AI at the Wheel: The Effectiveness of Advanced Driver-Assistance Systems

Vikram Maheshri *University of Houston*
Clifford Winston *Brookings Institution*
Yidi Wu *Cornerstone Research*

Abstract

Has automakers' use of artificial intelligence (AI) in advanced driver-assistance systems (ADASs) 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 ADASs 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 ADASs have 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 transportation technologies using AI.

1. Introduction

Since the 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.[1]

We are grateful to Wesley Blundell, Michael French, Edward Glaeser, Fred Mannering, Ashley Nunes, Alex Oyler, Marc Scribner, participants in the spring 2024 National Bureau of Economic Research 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] Government highway expenditures also have been used to improve the safety of the road system.

[*Journal of Law and Economics,* vol. 69 (May 2026)]

0022-2186/2026/6902-0013$10.00

Beginning in the late 2000s, automakers took an important step forward in improving safety by steadily equipping their vehicles with advanced driver-assistance systems (ADASs) based on artificial intelligence. An ADAS consists of a suite of safety features that assist in both the forward dimension (automatic emergency braking [AEB] and adaptive cruise control) and the lateral dimension (lane departure warning and blind-spot collision prevention).[2] An ADAS is standard for some vehicle makes, models, and trims; can be purchased as an option for others; and is unavailable for purchase for the remaining makes, models, and trims.[3] According to the American Automobile Association, at least one ADAS feature was available in 92.7 percent of new vehicle models in the United States in 2018 (Slovick 2020).

An ADAS distinguishes itself from other automobile safety features because it assists the driver by performing functions 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 an ADAS enhances safety by substituting for a driver's action or alerting a driver to help prevent an accident from occurring.

The recent adoption of ADASs in the United States motivates our interest in assessing their effectiveness at reducing accident risk. As is 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 its intended effects. In contrast, an engineering analysis, which is the basis for current estimates of the effectiveness of ADASs, does not account for the potential bias of motorists who self-select into vehicles equipped with ADASs and for motorists' potential behavioral response to drive more aggressively in ADAS-equipped vehicles.

To date, ADASs have been voluntarily installed in certain vehicles by automakers and selected by motorists through their choices of vehicle and trim. But recently, the federal government issued one of the most significant changes to safety standards in years by requiring that all new passenger cars and light trucks be equipped with AEB systems, an important component of ADASs. Automakers have until fall 2029 to ensure that the AEB systems on their 2030 vehicles comply with federal safety standards.[4] Thus, our assessment is further motivated by policy-

[2] Except for adaptive cruise control, features of advanced driver-assistance systems (ADASs) engage autonomously because they are often enabled by default.

[3] In addition to safety technology, a vehicle's trim includes powertrain options, aesthetic features, and comfort amenities.

[4] 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 62 mph; vehicles traveling as fast as 45 mph 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 (49 C.F.R 571, 595, and 596 [2025]). Currently, no commercially available automatic emergency braking technologies satisfy these stringent technical requirements.

makers' mandating the adoption of a component of ADASs without systematic evidence of its effectiveness in improving safety in real-world situations.

As noted, the availability of ADASs 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 ADASs by comparing the safety performance of extremely similar vehicles with and without ADASs. To the best of our knowledge, our paper is the first to use modern data-collection methods to extract the detailed information 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; 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 prevented.[5] In contrast, our approach enables us to identify the unconditional effectiveness of ADASs' 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 comprising 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, which enables us to construct the accident history of the universe of registered vehicles.[6] 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.[7]

[5] 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 nonfatal accidents. However, because we are concerned with the effectiveness of a vehicle attribute, we aggregate our data by vehicle type.

