THE COST OF CONVENIENCE: RIDEHAILING AND TRAFFIC FATALITIES*

John M. Barrios
Washington University in St. Louis & NBER

Yael V. Hochberg
Rice University & NBER

Hanyi Yi
Boston College

Current Version: April 2022

We examine the effect of the introduction of ridehailing in U.S. cities on fatal traffic accidents. The arrival of ridehailing is associated with an increase of approximately 3% in the number of fatal accidents, for both vehicle occupants and pedestrians. Consistent with ridehailing increasing road usage, we find that introduction is associated with an increase 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 a variety of operational as well as policy prescriptions for regulation of ridehailing operations that may help limit such externalities.

* We thank the Editor, Associate Editor, and referees, 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, John Gray, Shane Greenstein, Jonathan Hall, Sharique Hasan, Susan Helper, 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. Corresponding Author: John Barrios (john.barrios@wustl.edu), Washington University in St Louis, Olin School of Business.

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 major traffic accidents and traffic fatalities 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 and 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, 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 and Mishra, 2017).[1] This evidence runs counter to the naïve view that utilization remains fixed, and suggests that RH will increase vehicle miles travelled (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 US will lead to an increase in fatal accidents and fatalities.

Using the staggered introduction of ridehailing across U.S. 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 generalized difference-in-differences (DD) specification with fixed effects for location and time (quarter-year) and location-specific linear and quadratic trends. Our DD specification allows us to control for

[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 (Clewlow and Mishra, 2017; Henao, 2017; Circella et al., 2018).

[2] From a supply perspective, a local report that examines detailed ridehailing data in New York City suggests that ridehailing 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 ridehailing has added 5.7 billion miles of annual driving in the Boston, Chicago, Los Angeles, Miami, New York, Philadelphia, San Francisco, Seattle, and Washington, D.C., metro areas (Schaller, 2018). While pooling services, such as UberPool and LyftLine, can reduce the overall increase in vehicle miles, these modes of ridehailing 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.

macroeconomic changes, such as the Great Recession, fuel costs, and improvements in vehicle technology, city-specific conditions such as average weather patterns and city topology, a variety of city-level controls, as well as location-specific pre-trends in accidents that existed prior to the arrival of RH.

We report several findings. First, we document a 2% to 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 ridehailing. 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 ridehailing 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 understanding of the societal impacts of RH platforms. While prior OM studies on on-demand service platforms have focused on RH price optimization problems (Bai et al. 2019; Banerjee et al., 2015; Banerjee et al., 2017; Besbes et al. 2021), market competition (Besbes et al. 2021; Yu et al. 2020), and the matching of supply and demand (Feng et al. 2020; Ozkan and Ward, 2020). There is little literature on the societal impact and sustainability of on-demand RH services. Our study provides one of the first large-sample investigations of the potential negative externalities of ridehailing. 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 and Hu 2020, He et al. 2017) suggest that platforms of this nature have the potential to alleviate the last mile problem of city commutes. Our work suggests 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). 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). 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. Our analysis provides unique insights into the impacts of new app-based platform services that were conceived with the aim of increasing 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 literature on sustainable operations research (Bouchery et al. 2017, Kleindorfer et al. 2005) by providing empirical evidence for the impact of ridesharing services on traffic congestion and safety. The theoretical potential effects of ridesharing services on the smart-city movement and sustainability have been discussed in the literature (Qi and Shen 2019, Mak 2022), 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 create a significant negative impact on the smart-city movement and 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 and Kirk, 2016; Martin-Buck, 2017; Greenwood and Wattal, 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 the sample to a single state, and the unit of observation, which confounds urban (where our effects are concentrated) and rural areas in the same observations, likely drives the difference between their results and our findings. In contrast to Greenwood and Wattal (2017), our analysis of alcohol-related traffic fatalities over the entire US and a longer sample period suggests no consistent significant change in drunk-driving related fatal

[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.

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. The advent of such on-demand service platforms has 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. (2012) 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), using a theoretical queueing model, 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. (2019a) 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, and show that it may be optimal to strategically reject demand at low demand locations to induce repositioning of 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 that under traditional street-hailing systems. They find that the on-demand matching mechanism can result in lower efficiency. Ozkan 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 nonmonotonic in the labor pool size. While these studies focus on the optimization of a single platform's operation from the perspective of the company and its customers, our theoretical framework incorporates the interaction of ride-hailing companies and the general public's safety.

Our paper contributes to this literature by proposing 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 (1969), and persists even today. This literature approaches tradeoffs from a number of perspectives, including that of whether such tradeoffs exist (e.g. Ferdows and De Meyer, 1990; Rosenzweig and Easton, 2010), and how managers should balance competing objectives (e.g. Smith and Lewis, 2011; Birkinshaw and Gupta, 2013; Ashforth and Reingen 2014; Pagell et al, 2015). 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, operation 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 U.S. was close to $121 billion in 2011 (Schrank et al., 2012). 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 (Hennessy and Wiesenthal, 1999; Gee and Takeuchi, 2004; Gottholmseder et al., 2009). 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 traffic reductions due to the introduction of electronic toll collection (E-ZPass) reduce vehicle emissions near highway toll plazas, which reduces prematurity and low birth weight among mothers near the toll plaza. Beland and Brent (2018) examine the relationship between 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 one of the many public goods that rely on well-functioning road infrastructure. Ossokina and Verweij (2015) exploit a quasi-experiment that reduces traffic congestion on certain streets in the Netherlands and find that the decrease in traffic leads to an increase in housing prices. Other related work includes Duranton and Turner (2011), who explore what they term "the fundamental law of road congestion," concluding that the provision of 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. Nagin et al., 2002; Di Tella and Schargrodsky, 2003; Banerjee et al. 2018; Duflo et al. 2012). 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 and Hanowski, 2011; Miller at el. 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.

2.3. The Effects of Ridehailing on Accidents

Before the advent of RH, private for-hire transportation consisted primarily of taxi services, limousines, or bus and van services. All of these services, with the exception of taxis, required advance reservations. All were relatively costly, with the number of vehicles available for hire varying substantially across locations. The taxi industry was heavily regulated by local municipalities, which placed restrictions on the number of vehicles available for hire, the rates charged, and the licensing and insurance requirements for drivers and vehicles. Quantity restrictions, in particular, were thought to lead to inconveniences for riders given the shortage of taxis during periods of high demand.

Uber was the first RH firm in the United States, launching in San Francisco in May 2010, and was followed two years later by Lyft and Sidecar. RH then expanded rapidly across the country. By the end of 2014, RH firms operated in 80% of U.S. cities with a population of 100,000 or more. Much of the spread in RH was driven by the convenience for users, stemming from new technology easing the matching of riders and drivers and enabling seamless payment through an app. RH firms' exemption from (or willful disregard for) taxi and livery regulations allowed them to expand supply during high demand periods, and adjusted prices in real time to encourage more riders and drivers to participate in the market. This has in turn engendered backlash from advocacy groups and policymakers concerned with the effects of RH technology in their cities.[5]

We begin by outlining why RH may lead to an increase in fatal traffic accidents and fatalities. 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] Here, we verbally explain the intuitions behind the model and state the associated hypotheses we test.

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). The number of VMTs can further be broken down into three sub-categories:

i. the number of VMTs generated by people driving themselves from origin to destination
ii. the number of VMTs generated by RH drivers carrying passengers from origin to destination
iii. the number of VMTs generated by RH drivers while driving in-between passenger pickups.

Even if the VMT generated by a potential rider driving themselves to a destination is offset by the VMT generated if a RH driver drives them as people move toward RH vehicles. RH still introduces "between driving" (between fares, waiting for fares, going from fare location to fare location) in a city. Consequently, while VMT generated by people driving themselves is almost certainly decreased by the introduction of RH, the technology leads to the introduction of additional vehicle miles in the form of

[5] In many ways, ridehailing has become the modern poster child for the classic battle between what are argued to be outdated regulations, supported by rent-seeking incumbents, and the introduction of a welfare-enhancing technology. Many new technologies face frictions that slow diffusion (Grubler, 1991). Parente and Prescott (1994) argue that one such friction is resistance on the part of sectoral interests. Indeed, emphasizing barriers to technology adoption, economic historians, such as Rosenberg et al. (1986), argue that the reason why the West grew rich before the rest of the world was that active resistance to technology adoption was weaker there. Most economic histories of technological adoptions provide cases in which adoption was met with fierce resistance (Mokyr 1990).

[6] 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.

VMT driven by RH drivers and the VMT generated by RH drivers driving from fare to fare or between potential fare-generating neighborhoods. Thus, the effect of the introduction of RH in a city on the number of VMTs on the road depends on whether the decrease in personal driving is more than offset by the RH VMT and between-fare VMT generated by the introduction of RH.

Taking a naïve view of RH, each person who no longer chooses to drive themselves is now driven by a RH driver, thus precisely offsetting the effect on the overall VMT. However, unless there are absolutely no between-fare miles driven by a ride-sharing driver, we would expect to see an increase in overall VMTs after RH arrives. The limited evidence to date suggests that there is considerable between-fare travel by drivers. 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. Furthermore, RH companies initially offered subsidies designed to induce drivers to spend more time out on the road active in the app, so as to decrease wait time for passengers. Finally, while not the focus of their study, the analysis of Chen et al. (2017) is consistent with a mismatch between rider demand and the supply of drivers, particularly given the flexibility afforded to the drivers.

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 and 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 and Burtch 2020).

Of course, in some cities, at later dates, the option to "carpool" via RH was introduced, in the form of Uber Pool and Lyft Line. With the introduction of these services, the reduction in own drive car hours may not be fully offset by RH drive hours, as multiple people may be substituting away from driving themselves into a single RH car. While Uber and Lyft have both heavily invested in promoting their shared services, Uber reports that UberPool accounts for only 20% of trips in cities where it is offered, and Lyft reports that 37% of users in cities with LyftLine request a Line trip, and many trips are not matched, thus leaving a single rider (Schaller, 2018). Pooled rides are also cheaper, potentially inducing more substitution from other modes of transport. It is not clear what fraction of rides must be pooled to counteract the VMT driven between fares, but Schaller (2018) suggests that, even if half of the rides were pooled, total VMT would still increase. Furthermore, stepping away from the naïve model, survey evidence suggests that many riders are substituting away not from driving themselves but rather from other forms of transportation, including walking, biking, and, more importantly, public transportation (Clewlow and Mishra, 2017). Thus, it is likely that pooled ride adoption would need to be extremely high to offset such substitution effects.

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 ridehailing 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. Put another way, the introduction of RH makes it less costly to have someone else drive you, but also makes the gains from getting out on the road as a driver greater (as you can make money by doing so). Lower quality drivers, who in the absence of compensation may not have driven, now have an incentive to drive. Moreover, more affluent people are more likely to use RH (Pew Research Center, 2016), and the less affluent are more likely to become RH drivers. To the extent that this substitution leads to more vehicle miles driven in older or cheaper cars with fewer safety features, this may positively affect accident rates.[7] Ridehailing VMT is also different from self-driving VMT, as ridehailing 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 ridehailing vehicle).[8] Furthermore, RH drivers connect to passengers via apps on mobile phones, and navigate to both 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 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 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.[9]

Many indicators suggest that both total VMT and driver quality 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. A driver with more than 2,500 lifetime trips completed earns 14% more per hour than a driver who has completed fewer than 100 lifetime trips, in part because he learns where and when to drive, which may decrease the extra VMT driven between fares. 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 behavior, on the part of certain

[7] RH platforms generally allow vehicles up to ten years old, allowing for considerable differences in technology advancement within the eligible pool.

[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] Of course, RH drivers, especially those with more experience from more hours driven, may in fact represent improved quality. To the extent that the substitution goes the other way and lower quality drivers are substituted by better drivers, this may reduce accident rates if the increase in quality offsets the increase in VMT. In general, better (worse) drivers should reduce (increase) accidents, all else equal. The effect of RH on the quality of the average driver on the road, however, is ambiguous. If the quality of the average driver increases, this could offset the quantity effect above. If it decreases or does not change, the quantity effect will prevail. Which effect dominates, of course, is an empirical question.

driver populations, that directly or indirectly increases the probability of accidents, such as gaming time-and-distance pay algorithms by taking longer routes, speeding, etc. (Cook et al., 2018). As a result, we do not have a clear hypothesis regarding the persistence of the effects over time.

We expect that the effect of RH on accidents to also vary depending on the urban or rural nature of a location. Specifically, we expect that the effects of RH to be more heightened in urban areas than in suburban or rural areas. First, in dense urban areas, the number of people per mile and per mile of road is higher. As a result, traffic congestion is typically elevated to begin with, and 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 pullover 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.

Second, people are more likely to substitute public transit trips with point-to-point RH trips when they require less waiting time, are faster, and more convenient (Babar and 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 and Hamidi, 2014). Additionally, since the cost per mile for RH is higher than that for public transit, the shift from public transit to RH is more likely to occur for short-distance trips (Babar and Burtch 2020, Hall et al. 2018), which are more common in urban areas. Thus, substituting short-distance public transit trips with RH in urban areas may further increase RH VMT and between ride VMT, increasing congestion and the potential for accidents and fatalities.

Finally, in rural or more sparsely populated areas, road usage is typically lower, and the first and last mile problem is more serious, given the absence of large transit centers and the significant separation of land uses (Qi et al. 2018). As a result, RH may serve primarily to complement existing public transit options (e.g. airports, train stations) to address the (first or) last mile problem (Babar and Burtch 2020, Fahnenschreiber et al. 2016, Qi et al. 2018, Song and Huang 2020), resulting 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.

2.4. Formal Hypotheses

Overall, given that increased VMT should result in an increase in fatal accidents, and given that a decrease in average driver quality would also lead to an increase in fatal accidents, we hypothesize that:

> **H1A:** the introduction of RH in a city will lead to an increase in fatal accidents and fatalities

and that an increase in the use or adoption of RH should similarly leads to a higher number of accidents:

> **H1B:** fatal accidents and fatalities post RH introduction should be increasing in RH adoption intensity.

We expect a moderating role of an area's urban or rural nature on the relation between ridehailing and accidents. We posit that:

> **H2:** the effects of the introduction of RH on accidents will be larger in more urban areas.

We then focus on the mechanisms for the increase in fatal accidents: the quantity and quality channels. On the quantity side, we posit that the introduction of RH will lead to more road and car usage:

> **H3A:** the introduction of RH in a city will lead to an increase in VMT (extensive margin)

> **H3B:** the introduction of RH in a city will lead to an increase in new car registrations (intensive margin)

> **H3C:** 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.

On the quality side, our conceptual framework suggests that:

> **H4:** the introduction of RH in a city will lead to a decrease in driving quality.

3. DATA AND VARIABLES

Our sample consists of all incorporated "places"[10] in the continental United States with population greater than or equal to 10,000 in 2010.[11] Our sample covers the period 2001 to 2016; all results are robust to employing shorter pre-RH sample windows. Our list of incorporated places is obtained from the Census Bureau and covers all self-governing cities, boroughs, towns, and villages in the United States.[12] (For ease of interpretation, we interchangeably refer to these as "cities" or "locations" throughout the text.) Our observations are measured at the quarterly level. Thus, the sample contains 190,080 quarterly observations on 2,970 "places" from 2001 to 2016, among which 1,199 adopt RH prior to 2017. Figure 2 shows the diffusion of RH across the United States, by cities and population. Diffusion of RH across U.S. cities began slowly, accelerating rapidly after 2013. Diffusion by population follows a standard S-curve, consistent with general historical patterns of new technology diffusion.[13]

3.1. 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 not available, we use the city and state identifier codes to assign observations to the appropriate place. Geographic coordinates are present in 98% of FARS's observations, 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

---

[10] We use incorporated places, rather than Census Designated Places (CDPs), because CDP annual population estimates are not readily available, except by individual place download, whereas population data is available for incorporated places for mass download through the census.

[11] Some places in our sample had lower populations than 10,000 during the sample period, most notably during the period of 2001–2010. We impose the cutoff on population as measured in 2010. As an example, consider Hutto, Texas, a suburb of the Austin-RoundRock metro area. In 2001, Hutto had a population of 3,030, the lowest in our sample. By 2010, it had grown to over 14,000, mimicking the growth of the Austin metro area. As it has population above 10,000 in 2010, it is included in our sample. Our results are robust to permutations to this cutoff.

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

[13] In the Online Appendix, we further demonstrate the robustness of our results to using shorter pre-sample periods.

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.[14]

3.2. Ridehailing Launch and Driver Enrollment Intensity

Data on RH launch dates for each city are obtained directly from Uber and Lyft.[15] The companies provided dates of service launch for each type of service launched: (i) UberBlack/UberTaxi, which allows customers to hail a livery or taxi vehicle; (ii) UberX/Lyft, which allow customers to hail regular cars driven by driver-partners; and (iii) UberPool/Lyft Line, which allow customers to share a hailed vehicle with others. We merge these dates with Census Bureau's incorporated place directory in 2010.

While Uber and Lyft declined to provide data on driver enrollment and usage for this project, other researchers have shown a strong correlation between Google trends for searches for RH keywords and actual driver uptake (Cramer and Krueger, 2016). To measure the intensity of driver adoption, we thus follow the spirit of the work of Cramer and Krueger (2016) and Hall et al. (2018) and use Google searches for the terms "Uber," "Lyft," and "rideshare."[16] We note that while prior work 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.

Rather than using the standard Google Trends index, which scales results from 0 to 100 based on the most popular term entered, and that does not easily allow comparisons across geographic areas and time periods, we use data from the Google Health Trends API, which describes how often a specific search term is entered relative to the total search volume on Google search engine within a geographic region and time range, and returns the probability of a search session that includes the corresponding term, which makes comparisons across locations and time feasible.[17] We track trends for searches for these terms using the Google Health Trends API for all Nielsen Designated Market Areas (DMAs) at monthly frequency

[14] 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.

[15] Here, we use the exact cities indicated by Uber and Lyft, even if we suspect or believe that the launch covered adjacent cities as well (e.g., San Francisco launched in 2010, and there is no separate launch date for San Jose or Palo Alto). Since this means some places we include in our control may in fact be treated in later years in the sample as service expands slowly out beyond original boundaries, we are biasing against finding an effect of treatment.

