

**AI AT THE WHEEL: THE EFFECTIVENESS OF
ADVANCED DRIVER-ASSISTANCE SYSTEMS**

Vikram Maheshri
University of Houston

Clifford Winston
Brookings Institution

Yidi Wu
Georgetown University ¹

Abstract. Has automakers' use of artificial intelligence in advanced driver-assistance system (ADAS) technologies improved automobile safety? We address this question with a first of its kind trim-level dataset of the universe of registered automobiles and accidents in Texas over a 9-year period. We find that ADAS technologies reduce the risk of a motorist getting in any type of accident by 11 to 14 percent and reduce the risk of a motorist getting in a single vehicle fatal accident by roughly one third. Our finding that ADAS has improved automobile safety is especially important because it provides early evidence of the benefits of vehicle automation in actual travel environments. Hopefully, it will spur greater interest in the development and widespread adoption of fully autonomous vehicles and in the potential benefits of other AI transportation technologies.

January 2026

JEL Codes: O33 Technological Change Consequences; R41 Transportation Safety; L51 Economics of Regulation.

¹ Maheshri: Department of Economics, University of Houston, vmaheshri@uh.edu. Winston: Economic Studies: The Brookings Institution, cliffordmwinston@gmail.com. Wu: Department of Economics, Georgetown University, yw648@georgetown.edu. We are grateful to Wesley Blundell, Michael French, Edward Glaeser, Fred Mannering, Ashley Nunes, Alex Oyler, Marc Scribner, participants in the Spring 2024 NBER Transportation Economics conference, various seminar participants, and a referee for helpful comments. We particularly acknowledge many helpful comments and encouragement provided by Sam Peltzman.

1. Introduction

Since Ford Motor Company mass produced the Model T more than a century ago, the US automobile industry has gradually introduced notable vehicle safety improvements including headlights, automatic windshield wipers, shatterproof glass, improved braking, advances in body structure, collapsible steering columns, and occupant safety devices. Government policies also have sought to improve automobile safety by requiring motorists to have a valid driver's license, prohibiting driving under the influence of alcohol or drugs, setting and enforcing speed limits, and requiring vehicles to satisfy National Highway Traffic Safety Administration (NHTSA) safety standards.²

Beginning in the late 2000s, automakers took an important step forward to improve safety by steadily equipping their vehicles with advanced driver-assistance systems (ADAS) based on artificial intelligence. ADAS consists of a suite of safety features that assist in both the forward dimension (automatic emergency braking and adaptive cruise control), and the lateral dimension (lane departure warning and blind spot collision prevention).³ ADAS is standard for some vehicle makes, models, and trims, can be purchased as an option for other makes, models, and trims, and is unavailable for purchase at this time for the remaining makes, models, and trims.⁴ According to the American Automobile Association, at least one ADAS feature was available in 92.7% of new vehicle models in the United States in 2018.⁵

ADAS distinguishes itself from other automobile safety features because it assists the driver by making its own decisions in response to safety threats in real highway travel settings; for example, it may brake automatically to avoid a collision. Other safety features, such as airbags, enhance safety by reducing the severity of an injury if an accident occurs, but ADAS enhances safety by substituting for a driver's attention to prevent an accident from occurring.

The recent adoption of ADAS in the US motivates our interest in assessing its effectiveness at reducing accident risk. As appropriate for assessing the performance of a new technology, we account for the people who select the technology as well as for how they use it in practice because those choices can reinforce or compromise the intended effects of the technology. In contrast, an

² Government highway expenditures also have been used to improve the safety of the road system.

³ Except for adaptive cruise control, ADAS features engage autonomously because they are often enabled by default.

⁴ In addition to safety technology, a vehicle's trim includes powertrain options, aesthetic features, and comfort amenities

⁵ <https://www.electronicdesign.com/markets/automotive/article/21126132/how-technology-is-driving-the-democratization-of-adas>

engineering analysis, which is the basis for current estimates of the effectiveness of ADAS, does not account for the potential bias of motorists who self-select into vehicles equipped with ADAS and for motorists' potential behavioral response to drive more aggressively in ADAS equipped vehicles.

To date, ADAS has been voluntarily installed in certain vehicles by automakers and selected by motorists through their choice of vehicle and trim. But recently, the federal government issued one of the most significant changes to car safety standards in years by requiring that all new passenger cars and light trucks be equipped with automatic emergency braking (AEB) systems, an important component of ADAS. Automakers have until the fall of 2029 to ensure that the AEB systems on their 2030 vehicles comply with federal safety standards.⁶ Thus, our assessment is further motivated by policymakers' mandating the adoption of a component of ADAS without systematic evidence of its effectiveness in improving safety in real world situations.

As noted, the availability of ADAS varies at the model year-make-model-trim level. Thus, to execute our analysis, we aggregate data on accidents and driving to this level, which allows us to assess the effectiveness of ADAS by comparing the safety performance of extremely similar *vehicles* with and without ADAS. To the best of our knowledge, our paper is the first to use modern data collection methods to extract the detailed information that is necessary to conduct a trim-level analysis of automobile safety.

Generally, our approach stands in contrast to the vast empirical safety literature that conducts analyses at the incident level in an attempt to identify the determinants of automobile accidents (for example, Haghani and Bliemer (2023) and Anderson and Auffhammer (2014)). By doing so, that literature is subject to selectivity bias because precisely who chooses to drive what vehicles at what specific times of day under which specific personal and driving conditions is unlikely to be random. Circumventing this bias requires, at a minimum, researchers to identify only treatment effects on accident severity conditional upon an accident occurring (Maheshri and Winston, 2025). However, a conditional analysis cannot account for accidents that have been

⁶ Under the rule, all new vehicles would be required to have a version of automatic emergency braking that is "much more effective at much higher speeds." Specifically, all cars would need to be able to stop and avoid contact with a vehicle in front of them when traveling up to 62mph; vehicles traveling as fast as 45mph would need to come to a complete stop to avoid hitting pedestrians; and braking systems would be required to detect pedestrians and cyclists at night. https://www.nhtsa.gov/sites/nhtsa.gov/files/2024-04/final-rule-automatic-emergency-braking-systems-light-vehicles_web-version.pdf Currently, no commercially available automatic emergency braking technologies satisfy these stringent technical requirements.

prevented.⁷ In contrast, our approach enables us to identify the unconditional effectiveness of ADAS availability with an identification strategy that can address the various selection issues that arise. To implement this approach, we require information on the universe of vehicles on the road in a given geographical area, including those vehicles that have not been involved in accidents.

We fulfill this information requirement by constructing a panel dataset comprised of all registered vehicles in Texas from 2010 to 2018. We link the dataset to a record of all accidents for which a police report was filed in Texas during this period, enabling us to construct the accident history of the universe of registered vehicles.⁸ Although we include the universe of vehicles registered in Texas, some of which get into accidents (and most of which do not), we exclude vehicles not registered in Texas in our initial analysis because we do not observe out of state vehicles that are not involved in accidents.⁹

Finally, we combine the vehicle accident histories with a panel dataset that we construct that identifies the availability of ADAS-related safety features on each trim of every vehicle that was registered during the sample period. The two data sources are merged via a specialized matching procedure using machine learning techniques that decodes the precise trim level of a vehicle from its Vehicle Identification Number (VIN). In sum, our data set consists of the universe of vehicles in Texas, including vehicles that have and have not been involved in accidents and vehicles that are equipped and not equipped with ADAS-related safety features. Our findings should be reasonably representative of the US as a whole because Texas is a large state with a diverse population and fleet of vehicles, which is driven in urban and rural environments that vary in geography, weather, and traffic density.

⁷ Some analyses of auto safety (for example, Edlin and Karaca-Mandic (2006)) aggregate data by driver type to estimate the determinants of drivers who are involved in fatal and non-fatal accidents. However, because we are concerned with the effectiveness of a vehicle attribute, we aggregate our data by vehicle type.

⁸ Texas takes all automobile accidents seriously. Even in a minor accident with no injuries, drivers who leave the scene of the accident without calling the police could be charged subsequently with a misdemeanor. State law mandates that drivers involved in an accident causing injury, death, or property damage exceeding \$1,000 must report the incident to law enforcement. This said, it is still possible for non-fatal accidents to be under-reported in the Texas police data, which may conceivably affect our findings on the effect of ADAS on non-fatal accidents but have no effect on our findings on the effect of ADAS on fatal accidents. We are not aware of any external evidence that documents the extent of under-reporting of non-fatal accidents in Texas and that assesses the extent of any bias associated with this under-reporting.

⁹ As part of our sensitivity analyses, we explore later how our results would be affected if we included both in-state and out of state vehicles in the accident data. As expected, our main findings are not affected because out of state vehicles account for a very small share of accidents in Texas.

Given this data, we meet the fundamental challenge of identifying the causal effect of a new technology on automobile safety—the adoption of the technology is generally non-random—by exploiting plausibly exogenous variation in the availability of ADAS on different vehicles over time. As the new technology (treatment) varies by vehicle type, we compare the aggregate safety performance of vehicles with and without ADAS, as opposed to comparing the disaggregate safety performance of individual drivers. Using the latter approach, the effect of ADAS would be identified only under the dubious assumption that a driver’s propensity to purchase an ADAS-equipped vehicle was uncorrelated to their attitudes toward safety and their driving abilities (perhaps conditional on some small set of observable driver characteristics).

In our approach, we leverage the fact that ADAS became available at different times for different trim levels—notably within vehicles of the same make and model.¹⁰ We therefore identify the causal effect of ADAS on accidents under the weaker assumption that drivers did not systematically opt for higher trim level vehicles solely because of the availability of ADAS. Of course, drivers of higher trim vehicles are likely to differ from drivers of lower trim vehicles in some respects. However, vehicles of different trim levels vary in multiple dimensions by offering dozens of appealing features, many of which are related to comfort and aesthetics and not to safety. This fact lends credence to our identifying assumption, which relies on a combination of the choice of higher trim versus lower trim *and* the timing of ADAS availability.

We also recognize that the effects of self-selection in influencing the effectiveness of ADAS may be reflected in drivers’ risk preferences that are manifested in several ways. That is, those preferences may affect when drivers decide to adopt a vehicle with ADAS; whether and how they drive in different highway conditions; and whether and how they drive in different types of vehicles.

Fortunately, data are available that enable us to test directly for systematic patterns related to drivers’ risk preferences that suggest whether they self-select into vehicles equipped with ADAS technologies. We assembled a large sample of Texas households who owned vehicles, and we obtained from Acxiom, a database marketing company, many of those households’ socioeconomic characteristics. We also collected data on the households’ vehicle safety records. Using these data, we provide evidence that the type of motorists who purchase vehicles with higher trim is not

¹⁰ Wählberg and Dorn (2023) assess the effectiveness of vehicle electronic stability control (ESC) on fatal crash rates, but they do not compare cars’ safety performance with and without ESC.

systematically affected by the availability of ADAS in higher trim vehicles. We also provide evidence that the evolution of the crash rate of drivers who never purchased ADAS equipped vehicles is similar to the evolution of the crash rate of drivers who eventually purchase ADAS equipped vehicles. Finally, we explore heterogeneity in the effects of ADAS across a wide range of vehicle characteristics, such as price and size, and highway travel conditions, such as clear weather, which are known to be positively correlated with the purchase and driving behavior of safer drivers. We find that the effects of ADAS on accident outcomes do not vary systematically across vehicle characteristics and travel conditions.

