Received: 15 September 2021 | Revised: 17 August 2022 | Accepted: 19 August 2022

DOI: 10.1002/joom.1221

RESEARCH ARTICLE

WILEY

The cost of convenience: Ridehailing and traffic fatalities

John M. Barrios[1] | **Yael V. Hochberg**[2] | **Hanyi Yi**[3]

[1]Olin School of Business, Washington University in St. Louis & NBER, St. Louis, Missouri, USA

[2]Rice University & NBER, Houston, Texas, USA

[3]Carroll School of Management, Boston College, Newton, Massachusetts, USA

Correspondence
John M. Barrios, Olin School of Business, Washington University in St. Louis & NBER, St. Louis, MO, USA.
Email: john.barrios@wustl.edu

Funding information
University of Chicago Booth School of Business

Handling Editor: John Gray

Abstract

We examine the effect of the introduction of ridehailing in US cities on fatal traffic accidents. The arrival of ridehailing is associated with an approximately 3% increase in the number of fatal accidents, for both vehicle occupants and pedestrians. Consistent with ridehailing increasing road usage, we find that its introduction is associated with increases in proxies for traffic congestion and with new car registrations. Consistent with a driver quality channel, accident increases are concentrated in ridehailing-eligible vehicles and those with passenger configurations suggestive of ridehailing. Back-of-the-envelope estimates of the annual cost in human lives range from $5.33B to $13.24B. We propose various operational and policy prescriptions for the regulation of ridehailing operations that may help limit such externalities.

1 | INTRODUCTION

The adoption of mobile apps on GPS-enabled smart devices has recently led to a rise in the number and popularity of multi-sided platforms for the sharing economy. Prime examples of such platforms are online car-hailing platforms Uber and Lyft, which quickly match drivers to customers seeking a ride and have seen accelerated growth in recent years. These ridehailing (RH) services have fundamentally changed how many individuals are transported in cities and towns across the United States. On the one hand, the advent of these platforms such as Uber and Lyft have brought significant operational efficiency and convenience to consumers seeking cheap and flexible door-to-door transportation. On the other hand, the popularity of these services, and the reliance of drivers on mobile devices to both accept RH requests and navigate to the destination, present potential concerns regarding off-setting negative effects such as increases in traffic congestion and car-exhaust pollution, and distracted driving. In this paper, we consider whether RH services in fact impose negative safety externalities, and if so, how policy makers and operators of such platforms should account for them. We present evidence suggesting that such costs exist, are not trivial, and can be measured in human lives—precisely, in increased traffic accidents and traffic fatality rates.

Mass adoption of RH platforms may have multiple implications for road safety. A naïve view of the effects of RH sees it as removing drivers who would have driven themselves with their cars and replacing them with RH drivers for the same mileage. Moreover, one might argue that many of the users who substitute for being driven are often doing so because they are (or will be) inebriated or otherwise impaired (Greenwood & Wattal, 2017), such that the substitution potentially increases the quality of driving, while (in theory) holding car utilization fixed. Under this view, there will be no increase in the vehicle miles traveled (VMT), and a possible increase in driver quality. Consequently, there should be no increase in accident rates—in fact, there might even be a reduction.

This naïve view, however, ignores additional considerations presented by the substitution. First, the advent of RH platforms transforms vehicles into productive assets for individuals who now may find it lucrative to provide services as drivers. Moreover, these RH drivers have riders in their car for only a fraction of the time that they are on the road: they must drive from fare to fare, and they drive from location to location in the city seeking better fare prospects, as there is not always a fare available at their location. Furthermore, RH companies often subsidize drivers to stay on the road, even when utilization is low, to ensure that supply is quickly available. The naïve view also assumes that only those who would have otherwise driven themselves utilize RH services, which is unlikely. The convenience and lower pricing of RH apps suggest that there may be a significant number of additional riders substituting away from other modes of transportation, such as subways, buses, biking, or walking. Surveys report that fewer than half of RH rides in major metro areas actually substitute for a trip that someone would have made in a car (Schaller, 2018), and that only 39% of respondents in major metro areas would have used car-based transportation (drive themselves, carpool, or take a taxi) if RH had not been available. The rest substitute from rail, biking, walking or not traveling at all (Clewlow & Mishra, 2017).[1] This evidence runs counter to the naïve view that utilization remains fixed, and suggests that RH will increase VMT overall.[2]

Increased VMT is not the only potential effect of RH services. RH drivers connect to passengers via mobile phone apps and navigate to pickup and drop-off locations using GPS mapping capabilities built into the RH platform apps. Drivers often pick up and drop off passengers in unfamiliar locations, and on streets where no safe areas are available to pull off. Monitoring of mobile devices may lead to distracted driving and attempts to follow algorithmic routes may lead to the need to move quickly and unexpectedly across lanes of traffic, or to undertake unexpected turns (Dingus et al., 2016; Owens et al., 2018). Active interaction with passengers may also contribute to distraction. Long hours or long stretches of continuous driving may also contribute to fatigue and distraction (Dingus et al., 2016; see Scott et al., 2021 for a discussion of this issue in the trucking industry and the effects of monitoring efforts to address it). Depending on the driving quality of the average RH driver, as compared to the driving quality of former drivers who now ride, this may also lead to a change in the average quality of drivers on the road. Overall, we theorize that the introduction of RH services in the United States will lead to an increase in fatal accidents.

Using the staggered introduction of RH across US cities, we show that RH introduction in a metropolitan area leads to an economically meaningful increase in overall motor vehicle fatalities. We define the entry of RH into cities using rollout dates obtained from Uber and Lyft. We use the launch date of the first service in each city to determine the first quarter of treatment. Our outcome measures are a variety of fatal traffic accident-related measures from the Fatal Accident Reporting System (FARS) maintained by the National Highway Traffic Safety Administration (NHTSA).[3] We utilize a two-way fixed effect (city, time) generalized difference-in-differences (DD) model. Our models include location-specific linear and quadratic trends. The structure of the model allows us to control for changes in the macroeconomy, fuel costs, and improvements in vehicle technology. It also addresses city-specific conditions such as average weather patterns and city topology, a variety of city-level controls, and location-specific pre-trends in accidents that existed prior to the arrival of RH.

We report several findings. First, we document a 2%–4% increase in the number of fatal accidents and fatalities that persists throughout the week, on weekends, at night, and on weekend nights. Fatal accidents and fatalities also increase with proxies for the intensity of driver adoption of RH. When we separate accidents into those involving car occupants versus non-occupants (pedestrians, bicycle riders, etc.), we find similar increases in the number of fatal accidents, the number of pedestrians and bike riders involved in these accidents, and the number of fatalities in such accidents. Second, we provide evidence on both the quantity and quality mechanisms discussed above. We document that at the intensive margin, VMT, excess gas consumption, and annual hours of delay in traffic all rise following the entry of RH. At the extensive margin, we find a 3% increase in new car registrations, consistent with RH services creating a new productive use for vehicle ownership. Yet our original findings persist even when controlling for this increase in VMT, suggesting a quality channel as well. We further show that the increase in fatal accidents and fatalities following the introduction of RH is concentrated in vehicles eligible to serve as RH vehicles. Finally, we document that there appears to be no decrease in accidents and fatalities related to drunk driving post-RH, such that there does not appear to be a large quality improvement. Our back of the envelope calculation suggests a potential annual cost of just under $10 billion for these fatalities alone. We end the paper with a discussion of potential platform operations changes and possible regulation of RH operations that may serve to reduce the accident externality.

Our analysis offers several contributions to the operations management literature. First, we provide a more holistic view of the societal effects of RH platforms. While prior literature on on-demand service platforms in OM focused on RH price optimization problems (Bai et al., 2019; Banerjee et al., 2017; Besbes et al., 2021; Siddartha et al., 2015), market competition (Besbes et al., 2021; Yu et al., 2020), or the matching of supply and demand (Feng et al., 2020; Özkan & Ward, 2020). There is little research on the sustainability and societal impact of on-demand RH services. Our study provides one of the first large-sample investigations of the potential negative externalities of RH. Second, our study also spotlights the potential cost to the benefit of alleviating the last mile problem by RH platforms. The last mile problem is a primary concern for urban logistics and operations management (e.g., Liu et al., 2021; Qi et al., 2018). Recent studies on vehicle sharing (e.g., Benjaafar & Ming, 2020; He et al., 2017) suggest that such platforms have the potential to alleviate the last mile problem. Our work implies there could also be associated costs.

Our focus on the negative externalities of RH platforms is similar in spirit to the debate in the operations management literature on how to account for both safety and operational effectiveness in designing production systems (Pagell et al., 2015; Pagell et al., 2020). Our findings suggest the need for models of optimal operations of RH and similar driving-gig platforms that account for the externalities imposed by such operations.[4] If negative externalities are not accounted for, even if the private costs are exceeded by the private benefits for the individual user, the social costs may not be (see e.g., Beland and Brent (2018), who examine the effects of traffic congestion on first responders, and Scott et al. (2021) who explore unintended responses to IT enabled monitoring). Considering such externalities should be an important input into operations research models that analyze optimal operations (such as pricing schemes) of new technological solutions. Our analysis provides unique insights into the impacts of new app-based platform services that were conceived to increase rider convenience. While we study RH in particular, our findings likely generalize to food delivery platform services and other apps that provide services to app-based customers utilizing driver matching. We discuss a number of ways that RH operations managers can address these concerns.

Third, we contribute to the sustainable operations research literature (Bouchery et al., 2016; Kleindorfer et al., 2005) by providing empirical evidence for the impact of RH services on traffic congestion and safety. The potential theoretical effects of ridesharing services on the smart-city movement and sustainability have been discussed in the literature (Mak, 2022; Qi & Shen, 2019), but, to date, there has been little empirical evidence examining these theoretical conjectures. Our results suggest that, given RH's impact on congestion and safety, RH services may generate significant negative impacts on the smart-city movement and its sustainability in urban areas. We propose several design policies to help alleviate these negative externalities that operations leaders at RH companies and policymakers may consider.

Finally, we provide a considerable contribution and improvement on existing studies on RH's effect on traffic accidents. This literature has focused on measures of alcohol-related fatal accidents, fatalities, and citations for driving while under the influence of alcohol (Brazil & Kirk, 2016; Greenwood & Wattal, 2017; Martin-Buck, 2017), and finds mixed evidence of either a reduction or no significant change in drunk accidents/fatalities and DUIs. While at first glance, our results may appear to contradict to these studies, many of the results in these studies are driven by the selected sample period and are sensitive to accounting for pre-existing trends in accident rates, due to a 2007 change in how FARS classified "drunk accidents." The exception to this is Greenwood and Wattal (2017), who explore drunk driving fatalities in California counties after the introduction of certain Uber services. The limitation of a single state sample, and the unit of observation, which confounds urban (where our effects are concentrated) and rural areas in the same observations, likely drive the difference between their results and ours. In contrast to Greenwood and Wattal (2017), our analysis of alcohol-related traffic fatalities over the entire United States and a longer sample period suggests no consistent significant change in drunk-driving related fatal accidents as a result of RH. Notably, our research design does not solely focus on fatalities resulting from drunk driving or alcohol consumption, as RH is likely to affect accident fatalities more generally. Rather, we focus on total fatal accidents, using a broad sample and exploiting the introduction of both Uber and Lyft as well as their various services. While RH may indeed displace some drunk drivers, our estimates suggest that overall accident rates and fatalities increase in the wake of RH introduction.

The remainder of this paper is structured as follows. Section 2 summarizes the relevant literature and develops a set of formal hypotheses. Section 3 describes our sample, data, and variables. Section 4 describes the research design and causal identification of effects. Section 5 presents the main analysis and empirical findings. Section 6 concludes with a discussion of back-of-the envelope costs, managerial and policy implications, theoretical contributions, limitations, and directions for future research.

2 | LITERATURE REVIEW AND HYPOTHESES

2.1 | On-demand service platforms and ridehailing

Facilitated by the development of internet technology, the past decade has seen the rise of two-sided markets for sharing economies/platforms. On-demand ride-hailing services like Uber and Lyft leverage information and mobile technologies to coordinate and connect independent drivers to passengers in real-time. These on-demand ride services benefit from scale economies allowing them to expand their operations in a time-efficient and cost-effective manner owing to network effects and because they can operate without owning physical assets and regular drivers. Such on-demand service platforms have led to a burgeoning literature in operations management.

A growing literature in this arena explore RH platforms in particular. The bulk of this work focuses on optimal design, operations, and pricing for RH platforms. For example, Besbes et al. (2021) examine a two-dimensional framework in which a platform selects prices for different locations, and drivers respond by choosing where to relocate based on prices, travel costs, and driver congestion levels. They derive an optimal solution where the platform applies different treatments to different locations. In some locations, prices are set so that supply and demand are perfectly matched; over congestion is induced in other locations, and some less profitable locations are indirectly priced out. Banerjee et al. (2015) argue that static pricing strategies can perform well in the RH setting. In contrast, Cachon et al. (2017) show that surge pricing benefits both providers and consumers when providers have a high opportunity cost. Hu and Zhou (2020) examine matching issues between trip prices and driver wages. They find that a commission contract can be optimal when both demand and supply curves are affine functions having common price and wage sensitivity. Bimpikis et al. (2019) examine the spatial transition of a ride-sharing network and describe the value of spatial price discrimination for the platform. Specifically, they show that the pricing policy that uses a fixed commission rate could result in a loss of profit in case of heterogeneous demand patterns across different locations.

Another part of the literature speaks to optimization of specific performance metrics for RH platforms and alternatives to pricing that can be used to match supply and demand. Afeche et al. (2018) study the problem of matching self-interested drivers with riders, showing that it may be optimal to purposefully reject demand at low demand locations to induce repositioning rivers to high-demand locations. Banerjee et al. (2018) explore optimal assignment policies in closed queuing network models similar to RH systems. Braverman et al. (2019) demonstrate the benefit of using fluid-based optimal routing platforms for empty cars in the network to maximize network utility. Feng et al. (2020) compare customers' average waiting time under the on-demand matching mechanism relative to traditional street-hailing systems. They find that the on-demand matching mechanism can result in lower efficiency. Özkan and Ward (2020) propose dynamic matching policies based on a continuous linear program to improve the efficiency of matches. Finally, Benjaafar et al. (2022) examine the labor welfare of on-demand service platforms, showing that both average labor welfare and agent workload are non-monotonic in the labor pool size. While these studies focus on optimizing a single platform's operation from the company's and its customers' perspective, our theoretical framework incorporates the interaction of ride-hailing companies and the general public's safety.

Our paper contributes to this literature by advocating the need for models of optimal operations in RH platform should account for the externalities imposed by such operations. An ongoing discourse in the management literature revolves around how managers and policy makers should respond to contradictory demands. In the operations management literature, discussion of tradeoffs dates back to Skinner (1996) and persists even today. This literature approaches tradeoffs from a number of perspectives, including that of whether such tradeoffs exist (e.g., Ferdows & De Meyer, 1990; Rosenzweig & Easton, 2010), and how managers should balance competing objectives (e.g., Ashforth & Reingen, 2014; Birkinshaw & Gupta, 2013; Pagell et al., 2015; Smith & Lewis, 2011). Similarly, while discussion of externalities of new technologies often focuses on positive externalities that benefit society, some technologies may also impose negative externalities. When new technologies are introduced in markets that account for these externalities, they often induce competition with existing products and services that enhance welfare. If these negative externalities are not accounted for, even if the private costs are exceeded by the private benefits for the individual user, the social costs may not be (see e.g., Beland and Brent (2018), who examine the effects of traffic congestion on first responders, and Scott et al. (2021) who explore unintended responses to IT enabled monitoring). It is the sum of social and private costs as compared to the sum of social and private benefits that is key to welfare effects. The consideration of such externalities should be an important input into operations research models that analyze optimal operations (such as pricing schemes) of new technological solutions that try to make firm operations sustainable.

Our work also provides insights that may inform operations managers at on-demand platform companies of the sort we examine and local policy makers seeking to reduce traffic congestion and fatalities. On the one hand, on-demand platforms for RH present multiple benefits for multiple stakeholders. RH services are often more convenient and cheaper for consumers than alternatives such as taxis. For drivers, RH platforms offer an opportunity for flexible gig employment. These platforms, however, also raise concerns regarding social costs in arenas such as labor laws, safety, traffic congestion, and effects on the taxi industry (Yu et al., 2020). Optimal regulation of such services remains unclear. Much of the literature exploring regulation of on-demand platforms to date has concentrated in the legal literature (e.g., Posen, 2015). Recently, operations management researchers have begun to explore aspects of such regulation. Yu et al. (2020) explore the impacts of a Chinese regulatory change limiting the number of registered RH drivers. They utilize a multi-stakeholder, two-period framework to evaluate the impact of such policies on a variety of objectives associated with different RH stakeholders, concluding that in the absence of regulation, RH may have the effect of driving the taxi industry out of the market. In the last section of this paper, we discuss a number of potential approaches for RH operations managers and local policy makers that may help reduce the negative externalities arising from potential increased VMT and decreased driver quality brought about by RH services.