[16] We use the freebase identifiers for term "Uber" (/m/0gx0wlr) and "Lyft" (/m/0wdpqnj). Freebase identifiers denote all searches that were classified to be about this topic.

[17] These probabilities are calculated on a uniformly distributed random sample of 10%-15% of Google web searches. Mathematically, the numbers returned from the Google Trend API can be officially written as:

$$Value_{[time, term\ restriction]} = P(term - restriction | time\ and\ geo - restriction) * 10M$$

This probability is multiplied by 10 million in order to be more human readable.

from January 2004 to December 2016. We aggregate the data to the quarter level and match the DMAs to Census incorporated places using a crosswalk provided by Nielsen. Thus, in the accident specifications that use log search volume as a proxy for driver enrollment intensity, we interpret the coefficients in terms of percentage change in search volume.

3.3. Other Data

We use several measures to explore heterogeneity by city characteristics and as control variables in our models. We obtain annual city population estimates and population density from the U.S. Census and annual county income per capita from the Bureau of Economic Analysis. Household vehicle ownership and means of transportation to work at the city level are gathered from the 2010 American Communities Survey. Controlling for population and per capita income which vary by time and location are of first order importance as they provide a proxy for specific concerning confounders. For example, 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.[18]

To explore mechanisms that may drive any change in accident rates upon arrival of RH, we use a variety of data sources. We obtain data on new car registrations by zip code on a monthly level from Polk Automotive. We aggregate the data at city and quarter level using UDS Mapper's zip code-to-ZCTA crosswalk[19] and Census' ZCTA-to-place crosswalk. We obtain estimates of city and freeway vehicle miles traveled, 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 U.S. 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.[20]

3.4. Summary Statistics

Table 1 presents summary statistics for our sample. The places average 54,500 in population, have an income per capita of $39,720, and a population density of roughly 3,000 people per square mile. Prior to the arrival of RH, 2.96% of residents in our average city/place used public transportation to commute,

---

[18] 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/
[19] 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.
[20] We discuss the TAMU data construction methodology in further detail in the Online Appendix.

10.6% commuted by carpool, and 33% owned vehicles. The average city in our sample had 670 new car registrations per year. As can be seen from the distributional statistics in the table, there is wide variation across all these characteristics across the sample. The table further presents summary statistics on rate (per 100,000 population) of accidents for the cities in our sample over the sample period. 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.

4. 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 did not arrive in all U.S. cities at the same time, but rather was rolled out gradually, at different times, in different geographies. Of course, RH companies' choice of cities to launch in first is unlikely to be random. However, it is unlikely that Uber and Lyft were specifically selecting cities to roll out services 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.[21] 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 employ a standard generalized difference-in-differences approach that takes advantage of the fact that RH was rolled out across US cities at different times. In contrast to standard difference-in-differences frameworks, which compare 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 and Mullinathan 2001). Formally, we index cities by c and time by t, and estimate models of the form:

$$\log(1 + accidents_{t,c}) = \propto_c + \gamma_t + \beta' X_{t,c} + \theta_c t + \delta POSTxTREATED_{t,c} + \varepsilon_{t,c}, \qquad (1)$$

[21] 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). While both population and per capita income load positively and significantly in predicting earlier entry, 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 and similarly show that while population and income strongly predict entry timing, there is no statistically significant loading on trends in fatal accidents.

where $accidents_{t,c}$ is our measure of accidents in city c in quarter t, $POSTxTREATED_{t,c}$ is a dummy variable that equals one if 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. We use robust standard errors, clustered at the city level to address concerns that the errors are serially correlated at the city level. Our observations are at the quarterly level, and cover the first quarter of 2001 through the fourth quarter of 2016.

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.[22] Our DD fixed effects specification allows us to capture macroeconomic changes, such as the Great Recession, fuel costs, and improvements in vehicle technology, as well as city-specific conditions such as average weather patterns and city topology. The location-specific time trend captures location-specific pre-trends in accidents that existed prior to the arrival of ridehailing. To capture potential time-and-city varying confounders, such as population changes or increases in employment or income, we further control for population level and per capita income.

A number of additional checks must be performed to validate that the assumptions underlying the generalized DD framework hold true. 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 ten years preceding and four 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 (Goodman-Bacon, 2018; De Chaisemartin, and d'Haultfoeuille, 2020). In theory, negative weights on our treated city-

[22] For robustness, we also estimate all our models with the inclusion of a location-specific quadratic trend as well, with qualitatively similar results.

quarter 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 we use is correctly specified because the dose-response relationship between the residualized outcome variable and the residualized treatment variable is linear (Goodman-Bacon 2018, De Chaisemartin and d'Haultfoeuille 2020).

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 and 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, in our models, we also have a large group of never-treated cities. As explained in Goodman-Bacon (2018), 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 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.

5. EMPIRICAL RESULTS

5.1. Main Analysis

We begin our analysis by examining changes in the level of accidents in the treated cities around the introduction of RH. Figure 1 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 difference-in-differences (DD) specification with fixed effects for location and time (quarter-year) and location-specific linear and quadratic trends. In our estimations, we use a number of measures for $accidents_{t,c}$. 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 percent the sample (92,396 out of 190,080 observations). The first column of each set reports estimates without the inclusion of 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. For brevity, we report only the coefficient on the variable of interest— $POST * TREATED_{t,c}$ in the table. Our results remain robust to the use of 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 1, and supporting Hypothesis H1A, we observe a consistently positive and significant coefficient on the $POST * TREATED_{t,c}$ 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 $POST * TREATED_{t,c}$, 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 5^{th} percentile to (Appendix Table B2) or employing the inverse hyperbolic sine transformation (Appendix Table B3).

While we have no clear prediction on how long or whether the effect of RH on fatal accidents will last, we can examine this empirically. In the appendix, we break the post-RH variables into quarters to examine the dynamics of the effect up to two 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.[23]

5.1.1. 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 Appendix A, we formally examine whether fatal accidents predict RH entry order. Appendix Table A1 presents multinomial logits for early, mid, and late entry while 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.

Angrist and Pischke (2008) state that the key assumption of the staggered entry DD model is the parallel trend assumption, which states that in absence of the treatment (RH entry), the dependent variable of both the treatment and the control groups should exhibit the same trend. We thus assess graphically whether the parallel trends assumption holds in Figure 3. Specifically, Figure 3 graphically presents the DD estimates (with each dot representing annual-coefficients) for the ten years preceding and four 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, providing support for our inferences (parallel trends in the pre-period). Post-RH, we see a clear increasing treatment effect.[24]

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

[23] 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.
[24] In Online Appendix Figure A2, we present versions of these graphs using quarterly estimates, with similar conclusions.

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).[25] Using this matched sample, we re-estimate our models and present the estimates in Table 2 Panel B. The estimates are consistent with a 2-4% increase in traffic accidents and fatalities following RH entry. These effects are robust to a different matching approach (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 DD regressions. As Goodman-Bacon (2018) notes, the two-way fixed-effect estimator is a weighted average of all potential 2x2 DD estimates, 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 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 states to never treated states ("Never vs. Timing," which compares states that adopt the policy at some point during the sample period and those that did not). Overall, the decomposition strengthens our causal interpretation of the results.

5.1.2. 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

[25] We did not utilize propensity score matching as there have been some concerns raised about this technique (King & Nielsen 2019).

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 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, in each specification 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 MSA or 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 standard errors 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 standard errors 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 city growth to be reflected in wages. However, we observe no evidence of a statistically significant growth in wages post-RH.

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 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.

5.2. Adoption Intensity

We next provide evidence in support of Hypothesis H1B. In Table 5, we explore the effect of the intensity of service adoption or popularity, proxied for by Google searches for RH keywords. In the main models presented in Table 2, we employ the first launch of a RH service as our treatment date. Take up of these services, however, is likely to intensify over time. We interact our $TREATMENT$ indicator with the intensity of Google searches measure and re-estimate our models. 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 are consistent with an increase in accidents following an increase in our Google Trends intensity measure. For all six models, the coefficient estimate on $POST * INTENSITY$ is positive and statistically significant. Thus, as our proxy for adoption/popularity intensity (Google trends search intensity) increases, so do fatal accidents.

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.[26]

5.3. 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 nonvehicle 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

[26] An ideal additional test would be to look at U.S. cities where RH was introduced and then withdrawn. Unfortunately, these cities are few, and the circumstances do not allow for the types of tests we would want. For example, Uber and Lyft both withdrew from the Austin market at one point in 2016 in a regulatory dispute, but at least five other RH services were still operating and took up the slack. Uber and Lyft then returned to the Austin market within a year, after Texas passed HB100, creating looser statewide rules that superseded Austin's (their return led to immediate massive drops in volume for the competitors that sprung up in their absence). In Las Vegas, the other city we are aware of, RH was introduced, then outlawed after only one month of service.

indicates they involve persons that are not motor vehicle occupants or motorcycle riders.[27] Thus, "pedestrian" in our context refers to both pedestrians as well as bicycle, skateboard and scooter riders, etc.

In Table 6, we present the estimates from models similar to those in Table 2, 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 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 ten years preceding and four 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 A2, we present versions of these graphs using quarterly estimates, with a similar conclusion.

5.4. Urban Versus Rural Areas

Hypothesis H2 posits that the effects of the introduction of RH on accidents will be larger in more urban areas. In dense urban areas, the number of people per mile and per mile of road is higher, traffic congestion is typically elevated to begin with, and any increases in VMT usage may have a bigger impact on congestion and accidents overall as a result. Moreover, in urban areas, and in areas with high usage of cars on the roads or car ownership, 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 also have fewer locations for safe pullover 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.

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. For each characteristic, we divide cities into quartiles and re-estimate our models, interacting $POST * TREATMENT_{c,t}$ with the four quartile indicators for the city characteristic. As before, all models include location and year-quarter fixed effects, a location-specific linear time trend, and control variables.

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

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

cities (fourth quartile), with an estimate for $POST * TREATED * Q4$ 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 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.

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 H2.

5.5. 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 and Aydinonat 2014), we explore whether the empirical evidence is consistent with one or both of the theoretical mechanisms outlined by our conceptual framework.

5.5.1. 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 (in contrast to highway VMT, which are readily available from the Department of Transportation). To examine this channel, first, on the intensive margin, we use annual estimates of arterial vehicle miles traveled, excess gas consumption,[28] 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 H3A. 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

[28] 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 p.m. to 5 a.m.) for each roadway section and hour of the week. We describe these measures in detail in the Online Appendix Section 3.

aggregate locations up to the urban area.[29] Moreover, we can estimate only for the years up to 2014, for these 99 urban areas, leaving us with 1,386 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 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.[30] We present estimates graphically over time in Figure 7. 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 H3B. 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 IV.B. The estimates suggest that new car registrations increase with the intensity of Google searches for Uber/Lyft/ridesharing. This relationship intensifies when RH begins

[29] 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.

[30] 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 U.S. cities (Boston, San Francisco, Washington DC, Chicago, and Los Angeles). https://drive.google.com/file/d/1FIUskVkj9lsAnWJQ6kLhAhNoVLjfFdx3/view

in a treated city. These results suggest that new vehicle purchases increase as RH services become more intensely used.

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 IV.D: 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* quartile of ex-ante carpool usage. Figure 6 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 findings regarding VMT (the intensive margin).

Hypothesis H3C 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. People are more likely to substitute public transit trips with point-to-point RH trips when they require less waiting time, are faster, and more convenient (Babar and 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 and Hamidi, 2014). To the extent that riders substitute into RH and away from non-car modes of transportation, such as including walking, biking, and, more importantly, public transportation (Clewlow and Mishra, 2017), the VMT generated simply by the rides themselves might also exceed the personal driving it displaces, before we even account for between fare VMT. Thus, the safety 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 service.

In Table 10 we test this moderating hypothesis. We break out our results across quartile of ex ante public transportation usage and ex ante carpool usage as reported by the American Community Survey. Once again, for each characteristic, we divide cities into quartiles and re-estimate our models, interacting $POST * TREATMENT_{c,t}$ with the four quartile indicators for the city characteristic. As before, all models include location and year-quarter fixed effects, a location-specific linear time trend, and control variables. 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 H3C.

5.5.2. *Driver Quality*

Examining driver quality is challenging given the nature of the available data, and thus, providing evidence related to Hypothesis H4 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 (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 vehicles to be four door, and typically less than ten 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 x Year*, *City x Day of Week* and *City x 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, but where the dependent variable is the log of one plus the number 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, as each RH vehicle involves 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) to (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 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 Appendix Table C2, we show that the coefficient on $POST\ X\ TREATED_{c,t}$ 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), and are consistent with hypothesis.

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, 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 (see e.g. https://fortune.com/2017/04/13/uber-drunk-drivers/), 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 U.S. cities (see e.g., Brazil and Kirk, 2016; Greenwood and 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 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 Appendix Section C.3 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 to 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. 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.[31] We present estimates graphically over time in Figure

[31] 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 8/2014 to 8/2015 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

8. 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.

6. DISCUSSION

6.1. Back of the Envelope Cost Calculation

Up until this point, our study has documented a societal cost associated with the introduction of RH. To say something about welfare, we must also consider its benefits. Benefits come from, for example, the consumer surplus provided by convenience. Cohen et al. (2016) use a combination of Uber's "surge" pricing algorithm and individual-level data to estimate demand elasticities at several points along the demand curve. They then use these elasticity estimates to estimate consumer surplus. They estimate that, in 2015, the UberX service generated about $2.9 billion in consumer surplus in the four U.S. cities they examine. Moreover, their calculations suggest that 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 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 U.S.[32] The 3% annual increase associated with the introduction of RH in fatalities represents an additional 987 lives lost each year.[33] The U.S. 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.[34] 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.'s (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. Our estimates, moreover, do not include the costs imposed by nonfatal accidents, for which data are not readily available. We can assume that the pattern for fatal accidents is repeated for nonfatal accidents, leading to costs in material and healthcare that may dwarf

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.

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

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

[34] 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

these VSL estimates. 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.

6.2. Managerial and Policy Implications and Suggestions

Our study provides insights that may be relevant to multiple stakeholders in the RH ecosystem. First, RH companies themselves 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 certain 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.

6.2.1. 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 operation efficiency against the safety implications of their operational choices. A number of possible operational changes that help to address both the quantity and congestion channel and driver quality channel come to mind.

First, from a quantity perspective, voluntarily limiting operations under specific conditions may help reduce the congestion effects of RH services. Operations managers could choose to reduce surge incentives in areas or at times of high traffic congestion using current technological advances. Similarly, RH companies could choose to limit service or eliminate it entirely in certain areas during periods of high road congestion. Finally, operations managers could choose to eliminate low pricing for Pool rides that ultimately only have a single rider, reducing substitution from other modes of transportation.

Second, RH operations managers could also take a number of steps with an eye towards improving driver quality and reducing 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.[35] Driving record background checks, limitation of driver eligibility to those with no record of driving violations in the last five 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

[35] Uber enacted one such limit a few years ago: "Starting on Wednesday, 2 May 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 hours of driving. You'll then have to be offline for an entire 6 hour period before the timer resets." See https://www.uber.com/en-ZA/blog/driving-hours-limit/.

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.

At the platform level, operations managers could consider modifications to navigation algorithms that would 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.

6.2.2. 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 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. Enhanced 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 limits on the number of hours a driver may work on gig platforms in a given 24-hour 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-hour 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.

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.

Importantly, advances in big data processing and machine learning, combined with requiring unrestricted access to the more granular data held by RH companies, insurers, and federal and local governments, could allow researchers to better evaluate and design optimal implementations of such policies.

6.3. 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, we our analysis is conducted on the initial years of ridehailing 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.

6.4. 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 driving to work or to a social event, it can impose much greater costs on first responders and those being assisted by them (Beland and Brent, 2018). For illustration, suppose there are 100 heart attack victims transported to the emergency room each day by car or ambulance. These individuals face much higher costs from congestion and road delays. As congestion increases, a higher proportion of these 100 cases 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, 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. In sum, our findings suggest more research will be needed to better quantify both the societal cost and benefit of RH. More generally, our work points to the need for better consideration of societal costs and externalities associated with the introduction of new technologies.

Finally, 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.

6.5. 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 U.S. 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 percent higher than it was 10 years ago, reversing a longstanding trend of decline in pedestrian deaths from motor vehicle crashes.

In this paper, we provide evidence consistent with RH imposing an increase in fatal accidents and fatalities on the motor vehicle occupants and pedestrians in the cities it serves. 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. And 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.

REFERENCES

Afeche, Philipp, Zhe Liu, and Costis Maglaras. "Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance." Rotman School of Management Working Paper 3120544 (2018): 18-19.

Anderson, Michael L. "Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion." American Economic Review 104, no. 9 (2014): 2763-96.

Angrist, Joshua D., and Jörn-Steffen Pischke. "Mostly harmless econometrics." In Mostly Harmless Econometrics. Princeton university press, 2008.

Ashforth, Blake E., and Peter H. Reingen. "Functions of dysfunction: Managing the dynamics of an organizational duality in a natural food cooperative." Administrative Science Quarterly 59, no. 3 (2014): 474-516.

Autor, David H. "Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing." Journal of labor economics 21, no. 1 (2003): 1-42.

Babar, Y., G. Burtch. 2020. Examining the heterogeneous impact of ride-hailing services on public transit use. Inform. Syst. Res. 31(3): 820–834.

Bai, Jiaru, Kut C. So, Christopher S. Tang, Xiqun Chen, and Hai Wang. "Coordinating supply and demand on an on-demand service platform with impatient customers." Manufacturing & Service Operations Management 21, no. 3 (2019): 556-570.

Banerjee Siddartha, Ramesh Johari and Carlos Riquelme. "Pricing in ride-sharing platforms: A queueing-theoretic approach." In Proceedings of the 16th ACM Conference on Economic Computing (2015), 639–639.

Banerjee, Siddhartha, Daniel Freund, and Thodoris Lykouris. "Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework." In Proceedings of the 2017 ACM Conference on Economics and Computation, pp. 517-517. 2017.