Importantly, we find that ADAS is highly effective at improving automobile safety even after accounting for drivers' behavioral responses to its availability and installation. Specifically, ADAS technologies reduce the risk of a motorist getting in any type of accident by 11 to 14 percent and reduce the risk of a motorist getting in a single vehicle fatal accident by roughly one third. ADAS has a small and statistically imprecise effect on reducing the risk of a motorist getting in a multivehicle fatal accident, but we suggest that ADAS is likely to reduce the fatality risk of those types of accidents as a greater share of the nation's vehicle fleet is equipped with autonomous vehicle safety features.

In this respect, ADAS represents a significant advance in automobile safety by substituting effectively for a driver's attention and by providing credible evidence that fully autonomous vehicles will notably improve highway transportation safety. From a policy perspective, our evidence of some of the benefits of vehicle automation in actual travel environments should be interpreted as an essential component of a comprehensive cost-benefit assessment of whether automakers should be required to install ADAS in all their vehicles.

2. Estimating the Efficacy of ADAS

The staggered rollout of the availability of ADAS over time and across different automobile makes, models and trims generates temporal and cross-sectional variation in registered vehicles' safety features that enables us to identify the causal effect of ADAS on accident risk. We construct aggregate versions of the key variables, ADAS availability and accidents, to execute the empirical analysis, and then explain our specification to estimate the effect of ADAS availability on accident risk.

ADAS Availability

In most safety analyses, a vehicle type, which we index by i , is defined as a combination of make and model. However, within a make-model combination in our analysis, some vehicles (e.g., luxury editions) may have ADAS and others (e.g., standard editions) may not. We therefore expand the definition of vehicle type as a combination of make, model and trim, where trim levels, defined in the data section, are indexed separately by j .

Crucial to our analysis is that the availability of ADAS for a given vehicle make and model may vary over time because it is not available in earlier model years of some vehicles, but it is available in later model years. Moreover, some vehicle makes and models may never have ADAS available during the sample period. Let y index the model year of a given vehicle type. Then, our treatment variable, the availability of ADAS, which we denote by the dummy variable S_{yij} , varies at the model year y , make-model i , and trim level j .

Accidents

For each vehicle in each calendar year of our sample, we observe the vehicle's model year, type (make-model), trim level, whether it was involved in an accident, and if so, the accident severity (ranging from property damage only to a fatal accident). We denote by t the calendar year, which will generally differ from the model year, of a specific year in a vehicle's accident history. Because we are interested in the effect of a treatment that occurs at the vehicle level, we aggregate accident outcomes to the model year-type-trim-calendar year level and denote by A_{yijt} the total number of accidents of a given severity that vehicles yij had in year t .

The temporal and cross-section variation in our panel is distinctive because it contains two different temporal dimensions: a vehicle's model year y and calendar year t . Although the accident outcome varies over the calendar year dimension, the treatment varies only over the model year dimension y —older models of a vehicle type that were untreated remain untreated even if newer models of that type are treated. Hence, a different treatment variable may be observed at a given $t \geq y$. We exploit the variation in the treatment variable within vehicle type and across trim, model years and calendar years to identify the causal effect of ADAS on accidents.

In table 1, we illustrate the organization of our data for a single vehicle type, the Acura MDX, using the calendar year as the primary temporal dimension for the 2000 to 2019 sample

period.¹¹ The Acura has three trim levels that we denote as Low (*L*), Medium (*M*), and High (*H*), each associated with the period, if any, that they were equipped with ADAS.¹² Vehicles with a low trim level were never equipped with ADAS during our sample period; vehicles with a medium trim level were equipped with ADAS in model year 2018 but not before that calendar year; and vehicles with a high trim level were equipped with ADAS in 2015 but not before that calendar year. The three different trim levels of Acura MDX's on the road during our sample period enable us to define the treated vehicles as Acura MDX's with high and/or medium trim levels that include ADAS. Our untreated or control vehicles are Acura MDX's that do not include ADAS. We define all of the other treated vehicle types and control vehicles in our sample in the same way.

Specification

Previous safety research (e.g., Maheshri and Winston (2024)) has specified vehicle accidents A_{yijt} in a Poisson regression framework because accidents take on small, discrete, non-negative values (Cameron and Trivedi (1998)). Although some empirical analyses in the transportation literature estimate accident equations using a negative binomial model, this is inadvisable because, as pointed out by Wooldridge (1999), the negative binomial estimator is a non-robust estimator of conditional mean parameters, and this weakness is exacerbated when using fixed effects.¹³ Negative binomial regressions may be appropriate if the objective is simply to maximize the fit of the model and there is overdispersion in the dependent variable. However, we are interested in a mean causal effect; thus, it is appropriate to use the Poisson QMLE estimator discussed by Gourierieux et. al. (1984) because it yields consistent estimates of the effect of interest without the distractions of variance assumptions.

¹¹ Note the model year for vehicles manufactured up to June 2018 will be 2018, but the model year for any of the vehicles in our sample manufactured from July through December in each year (for example, 2015) can be advertised as the next model year. Hence even though our sample period corresponds to 2010-2018, it includes some model year 2019 vehicles.

¹² Manufacturers distinguish trims by a large number of features and frequently change the names of different trims for marketing purposes. For example, the 2018 Acura MDX was offered in 5 trim configurations that were marketed in three trim levels: Standard, Technology and Advance. Meanwhile, the 2015 Acura MDX was marketed in 4 trim levels: Base, Advance/Entertainment, Tech, and Tech/Entertainment. Because we are analyzing the effects of ADAS on safety and we wish to maintain a consistent treatment of makes and models over time, we aggregated all vehicles of a given make and model that introduced ADAS in the same model year as a single trim (in this case, Low, Medium or High).

¹³ As Gourierieux et. al. (1984) show, the negative binomial estimator requires both mean and variance to be correctly specified, whereas the Poisson estimator only requires the mean to be correctly specified. Moreover, as Wooldridge (1999) has noted, the negative binomial estimator only “pretends” to solve the problem of over- or under-dispersion of data, but this is only true if the variance is correctly specified; otherwise, misspecification bias will be inherited. Perhaps most critically for our application, negative binomial estimators suffer from the incidental parameters problem and also are scale dependent.

We therefore specify our models of accidents and fatal accidents as:

$$A_{yijt} = \exp(\beta S_{yijt} + \lambda_{ijt} + \lambda_{iyt} + \epsilon_{yijt}), \quad (1)$$

where S_{yijt} is a dummy variable equal to one if ADAS was available either as standard equipment or purchased through an optional package on vehicle yij in year t and zero otherwise; λ_{ijt} are make-model-trim-calendar year fixed effects; λ_{iyt} are make-model-model year-calendar year fixed effects; and ϵ_{yijt} is an error term.¹⁴

The key identifying assumption that enables parameter β to be interpreted as the causal effect of the availability of ADAS on selected vehicles on the total number of accidents is that $cov(S_{yijt}, \epsilon_{yijt} | \lambda_{ijt}, \lambda_{iyt}) = 0$. That is, motorists who purchase higher trim vehicles during the first model year that ADAS is made available in those vehicles are not systematically different from the motorists who purchase higher trim vehicles of other model years. We provide several pieces of empirical evidence below to support this assumption.

Based on the early experiences with autonomous vehicles in controlled testing environments (Blanco, et. al., 2016, Mosquet, Andersen, and Arora, 2016), we expect the availability of ADAS to reduce accidents. But as discussed in Maheshri and Winston (2025), the findings in controlled testing environments should be viewed with caution because they are likely to be biased by not being based on a random sample of drivers.

3. Data

We constructed a data set consisting of all the registered vehicles in Texas from 2010 to 2018 along with their trim, which we used to identify whether a given vehicle is equipped with ADAS. We used leading vehicle data aggregators that describe the available safety features in all new vehicle trims, including ADAS, to identify the vehicles equipped with ADAS during the sample period.¹⁵ Then, for each vehicle we merged information from the universe of Texas police

¹⁴ Data on specific vehicles that were purchased with ADAS as an optional package are not available. However, when a vehicle, defined by make and model, offers ADAS features as an option instead of as standard, most consumers who select that vehicle also are likely to purchase the optional ADAS features. The reason is that the entire trim package of a vehicle that offers optional ADAS features is usually more expensive than the entire trim package of the same or similar vehicle that does not offer ADAS as an option. Thus, consumers who do not want the optional ADAS features would, in all likelihood, decide to reduce their costs by simply choosing a similar vehicle that does not contain a trim package that gives them the opportunity to purchase ADAS as an option. Anecdotal evidence obtained from car dealers was consistent with this characterization of consumer behavior.

¹⁵ Vehicle data aggregators use automotive data aggregation platforms, which are centralized systems designed to collect, organize, and process data generated by vehicles within the automotive ecosystem.

accident reports to construct its detailed accident history. To the best of our knowledge, this is the first data set at the vehicle trim level that has been used to assess the efficacy of vehicle safety features.

Vehicle Identification Numbers. Extracting safety features from Vehicle Identification Numbers (VINs) and then aggregating vehicles to the trim level is a formidable task. It requires matching unstructured descriptive data from auto manufacturers to vehicle trims that are only partially identified by their VIN and then classifying vehicles into meaningful groupings of trims in a consistent manner. We decoded the VIN of every vehicle in our sample using a commercially available VIN decoder. The decoder identified each vehicle down to the trim level, which is critical for our analysis because different versions of the same vehicle make and model have different safety features.¹⁶ We then collected detailed information from data aggregators, such as TrueCar and MotorTrend, by scraping their websites and employing string manipulation techniques to verify the availability of ADAS for each vehicle. We provide a detailed description of how we extracted safety features from Vehicle Identification Numbers (VINs) and aggregated vehicles to the trim level in Appendix A. In all, we constructed a panel of annual and fatal accidents from 2010 to 2018 for 6,268 distinct vehicle types defined as a unique model year-make-model-trim combination.

In Figure 1, we present the evolution of ADAS availability for vehicles in our sample during calendar years 2010-2018. Panel A shows that vehicle trims equipped with ADAS safety technology were relatively scarce for most of our sample period, though they gradually became more common after 2015. Indeed, Panel B shows that while the number of vehicles equipped with ADAS safety technology represents a small (less than 20%) share of all vehicles throughout our sample period, newly registered vehicles are increasingly more likely to be equipped with ADAS towards the last few years of our sample period. This pattern reflects the fact that vehicles are infrequently purchased durable goods; hence, there is a considerable delay between the availability of a new safety innovation and its adoption by motorists.