[6] 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 a driver involved in an accident causing injury, death, or property damage exceeding $1,000 must report the incident to law enforcement. This being said, it is still possible for nonfatal accidents to be underreported in the Texas police data, which may conceivably affect our findings on the effect of ADASs on nonfatal accidents but would have no effect on our findings on the effect of ADASs on fatal accidents. We are not aware of any external evidence that documents the extent of underreporting of nonfatal accidents in Texas and that assesses the extent of any bias associated with this underreporting.

[7] 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.

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 decode the precise trim level of a vehicle from its vehicle identification number (VIN). In sum, our dataset 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 are not equipped with ADAS-related safety features. Our findings should be reasonably representative of the United States as a whole because Texas is a large state with a diverse population and fleet of vehicles driven in urban and rural environments that vary in geography, weather, and traffic density.

Given these 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 nonrandom—by exploiting plausibly exogenous variation in the availability of ADASs 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 ADASs, as opposed to comparing the disaggregated safety performance of individual drivers. Using the latter approach, the effect of ADASs would be identified only under the dubious assumption that a driver's propensity to purchase an ADAS-equipped vehicle was uncorrelated with their attitudes toward safety and their driving abilities (perhaps conditional on some small set of observable driver characteristics).

In our approach, we leverage that ADASs became available at different times for different trim levels—notably, within vehicles of the same make and model.[8] We therefore identify the causal effect of ADASs on accidents under the weaker assumption that drivers did not systematically opt for higher-trim-level vehicles solely because of the availability of ADASs. 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 trim (higher versus lower) and the timing of ADASs' availability.

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

[8] Af Wåhlberg and Dorn (2023) assess the effectiveness of vehicle electronic stability control on fatal crash rates, but they do not compare cars' safety performance with and without this technology.

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 ADASs. 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 motorist who purchases a vehicle with a higher trim is not systematically affected by the availability of ADASs 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 ADASs 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 ADASs on accident outcomes do not vary systematically across vehicle characteristics and travel conditions.

Importantly, we find that ADASs are highly effective at improving automobile safety even after accounting for drivers' behavioral responses to their availability and installation. Specifically, ADASs 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. The systems have a small and statistically imprecise effect on reducing the risk of a motorist getting in a multi-vehicle fatal accident, but we suggest that ADASs are 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, ADASs represent a significant advance in automobile safety by substituting effectively for a driver's action or alerting a driver and by providing credible evidence that fully autonomous vehicles could notably improve highway transportation safety. From a policy perspective, our evidence of some 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 an ADAS in all their vehicles.

2. Estimating Efficacy

The staggered rollout of the availability of ADASs 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 ADASs on accident risk. We construct aggregate versions of the key variables, ADAS availability and accidents, to execute the empirical analysis, and then we explain our specification to estimate the effect of ADAS availability on accident risk.

2.1. 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 (for example, luxury editions) may have an ADAS and others (for example, standard editions) may not. We therefore expand the definition of vehicle type to be a combination of make, model, and trim, where trim levels, defined in Section 3, are indexed separately by j.

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

2.2. Accidents

For each vehicle in each calendar year of our sample, we observe its 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 vehicle type y_{ij} had in year t.

The temporal and cross-sectional 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 t, 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 trims, model years, and calendar years to identify the causal effect of ADASs 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 2000–2019 sample period.[9] The Acura has three trim levels, which we denote low (L), medium (M), and high (H), and each is associated with the period, if any,

[9] 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 2010–18, it includes some model year 2019 vehicles.

Table 1

Example of the Data Structure for the Acura MDX, by Calendar Year

	2013	2014	2015	2016	2017	2018	2019
Treated vehicles			2015 H	2015–16 H	2015–17 H	2018 M 2015–18 H	2018–19 M 2015–19 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

Note. There are three trim levels for the MDX: L, M, and H. Trim level H received an advanced driver-assistance system in model year 2015. Trim level M received it in 2018.

that it was equipped with an ADAS.[10] Vehicles with a low trim level were never equipped with ADASs during our sample period; vehicles with a medium trim level were equipped with ADASs in model year 2018 but not before that calendar year; and vehicles with a high trim level were equipped with ADASs in 2015 but not before that calendar year. The three different trim levels of Acura MDXs on the road during our sample period enable us to define these particular treated vehicles as Acura MDXs with high and/or medium trim levels that include ADASs. The untreated or control vehicles are Acura MDXs that do not include ADASs. We define all of the other treated vehicle types and control vehicles in our sample in the same way.