Banerjee, Siddhartha, Yash Kanoria, and Pengyu Qian. "State dependent control of closed queueing networks." ACM SIGMETRICS Performance Evaluation Review 46, no. 1 (2018): 2-4.

Beland, Louis-Philippe, and Daniel Brent. "Traffic Congestion and the Performance of First Responders: Evidence from California Fire." (2018).

Benjaafar, Saif, and Ming Hu. "Operations management in the age of the sharing economy: What is old and what is new?." Manufacturing & Service Operations Management 22, no. 1 (2020): 93-101.

Benjaafar, Saif, Jian-Ya Ding, Guangwen Kong, and Terry Taylor. "Labor welfare in on-demand service platforms." Manufacturing & Service Operations Management 24, no. 1 (2022): 110-124.

Bertrand, Marianne, and Sendhil Mullainathan. "Are CEOs rewarded for luck? The ones without principals are." The Quarterly Journal of Economics 116, no. 3 (2001): 901-932.

Besbes, Omar, Francisco Castro, and Ilan Lobel. "Surge pricing and its spatial supply response." Management Science 67, no. 3 (2021): 1350-1367.

Bhargava, Saurabh, and Vikram S. Pathania. "Driving under the (cellular) influence." American Economic Journal: Economic Policy 5, no. 3 (2013): 92-125.

Bimpikis, Kostas, Ozan Candogan, and Daniela Saban. "Spatial pricing in ride-sharing networks." Operations Research 67, no. 3 (2019): 744-769.

Birkinshaw, Julian, and Kamini Gupta. "Clarifying the distinctive contribution of ambidexterity to the field of organization studies." Academy of Management Perspectives 27, no. 4 (2013): 287-298.

Blackwell, Matthew, Stefano Iacus, Gary King, and Giuseppe Porro. "cem: Coarsened exact matching in Stata." The Stata Journal 9, no. 4 (2009): 524-546.

Borusyak, Kirill, and Xavier Jaravel. "Revisiting event study designs." Available at SSRN 2826228 (2017).

Bouchery, Yann, Charles J. Corbett, Jan C. Fransoo, and Tarkan Tan, eds. Sustainable supply chains: A research-based textbook on operations and strategy. Vol. 4. Springer, 2016.

Braverman, Anton, Jim G. Dai, Xin Liu, and Lei Ying. "Empty-car routing in ridesharing systems." Operations Research 67, no. 5 (2019): 1437-1452.

Brazil, Noli, and David S. Kirk. "Uber And Metropolitan Traffic Fatalities in the United States," American Journal of Epidemiology, 184 (2016), 192–198.

Buchak, Greg. "Financing the Gig Economy." University of Chicago Booth Working Paper, 2018.

Burgdorf, Jacob, Conor Lennon, and Keith Teltser. "Do Ridesharing Services Increase Alcohol Consumption?." Working Paper, 2019.

Cachon, Gerard P., Kaitlin M. Daniels, and Ruben Lobel. "The role of surge pricing on a service platform with self-scheduling capacity." Manufacturing & Service Operations Management 19, no. 3 (2017): 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.

Castillo, Juan Camilo, Dan Knoepfle, and Glen Weyl. "Surge pricing solves the wild goose chase." In Proceedings of the 2017 ACM Conference on Economics and Computation, pp. 241-242. 2017.

Chen, M, Keith, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen, “The Value of Flexible Work: Evidence from Uber Drivers,” NBER Working Paper No. w23296, 2017.

Clewlow, Regina R., and Gouri Shankar Mishra, “Disruptive Transportation: The Adoption, Utilization and Impacts of Ride-Hailing in the United States,” Institute of Transportation Studies, University of California, Davis, Working Paper, 2017.

Circella, Giovanni, Farzad Alemi, Kate Tiedeman, Susan Handy, and Patricia Mokhtarian, “The Adoption of Shared Mobility in California and Its Relationship with Other Components of Travel Behavior,” Institute of Transportation Studies, University of California, Davis Working Paper, 2018.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, “Using Big Data to Estimate Consumer Surplus: The Case of Uber,” NBER Working Paper No. w22627, 2016.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John A. List, and Paul Oyer, “The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers,” Stanford Graduate School of Business Working Paper, 2018.

Cramer, Judd, and Alan B. Krueger, “Disruptive Change in the Taxi Business: The Case of Uber,” American Economic Review, 106 (2016), 177–182.

Currie, Janet, and Reed Walker. "Traffic congestion and infant health: Evidence from E-ZPass." American Economic Journal: Applied Economics 3, no. 1 (2011): 65-90.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. "Two-way fixed effects estimators with heterogeneous treatment effects." American Economic Review 110, no. 9 (2020): 2964-96.

Di Tella, Rafael, and Ernesto Schargrodsky. "The role of wages and auditing during a crackdown on corruption in the city of Buenos Aires." The Journal of Law and Economics 46, no. 1 (2003): 269-292.

Dills, Angela K., and Sean E. Mulholland, “Ride‐Sharing, Fatal Crashes, and Crime,” Southern Economic Journal, 84 (2018), 965–991.

Dingus, Thomas A., Feng Guo, Suzie Lee, Jonathan F. Antin, Miguel Perez, Mindy Buchanan-King, and Jonathan Hankey. "Driver crash risk factors and prevalence evaluation using naturalistic driving data." Proceedings of the National Academy of Sciences 113, no. 10 (2016): 2636-2641.

Duflo, Esther, Rema Hanna, and Stephen P. Ryan. "Incentives work: Getting teachers to come to school." American Economic Review 102, no. 4 (2012): 1241-78.

Duranton, Gilles, and Matthew A. Turner. "The fundamental law of road congestion: Evidence from US cities." American Economic Review 101, no. 6 (2011): 2616-52.

Edlin, Aaron S., “Per-Mile Premiums for Auto Insurance.” In Economics for an Imperfect World: Essays in Honor of Joseph E. Stiglitz, edited by Richard Arnott, Bruce Greenwald, Ravi Kanbur, and Barry Nalebuff. Cambridge, MA: MIT Press, 2003.

Edlin, Aaron S., and Pinar Karaca-Mandic, “The Accident Externality from Driving.” Journal of Political Economy, 114 (2006), 931-955.

Ewing, R., S. Hamidi 2014. Measuring Urban Sprawl and Validating Sprawl Measures. National Institutes of Health and Smart Growth America, Washington, DC.

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

Feng, Guiyun, Guangwen Kong, and Zizhuo Wang. "We are on the way: Analysis of on-demand ride-hailing systems." Manufacturing & Service Operations Management (2020).

Ferdows, Kasra, and Arnoud De Meyer. "Lasting improvements in manufacturing performance: in search of a new theory." Journal of Operations management 9, no. 2 (1990): 168-184.

Gee, Gilbert C., and David T. Takeuchi. "Traffic stress, vehicular burden and well-being: a multilevel analysis." Social science & medicine 59, no. 2 (2004): 405-414.

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

Gottholmseder, Georg, Klaus Nowotny, Gerald J. Pruckner, and Engelbert Theurl. "Stress perception and commuting." Health economics 18, no. 5 (2009): 559-576.

Greenwood, Brad N., and Sunil Wattal, “Show Me the Way to Go Home, An Empirical Investigation of Ride-Sharing and Alcohol Related Motor Vehicle Fatalities,” MIS Quarterly, 41 (2017), 163–187.

Grübler, Arnulf, “Diffusion: Long-Term Patterns and Discontinuities,” Technological Forecasting and Social Change 39 (1991), 159–180.

Haggag, Kareem, Brian McManus, and Giovanni Paci, “Learning by Driving: Productivity Improvements by New York City Taxi Drivers,” American Economic Journal: Applied Economics, 9 (2017), 70–95.

Hall, Jonathan D., Craig Palsson, and Joseph Price. "Is Uber a Substitute or Complement for Public Transit?," Journal of Urban Economics, 108 (2018), 36-50.

He, Long, Ho-Yin Mak, Ying Rong, and Zuo-Jun Max Shen. "Service region design for urban electric vehicle sharing systems." Manufacturing & Service Operations Management 19, no. 2 (2017): 309-327.

Henao, Alejandro, “Impacts of Ridesourcing – Lyft and Uber – on Transportation including VMT, Mode Replacement, Parking, and Travel Behavior,” Working Paper, 2017.

Hennessy, Dwight A., and David L. Wiesenthal. "Traffic congestion, driver stress, and driver aggression." Aggressive Behavior: Official Journal of the International Society for Research on Aggression 25, no. 6 (1999): 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, Paul W. "Statistics and causal inference." Journal of the American statistical Association 81, no. 396 (1986): 945-960.

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

Hu, Ming, and Yun Zhou. "Dynamic type matching." Manufacturing & Service Operations Management (2021).

Imbens, Guido W., and Jeffrey M. Wooldridge. "Recent developments in the econometrics of program evaluation." Journal of economic literature 47, no. 1 (2009): 5-86.

King, Gary, and Richard Nielsen. "Why propensity scores should not be used for matching." Political Analysis 27, no. 4 (2019): 435-454.

Kleindorfer, Paul R., Kalyan Singhal, and Luk N. Van Wassenhove. "Sustainable operations management." Production and operations management 14, no. 4 (2005): 482-492.

Knittel, Christopher R., Douglas L. Miller, and Nicholas J. Sanders. "Caution, drivers! Children present: Traffic, pollution, and infant health." Review of Economics and Statistics 98, no. 2 (2016): 350-366.

Kong, H., X. Zhang, J. Zhao. 2020. How does ridesourcing substitute for public transit? A geospatial perspective in Chengdu, China. J. Transp. Geogr. 86:102769.

Liu, Sheng, Long He, and Zuo-Jun Max Shen. "On-time last-mile delivery: Order assignment with travel-time predictors." Management Science 67, no. 7 (2021): 4095-4119.

Mak, Ho-Yin. "Enabling smarter cities with operations management." Manufacturing & Service Operations Management 24, no. 1 (2022): 24-39.

Mangrum, Daniel, and Alejandro Molnar. “The marginal congestion of a taxi in New York City.” Vanderbilt University Working Paper, 2017.

Martin-Buck, Frank, “Driving Safety, An Empirical Analysis of Ridesharing’s Impact on Drunk Driving and Alcohol-Related Crime,” University of Texas at Austin Working Paper, 2017.

Metropolitan Area Planning Council, “Fare Choices: A Survey of Ride-Hailing Passengers in Metro Boston,” February 2018.

Miller, Jason, John P. Saldanha, Manus Rungtusanatham, A. Michael Knemeyer, and Thomas J. Goldsby. "How does electronic monitoring affect hours-of-service compliance?." Transportation journal 57, no. 4 (2018): 329-364.

Mokyr, Joel, “Punctuated Equilibria and Technological Progress,” The American Economic Review, 80 (1990), 350–354.

New York City Department of Transportation, “NYC Mobility Report,” June 2018.

Nagin, Daniel S., James B. Rebitzer, Seth Sanders, and Lowell J. Taylor. "Monitoring, motivation, and management: The determinants of opportunistic behavior in a field experiment." American Economic Review 92, no. 4 (2002): 850-873.

Naumov, Sergey, David R. Keith, and Charles H. Fine. "Unintended consequences of automated vehicles and pooling for urban transportation systems." Production and Operations Management 29, no. 5 (2020): 1354-1371.

Ostrovsky, Michael, and Michael Schwarz. "Carpooling and the economics of self-driving cars." In Proceedings of the 2019 ACM Conference on Economics and Computation, pp. 581-582. 2019.

Owens, Justin M., Thomas A. Dingus, Feng Guo, Youjia Fang, Miguel Perez, and Julie McClafferty. "Crash risk of cell phone use while driving: A case-crossover analysis of naturalistic driving data." (2018).

Özkan, Erhun, and Amy R. Ward. "Dynamic matching for real-time ride sharing." Stochastic Systems 10, no. 1 (2020): 29-70.

Pagell, Mark, Robert Klassen, David Johnston, Anton Shevchenko, and Sharvani Sharma. "Are safety and operational effectiveness contradictory requirements: The roles of routines and relational coordination." Journal of Operations Management 36 (2015): 1-14.

Parente, Stephen L., and Edward C, Prescott, "Barriers to Technology Adoption and Development," Journal of Political Economy, 102 (1994), 298–321.

Peltzman, Sam, "The Effects of Automobile Safety Regulation," Journal of Political Economy, 83 (1975), 677–725.

Pew Research Center, May, 2016, "Shared, Collaborative and On Demand: The New Digital Economy."

Posen, Hannah A. "Ridesharing in the sharing economy: Should regulators impose Uber regulations on Uber." Iowa L. Rev. 101 (2015): 405.

Qi, W., L. Li, S. Liu, Z. J. M. Shen. 2018. Shared mobility for last-mile delivery: Design, operational prescriptions, and environmental impact. Manuf. Serv. Oper. Manag. 20(4): 737–751.

Qi, Wei, and Zuo‐Jun Max Shen. "A smart‐city scope of operations management." Production and Operations Management 28, no. 2 (2019): 393-406.

Qi, Wei, and Zuo‐Jun Max Shen. "A smart‐city scope of operations management." Production and Operations Management 28, no. 2 (2019): 393-406.

Rosenberg, Nathan, Luther E. Birdzell, and Glenn W. Mitchell, How the West Grew Rich, Mumbai: Popular Prakashan, 1986.

Rosenzweig, Eve D., and George S. Easton. "Tradeoffs in manufacturing? A meta‐analysis and critique of the literature." Production and operations management 19, no. 2 (2010): 127-141.

Schaller, Bruce, "The New Automobility: Lyft, Uber and the Future of American Cities," Schaller Consulting Report, 2018.

Schrank, David, Tim Lomax, and Bill Eisele. "2011 urban mobility report." (2011).

Scott, Alex, Andrew Balthrop, and Jason W. Miller. "Unintended responses to IT‐enabled monitoring: The case of the electronic logging device mandate." Journal of Operations Management 67, no. 2 (2021): 152-181.

Skinner, Wickham. "Manufacturing strategy on the "S" curve." Production and operations management 5, no. 1 (1996): 3-14.

Smith, Wendy K., and Marianne W. Lewis. "Toward a theory of paradox: A dynamic equilibrium model of organizing." Academy of management Review 36, no. 2 (2011): 381-403.

Song, Y., Y. Huang. 2020. Investigating complementary and competitive relationships between bikeshare service and public transit: a spatial-temporal framework. Transp. Res. Rec. 2674 (1): 260–271.

Vickrey, William, "Automobile Accidents, Tort Law, Externalities, and Insurance: An Economist's Critique." Law and Contemporary Problems 33 (1968), 464–87.

Ylikoski, Petri, and N. Emrah Aydinonat. "Understanding with theoretical models." Journal of Economic Methodology 21, no. 1 (2014): 19-36.

Yu, Jiayi Joey, Christopher S. Tang, Zuo-Jun Max Shen, and Xiqun Michael Chen. "A balancing act of regulating on-demand ride services." Management Science 66, no. 7 (2020): 2975-2992.

TABLES

Table 1

Summary Statistics

	Mean	Median	Std. Dev.
Population (thousands)	54.50	23.51	199.99
Income per capita (thousands $)	39.72	37.47	12.17
Population density	2,999.28	2,169.70	3,159.57
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	2,340
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

Notes: The sample contains 190,080 quarterly observations on 2,970 census incorporated places from 2001 to 2016. Population density measures 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.

Table 2

Effect of Ridehailing on Traffic Safety

Panel A: Overall Effect

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	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)
$Post\ X\ Treated_{t,c}$	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)	(2)	(3)	(4)	(5)	(6)
	Total Accidents	Total Fatalities	Total Accidents	Total Fatalities	Accident Rate	Fatality Rate
$Post\ X\ Treated_{t,c}$	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

Notes: This table presents results from generalized difference-in-difference regressions. The dependent variables are listed at the top of each column. $Post\ * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Panel A presents the overall effect of ridehailing on two traffic safety measures. Total accidents is the number of fatal accidents according to the definition used by NHTSA. Total fatalities is 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). Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors—clustered at city level—are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

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

Notes: This table presents the dynamic effects of ridehailing on traffic safety. The dependent variables are listed at the top of the columns. Total accidents is the number of fatal accidents, according to the definition used by NHTSA. Total fatalities is 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. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4

Robustness Tests and Inferences

Robustness Type	Robustness Detail	Inferences	Location of Detailed Analysis in Appendix
Method	Populationg 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 standard errors 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 Weekend Night	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
	Shorter sample (2007-2016)	Unchanged	Table B10

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

Table 5

Variation of Ridehailing Service and Adoption Intensity

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Post X Treated}_{t,c} \text{ X Log Rideshare} - \text{Related Google Search Volume}_{ct}$	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

Notes: This table shows how the effect of ridehailing on traffic safety varies with the intensity of service. In all panels, the dependent variables are listed at the top of the columns. Total accidents is the number of fatal accidents, according to the definition used by NHTSA. Total fatalities is 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 volume for the terms "Uber," "Lyft," and "rideshare." Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

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)
$Post\ X\ Treated_{t,c}$	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

Notes: This table presents results from generalized difference-in-difference regressions. The dependent variables are listed 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. $Post\ X\ Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. City-specific linear trends are included in all columns. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7

Effects of Ridehailing on Traffic Fatalities: Urban and Rural Cities

	Log (1+ Total Accidents)		
	(1)	(2)	(3)
	Population	Population Density	Ex Ante Vehicle Ownership
$Post\ X\ Treated_{t,c}\ X\ Q4$	0.0746***	0.0354***	0.0769***
	(0.0113)	(0.0109)	(0.0111)
$Post\ X\ Treated_{t,c}\ X\ Q3$	0.0049	0.0323**	0.0024
	(0.0136)	(0.0136)	(0.0145)
$Post\ X\ Treated_{t,c}\ X\ Q2$	0.0041	0.0504***	-0.0054
	(0.0154)	(0.0163)	(0.0148)
$Post\ X\ Treated_{t,c}\ X\ Q1$	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

Notes: This table presents heterogeneous effects of ridehailing on traffic safety by population, population density, and ex ante vehicle ownership. 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 is the number of fatal accidents according to the definition used by NHTSA. $Post\ X\ Treated_{t,c}$ is a dummy variable that equals one if 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. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 8