Vehicle Miles Travelled. The introduction of ADAS could affect driving intensity, as measured by vehicle miles travelled (VMT), which could confound our findings on the safety effects of ADAS. To explore this possibility, we constructed a subsample from our main sample of

¹⁶ Using the example in table 1, the Acura MDX high level trim is called the Type S Advance, which made ADAS available in model year 2015. The low level is the base trim, which has not made ADAS available.

vehicles that contained the VMT for each vehicle in the subsample that was obtained from the Texas Commission on Environmental Quality (TCEQ). The 14 largest counties in Texas require each vehicle to be subjected to emissions testing annually prior to being registered. Among other information, TCEQ collects the VINs of the vehicles and the exact annual odometer readings for each registered vehicle in 14 counties, from which we are able to construct average annual measures of VMT for each vehicle type.

Self-Selection. Finally, we collected data to explore whether the availability of ADAS safety features led consumers to self-select systematically into ADAS-equipped and non-equipped trims. We obtained data from Acxiom for more than 200,000 randomly selected registered vehicle owners in Texas from our main sample containing information about their race, income, marital status, household size, and propensity to adopt new technologies. We use the data to construct the average demographic characteristics of owners of each vehicle type in our sample. We then estimate whether those characteristics when interacted with the availability of ADAS features have distinct effects on accident outcomes, which could be interpreted as suggesting that consumers are self-selecting into ADAS-equipped and non-equipped trims based on observed influences.

In figure 2, we present a flow chart to summarize the sequence of the data collection and construction process, the variables collected, and the key summary statistics for the variables. We use the final data set for our estimations. As noted, Appendix A provides a detailed description of the process of linking the data from the vehicle registrations, police accident reports, and trim level attributes to identify whether ADAS safety technology was available on each vehicle registered in Texas.

4. Estimation Results

Table 2 presents the effects of the availability of ADAS on all accidents and on fatal accidents. It is standard to facilitate interpretation of Poisson regression estimates by presenting them as incidence risk ratios (IRRs) because the nonlinear mathematical structure of the model specifies the dependent variable as the log of the expected number of events for observation i . An IRR has a clear qualitative and quantitative interpretation. An IRR greater than 1 corresponds to a positive effect on vehicle accidents and an IRR of 1.25 means the expected count increases by 25%. An IRR less than 1 corresponds to a negative effect on vehicle accidents and an IRR of 0.80

means the expected count decreases by 20%. We also present marginal effects in table 2 as if we were presenting standard regressions results, but we stress that the IRRs are easier to interpret.

We did not specify accident and fatality rates per vehicle mile of travel because the adoption of ADAS is likely to simultaneously influence VMT as well as accidents and fatalities. Even if one of those influences were small, it would still prevent us from determining the distinct effects of ADAS on accidents and fatalities. However, we directly estimate the effect of the adoption of ADAS on VMT in order to provide further context for our results.

In each regression, we restrict our sample to model year-make-model-trim combinations with at least ten registered vehicles to ensure that the results are not affected by rare vehicles, such as Ferraris.¹⁷ We also restricted the sample to vehicle types that are equipped with ADAS at some point during the sample period to ensure that the results are not affected by variation among never treated vehicles; when we relaxed this assumption and included never treated vehicle types, our findings were unchanged.

Basic Findings

Based on the estimated IRRs in table 2, the availability of ADAS reduces the total number of accidents of a given vehicle and trim type by 14% in single vehicle accidents (column 1) and by 11% in multivehicle accidents (column 2), and the effects are statistically significant. The associated marginal effects also are statistically significant and can be shown to be quantitatively consistent with the estimates of the IRRs.¹⁸

Bear in mind that only a small share of the entire vehicle fleet is equipped with autonomous vehicle safety features. It therefore plausible that ADAS is less effective in reducing multivehicle compared with reducing single vehicle accidents because the other vehicle involved in a multivehicle accident may not be equipped with autonomous vehicle safety features. As more of

¹⁷ If we eliminate the assumption, our standard errors increase but the coefficients obtained with and without the assumption are not statistically indistinguishable

¹⁸ Wooldridge (2023) proposes a method to estimate treatment heterogeneity by using a robust two way fixed effects estimator for a Poisson regression. We used his estimator here as a robustness test of our findings and obtained slightly larger but less precise average treatment effects of ADAS. The loss in statistical precision arises because the estimator proposed by Wooldridge is less efficient than the simpler estimator that we used here. In any case, our quantitative estimates of the effects of ADAS on automobile safety appear to be consistent with those obtained by using the more sophisticated estimator.

the nation's vehicle fleet is equipped with autonomous vehicle safety features, ADAS may become equally effective in reducing single and multivehicle accidents.¹⁹

ADAS technologies are even more effective at reducing single vehicle fatal accidents than they are at reducing non-fatal accidents. We find that the total number of fatal accidents involving a given vehicle and trim type is reduced by roughly one third and the effect is statistically significant (column 3). For example, a lane departure warning could wake up a drowsy driver who is on the road by herself and prevent a fatal accident from occurring because the driver's vehicle crashed into a retaining wall or went over an embankment at high speed. Thus, the lane departure warning effectively substitutes for a driver's attention to prevent a fatal single vehicle accident.

In contrast, we find that ADAS technologies have a small, nearly zero effect, on reducing multivehicle fatal accidents (column 4). But the standard error is large, and a 95% confidence interval indicates that ADAS could reduce auto fatalities in multivehicle accidents by as much as 22% or could increase auto fatalities in multivehicle accidents by as much as 28%; thus, there is considerable uncertainty about the effect of ADAS in multivehicle fatal accidents.

Notwithstanding this uncertainty, the upper bound of ADAS's effectiveness in reducing multivehicle fatal accidents is notably less than ADAS's mean effect in reducing single vehicle fatalities. This is plausible because ADAS's features are effectively designed to be more effective in reducing single vehicle versus multivehicle fatal accidents. That is, they can prevent a vehicle equipped with ADAS features from, for example, crashing into another vehicle, which could result in a fatality. However, they are much less able to prevent a vehicle equipped with ADAS features from being hit by another vehicle, which also could result in a fatality, especially in the likely event that the vehicle is not equipped with ADAS features during our sample period.²⁰

¹⁹ The introduction of more advanced autonomous driving technologies that allow for communication between vehicles may generate greater safety benefits for multivehicle accidents, but such technologies are not likely to be developed in the near future.

²⁰ Importantly, we acknowledge that it is not clear why the estimate of ADAS's effect on multivehicle fatal accidents is less precise than the estimate of ADAS's effect on multivehicle non-fatal accidents and the estimate of its effect on single vehicle fatalities. We explored whether there were vehicle correlations for multivehicle accidents that result in fatalities that are distinct from vehicle correlations for other types of accidents. For example, multivehicle fatal accidents may often involve heavy cars, so single vehicle fatal accidents and multivehicle accidents have more variation in the data and smaller standard errors. However, we did not find evidence that multivehicle fatal accident patterns are systematically different from other types of accidents. A more reliable estimate of ADAS's effect on reducing fatalities in multivehicle accidents is likely to be possible in future research as the adoption of ADAS-equipped vehicles increases.

Finally, as a check that our results do not simply reflect changes in driving intensity, we present the effects in Table 3 of ADAS availability on vehicle miles travelled. We are unable to obtain statistically significant effects but based on the estimates in specification (2), we can rule out that the availability of ADAS technology will not *reduce* VMT by more than 4%, which implies that changes in VMT cannot explain the sizeable effects that ADAS availability has on both all and fatal accidents.²¹

Reconciling Our Findings with Fatalities During the Sample Period

To reconcile our findings with fatalities during the sample period, it is useful to compare our estimated effects of ADAS adoption with external estimates of the effects of other important influences on auto fatalities that are available in the automobile safety literature. The estimates summarized in Appendix B table 1 indicate that, all else constant, the observed pattern of ADAS adoption would have resulted in a 1.89% *decrease* in auto fatalities in Texas from 2010 to 2018. In comparison, *ceteris paribus*, the observed increase in vehicle weight would have resulted in a 6.39% increase in fatalities, and the observed decrease in alcohol use while driving (as proxied by DUI arrests) would have resulted in a 6.55% decrease in auto fatalities. Hence, the modest changes in ADAS availability during the sample period were roughly one third as important as changes in vehicle weight and driving behavior in explaining auto fatalities.²²

The observed increase in auto fatalities in Texas of 19.5% over the sample period appears to be largely attributable to a 20.39% increase in total vehicle miles travelled. Of course, other less important effects on highway safety are likely to be changing during this period as well. In sum, our findings imply that motorists' adoption of ADAS measurably reduced auto fatalities over the sample period. Continuing adoption of these technologies has the potential to improve safety even more significantly because ADAS was not available in roughly 85% of the registered vehicles at the end of our sample period.

Heterogeneity

In Figure 3, we show the extent that the effects of ADAS on accidents vary in any systematic way with vehicle characteristics, which could have implications for consumers'

²¹ For sensitivity purposes, we also explored whether any of the preceding findings were affected when we included both in-state and out of state vehicles in the accident data. As we expected, none of our findings were materially affected because based on our data on total accidents in Texas, out of state vehicles account for only about 5% of all the vehicles.

²² It is potentially misleading to add the implied changes together to estimate their joint contribution because each change is based on a different counterfactual.

selective purchase behavior of vehicles when ADAS is available. Generally, we find little heterogeneity in the effects of ADAS in reducing accidents based on vehicle size, price, and manufacturer nationality, with the puzzling exception that its availability appears to be more effective in lighter than heavier vehicles.²³ Because safer drivers tend to purchase larger, heavier, and more expensive vehicles, those findings are broadly consistent with our maintained assumption that safer drivers do not systematically switch into safer vehicles. We find no systematic heterogeneity in the effects of ADAS technologies on fatal accidents.²⁴

The effectiveness of ADAS on accidents also may vary by driving behavior and conditions when an accident occurs. These findings also could reflect selectivity to the extent that more risky drivers tend to purchase vehicles with ADAS features because they are more likely than less risky drivers to drive more dangerously and to drive in more dangerous conditions. ADAS could therefore possibly offset risky drivers' choices of how, when, and where to drive. Figure 4 shows, however, that there is no evidence of heterogeneity in the effects of ADAS on all accidents by speed of crash, roadway type, roadway conditions, weather conditions, and day or week or time of day. If anything, ADAS appears to compress drivers' risk profiles by sufficiently offsetting drivers' choices that may increase accident risk. Although we find some heterogeneity for accident conditions in our point estimates for fatal accidents, those differences are not statistically significant.