2.3. Specification

Previous safety research (see, for example, Maheshri and Winston 2024) has specified vehicle accidents A_{yijt} in a Poisson regression framework because accidents take on small, discrete, nonnegative 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 nonrobust estimator of conditional mean parameters, and this weakness is exacerbated when using fixed effects.[11] 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 quasi-maximum likelihood estimator discussed by Gourierieux, Monfort, and Trognon (1984) because it yields consistent estimates of the effect of interest without the distractions of variance assumptions.

We therefore specify our models of accidents and fatal accidents as

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

where S_{yijt} is a dummy variable equal to one if an ADAS was available either as standard equipment or purchased through an optional package on vehicle y_{ij} in

[10] 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 five trim configurations that were marketed at three trim levels: standard, technology, and advance. Meanwhile, the 2015 Acura MDX was marketed at four trim levels: base, advance/entertainment, tech, and tech/entertainment. Because we are analyzing the effects of ADASs 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 ADASs in the same model year as a single trim (in this case, low, medium, or high).

[11] As Gourierieux, Monfort, and Trognon (1984) show, the negative binomial estimator requires both mean and variance to be correctly specified, whereas the Poisson estimator requires only 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 underdispersion 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.

year t and zero otherwise, λ_{ijt} is a group of make-model-trim-calendar year fixed effects, λ_{iyt} is a group of make-model-model year-calendar year fixed effects, and ε_{yijt} is an error term.[12]

The key identifying assumption that enables parameter β to be interpreted as the causal effect of the availability of ADASs on selected vehicles on the total number of accidents is that $\mathrm{cov}(S_{yijt}, \varepsilon_{yijt} \mid \lambda_{ijt}, \lambda_{iyt}) = 0$. That is, motorists who purchase higher trim vehicles during the first model year that ADASs are 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.

On the basis of early experiences with autonomous vehicles in controlled testing environments (Blanco et al. 2016; Mosquet, Andersen, and Arora 2016), we expect the availability of ADASs 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 dataset consisting of all vehicles registered in Texas from 2010 to 2018 along with their trims, which we used to identify whether a given vehicle is equipped with an ADAS. We used leading vehicle data aggregators that describe the available safety features in all new vehicle trims, including ADASs, to identify the vehicles equipped with ADASs during the sample period.[13] Then, for each vehicle we merged information from the universe of Texas police accident reports to construct its detailed accident history. To the best of our knowledge, this is the first dataset at the vehicle-trim level that has been used to assess the efficacy of vehicle safety features.

Vehicle Identification Numbers. Extracting safety features from 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 VINs 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

[12] Data on specific vehicles that were purchased with an ADAS as an optional package are not available. However, when a vehicle, defined by make and model, offers an ADAS as an option instead of as standard, most consumers who select that vehicle also are likely to purchase the optional ADAS. The reason is that the entire trim package of a vehicle that offers an optional ADAS is usually more expensive than the entire trim package of the same or similar vehicle that does not offer an ADAS as an option. Thus, consumers who do not want the optional ADAS 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 an ADAS as an option. Anecdotal evidence obtained from car dealers was consistent with this characterization of consumer behavior.

[13] 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.