Effect of Ridehailing on Road Utilization and Congestion

	(1)	(2)	(3)
	Log Arterial Street VMT	Log Excess Fuel Consumption	Log Hours of Delay in Traffic
$Post\ X\ Treated_{t,u}$	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	1,386	1,386	1,386
R2	0.998	0.999	0.999

Notes: The sample contains 1,386 annual observations on 99 urban areas from 2001 to 2014. 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 vehicle miles traveled 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 a dummy variable that equals one if urban area u adopted at least one ridehailing service at year t. Urban area-specific linear trends are included in all regressions. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the urban area level, are reported in parentheses. 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. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 9

The Effect of Ridehailing on New Car Registrations

Panel A: Main Effect

	Log (1+New Car Registrations)		
	(1)	(2)	(3)
$Post\ X\ Treated_{t,c}$	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		
$Post*Treated_{t,c}*$ 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
$Post\ X\ Treated_{t,c}\ X\ Q4$	0.0842***	0.0926***	0.0620***	0.1572***	0.0809***
	(0.0096)	(0.0118)	(0.0110)	(0.0144)	(0.0093)
$Post\ X\ Treated_{t,c}\ X\ Q3$	0.0347**	0.0554***	0.0673***	0.0674***	0.0427***
	(0.0142)	(0.0110)	(0.0123)	(0.0108)	(0.0132)
$Post\ X\ Treated_{t,c}\ X\ Q2$	0.0149	0.0168	0.0299**	0.0195*	-0.0001
	(0.0168)	(0.0118)	(0.0142)	(0.0109)	(0.0195)
$Post\ X\ Treated_{t,c}\ X\ Q1$	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
R2	0.97	0.97	0.97	0.97	0.97

Notes: This table presents the effect of ridehailing on new car registrations. In all panels, the dependent variables are the natural logarithm of one plus new car registrations. $Post\ X\ Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Panel A presents results from generalized difference-in-difference regressions without city time trend, with linear trend, and with quadratic trend. Panel B shows how the effect varies with the intensity of ridehailing service. Log RideHail Google Search Volume is the natural logarithm of Google search volume 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. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 10

Effects of Ridehailing on Traffic Fatalities: The Role of Public Transport

	Log (1+Total Accidents)	
	(1)	(2)
	Ex Ante Public Transportation Usage	Ex Ante Car Pool Usage
$Post\ X\ Treated_{t,c}\ X\ Q4$	0.0381***	0.0448***
	(0.0112)	(0.0151)
$Post\ X\ Treated_{t,c}\ X\ Q3$	0.0537***	0.0633***
	(0.0123)	(0.0131)
$Post\ X\ Treated_{t,c}\ X\ Q2$	0.0175	0.0174
	(0.0155)	(0.0129)
$Post\ X\ Treated_{t,c}\ X\ Q1$	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

Notes: This table presents heterogeneous effects of ridehailing on traffic safety by public transport usage. 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 is the number of fatal accidents according to the definition used by NHTSA. $Post\ X\ Treated_{t,c}$ 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. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 11

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
$Post\ X\ Treated_{t,c}$	0.0293***	0.0246**	0.0021
	(0.0099)	(0.0116)	(0.0057)
City X Year FE	Yes	Yes	Yes
City X Day of Week FE	Yes	Yes	Yes
City X 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)
$Post\ X\ Treated_{t,c}$	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

Notes: This table presents results from generalized difference-in-difference regressions where the dependent variables for each regression are listed at the top of the columns. 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. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 12

The Effect of Ridehailing on Alcohol-Involved Accidents and Fatalities

	Log (1+Drunk Accidents)			Log (1+Drunk Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post * Treated_{t,c}$	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

Notes: This table presents coefficient estimates from generalized difference-in-difference regressions. 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. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

FIGURES

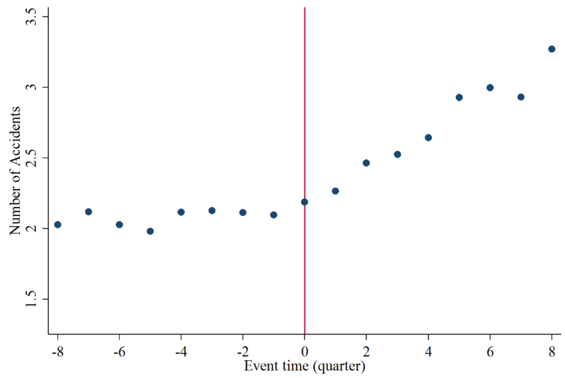

Figure 1

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.

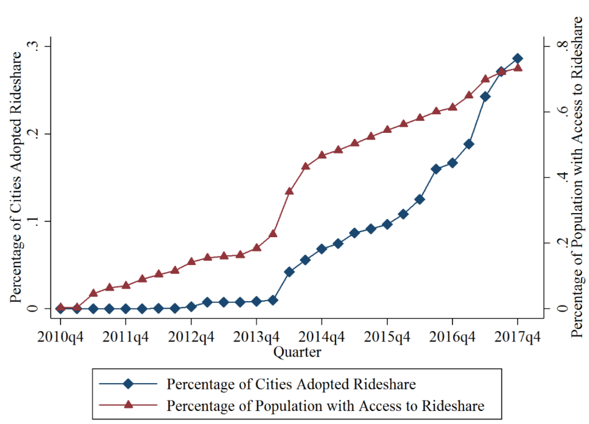

Figure 2

Ridehailing Diffusion

This figure shows the diffusion of ridehailing across the U.S. by cities and population. The sample consists of all census incorporated places in the United States. The navy (red) line graphs the percentage of cities (population) that adopted ridehailing in each quarter between the fourth quarter of 2010 and the fourth quarter of 2017.

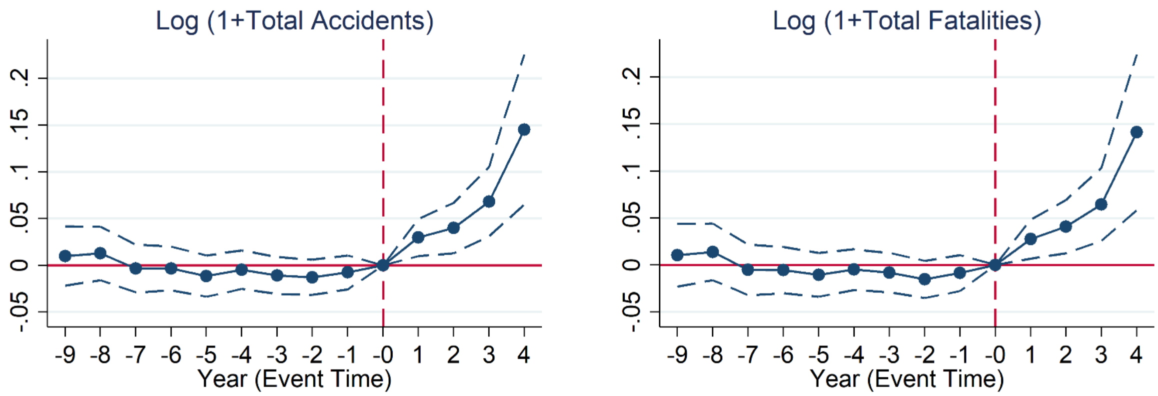

Figure 3

Difference-in-Differences Estimators: Accidents and Fatalities

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To map out the pattern in the counterfactual treatment effects, we regress the various outcome measures on lag and lead indicators (bunched by four quarters) for the entry of ridehailing. The vertical red line indicates the year of entry. We provide a description of the variables in section 3.

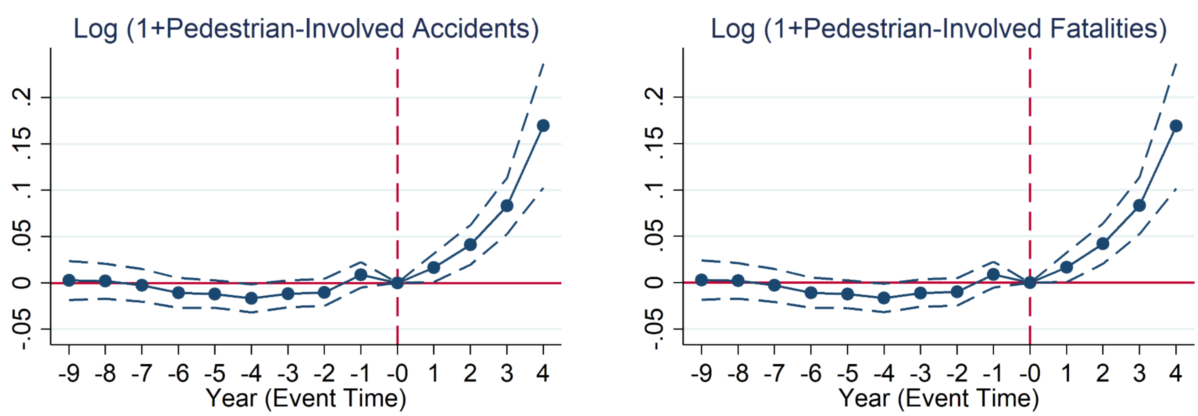

Figure 4

Difference-in-Differences Estimators: Pedestrian-Involved Accidents and Fatalities

This figure graphs the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To graph the dynamics in the counterfactual treatment effects, we regress our various outcome measures on lag and lead indicators (bunched by four quarters) for the entry of ridehailing. The vertical red line indicates the year of entry. A description of the variables are provided in section 3.

Panel A: City Characteristics

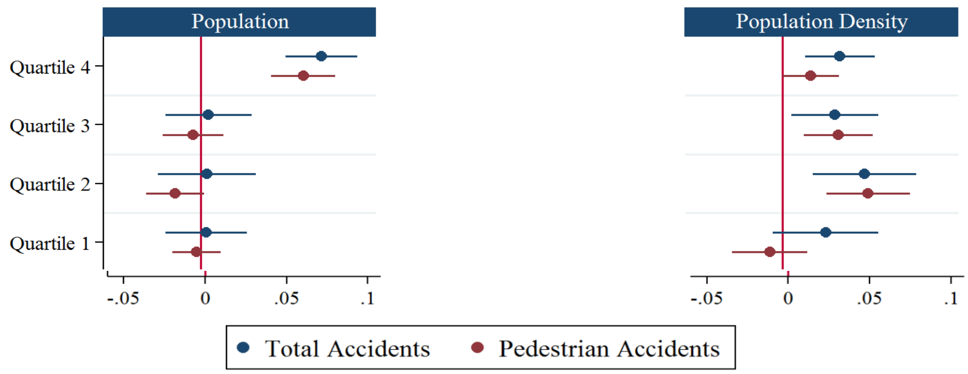

Panel B: Ex-Ante Behaviors

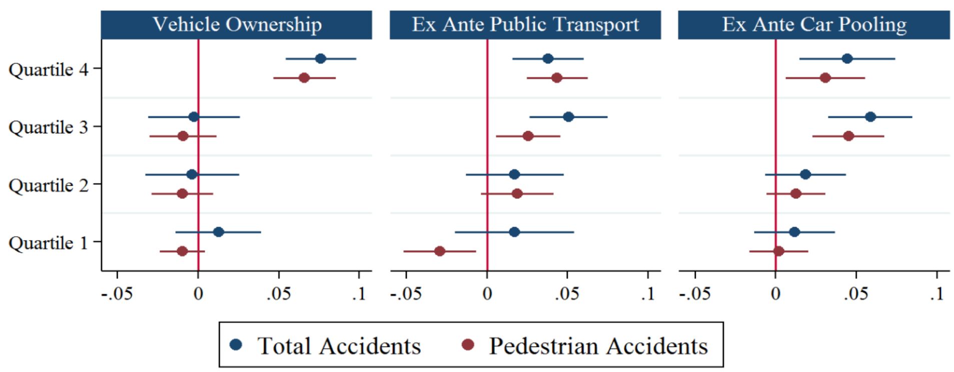

Figure 5

Heterogeneity by City Characteristics

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level, 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 variables for the regressions are listed at the bottom of each figure. Section 3 provide a description of the variables.

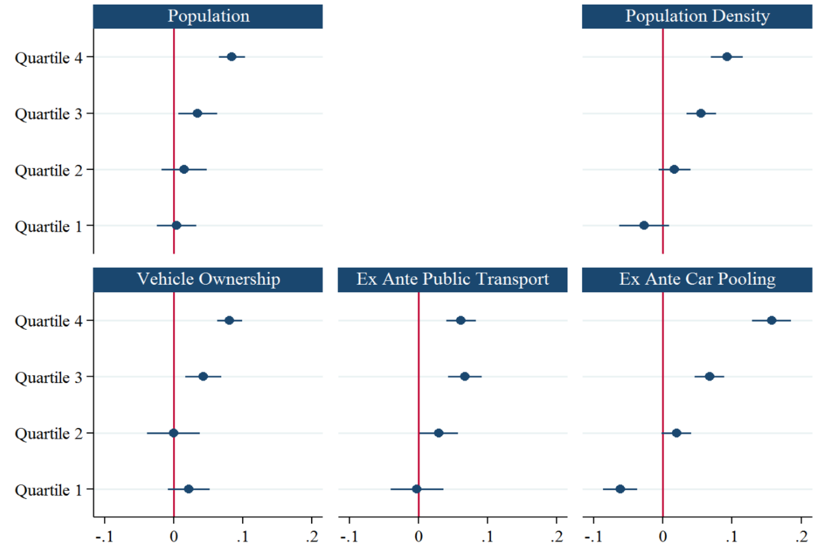

Figure 6

Heterogeneous Effect on Vehicle Registration by City Characteristics

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level, 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.

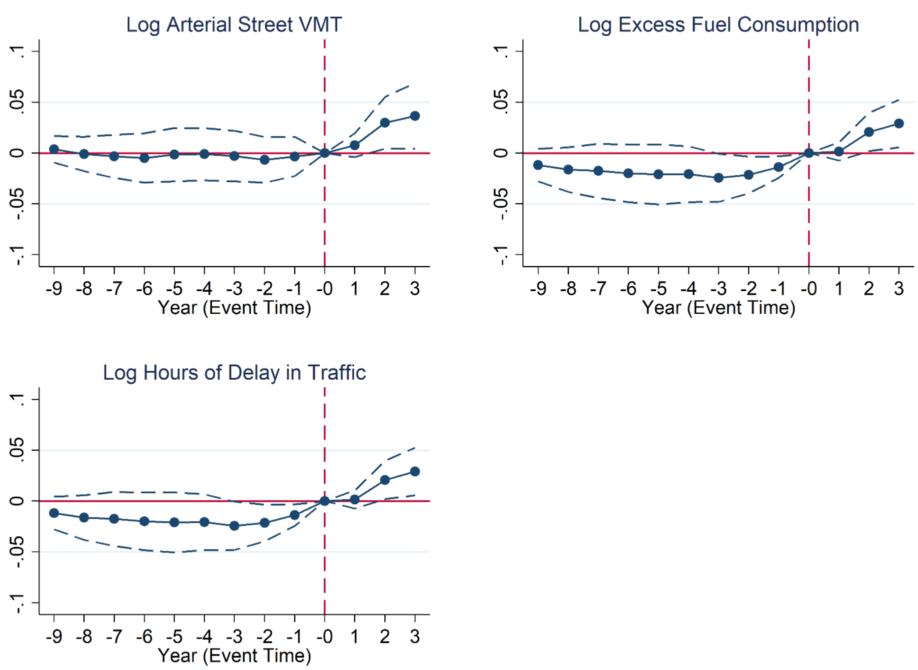

Figure 7

Difference-in-Differences Estimators: Congestion

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To map out the pattern in the counterfactual treatment effects, we regress 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.

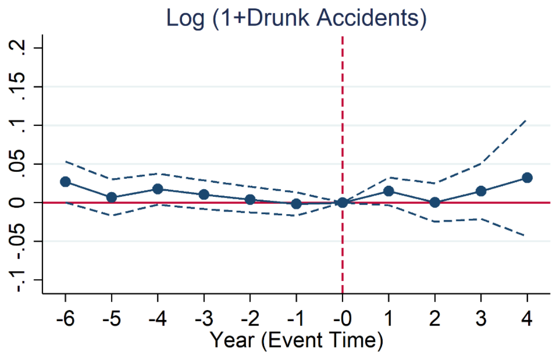

Figure 8

Difference-in-Differences Estimators: Drunk Accidents and Fatalities

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To map out the pattern in the counterfactual treatment effects, we regress our drunk accident measures on lag and lead indicators (bunched by four quarters) for the entry of ridehailing. The vertical red line indicates the year of entry. We provide a description of the variables in section 3.

ONLINE APPENDIX

BARRIOS, HOCHBERG, AND YI (2022)

THE COST OF CONVENIENCE: RIDEHAILING AND TRAFFIC FATALITIES

ONLINE APPENDIX

THE COST OF CONVENIENCE: RIDEHAILING AND TRAFFIC FATALITIES

This Online Appendix provides supplementary analysis to the main manuscript.

A. DD ASSUMPTIONS AND CAUSAL IDENTIFICATION

In the table below (A1), we run multinomial logit models to predict city entry of ridehailing services. 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). While both population and per capita income load positively and significantly in predicting earlier entry, the change in accident rates over the 3, 5 or 10 years prior to entry does not load significantly.