5. Potential Sources of Bias That Could Affect the Interpretation of Our Findings

We have assessed important potential sources of bias to our estimates of the effects of ADAS technologies on motorists' safety. We now explore in depth the three primary potential sources of bias to our estimates that could affect the *interpretation* of our finding that ADAS has significant effects on reducing all accidents and fatal accidents: selection bias, offsetting behavior, and contamination of the control group.

²³ In this initial analysis, it is inappropriate to strongly speculate about possible explanations for this finding, especially because we find that the difference between the effects that we estimate for cars and trucks is not statistically significant.

²⁴ We are unable to estimate precise heterogeneous effects on ADAS on fatalities for many of the vehicle type/weight/MSRP/automaker categories, in all likelihood because of the infrequency of fatal accidents.

Selection Bias

We have stressed that drivers' decisions to self-select into treatment—that is, drive a vehicle with ADAS safety features—is the main source of bias in our analysis because it would indicate that instead of being random, drivers' adoption decisions may be strongly correlated with their safety preferences and behavior. If, for example, safer drivers were systematically more likely to adopt ADAS than were riskier drivers, then our estimates of the effects of ADAS on automobile safety would be biased upwards. Conversely, our estimates of the effects of ADAS on automobile safety would be biased downward if riskier drivers were more likely to adopt ADAS than safer drivers. The latter behavior would be more relevant in the case of a safety feature like ADAS that can compensate for a driver's riskiness, instead of a safety feature like airbags that does not compensate for a driver's riskiness but engages *after* a vehicle is involved in a collision.

Either of these issues would constitute a failure of the parallel trends assumption underlying our identification strategy. We clarify how this is less of a concern here by respecifying our empirical model of accidents given in equation (1) as:

$$A_{yijt} = \exp\left(\sum_{\tau=-3,-1,1,3,4} \beta^{\tau} \times 1(\tilde{y}_{ij} - y = \tau - 1) + \text{controls} + \lambda_{ijt} + \lambda_{iyt} + \epsilon_{yijt}\right) \quad (2)$$

where \tilde{y}_{ij} denotes the model year in which vehicle ij is first equipped with ADAS and $1(\cdot)$ represents the indicator function. The coefficients β^{τ} correspond to the effect of ADAS in the τ vehicles equipped with ADAS. Finally, we include $1(\tilde{y}_{ij} - y < -3)$ and $1(\tilde{y}_{ij} - y > 4)$ as controls to normalize all effects relative to the model year just prior to treatment (e.g., 2014 for the high trim Acura MDX available in 2015). As before, we estimate the model using the sample of vehicles that were equipped with ADAS at some point during the sample period; we expect the IRR associated with β^{τ} for $\tau < 0$ to be equal to 1 if our estimates did not suffer from self-selection, i.e., there should be no treatment effect in model years prior to treatment.

We present regression results in an event-study style plot in figure 5 showing that the effectiveness of ADAS at reducing accidents and especially at reducing fatalities increases with the model years that ADAS technologies are available on higher trims.²⁵ There are two potential

²⁵ We refer to the plot as “event-study style” because our data is organized along two time dimensions, calendar year and model year. Accordingly, a given make-model-trim vehicle will contribute different numbers of observations to the estimation of each effect shown in Figure 2. For instance, the 2014 Acura MDX contributes 5 observations to the estimation of the point with -1 model years because ADAS was available in the higher trim calendar years 2014-2018, but the 2016 Acura MDX contributes only 3 observations to the estimation of the point with +1 model years because ADAS was available in higher trim calendar years 2016-2018.

explanations for this pattern: (1) ADAS reduces the prevalence of accidents by an amount that is quantitatively consistent with the parameter estimates in column 2 of table 2, which are based on our original specification in equation (1), or (2) drivers systematically switch to ADAS equipped trims only when they are made available, and they avoid higher level trims in earlier model years when ADAS was not available.

We reject the second explanation because higher trim vehicles differ from their lower trim counterparts in a variety of important dimensions, not just in the availability of ADAS. Those dimensions include non-ADAS vehicle safety features, such as side curtain and seat mounted side impact airbags, as well as non-safety features, such as a premium leather collection. In Appendix C, we report a complete list of the 15 non-ADAS and non-safety related trim features that were available for vehicles with high trim, but not for vehicles with low trim. The fact that trim choice is influenced by more than just the availability of ADAS lends credence to our first explanation that the pattern of results is credibly aligned with the estimates of the effect of ADAS on all accidents and fatal accidents.

We find that the effects of ADAS are smaller in the first model year when automakers introduce ADAS than after subsequent periods because over time automakers may improve the integration of ADAS technologies in a vehicle's safety performance or drivers might use ADAS features more effectively or both. We cannot offer a more precise explanation because the estimates we obtain for each year are obtained using different samples of drivers. We leave it to future research to either challenge or provide a more precise explanation for this finding.

We also provide direct evidence against the claim that the findings are influenced by systematic self-selection of safer drivers into ADAS equipped vehicles by estimating the effect of ADAS availability on the demographic characteristics of adopting households. We accomplish this by replacing the dependent variable in our main specification (equation (1)) with a new dependent variable to obtain:

$$X_{yij} = \exp(\beta_X S_{yij} + \lambda_{ijt} + \lambda_{iyt} + \epsilon_{yij}), \quad (3)$$

where X_{yij} is a demographic characteristic of owners of vehicle yij .²⁶ Recall that we obtained this variable from a survey by Acxiom, which was conducted for a given calendar year and did not

²⁶ We continue to estimate the specification using Poisson regression because the demographic variables are non-negative with a small number of values. Using Poisson regression in such instances has become common practice.

vary by calendar year. The parameter β_X represents the effect of ADAS availability on an average demographic characteristic of owners. If, for example, drivers differentially sort into ADAS enabled vehicles *because ADAS is newly available*, we would expect to find that our estimate of β_X would be statistically significantly different from zero because the adoption of ADAS was associated with drivers that had certain demographic characteristics.

We present estimates of β_X for a variety of household demographics in figure 6, which again suggests that drivers do not systematically select into ADAS equipped vehicles.²⁷ Drivers who switch to ADAS equipped trims immediately after the availability of ADAS are slightly more likely to be male, educated, and married, but they are of similar age, from similar sized households, and earn similar incomes. Not surprisingly, these drivers have a higher propensity to own more vehicles and to adopt new technologies. To the extent that there are statistically significant differences between drivers who opt into higher trims when ADAS becomes available, the effects for virtually any characteristic are very small (less than a 5% change).²⁸ In sum, the evidence bolsters our claim that the findings of ADAS's efficacy in improving safety cannot be explained by selection.

Finally, Figure 7 provides additional circumstantial evidence against the presence of selection bias by showing that over time the safest drivers did not disproportionately switch to vehicles equipped with ADAS when those safety features were first made available. If this were the case, we would expect the earliest adopters of ADAS equipped vehicles to have fewer accidents (pre-adoption) than later adopters of ADAS vehicles (pre-adoption). However, the pre-adoption trends of the crash rates for all groups of drivers are roughly parallel and at similar levels. As we would expect, given the effectiveness of ADAS, the crash rate of drivers who never switched into ADAS equipped vehicles is somewhat higher than the crash rate of drivers who switched into ADAS equipped vehicles at some point during our sample period.

²⁷ We also attempted to construct a demographic “riskiness index” by predicting accident rates using all of the demographic variables that were available. We did not find that this index changed systematically when ADAS became available, but the fit of the prediction was poor. Independently, we did not find any increase in the share of higher trim level vehicles within a model year-make-model when ADAS was made available.

²⁸ The only exception is that Black drivers adopt ADAS equipped trims when they are first made available at close to a 20% change. But the fact that their response is not accompanied by any other demographic shift among all drivers strongly suggests that it is not correlated with the safety preferences of Black drivers only.

Offsetting Behavior

A second potential source of bias to our estimates is that the adoption of ADAS might affect a driver's *behavior* on the road. For example, a driver with ADAS might take more risks while driving, like texting and paying less attention to traffic conditions, which would offset the safety benefits of ADAS.²⁹ Alternatively, because ADAS features include auditory and visual warnings to drivers when other vehicles are approaching, ADAS may induce drivers to make a safety augmenting response. In any case, given that our interest is to estimate the effect of ADAS on automobile safety in *actual* driving conditions instead of the controlled environments typically studied by engineers, it is appropriate for any change in drivers' behavior in response to the adoption of ADAS to be incorporated in our estimates. That is, in this case, offsetting behavior does not bias our estimates and need not be addressed. Note that our estimates of the heterogeneous effects of ADAS by vehicle characteristics did not suggest that drivers' risk preferences led them to systematically change their behavior in response to adopting ADAS.

Contamination and Externalities

A final potential source of bias could be caused by *contamination* of the control group, which could occur because treated and untreated vehicles may periodically be involved in accidents with each other. Thus, any safety improvement in the treated vehicles, for example, due to the adoption of ADAS, also may improve the safety of untreated vehicles and cause an estimate of the effectiveness of ADAS safety features—or any other safety features—to be biased downward because it does not account for the positive spillover of safety accruing to vehicles that are not equipped with those safety features.

All observational analyses of accident data that are generated when treated and untreated vehicles share the same roadways will be susceptible to contamination bias, but the bias is mitigated in our analysis for two reasons. First, the vast majority of vehicles (new and used) on the road during our sample period did not have ADAS available as an option at the time of manufacture.³⁰ Second, nearly 50% of the fatal accidents in our sample were single-vehicle accidents.

²⁹ Peltzman (1975) and Winston, Maheshri, and Mannering (2006) found evidence that drivers' offsetting behavior reduced the overall safety benefits of seatbelts and airbags and antilock brakes, respectively.

³⁰ Slightly more than 25% of all the vehicle models in our sample have ADAS, while the share of the total number of vehicles on the road that have ADAS is much smaller.

Generally, the potential for large externalities from ADAS adoption is significantly reduced because fatal accidents are the main contributor to the costs of automobile accidents and because we find ADAS to be effective only at reducing single vehicle fatal accidents. Given those accidents have no scope for externalities, then although external benefits may arise from reductions in fatal multivehicle and pedestrian/cyclist accidents, a plausible conclusion to be drawn from our findings is that they imply modest externalities.

6. Final Comments

Historically, automakers' introduction of a new safety feature has spurred controversy over its effectiveness at reducing the probability of fatal and severe injuries, accounting for drivers' behavior in response to the safety feature. We have addressed this issue empirically in the context of automakers' introduction of ADAS safety features. We have presented causal evidence that ADAS has improved automobile safety by significantly reducing the probability of motorists being involved in fatal and nonfatal accidents, appropriately accounting for the change in drivers' behavior in response to the installation of those safety features in their vehicle. We also have tested for the possibility that our finding could be compromised by selectivity bias that could appear in multiple contexts, and we have consistently rejected that possibility.