Finally, the literature in operations has explored the effects of other types of transportation interventions that are poised to come online in coming years. Ostrovsky and Schwarz (2019) explore the economics of autonomous vehicles and the interplay with carpooling and road pricing, and Naumov et al. (2020) explore the potential unintended consequences for public transport systems and traffic congestion from the use of automated vehicles.

2.2 | Transportation safety, vehicular accidents, and traffic congestion

A growing literature in operations management, economics, and public policy, explores issues related to transportation, vehicular accidents, and traffic congestion. Traffic congestion is one of the most prominent issues plaguing major metro areas. It is estimated that the value of time and fuel expenditures associated with congestion in the United States was close to $121 billion in 2011 (Schrank et al., 2011). While the direct costs of traffic are primarily due to lost time and reliability, there is research using survey data that also links traffic congestion to adverse mental health outcomes, including stress and aggression (Gee & Takeuchi, 2004; Gottholmseder et al., 2009; Hennessy & Wiesenthal, 1999). For example, pollution from traffic also negatively affects children's contemporaneous health (Knittel et al., 2016), and has been documented to have a long-run effect on mortality among the elderly (Anderson, 2014). A robust literature further discusses potential mitigation strategies for several other traffic congestion externalities, including pollution, health, and housing prices. Currie and Walker (2011) show that the traffic reduction stemming from the introduction of an electronic toll (E-ZPass) reduced vehicle emissions near highway toll plazas, which reduces prematurity and low birth weight among mothers near the toll plaza. Beland and Brent (2018) examine traffic congestion and emergency response times. They find that traffic slows down fire trucks arriving at the scene of an emergency and increases the average monetary damages from fires, highlighting the negative effect of traffic on a critical public good that relies on well-functioning roads. Ossokina and Verweij (2015) exploit a quasi-random experiment that reduced traffic congestion on certain streets in the Netherlands and find that the decrease leads to an increase in house prices. Other related work includes Duranton and Turner (2011), who explore "the fundamental law of road congestion," concluding that providing additional roads is unlikely to relieve congestion, documenting the positive relationship between additional lane kilometers and subsequent VMT.

A second related literature examines the relationship between cellphone and smartphone usage and traffic accidents. RH drivers interact with smartphone-based apps regularly, which may affect the quality of their driving. Dingus et al. (2016) show that distracted driving attributable to smartphones is increasingly the cause of motor vehicle crashes. Owens et al. (2018) show that visual-manual smartphone tasks (such as dialing, browsing, texting, reaching for phone, and interacting with apps) are associated with higher incidence of crashes. These studies suggest the possibility for distracted driving by RH drivers, which may lay a foundation for the quality effects we observe in our analysis.

Finally, a third related literature in operations explores how monitoring technologies may impact safety. A number of studies demonstrate that monitoring can decrease illicit or unwanted behavior (e.g., Banerjee et al., 2018; Di Tella & Schargrodsky, 2003; Duflo et al., 2012; Nagin et al., 2002). To the extent that measures to address the patterns we uncover in our analysis may consist of implementation of monitoring technologies, consideration must be given to the documented displacement of illicit behavior raised in the operations management literature. For example, while monitoring technology has been shown to affect safety in the trucking industry (Cantor et al., 2009; Hickman & Hanowski, 2011; Miller et al., 2018), more recent work suggests that electronic monitoring may lead to other unwanted behaviors that call into question the cost-benefit calculations regarding these technologies (Scott et al., 2021). This literature relates closely to certain policy prescriptions in the final section of this paper, where we propose a number of possible measures to mitigate the effects we document.

3 | HYPOTHESIS DEVELOPMENT

We begin by outlining why RH may lead to an increase in fatal traffic accidents and fatalities. Here, we verbally explain the intuitions and state the associated hypotheses we test.[5] Fatal accidents in a given location can be thought of as being affected by two primary forces: (i) the number of VMT in the location (quantity channel), and (ii) the driving quality of the average driver on the road (quality channel). We hypothesize that:

> **Hypothesis 1A.** The introduction of RH in a city will lead to an increase in fatal accidents and fatalities.

While the VMT generated by a potential rider driving themselves to a destination may be offset by the VMT generated if a RH driver drives them, RH still introduces "between driving" (between fares, waiting for fares, going from fare location to fare location) in a city. Consequently, RH leads to the introduction of additional vehicle miles in the form of VMT driven by RH drivers driving from fare to fare or between potential fare-generating neighborhoods.[6] Unless there is an improvement in driving quality, this increase in VMT is likely associated with a corresponding increase in accidents.[7] Furthermore, this increase should be increasing in the use or adoption of RH, and thus:

> **Hypothesis 1B.** Fatal accidents and fatalities post RH introduction should be increasing in RH adoption intensity.

We expect that the effects of RH will be larger in urban areas than in suburban or rural areas. In dense urban areas, the number of people per mile and per mile of road is higher, and traffic congestion is typically elevated to begin with. Any increases in VMT usage may have a bigger impact on congestion and accidents overall, due to the high starting level. Moreover, in urban areas and in areas with high usage of cars, the effects of driving behavior such as pulling across lanes, stopping suddenly to pick up or drop off a passenger, and so forth, may be more likely to have a higher chance of causing an accident. Urban areas may have fewer locations for safe pull-over on city streets for pickup and dropoff, and may have higher numbers of pedestrians and bicycle riders, who may be impacted by what happens on the roads.

Additionally, people are more likely to substitute public transit trips with point-to-point RH trips when they require less waiting time, are faster, and are more convenient (Babar & Burtch, 2020; Kong et al., 2020). This substitution increases as trip length decreases which is more likely in urban areas where the average RH trip length is shorter (Ewing & Hamidi, 2014). Furthermore, as the cost per mile of RH is higher than that of public transit, the shift from public transit to RH is more likely to occur for shorter trips (Babar & Burtch, 2020; Hall et al., 2018), which are more common in urban areas. This substitution in urban areas may further increase RH VMT and between ride VMT, thus, increasing congestion and the potential for accidents and fatalities.

Finally, road usage is typically lower in sparsely populated rural areas, and the first and last mile problem is more serious, given the absence of large transit centers and the significant land-use separation (Qi et al., 2018). As a result, RH may serve primarily to complement existing public transit options (e.g., airports, train stations), helping to address the first and last mile problem (Babar & Burtch, 2020; Fahnenschreiber et al., 2016; Qi et al., 2018; Song & Huang, 2020). This would result in less substitution from other modes of transport that might increase VMT in an urban setting. Moreover, due to the generally lower congestion and road usage, RH usage, and its impact on accidents, should be moderated in these settings. We thus hypothesize that:

> **Hypothesis 2.** The effects of the introduction of RH on accidents will be larger in more urban areas.

Next, we focus on the mechanisms for the increase in fatal accidents. Given our discussion above, on the quantity side, we posit that the introduction of RH will lead to more road usage:

> **Hypothesis 3A.** The introduction of RH in a city will lead to an increase in VMT (extensive margin).

Moreover, the introduction of RH platforms also allows individuals to transform their vehicles into productive assets. Thus, individuals may find it more lucrative to own a vehicle and provide services as drivers. Therefore, we posit that:

> **Hypothesis 3B.** The introduction of RH in a city will lead to an increase in new car registrations (intensive margin).

Moreover, to the extent that riders also substitute away from non-car modes of transportation, such as including walking, biking, and, more importantly, public transportation (Clewlow & Mishra, 2017), the VMT generate simply by the rides themselves might also exceed the personal driving it displaces, before we even account for between fare VMT. This suggests another moderating hypothesis: that the effect may be larger in cities where ex ante usage of public transportation or carpooling is highest, as riders substitute away from these modes of transport to RH services (Babar & Burtch, 2020). Thus:

> **Hypothesis 3C.** The effects of the introduction of RH in a city will be higher in cities with high ex ante use of public transportation and carpooling.

Assessing the effects of the introduction of RH on the quality of the average driver on the road is less straightforward. On the one hand, the people substituting into a RH vehicle, rather than driving themselves, may be low quality drivers (impaired or unskilled or may just prefer not to drive), but they may be high quality drivers who simply dislike driving. On the other hand, there is no guarantee that the driver who substitutes for them is of higher quality. For example, lower quality drivers, who in the absence of compensation may not have driven, now have an incentive to drive. RH VMT is also different from self-driving VMT, as RH drivers often stop in random locations mid-street to pick up or drop off a rider. This type of haphazard dropoff and pickup stoppage may also lead to additional accidents (even if the accidents do not involve the RH vehicle).[8] Furthermore, RH drivers connect to passengers via mobile phone apps and navigate to pickup and dropoff locations using GPS mapping capabilities built into the RH platform apps. Drivers are often picking up and dropping off passengers in unfamiliar locations, and on streets where no safe areas are available to pull off. Monitoring mobile devices may lead to distracted driving, and attempts to follow algorithmic routes may lead to the need to move quickly and unexpectedly across lanes of traffic, or to undertake unexpected turns (Dingus et al., 2016; Owens et al., 2018). Active interaction with passengers may also contribute to distraction. Long stretches of driving may also contribute to fatigue and distraction (Dingus et al., 2016). Overall, we would thus expect a reduction in average driver quality:

> **Hypothesis 4.** The introduction of RH in a city will lead to a decrease in driving quality.

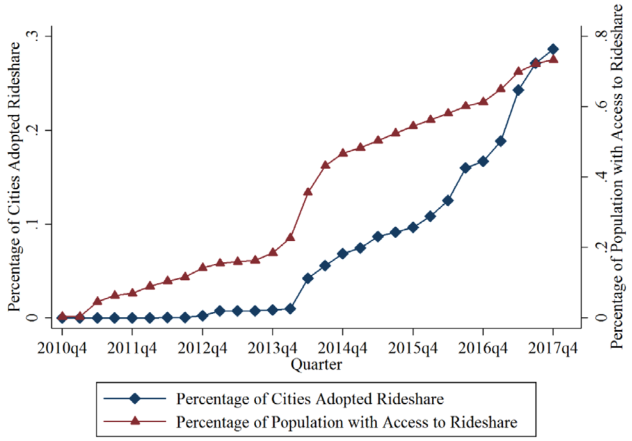

FIGURE 1 Ridehailing diffusion. The figure graphs the introduction of ridehailing across US cities and population. The sample contains all census incorporated places in the United States. The red (navy) line represents the percentage of population (cities) that adopted ridehailing in each quarter between the fourth quarter of 2010 and the fourth quarter of 2017.

4 | DATA AND VARIABLES

We construct a sample which includes all US incorporated places that had a population of 10,000 or more in 2010. Our sample runs from 2001 to 2016. Our results are all robust to the use of alternative sample windows. We utilize Census Bureau data on incorporated places. The incorporated places include all US self-governing cities, towns, villages, and boroughs.[9] Throughout, we refer to these types of places interchangeably as "cities" or "locations." We measure observations at the quarterly level. The final sample consists of 190,080 quarterly observations on 2970 "places," 1199 of which adopt RH prior to 2017. RH diffusion is depicted graphically in Figure 1. Diffusion begins slowly and accelerates rapidly after 2013.

4.1 | Fatal accidents

We obtain data on accidents involving at least one fatality ("fatal accidents") from the National Highway Traffic Safety Administration (NHTSA) Fatal Accident Reporting System (FARS). To qualify as a FARS case, a crash must involve a motor vehicle traveling on a traffic way customarily open to the public and must have resulted in the death of a motorist or a nonmotorist within 30 days of the crash. Importantly, the data identify whether any drivers involved are under the influence of alcohol. We aggregate the incident-level FARS data into quarterly totals for each place/city. When the data contain geographic coordinates, we use Google Map's Geocoding API service to determine the corresponding place/city. When the coordinates are unavailable, we use the city and state identifier codes to assign observations to the appropriate place. In 98% of FARS's observations, geographic coordinates are present, and we successfully match more than 99% to incorporated places.

We construct various measures of accident volume from the FARS data. Total Accidents is the total number of fatal accidents according to the definition used by NHTSA. Total Fatalities is the total number of fatalities across all fatal accidents. We further classify accidents based on their time of occurrence: (i) weekday: Monday through Thursday; (ii) weekend: Friday through Sunday; (iii) night: after 5 PM and before 2 AM; (iv) Friday and Saturday night: after 5 PM and before 6 AM on Friday and Saturday. Additionally, we further separate out accidents involving pedestrians and calculate three measures of pedestrian-involved accidents. Pedestrian-Involved Accidents is the number of fatal accidents involving at least one pedestrian. Pedestrian-Involved Fatalities is the total number of fatalities in all accidents involving at least one pedestrian. Finally, Pedestrians Involved in Fatal Accidents is the total number of pedestrians involved in fatal accidents. When we refer to accident "rates," these are defined as the number of accidents per 100,000 people or the number of accidents per billion city VMT, as indicated.[10]

4.2 | Ridehailing launch and driver enrollment intensity

We obtain city-level RH launch dates from Uber and Lyft. These launch dates are merged with the Census Bureau's incorporated place data from 2010. Uber and Lyft declined to provide data on driver enrollment and usage for this project. Prior literature, however, has shown that actual driver uptake and Google trends for searches for RH keywords are strongly correlated (Cramer & Krueger, 2016). Accordingly, Google searches for the terms "Uber," "Lyft," and "rideshare" serve as our proxies for driver adoption. While prior research interprets these search trends as proxies for driver adoption, another way of looking at these measures is as a measure of popularity of RH services. Either interpretation would be reasonable, and would not change our inferences.

Google searches are performed using the Google Health Trends API. Google Health Trends provides the search frequency for specific search terms relative to total search volume in a specific geographic area over a specific time period. The result is probability of a search for the corresponding term. This allows us to conduct comparisons across time and locations. We compute these probabilities for each Nielsen Designated Market Area (DMA) at a monthly frequency from 2004 to 2016. Data is aggregated by quarter. We use a Nielson-provided crosswalk to match to our sample of places.

4.3 | Other data

We use yearly city population and population density estimates from the US Census. County income per capita at the annual level is obtained from the Bureau of Economic Analysis (BEA). Household vehicle ownership and means of transportation to work at the city level are gathered from the 2010 American Communities Survey. A reasonable concern might be that Uber and Lyft specifically chose cities to enter based on smartphone adoption trends, and any increase in accidents we document could be due to increased levels of smartphone adoption in those specific cities relative the ones not entered by RH companies, with smartphone usage leading to distracted driving, which in turn leads to increased accidents. While data on smartphone adoption by city is not publicly available, smartphone adoption is known be highly correlated with per capita income, which we thus include as a control variable in our models.[11]

To explore mechanisms that may drive any change in accident rates upon arrival of RH, we use a variety of data sources. We obtain monthly data on new car registrations by zip code from Polk Automotive. We aggregate the data at city and quarter level using UDS Mapper's zip code-to-ZCTA crosswalk[12] and Census' ZCTA-to-place crosswalk. We obtain estimates of city and freeway VMT, total annual excess fuel consumption, and total annual hours of traffic delay for a sample of 101 urban areas from the Texas A&M Transportation Institute Urban Mobility Scorecard, covering the period of 1982–2014. Of the 101 urban areas covered by TAMU in their report, 99 fall into our sample of continental US cities. For a set of tests regarding road use and driver quality, we use the census's urban area-to-place crosswalk to aggregate our main sample at urban area and annual level to merge the information with TAMU's dataset.[13]

4.4 | Summary statistics

The summary statistics for our sample are in Table 1. Sample places have average population of 54,500, have an income per capita of $39,720, and a population density of roughly 3000 people per square mile. Prior to the arrival of RH, 2.96% of residents in our average city/place used public transportation to commute, 10.6% commuted by carpool, and 33% owned vehicles. The average city in

TABLE 1 Summary statistics

	Mean	**Median**	***SD***
Population (thousands)	54.50	23.51	199.99
Income per capita (thousands $)	39.72	37.47	12.17
Population density	2999.28	2169.70	3159.57
Carpool usage	10.62	10.05	3.98
Public transportation usage	2.96	1.19	4.96
Household vehicle ownership (thousands)	32.72	15.45	80.62
New car registration	670	264	2340
Accident rate	3.49	0.96	5.66
Fatality rate	3.84	0.98	6.51
Pedestrian-involved accident rate	0.58	0.00	1.80
Pedestrian-involved fatality rate	0.59	0.00	1.86

Note: Our sample includes 190,080 observations from 2001 to 2016 on 2970 census incorporated places. The sample is at a quarterly level. Population density is defined as population per square mile. Carpool usage measures the percentage of population commuting to work using a carpool. Public transportation usage measures the percentage of population commuting to work using public transportation. Household vehicle ownership measures the total number of available vehicles in households. New car registration measures the total number of new vehicle registrations. All rates are measured per 100,000 population. Accident is the number of fatal accidents, according to the definition used by NHTSA. Fatality is the total number of fatalities across all fatal accidents. Pedestrian-involved accident is the number of fatal accidents involving at least one pedestrian. Pedestrian-involved fatalities is the total number of fatalities in all accidents involving at least one pedestrian.

our sample had 670 new car registrations per year. The table further presents summary statistics on rate (per 100,000 population) of accidents. We present statistics for total accidents and fatalities and total pedestrian-related accidents and fatalities. Pedestrian accidents and fatalities are approximately 20% of the total.