	(1)	(2)	(3)
Early Entry			
Log (Population)	0.8410***	0.8396***	0.8363***
	(0.1329)	(0.1328)	(0.1320)
Log (Income)	2.8368***	2.8486***	2.8177***
	(0.6290)	(0.6291)	(0.6317)
3-Year Accidents Trend	0.3146		
	(0.7033)		
5-Year Accidents Trend		0.5028	
		(1.1834)	
10-Year Accidents Trend			0.9669
			(2.0542)
Late Entry			
Log (Population)	-0.3277***	-0.3284***	-0.3279***
	(0.0623)	(0.0624)	(0.0624)
Log (Income)	-0.8070***	-0.8114***	-0.8366***
	(0.2642)	(0.2646)	(0.2656)
3-Year Accidents Trend	0.0305		
	(0.1603)		
5-Year Accidents Trend		0.3406	
		(0.2595)	
10-Year Accidents Trend			0.6513
			(0.4620)
Observations	1199	1199	1199

Table A1. Multinomial Logit Estimation of Ridehailing Entry

Notes: This table presents multinomial logit estimates of ridehailing entry decisions. The outcome variable is defined as 0 if the city adopted ridesharing in 2010 through 2012, 1 if the city adopted ridesharing in 2013 and 2014 (the base outcome), and 2 if the city adopted ridesharing in 2015 or 2016. 3,5, or 10-Year Accidents Trends are the 3,5, or 10-year average quarterly change in accidents before ridesharing entry. Standard errors are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

In following table (A2), we estimate Cox proportional hazard models for ridehailing entry into cities. The reported coefficient estimates are hazard ratios. We collapse observations at the city-year level to calculate annual percentage changes in fatal accidents and fatalities, population, and income. As can be seen from the table, while population and income strongly predict entry timing, there is no statistically significant loading on trends in fatal accidents or fatalities.

	(1)	(2)	(3)	(4)
	All Cities	Rideshare Cities	All Cities	Rideshare Cities
Annual % Change in Accidents	0.9917	1.0110		
	(0.0298)	(0.0294)		
Annual % Change in Fatalities			0.9776	0.9988
			(0.0295)	(0.0296)
Annual % Change in Pop	1.1119***	0.9408	1.1119***	0.9411
	(0.0213)	(0.0453)	(0.0213)	(0.0453)
Annual % Change in Income	1.2343***	1.2026***	1.2348***	1.2036***
	(0.0531)	(0.0657)	(0.0532)	(0.0657)
Log Pop	1.8447***	1.4364***	1.8455***	1.4373***
	(0.0510)	(0.0378)	(0.0510)	(0.0378)
Log Income Per Capita	1.0938***	1.0907***	1.0931***	1.0898***
	(0.0372)	(0.0360)	(0.0372)	(0.0360)
Observations	35,395	20,428	35,395	20,428

Table A2. Modeling Ridehailing Adoption: Hazard Models

Notes: This table presents results from proportional cox hazard model estimations. The reported coefficient estimates are hazard ratios. We collapse observations at the city-year level to calculate annual percentage changes in total accidents, total fatalities, population, and income percapita. All variables are standardized to have a mean of 0 and standard deviation of one to facilitate comparison between estimated hazard ratios. In Columns (1) and (3), we include all cities in our sample. In Columns (2) and (4), we limit the analysis to cities that adopted rideshare during our sample period. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

We have also addressed several concerns about the staggered entry difference-in-difference specification brought up in Borusyak and Jaravel (2017) and the robustness tests they propose, which include a random effects model, manually detrending our outcomes, and re-running the parallel trends omitting the most and least recent time-period. We also show that the results remain similar when estimated solely off the sample of ever-treated cities.

Panel A: Random Effects Model

	(1) Log (1+Total Accidents)	(2) Log (1+Total Fatalities)
Post * Treated	0.0184***	0.0175***
	(0.0063)	(0.0066)
Quarter Fixed Effects	Yes	Yes
Control Variables	Yes	Yes
Observations	190080	190080
R^2	0.3235	0.3135

Panel B: Estimate Time Effects Solely from Control Cities

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
Post * Treated	0.0241***	0.0696***	0.0718***	0.0234***	0.0718***	0.0737***
	(0.0059)	(0.0066)	(0.0085)	(0.0061)	(0.0069)	(0.0089)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Fixed Effects	No	No	Yes	No	No	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	190080	190080	190080	190080	190080	190080
R^2	0.606	.616	0.623	0.591	0.601	0.609

Table A3. Staggered Difference-in-Differences Sensitivity Analysis

Notes: In this table, we conduct several sensitivity analyses for staggered difference-in-differences models. Panel A presents results from random effect models. Panel B presents results from generalized difference-in-difference regressions, where all outcomes and control variables are manually detrended using the means of the variables estimated from the control cities. Hence, the time effects are identified solely from the control group. We provide a description of the variables in Section 3. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log (Total Accidents)			Log (Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0178**	0.0281***	0.0283***	0.0176**	0.0274***	0.0279***
	(0.0073)	(0.0083)	(0.0097)	(0.0076)	(0.0087)	(0.0102)
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	110272	110272	110272	110272	110272	110272
R2	0.679	0.687	0.693	0.665	0.673	0.68

Table A4. Estimating Main Effects from Treated Cities Only

Notes: This table presents results from generalized difference-in-difference regressions when we drop all cities that never adopted ridehailing. The dependent variables are listed at the top of each column. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors—clustered at city level—are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

In Figure A1, we further follow the suggestions of Borusyak and Jaravel (2017) and use the first event year and the last event year before treatment as the base periods for estimation. To map out the pattern in the counterfactual treatment effects, we regress the various outcome measures on the annual lag and lead indicators for the entry of ridehailing.

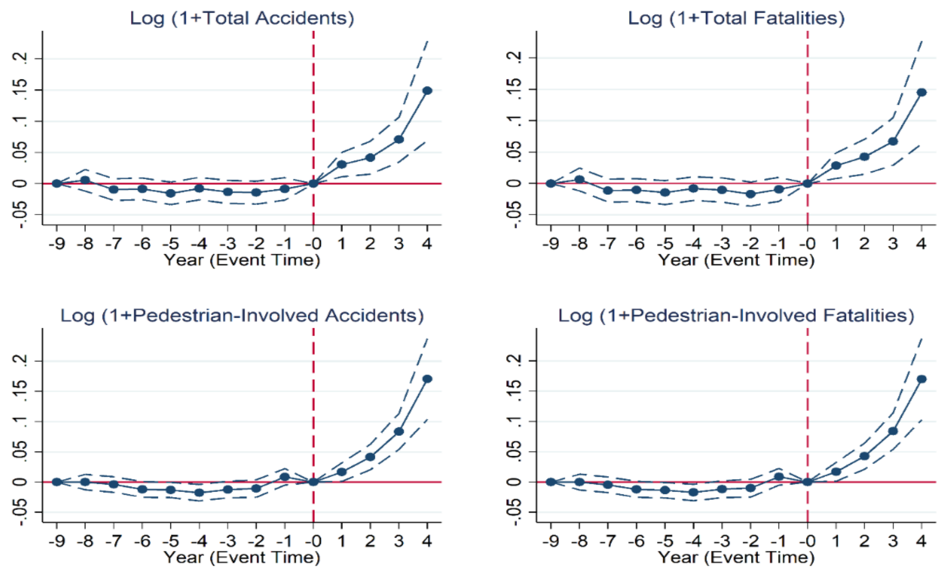

Figure A1. Annual Difference-in-Differences Estimators: Changing Base Years

This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. We follow the suggestions of Borusyak and Jaravel (2017) and use the first event year and the last event year before treatment as the base periods for estimation. To map out the pattern in the counterfactual treatment effects, we regress the various outcome measures on the annual lag and lead indicators for the entry of ridehailing. The vertical red line indicates the quarter of entry. We provide a description of the variables in section 3.

We then plot our parallel trend graphs using quarterly data to corroborate our main figures in the paper (which are done at the annual level).

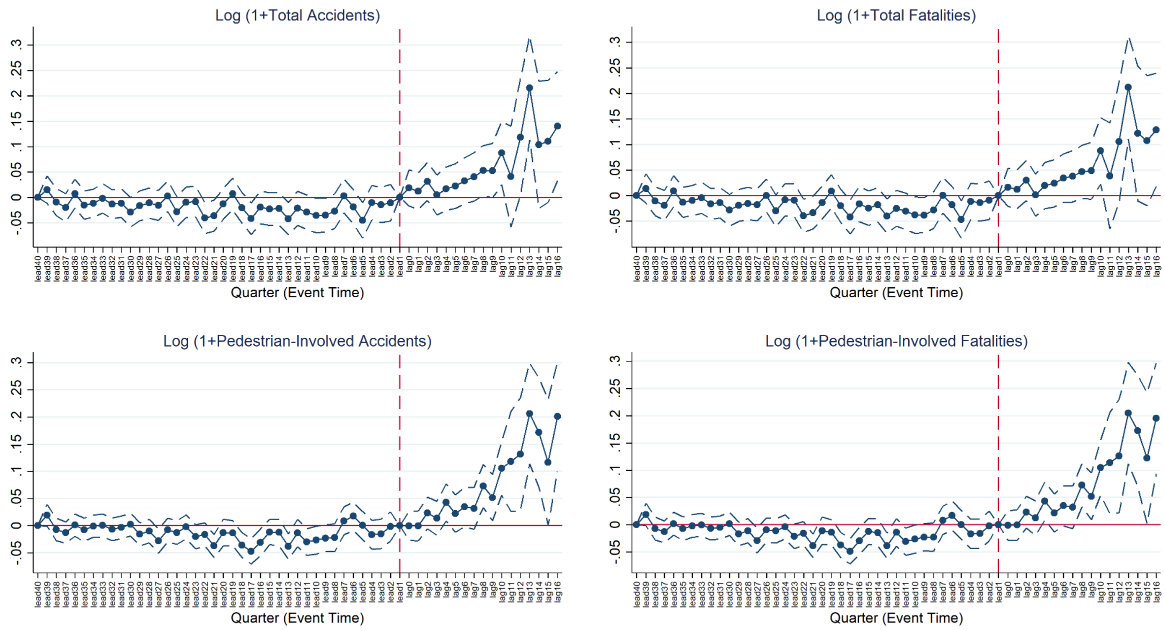

Figure A2. Quarterly Difference-in-Difference Estimators

Notes: This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To map out the pattern in the counterfactual treatment effects, we regress the various outcome measures on the quarterly lag and lead indicators for the entry of ridehailing. The vertical red line indicates the quarter of entry. We provide a description of the variables in section 3.

In Table A5 below, we repeat our analysis in Table 2 Panel B, and conduct a coarsened exact matching for treated and control cities based on income and population. We focus on pre-treatment income and population for every city (before 2010, the year in which the first treatment occurred).

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0196**	0.0367***	0.0433***	0.0189**	0.0377***	0.0428***
	(0.0081)	(0.0104)	(0.0135)	(0.0083)	(0.0114)	(0.0143)
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	155,008	155,008	155,008	155,008	155,008	155,008
R2	0.56	0.57	0.58	0.55	0.56	0.57

Table A5. Coarsened Exact Matching Using Pre-treatment Characteristics

Notes: This table presents results from difference-in-difference estimators on a matched sample. We repeat our analysis in Table 2 Panel B, and conduct a coarsened exact matching for treated and control cities based on income and population. We focus on pre-treatment income and population for every city (before 2010, the year in which the first treatment occurred). Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

We also conduct a placebo test to examine whether the magnitude of our inferences are be spurious relative to estimates generated from a simulation of random post periods in the sample. We find no discernable patterns of accidents and drunk accidents in event time before and after placebo RH entry. We use locations that did not adopt RH (to eliminate contamination from cities that adopted), and for each run, we 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 the 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.

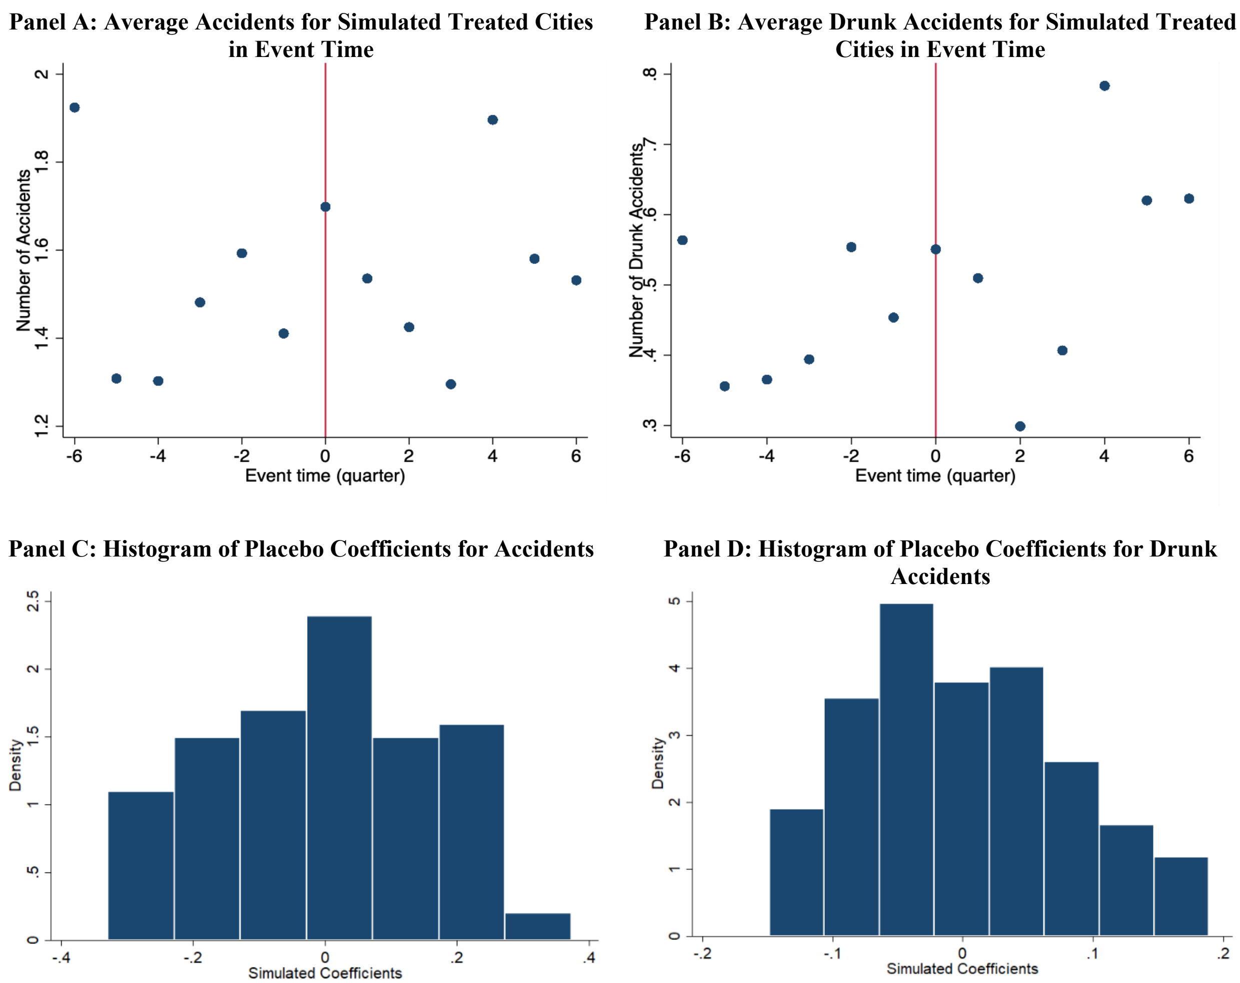

Figure A3. Placebo tests on Rideshare Entry

Notes: This figure illustrates the simulated ridehail entry effects on fatal accidents (Panel A and C) and drunk accidents (Panel B and D) in event time. We conduct a placebo test using only locations that did not adopt ridehail during the sample period to eliminate contamination from cities that adopted. We use various cutoffs for the city population. For each run, we simulate 100 cities adopting ridehail by assigning a random adoption date from within the set of actual launch dates in the cities that adopted. We run the simulation 100 times. We then plot the average number of accidents and drunk accidents per 100,000 population in simulated event time, and the histogram of coefficients on $Post_t * Treated_c$ from placebo regressions.

To help give more context to the empirical design, we also run the Goodman-Bacon decomposition and analyze the weights underlying our staggered DiD regressions. As Goodman-Bacon (2018) notes, the two-way fixed-effect estimator is a weighted average of all potential 2x2 DiD estimates, 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 and dig deeper into the comparisons that contribute to the coefficient in our main table.

Table A6 and Figure A4 below show the results of the decomposition for the DD estimates. Most of the variation used to estimate the DD coefficients results from the cleanest comparison of treated states to never treated states ("Never vs Timing", which compares states that adopted the policy at some point during the sample period and those that did not). The average estimate derived from this source of variation is 0.018 for accidents, and 0.019 for fatalities, and both have a weight of 83%.

Panel A: Accidents		
DD Comparison	Beta	Weight
Timing Groups	0.011	0.166
Never vs Timing	0.019	0.833
Within	1.109	0.001
Panel B: Fatalities		
DD Comparison	Beta	Weight
Timing Groups	0.012	0.166
Never vs Timing	0.018	0.833
Within	1.289	0.001

Table A6. Goodman-Bacon Decomposition

Notes: This table displays the decomposition of the staggered difference-in-differences coefficients using the method suggested by Goodman-Bacon (2018). The outcome variable in Panel A is the natural logarithm of one plus accidents. The outcome variable in Panel B is the natural logarithm of one plus fatalities.

Panel A: Accidents

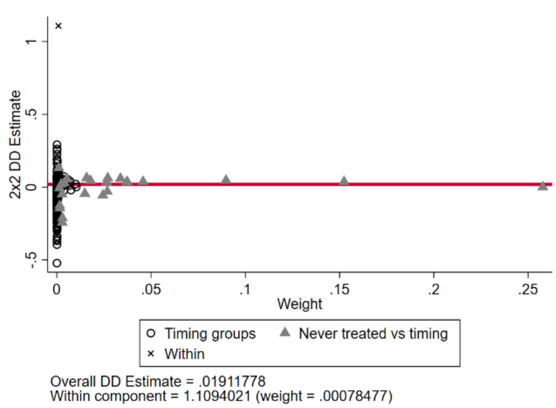

Panel B: Fatalities

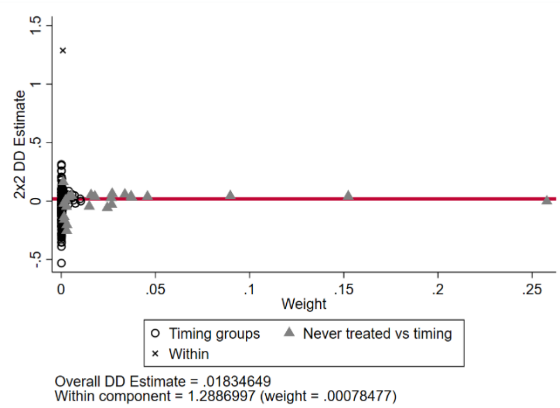

Figure A4. Goodman-Bacon Decomposition

Notes: This figure displays the decomposition of the staggered difference-in-differences coefficients using the method suggested by Goodman-Bacon (2018). The outcome variable in Panel A is the natural logarithm of one plus accidents. The outcome variable in Panel B is the natural logarithm of one plus fatalities.