As an important qualification regarding public policy, although we have found that ADAS has measurably improved automobile safety, it is premature for two reasons to conclude that policymakers should require automakers to equip all new passenger cars and light trucks with ADAS or even just automatic emergency braking systems, which is required by the fall of 2029. First, the benefits to motorists of autonomous safety technologies must be compared with their substantial installation cost. Second, the share of the total number of vehicles on the road that have ADAS or even AEB systems is small but increasing. Thus, our evidence of some of the benefits of vehicle automation in actual travel environments should be interpreted as an essential component of an initial comprehensive assessment of whether automakers should be required to install ADAS in all their vehicles. Other potentially informative evidence from insurance markets is not yet available.³¹

³¹ We conducted some interviews with insurance industry personnel to see if the adoption of ADAS was reducing insurance rates and we were told that it was premature for auto insurers to be willing to offer discounts for ADAS-equipped vehicles. In the future, when the vehicle capital stock has turned over sufficiently to be comprised of a large share of ADAS-equipped vehicles, it would be useful to estimate the effect of the staggered adoption of ADAS-

Looking further ahead, hopefully our findings will spur greater interest among policymakers and the public in the development and widespread adoption of fully autonomous vehicles and in the potential benefits of other AI transportation technologies (Winston and Karpilow, 2020, Winston, Yan, and Associates, 2024).

equipped vehicles on the nation's insurance costs, which will reflect a tradeoff between the lower claims caused by ADAS's reduction in accidents and the higher claims caused by ADAS's increase in the cost of a car and repairs.

Appendix A. Summary of the Data

As noted, we constructed a data set to analyze the effects of ADAS on all automobile accidents that occurred in Texas from 2010 to 2018 by combining information from three main sources: 1) all registered vehicles from the Texas Department of Transportation, 2) police accident reports from the Texas Department of Public Safety, and 3) trim level vehicle attributes from leading vehicle data aggregators. The main challenge to constructing our data set is to link the data from all three sources to identify whether ADAS safety technology is available on each vehicle registered in Texas.

The registration data contain the Vehicle Identification Number (VIN) of registered vehicles in Texas for the years 2010-2018. The VINs account for all in-state vehicles that were on the road during our sample period. Importantly, we are able to observe *all* the vehicles that were equipped with ADAS safety features in a given calendar year, regardless of whether they were involved in an accident.

The Texas police accident reports record all single and multi-vehicle auto accidents in the state of Texas involving motorists and pedestrians for the years 2010-2018. Importantly, these accident reports also include the VIN of all in-state and out of state vehicles involved in each accident, as well as accident severity, which ranges from vehicle damage only to fatalities. Because we have no information on out of state vehicles that were not involved in an accident, our initial estimates included only in-state vehicles involved in accidents. We then performed a sensitivity analysis including in-state and out of state vehicles in accidents in Texas.

We obtained vehicle attributes down to the trim level by web scraping multiple leading vehicle data aggregators, including TrueCar, Inc., MotorTrend, and Kelly Blue Book. The attributes data are at the detailed model year-make-model-trim level, which enables us to identify the specific safety features of a vehicle that vary at both the model year and the trim level.

To link the data from the three sources, we proceeded as follows. First, for a given VIN in the registration data, we used a commercially available VIN decoder to obtain a string that describes its model year, make, model, and trim (henceforth nameplate).³² Given the large number of registered VINs, simply decoding one-by-one was prohibitively time consuming.³³ Although manufacturers have a certain level of flexibility in terms of what information is encoded in their VINs, the make and model year of a VIN is always encoded in the first three and the tenth digit, respectively.³⁴ We therefore took a representative sample of more than two million VINs weighted by make and model year for decoding. After successfully decoding this sample of VINs and

³² We should point out that not all VIN decoders can decode a VIN to the trim level; most can decode only to the model year-make-model level. For example, NHTSA provides a free VIN decoder that does not decode to the trim level. See <https://www.nhtsa.gov/vin-decoder>.

³³ Decoding the entire registration data would take years.

³⁴ See <https://www.federalregister.gov/documents/2022/03/09/2022-04030/vehicle-identification-number-vin-requirements-manufacturer-identification-certification-replica>.

obtaining their nameplate, we trained a random forest model to classify the remaining VINs.³⁵ This process allowed us to identify the nameplate of each VIN in our registration data.³⁶

Second, the vehicle attributes data contain detailed information on all the features available for a given nameplate, including the safety related features of interest. We provide a summary of the trim level features we were able to obtain, including both safety and non-safety related features, at the end of this appendix. ADAS safety features, however, are marketed under various names by different auto manufacturers with no standardization. For example, Adaptive Cruise Control is called “Intelligent Cruise Control” by Nissan and “Radar Cruise Control with Stop and Go” by Mazda, even though both correspond to the same underlying technology. We therefore used various string manipulation techniques coupled with manual inspection to correctly identify each ADAS safety feature for a given nameplate.

Third, although the decoder provides a nameplate string for a given VIN, this string rarely matches the string we were able to obtain from the web scraped vehicle attributes data, which prohibits a direct merge. For example, the VIN “5J8YD4H05LL024902” is decoded as “2020, Acura, MDX, A-SPEC.” Its counterpart in the attributes data is “2020 Acura MDX Technology and A-Spec Package,” even though they represent the same nameplate. We therefore used fuzzy string match techniques to link the two nameplates. This process allowed us to identify whether ADAS safety technology was available for each VIN in our sample.³⁷

Lastly, because the availability of ADAS safety technology varies at both the model year and trim level, we aggregate nameplates to the trim and the first model year that it received ADAS. That is, for each vehicle make-model (denoted i), we define vehicle type (denoted j) as a combination of trim and the first model year in which ADAS became available. Specifically, for each type j , we aggregate all trims of this type by taking averages of other attributes (such as MSRP). Using the example in Table 1, all Acura MDX trims that began to receive ADAS in model year 2015 are aggregated into the “high” trim; all MDX trims that began to receive ADAS in model year 2018 but not before are aggregated into the “medium” trim; and the remaining MDX trims that never received ADAS are aggregated into the “low” trim.³⁸ In the end, our data construction process resulted in 6,268 unique vehicle types (i.e., aggregated trims) that vary at the yij level. This allowed us to define our treatment variable, the availability of ADAS (denote by the dummy variable S_{yij}) to also vary at the yij level. Specifically, for vehicle type yij , S_{yij} has value 1 if the aggregated trim j received ADAS in model year y ; 0 otherwise.

³⁵ We also experimented with other algorithms such as Decision Tree and Naïve Bayes. Random Forest was our preferred algorithm given its robustness to overfitting and our large sample size. To further reduce computational complexity in this process, we drop the last six digits of each VIN, which only contains a vehicle’s serial number.

³⁶ We also filtered out VINs that either have model years older than 2000 or pertain to irrelevant vehicle categories, such as motorcycles and heavy-duty trailer trucks.

³⁷ The VIN decoder also provides attribute information, such as MSRP, body type, fuel type. We cross-checked attributes from both the decoder and our web scraped data and found that they generally agreed for each nameplate.

³⁸ Once a trim receives ADAS in a model year, it continues to have ADAS in all subsequent model years. This process results in three distinct groups of Acura MDX trims, each with a unique set of aggregated attributes.

Appendix B.

Table 1. Estimates of Influences on Automobile Fatalities During the Sample Period

	2010	2018	Total Change	Implied Change in total fatalities, <i>ceteris paribus</i>
VMT (billion)	234.261	282.037	+20.39%	+20.39% ^a
Vehicle Weight	4001lbs	4137lbs	+136lbs	+6.39% ^b
DUI Arrests	94,440	73,978	-21.67%	-6.55% ^c
ADAS Availability	~0%	~15%	+15%	-1.89% ^d
Observed Fatalities	3,060	3,657	+19.5% ^e	

Notes:

^a The implied change in fatalities was obtained under the plausible *ceteris paribus* assumption that the fatality rate (i.e., fatalities/VMT) is unchanged over this period. This assumption is less plausible during macroeconomic shocks, which change the mix of risky and less risky drivers. Data on VMT were obtained from the Texas Department of Transportation.

^b The implied change in fatalities was obtained by relying on a finding by Anderson and Auffhammer (2014) that a 1000 pound increase in curb weight led to a 47% increase in the probability of a fatal accident. Thus, we multiply (136/1000) by 0.47 to obtain 6.39%. Data on vehicle weights were obtained from national estimates from new car sales by the Environmental Protection Agency.

^c The implied change in fatalities was obtained as follows. According to Levitt and Porter (1999) drivers using alcohol are 8.04 times as likely to be in a fatal accident than sober drivers. Call this risk ratio r . According to the Texas Department of Transportation, 34.5% of accidents in Texas in 2010 involved alcohol. Call this f . Let a_{2010} be the baseline 2010 rate of driving while using alcohol.

Note

$$f = \frac{a \cdot r}{a \cdot r + (1 - a)}$$

which implies $a_{2010} = 6.15\%$. The relative change in risk from 2010 to 2018 can be written as

$$\frac{r_{2018} - r_{2010}}{r_{2010}} = \frac{1 + (r - 1)a_{2018}}{1 + (r - 1)a_{2010}} - 1$$

Assuming that DUI arrests proxy for the fraction of drivers who use alcohol, then according to statistics from the Texas Department of Public Safety, $a_{2018} = 0.7833a_{2010}$. Plugging this into

the previous equation yields an implied decrease of 6.55% of fatalities due to the reduction in alcohol use while driving, *ceteris paribus*.

^d The implied change in fatalities, -1.89%, was obtained by calculating $\frac{0.15 \cdot (-0.32) \cdot 1204}{3060}$, where 3,060 is observed fatalities at the start of the sample period, 0.15 is the increase in the prevalence of ADAS features, -0.32 is our model's estimated reduction in fatal accidents due to ADAS availability, and 1,204 is the number of single vehicle fatal accidents in Texas in 2010 according to the Texas Department of Transportation.

^e Data on fatalities were obtained from the Texas Department of Transportation.

