance_{ij}^{-1}$ (appendix tables A2 and A3).

⁹There is rarely more than one fatal crash in a tract-month observation, and there is typically one fatality per fatal crash. Average total fatalities in a tract-month observation conditional on any fatal crash occurring are 1.13.

⁴According to the FARS codebook, a driver charged with an alcohol violation does not count as drinking unless the driver also has a positive BAC or the police report alcohol involvement.

⁵Census tracts are based on 2010 Census definitions.

⁶To protect their proprietary data, Uber provided us with trip counts normalized to an arbitrary level. Seattle and NYC, however, publish aggregate information about ridesharing trips. Thus, if our data included Census tracts in these two cities, we could back out raw trip counts. Without tract-level information for Seattle or NYC, we are unable to include either city in our analysis.

TABLE 1.—THE EFFECT OF UBER ON TRAFFIC FATALITIES, MAIN RESULTS

	(1)	(2)	(3)	(4)	(5)
	Any death				
Rideshare index	-0.129 (0.022) <i>0.001</i>	-0.127 (0.022) <i>0.001</i>	-0.147 (0.031) <i>0.001</i>	-0.131 (0.037) <i>0.002</i>	-0.103 (0.039) <i>0.013</i>
Observations	2,664,480	2,186,460	1,091,760	657,840	439,140
Max tract rideshare activity	Nonzero	Top 50 pct	Top 25 pct	Top 15 pct	Top 10 pct
Tracts	44,408	36,441	18,196	10,964	7,319
Mean dependent variable	2.675	2.451	2.243	2.130	1.982
Mean index	0.119	0.145	0.279	0.432	0.598

This table reports coefficient estimates from five separate least squares regressions that progressively restrict the sample to locations with higher ridesharing activity by the end of our sample period. The unit of observation is the Census tract by month. The dependent variable in all regressions is an indicator for any fatalities, multiplied by 100. The independent variable is the weighted average of rideshare activity originating within a 10-mile radius of Census tract i in month t , with weights equal to $0.2 \times \text{distance}^{-1}$, normalized such that a value of 1 corresponds to the average value of the index in 2019. All regressions include Census-tract and month-of-sample fixed effects. Parentheses contain standard errors clustered by CBSA. FDR-control q -values in *italics*.

results we multiply y_{it} by 100 for coefficient readability. The independent variable of interest, Rideshare_{it} is the rideshare activity index for tract i in month t , constructed as described in section II. We rescale the independent variable such that a value of 1 corresponds to the average value of the index in 2019 in our main analytic sample (a value that lies well within the support of our data).¹⁰ The coefficient of interest thus approximates the average effect of ridesharing in 2019 on the probability of a fatal crash, measured in percentage points.

All specifications include month-of-sample (δ_t) and Census-tract (γ_i) fixed effects, controlling for any factors that change uniformly over time or differ across tracts in a time-invariant manner. The estimation sample runs from 2012 to 2016 and is restricted to include only tracts where Uber entry occurred during this period. We exclude other tracts because they are generally rural and unlikely to serve as a valid counterfactual. Conceptually this strategy is similar to a “staggered adoption” design, but with a continuous treatment.¹¹ The identifying assumption is that, after controlling for tract and month fixed effects, ridesharing growth

is uncorrelated with other tract-specific, time-varying factors that affect traffic fatalities. We probe this assumption in section IIIB.

To account for dependence over time and across tracts, we cluster standard errors at the Core-Based Statistical Area (CBSA) level.¹² Since we explore a variety of specifications and samples, we also report false-discovery-rate (FDR) adjusted “ q -values” (Benjamini & Hochberg, 1995). Briefly, the FDR represents the expected proportion of rejections that are false discoveries; controlling FDR at $q < 0.1$ indicates that 90% of rejections should be true rejections (Anderson, 2008b).¹³

Table 1 reports estimates of β for a range of sample restrictions. The first column estimates equation (1) using all tracts with nonzero ridesharing activity by the end of 2016 (approximately 44,000 tracts). A one-unit increase in ridesharing activity reduces the probability of a fatal crash by 0.129 percentage points ($t = -5.9$). This corresponds to a 4.8% decrease in fatalities. Columns 2 through 5 estimate equation (1) when restricting the sample to include tracts whose maximum rideshare activity falls within the top 50%, 25%, 15%, and 10% of all tract maximums respectively. The estimate of β in column 2 implies that a one-unit increase in ridesharing activity reduces fatalities by 0.127 percentage points ($t = -5.8$), or 5.2% of the mean. The estimates of β in columns 3 and 4 are similar in magnitude and remain highly significant. In column 5 the implied effect (-0.103 percentage points) is 5.2% of the mean and is statistically significant ($t = -2.6$), despite dropping 90% of tracts. Overall, the evidence in table 1 implies a significant negative relationship between ridesharing and traffic fatalities.¹⁴

¹⁰COVID-19 heavily impacted Uber activity post-2019.