5 | RESEARCH DESIGN AND CAUSAL IDENTIFICATION

We are interested in estimating the causal effect of introducing RH platforms into a city on road safety. Ideally, we would randomly assign RH to cities and then compare outcomes between treated and control cities. Instead, we utilize the fact that RH was rolled out gradually at different times in different geographies. Of course, the cities in which RH companies' choose to launch first are unlikely to be randomly selected. However, it is also unlikely that RH platforms were specifically selecting new cities based on trends in fatal accident occurrence. Rather, in the Online Appendix, we show that RH companies appear to have been selecting cities with higher populations and higher per capita income, and empirically, we observe no significant relationship between trends in fatal accidents and entry order.[14] The lack of relationship between RH entry order and accident trends addresses the most critical concern for identification of effects in a difference-and-differences specification, which then allows us to implement standard models for estimating difference-in-differences effects with staggered adoption of treatment.

Specifically, we use a generalized staggered entry difference-in-differences (DD) approach that takes advantage of the staggered rollout of RH. In contrast to the standard DD framework, which compares a set of treated and untreated units to each other pre- and post- a single treatment event, the generalized DD methodology compares a treated city post-treatment to itself during the pre-treatment period, as well as to other not-yet-treated cities (even if they will be treated eventually) in that period. Stated differently, our control group is not restricted to cities that never adopt RH platforms. In fact, for any given time period t, our model implicitly takes as the control group all cities in that period that did not also adopt RH in that particular period (Bertrand & Mullainathan, 2001). We denote cities by c and denote time by t. We estimate variation on the following model:

$$\log(1+\text{accidents}_{t,c}) = \propto_c + \gamma_t + \beta' X_{t,c} + \theta_c t + \delta \text{POST} \times \text{TREATED}_{t,c} + \varepsilon_{t,c}, \quad (1)$$

where $\text{accidents}_{t,c}$ is our measure of accidents in city c in quarter t, $\text{POST} \times \text{TREATED}_{t,c}$ is a dummy variable for whether the city c has adopted RH service by time t, $\propto_c$ is a city fixed effect, γ_t is quarter-year fixed effect, $X_{t,c}$ is a vector of time-varying, city specific control variables, and $\theta_c t$ is a city-specific linear time trend. Our *SE*s are robust and clustered at the city level.

This design fully controls for fixed differences in accident incidence between ever-treated and never-treated cities via the city fixed effects. The time fixed effects control for aggregate fluctuations in accidents over time. As a result of the inclusion of the city and time dummies, the typical main Post_t and Treated_c are absorbed by the fixed effects structure. Specifically, the variable equivalent to Post_t in a standard differences-in-differences model is absorbed here by the quarter-year time fixed effects, and Treated_c is absorbed by the city fixed effects. We further improve the estimation by including city-specific time trends (e.g., Autor, 2003), to capture city-specific dynamics in accident incidence over time unrelated to the treatment.

A number of additional checks must be performed to validate the assumptions underlying the generalized DD framework. For example, a critical assumption is that the control group acts as a valid counterfactual—that in the absence of treatment, the treatment group would have followed the same trend as the control group, known as the "parallel trends" assumption (Angrist & Pischke, 2008). While this assumption involves a counterfactual and is therefore impossible to test explicitly (Holland, 1986; Imbens & Wooldridge, 2009), we can observe whether treatment and control groups moved in parallel prior to the mandate. For generalized DD models, the standard test consists of graphically presenting the difference-in-differences estimators pooled in event time relative to the period before the treatment is enacted (with each dot representing annual-coefficients) for the 10 years preceding and 4 years following RH adoption for total accidents and total fatalities. In both panels, we expect to see the counterfactual treatment effects in the pre-RH periods are statistically indistinguishable from zero, providing support for our inferences (parallel trends in the pre-period).

An additional check addresses the fact that in staggered adoption differences-in-differences models, especially later in the treatment periods, negative weights can be assigned to treatment periods (De Chaisemartin & d'Haultfoeuille, 2020; Goodman-Bacon, 2018). In theory, negative weights on our treated observations are not a cause for concern when the treatment effects (dose response) are homogeneous. When the treatment effects are homogeneous, the two-way fixed effects model is correctly specified because the dose–response relationship between the residualized outcome variable and the residualized treatment variable is linear (De Chaisemartin & d'Haultfoeuille, 2020; Goodman-Bacon, 2018).

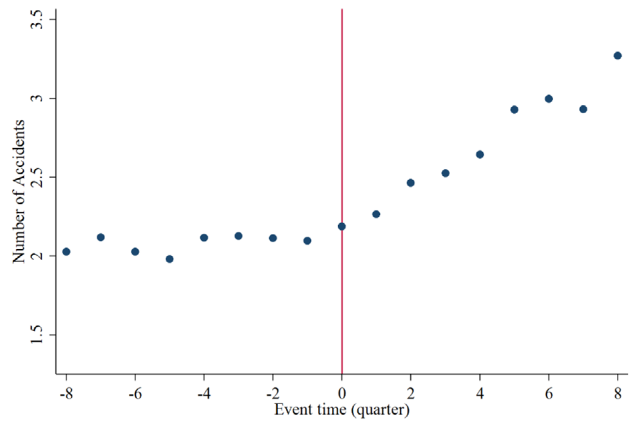

FIGURE 2 Average accidents for treated cities in event time. This figure graphs the average level of accidents for treated cities in event time. The quarter of ridehailing entry is indicated by the red vertical line at event time zero.

We have no water-tight way of testing for homogeneous treatment effects here, however, and the two-way fixed effects estimator can be biased when treatment effects are not homogeneous. As explained in Goodman-Bacon (2018) and De Chaisemartin & d'Haultfoeuille, 2020, negative weights occur when already-treated units act as controls, leading to changes in their treatment effects over time getting subtracted from the DD estimate. This negative weighting only arises when treatment effects vary over time. This would then typically bias the regression DD estimates away from the sign of the true treatment effect. In our setting, negative weights would thus lead to an attenuation bias in the estimate, if anything. Moreover, our models contain an extensive group of never-treated cities. As Goodman-Bacon (2018) explained, having a sufficiently sizeable never-treated group combined with a long enough pre-treatment period helps reduce the importance of negative weights in the treatment group, because the models utilize more variation from the never-treated observations.

Nonetheless, further testing to help assuage concerns about negative weights is warranted. In models such as ours, negative weights will tend to occur in early-adopter cities (where the city-level treatment mean is high) and in later quarters (when the quarter-level treatment mean is also high). We can therefore assess the sensitivity of our inferences by using various sample cuts that change the variation in these later adopters to determine whether our inferences would change. To do this, we employ a variety of sample periods in our robustness tests.

Second, we can formally assess negative weighting using the Goodman-Bacon decomposition, which allows us to examine the treatment effects when only using early adopters. Since the betas in staggered entry difference-in-differences models are weighted by a combination of the

TABLE 2 Hypothesis testing: Hypothesis 1A: Effect of ridehailing on traffic safety.

Panel A: Overall effect

	Log (1 + Total accidents)			Log (1 + Total fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated$_{t,c}$	0.0191***	0.0360***	0.0332***	0.0183***	0.0360***	0.0335***
	(0.0062)	(0.0074)	(0.0091)	(0.0065)	(0.0077)	(0.0095)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	190,080	190,080	190,080
R2	0.61	0.62	0.63	0.60	0.60	0.61

Panel B: Coarsened exact matching

	Log (1 + Total accidents)			Log (1 + Total fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated$_{t,c}$	0.0214**	0.0319***	0.0398***	0.0195*	0.0321***	0.0383***
	(0.0099)	(0.0108)	(0.0129)	(0.0102)	(0.0112)	(0.0137)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	125,888	125,888	125,888	125,888	125,888	125,888
R2	0.57	0.58	0.59	0.56	0.57	0.58

Panel C: Count models, and accident/fatality rate

	Poisson		Negative binomial		Accident and fatality rate	
	(1) Total accidents	(2) Total fatalities	(3) Total accidents	(4) Total fatalities	(5) Accident rate	(6) Fatality rate
Post × Treated$_{t,c}$	0.0485***	0.0510***	0.0481***	0.0479***	0.2219***	0.2222***
	(0.0148)	(0.0152)	(0.0155)	(0.0139)	(0.0669)	(0.0749)
Empirical model	Poisson	Poisson	Negative Binomial	Negative Binomial	OLS with city linear trend	OLS with city linear trend
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	188,864	188,864	188,864	188,864	190,080	190,080
R2	N/A	N/A	N/A	N/A	0.40	0.38

Note: The table reports estimates from generalized difference-in-difference regressions. *SE*s, clustered by city, are reported in parentheses. We list the dependent variables at the top of each column. Post × Treated$_{t,c}$ is an indicator that takes the value of one if city c adopted at least one ridehailing service at time *t*, and 0 otherwise. Panel A presents the overall effect of ridehailing on two traffic safety measures. Total accidents are the number of fatal accidents according to the definition used by NHTSA. Total fatalities are the total number of fatalities across all fatal accidents. Panel B presents effects estimated on a sample of matched cities. We use coarsened exact matching to match treated and control cities based on two dimensions: income and population. In Panel C columns (1) through (4), we run count models (Poisson and negative binomial) rather than OLS regressions. In Panel C columns (5) and (6), we replace the outcome variables with accident and fatality rates (per 100,000 people). We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

treatment timing and number of pre-periods, the Goodman-Bacon decomposition into early versus late adopters and early-treated versus never-treated allows us to look at the treatment effect without being influenced by negative weights.

6 | EMPIRICAL RESULTS

6.1 | Main analysis

We begin our analysis by examining changes in the level of accidents in the treated cities around the introduction of RH. Figure 2 plots the raw quarterly average accidents over event time in RH cities. At the time of RH initiation—time zero—we see a distinct break in the trend of accident incidence in the RH cities: accident numbers begin to rise sharply, relative to the pre-event time trend. In other words, we observe a clear structural break in the pattern in the raw data for treated cities at the time of treatment.

We investigate this increase formally using our staggered entry generalized DD specification. In our estimations, we use several measures for $\text{accidents}_{t,c}$. In Panel A of Table 2, we employ our two main measures of total fatal accidents. Columns (1), (2) and (3) use total accidents, and columns (4), (5) and (6) use total fatalities. These variables take non-zero values for about 48% the sample (92,396 out of 190,080 observations). The first column of each set reports estimates without including the city-specific linear time trend or quadratic trend, the second column of each pair adds the city-specific linear trend, and the third column further adds a city-specific quadratic trend. The table reports the coefficient on $\text{POST} \times \text{TREATED}_{t,c}$, which is our variable of interest. The results are robust to using count models (Poisson and negative binomial) which account for zeros in our dependent variable, as shown in Panel C.

Consistent with the raw data plotted in Figure 2, and supporting Hypothesis 1A, we observe a consistently positive and significant coefficient on the $\text{POST} \times \text{TREATED}_{t,c}$ variable for both total accidents and total fatalities, regardless of specification. Before accounting for the location-specific time trend, the effect ranges in magnitude from an increase of 1.83% in total fatalities (column [4]) to 1.91% increase in total fatal accidents (column [1]). Once we include the location-specific time trend, the magnitudes of the increase are approximately 3.6% for both measures of accidents. The magnitude of the effect decreases slightly, to a little over 3.3%, once we include the quadratic trend. In the Online Appendix, we further break out the effect of introducing each element of our main specification in turn (Table B1). Figure B1 graphs the coefficient and associated confidence interval for the variable $\text{POST} \times \text{TREATED}_{t,c}$, first itself, then adding year fixed effects, city fixed effects, and the city-specific linear trend, and city-specific quadratic trend, each in turn. For brevity, in much of the remainder of the reported analysis, we report only the DD models with city-specific linear trends, but all results remain robust to adding quadratic trends as well, and all model specification estimates are available upon request.

Our dependent variable is characterized a by a large proportion of zero observation (about 52% of city-quarter observations have zero accidents). While the use of the $\log(1+y)$ transformation is common in the literature to address this issue, it has shortcoming, especially when the proportion of zero observations is high. We take several steps to address this concern. First, we use count models to verify our inferences. Our results are robust to both Negative Binomial and Poisson models, and we report these results in Table 2 Panel C Columns (1) through (4). Second, we use accident and fatality rates (defined as per 100,000 population) as outcome variables, and still find a statistically significant and economically meaningful increases in accidents and fatalities, reported in Table 2 Panel C Columns (5) and (6). Third, our results are robust to other alternative formulations of the outcome variable that reduces sparsity, such as bottom coding the outcome variables at the 5th percentile to (Online Appendix Table B2) or employing the inverse hyperbolic sine transformation (Online Appendix Table B3).

Many indicators suggest that our documented effects may adjust over time. Cook et al. (2018) note that, even in the relatively simple production of a passenger's ride, experience is valuable for drivers. Similarly, Haggag et al. (2017) show that experience is important for taxi drivers. At the same time, not all learning-by-doing is necessarily good for accident rates: for example, learning by doing to maximize earnings could lead to driving behavior that directly or indirectly increases the probability of accidents, such as gaming time-and-distance pay algorithms by taking longer routes, speeding, and so forth. (Cook et al., 2018). As a result, we do not have a clear hypothesis regarding the persistence of the effects over time. We can, however, examine the effects over time empirically. In the Online Appendix, we break the post-RH variables into quarters to examine the dynamics of the effect up to 2 years after the introduction of RH in our sample cities. Table 3 shows the results. The table shows that RH's increase in accidents and fatalities persists and, in fact, appears to be increasing six quarters after introduction in the city, consistent with a time gap between launch and widespread adoption in a location.[15]

TABLE 3 Dynamic effect of ridehailing on traffic safety.

	Log (1 + Total accidents)			**Log (1 + Total fatalities)**		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Ridehailing tenure						
1–2 Quarters	0.0246***	0.0346***	0.0323***	0.0242**	0.0348***	0.0325***
	(0.0094)	(0.0100)	(0.0109)	(0.0098)	(0.0104)	(0.0114)
3–4 Quarters	0.0255**	0.0371***	0.0350***	0.0237**	0.0356***	0.0337**
	(0.0105)	(0.0112)	(0.0125)	(0.0110)	(0.0117)	(0.0131)
5–6 Quarters	0.0215*	0.0343***	0.0279*	0.0241**	0.0372***	0.0305*
	(0.0114)	(0.0125)	(0.0149)	(0.0120)	(0.0132)	(0.0157)
7–8 Quarters	0.0189	0.0394***	0.0295*	0.0187	0.0402***	0.0293*
	(0.0121)	(0.0130)	(0.0164)	(0.0127)	(0.0137)	(0.0173)
9–10 Quarters	0.0060	0.0350**	0.0213	0.0007	0.0309*	0.0164
	(0.0141)	(0.0154)	(0.0190)	(0.0147)	(0.0161)	(0.0200)
11–12 Quarters	−0.0004	0.0452**	0.0243	−0.0014	0.0460*	0.0253
	(0.0216)	(0.0228)	(0.0272)	(0.0226)	(0.0239)	(0.0284)
>12 Quarters	−0.0198	0.0839**	0.0109	−0.0227	0.0826**	0.0074
	(0.0280)	(0.0350)	(0.0424)	(0.0294)	(0.0359)	(0.0444)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	190,080	190,080	190,080
R2	0.61	0.62	0.63	0.60	0.60	0.61

Note: The table reports estimates of the dynamic effects of ridehailing on traffic safety. *SE*s, clustered by city, are reported in parentheses. The dependent variables are listed at the top of the columns. Total accidents are the number of fatal accidents, according to the definition used by NHTSA. Total fatalities are the total number of fatalities across all fatal accidents. Ridehailing tenure variables are dummy variables that take the value of one if ridehailing has been in effect for the specified periods of time. We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

6.1.1 | Evaluating DD assumptions and causal inference

The validity of our causal inference rests on several assumptions of the generalized DD design. First, a common concern with the staggered entry design is that the treatment may not be exogenous to the outcome of interest. Here, the concern is that RH entry is endogenous and that the choice of cities to enter was affected by trends in fatal accidents. While discussions with RH companies suggest this was not the case, in the Online Appendix A, we formally examine whether fatal accidents predict RH entry order. Online Appendix Table A1 presents multinomial logits for early, mid, and late entry while Online Appendix Table A2 presents a Cox proportional hazard rate model for RH entry into cities. In both models, we observe that accident rates in the city prior to entry do not appear to predict RH arrival. Entry, however, does load positively on population and income. Since we control for these variables directly in each of our specifications, this should not be a concern for causal inference, Overall, the entry order tests support a causal interpretation for our findings.