Appendix C. Summary of Non-ADAS and Non-Safety Related Trim Features

The non-ADAS and non-safety related trim features that were available for vehicles with high trim, but not for vehicles with low trim are as follows:

Non-ADAS Vehicle Safety Features

Rear and side view with simulated aerial camera
360 Degree Surround Camera
Panoramic View Monitor
Digital Backup Sensors
Active Blind Spot w/Front Park Sensor
Adaptive Light Control
Auto-Dimming Rearview Mirror
Bi-Xenon Cornering Headlights
Black Out LED Daytime Running Lights
Enhanced Active Park Assist w/Forward Sensing System
Inflatable Rear-Seatbelts
Night View Assist PLUS w/Pedestrian Detection
Side Curtain and Seat Mounted Side Impact Airbags
Trailer Tow Camera System
Heated Sideview Mirrors

Non-Safety Related Trim Features

Rear power outlet(s)
Cargo area power outlet(s)
Anti-Theft Alarm System w/Immobilizer
Intrusion Sensor
Heated Windshield Washer Reservoir (SPC)
Keyless Entry w/Hands-Free Tailgate Opening
Headlamp Washers
Premium Leather Collection
Heated Rear Seats

References

- Alarfaj, Abdullah F., W. Michael Griffin, and Constantine Samaras. 2020. "Decarbonizing US Passenger Vehicle Transport Under Electrification and Automation Uncertainty Has a Travel Budget," *Environmental Research Letters*, volume 15, pp. 1-12.
- Anderson, Michael L., and Maximilian Auffhammer. 2014. "Pounds that kill: The external costs of vehicle weight." *Review of Economic Studies*, volume 81, no. 2, pp. 535-571.
- Blanco, Myra, Jon Atwood, Sheldon Russell, Timothy Trimble, Julie McClafferty, and Miguel Perez. 2016. *Automated Vehicle Crash Rate Comparison Using Naturalistic Data*. Final Report, Blacksburg, Virginia: Virginia Tech Transportation Institute.
- Cameron, A. C. and P. K. Trivedi. 1998. *Regression Analysis of Count Data*. Cambridge University Press, Cambridge.
- Edlin, Aaron S., and Pinar Karaca-Mandic. 2006. "The Accident Externality from Driving," *Journal of Political Economy*, volume 114, October, pp. 931-955.
- Gourieroux, Christian, et al. 1984. "Pseudo Maximum Likelihood Methods: Applications to Poisson Models." *Econometrica*, vol. 52, no. 3, pp. 701-720.
- Haghani, Milad, and Michiel C.J. Bliemer. 2023. "Emerging Trends and Influential Outsiders of Transportation Science," *Transportation Letters*, volume 15, no. 5, pp. 386-422.
- Harper, Sam, Thomas J. Charters, and Erin C. Strumpf. 2015. "Trends in Socioeconomic Inequalities in Motor Vehicle Accident Deaths in the United States, 1995-2010," *American Journal of Epidemiology*, volume 182, October, pp. 606-614.
- Levitt, Steven and Jack Porter. 1999. "Estimating the Effect of Alcohol on Driver Risk Using Only Fatal Accident Statistics," NBER Working Paper 6944, February.
- Maheshri, Vikram, and Clifford Winston. 2025. "A Critique of Automobile Safety Methodologies: Implications for Policymakers as We Enter the Autonomous Era," working paper.
- Maheshri, Vikram, and Clifford Winston. 2024. "The Effect of Reductions in Vehicle Miles Traveled on Highway Fatalities and Congestion with Heterogeneous Motorists," in Clifford Winston, Jia Yan, and Associates, *Revitalizing a Nation: Competition and Innovation in the US Transportation System*, Brookings Institution Press, pp. 11-25.
- Mosquet, Xavier, Michelle Andersen, and Aakash Arora. 2016. "A Roadmap to Safer Driving Through Advanced Driver Assistance Systems," *Auto Tech Review*, volume 5, pp. 20-25.
- National Highway Traffic Safety Administration. 2023a. Title of the Regulation: Federal Motor Vehicle Safety Standards: Automatic Emergency Braking Systems for Light Vehicles. Federal Register, US Department of Transportation.
- National Highway Traffic Safety Administration. 2023b. *The Economic and Societal Impact of Motor Vehicle Crashes, 2019* (revised), US Department of Transportation, February.

Peltzman, Sam. 1975. "The Effects of Automobile Safety Regulation," *Journal of Political Economy*, volume 83, August, pp. 677-725.

Wählberg, A.E., and L. Dorn. 2023. "The Effects of Electronic Stability Control (ESC) on Fatal Crash Rates in the United States," *Journal of Safety Research*, December, pp. 1-13.

Winston, Clifford, and Quentin Karpilow. 2020. *Autonomous Vehicles: The Road to Economic Growth?*, Brookings Institution Press, Washington, DC.

Winston, Clifford, Vikram Maheshri, and Fred Mannering. 2006. "An Exploration of the Offset Hypothesis Using Disaggregate Data: The Case of Airbags and Antilock Brakes," *Journal of Risk and Uncertainty*, volume 32, March, pp. 83-99.

Winston, Clifford, Jia Yan, and Associates. 2024. *Revitalizing a Nation: Competition and Innovation in the US Transportation System*, Brookings Institution Press.

Wooldridge, J.M., 2023. "Simple Approaches to Nonlinear Difference-in-Differences with Panel Data," *The Econometrics Journal*, volume 26, March, pp.C31-C66.

Table 1. Example of the Data Structure for the Acura MDX

Treated Vehicles			2015 H	2015-2016 H	2015-2017 H	2018 M 2015-2018 H	2018-2019 M 2015-2019 H
Untreated Vehicles	2000-2013 L 2000-2013 M 2000-2013 H	2000-2014 L 2000-2014 M 2000-2014 H	2000-2015 L 2000-2015 M 2000-2014 H	2000-2016 L 2000-2016 M 2000-2014 H	2000-2017 L 2000-2017 M 2000-2014 H	2000-2018 L 2000-2017 M 2000-2014 H	2000-2019 L 2000-2017 M 2000-2014 H

Calendar Year	2013	2014	2015	2016	2017	2018	2019
---------------	------	------	------	------	------	------	------

Notes: There are three trim levels for the MDX: L, M and H. Trim level H received ADAS safety features in model year 2015. Trim level M received ADAS safety features in 2018.

Table 2. Effects of ADAS on Accidents and Fatalities


Dependent Variable	All Accidents		Fatal Accidents	
	Single Veh.	Multi Veh.	Single Veh.	Multi Veh.
	(1)	(2)	(3)	(4)
ADAS Safety Features Dummy	0.86***	0.89**	0.68**	0.99
	[-1.31]***	[-8.74]***	[-0.12]**	[-0.00]
	(0.04)	(0.04)	(0.12)	(0.12)
Make-Model-Trim-Calendar Year (ijt) FEs?	Y	Y	Y	Y
Make-Model-Model Year-Calendar Year (yt) FEs?	Y	Y	Y	Y
Pseudo R-squared	0.77	0.87	0.21	0.24
Number of observations	4,776	4,983	1,643	2,315

Notes: Incidence Risk Ratios are presented from Poisson maximum likelihood regressions. Marginal effects are presented below in hard brackets []. Heteroskedasticity robust standard errors for the model parameters clustered by model year-make-model are presented in parentheses. Vehicle trims that are never equipped with ADAS are excluded. *** 99% significance, ** 95% significance, * 90% significance.

Table 3. Effects of ADAS on Driving

Dependent Variable	VMT	log(VMT)
	(1)	(2)
ADAS Safety Features Dummy	-31.56 (138.07)	0.04 (0.04)
Make-Model-Trim-Calendar Year (ijt) FEs?	Y	Y
Make-Model-Model Year-Calendar Year (yt) FEs?	Y	Y
Pseudo R-squared	0.74	0.77
Number of observations	6,565	6,464

Notes: Heteroskedasticity robust standard errors clustered by model year-make-model presented in parentheses. Vehicle trims that are never equipped with ADAS are excluded. For registrations, vehicles in their first*** 99% significance, ** 95% significance, * 90% significance.

Figure 1. Evolution of the Rollout of ADAS**Panel A: Number of Unique Vehicle Trims over Time**

Alt text: Line charts showing the evolution of vehicle trims with and without ADAS over time in absolute numbers and relative proportions.