¹¹Recent work on two-way fixed effect designs suggests that in some cases it can be important to consider how dynamic treatment effects and treatment effect heterogeneity may interact with the implicit regression weights in scenarios with staggered adoption of a “treatment”. Most concerning is the possibility that a two-way fixed effect regression might reject the null hypothesis even if the average treatment effect is zero. Much of this work pertains to binary treatments and is not directly relevant, but Callaway et al. (2024) addresses the case of a continuous treatment and suggests limiting comparisons to always involve never-treated units. Such comparisons are undesirable in our context, however, because rural (i.e., never-treated) areas are not credible controls for urban areas. Regardless, dynamic effects and treatment effect heterogeneity cannot plausibly explain our regression estimates. Dynamic effects are unrealistic in this context, as Uber rides in month t are unlikely to have large impacts on fatalities in month $t + s$. Furthermore, treatment effect heterogeneity, while likely present in some form, would be unlikely to generate our results if the true average treatment effect were zero. For us to estimate a negative coefficient if the true average treatment effect were zero, cities with low ridership growth, which implicitly serve as controls for those with high ridership growth, would need to have large, positive Uber treatment effects (i.e., Uber ridership would need to increase traffic fatalities in these cities). Such a scenario is unrealistic since the obvious mechanism for a positive effect of Uber on traffic fatalities would be through raw increases in vehicle-miles traveled due to Uber activity, which by definition will be minimal in cities with low ridership growth. Nevertheless, we check in section IIIB that a fully-interacted version of our main specification, esti-

imating a separate regression coefficient for each Census tract, generates an average point estimate that is similar to our main specification.

¹²Clustering at the county or state level yields qualitatively similar results.

¹³For FDR control we define the family of hypotheses tested to include all main-body and appendix tables, with the exception of explicit placebo tests, which form their own family. The main family of hypotheses includes tables 1, A2, A3, 2, 4, A4, and columns 3 and 4 of table 3 (29 tests). The family of placebo hypotheses includes columns 1 and 2 of table 3 (2 tests).

¹⁴Appendix tables A2 and A3 reproduce table 1 using an independent variable that applies a weight proportional to $\text{distance}^{-0.9}$ or $\text{distance}^{-1.1}$ respectively (as compared to distance^{-1}). The statistical significance and coefficient magnitudes are broadly similar to table 1. If anything the

TABLE 2.—THE EFFECT OF UBER ON TRAFFIC FATALITIES, STATE-BY-MONTH FES

	(1)	(2)	(3)	(4)	(5)
	Any death				
Rideshare index	−0.105 (0.020) <i>0.001</i>	−0.101 (0.021) <i>0.001</i>	−0.088 (0.021) <i>0.001</i>	−0.071 (0.024) <i>0.007</i>	−0.065 (0.023) <i>0.009</i>
Observations	2,664,480	2,186,460	1,091,760	657,660	439,080
Max tract rideshare activity	Nonzero	Top 50 pct	Top 25 pct	Top 15 pct	Top 10 pct
Tracts	44,408	36,441	18,196	10,961	7,318
Mean dependent variable	2.675	2.451	2.243	2.131	1.982
Mean index	0.119	0.145	0.279	0.432	0.598

This table reports coefficient estimates from five separate least squares regressions. Specifications are identical to table 1 except that all regressions include state-by-month-of-sample fixed effects. Parentheses contain standard errors clustered by CBSA. FDR-control q -values in *italics*.

B. Threats to Identification

We anticipate two primary threats to identification in equation (1): differential trends in factors specifically affecting drunk driving and differential trends in factors affecting general traffic fatality rates. Examples of the first threat include alcohol prices, drunk-driving penalties, and enforcement. Examples of the second threat include changes in vehicle miles traveled, infrastructure improvements, vehicle-fleet technology, driver characteristics, and health care systems.