We assess the parallel trends assumption in Figure 3. Specifically, Figure 3 graphically presents the DD estimates (with each dot representing annual-coefficients) for the 10 years preceding and 4 years following RH adoption for total accidents and total fatalities. In both panels, the counterfactual treatment effects in the pre-RH periods are statistically indistinguishable from zero, supporting our inferences (parallel trends in the pre-period). Post-RH, we see a clear increasing treatment effect.[16]

Next, we address several concerns regarding staggered entry DD specifications brought up in Borusyak and Jaravel (2017) and conduct their proposed robustness tests in

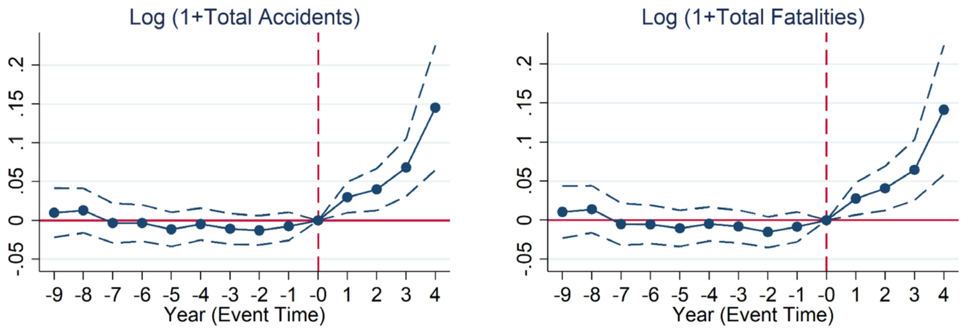

FIGURE 3 Difference-in-differences estimators: Accidents and fatalities. The figure graphs the regression coefficient estimates and their 95% confidence intervals. The confidence intervals are constructed based on city level clustered *SE*s. To generate the coefficients, we run regressions of the various outcome measures on lag and lead indicators (bunched by four quarters) for ridehailing entrance in the location. The vertical red line specifies the year of entry. We provide a description of the variables in Section 3.

the Online Appendix, which include a random effects model, manually detrending our outcomes, and re-running the parallel trends omitting the most and least recent time-period (Table A3 and Figure A1). We also show that the results are similar when our base specification is estimated solely off the sample of ever-treated cities (Table A4).

Given that RH rollout appears to have been related to population and income, we can further strengthen our empirical design by combining the DD specification with a matching strategy. Specifically, we conduct coarsened exact matching based on city population and income during the sample period, employing the default binning algorithm in Blackwell et al. (2009). We re-run our regression model on the matched sample and present the new estimates in Table 2 Panel B. The estimates suggest a 2%–4% increase in traffic accidents and fatalities following RH entry. These effects are robust to a different matching approach (Online Appendix Table A5).

Importantly, we also conduct a placebo test to further validate that the relationship between RH and fatalities is causal. We examine whether the magnitude of our inferences is spurious relative to estimates generated from a simulation of random post periods in the sample (Figure A3). Specifically, we use locations that did not adopt RH (to eliminate contamination from cities that adopted) and simulate 100 cities adopting RH by assigning a random adoption date from within the set of actual launch dates in the cities that adopted. We run this simulation 100 times. We then plot the number of accidents per 100,000 population in simulated event time. We observe no discernable patterns of accidents and drunk accidents in event time before and after the placebo RH entry. This provides further support for a causal interpretation of our results.

Finally, we run the Goodman-Bacon decomposition to provide more context to our staggered DD design, especially the potential for negative weights in our staggered-entry DD models. As noted by Goodman-Bacon (2018), the two-way fixed-effect estimator is a weighted average of all potential 2×2 DD combinations, where the weights are determined by both the size of the treated group and the timing of the treatment. In running the decomposition, we open the black box of the two-way fixed-effect estimator, digging deeper into the comparisons that contribute to our coefficient of interest in the main table. Table A6 and Figure A4 in the Online Appendix show the decomposition results for the DD estimates. Most of the variation used to estimate the DD coefficients results from the cleanest comparison of treated cities to never treated cities ("Never vs. Timing," which compares cities that never adopt the treatment to those that adopt the treatment at some point during the sample period). Overall, the decomposition strengthens our causal interpretation.

6.1.2 | Robustness tests

We conduct a number of robustness tests to further strengthen our results and rule out alternative hypotheses. We provide a detailed list of such robustness tests in Table 4, with references to the detailed analysis and estimates' location in the Online Appendix.

Our first set of robustness tests show robustness to several variations to our empirical methodology. First, we run weighted OLS regressions, in which each observation is weighted by city population or the square root of city population. Such weighting allows us to give higher weights to larger cities, which are less likely to have zero accidents and fatalities in the sample. We continue to find similar effects. We also further assuage concerns with functional form issues in our dependent variable and run both an extensive margin regression where the

TABLE 4 Robustness tests and inferences

Robustness type	Robustness detail	Inferences	Location of detailed analysis in appendix
Method	Population weighting	Unchanged	Table B4
	Coarsened exact matching	Unchanged	Table 2 Panel B and Table A5
	Poisson	Unchanged	Table 2 Panel C
	Negative binomial	Unchanged	Table 2 Panel C
	Different model specifications	Unchanged	Table B1 And Figure B1
	Bottom coding outcomes at the 5th percentile	Unchanged	Table B2
	Inverse sine transformation	Unchanged	Table B3
	Extensive and intensive margin	Effects are mainly concentrated on the intensive margin	Table B5
	DiD sensitivity analysis	Unchanged	Table A3
	Estimating main effects from treated cities	Unchanged	Table A4
	Annual DiD estimators when changing base years	Unchanged	Figure A1
	Quarterly DiD estimators	Unchanged	Figure A2
	Goodman-Bacon decomposition	Unchanged	Table A6 and Figure A4
Measure	Treatment measured at CBSA-level	Unchanged	Table B6
	Cluster *SE*s by CBSA	Unchanged	Table B6
	Placebo tests on RH entry	We observe no discernable pattern using a set of placebo treatment dates. More detailed discussions can be found in Section 5.1.1.	Figure A3
	Accident and fatality rate	Unchanged	Table 2 Panel C
	Alternative hypothesis: City Growth	Effects not driven by city growth	Table B7
Control	Population growth rate	Unchanged	Table B8 and Figure B2
	Unemployment rate	Unchanged	Table B8 and Figure B2
	Retail gas price	Unchanged	Table B8 and Figure B2
	Retail gas price change	Unchanged	Table B8 and Figure B2
	Cellphone usage	Unchanged	Table B9
Sample	Weekday	Effects are larger on weekends than on weekdays, and at night. More detailed discussions can be found in Section 5.1.2.	Table B11 and Figure B3
	Weekend		
	Night		
	Shorter sample (2007–2016)	Unchanged	Table B10

Note: This table summarizes the robustness tests conducted in the paper, our inferences from such tests, and the corresponding location of detailed analysis in the Online Appendix. We break down robustness tests into four categories: methodological robustness, robustness to measurements, robustness to additional control variables, and robustness to sample variations.

dependent variable is a dummy for the outcome variable being strictly positive, and an intensive margin regression where the dependent variable is the log of the outcome variable, with the sample restricted to positive values of the outcome. Again, we find similar results.

Our second set of robustness tests address concerns about the measurement of treatment. In many cases, RH was rolled out at the core-based statistical area (CBSA) level rather than to an individual city within the CBSA. We demonstrate robustness to variations in fixed effects and city trends, as well as clustering *SE*s at the CBSA level. We also demonstrate that our estimates remain robust and statistically significant when re-running our models at the CBSA level rather than the place-level, utilizing the earliest adoption date within the CBSA and clustering *SE*s at the CBSA level. We observe a somewhat larger magnitude of the effect, at roughly 5% (vs. 3.6% in main specification). Another competing explanation for our observed results is that the effect is merely driven by city growth as opposed to RH. If our fatal accident measures are merely picking up city growth that is reflected in additional VMT, we would expect for this growth to be reflected in wages in the city. However, we observe no evidence of a statistically significant growth in wages post-RH (untabulated).

Third, our results are robust to additional control variables to address specific concerns that could confound inference. First, we demonstrate robustness to controlling for the population growth rate, retail gas prices, the change in retail gas prices, and the unemployment rate. We also rule out the explanation that RH merely coincides with smartphone adoption patterns, and that our results are spuriously pick up the increase accidents driven by increased smartphone usage. We directly address this concern by controlling for proxies of smartphone adoption. We utilize data from PowerAnalytics, which provides aggregate gross receipts, employment and number of establishments for NAICS Code 517312: Wireless Telecommunications Carriers (Except Satellite) at the MSA level on an annual basis. As an additional proxy for smartphone adoption, we also use Google Health Trends API search volume for smartphone related keywords ("iPhone," "Android," "Samsung Galaxy," "Smartphone," and "cellphone"). We demonstrate the robustness of our main finding to the inclusion of these four control variables. While accidents load positively on the cellphone sales and employment, their addition, however, do not affect the significance or the magnitude of the main effect, which remains robust in all specifications.

Our last set of robustness tests demonstrate the robustness and variations in our results across different subsamples. When we use a shorter sample (2007–2016), our inferences remain unchanged. Furthermore, we break out weekend accidents, nighttime accidents, weekday accidents, and weekend night accidents for total fatal accidents and total fatalities. We observe similar patterns to those exhibited in the models in Table 2. Accident and fatality increases are lowest on weekend nights (Friday and Saturday, after 5 PM, and before 6 AM) at 2.50% and 2.69% respectively, consistent with weekend nights being the most likely period in which RH may be reducing the number of impaired drivers. For total weekend and nighttime accidents and fatalities, the magnitudes of the estimated increases are between 3% and 4%. We graph these estimates in Online Appendix Figure B3. Panel A presents the estimates and confidence intervals for total fatal accidents and total fatalities on weekends and nights. Panel B further splits the sample into large (highest quartile) and small (lowest quartile) cities by population, and graphs the estimates for accidents and drunk accidents on weekends and nights for each. The panel hints at what we will see in the heterogeneity estimations: that the effects of RH appear to be larger in larger cities.

6.2 | Adoption intensity

We next provide evidence in support of Hypothesis 1B. In Table 5, we explore the effect of the intensity of service adoption or popularity, proxied for by Google searches for RH keywords. In our main models, we define treatment as the first launch of a RH service. Utilization of RH over time, however, is likely to increase. We re-estimate the models, including the interaction of TREATMENT with our Google search intensity measure. We present the results of this estimation where the outcome variable is related to total accidents in columns (1), (2), and (3) and total fatalities in columns (4), (5), and (6). The estimates suggest that accidents increase with increases in the intensity measure. In each model, the coefficient on POST × INTENSITY is positive and statistically significant.[17]

6.3 | Pedestrians versus vehicle occupants

An important question is whether the increase in accidents and fatalities suggested by the estimates in Table 2 are concentrated among vehicle occupants versus the alternative of potentially imposing an externality on non-vehicle occupants (bystanders). The FARS data allow us to separate out accidents in which pedestrians were involved. We code accidents as pedestrian-involved if the FARS database indicates they involve persons that are not motor vehicle occupants or motorcycle riders.[18] Thus, "pedestrian" in our context refers to both pedestrians as well as bicycle, skateboard and scooter riders, and so forth.

In Table 6, we present the estimates from models similar to those described in Equation (1), substituting our measures of total fatal accidents with similar measures that solely count accidents in which a pedestrian was involved. Our *accidents* measure in columns (1) and (2) is

TABLE 5 Hypothesis testing: Hypothesis 1B: Variation of ridehailing service and adoption intensity.

	Log (1 + Total accidents)			Log (1 + Total fatalities)		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Post × $Treated_{t,c}$ × Log rideshare – Related Google Search $Volume_{ct}$	0.0027***	0.0049***	0.0028**	0.0026***	0.0049***	0.0028**
	(0.0007)	(0.0009)	(0.0012)	(0.0008)	(0.0010)	(0.0013)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	154,440	154,440	154,440	154,440	154,440	154,440
R2	0.61	0.62	0.63	0.60	0.61	0.62

Note: The table reports estimates of the effect of ridehailing on traffic safety varied by the intensity of service. *SE*s, clustered by city, are reported in parentheses. In all panels, the dependent variables are listed at the top of the columns. Total accidents are the number of fatal accidents, according to the definition used by NHTSA. Total fatalities are the total number of fatalities across all fatal accidents. In Panel A, single (pooled) ride service is a dummy variable that takes the value of one if any single (pooled) ride service is adopted. In Panel B, Log RideHail Google Search Volume is the natural logarithm of Google search intensity for the terms "Uber," "Lyft," and "rideshare." We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

TABLE 6 Externality of ridehailing on non-vehicle occupants.

	(1) Log (1 + Pedestrian-involved accidents)	**(2) Log (1 + Pedestrian-involved fatalities)**	**(3) Log (1 + Pedestrians involved in fatal accidents)**
Post × $Treated_{t,c}$	0.0245***	0.0245***	0.0277***
	(0.0058)	(0.0058)	(0.0062)
City and quarter fixed effects	Yes	Yes	Yes
City linear trend	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
Observations	190,080	190,080	190,080
R2	0.54	0.54	0.55

Note: The table presents reports estimates from generalized difference-in-difference regressions. *SE*s, clustered by city, are reported in parentheses. We list the dependent variables at the top of each column. Pedestrian is defined as any person not in or upon a motor vehicle or other vehicle. Pedestrian-involved accident measures the number of fatal accidents involving at least one non-vehicle occupants. Pedestrian-involved fatalities measures the total number of fatalities in all accidents involving at least one non-vehicle occupants. Pedestrians involved in fatal accidents measures the total number of non-vehicle occupants involved in fatal accidents. Post × $Treated_{t,c}$ is an indicator variable that takes the value of one if city c adopted one or more ridehailing services at time *t*, and zero otherwise. City-specific linear trends are included in all columns. We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

the total number of accidents in which a pedestrian was involved; in columns (3) and (4), it is the total number of fatalities in accidents that involved a pedestrian; and in columns (5) and (6), it is the number of pedestrians involved in fatal accidents. The estimates from these models follow the same pattern as the estimates of our main models, suggesting that the increase in accidents, following RH entry, imposes an externality on nonvehicle occupants. The magnitudes of these increases mirror those in our main models, ranging from a 2.45% increase in total fatal accidents involving a pedestrian and in fatalities in accidents involving a pedestrian, to an increase of 2.77% in the number of pedestrians who are involved in fatal accidents.

Figure 4 graphically presents the DD estimates (with each dot representing two quarter-coefficients) for the 10 years preceding and 4 years following RH adoption for pedestrian accidents. As in our main models, the counterfactual treatment effects in the pre-RH periods are statistically indistinguishable from zero, providing further support for our inferences. In Online Appendix

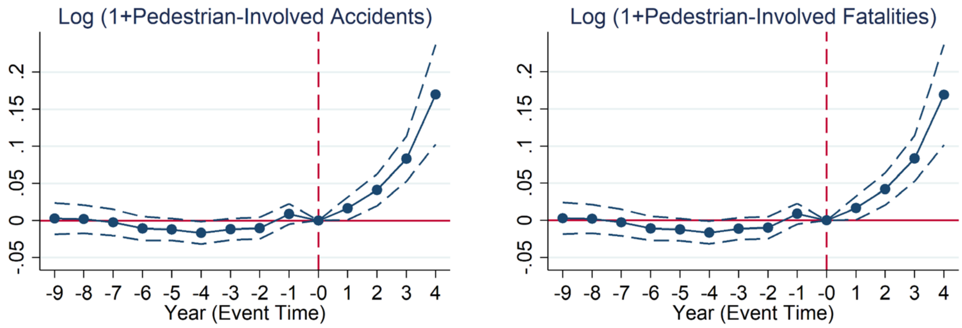

FIGURE 4 Difference-in-differences estimators: Pedestrian-involved accidents and fatalities. The figure graphs the regression coefficient estimates and their 95% confidence intervals. The confidence intervals are constructed based on city level clustered *SE*s. To generate the dynamics of the treatment effects, we run regressions of our various outcome measures on lag and lead indicators (bunched by four quarters) for ridehailing entrance. The vertical red line specifies the year of entry. A description of the variables is provided in Section 3.

Figure A2, we present versions of these graphs using quarterly estimates, with a similar conclusion.