Panel B: Share of ADAS Equipped Vehicles over Time

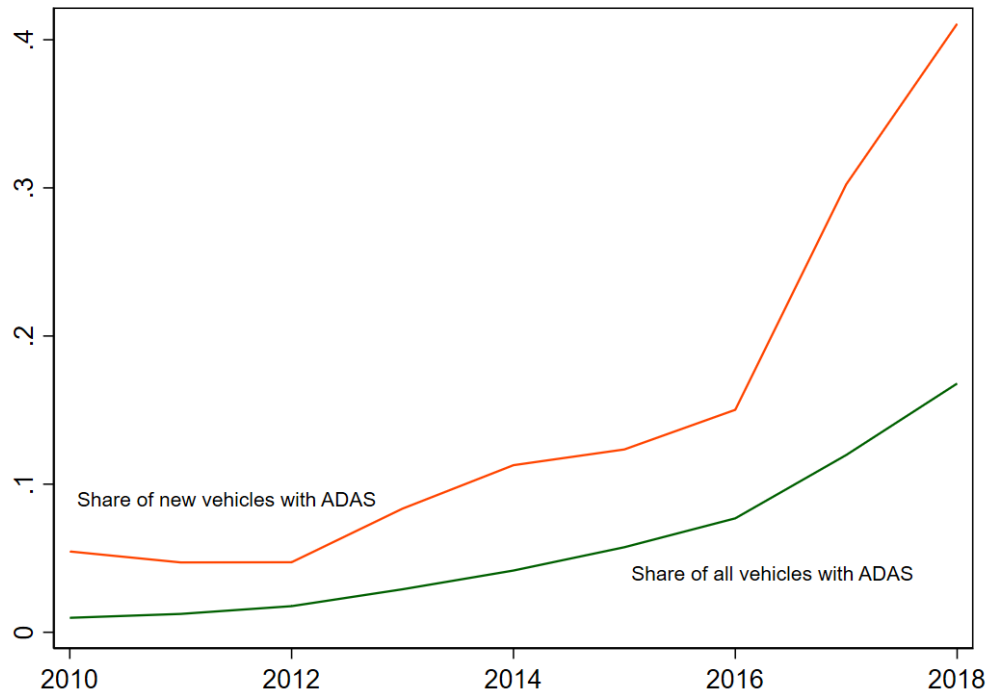
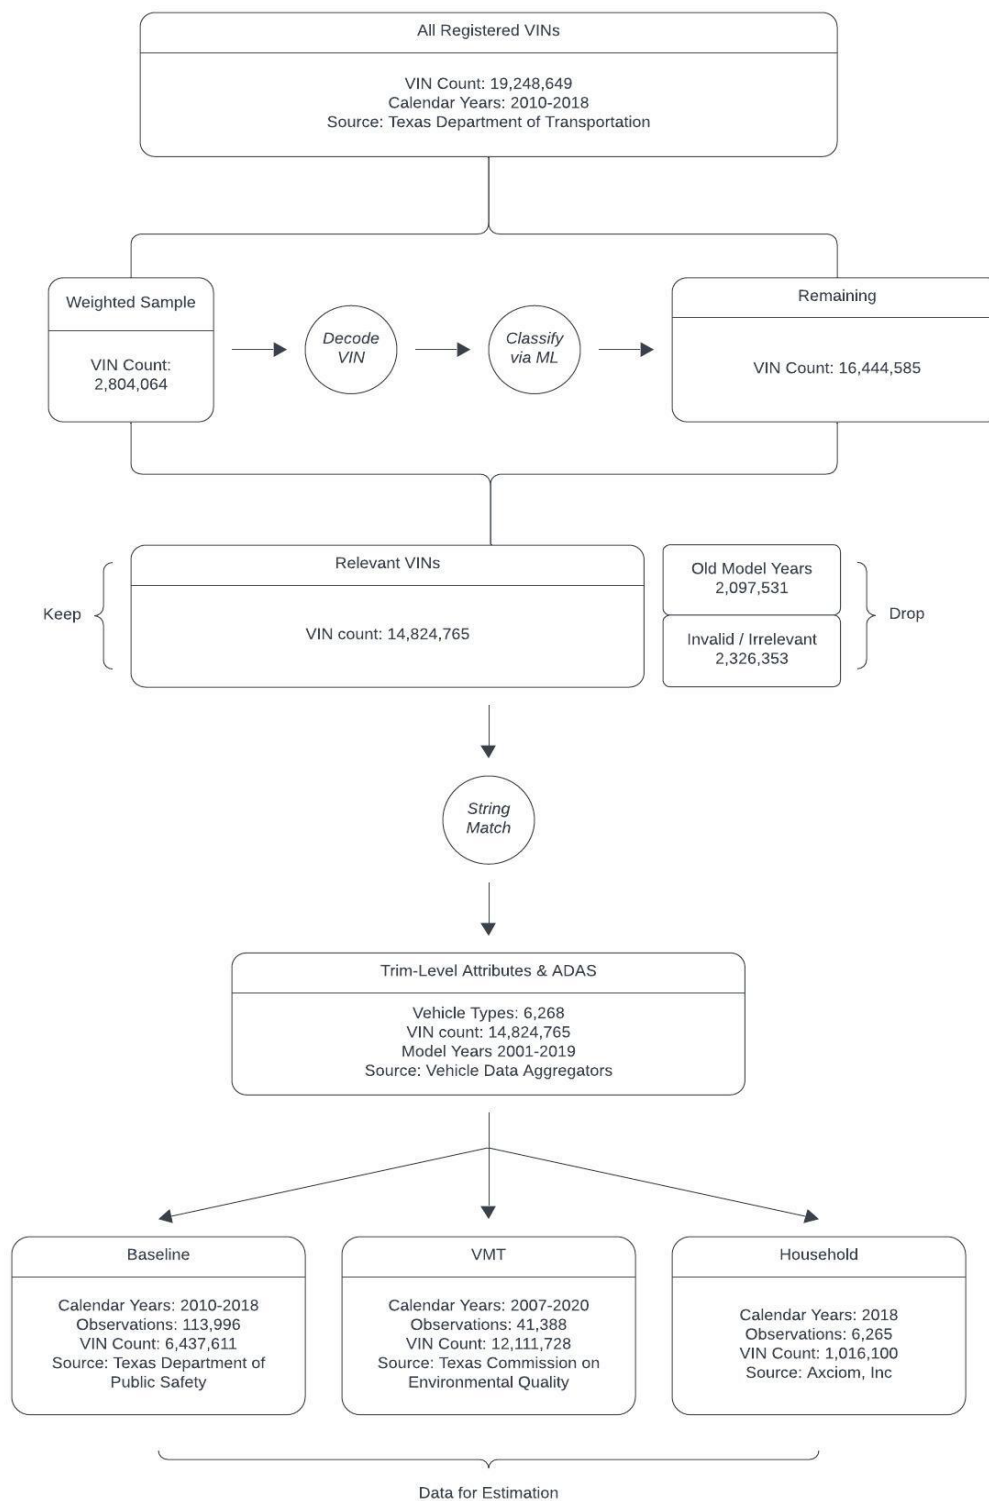
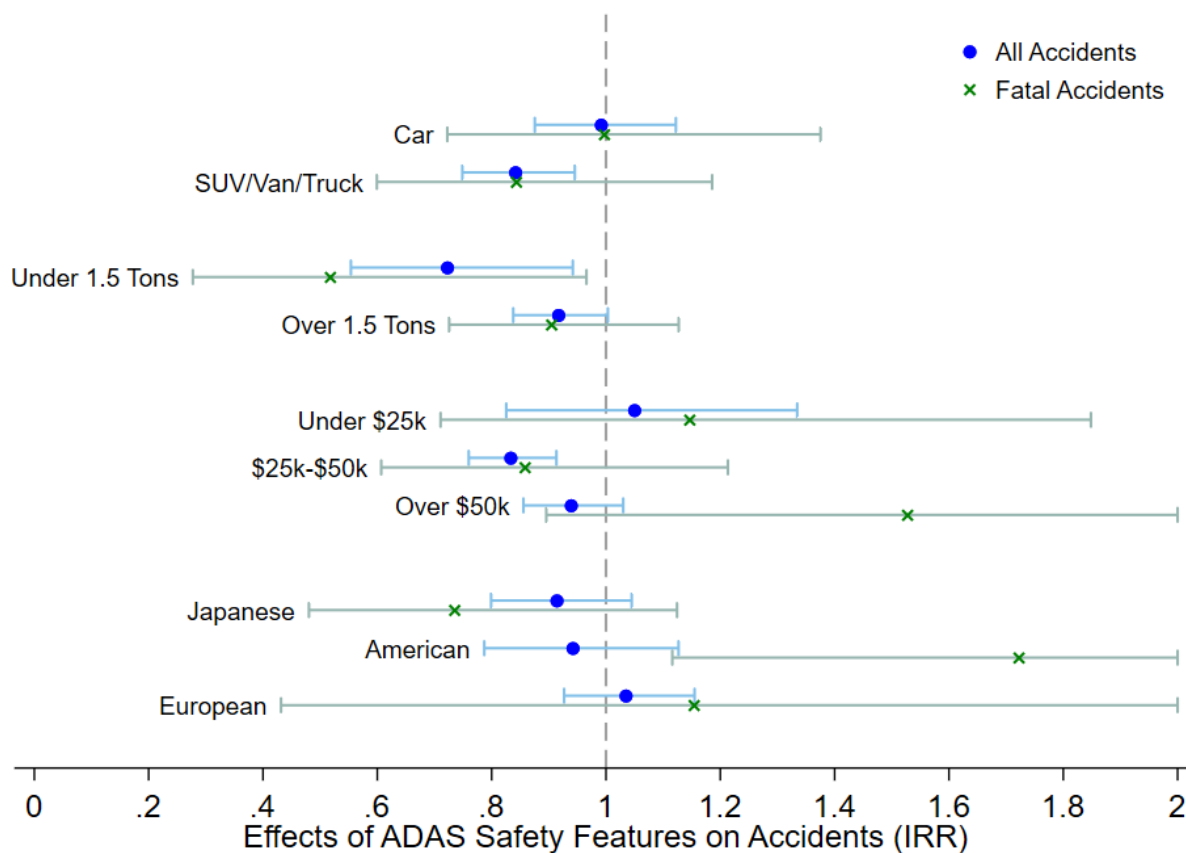


Figure 2. Overview of the Data Construction



Alt text: Flow chart describing the dataset construction process.

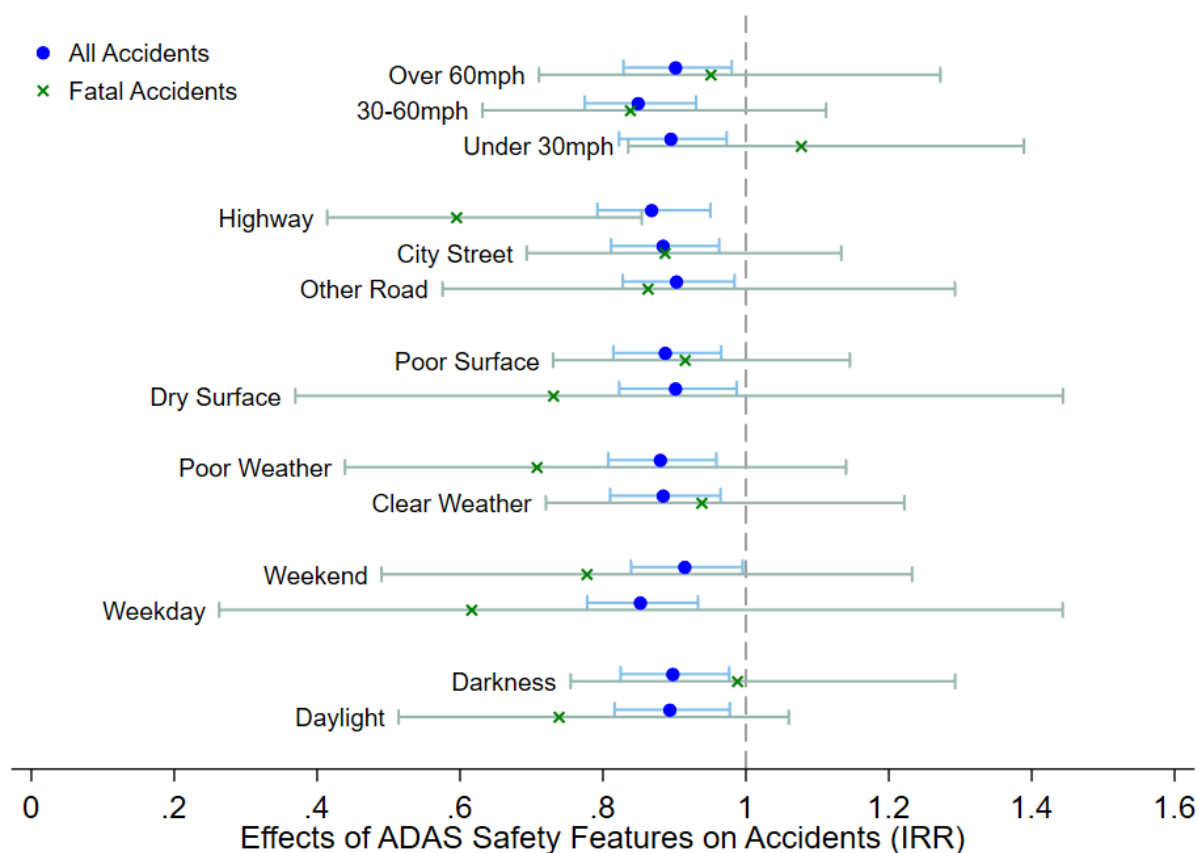
Figure 3. Heterogeneous Effects of ADAS Safety Features on Accident Rate by Vehicle Type



Note: Incidence Risk Ratios are presented from Poisson maximum likelihood regressions with 95% confidence intervals formed from heteroskedasticity robust standard errors clustered by model year-make-model. Parameter estimates and confidence intervals are truncated at 2 for clarity.

Alt text: Chart showing the effects of ADAS availability on the accident rates of different vehicles (differing by body type, weight, price, and country of automaker).