To address the first threat, we modify equation (1) to include state-by-month fixed effects, δ_{st} :

$$y_{ist} = \beta \cdot \text{Rideshare}_{ist} + \delta_{st} + \gamma_i + \varepsilon_{ist}. \quad (2)$$

Ridesharing growth in our sample occurs over a period of only several years, and most policies that could quickly affect alcohol prices or drunk driving—including alcohol taxes, drunk-driving penalties, and highway-patrol enforcement—vary primarily at the state level. State-by-month fixed effects shift the identifying variation away from between-CBSA growth in ridesharing activity and towards within-CBSA growth in ridesharing activity, thus absorbing many potential confounders of interest. This shift, however, could exacerbate attenuation bias from measurement error in ridesharing activity.¹⁵

Table 2 reports results from estimating equation (2) with traffic fatalities as the dependent variable. Like table 1, the first column uses all tracts with nonzero ridesharing activity by the end of 2016, while columns 2 through 5 restrict the sample to include tracts whose maximum rideshare activity falls within the top 50%, 25%, 15%, and 10% of all tract maximums respectively.¹⁶ In all columns the estimates of β are attenuated in magnitude relative to the analogous estimates from table 1 (between 19% and 46% smaller) but they

patterns in appendix table A2 suggest that using a distance weight that declines less steeply might yield larger estimates. To remove researcher discretion, however, we prefer the simple inverse-distance weight.

¹⁵Identification from within-CBSA growth in ridesharing activity may be sensitive to measurement error in our inverse-distance-weighted average of nearby trip activity. To the degree that we cannot perfectly assign multitract trips to all affected tracts, analyses that focus more heavily on within-CBSA comparisons of different Census tracts (i.e., models that include state-by-month fixed effects) could suffer from attenuation bias.

¹⁶Sample sizes may vary slightly across tables due to singletons that arise upon inclusion of state-by-month fixed effects.

remain highly significant with t -statistics between -2.8 and -5.3 .

To address the second threat, we use daylight hours to execute a “triple-differences” design. If deterring drunk-driving is the primary mechanism through which ridesharing decreases fatalities, then we expect the effects of ridesharing activity to concentrate during evenings and nights for two reasons. First, drunk-driving fatalities concentrate during evenings and nights; in our analytic sample, approximately 75% of identified alcohol-involved fatalities occur between 8 pm and 6 am (see appendix figures A2 and A3). Second, daytime drunk-driving fatalities are less plausibly related to dining and entertainment trips, and more plausibly related to alcohol abuse. To the extent that ridesharing is a better substitute for driving for dining and entertainment trips, ridesharing activity should have stronger effects during night than day.

Most other factors that affect traffic fatality rates, however, should impact both daytime and nighttime fatalities. Testing for effects during daytime hours thus represents a falsification test for our research design. Table 3 reports estimates of equation (2) when limiting measurement of the dependent variable to daytime hours (8 am to 5 pm). Columns 1 and 2 report estimates of β for tracts whose maximum rideshare activity is within the top 50% or the top 25% (of maximums) respectively. In both cases the estimate of β is close to zero, precisely estimated, and statistically insignificant. These results imply the absence of differential trends in traffic fatalities in high rideshare-growth tracts.

We formally estimate the triple-differences design with the regression:

$$y_{itd} = \beta \cdot \text{Rideshare}_{it} \cdot \mathbf{1}(d = 1) + \delta_{itd} + \gamma_{it} + \phi_{id} + \alpha_{std} + \varepsilon_{itd}, \quad (3)$$

where y_{itd} is an indicator for any fatalities in tract i in month t during time-of-day d . We let $d = 0$ correspond to hours between 8 am and 5 pm (daytime), and $d = 1$ correspond to all other hours. The coefficient of interest, β , is on the interaction between Rideshare_{it} and an indicator for nondaytime hours. To operationalize the triple-differences design equation (3) includes month-by-time-of-day (δ_{itd}), tract-by-month (γ_{it}), and tract-by-time-of-day (ϕ_{id}) effects, absorbing all but the itd -level variation. For completeness

TABLE 3.—THE EFFECT OF UBER ON TRAFFIC FATALITIES, DAYTIME HOURS AND TRIPLE DIFFERENCES

	(1) Any daytime death	(2) Any daytime death	(3) Any death	(4) Any death
Rideshare index	-0.011 (0.016) <i>0.984</i>	-0.000 (0.016) <i>0.984</i>		
Rideshare index * Night			-0.079 (0.026) <i>0.005</i>	-0.087 (0.021) <i>0.001</i>
Observations	2,186,460	1,091,760	4,372,920	2,183,520
Max tract rideshare activity	Top 50 pct	Top 25 pct	Top 50 pct	Top 25 pct
Tracts	36,441	18,196	36,441	18,196
Mean dependent variable	0.731	0.619	1.226	1.122
Mean index	0.145	0.279	0.0726	0.139