6.4 | Urban versus rural areas

Hypothesis 2 posits that the effects of the introduction of RH on accidents will be larger in more urban areas. We explore this formally by breaking out our results across population, population density, and ex-ante vehicle ownership, as reported by the American Community Survey. We divide cities into quartiles for each characteristic, and re-estimate our models, interacting $\text{POST} \times \text{TREATED}_{t,c}$ with the quartile indicators. Table 7 Column (1) presents the estimates where the city characteristic of interest is city population. The estimates suggest that the increase in accidents observed in our main models is concentrated in large cities (fourth quartile), with an estimate for $\text{POST} \times \text{TREATED}_{t,c} \times \text{Q4}$ of 7.4% that is statistically significant; in contrast, the estimates for the bottom three quartiles of city population are an order of magnitude smaller and insignificant at conventional levels. This is consistent with the larger magnitude estimates that arise in the population-weighted analyses presented in the Online Appendix. It is also consistent with findings in Edlin and Karaca-Mandic (2006) on accident externalities from driving more generally: high traffic locations have economically large externalities, while in contrast, the accident externality per driver in low-traffic locations appears to be quite small (Figure 5).

Column (2) repeats the exercise using population density instead of population. We similarly find that the effect is concentrated in higher quartiles of density; in contrast, we observe no effect in the lowest quartile of population density. In column (3), we turn to a measure of ex-ante vehicle ownership. We observe that the increase in accidents following the launch of RH services appears to be concentrated in cities in the top quartile of ex-ante vehicle ownership. Overall, the results across all three characteristics are consistent with Hypothesis 2.

6.5 | Mechanisms

Having established a robust pattern of estimates consistent with an increase in fatal accidents and fatalities following the launch of RH services in a city, we now consider the two mechanisms discussed in our hypothesis development: increases in quantity (road utilization in the form of VMT) and changes in driver quality. Recognizing that providing definitive proof as to the existence of a particular mechanism is infeasible (Ylikoski & Emrah Aydinonat, 2014), we explore whether the empirical evidence is consistent with one or both of the theoretical mechanisms outlined by our conceptual framework.

6.5.1 | Quantity

We begin with an exploration of the effects of RH on measures of road congestion. Road-utilization and congestion data for city roads are not readily available for most cities. To examine this channel, first, on the intensive margin, we use annual estimates of arterial VMT, excess gas consumption,[19] and hours delay in traffic for 99 urban areas reported by the TAMU Transportation Institute for the years 2000–2014.

In Table 8, we provide empirical evidence consistent with Hypothesis 3A. We estimate similar models to our main specification, replacing the *accidents* variable as our dependent variable with arterial street daily VMT (column [1]), annual excess fuel consumption (column [2]), and annual hours of delay (column [3]). Due to the limited availability of data relative to the full sample, the models in Table 8 aggregate locations up to the urban area.[20] Moreover, we can estimate only for the years up

TABLE 7 Hypothesis testing: Hypothesis 2: Effects of ridehailing on traffic fatalities: Urban and rural cities

	Log (1 + Total accidents)		
	(1) Population	(2) Population Density	(3) Ex Ante Vehicle Ownership
Post × $\text{Treated}_{t,c}$ × Q4	0.0746***	0.0354***	0.0769***
	(0.0113)	(0.0109)	(0.0111)
Post × $\text{Treated}_{t,c}$ × Q3	0.0049	0.0323**	0.0024
	(0.0136)	(0.0136)	(0.0145)
Post × $\text{Treated}_{t,c}$ × Q2	0.0041	0.0504***	−0.0054
	(0.0154)	(0.0163)	(0.0148)
Post × $\text{Treated}_{t,c}$ × Q1	0.0035	0.0265	0.0136
	(0.0128)	(0.0165)	(0.0136)
City and quarter fixed effects	Yes	Yes	Yes
City linear trend	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
Observations	190,080	190,080	190,080
R2	0.62	0.62	0.62

Note: The table reports estimates of heterogeneous effects of ridehailing on traffic safety by population, population density, and ex ante vehicle ownership. *SE*s, clustered by city, are reported in parentheses. Population measures annual city population. Population density measures population per square mile. Vehicle ownership measures the total number of available vehicles in households. Total accidents are the number of fatal accidents according to the definition used by NHTSA. Post × $\text{Treated}_{t,c}$ is an indicator variable that takes the value of one if city c adopted one or more ridehailing services at time t, and zero otherwise. Q4 is an indicator for whether a city's population, population density, or ex ante vehicle ownership falls in the fourth quartile of the distribution of the variable. Similarly, Q1 is an indicator for whether a city's population, population density, or ex ante vehicle ownership falls in the first quartile of the distribution of the variable. All columns include city-specific linear trends. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

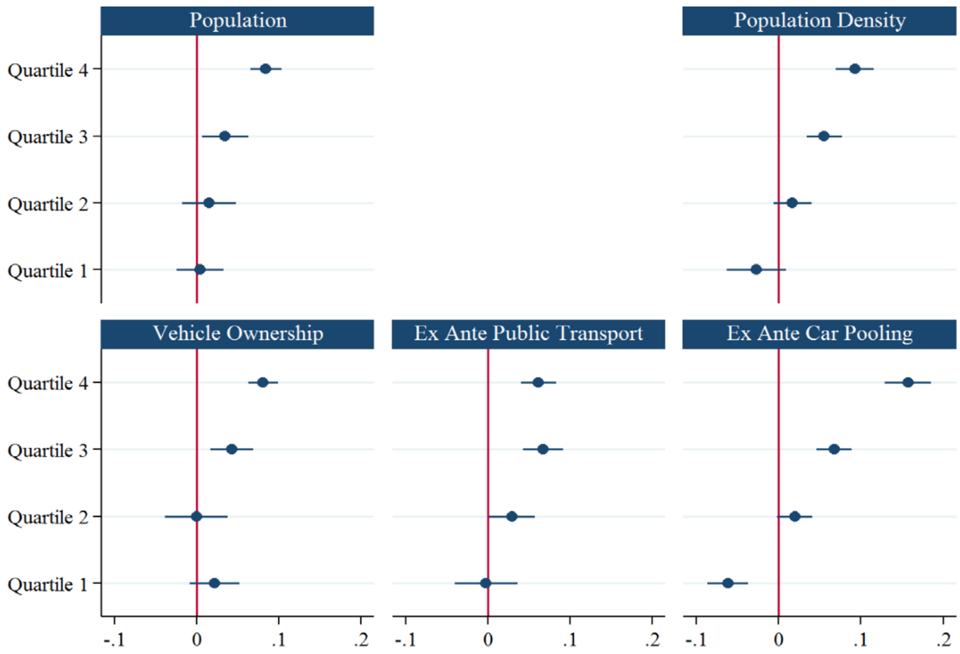

FIGURE 5 Heterogeneous effect on vehicle registration by city characteristics. The figure graphs the regression coefficient estimates and their 95% confidence intervals. The confidence intervals are constructed based on city level clustered *SE*s. The coefficient estimates are broken down by quartiles for five city characteristics: population, population density, vehicle ownership, ex-ante usage of public transportation, and ex-ante use of carpooling. The outcome variable in all regressions is the natural logarithm of one plus new vehicle registrations. We provide a description of the variables in Section 3.

to 2014, for these 99 urban areas, leaving us with 1386 observations (as compared to 190,080 in our other models). Still, for all three models, we obtain a positive and significant estimate for the coefficient on our variable of interest, POST × TREATMENT, though with lower statistical significance levels. The economic magnitudes for both measures are roughly on the order of a 1.6% increase: approximately one half the

TABLE 8 Hypothesis testing: Hypothesis 3A: Effect of ridehailing on road utilization and congestion.

	(1) Log arterial street VMT	(2) Log excess fuel consumption	(3) Log hours of delay in traffic
Post × $Treated_{t,u}$	0.01799**	0.01562**	0.01563**
	(0.00834)	(0.00715)	(0.00715)
Urban area and year fixed effects	Yes	Yes	Yes
Urban area linear trend	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
Observations	1386	1386	1386
R2	0.998	0.999	0.999

Note: The sample contains 1386 annual observations on 99 urban areas from 2001 to 2014. *SE*s, clustered by urban area, are reported in parentheses. The dependent variables are the natural logarithm of congestion-related measures listed at the top of each column. Arterial Street VMT measures the total number of VMT on arterial streets in an urban area. Excess fuel consumption measures the extra fuel consumed, due to inefficient operation in slower stop-and-go traffic. Hours of delay measures the amount of extra time spent traveling, due to congestion. Post × $Treated_{t,u}$ is an indicator variable that takes the value of one if urban area u adopted one or more ridehailing services at time *t*, and zero otherwise. Urban area-specific linear trends are included in all regressions. We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. For more detailed information on the dependent variables, please refer to https://static.tti.tamu.edu/tti.tamu.edu/documents/mobility-scorecard-2015-wappx.pdf. We use ***, **, and * to denote statistical significance at the 1%, 5%, and 10% levels, respectively.

magnitude of the effect of RS on total fatal accidents, and suggesting that quantity effects alone are unlikely to fully explain the increase in accidents estimated in the main models.[21] We present estimates graphically over time in Figure 6. Similar to our main models, the figures suggest no evidence of pre-trends in the period prior to the introduction of RH, with a significant positive increase post–RH.

Next, in Table 9, we provide empirical evidence consistent with Hypothesis 3B. We examine the extensive margin in usage by estimating similar models where the dependent variable is the logarithm of new car registrations as reported by Polk Automotive. Both Lyft and Uber often report numbers from surveys of users, suggesting some of their riders forgo owning their own cars, and thus argue that they are removing vehicles from the road. These surveys, however, do not account for the possibility that, while some of the rider population is forgoing owning a vehicle, others may be purchasing vehicles precisely in order to work as RH drivers. While the advent of RH may reduce personal car usage for some, it also transforms cars into a productive asset, as RH now makes it lucrative to drive. Both Uber and Lyft offer programs subsidizing the purchase or leasing of vehicles for those willing to become driver-partners on their platform. Consistent with this notion, Buchak (2018), in contemporaneous analysis, documents that RH entry coincides with sharp increases in auto loans, auto sales, employment and vehicle utilization among low-income individuals. Thus, while RH may enable reductions in vehicles purchased on the rider (demand) side, it also provides strong incentives for the purchase of more vehicles on the supply (driver) side. Which effect dominates is an empirical question.

Panel A of Table 9 reports the estimates from models with and without the location-specific trend. The estimates suggest that the initiation of RH leads to an increase in new car registrations, rather than an overall decrease. This increase is in the range of 3–5% when including the location-specific time trend. A caveat to the magnitudes is that we cannot separate out regular car purchases from RH-service intended vehicle registrations. In Panel B Column (1), we advance the intuition of this extensive margin effect by examining how new car registrations respond to the interaction of RH intensity, as proxied by the Google search share variable used in Section IVB. The estimates suggest that new car registrations increase with the intensity of Google searches for Uber/Lyft/RH. This relationship intensifies when RH begins in a treated city. These results suggest that new vehicle purchases increase as RH services adoption intensifies.

Turning to Panel C of the table, the heterogeneity in this increase along city characteristics lines up with the heterogeneity in the increase in accidents documented in Section IVD: the new car registrations are concentrated in cities with above median population and in cities with above median ex-ante vehicle ownership. Moreover, the increase in new car registrations is larger in cities with high ex-ante public transport usage and car pool usage. They are decreasing only in the cities with the lowest

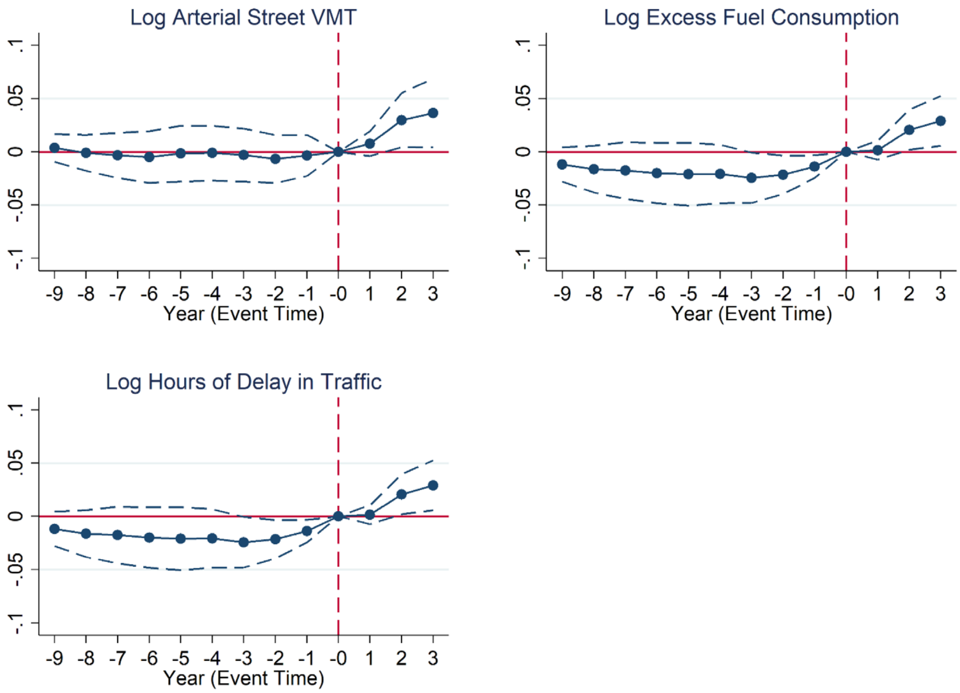

FIGURE 6 Difference-in-differences estimators: Congestion. The figure graphs the regression coefficient estimates and their 95% confidence intervals. The confidence intervals are constructed based on city level clustered *SE*s. We map the counterfactual treatment effects by running regressions of congestion measures on lag and lead indicators for the entry of ridehailing. The vertical red line indicates the year of entry. We provide a description of the variables in Section 3.

quartile of ex-ante carpool usage. Figure 5 graphs the estimates and confidence intervals.

Interestingly, the estimates in Panel C of Table 9 suggest that the increase in new car registrations is higher in cities with high population density: the estimates imply a 9% increase in new registrations in the cities in the highest quartile, a 6% increase in cities in the second quartile, a 2% increase for cities in the third quartile, and a statistically insignificant 3% decrease in cities in the lowest quartile. Overall, this fact pattern suggests increases in congestion prompted by RH. The increase in new car registration also appears to be concentrated in cities with higher population levels in general, consistent with our VMT findings.

Hypothesis 3C posits that the effects of the introduction of RH in a city will be higher in cities with high ex ante use of public transportation and carpooling. Table 10 tests this moderating hypothesis, breaking out our results across quartile of ex ante public transportation usage and ex ante carpool usage. We then interact POST $\times$ TREATMENT$_{t,c}$ with quartile indicators for each characteristic. We observe that the increase in accidents is concentrated in cities with higher ex-ante usage of public transportation (top two quartiles), and in cities that had above-median carpool usage. These estimates are consistent with a substitution effect to RH and away from public transport and carpooling, and are also consistent with RH serving to fill last-mile transport needs, complementing public transport use (Hall et al., 2018) in nonurban areas. Both sets of estimates are consistent with Hypothesis 3C.

6.5.2 | Driver quality

Examining driver quality is challenging given the nature of the available data, and thus, providing evidence related to Hypothesis 4 must by necessity be indirect. As seen in the previous analyses, the magnitude of the association between RH arrival and VMT is only about half the magnitude of the association with fatal accident measures. This difference would be consistent with the presence of decreased quality overall. Moreover, in additional analysis, we utilize the subsample of 99 urban areas for which the TAMU VMT data is available, and re-estimate our difference-in-difference models, this time additionally controlling for VMT, to see whether the increased VMT (quantity channel) absorbs the entire main effect (Online Appendix Table C1). When estimating at the aggregated Urban Area level, the coefficient of interest remains positive, even controlling for the increase in VMT. When we estimate the models at the place level, we observe positive and statistically significant coefficients, controlling for VMT. Thus, the evidence suggests that part of the effect we document comes from outside the quantity channel, consistent with a reduction in overall driver quality.

Furthermore, if the effect is primarily driven by RH drivers, as opposed to just a VMT increase overall, the effect should be more prominent for vehicles more likely to be RH vehicles. RH companies restrict the types and ages of vehicles that can be used by driver-partners. While these restrictions vary somewhat from company to company, generally speaking, all RH companies require

TABLE 9 Hypothesis testing: Hypothesis 3B: The effect of ridehailing on new car registrations.