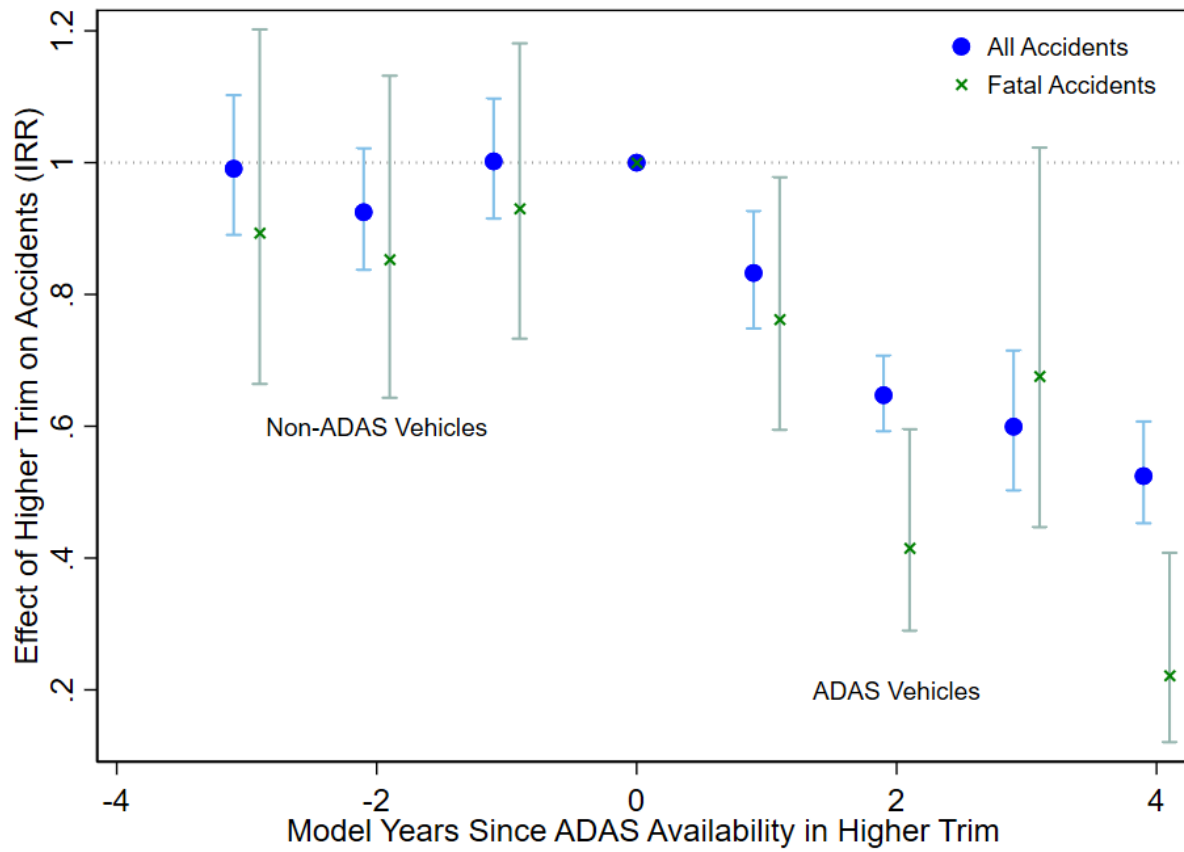
Figure 4. Heterogeneous Effects of ADAS Safety Features on Accident Rate by Accident Conditions



Note: Incidence Risk Ratios are presented from Poisson maximum likelihood regressions with 95% confidence intervals formed from heteroskedasticity robust standard errors clustered by model year-make-model.

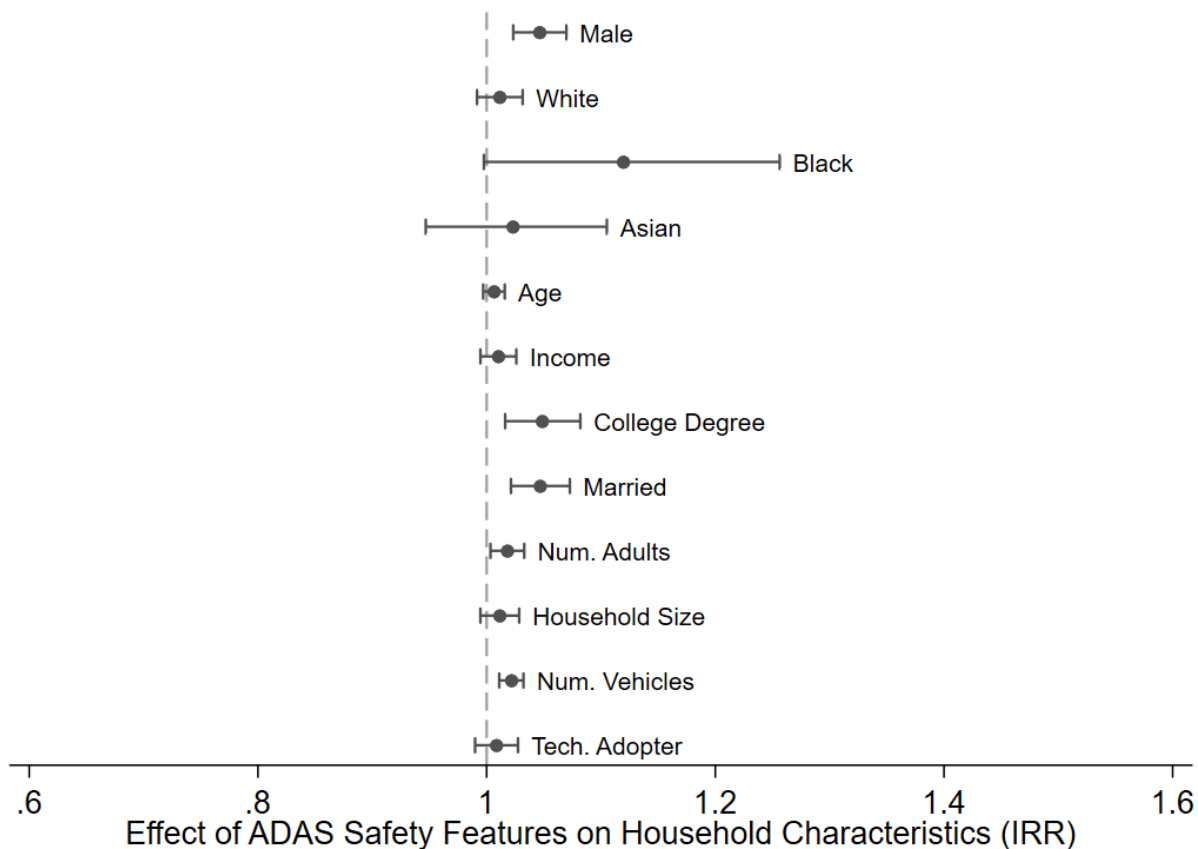
Alt text: Chart showing the effects of ADAS availability on accident rates in different driving conditions (differing by speed, road type, surface and weather conditions, day of week and time of day).

Figure 5. Event Study Style Plot of the Effects of ADAS



Alt text: An event-study style chart showing the effects of ADAS availability on accident rates by number of model years since ADAS was available in a higher trim level.

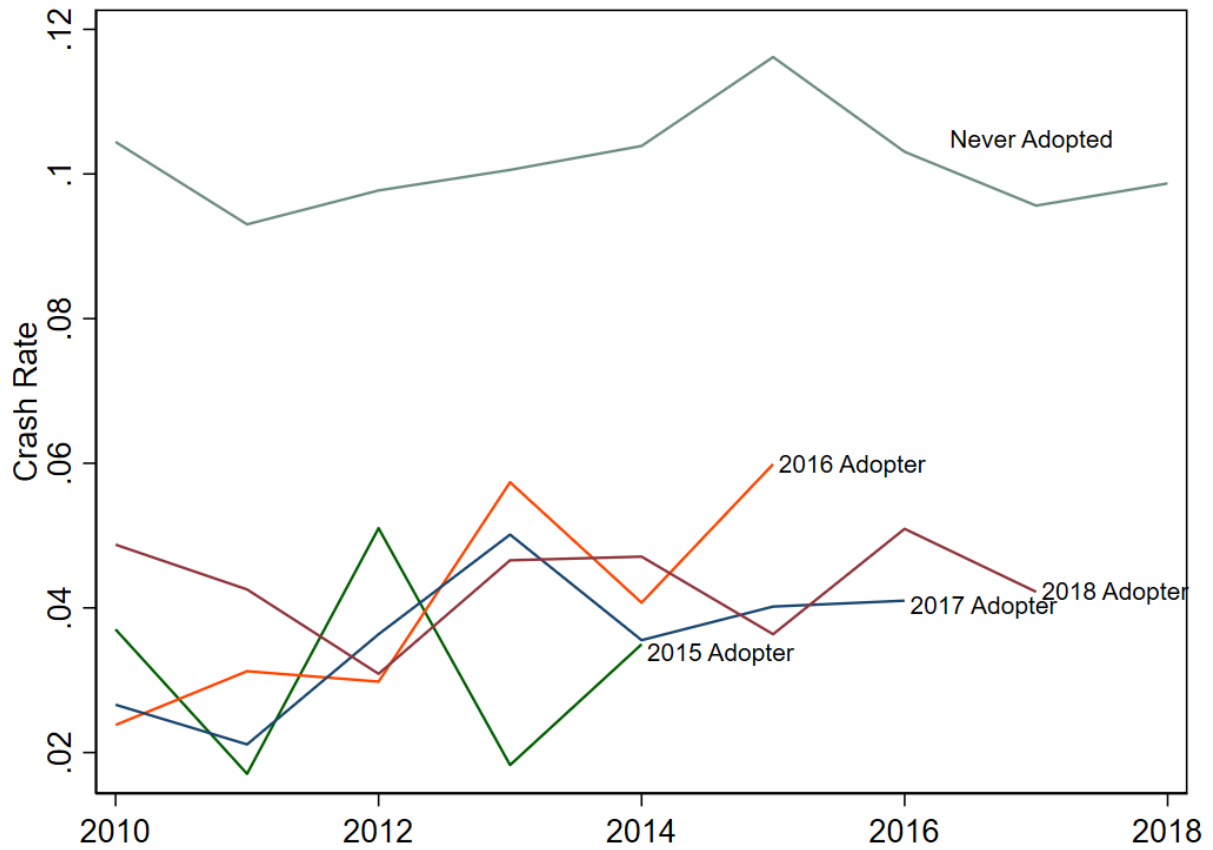
Figure 6. Effects of ADAS by Household Characteristics (IRR)



Note: Incidence Risk Ratios are presented from Poisson maximum likelihood regressions with 95% confidence intervals formed from heteroskedasticity robust standard errors clustered by model year-make-model.

Alt text: Chart showing the effects of ADAS availability on various household characteristics (gender, race, age, income, education, marital status and tech adoption propensity of driver; and the size of and number of vehicles owned by the household.)

Figure 7. Crash Rate over Time by Household ADAS Adoption Year



Notes: This figure shows the evolution of crash rates of households over time split up by when each household first adopted an ADAS enabled vehicle.

Alt text: Same as notes.