This table reports coefficient estimates from four different least squares regressions. The unit of observation is the Census tract by month (columns 1 and 2) or Census tract by month by time-of-day (columns 3 and 4). The dependent variable in columns 1 and 2 is an indicator for any fatalities during daytime hours (8 am to 5 pm), multiplied by 100. The dependent variable in columns 3 and 4 is an indicator for any fatalities in either daytime or non-daytime hours, multiplied by 100. In all regressions, the independent variable is the weighted average of rideshare activity originating within a 10-mile radius of Census tract i in month t , with weights equal to $0.2 \times \text{distance}^{-1}$, normalized such that a value of 1 corresponds to the average value of the index in 2019. All regressions include Census-tract and state-by-month-of-sample fixed effects; columns 3 and 4 include tract-by-month-of-sample, time-of-day-by-month-of-sample, tract-by-time-of-day, and state-by-month-of-sample-by-time-of-day fixed effects. Parentheses contain standard errors clustered by CBSA. FDR-control q -values in *italics*.

we include state-by-month effects, now interacted with time-of-day (α_{std}).

Columns 3 and 4 report estimates of β from equation (3), using the same estimation samples as columns 1 and 2, respectively. In both cases, the estimate of β is approximately similar in magnitude to the corresponding estimate from table 2 and statistically significant ($t = -3.0$ and $t = -4.1$ respectively), as expected given the null effects in columns 1 and 2.

We also provide visual evidence of the effects of ridesharing on fatalities. Our continuous treatment measure precludes a standard event-study design, but we approximate an event-study design by splitting the sample into two groups based on ridesharing growth. We define the “treated” group as any tract that experiences ridesharing growth of approximately two standard deviations (or more) during the sample period. This split is analogous to the implicit treatment-control split in a conventional event-study design with a treated population of roughly 50%.¹⁷ To generate the figure we residualize the data with respect to treatment group and aggregate time indicators and then compute group-month means of the residualized data. We normalize the initial period to zero for both groups and, following Weill et al. (2021), we plot three-month moving averages of the group-month means. We cluster bootstrap confidence intervals at the CBSA level.

Figure 1 reveals that fatalities in “treated” Census tracts fall markedly relative to “control” tracts. The control time series is flat with minimal variation—this pattern occurs because most tracts are control tracts, so the aggregate time indicators remove most of the variation.¹⁸ The treatment group time series declines by about 30% (0.6 percentage points)

¹⁷The standard deviation of a Bernoulli random variable with $p = 0.5$ is 0.5. A change in treatment from 0 to 1 thus represents a two standard deviation change when $p = 0.5$ and a greater than two standard deviation change when $p \neq 0.5$.

¹⁸In a standard event study figure, the control group would be mechanically constrained to follow the $y = 0$ axis, as the plotted treatment group coefficients would represent treatment group means relative to the control group. Appendix figures A4 and A5 plot versions of figure 1 in which we

relative to the baseline treatment group mean, and the decline becomes statistically significant by mid-2014. Notably, 90% of ridesharing activity growth in the treated group occurs from 2014 onwards. Overall, figure 1 corroborates our regression results.

As an additional robustness check, we randomly permute ridesharing growth across Census tracts and estimate “placebo” regressions using permuted ridesharing growth as the regressor of interest. Appendix figure A6 plots the distribution of coefficients from 200 placebo regressions. The distribution of placebo coefficients is centered at zero, as expected, and the most negative placebo coefficient is still smaller in magnitude than the actual regression coefficient.

Finally, to address possible treatment effect heterogeneity, we confirm that a fully-interacted version of equation (1) yields a similar point estimate as the homogeneous version of equation (1). For this exercise, we remove aggregate time trends from the outcome and treatment and then estimate 36,446 regressions, one for each Census tract in column 2 of table 1. The average of the 36,446 coefficients is -0.161 percentage points, which is approximately the same magnitude as the single-coefficient estimate of -0.127 percentage points.¹⁹