Panel A: Main effect

	Log (1 + New car registrations)		
	(1)	**(2)**	**(3)**
Post × $\text{Treated}_{t,c}$	0.0291***	0.0490***	0.0296***
	(0.0079)	(0.0069)	(0.0065)
City and quarter fixed effects	Yes	Yes	Yes
City linear trend	No	Yes	Yes
City quadratic trend	No	No	Yes
Control variables	Yes	Yes	Yes
Observations	190,080	190,080	190,080
R2	0.94	0.97	0.98

Panel B: Intensity

	Log (1 + new car registrations)	
	(1)	**(2)**
Google Search Vol		
Post × $\text{Treated}_{t,c}$ × Log Rideshare–Related Google Search Vol_{ct}	0.0075***	
	(0.0009)	
Rideshare service type		
Single ride service (UberBlack/Taxi/X, Lyft)		0.0463***
		(0.0067)
Pooled ride service (UberPool, Lyft Line)		0.0300***
		(0.0107)
City and quarter fixed effects	Yes	Yes
City linear trend	Yes	Yes
Control variables	Yes	Yes
Observations	154,440	190,080
R2	0.97	0.97

Panel C: Heterogeneous effects

Dependent Var: Log (1 + New car registrations)	**(1) Population**	**(2) pop density**	**(3) Public transport**	**(4) Carpool**	**(5) Vehicle ownership**
Post × $\text{Treated}_{t,c}$ × Q4	0.0842***	0.0926***	0.0620***	0.1572***	0.0809***
	(0.0096)	(0.0118)	(0.0110)	(0.0144)	(0.0093)
Post × $\text{Treated}_{t,c}$ × Q3	0.0347**	0.0554***	0.0673***	0.0674***	0.0427***
	(0.0142)	(0.0110)	(0.0123)	(0.0108)	(0.0132)
Post × $\text{Treated}_{t,c}$ × Q2	0.0149	0.0168	0.0299**	0.0195*	−0.0001
	(0.0168)	(0.0118)	(0.0142)	(0.0109)	(0.0195)
Post × $\text{Treated}_{t,c}$ × Q1	0.0041	−0.0274	−0.0020	−0.0619***	0.0219
	(0.0147)	(0.0186)	(0.0194)	(0.0127)	(0.0155)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes
City linear trend	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	190,080	190,080

TABLE 9 (Continued)

Panel C: Heterogeneous effects					
Dependent Var: Log (1 + New car registrations)	**(1) Population**	**(2) pop density**	**(3) Public transport**	**(4) Carpool**	**(5) Vehicle ownership**
R2	0.97	0.97	0.97	0.97	0.97

Note: The table presents estimates of the effect of ridehailing on new car registrations. *SE*s, clustered by city, are reported in parentheses Post × Treated$_{t,c}$ is an indicator variable that takes the value of one if city *c* adopted one or more ridehailing services at time t, and zero otherwise. Panel A displays estimates from generalized difference-in-difference regressions without city time trend, with linear trend, and with quadratic trend. Panel B demonstrates how the effect varies with the intensity of ridehailing service. Log RideHail Google Search Volume denotes the natural logarithm of Google search intensity for the terms "Uber," "Lyft," and "rideshare". Single (Pooled) Ride Service is a dummy variable that takes the value of one if any single (pooled) ride service is adopted. Panel C breaks out results across a variety of city characteristics and ex-ante behaviors. The variable used for sample cut is listed at the top of each column. Population measures annual city population. Population density measures population per square mile. Pop Density measures the population per square mile. Public Transport measures the percentage of the population commuting to work using public transportation. Carpool measures the percentage of the population commuting to work using carpools. Vehicle Ownership measures the total number of available vehicles in households. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

vehicles to be four door, and typically less than 10 years old (less than 5 for some categories of service). We take advantage of this to provide further evidence in support of a quality channel. In Table 11, we utilize the vehicle level dataset provided by NHTSA to explore this notion. Panel A columns (1) and (2) presents estimates from models at the crash-vehicle level. We use linear probability models where the dependent variables are indicators for whether the crash vehicle has RH-eligible characteristics. We define a RH-eligible vehicle as a four-door sedan, SUV or minivan that is less than 5 or 10 years old. We include city × year, city × day of week and city × hour of day fixed effects in each model to account for differences in passenger flows due to work days and rush hours. The estimates demonstrate that the effect holds at similar magnitudes for accidents involving at least one RH-eligible vehicle. Column (3) presents the estimates for a placebo test where we estimate the same model for the natural logarithm of one plus the count of accidents involving a two-door vehicle (which are not eligible for RH). Here, we observe no significant effect post-RH (either economically or statistically).

Furthermore, if accidents post-RH entry are primarily RH vehicle driven (which would support the lower quality argument), we should observe more passengers per accident on average, as during the RH rides, RH vehicles involve a driver and at least one passenger, whereas before RH, there could be a single driver only. We demonstrate that this is the case in Table 11 Panel B, where we estimate our previous DD model, but where the dependent variable is log of one plus the number of passengers involved in accidents. As in our main models, we estimate three version of each model, once with no trends, once with a location-specific linear trend, and once with both linear and quadratic location-specific trends. In all three models (columns [1]–[3]) we a significant increase of between 2.5% and 6.6% in the number of passengers involved in accidents post RH entry. In columns (4) through (6), we estimate models where the dependent variable is log number of passengers per vehicle involved in accidents. In order to calculate number of passengers per vehicle, we must restrict the sample solely to observations where there is at least one accident, resulting in a loss of some statistical power. Here too the estimates indicate an increase in the number of passengers per vehicle involved in accidents.

In the Online Appendix, we conduct a number of further analyses utilizing data at the accident level. We estimate linear probability models using data from NHTSA's person-level file which details the seating position of vehicle occupants involved in accidents. In Online Appendix Table C2, we show that the coefficient on POST × TREATED$_{t,c}$ is positive and significant when the dependent variable is an indicator for an accident involving at least one back row passenger and at least one back row adult passenger, but not when the dependent variable is an indicator for an accident involving at least one back row passenger but where all back-row passengers are children. The former two scenarios are more likely to be accidents that involve RH vehicles, while the latter scenario is unlikely to be an accident that involves a RH vehicle. These patterns are consistent with the mechanism for our findings being at least partially a quality effect associated with RH drivers (as opposed to non-RH drivers).

All this said, many of the arguments supporting the notion of improved quality concentrate primarily on reduction of drunk driving specifically. To the extent that the substitution of inebriated drivers with sober RH drivers leads to improved quality along this dimension,

TABLE 10 Hypothesis testing: Hypothesis 3C: Effects of ridehailing on traffic fatalities: The role of public transport.

	Log (1 + Total accidents)	
	(1) Ex ante public transportation usage	(2) Ex ante car pool usage
Post × $\text{Treated}_{t,c}$ × Q4	0.0381***	0.0448***
	(0.0112)	(0.0151)
Post × $\text{Treated}_{t,c}$ × Q3	0.0537***	0.0633***
	(0.0123)	(0.0131)
Post × $\text{Treated}_{t,c}$ × Q2	0.0175	0.0174
	(0.0155)	(0.0129)
Post × $\text{Treated}_{t,c}$ × Q1	0.0181	0.0120
	(0.0190)	(0.0128)
City and quarter fixed effects	Yes	Yes
City linear trend	Yes	Yes
Control variables	Yes	Yes
Observations	190,080	190,080
R2	0.62	0.62

Note: The table reports estimates for the heterogeneous effects of ridehailing on traffic safety by public transport usage. *SE*s, clustered by city, are reported in parentheses Public transportation usage measures the percentage of population commuting to work using public transportation. Carpool usage measures the percentage of population commuting to work using carpool. Total accidents are the number of fatal accidents according to the definition used by NHTSA. Post × $\text{Treated}_{t,c}$ is a dummy variable that equals one if a city c adopted at least one ridehailing service at time t. Q4 is an indicator for whether a city's population, population density, or ex ante vehicle ownership falls in the fourth quartile of the distribution of the variable. Similarly, Q1 is an indicator for whether a city's population, population density, or ex ante vehicle ownership falls in the first quartile of the distribution of the variable. All columns include city-specific linear trends. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

we would expect to see a reduction in the rate of drunk fatal accidents and fatalities following the introduction of RH. However, any such reduction may be limited or somewhat offset by incidence of drunk driving by RH drivers themselves,[22] or increases in alcohol consumption more generally due to behavioral responses stemming from risk compensation (Burgdorf et al., 2019).

Several prior studies have made attempts to examine the effect of RH on fatal accidents in US cities (see e.g., Brazil & Kirk, 2016; Greenwood & Wattal, 2017; Martin-Buck, 2017). Most of these studies rely on data on fatal alcohol-related auto accidents from the NHTSA over the period of 2000–2014, and a number conclude that the advent of RH is associated with a significant reduction in alcohol related accidents. These studies employ a variety of specifications, including difference-in-differences specifications. These studies, suffer from a number of deficiencies that put the interpretation of their results in question. First and foremost, these studies do not appear to account for an important change in how accidents were classified as involving alcohol impairment, which took place in the beginning of 2008. Specifically, in the years prior to 2008, alcohol-related accidents were recorded by the NHTSA as any fatal accident involving at least one person—vehicle occupant (driver or non-driver) or pedestrian—having blood alcohol levels above the legal threshold for impaired driving. In other words, prior to 2008, an accident in which a sober taxi driver kills a drunk pedestrian would be classified as a drunk accident, and a sober driver driving a passenger who was impaired who gets into an accident that results in a death would also be classified as a drunk accident. From 2008 onwards, alcohol-related accidents are recorded as “drunk accidents” only if the driver himself was impaired. This definitional change leads to a massive mechanical decrease in accidents classified as “alcohol-related” and to a corresponding increase in non-drunk accidents. The mechanical drop (see Online Appendix Figure C1) has clear effects when estimating models that include the pre-change years. We present evidence on the effects of the definitional change on the results of prior work in Online Appendix Section C3 where we also discuss issues related to sample coverage and model specification that affect the interpretation of these studies.

Here, we instead break the sample of accidents into those involving a drunk driver and those that do not involve a drunk driver. We define outcome measures as follows: Total Drunk Accidents is the total number of fatal accidents involving any drunk drivers. Total Drunk Fatalities is the total number of fatalities in all drunk-driver accidents. We restrict our models to the period of 2008–2016 for the estimation, to avoid the impact of the NHTSA definitional change. In the Online Appendix Figure C2, we show the sensitivity of the results when including the years prior to the change in the estimation: the models document a negative coefficient when linear and quadratic trends are not included, similar to the results that have been found in past studies that did not account for trends and used the prior years; the coefficients flips to a positive once when we remove the years prior to the definition change, or when we add linear (and/or quadratic) trends, as would be expected. In Figure C3, we show this pattern does not exist for total fatalities, which are not affected by the definitional change.

Table 12 presents the estimates from our DD models for these measures for the post-definition change period.

TABLE 11 Hypothesis testing: Hypothesis 4: Ridehailing eligible and non-eligible vehicles

Panel A: Ridehailing-eligible vehicles			
	Characteristics of RH-eligible vehicles		**Placebo (non-eligible)**
	(1)	**(2)**	**(3)**
	5-year or newer (4 door)	**10-year or newer (4 door)**	**2-Door**
Post × $\text{Treated}_{t,c}$	0.0293***	0.0246**	0.0021
	(0.0099)	(0.0116)	(0.0057)
City × Year FE	Yes	Yes	Yes
City × Day of week FE	Yes	Yes	Yes
City × Hour of day FE	Yes	Yes	Yes
Observations	397,839	397,839	397,839
R2	0.23	0.24	0.22

Panel B: Passengers involved in accidents						
	Log (1 + passengers)			**Log (passengers per vehicle)**		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Post × $\text{Treated}_{t,c}$	0.0253**	0.0657***	0.0499***	0.0178*	0.0419***	0.0185
	(0.0106)	(0.0125)	(0.0151)	(0.0093)	(0.0110)	(0.0132)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	97,599	97,599	97,599
R2	0.52	0.54	0.54	0.08	0.12	0.16

Note: In this table, we report generalized difference-in-difference regression estimates, and list the dependent variables at the top of the columns. We report robust *SE*s clustered at the city level in parentheses. In Panel A, the unit of observation is crash-vehicle, and the outcomes are indicators for whether the vehicle involved in the accident is newer than 5 years old (4 door vehicle), newer than 10 years old (4 door vehicle), and 2-door respectively in columns (1), (2), and (3). In Panel B, Passengers is the total number of persons in motor vehicles involved in fatal accidents. Passenger Per Vehicle is the average number of persons per motor vehicle involved in fatal accidents. We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. We use ***, **, and * to denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Column (1) presents estimates for the model without city-specific trends: we observe a positive, but statistically insignificant coefficient. Once we include linear or linear and quadratic trends, we observe a positive and weakly statistically significant coefficient, on the order of roughly 1.5% increase in fatal accidents post RH introduction. Thus, the estimates suggest that any quality improvements from RH in the form of removal of drunk drivers from the road do not seem to swamp the quantity effect or other quality effects which we cannot directly observe.[23] We present estimates graphically over time in Figure 7. The figure supports the conclusion that there is no statistically significant reduction in drunk-driving-involved fatal accidents after the introduction of RH. While we cannot offer evidence on whether some of the effects we document are driven directly by a reduction in overall driver quality, the results above suggest that there is not a significant quality improvement post-RH, at least in terms of removal of drunk drivers, to reduce the number of accidents overall.

7 | DISCUSSION

7.1 | Back of the envelope cost calculation

To assess the overall impact of RH, we must also consider its benefits, such as the consumer surplus provided by convenience: Cohen et al. (2016) estimate that in 2015, the overall consumer surplus generated by UberX in the United States in 2015 was $6.8 billion. Our estimates allow us to attempt to quantify the cost of the RH's increase in fatal accidents, using estimates of the value of

TABLE 12 Hypothesis testing: Hypothesis 4: The effect of ridehailing on alcohol-involved accidents and fatalities.

	Log (1 + Drunk accidents)			Log (1 + Drunk fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × $\text{Treated}_{t,c}$	0.0059	0.0146**	0.0166*	0.0052	0.0148**	0.0173*
	(0.0050)	(0.0070)	(0.0087)	(0.0054)	(0.0074)	(0.0093)
City and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City linear trend	No	Yes	Yes	No	Yes	Yes
City quadratic trend	No	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	106,920	106,920	106,920	106,920	106,920	106,920
R2	0.46	0.48	0.50	0.45	0.47	0.49

Note: The table reports coefficient estimates from generalized difference-in-difference regressions for the post-definition change period. *SE*s, clustered by city, are reported in parentheses. The dependent variables are listed at the top of each column. Drunk accidents is the number of fatal accidents involving any drunk drivers. Drunk fatalities is the total number of fatalities in all drunk accidents. Post × $\text{Treated}_{t,c}$ is an indicator taking the value one if city *c* adopted at least one ridehailing service at time *t*. We include controls for the natural logarithm of population and the natural logarithm of income per capita in all models. *** (1%), ** (5%), * (10%) indicate the level of statistical significance.

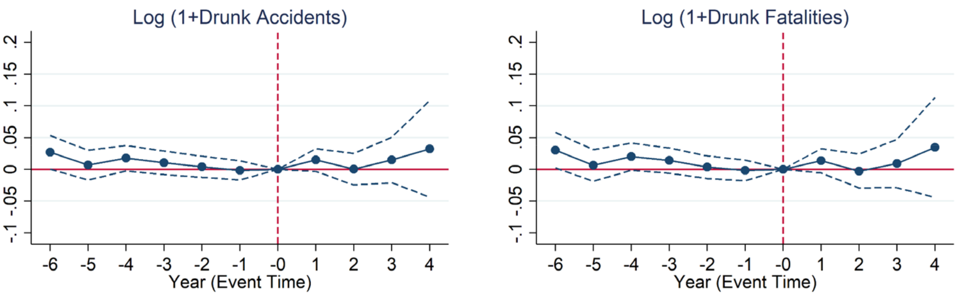

FIGURE 7 Difference-in-differences estimators: Drunk accidents and fatalities. The figure graphs the regression coefficient estimates and their 95% confidence intervals. The confidence intervals are constructed based on city level clustered *SE*s. We run regressions of our drunk accident measures on lag and lead indicators for the entrance of ridehailing. We bunch lags and leads by four quarters. The vertical red line indicates the year of entry. We provide a description of the variables in Section 3.

a statistical life. Assuming RH services are eventually made available across the entire United States, we can do a back-of-the-envelope calculation of the costs of the increase in accidents we document holding all else constant. In 2010, the year before RH began, there were 32,885 motor vehicle fatalities in the United States.[24] The 3% annual increase associated with the introduction of RH in fatalities represents an additional 987 lives lost each year.[25] The US Department of Transportation estimates the value of a statistical life (VSL) at $9.6 million for 2015; the DOT recommends analysts use a test range of $5.4 million (low) to $13.4 million (high) in 2015 dollars.[26] Applying the VSL and assuming an annual increase of 987 lives lost per year, the annual cost of the increase in fatalities associated with RH can be estimated as roughly $9.48 billion per year, with a range of $5.4 billion to $13.24 billion. A comparison of our cost estimate with Cohen et al. (2016) estimates of consumer surplus generated by RH services suggests that the costs from the increase in fatal accidents match or surpass the benefits of convenience to direct consumers of RH. The incremental cost derives from the externalities associated with driving and traffic congestion, where riders of RH do not bear the full cost of being on the road—some of this cost is borne by pedestrians, as we document above. Our estimates, moreover, do not include the costs imposed by nonfatal accidents, for which data are not readily available.