B. ROBUSTNESS TESTS

In this section, we show that the effects of ridehailing on traffic accidents and fatalities are robust to variations in empirical methodology, measures, control variables, and subsamples.

B.1. Variations in empirical methodology

Below, we demonstrate the effects of adding our control variable, fixed effects, and location-specific trends one-by-one to the model. We graph the coefficients of interest (*POST x TREATED*) in the corresponding Figure B1.

Panel A: Total Accidents

	Log (1+Total Accidents)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Post\ X\ Treated_{t,c}$	0.1832***	0.1341***	0.0405***	0.0803***	0.0191***	0.0360***	0.0332***
	(0.0204)	(0.0105)	(0.0059)	(0.0125)	(0.0062)	(0.0074)	(0.0091)
Control Variables	No	Yes	Yes	Yes	Yes	Yes	Yes
City FE	No	No	Yes	No	Yes	Yes	Yes
Year-Quarter FE	No	No	No	Yes	Yes	Yes	Yes
City Linear Trend	No	No	No	No	No	Yes	Yes
City Quadratic Trend	No	No	No	No	No	No	Yes
Observations	190,080	190,080	190,080	190,080	190,080	190,080	190,080
R2	0.038	0.346	0.605	0.352	0.610	0.619	0.626

Panel B: Total Fatalities

	Log (1+Total Fatalities)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Post\ X\ Treated_{t,c}$	0.1814***	0.1360***	0.0411***	0.0797***	0.0183***	0.0360***	0.0335***
	(0.0208)	(0.0108)	(0.0061)	(0.0128)	(0.0065)	(0.0077)	(0.0095)
Control Variables	No	Yes	Yes	Yes	Yes	Yes	Yes
City FE	No	No	Yes	No	Yes	Yes	Yes
Year-Quarter FE	No	No	No	Yes	Yes	Yes	Yes
City Linear Trend	No	No	No	No	No	Yes	Yes
City Quadratic Trend	No	No	No	No	No	No	Yes
Observations	190,080	190,080	190,080	190,080	190,080	190,080	190,080
R2	0.036	0.335	0.590	0.340	0.595	0.604	0.612

Table B1. Effect of Ridehailing on Traffic Safety in Different Model Specifications

Notes: This Table illustrates the changes in the generalized difference-in-difference regression coefficient estimates when using different model specifications. The dependent variables are in Panel A are the natural logarithm of one plus total fatal accidents, and in Panel B are the natural logarithm of one plus total fatalities. We provide a detailed description of these variables in section 2. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables include the natural logarithm of population and the natural logarithm of income per capita. The lower order interaction term, *Treated*, is included in the regressions, and is absorbed when city fixed effects are included. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

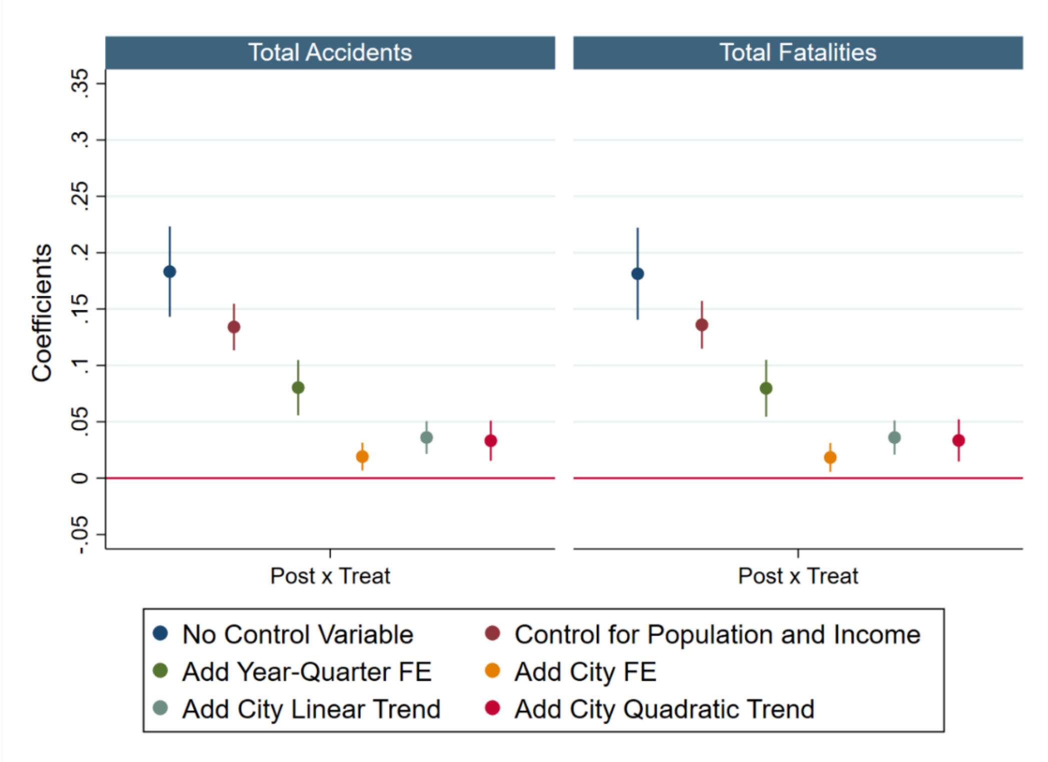

Figure B1. Effect of Ridehailing on Traffic Safety in Different Model Specifications

Notes: This figure shows the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level. To demonstrate how the point estimates and standard errors change across different model specifications, we incrementally add control variables, fixed effects, and city-specific linear and quadratic time trends into the difference-in-difference regression specification. The outcome variables are displayed at the top of each graph. We provide a description of the variables in Section 3.

As further robustness, we bottom code the dependent variable at the 5th percentile of non-zero values, using the log of the bottom coded variable as the dependent variable. We find consistent results to those of our main specification.

	Log (Total Accidents)			Log (Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0242***	0.0526***	0.0473***	0.0235***	0.0519***	0.0473***
	(0.0058)	(0.0069)	(0.0084)	(0.0061)	(0.0074)	(0.0090)
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	190080	190080	190080	190080	190080	190080
R2	0.67	0.678	0.685	0.64	0.649	0.656

Table B2. Bottom Coding at the 5th Percentile

Notes: This table presents robustness of our main results in Table 2 Panel A, when using a different way to construct our outcome variables. We bottom-code our outcome variables, Total Accidents and Total Fatalities, at the 5th percentile of non-zero values, and use the log of the bottom coded variables as the dependent variables. Total accidents is the number of fatal accidents according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, clustered at city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

In Table below we demonstrate that our results are robust to an alternative formulation of the outcome variable. Specifically, we demonstrate that similar estimates are obtained when employing the inverse hyperbolic sine instead of log(1+accidents) and log(1+fatalities).

	(1) $\sinh^{-1}(Accidents)$	(2) $\sinh^{-1}(Fatalities)$
$Post\ X\ Treated_{t,c}$	0.0443***	0.0439***
	(0.0095)	(0.0099)
City and Quarter Fixed Effects	Yes	Yes
City Linear Trend	Yes	Yes
Control Variables	Yes	Yes
Observations	190,080	190,080
R2	0.61	0.59

Table B3. Robustness to inverse hyperbolic sine transformation

Notes: This table presents robustness of our main results in Table 2 Panel A, when using a different way to construct our outcome variables. We use the inverse hyperbolic sine transformation of the two variables, Total Accidents and Total Fatalities, as the dependent variables. Total accidents is the number of fatal accidents according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, clustered at city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

In Table B4 below, we show effects from weighted regressions. In Columns (1) and (3), each observation is weighted by the population of the city. In Columns (2) and (4), each observation is weighted by the square root of the city population. The effects are statistically significant and larger in magnitudes than our main effects.

	Log (Total Accidents)		Log (Total Fatalities)	
	(1)	(2)	(3)	(4)
	W = Population	W = Sq. rt (Population)	W = Population	W = Sq. rt (Population)
$Post\ X\ Treated_{t,c}$	0.0610***	0.0466***	0.0624***	0.0474***
	(0.0128)	(0.0083)	(0.0125)	(0.0087)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes	Yes
Observations	190,080	190,080	97,669	190,080
R2	0.92	0.79	0.91	0.78

Table B4. Population Weighting

Notes: This table presents results from weighted difference-in-difference regressions. In Columns (1) and (3), the weight for each observation is the city's population. In Columns (2) and (4), the weight for each observation is the square root of the city's population. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

To further assuage concerns with functional form issues in our dependent variable, we then run both an extensive margin regression where the 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 (fatal accidents).

Panel A: Intensive Margin

	Log (Total Accidents)			Log (Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0374***	0.0728***	0.0587***	0.0374***	0.0713***	0.0586***
	(0.0093)	(0.0113)	(0.0142)	(0.0098)	(0.0119)	(0.0150)
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	97669	97669	97669	97669	97669	97669
R2	0.639	0.65	0.661	0.601	0.615	0.626

Panel B: Extensive Margin

	I (Total Accidents>0)			I (Total Fatalities>0)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0036	-0.0045	-0.0012	0.0036	-0.0045	-0.0012
	(0.0053)	(0.0065)	(0.0083)	(0.0053)	(0.0065)	(0.0083)
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	190080	190080	190080	190080	190080	190080
R2	0.322	0.335	0.346	0.322	0.335	0.346

Table B5. Extensive Margin and Intensive Margin Regressions

Notes: In this table, we decompose our main results in Table 2 Panel A by separately running extensive margin and intensive margin regressions. In Panel A, the dependent variables are the natural log of total accidents and total fatalities. Naturally, the sample only includes city-quarters with positive values of accidents and fatalities. In Panel B, the dependent variables are indicator variables for whether total accidents and fatalities are above zero. Total accidents is the number of fatal accidents according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, clustered at city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

B.2. Variations in measurement

In Table below, we demonstrate robustness to clustering standard errors at the CBSA level (Columns (1) and (2)), and conduct difference-in-difference tests in which we define treatment the

CBSA level based on the earliest treatment date within a CBSA (Columns (3) and (4)). Our inferences remain unchanged in both cases.

In many cases, RH was rolled out at the MSA or CBSA level rather than to an individual city within the CBSA. In Table below, we also conduct difference-in-difference tests in which we define treatment the CBSA level based on the earliest treatment date across cities within a CBSA (Columns (3) and (4)). Our inferences remain unchanged in both cases.

	Cluster Standard Errors at CBSA Level		Diff-in-Diff at CBSA Level	
	(1)	(2)	(3)	(4)
	Log (1+Total Accidents)	Log (1+Total Fatalities)	Log (1+Total Accidents)	Log (1+Total Fatalities)
$Post\ X\ Treated_{t,c}$	0.0360***	0.0360***	0.0524***	0.0579***
	(0.0095)	(0.0097)	(0.0148)	(0.0153)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes	Yes
Observations	190,080	190,080	52,928	52,928
R2	0.62	0.60	0.81	0.79

Table B6. CBSA-level analysis and standard error cluster

Notes: In columns (1) and (2), we run the same specification as in Table 2 Panel A, but cluster standard errors at the CBSA level instead. In columns (3) and (4), the difference-in-difference specifications are run at CBSA level, utilizing the earliest adoption date within the CBSA. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors—clustered at CBSA level—are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

A competing explanation for our observed results is that the effect is merely driven by city growth as opposed to RH. While our difference-in-differences estimation approach should address this concern, we do one set of additional analyses to further bolster our inference. Table B7 below presents estimates from models similar to those in our main specifications, but where we replace the LHS variable with a measure of city economic conditions instead of fatal accidents: the natural logarithm of average weekly wage in a city in a given quarter. If our fatal accident measures are merely picking up city growth that is reflected in additional VMT, we would expect to this city growth reflected in wages. As can be seen from the estimates in the table, however, we observe no evidence of a statistically significant growth in wages in post-RH. Rather, the estimates on $Post\ X\ Treated_{t,c}$ are negative and statistically significant, suggesting that we are not simply picking up economic growth in the city.

	Log Average Weekly Wage			
	(1) >2000	(2) >2005	(3) Treat = 1	(4) >2005 & Treat=1
$Post\ X\ Treated_{t,c}$	-0.0093***	-0.0045***	-0.0079***	-0.0032***
	(0.0009)	(0.0009)	(0.0011)	(0.0009)
Log Pop	0.0262***	0.0576***	0.0250**	0.0646***
	(0.0095)	(0.0132)	(0.0123)	(0.0157)
Log Income Per Capita (lag)	0.2597***	0.2202***	0.2603***	0.1865***
	(0.0140)	(0.0131)	(0.0144)	(0.0114)
Unemployment Rate (lag)	-0.0033***	-0.0029***	-0.0022***	-0.0021***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)
Observations	198,238	142,017	115,024	82,401
R-squared	0.9818	0.9793	0.9813	0.9785
City FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes

Table B7. Are results driven by city growth?

Notes: This table presents results from generalized difference-in-difference regressions. The dependent variable is the natural logarithm of average weekly wage in a city in a given quarter. $Post\ X\ Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables in the regressions include the natural logarithm of population, income per capita (lagged one quarter), and unemployment rate (lagged one quarter). Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

B.3. Additional controls

We next demonstrate the robustness of the results to controlling for additional variables, such as population growth (instead of population level), unemployment, gas prices and gas taxes. Coefficients are reported below.

	Log (1+Total Accidents)			
	(1)	(2)	(3)	(4)
$Post\ X\ Treated_{t,c}$	0.0361***	0.0292***	0.0335***	0.0361***
	(0.0074)	(0.0074)	(0.0074)	(0.0074)
Population Growth Rate	0.1007			
	(0.1012)			
Unemployment Rate		-0.0130***		
		(0.0013)		
Retail Gas Price			0.1163***	
			(0.0201)	
Retail Gas Price Change (%)				0.0012***
				(0.0004)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes	Yes
Observations	187,110	190,065	190,080	190,080
R2	0.62	0.62	0.62	0.62

Table B8. Robustness to Additional Controls

Notes: This table presents results from generalized difference-in-difference regressions when controlling for various additional control variables. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Population growth rate is the annual percentage growth in city population? Unemployment rate is the quarterly average county unemployment rate. Retail Gas Price is average quarterly retail gasoline price (dollars per gallon). Retail Gas Price Change is the quarterly percentage change in retail gasoline price. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

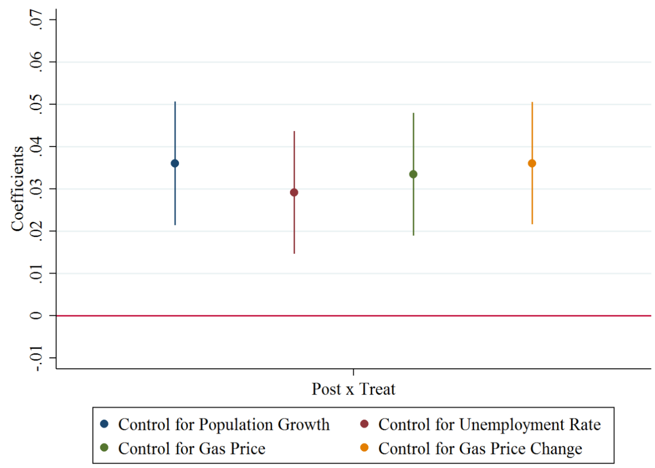

Figure B2. Robustness to Additional Control Variables

Notes: This figure illustrates how the regression coefficients vary when we include additional control variables (See Table B8 for details). The outcome measures for all regressions are Log (1+Total Accidents). Two-tailed 95% confidence intervals based on standard errors clustered at the city level are displayed in the figure. We provide a description of the variables in Section 3.

Of course, another natural concern is that our models are somehow merely picking up increases in smartphone usage during driving, which leads to distracted driving and in turn to additional accidents. By necessity, RH involves smartphone usage, and thus, one possibility is that in fact RH leads to increased accidents through RH drivers' adoption and usage of smartphones and the associated distracted driving. This, however, would still mean the increase in accidents is as a result of RH entry. The alternative we would like to rule out, however, is that RH merely coincides perfectly with smartphone adoption patterns unrelated to RH itself, and thus that our results spuriously pick up this unrelated increase in smartphone usage. While we expect that the structure of the DD models and FE should absorb much of these types of trends, we attempt to more directly address this concern by controlling in our models for proxies for 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"). Table B9 below demonstrates the robustness of our main finding to the inclusion of these four control variables. Accidents load positively on the cellphone sales and employment, but do not affect the significance or the magnitude of the main effect, which remains robust in all specifications.