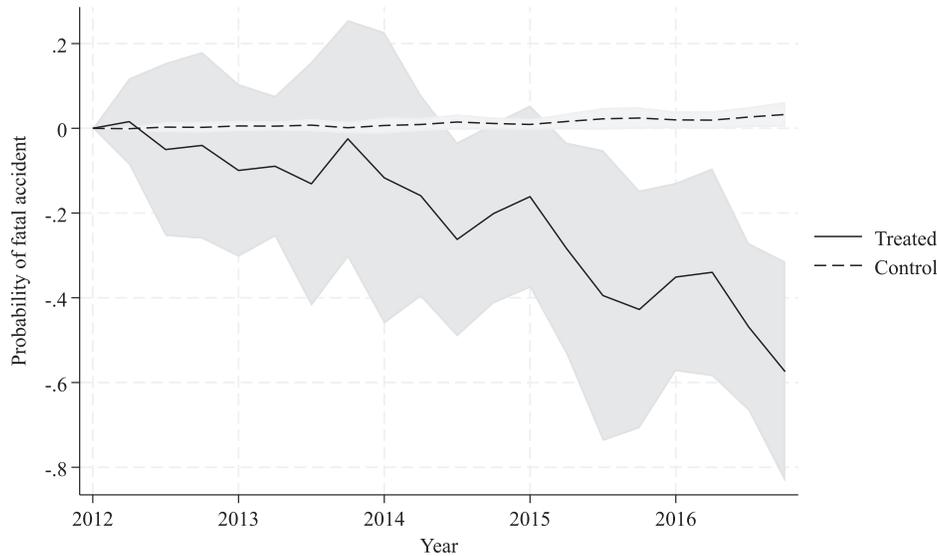
C. Alcohol-Involved Fatalities

While our main analysis focuses on all traffic fatalities, we also explore using fatalities coded as alcohol-involved as an alternative dependent variable. To measure the effect of rideshare activity on alcohol-involved traffic fatalities, we estimate versions of equations (2) and (3) that specify the

decrease or increase the threshold for defining the “treatment” group by 30%; they look broadly similar to figure 1.

¹⁹The somewhat larger average coefficient conforms with our expectations. The standard two-way fixed effects regression should weight high-growth tracts more heavily, as these tracts have higher residual treatment variation. If the safety benefits of ridesharing are concave with respect to ridership—for example, if drunk-driving trips are some of the first ones substituted away—then reweighting the effect for the average tract should increase the estimate.

FIGURE 1.—FATAL ACCIDENT TIME SERIES OF “TREATED” AND “CONTROL” TRACTS



Probabilities in percentage points. Treated tracts ($N = 2,385$) have approximately two standard deviations higher ridesharing growth than control tracts ($N = 42,035$). Gray shading indicates CBSA cluster-bootstrapped 95% confidence interval.

TABLE 4.—THE EFFECT OF UBER ON ALCOHOL-RELATED TRAFFIC FATALITIES

	(1) Drunk death	(2) Drunk death	(3) Drunk death	(4) Drunk death
Rideshare index	-0.051 (0.013) <i>0.001</i>	-0.039 (0.012) <i>0.003</i>		
Rideshare index * Night			-0.040 (0.018) <i>0.034</i>	-0.042 (0.018) <i>0.026</i>
Observations	2,186,460	1,091,760	4,372,920	2,183,520
Max tract rideshare activity	Top 50 pct	Top 25 pct	Top 50 pct	Top 25 pct
Tracts	36,441	18,196	36,441	18,196
Mean dependent variable	0.702	0.609	0.351	0.305
Mean index	0.145	0.279	0.0726	0.139

This table reports coefficient estimates from four separate least squares regressions. Specifications are identical to table 2, columns 2 and 3, and table 3, columns 3 and 4, except that the dependent variable is an indicator for alcohol-related fatalities (rather than any fatalities), multiplied by 100. Parentheses contain standard errors clustered by CBSA. FDR-control q -values in *italics*.

dependent variable as an indicator for any fatal crash with reported alcohol involvement. Columns 1 and 2 in table 4 estimate equation (2) with y_{it} specified as an indicator for any alcohol-involved fatal crash. The estimates of β are highly significant ($t = -3.9$ and $t = -3.3$). While the raw point estimates are roughly half as large as the analogous all fatalities estimates in table 2, in proportional terms they are 60% to 80% larger (7.3% and 6.4% of the dependent variable means respectively). Columns 3 and 4 in table 4 estimate equation (3), the triple-differences design, with y_{itd} specified as an indicator for any alcohol-involved fatal crash. The estimates of β are of similar magnitude as in columns 1 and 2 and remain statistically significant. Again, the raw point estimates are approximately half as large as the analogous any fatalities estimates in table 3, but in proportional terms they are almost twice as large. Overall the results are consistent with deterrence of drunk-driving being the primary mechanism through which ridesharing reduces traffic fatalities.