7.2 | Managerial and policy implications and suggestions

Our study provides insights that may be relevant to multiple stakeholders in the RH ecosystem. First, RH companies can choose to implement policies and interventions that change driver behavior to reduce the effects we document. Second, policy makers can enact regulation to prevent scenarios which may be contributing to the effect. In this section, we discuss possible implications and solutions for operations managers at RH companies and for policy makers seeking to improve road safety in their regions.

7.2.1 | Suggestions for RH operations managers

Our study highlights that efficiencies associated with optimization of RH platforms may also come with associated social costs. Operations managers at RH companies may wish to balance operational efficiency against the safety implications of their operational choices. A number of possible changes that help to address both the quantity and congestion channel and driver quality channel come to mind.

First, RH operations managers could also take a number of steps to improve driver quality and reduce distracted driving. At the driver level, mandating additional certification, training, and monitoring for RH drivers could serve to address quality problems, particularly for less active or newer drivers and for those who drive for long periods of time without breaks. Limiting hours working on the app, or mandating periodic rest periods, may also serve to avoid fatigued driving.[27] Driving record background checks, limitation of driver eligibility to those with no record of driving violations in the last 5 years, and periodic rechecking of driver records could help eliminate lower quality drivers from the labor pool. Periodic drug and blood alcohol level tests might also prove fruitful in reducing accidents. Active monitoring of drivers might also facilitate an improvement in driving habits. For example, monitoring of the sort currently provided (on a voluntary basis) by insurance companies such as Allstate and Root, and which records speeding, hard braking, and other dangerous driving behaviors, combined with rewards (penalties) for good (bad) driving behavior, could incentivize drivers to be less aggressive on the road. Electronic monitoring is common in the trucking industry, for example, to address both driving quality and time spent on the road, which may lead to fatigue.

Operations managers could consider modifications to navigation algorithms at the platform level to promote better driving safety. For example, the in-app navigation tools could route drivers in a manner that would reduce contact with pedestrians and avoid certain driving aspects that could lead to higher chances of an accident. This could include avoidance of streets with high pedestrian and bicycle usage, avoidance of left turns immediately following a pickup on the right hand side of the street, or left turns in general, avoidance of routes that require a driver to cross many lanes of traffic, and directing drivers around congested neighborhoods and roads, even at the expense of a longer ride. Such changes would likely need to be accompanied by a requirement that drivers follow the native in-app navigation and penalize drivers who deviate from the in-app provided routing.

7.2.2 | Suggestions for policy makers

Policy and regulation could also serve to both encourage the above types of actions by RH operators and provide other opportunities for reducing risk of accidents or internalizing the negative safety externality of RH. More generally, while RH companies could voluntarily alter operations to reduce the risk of increased congestion and accidents, to the extent such changes affect revenues and profitability, they may optimally choose not to. Moreover, for some types of interventions, such as driving hours limits, rules that apply to a single platform may not solve the problem: drivers reaching the hour limits on Uber may simply switch over to Lyft for the remainder of their driving day. As a result, regulation of RH operations likely offers the best solutions to address the congestion and accident externalities.

One straightforward possibility is the imposition through regulation of restrictions on the amount of RH activity permitted during particularly high congestion times, such as rush hour. Alternatively, a time-varying congestion tax applied to RH trips that is dependent on traffic congestion conditions (either in real time or based upon historical patterns) would likely reduce the demand for RH services during times when the largest number of vehicles are already on the road, reducing congestion and accordingly, reducing the likelihood of vehicular accidents. Current optimal pricing approaches in the operations literature focus on optimal matching of riders and drivers without accounting for the potential need to underserve demand in certain situations in order to minimize externalities. In similar vein, options such as disallowing the provision of low pricing for pooled rides if there is only a single party in the vehicle rather than a true carpool of multiple parties may be useful in addressing potential added demand for RH that stems from low pricing versus alternatives such as public transport, biking, or walking.

Policy makers may also wish to consider changes that seek to prevent accidents caused by RH vehicles stopping to pick up or dropoff riders in the middle of a busy street.

For example, cities may wish to define designated pickup and dropoff zones on each street that allow for safe pull over by RH vehicles, and mandate that RH vehicles pull over into designated zones for pick up and dropoff. Enhancing enforcement of ticketing and fines for RH drivers who block crosswalks and bicycle lanes, or who engage in unsafe dropoff or pickup behavior, might also prove beneficial. Limitations of the type placed on commercial drivers may also be beneficial. For example, many states regulate the number of hours commercial drivers may work consecutively, after which a long rest period is mandated. Placing similar limits on the number of hours a driver may work on gig platforms in a given 24-h period, mandating periodic rest periods, and requiring long breaks after long driving sessions may be beneficial in reducing fatigued driving. Importantly, policy makers must draft policy in a manner that addresses the fact that many gig drivers drive for multiple platforms in a given 24-h period; mandating that each platform separately place limits may still allow drivers to drive for many hours in a row by simply switching platforms when they hit an individual platforms time limit. Mandates should cover hours of gig driving more broadly, and may need to require platforms to report working hours for drivers to a centralized database, to allow for enforcement and penalization of drivers who violate the driving limits. In similar vein, policy makers may opt to prohibit simultaneous concurrent use of multiple platforms by drivers.

Additionally, limiting RH platform operations under specific conditions may help reduce the congestion effects of RH services. Policy makers could choose to limit surge incentives in certain areas or at times of high traffic congestion. Similarly, policy makers could choose to limit service or eliminate it entirely in certain areas during periods of high road congestion. Finally, policy makers could choose to impose taxation schemes that would effectively limit low pricing for Pool rides that ultimately only have a single rider, reducing substitution from other modes of transportation.

Finally, one potential suggested solution to the allocation problem imposed by the increase in accidents associated with RH is insurance. Insurance alters the economic incentives of agents by internalizing some of the costs. Here, however, the externalities will only be internalized in the presence of a compulsory insurance requirement. To date, RH companies such as Uber self-insure, and as a result, may not fully internalize this cost in the price of their service. Moreover, IPO and public market pressures may cause RH firms to under-insure, as provision for this cost may deteriorate financial performance. Furthermore, some of the costs involved under the current regime will be imposed on non-RH users/drivers, through increases in insurance premiums that result from the overall increase in accident probability. Insurance mandates can help force RH companies to internalize the full externality.

7.3 | Limitations

This study has a number of limitations. First, we cannot directly observe RH adoption in cities or the geographical distribution within cities. Instead, we infer adoption using the date of entry in the city and infer intensity using a proxy (google searches) that has been validated by prior literature. Second, our analysis is conducted on the initial years of RH operations. More specifically, we observe the cities in the sample for up to 5 years after the introduction of ride hailing. Our results are thus short to medium term in nature. Future research will be necessary to examine the long-term effects of these services as cities respond to their introductions. Third, although our estimates suggest there was no selection on accident incidence in RH entry into cities, it is important to note that the results of this work are not based on a randomized trial. Fourth, to the degree that limited information is available about the drivers of vehicles involved in the crashes, we are unable to uncover which populations and subpopulations are influenced to the greatest degree based on race, gender, age, or socioeconomic status. Future research may be able to link data to available information on such factors, providing further insights on where the impacts of RH on safety are strongest. Finally, while we document the traffic congestion and safety negative externalities resulting from the introduction of RH, we do not attempt to measure other potential adverse effects (e.g., fair wages, employment issues, or direct patron safety). Additionally, we do not attempt to quantify the potential positive externalities that may emerge from the introduction of ride-sharing platforms.

7.4 | Contributions and directions for future research

As noted above, our study documents only one particular social cost associated with RH, much as Cohen et al. (2016) documents a particular type of surplus. Our findings, however, suggest significant additional costs beyond the loss of life associated with increased traffic fatalities. Nationally, the number of traffic accidents in which individuals are injured is an order of magnitude higher than the number of those in which there is a fatality. Detailed data on such accidents is generally unavailable, but the economic and societal cost of all accidents in the United States in 2010 totaled $836 billion (https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812013). An increase in non-fatal accidents is also likely to be present here, with large associated societal costs.

Additionally, our findings suggest an increase in road utilization and congestion that imposes additional costs on society. While an increase in congestion may impose incremental costs on individuals, it can impose much greater costs on first responders and those being assisted by them (Beland & Brent, 2018). For example, as congestion increases, a higher proportion of ambulance transporters may encounter a delay in receiving life-saving medical attention. The disutility of the externality imposed by congestion is heterogeneous, however, unlike, say, the case of congestion in broadband telecommunication services, and thus it is not straightforward to solve this with differential pricing. Other costs, such as pollution, also exist, and are not solved by potential offered solutions such as insurance.

That said, advances in big data processing and machine learning, combined with the potential for unrestricted access to more granular data held by RH companies, insurers, and federal and local governments, could allow researchers to better evaluate the implementation of RH operations on society as well as design optimal implementations of such policies incorporating the safety externalities in said operations.

7.5 | Conclusion

Beginning in the mid-1980s the United States experienced a dramatic decrease in fatal accidents per capita and per vehicle mile driven. In 2010, 32,885 people died in motor vehicle traffic crashes in the United States—the lowest number of fatalities since 1949 (NHTSA, 2012). This decline halted and then reversed shortly after the introduction of RH into US cities. This increase has not been restricted to occupants of motor vehicles; the Governors Highway Safety Association recently noted that the 2018 pedestrian fatality pedestrian fatality figure was at its highest since 1990 and 35% higher than it was 10 years ago, reversing a longstanding trend of decline in pedestrian deaths from motor vehicle crashes.

While our documented effects alone are unlikely to fully explain the reversal of accident rate trends in recent years, they are worth further investigation and discussion. Moreover, while RH appears associated with more motor vehicle deaths, it does also bring many benefits. These include improved mobility for the disabled and minorities, flexible job opportunities that are especially valuable to those otherwise at high risk of unemployment, and customer convenience and resulting consumer surplus.

Still, the annual cost in human lives is nontrivial, and it is higher than estimates for annual consumer surplus generated. On top of this, our estimates do not include the costs imposed by nonfatal accidents, for which data is not readily available. We can assume that the pattern for fatal accidents is also repeated for nonfatal accidents, leading to costs in material and healthcare that may dwarf the costs in human lives. An essential contribution of our study is to point to the need for further research and debate about the overall cost–benefit tradeoff of RH and the operational and policy approaches to either increase the benefits or reduce the costs so as to be socially effective. Further research on this issue will likely necessitate unrestricted access to private data generated by RH companies.

ACKNOWLEDGMENTS

We thank Manuel Adelino, Marianne Bertrand, Jonathan Bonham, Eric Budish, Erik Brynjolfsson, Hans Christensen, Will Cong, Rebecca Dizon-Ross, Aaron Edlin, Michael Ewens, Mara Faccio, Austen Goolsbee, Shane Greenstein, Jonathan Hall, Sharique Hasan, Susan Helper (Discussant), Jessica Jeffers, Steve Kaplan, Emir Kamenica, Elisabeth Kempf, Ed Lazear, Christian Leuz, John List, Paul Oyer, David Robinson, Paola Sapienza, Rob Seamans, Scott Stern, Tom Wollmann, Luigi Zingales, and workshop participants at the University of Chicago Booth School of Business, Carnegie Mellon University Tepper School of Management, Duke University Innovation and Entrepreneurship Symposium, Rice University, UC Berkeley Law and Economics, and the NBER Entrepreneurship and Economics of Digitization Working Group meetings for helpful conversations, comments and suggestions. Parts of this research were conducted while Hochberg was visiting faculty at the University of Chicago. All errors are our own. Barrios gratefully acknowledges the support of the Stigler Center and the Centel Foundation/Robert P. Reuss Fund at the University of Chicago Booth School of Business.

ORCID

John M. Barrios https://orcid.org/0000-0003-3879-7118

ENDNOTES

[1] Similar numbers emerge from studies conducted by the Boston Metropolitan Area Planning Council (MAPC, 2018), the New York Department of Transportation (NYDOT, 2018), and other researchers (Circella et al., 2018; Clewlow & Mishra, 2017; Henao, 2017).

[2] From a supply perspective, a local report that examines detailed RH data in New York City suggests that RH

companies put 2.8 new vehicle miles on the road for each mile of personal driving they eliminated (a 180% overall increase). Moreover, the same report suggests that RH has added 5.7 billion miles of annual driving in the Boston, Chicago, Los Angeles, Miami, New York, Philadelphia, San Francisco, Seattle, and Washington, DC, metro areas (Schaller, 2018). While pooling services, such as UberPool and LyftLine, can reduce the overall increase in vehicle miles, these modes of RH currently represent a relatively small (20%) share of overall rides.

[3] While this data does not distinguish accidents in which a RH driver-partner car was involved from those where one was not, from a policy perspective, this distinction is not critical, as we wish to explore how the introduction of RH as a phenomenon shapes total fatal accident rates.

[4] In the case of RH, whether there were associated negative externalities was not clear ex-ante. In a 2014 University of Chicago Initiatives on Global Markets (IGM) survey of a panel of 43 top academic economists, all the panelists either agreed or strongly agreed with the statement "Letting car services such as Uber or Lyft compete with taxi firms on equal footing regarding genuine safety and insurance requirements, but without restrictions on prices or routes, raises consumer welfare." Many commented on the contribution of competition to consumer welfare; none suggested any potential negative externalities (one Nobel Prize winner noted specifically that he did not see any externalities involved). The comments were consistent with the panelists considering private welfare, rather than social welfare. Here, we shed light on the potential social costs of RH.

[5] Our Online Appendix provides a simple conceptual model in which accident rates are a function of two elements that are impacted by the introduction of RH technology: the number of vehicle miles traveled (VMT) and the average quality of drivers.

[6] Henao (2017) reports statistics suggesting RH drivers only have passengers in the car 39% of the time and 59% of the miles they drive while active on the app. Schaller (2018), using detailed data from New York City, shows that RH drivers on average drive 2.8 miles while waiting for a fare, 0.7 miles to pick up the fare, and 5.1 miles with a passenger in the car, implying a 59% utilization rate.

[7] As noted by Vickrey (1968), Edlin (2003) and Edlin and Karaca-Mandic (2006) and others, with every mile driven by a driver, that driver exposes themselves and others to the risk of an accident. Notably, these effects are compounded by the congestion and pollution effects of driving; we leave this topic to future research.

[8] Examples of haphazard pickup and dropoff stoppage include blocking bike lanes and crosswalks, suddenly pulling over, not pulling over completely (blocking lanes), and similar. While taxis often engage in similar behavior, taxis are clearly labeled, such that other drivers and pedestrians may know to expect erratic driving.

[9] https://www.census.gov/content/dam/Census/data/developers/understandingplace.pdf

[10] In analysis in the Online Appendix, we further separate out accidents by whether or not a drunk individual was involved. Total Drunk Accidents is the total number of fatal accidents involving any drunk drivers. Total Drunk Fatalities is the total number of fatalities in all drunk-driver accidents. Total Non-Drunk Accidents is the total number of fatal accidents not involving any drunk drivers. Total Non-Drunk Fatalities is the total number of fatalities in all nondrunk-driver accidents. We discuss these further in the Online Appendix.

[11] http://www.pewinternet.org/fact-sheet/mobile/ and http://www.pewresearch.org/fact-tank/2017/03/22/digital-divide-persists-even-as-lower-income-americans-make-gains-in-tech-adoption/

[12] The crosswalk can be found at https://www.udsmapper.org/zcta-crosswalk.cfm. The crosswalk is recommended by Missouri Census Data Center, http://mcdc.missouri.edu/geography/zipcodes_2010supplement.shtml.

[13] We discuss the TAMU data construction methodology in further detail in the Online Appendix.

[14] In the Online Appendix, Table A1, we run a multinomial logit to predict entry, where the observation level is a city. The dependent variable takes the value of 0 if RH launched in the city in 2010, 2011 or 2012 (early entry), a value of 1 if RH launched in the city in 2013 or 2014 (middle entry) and a value of 2 if RH launched in the city in 2015 or 2016 (late entry). Population and per capita income load positively and significantly in predicting earlier entry. In contrast, the change in accident rates over the 3, 5 or 10 years prior to entry does not load significantly. In Table A2, we estimate a Cox proportional hazard rate model with similar conclusions.

[15] As we go out further in time from RH adoption, the number of observations available for estimation drops sharply, and thus, the estimates for longer time horizons are less precise.

[16] In Online Appendix Figure A2, we present versions of these graphs using quarterly estimates, with similar conclusions.