Panel A: Total Accidents

	(1)	(2)	(3)	(4)	(5)
$Post\ X\ Treated_{t,c}$	0.0304***	0.0307***	0.0301***	0.0307***	0.0400***
	(0.0091)	(0.0091)	(0.0091)	(0.0091)	(0.0078)
Log (Cellphone Store Gross Receipts)				0.0181*	
	0.0061			(0.0106)	
	(0.0082)				
Log (Cellphone Store Employment)		0.0340**		0.0405***	
		(0.0143)		(0.0147)	
Log (Number of Cellphone Store Establishments)			-0.0096	-0.0318**	
			(0.0107)	(0.0143)	
					-0.0004
Log (Smartphone-Related Google Search Vol)					(0.0023)
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	92,316	92,316	92,316	92,316	92,316
R2	0.67	0.67	0.67	0.67	0.63

Panel B: Total Fatalities

	(1)	(2)	(3)	(4)	(5)
$Post\ X\ Treated_{t,c}$	0.0309***	0.0312***	0.0307***	0.0313***	0.0406***
	(0.0096)	(0.0096)	(0.0096)	(0.0096)	(0.0082)
Log (Cellphone Store Gross Receipts)	0.0073			0.0178	
	(0.0087)			(0.0112)	
Log (Cellphone Store Employment)		0.0352**		0.0408***	
		(0.0152)		(0.0157)	
Log (Number of Cellphone Store Establishments)			-0.0064	-0.0285*	
			(0.0117)	(0.0154)	
Log (Smartphone-Related Google Search Vol)					-0.0001
					(0.0025)
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	92,316	92,316	92,316	92,316	92,316
R2	0.66	0.66	0.66	0.66	0.61

Table B9: Ridehailing Effects and Cellphone Usage Controls

Notes: This table presents results from generalized difference-in-difference regressions when controlling for various additional control variables related to cellphone adoption. Gross receipts, employment, and the number of establishments are measured for industry with NAICS 517312: Wireless Telecommunication Carriers (Except Satellite) by MSA and year from 2007-2016. Log Smartphone-Related Google Search Vol is the natural logarithm of Google search volume for the terms “iPhone”, “Android”, “Samsung galaxy”, “smartphone”, and “cellphone”. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

B.4. Subsamples

We run our main specification in a shorter sample in Table B10 below to show the robustness of our inferences.

	(1) Log (1+Total Accidents)	(2) Log (1+Total Fatalities)
$Post\ X\ Treated_{t,c}$	0.0309***	0.0315***
	(0.0084)	(0.0088)
City and Quarter Fixed Effects	Yes	Yes
City Linear Trend	Yes	Yes
Control Variables	Yes	Yes
Observations	118,800	118,800
R2	0.63	0.61

Table B10 Difference-in-Differences Estimates Using a Shorter Sample 2007-2016

Notes: This table presents results from generalized difference-in-difference regressions using a shorter sample than our main specification. The sample includes all quarters from the first quarter of 2007 to the last quarter of 2016. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

In Table B11 below, we break out weekend accidents, nighttime accidents, weekday accidents, and weekend night accidents for total fatal accidents (Panel A) and total fatalities (Panel B). 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 Figure 4. 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 see in the heterogeneity estimations: that the effects of RH appear to be larger in larger cities.

Panel A: Total Accidents

	(1) Log (1+Weekday Accidents)	(2) Log (1+Weekend Accidents)	(3) Log (1+Accidents at Night)	(4) Log (1+Accidents at Fri. and Sat. Night)
$Post\ X\ Treated_{t,c}$	0.0272***	0.0350***	0.0396***	0.0250***
	(0.0069)	(0.0065)	(0.0065)	(0.0054)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	190,080
R2	0.52	0.54	0.55	0.45

Panel B: Total Fatalities

	(1) Log (1+Weekday Fatalities)	(2) Log (1+Weekend Fatalities)	(3) Log (1+Fatalities at Night)	(4) Log (1+Fatalities at Fri. and Sat. Night)
$Post\ X\ Treated_{t,c}$	0.0276***	0.0356***	0.0407***	0.0269***
	(0.0072)	(0.0069)	(0.0069)	(0.0058)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes
Control Variables	Yes	Yes	Yes	Yes
Observations	190,080	190,080	190,080	190,080
R2	0.51	0.52	0.54	0.44

Table B11: Effect of Ridehailing on Traffic Safety by Day and Time

Notes: This table presents the effect of ridehailing on accidents and fatalities, respectively in Panel A and B, by day and time. The dependent variables are listed at the top of the columns. Total accidents is the number of fatal accidents according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. Weekday is defined as Monday through Thursday. Weekend is defined as Friday through Sunday. Night is defined as 5 pm through 2 am. Friday and Saturday Night is defined as 5 pm Friday through 6 am Saturday and 5 pm Saturday through 6 am Sunday. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Accidents and Fatalities on Weekend and at Nights

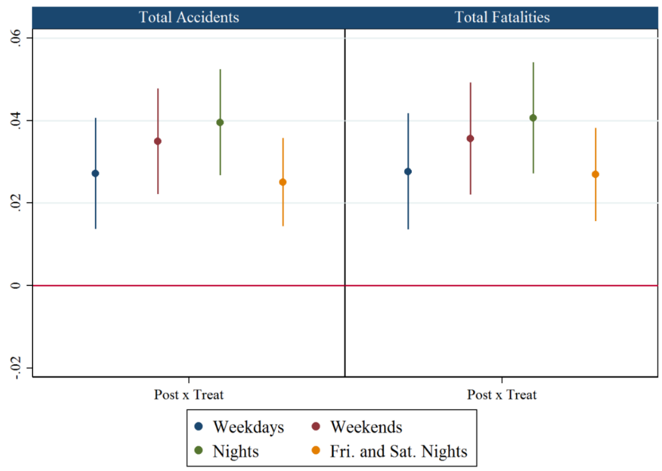

Panel B: Small vs. Large Cities

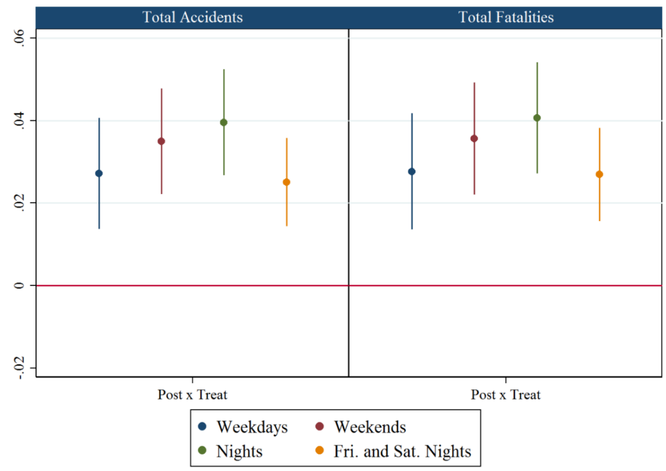

Figure B3: Nights and Weekends

Note: This figure displays the regression coefficient estimates and two-tailed 95% confidence intervals based on standard errors clustered at the city level, broken down by accidents at night and on the weekend. Panel A presents estimates for accidents and fatalities, while Panel B presents the coefficients for accidents separately for small and large cities. We provide a description of the variables in Section 3.

C. ADDITIONAL ANALYSIS

C.1. Further description of Texas A&M Mobility Scorecard data

In this section we describe the underlying data for the TAMU data for VMT, excess fuel consumption and hours of delay in traffic used in the paper. Our VMT measure comes from the Urban Mobility Scorecard which is produced annually from 1982-2014 by TAMU's Transportation Institute, one of the largest university-affiliated transportation research agencies in the U.S.

The base data for the TAMU's Urban Mobility Scorecard comes from two sources: INRIX and the U.S. Department of Transportation. INRIX is a private company that collects traffic speed data from a variety of sources including commercial vehicles, smart phones and connected cars with location devices. It uses a proprietary algorithm to filter inappropriate data (e.g., pedestrians walking next to a street) and provides TAMU with a dataset of average speeds for each road segment every 15 minutes. This dataset combined with the daily volume and roadway inventory data from Federal Highway Administration (FHWA)'s Highway Performance Monitoring System (HPMS) files are used to calculate travel mileage and delay statistics.

The key steps are the following. First, the daily traffic volume data are divided into the same time interval as the traffic speed data (15-minute intervals). The traffic volume is measured using average daily traffic (ADT) count. Second, because the geographic referencing systems are different for the speed and volume datasets, a geographic matching process utilizing Geographic Information Systems (GIS) tools are performed to assign traffic speed data to each HPMS road section. Third, daily vehicle-miles of travel are calculated as the ADT of a section of roadway multiplied by the length (in miles) of that section of roadway and is aggregated at the annual level. Last, in order to calculate congestion measures such as excess fuel consumption and hours of delay, a free-flow speed was estimated using the speed at low volume conditions (for example, 10 p.m. to 5 a.m.) for each roadway section and hour of the week. Excess fuel consumption (hours of travel) are calculated as the difference between the observed fuel consumption (hours of travel) and the free-flow fuel consumption (hours of travel). Any observed speed faster than the free-flow speed is changed to the free-flow speed so that congestion measures are capped at zero, rather than providing a negative value.

While much work is put into generating these estimates, there is likely a good deal of imputation present: for most roads in most states, real measurements are infrequent, every three years according to guidelines (and state compliance with the FHWA guidelines are not documented as far as we know). In principle, these potential VMT data issues would attenuate the coefficient towards not finding an effect, so the real effect on VMT might be larger than that which we document.

C.2. Additional Quality Evidence

One way to determine whether the main effect we document is driven primarily by the quantity channel is to re-estimate our models, adding various VMT (quantity) measures as additional RHS variable, and seeing whether the VMT control absorbs the effect. We do this in the table below. In Panel A, we run the specification on 99 urban areas. In Panel B, we disaggregate to the place (city) level within the urban area. In Panel A, controlling for VMT, we observe a positive coefficient on Post * Treated; the magnitude of the estimate is about one-quarter of the total accident increase we estimate in our main specifications. In Panel B, we have more power but are assuming VMT increases are spread evenly across incorporate places within the larger urban area. Here we also observe a positive coefficient of interest, controlling for VMT, this time statistically significant. Both panels are consistent with the documented increase in accidents being attributable to more than merely a quantity effect.

That said, there are caveats to this analysis. The sample size is significantly reduced, the VMT estimates are recorded only once a year, and considerable imputation is involved, as described above. Thus, we expect to have measurement error which should attenuate the effect and issues of power given that we are trying to identify a relatively small and precise effect.

Panel A: Urban Area Level

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,u}$	0.0044	0.0034	0.0042	0.0034	0.0027	0.0033
	(0.0322)	(0.0320)	(0.0321)	(0.0323)	(0.0321)	(0.0322)
Log Arterial Street Daily VMT	-0.0021			0.0134		
	(0.1503)			(0.1487)		
Log(Freeway Daily VMT)		0.0522			0.0454	
		(0.1459)			(0.1262)	
Log(Arterial Street and Freeway Daily VMT)			0.0110			0.0168
			(0.1826)			(0.1652)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes	Yes	Yes
City Quadratic Trend	No	No	No	No	No	No
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	990	990	990	990	990	990
R2	0.98	0.98	0.98	0.98	0.98	0.98

Panel B: Place Level

	Log (1+Total Accidents)		Log (1+Total Fatalities)		Log (1+Fatal Accidents Involving Pedestrians)		Log (1+Fatalities in Accidents Involving Pedestrians)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Post\ X\ Treated_{t,c}$	0.0431***	0.0428***	0.0441***	0.0437***	0.0274**	0.0275**	0.0281**	0.0282**
	(0.0140)	(0.0140)	(0.0145)	(0.0145)	(0.0126)	(0.0126)	(0.0128)	(0.0128)
Log Arterial Street Daily VMT	0.0379		0.0391		-0.0295		-0.0282	
	(0.0464)		(0.0486)		(0.0424)		(0.0427)	
Log(Arterial Street and Freeway Daily VMT)		0.0827		0.0906		-0.0341		-0.0355
		(0.0554)		(0.0583)		(0.0616)		(0.0616)
City and Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City Linear Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City Quadratic Trend	No	No	No	No	No	No	No	No
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	87,584	87,584	87,584	87,584	87,584	87,584	87,584	87,584
R2	0.72	0.72	0.71	0.71	0.65	0.65	0.64	0.64

Table C1. Difference-in-Difference Estimation Controlling for Quantity Channel

Notes: This table presents results from generalized difference-in-difference regressions when controlling for various proxies for vehicle miles traveled. In all panels, the dependent variables are listed at the top of the columns. Total accidents is the number of fatal accidents, according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. 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. In Panel A, we run the specification on 99 urban areas. $Post * Treated_{t,u}$ is a dummy variable that equals one if urban area u adopted at least one ridehailing service at year t. Urban area-specific linear trends are included in all regressions. In Panel B, the specifications are estimated at the city year level and $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. City-specific linear trends are included in all regressions in Panel B. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, clustered at the urban level in Panel A and clustered at the city level in Panel B, they are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

C.3. Drunk Accidents

While our main empirical question centers on the patterns for total fatal accidents, much has been made of the possible effects that RH has specifically on drunk fatal accidents, which comprise a quarter to a third of all fatal accidents. A reduction in alcohol-involved fatal accidents would not be inconsistent with our overall findings; drunk accidents could go down, but the decrease could still be swamped by increases in nondrunk accidents, leading to an overall increase.

In this section, we provide further insight into the patterns for accidents involving an impaired driver. We further discuss how our findings on drunk accidents and fatalities relate to prior work exploring the relationship between ridehailing and drunk accidents. The costs of drunk driving have been extensively documented. Drunk driving costs the U.S. tens of thousands of lives and billions of dollars in law enforcement, property damage, and lost productivity each year. The NHTSA estimates the direct economic cost of these accidents to be $44 billion and estimates the total societal costs at $201 billion.

As a result of these costs, extensive amounts of resources have been utilized to develop public policy to discourage intoxicated driving. Most policies focus on deterrence through fear of punishment and increasing the expected cost of a conviction (i.e. increasing the penalties for drunk driving convictions through higher fines, more jail time, and/or driver's license confiscation). Previous work on ridehailing services has proposed that ridehailing, by increasing the convenience of and decreasing the cost of alternative transportation, may in turn lead to fewer drunk drivers and thus reduce the level of fatalities due to impaired driving. Several previous studies have made attempts to examine the effect of RH on fatal accidents in U.S. cities (see e.g., Brazil and Kirk, 2016; Greenwood and Wattal, 2016; Martin-Buck, 2016; Dills and Mulholland, 2018). 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.

There are a number of reasons, however, why our secondary results looking at drunk accident and fatalities do not align with the conclusions of these prior studies. We discuss these in turn: (1) a definitional change in how alcohol-involved accidents are classified half way through our sample period; (2) smaller samples; (3) the specific coefficients that are interpreted by the studies.

Definition Changes for Alcohol Involvement

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 lead to a massive mechanical decrease in accidents classified as "alcohol-related" and to a corresponding increase in non-drunk accidents, as graphed below.

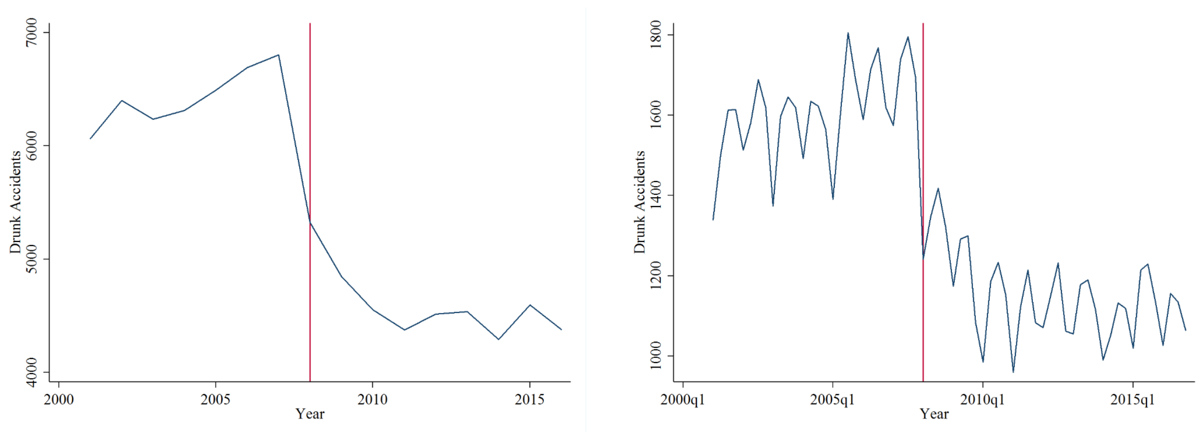

Panel A: Annual **Panel B: Quarterly**

Figure C1. Accidents Classified as Alcohol-Related, Annual and Quarterly

Notes: This figure displays the total number of drunk accidents through time. The vertical line indicates the year or quarter of the NHTSA drunk accidents categorization change. Specifically, prior to 2008, alcohol-related accidents were recorded as any fatal accident involving at least one vehicle occupant (driver or non-driver) or pedestrian being impaired (in the legal sense). From 2008 onwards, alcohol-related accidents are recorded as such only if the *driver* was impaired.

This mechanical drop has clear effects when estimating models that include the pre-change years. Figure A8 shows the coefficient of interest for the estimated "effect" of RH for drunk accidents when using different sample periods. The first four specifications include some years of the period prior to the definitional change. The second two include only years after the change, when a drunk accident clearly involved an impaired *driver*. These specifications are estimated in similar fashion to those in prior studies, such as Martin-Buck (2016), which do not control for location-specific linear and quadratic trends (Martin-Buck (2016) uses a sample period of 2000-2014). In the four specifications contaminated by the old definition of drunk accident, we observe an estimated negative coefficient. In the two uncontaminated specifications, we observe a positive coefficient, consistent with our main results.

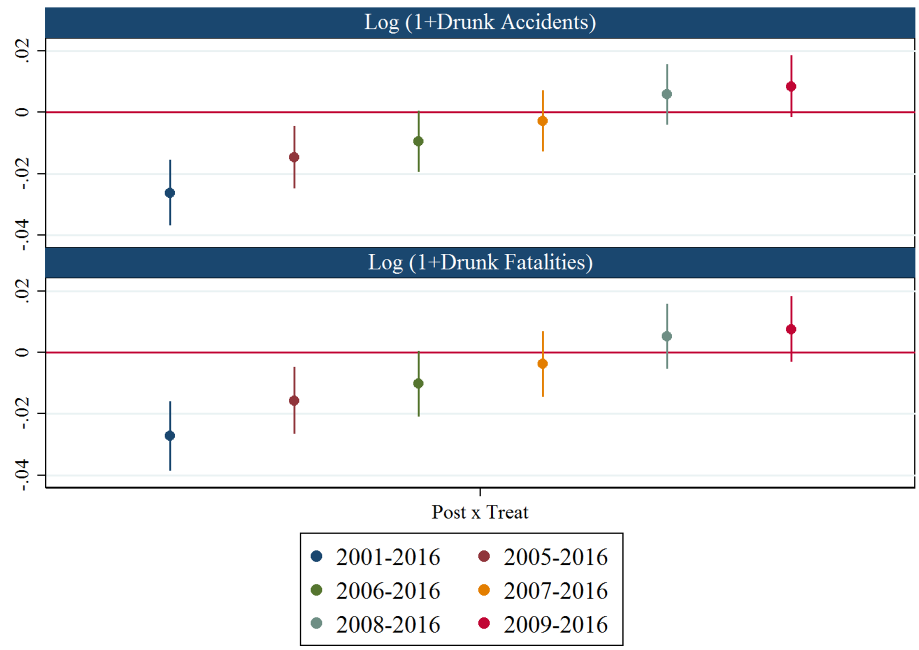

Figure C2. Drunk Accident and Fatalities Coefficient Estimates in Different Sample Periods

Notes: This Figure illustrates how the regression coefficient estimates of the effect of ridehailing on drunk accidents and fatalities vary with different sample time periods. The outcome measure for all each panel are displayed at the top of the panel. Two-tailed 95% confidence intervals based on standard errors clustered at the city level are displayed in the figure.