IV. Discussion

The coefficient estimates in table 1 represent the monthly effect of the average level of 2019 Uber activity in tracts contained in our main analytic sample (column 2 of table 1). To compute annual fatal crashes avoided due to Uber circa 2019, we multiply this coefficient by the number of sample tracts (36,446) times 12.²⁰ Then we multiply by the average number of fatalities conditional on any fatal crash occurring (1.13).

Our estimate implies that Uber saved 627 lives in 2019, a reduction of 5.2%.²¹ This calculation includes lives saved

²⁰Multiplying by 12 annualizes the estimate. We also divide by 100 since the tables multiply y_{it} by 100 for coefficient readability.

²¹Using the coefficient from column 3 of table 1 (top-25% tracts) generates similar estimates. We first normalize the index by the average level of Uber activity in these tracts. Doing so transforms the coefficient in column 3 to -0.279 . We then perform the same calculation as above and find that,

by Uber only; total lives saved by ridesharing would also include the impacts of competitors like Lyft.

To understand the economic magnitudes of these estimates, we apply the Department of Transportation VSL of \$10.9 million (\$2019, U.S. Department of Transportation, 2021). The annual life-saving benefits are \$6.8 billion. These benefits represent a mixture of internally and externally captured benefits. To approximate the share of benefits that are internal versus external, we estimate equation (1) separately for single-vehicle, multivehicle, and pedestrian/bicyclist-involved crashes. These results, reported in appendix table A4 suggest that single-vehicle, multivehicle, and pedestrian crashes account for approximately 32%, 59%, and 10% of fatalities respectively. If half the multivehicle fatalities and all of the pedestrian/bicyclist fatalities are external, then 40% of the life-saving benefits represent pure externalities.²² Of the remaining 60% of life-saving benefits, an indeterminate fraction are internalized. Many of the avoided fatalities are drunk-driving fatalities, some of which may not be the result of rational decisions.²³ Nevertheless, to be conservative we assume that drunk drivers understand and internalize the private safety benefits of using ridesharing. Under that assumption, 40% of the total life-saving benefits represent external benefits (\$2.7 billion), and the remaining 60% (\$4.1 billion) are internalized by riders.

We compare these external and internal benefits against two benchmarks: producer surplus captured by Uber shareholders and consumer surplus captured by Uber riders. Uber's market capitalization represents producer surplus captured by Uber shareholders. In 2019, Uber's average market capitalization was \$51.3 billion. After adjusting for the proportion of revenues outside the U.S. or not involving ridesharing, we estimate that domestic ridesharing accounted for \$19.2 billion of this market capitalization (see appendix A1.2 for details). To convert from stocks to flows we apply the price-earnings ratio of the S&P 500 during this period (22.7). We thus calculate that U.S. ridesharing producer surplus captured by Uber shareholders was equivalent to an annual stream of \$0.8 billion (\$19.2b/22.7). The external life-saving benefits of Uber (\$2.7 billion)—which represent pure welfare gains—thus exceed shareholder producer surplus.

in just the top 25% of tracts, Uber saved approximately 688 lives in 2019. Intuitively, the majority of lives saved are concentrated in the areas with the highest Uber ridership.

²²In principle, tort liability and insurance mandates could internalize these externalities. In practice, even when drivers are found liable, almost no drivers possess assets sufficient to cover a \$10 million VSL, and mandated levels of liability insurance are \$30,000 or less in all but one state (Insurance Information Institute, 2021).

²³While drunk individuals have impaired decision-making skills, they may choose whether to drive or rideshare at a point when they are still sober. If they fully understand the risks of drunk driving at that point, they should internalize their own safety benefits of ridesharing. The general consensus in the economic literature, however, is that the decision to drink and drive is not fully rational (Sloan, 2020); for example, drinker drivers exhibit time-inconsistent preferences (Sloan et al., 2014).

Estimating consumer surplus accruing from Uber's existence is challenging because it depends on the long-run price elasticity of demand. A recent experiment by Christensen and Osman (2021) estimates the medium-run price elasticity of demand for Uber by randomly assigning discounts of 25% to 50% to Uber users for three months. While this paper has limitations in our context—for example the three-month treatment is not truly long-run, and it was conducted in Cairo—it is to our knowledge the best available evidence on this parameter.²⁴ Christensen and Osman (2021) estimates price elasticities in the range of -4 to -8 .