[17] In unreported estimations, we perform a small falsification exercise, using only the sample of never-treated cities, and regress our accident measures on the Google trend search volume. We include our control variables, city and year-quarter fixed effects, and city-specific linear (and quadratic) trends. We observe no relationship between search volume for RH related terms and accident rates.

[18] FARS defines a pedestrian as "any person not in or upon a motor vehicle or other vehicle."

[19] Excess fuel consumption and excess hours of travel are calculated as the difference between the observed fuel consumption or hours of travel and the free-flow fuel consumption or hours of travel. The free-flow speed is estimated using the speed at low volume conditions (for example, 10 PM–5 AM) for each roadway section and hour of the week. We describe these measures in detail in the Online Appendix Section 3.

[20] TAMU uses the Department of Transportation (DOT) urban area boundaries. DOT urban areas were adopted from Census urban areas but have slight adjustments for transportation purposes. See https://www.fhwa.dot.gov/planning/census_issues/archives/metropolitan_planning/faqa2cdt.cfm#q24 and https://www.fhwa.dot.gov/legsregs/directives/fapg/g406300.htm.

[21] In early August 2019, over a year after the initial draft of this paper, Uber and Lyft released a joint study in which they admitted that RH cars contribute to increased overall congestion in six major US cities (Boston, San Francisco, Washington DC,

Chicago, and Los Angeles). https://drive.google.com/file/d/1FIUskVkj9lsAnWJQ6kLhAhNoVLjfFdx3/view

[22] See e.g. https://fortune.com/2017/04/13/uber-drunk-drivers/.

[23] The fact that there is some increase is not surprising or contradictory to our findings. There is a robust legal industry dedicated to lawsuits involving drunk RH drivers. For the period August 2014 to August 2015 in CA alone, Uber faced fines for 151 separate violations for drivers driving under the influence (and that is for drivers who get caught; the number that manage to avoid detection is likely higher). RH companies in some states are exempted from drug and alcohol screening programs imposed in other commercial driving operations. Media reports (e.g., https://www.sfgate.com/business/article/California-tells-Uber-it-s-sloppy-about-11069749.php) suggest that impaired driving by RH driver-partners is widespread. Moreover, we note that given the overall increase in VMT, there should also be an increased chance that a sober driver might hit or be hit by any remaining impaired drivers on the road.

[24] https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811552

[25] We round the estimated number of fatalities to the nearest whole number.

[26] Department of Transportation Revised Value of a Statistical Life Guidance (2016). Accessible at: https://www.transportation.gov/sites/dot.gov/files/docs/2016%20Revised%20Value%20of%20a%20Statistical%20Life%20Guidance.pdf

[27] Uber enacted one such limit a few years ago: "Starting on Wednesday, May 2, 2018, the driving hours limit feature will put a time limit on how long you can drive on Uber before having to take a break. The app will notify you of the need to take a break after 12 h of driving. You will then have to be offline for an entire 6 h period before the timer resets." See https://www.uber.com/en-ZA/blog/driving-hours-limit/.

REFERENCES

Afeche, P., Liu, Z., & Maglaras, C. (2018). Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Rotman School of Management Working Paper*, *3120544*, 18–19.

Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review*, *104*(9), 2763–2796.

Angrist, J. D., & Pischke, J.-S. (2008). Mostly harmless econometrics. In *Mostly harmless econometrics*. Princeton University Press.

Ashforth, B. E., & Reingen, P. H. (2014). Functions of dysfunction: Managing the dynamics of an organizational duality in a natural food cooperative. *Administrative Science Quarterly*, *59*(3), 474–516.

Autor, D. H. (2003). Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics*, *21*(1), 1–42.

Babar, Y., & Burtch, G. (2020). Examining the heterogeneous impact of ride-hailing services on public transit use. *Information Systems Research*, *31*(3), 820–834.

Bai, J., So, K. C., Tang, C. S., Chen, X., & Wang, H. (2019). Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management*, *21*(3), 556–570.

Banerjee, S., Freund, D., and Lykouris, T.. Pricing and optimization in shared vehicle systems: An approximation framework. Proceedings of the 2017 ACM Conference on Economics and Computation, 517, 2017.

Banerjee, S., Kanoria, Y., & Qian, P. (2018). State dependent control of closed queueing networks. *ACM SIGMETRICS Performance Evaluation Review*, *46*(1), 2–4.

Beland, L.-P., and Brent, D.. Traffic congestion and the performance of first responders: Evidence from California fire. (2018).

Benjaafar, S., Ding, J.-Y., Kong, G., & Taylor, T. (2022). Labor welfare in on-demand service platforms. *Manufacturing & Service Operations Management*, *24*(1), 110–124.

Benjaafar, S., & Ming, H. (2020). Operations management in the age of the sharing economy: What is old and what is new? *Manufacturing & Service Operations Management*, *22*(1), 93–101.

Bertrand, M., & Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. *The Quarterly Journal of Economics*, *116*(3), 901–932.

Besbes, O., Castro, F., & Lobel, I. (2021). Surge pricing and its spatial supply response. *Management Science*, *67*(3), 1350–1367.

Bimpikis, K., Candogan, O., & Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*, *67*(3), 744–769.

Birkinshaw, J., & Gupta, K. (2013). Clarifying the distinctive contribution of ambidexterity to the field of organization studies. *Academy of Management Perspectives*, *27*(4), 287–298.

Blackwell, M., Iacus, S., King, G., & Porro, G. (2009). Cem: Coarsened exact matching in Stata. *The Stata Journal*, *9*(4), 524–546.

Borusyak, K., and Jaravel, X.. Revisiting event study designs. Available at SSRN 2826228 (2017).

Bouchery, Y., Corbett, C. J., Fransoo, J. C., & Tan, T. (Eds.). (2016). *Sustainable supply chains: A research-based textbook on operations and strategy* (Vol. 4). Springer.

Braverman, A., Dai, J. G., Liu, X., & Ying, L. (2019). Empty-car routing in ridesharing systems. *Operations Research*, *67*(5), 1437–1452.

Brazil, N., & Kirk, D. S. (2016). Uber and metropolitan traffic fatalities in the United States. *American Journal of Epidemiology*, *184*, 192–198.

Buchak, G., Financing the gig economy. University of Chicago Booth Working Paper, 2018.

Burgdorf, J., Lennon, C., and Teltser, K.. Do ridesharing services increase alcohol consumption?, 2019.

Cachon, G. P., Daniels, K. M., & Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, *19*(3), 368–384.

Cantor, D. E., Corsi, T. M., & Grimm, C. M. (2009). Do electronic logbooks contribute to motor carrier safety performance? *Journal of Business Logistics*, *30*(1), 203–222.

Circella, G., Alemi, F., Tiedeman, K., Handy, S., and Mokhtarian, P., The adoption of shared mobility in California and its relationship with other components of travel behavior, Institute of Transportation Studies, University of California, Davis, 2018.

Clewlow, R. R., and Mishra, G. S., Disruptive transportation: The adoption, utilization and impacts of ride-hailing in the United States. Institute of Transportation Studies, University of California, Davis, 2017.

Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R., Using big data to estimate consumer surplus: The case of Uber, NBER Working Paper No. w22627, 2016.

Cook, C., Diamond, R., Hall, J., List, J. A., and Oyer, P., The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers, Stanford Graduate School of Business, 2018.

Cramer, J., & Krueger, A. B. (2016). Disruptive change in the taxi business: The case of Uber. *American Economic Review*, *106*, 177–182.

Currie, J., & Walker, R. (2011). Traffic congestion and infant health: Evidence from E-ZPass. *American Economic Journal: Applied Economics*, *3*(1), 65–90.

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

Di Tella, R., & Schargrodsky, E. (2003). The role of wages and auditing during a crackdown on corruption in the city of Buenos Aires. *The Journal of Law and Economics*, *46*(1), 269–292.

Dingus, T. A., Guo, F., Lee, S., Antin, J. F., Perez, M., Buchanan-King, M., & Hankey, J. (2016). Driver crash risk factors and prevalence evaluation using naturalistic driving data. *Proceedings of the National Academy of Sciences*, *113*(10), 2636–2641.

Duflo, E., Hanna, R., & Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review*, *102*(4), 1241–1278.

Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from US cities. *American Economic Review*, *101*(6), 2616–2652.

Edlin, A. S. (2003). Per-mile premiums for auto insurance. In R. Arnott, B. Greenwald, R. Kanbur, & B. Nalebuff (Eds.), *Economics for an imperfect world: Essays in honor of Joseph E. Stiglitz*. MIT Press.

Edlin, A. S., & Karaca-Mandic, P. (2006). The accident externality from driving. *Journal of Political Economy*, *114*, 931–955.

Ewing, R., & Hamidi, S. (2014). *Measuring urban sprawl and validating sprawl measures*. National Institutes of Health and Smart Growth America.

Fahnenschreiber, S.,¨undling, F. G., Keyhani, M. H., & Schnee, M. (2016). A multi-modal routing approach combining dynamic ride-sharing and public transport. *Transportation Research Procedia*, *13*, 176–183.

Feng, G., Kong, G., & Wang, Z. (2020). We are on the way: Analysis of on-demand ride-hailing systems. *Manufacturing & Service Operations Management*, *23*, 1237–1256.

Ferdows, K., & De Meyer, A. (1990). Lasting improvements in manufacturing performance: In search of a new theory. *Journal of Operations Management*, *9*(2), 168–184.

Gee, G. C., & Takeuchi, D. T. (2004). Traffic stress, vehicular burden and well-being: A multilevel analysis. *Social Science & Medicine*, *59*(2), 405–414.

Goodman-Bacon, A.. Difference-in-differences with variation in treatment timing. NBER Working Paper No. w25018, 2018.

Gottholmseder, G., Nowotny, K., Pruckner, G. J., & Theurl, E. (2009). Stress perception and commuting. *Health Economics*, *18*(5), 559–576.

Greenwood, B. N., & Wattal, S. (2017). Show me the way to go home, an empirical investigation of ride-sharing and alcohol related motor vehicle fatalities. *MIS Quarterly*, *41*, 163–187.

Haggag, K., McManus, B., & Paci, G. (2017). Learning by driving: Productivity improvements by New York City taxi drivers. *American Economic Journal: Applied Economics*, *9*, 70–95.

Hall, J. D., Palsson, C., & Price, J. (2018). Is Uber a substitute or complement for public transit? *Journal of Urban Economics*, *108*, 36–50.

He, L., Mak, H.-Y., Rong, Y., & Shen, Z.-J. M. (2017). Service region design for urban electric vehicle sharing systems. *Manufacturing & Service Operations Management*, *19*(2), 309–327.

Henao, A., Impacts of ridesourcing – Lyft and Uber – on transportation including VMT, mode replacement, parking, and travel behavior, 2017.

Hennessy, D. A., & Wiesenthal, D. L. (1999). Traffic congestion, driver stress, and driver aggression. *Aggressive Behavior: Official Journal of the International Society for Research on Aggression*, *25*(6), 409–423.

Hickman, J. S., & Hanowski, R. J. (2011). Use of a video monitoring approach to reduce at-risk driving behaviors in commercial vehicle operations. *Transportation Research Part F: Traffic Psychology and Behaviour*, *14*(3), 189–198.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, *81*(396), 945–960.

Hu, M., and Zhou, Y.. Price, wage, and fixed commission in on-demand matching. Available at SSRN 2949513 (2020).

Imbens, G. W., & Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, *47*(1), 5–86.

Kleindorfer, P. R., Singhal, K., & van Wassenhove, L. N. (2005). Sustainable operations management. *Production and Operations Management*, *14*(4), 482–492.

Knittel, C. R., Miller, D. L., & Sanders, N. J. (2016). Caution, drivers! Children present: Traffic, pollution, and infant health. *Review of Economics and Statistics*, *98*(2), 350–366.

Kong, H., Zhang, X., & Zhao, J. (2020). How does ridesourcing substitute for public transit? A geospatial perspective in Chengdu, China. *Journal of Transport Geography*, *86*, 102769.

Liu, S., He, L., & Shen, Z.-J. M. (2021). On-time last-mile delivery: Order assignment with travel-time predictors. *Management Science*, *67*(7), 4095–4119.

Mak, H.-Y. (2022). Enabling smarter cities with operations management. *Manufacturing & Service Operations Management*, *24*(1), 24–39.

Martin-Buck, F., Driving safety, an empirical analysis of ridesharing's impact on drunk driving and alcohol-related crime, University of Texas at Austin, 2017.

Miller, J., Saldanha, J. P., Manus Rungtusanatham, A., Knemeyer, M., & Goldsby, T. J. (2018). How does electronic monitoring affect hours-of-service compliance? *Transportation Journal*, *57*(4), 329–364.

Nagin, D. S., Rebitzer, J. B., Sanders, S., & Taylor, L. J. (2002). Monitoring, motivation, and management: The determinants of opportunistic behavior in a field experiment. *American Economic Review*, *92*(4), 850–873.

Naumov, S., Keith, D. R., & Fine, C. H. (2020). Unintended consequences of automated vehicles and pooling for urban transportation systems. *Production and Operations Management*, *29*(5), 1354–1371.

Ossokina, I. V., & Verweij, G. (2015). Urban traffic externalities: Quasi-experimental evidence from housing prices. *Regional Science and Urban Economics*, *55*, 1–13.

Ostrovsky, M., and Schwarz, M.. Carpooling and the economics of self-driving cars. Proceedings of the 2019 ACM Conference on Economics and Computation, 581–582. 2019.

Owens, J. M., Dingus, T. A., Guo, F., Fang, Y., Perez, M., and McClafferty, J. (2018). Crash risk of cell phone use while driving: A case-crossover analysis of naturalistic driving data.

Özkan, E., & Ward, A. R. (2020). Dynamic matching for real-time ride sharing. *Stochastic Systems*, *10*(1), 29–70.

Pagell, M., Klassen, R., Johnston, D., Shevchenko, A., & Sharma, S. (2015). Are safety and operational effectiveness contradictory requirements: The roles of routines and relational coordination. *Journal of Operations Management*, *36*, 1–14.

Pagell, M., Parkinson, M., Veltri, A., Gray, J., Wiengarten, F., Louis, M., & Fynes, B. (2020). The tension between worker safety and organization survival. *Management Science*, *66*(10), 4863–4878.

Posen, H. A. (2015). Ridesharing in the sharing economy: Should regulators impose Uber regulations on Uber. *Iowa Law Review*, *101*, 405.

Qi, W., Li, L., Liu, S., & Shen, Z. J. M. (2018). Shared mobility for last-mile delivery: Design, operational prescriptions, and environmental impact. *Manufacturing and Service Operations Management*, *20*(4), 737–751.

Qi, W., & Shen, Z.-J. M. (2019). A smart-city scope of operations management. *Production and Operations Management*, *28*(2), 393–406.

Rosenzweig, E. D., & Easton, G. S. (2010). Tradeoffs in manufacturing? A meta-analysis and critique of the literature. *Production and Operations Management*, *19*(2), 127–141.

Schaller, B. (2018). The new automobility: Lyft, Uber and the future of American cities. *Schaller Consulting Report*.

Schrank, D., Lomax, T., and Eisele, B. (2011). 2011 urban mobility report.

Scott, A., Balthrop, A., & Miller, J. W. (2021). Unintended responses to IT-enabled monitoring: The case of the electronic logging device mandate. *Journal of Operations Management*, *67*(2), 152–181.

Siddartha, B., Johari, R. and Riquelme, C.. Pricing in ride-sharing platforms: A queueing-theoretic approach. Proceedings of the 16th ACM Conference on Economic Computing (2015), 639.

Skinner, W. (1996). Manufacturing strategy on the “S” curve. *Production and Operations Management*, *5*(1), 3–14.

Smith, W. K., & Lewis, M. W. (2011). Toward a theory of paradox: A dynamic equilibrium model of organizing. *Academy of Management Review*, *36*(2), 381–403.

Song, Y., & Huang, Y. (2020). Investigating complementary and competitive relationships between bikeshare service and public transit: A spatial-temporal framework. *Transportation Research Record*, *2674*(1), 260–271.

Vickrey, W. (1968). Automobile accidents, Tort law, externalities, and insurance: An economist's critique. *Law and Contemporary Problems*, *33*, 464–487.

Ylikoski, P., & Emrah Aydinonat, N. (2014). Understanding with theoretical models. *Journal of Economic Methodology*, *21*(1), 19–36.

Yu, J. J., Tang, C. S., Shen, Z.-J. M., & Chen, X. M. (2020). A balancing act of regulating on-demand ride services. *Management Science*, *66*(7), 2975–2992.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Barrios, J. M., Hochberg, Y. V., & Yi, H. (2023). The cost of convenience: Ridehailing and traffic fatalities. *Journal of Operations Management*, *69*(5), 823–855. https://doi.org/10.1002/joom.1221