For contrast, in Figure C3 below we provide the same coefficients for total fatal accidents, which are not affected in any way by the definition change for drunk accidents, as they are the sum of drunk and nondrunk. As expected, these coefficients are consistent across all the sample periods, and do not vary in sign based on sample years, unlike the drunk accident estimates, which are affected mechanically by the definition change leading to a drop in accidents close to RH entry.

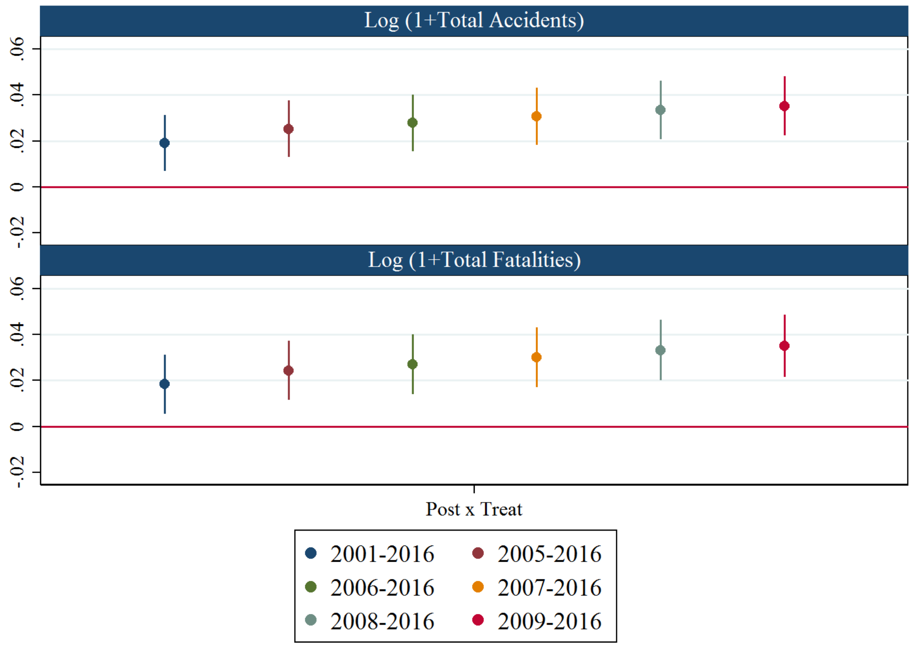

Figure C3. Total Accident and Fatalities Coefficient Estimates in Different Sample Periods

Notes: This Figure illustrates how the regression coefficient estimates of the effect of ridehailing on total accidents and fatalities vary with different sample time periods. The outcome measure for all each panel are displayed at the top of the panel. Two-tailed 95% confidence intervals based on standard errors clustered at the city level are displayed in the figure.

Sample Coverage

Other papers in this vein have significant differences in geographic sample coverage and size relative to this study. For example, Greenwood and Wattal (2017) use a sample period and data in which the definition of alcohol-involved accident is uniform throughout the sample. They look at fatalities in alcohol-related accidents, using the period 2009-2013, but with a much smaller sample than ours—540 townships in California, only 7% of which were treated with UberX service during the sample period. Perhaps unsurprisingly as a result, their models (DD models that do not account for location-specific linear or quadratic trends) demonstrate relatively low R-sq (0.035-0.041). Notably, in our analysis, the positive relationship between RH and accidents are clearly concentrated in larger cities and cities with high ex-ante use of public transport, which many CA cities do not have.

Model Specification

Finally, a number of other studies look at this relationship using the lens of alternative specifications which may be less appropriate for answering the question at hand. Brazil and Kirk (2016) use a negative binomial model rather than a difference-in-differences setup. They find no effect of RH on either accidents or drunk accidents using this specification. Dills and Mulholland (2018) use a difference-in-differences approach, looking at both total and drunk accidents in relationship to the launch of UberX at the county level. Their model estimates find an economically large positive coefficient on TREAT x POST, similar to our findings. However, they instead choose to draw conclusions from the coefficient on a different term included in the model: an interaction of TREAT x POST with a linear time trend (non-location-specific). Here they find a negative coefficient roughly $1/6^{th}$ to $1/10^{th}$ the magnitude of the main effect, which is statistically significant only in some specifications. They then interpret this coefficient as evidence that accidents are reduced post UberX entry, rather than as a reduction in the slope of the increase, which is what the specification would actually imply. Their models exhibit R-sq of 0.03-0.13, in contrast to our R-sq which is typically in the range of 0.4 to 0.6 or higher. Notably, their drunk accident sample also includes at least one year of data pre-definition change.

C.4. Accident Level Analysis on Passenger Seats

In this section we present accident level analysis on the effects of ridehailing on the number and location of passengers in the accidents. Specifically, 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 the table below, we show that the coefficient on POST X TREAT 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 further consistent with the mechanism for our findings being at least partially a quality effect associated with RH drivers (as opposed to non-RH drivers).

Panel A: At Least One Back Row Passenger

	(1)	(2)	(3)
	I (Back Row)	I (Back Row)	I (Back Row)
Post X Treat	0.0056	0.0094**	0.0101**
	(0.0036)	(0.0042)	(0.0051)
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	267920	267920	267920
R^2	0.017	0.028	0.038

Panel B: At Least One Back Row Adult

	(1)	(2)	(3)
	I (Back Row)	I (Back Row)	I (Back Row)
Post X Treat	0.003	0.0094***	0.0117***
	(0.0028)	(0.0036)	(0.0044)
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	267920	267920	267920
R^2	0.015	0.026	0.036

Panel C: At Least One Back Row Passenger, All Children (Placebo)

	(1)	(2)	(3)
	I (Back Row)	I (Back Row)	I (Back Row)
Post X Treat	0.0023	-0.0002	-0.0012
	(0.0025)	(0.0030)	(0.0036)
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	267920	267920	267920
R^2	0.017	0.028	0.038

Table C2. Accident Level Analysis

Notes: This table presents results from generalized difference-in-differences, where the unit of observation is accident. In Panel A, the outcome variable is an indicator for whether an accident involves a vehicle that has at least one passenger in the back row. In Panel B, the outcome variable is an indicator for whether an accident involves a vehicle that has at least one adult passenger (age greater than or equal to 18) in the back row. In Panel C, the outcome variable is an indicator for whether an accident involves a vehicle that has only child passengers in the back row. $Post * Treated_{t,c}$ is a dummy variable that equals one if city c adopted at least one ridehailing service at time t. Control variables include the natural logarithm of population and the natural logarithm of income per capita. Standard errors, adjusted for clustering at the city level, are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

C.5. Ridehailing and Smartphone Adoption

In the table below, we attempt to shed light on the notion that RH may drive smartphone adoption in places it was not prevalent before, and that smartphone adoption in turn (by RH drivers and passengers alike) may lead to higher accidents. We utilize the ex-ante stock of smartphones as proxied for by cumulative cellphone sales in the three years prior to RH entry in 2010. We interact this variable with *POST X TREATED*. In locations where cellphone adoption was higher pre-RH, the increase in accidents in lower, suggesting that in fact, the mechanism behind our findings may be that RH induces people who did not previously have smartphone to purchase smartphones (so that they can drive for the service), and as a result, there is an increase in smartphone-distracted driving; this increase, however, remains attributable to the entry of RH as a service in the market.

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post\ X\ Treated_{t,c}$	0.0215**	0.0387***	0.0478***	0.0185**	0.0381***	0.0479***
	(0.0084)	(0.0097)	(0.0118)	(0.0088)	(0.0103)	(0.0125)
$Post\ X\ Treated_{t,c}$ X Cumulative Cellphone Sales	-0.0063*	-0.0058	-0.0119**	-0.0053	-0.0055	-0.0113**
	(0.0033)	(0.0039)	(0.0050)	(0.0035)	(0.0042)	(0.0053)
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	152,576	152,576	152,576	152,576	152,576	152,576
R2	0.65	0.66	0.66	0.63	0.64	0.65

Table C3. Ridehailing and Smartphone Adoption

Notes: This table shows how the effects of ridehailing on traffic safety varies with ex-ante cellphone adoption. Cumulative cellphone sales are the MSA-level cumulative gross receipts for industry with NAICS 517312: Wireless Telecommunication Carriers (Except Satellite) from 2007-2010. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors are adjusted for city level clustering and are reported in parentheses. ∗∗∗, ∗∗, and ∗ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

C.6. Type of Service

In the table below, we separate out the treatment effect of the different types of services: those that are single rides (UberBlack/taxi/X, Lyft) versus pooled rides (UberPool, LyftLine). We pool UberBlack/taxi with UberX, due to the very small number of cities that have (had) UberBlack/taxi service. We thus report the treatment effect for pooled versus nonpooled service. The estimates in the table suggest that the rollout of pooled ride services does not reverse the overall treatment effect of nonpool RH. The coefficients for pool launch are negative, and roughly half the magnitude of those for single ride (nonpool) RH launch, but are not statistically significant at conventional levels. This is consistent with reported low adoption rates for pooled rides in cities that offer the service.

	Log (1+Total Accidents)			Log (1+Total Fatalities)		
	(1)	(2)	(3)	(4)	(5)	(6)
Single Ride Service (UberBlack/Taxi/X, Lyft)	0.0215***	0.0372***	0.0337***	0.0204***	0.0371***	0.0340***
	(0.0065)	(0.0075)	(0.0091)	(0.0067)	(0.0079)	(0.0096)
Pooled Ride Service (Uber Pool, Lyft Line)	-0.0233	-0.0141	-0.0191	-0.0199	-0.0130	-0.0191
	(0.0145)	(0.0150)	(0.0163)	(0.0153)	(0.0158)	(0.0172)
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

Table C4: Type of RH Service

Notes: This table shows how the effect of ridehailing on traffic safety varies with the type of service. Total accidents is the number of fatal accidents, according to the definition used by NHTSA. Total fatalities is the total number of fatalities across all fatal accidents. Single (Pooled) Ride Service is a dummy variable that takes the value of one if any single (pooled) ride service is adopted. Control variables in all regressions include the natural logarithm of population and the natural logarithm of income per capita. Standard errors are adjusted for city level clustering and are reported in parentheses. $***$, $**$, and $*$ represent statistical significance at the 1%, 5%, and 10% levels, respectively.

D. CONCEPTUAL MODEL

To better understand the expected effects of RH on accident rates, we develop 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) on roads and the average quality of drivers. 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.[36]

For notational purposes, we denote the accident rate for city i in period t as $A_{i,t}$ and the new technology (ridehailing) as θ. Accident rates can then be thought of as: $A_{i,t} = f(VMT_i(\theta); Q_{i,t}(\theta))$, where $VMT_i(\theta)$ is the number of vehicle miles traveled on the road in city i in period t (potentially a function of whether RH is available) and $Q_{i,t}(\theta)$ is the quality of the average driver on the road in city i in period t.

The number of VMTs can further be broken down into three sub-categories: (i) the number of VMTs generated by people driving themselves from origin to destination, denoted by VMT^{own}; (ii) the number of VMTs generated by RH drivers carrying passengers from origin to destination, denoted by VMT^{RH}; and (iii) the number of VMTs generated by RH drivers while driving in-between passengers, denoted by VMT^{btwnRH}. Thus, $VMT_i = VMT^{own} + VMT^{RH} + VMT^{btwnRH}$.

Note that, even if VMT^{own} and VMT^{RH} simply offset as people move from driving themselves to being driven in a RH vehicle, there is still "between driving" (between fares, waiting for fares, going from fare location to fare location) that is introduced by the advent of RH in a city. While VMT^{own} is almost certainly decreased by the introduction of RH, the technology leads to the introduction of additional vehicle miles in the form of VMT^{RH} and VMT^{btwnRH}. Thus, the effect of the introduction of RH in a city on the number of VMTs on the road depends on whether the decrease in VMT^{own} is more than offset by VMT^{RH} and VMT^{btwnRH} that are introduced with the technology. Taking the model naïvely (and ignoring for the moment the UberPool and LyftLine services), each person who no longer chooses to drive himself or herself is now driven by a RH driver, thus precisely offsetting the effect on the overall vehicle miles traveled. But unless there are absolutely no between-fare miles driven by a ride-sharing driver, we would expect to see an increase in overall VMTs after RH arrives.

The limited evidence to date suggests that there is considerable between-fare travel by drivers. 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

[36] Notably, these effects are compounded by the congestion and pollution effects of driving; we leave this topic to future research.

up the fare, and 5.1 miles with a passenger in the car, implying a 59% utilization rate. Furthermore, RH companies initially offered subsidies designed to induce drivers to spend more time out on the road active in the app, so as to decrease wait time for passengers. Finally, while not the focus of their study, the analysis of Chen et al. (2017) is consistent with a mismatch between rider demand and the supply of drivers, particularly given the flexibility afforded to the drivers.

More formally, we can write the first-order condition for the effects on accident rate A_i from the introduction of RH technology θ as:

$$\frac{\partial A_i}{\partial \theta} = \frac{\partial A_i}{\partial VMT_i}\frac{\partial VMT_i}{\partial \theta} + \frac{\partial A_i}{\partial Q_i}\frac{\partial Q_i}{\partial \theta}$$

where

$$\frac{\partial VMT_i}{\partial \theta} = \frac{\partial {VMT^{own}}_i}{\partial \theta} + \frac{\partial {VMT^{RH}}_i}{\partial \theta} + \frac{\partial {VMT^{btwnRH}}_i}{\partial \theta}.$$

Clearly, $\frac{\partial A_i}{\partial VMT_i}$ is positive, as every additional vehicle mile travelled will increase the likelihood of an accident and thus the overall accident rate. $\frac{\partial {VMT^{own}}_i}{\partial \theta}$ is negative. $\frac{\partial {VMT^{RH}}_i}{\partial \theta}$, however, will either equal or, more likely, due to substitution away from other forms of transport, be larger in absolute magnitude than $\frac{\partial {VMT^{own}}_i}{\partial \theta}$, and $\frac{\partial {VMT^{btwnRH}}_i}{\partial \theta}$ is positive. Thus the overall effect $\frac{\partial VMT_i}{\partial \theta}$ is positive: vehicle miles traveled increase with the introduction of RH.

Of course, in some cities, at later dates, the option to "carpool" via RH was introduced, in the form of Uber Pool and Lyft Line. With the introduction of these services, the reduction in own drive car hours may not be fully offset by RH drive hours, as multiple people may be substituting away from driving themselves into a single RH car. While Uber and Lyft have both heavily invested in promoting their shared services, Uber reports that UberPool accounts for only 20% of trips in cities where it is offered, and Lyft reports that 37% of users in cities with LyftLine request a Line trip, and many trips are not matched, thus leaving a single rider (Schaller, 2018). Pooled rides are also cheaper, potentially inducing more substitution from other modes of transport. It is not clear what fraction of rides must be pooled to counteract VMT^{btwnRH}, but Schaller (2018) suggests that, even if half of the rides were pooled, total VMT would still increase. Furthermore, stepping away from the naïve model, survey evidence suggests that $\frac{VMT^{RH}}{VMT^{own}} > 1$, as many riders are substituting away not from driving themselves but rather from other forms of transportation, including walking, biking, and, more importantly, public transportation (Clewlow and Mishra, 2017). Thus, it is likely that pooled ride adoption would need to be extremely high to offset such substitution effects.

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 ridehailing 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. Put another way, the introduction of RH makes it less costly to have someone else drive you, but also makes the gains from getting out on the road as a driver greater (as you can make money by doing so). Lower quality drivers, who in the absence of compensation may not have driven, now have an incentive to drive. Moreover, more affluent people are more likely to use RH (Pew Research Center, 2016), and the less affluent are more likely to become RH drivers. To the extent that this substitution leads to more vehicle miles driven by lower quality drivers or in lower quality cars, this may positively affect accident rates. Finally, ridehailing VMT is also different from self-driving VMT, as ridehailing 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 ridehailing vehicle).[37] Yet RH drivers, especially those with more experience from more hours driven, may in fact represent improved quality. To the extent that the substitution goes the other way and lower quality drivers are substituted by better drivers, this may reduce accident rates if the increase in quality offsets the increase in VMT.

Formally, $\frac{\partial A_i}{\partial Q_i}$ is negative: better drivers reduce accident rates, all else equal. The effect of RH on the quality of the average driver on the road, $\frac{\partial Q_i}{\partial \theta}$, however, is ambiguous. If the quality of the average driver increases, this could offset the quantity effect above. If it decreases or does not change, the quantity effect will prevail. Which effect dominates, of course, is an empirical question.

Many indicators suggest that both total VMT and driver quality 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. A driver with more than 2,500 lifetime trips completed earns 14% more per hour than a driver who has completed fewer than 100 lifetime trips, in part because he learns where and when to drive, which may decrease VMT^{btwnRH}. Similarly, Haggag et al. (2018) 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 behavior, on the part of certain driver populations, that directly or indirectly increases the probability of accidents, such as gaming time-and-distance pay algorithms by taking longer routes, speeding, etc.

[37] 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.