A long-run price elasticity of -4 implies that consumer surplus equals 33% of total spending (appendix A2). Total spending by domestic Uber riders in 2019 was approximately \$24.7 billion (see appendix A1.2 for details). An estimate of 2019 consumer surplus is thus $\$24.7\text{b} \cdot 0.33 = \8.2 billion. The internalized life-saving benefits alone would thus represent 50% of total consumer surplus. At a price elasticity of -2 the internalized life-saving benefits still represent 21% of consumer surplus.

V. Conclusion

Previous researchers have worked hard to learn as much as possible from publicly-available information about Uber, but their analyses have yielded conflicting conclusions. Using proprietary tract-level information on Uber ridership, we find robust, large, and statistically significant negative impacts on U.S. traffic fatalities.

In 2021 in the United States, there were almost 43,000 traffic fatalities, including over 13,000 that were alcohol-related. The total economic damages, applying a standard VSL, approach half a *trillion* dollars. Understanding the factors that contribute to these deaths continues to be an important question for economists and other researchers. Our results suggest that ridesharing can play an important role in reducing these deaths, and that these benefits may represent a meaningful fraction of the total consumer surplus from ridesharing.

²⁴Cohen et al. (2016) estimates the short-run price elasticity of demand for Uber using a clever natural experiment generated by surge pricing. The short-run price elasticity of demand is not our object of interest, however; as the study points out, "If...one wanted to know how consumers would be affected if Uber disappeared permanently, a long-run elasticity would be more appropriate." (p. 21) Similarly, Goldszmidt et al. (2020) estimate the value-of-time using a field experiment with Lyft. As part of this experiment they randomly vary the surge-pricing multiplier to identify price elasticities. While the manipulation can last up to eight weeks, it applies only to surge pricing, and all changes in quantities are measured conditional on launching the app (i.e., the relevant price variation is likely the difference between the price a user expects when launching the app and the actual displayed price). Thus the price elasticity measured is primarily short-run in nature, which works well for the authors' purposes but is less relevant to our welfare calculations. The results in both papers suggest that the short-run price elasticity of demand is likely inelastic.

REFERENCES

- Abouk, Rahi, and Scott Adams, "Texting Bans and Fatal Accidents on Roadways: Do They Work? Or Do Drivers Just React to Announcements of Bans?" *American Economic Journal: Applied Economics* 5:2 (2013), 179–199. 10.1257/app.5.2.179
- Anderson, Michael, "Safety for Whom? The Effects of Light Trucks on Traffic Fatalities," *Journal of Health Economics* 27:4 (2008a), 973–989. 10.1016/j.jhealeco.2008.02.001
- Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association* 103:484 (2008b), 1481–1495. 10.1198/016214508000000841
- Anderson, Michael L., and Maximilian Auffhammer, "Pounds That Kill: The External Costs of Vehicle Weight," *Review of Economic Studies* 81:2 (2014), 535–571. 10.1093/restud/rdt035
- Ashenfelter, Orley, and Michael Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy* 112:S1 (2004), S226–S267. 10.1086/379932
- Barrios, John M., Yael V. Hochberg, and Hanyi Yi, "The Cost of Convenience: Ridehailing and Traffic Fatalities," *Journal of Operations Management* 69:5 (2023), 823–855. 10.1002/joom.1221
- Benjamini, Y., and Y. Hochberg, "Controlling the False Discovery Rate," *Journal of the Royal Statistical Society: Series B* 57 (1995), 289–300. 10.1111/j.2517-6161.1995.tb02031.x
- Berger, Thor, Chinchih Chen, and Carl B. Frey, "Drivers of Disruption? Estimating the Uber Effect," *European Economic Review* 110 (2018), 197–210. 10.1016/j.euroecorev.2018.05.006
- Brazil, Noli, and David S. Kirk, "Uber and Metropolitan Traffic Fatalities in the United States," *American Journal of Epidemiology* 184 (2016), 192–198. 10.1093/aje/kww062
- Cairncross, John, Jonathan D. Hall, and Craig Palsson, "VancUber: The Long-Run Effect of Ride-Hailing on Public Transportation, Congestion, and Traffic Fatalities," *Canadian Journal of Economics/Revue canadienne d'économique* (2025).
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," NBER working paper 32117 (2024).
- Chen, M. K., Peter E. Rossi, Judith A. Chevalier, and Emily Oehlsen, "The Value of Flexible Work: Evidence from Uber Drivers," *Journal of Political Economy* 127:6 (2019), 2735–2794. 10.1086/702171
- Christensen, Peter, and Adam Osman, "The Demand for Mobility: Evidence from an Experiment with Uber Riders," NBER working paper 31330 (2021).
- Cohen, Alma, and Liran Einav, "The Effects of Mandatory Seat Belt Laws on Driving Behavior and Traffic Fatalities," this REVIEW 85:4 (2003), 828–843.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, "Using Big Data to Estimate Consumer Surplus: The Case of Uber," NBER working paper 22627 (2016).
- DeAngelo, Gregory, and Benjamin Hansen, "Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities," *American Economic Journal: Economic Policy* 6:2 (2014), 231–257. 10.1257/pol.6.2.231
- Dills, Angela K., and Sean E. Mulholland, "Ride-Sharing, Fatal Crashes, and Crime," *Southern Economic Journal* 84:4 (2018), 965–991. 10.1002/soej.12255
- Goldszmidt, Ariel, John A. List, Robert D. Metcalfe, Ian Muir, V. K. Smith, and Jenny Wang, "The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments," NBER working paper 28208 (2020).
- Greenwood, Brad N., and Sunil Wattal, "Show Me the Way to Go Home: An Empirical Investigation of Ride-Sharing and Alcohol Related Motor Vehicle Fatalities," *MIS Quarterly* 41:1 (2017), 163–187. 10.25300/MISQ/2017/41.1.08
- Hall, Jonathan D., Craig Palsson, and Joseph Price, "Is Uber a Substitute or Complement for Public Transit?" *Journal of Urban Economics* 108 (2018), 36–50. 10.1016/j.jue.2018.09.003
- Insurance Information Institute, "Background on: Compulsory Auto/Uninsured Motorists" (2021).
- Jacobsen, Mark R., "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior," *American Economic Journal: Applied Economics* 5:3 (2013), 1–26. 10.1257/app.5.3.1
- Levitt, Steven D., and Jack Porter, "Sample Selection in the Estimation of Air Bag and Seat Belt Effectiveness," this REVIEW 83:4 (2001), 603–615.
- Morrison, Christopher N., Sara F. Jacoby, Beidi Dong, M. K. Delgado, and Douglas J. Wiebe, "Ridesharing and Motor Vehicle Crashes in 4 US Cities: An Interrupted Time-Series Analysis," *American Journal of Epidemiology* 187:2 (2017), 224–232. 10.1093/aje/kwx233
- National Highway Traffic Safety Administration, "Fatality Analysis Reporting System (FARS) Analytical User's Manual 1975–2015" (2016).
- Peck, Jessica L., "New York City Drunk Driving after Uber," CUNY Working Paper 13 (2017).
- Peltzman, Sam, "The Effects of Automobile Safety Regulation," *Journal of Political Economy* 83:4 (1975), 677–725. 10.1086/260352
- Sloan, Frank A., "Drinking and Driving" (pp. 1–31), in *Handbook of Labor, Human Resources and Population Economics* (Cham: Springer, 2020).
- Sloan, Frank A., Lindsey M. Eldred, and Yanzhi Xu, "The Behavioral Economics of Drunk Driving," *Journal of Health Economics* 35 (May 2014), 64–81. 10.1016/j.jhealeco.2014.01.005
- Tarduno, Matthew, "The Congestion Costs of Uber and Lyft," *Journal of Urban Economics* 122 (2021), 103318. 10.1016/j.jue.2020.103318
- Teltser, Keith, Conor Lennon, and Jacob Burgdorf, "Do Ridesharing Services Increase Alcohol Consumption?" *Journal of Health Economics* 77 (2021), 102451. 10.1016/j.jhealeco.2021.102451
- U.S. Department of Transportation, "VSL Update 2021—Transmittal Memo" (2021).
- Weill, Joakim, Matthieu Stiller, Olivier Deschenes, and Michael Springborn, "Researchers' Degrees-of-Flexibility and the Credibility of Difference-in-Differences Estimates: Evidence from the Pandemic Policy Evaluations," NBER working paper 29550 (2021).
- White, Michelle J., "The 'Arms Race' on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety," *The Journal of Law and Economics* 47:2 (October 2004), 333–355. 10.1086/422979
- Zhou, You, "Ride-Sharing, Alcohol Consumption, and Drunk Driving," *Regional Science and Urban Economics* 85 (2020), 103594. 10.1016/j.regsciurbeco.2020.103594