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.[14] 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 an ADAS for each vehicle. We provide a detailed description of how we extracted safety features from VINs and aggregated vehicles to the trim level in the Online Appendix. 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 2010–18. Figure 1*A* shows that vehicle trims equipped with ADASs were relatively scarce for most of our sample period, though they gradually became more common after 2015. Indeed, Figure 1*B* shows that while the number of vehicles equipped with ADASs represents a small (less than 20 percent) share of all vehicles throughout our sample period, newly registered vehicles were increasingly more likely to be equipped with ADASs toward the last few years of our sample period. This pattern reflects 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 Traveled. The introduction of ADASs could affect driving intensity, as measured by vehicle miles traveled (VMT), which could confound our findings on the safety effects of ADASs. To explore this possibility, we constructed from our main sample a subsample that contains the VMT for each vehicle from the Texas Commission on Environmental Quality. The 14 largest counties in Texas require each vehicle to be subjected to emissions testing annually prior to being registered. Among other information, the commission collects the VINs of the vehicles and the exact annual odometer reading for each registered vehicle in those 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 safety features led consumers to self-select systematically into ADAS-equipped and non-ADAS-equipped trims. For more than 200,000 registered vehicle owners in Texas randomly selected from our main sample, we obtained data from Acxiom on 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-ADAS-equipped trims on the basis of observed influences.

[14] Using the example in Table 1, the Acura MDX high-level trim is called the type S advance, which made an ADAS available in model year 2015. The lowest level is the base trim, which did not make an ADAS available during the sample period.

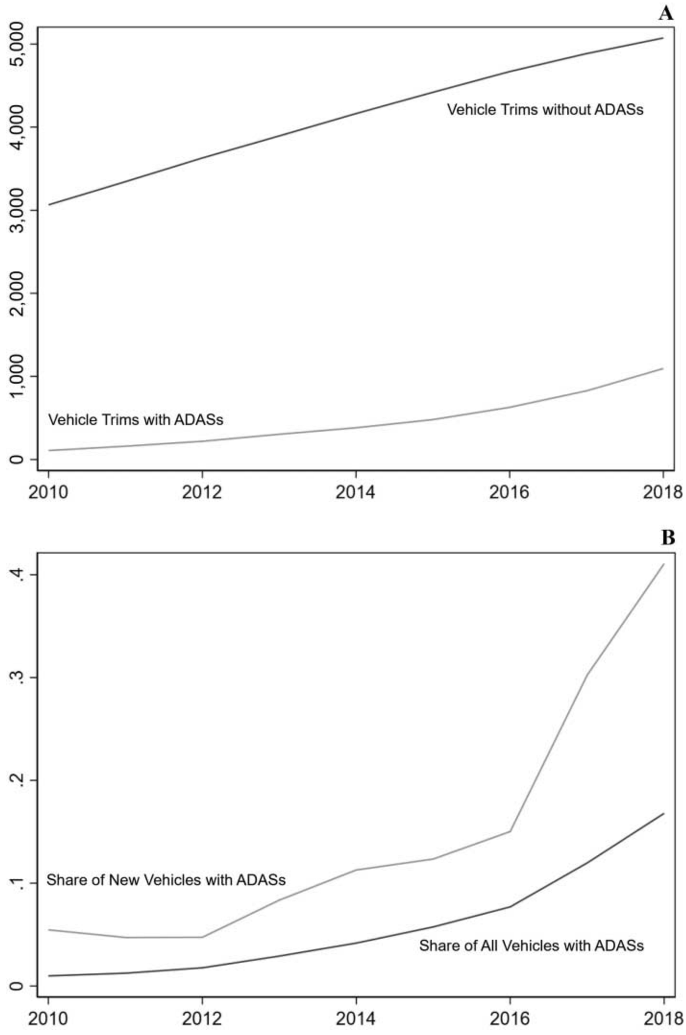

Figure 1. Evolution of the rollout of advanced driver-assistance systems (ADASs), 2010–18. *A*, Unique vehicle trims; *B,* share of ADAS-equipped vehicles.

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. We use the final dataset for our estimations. As noted, the Online Appendix 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 any ADAS-related safety technology was available on each vehicle registered in Texas.

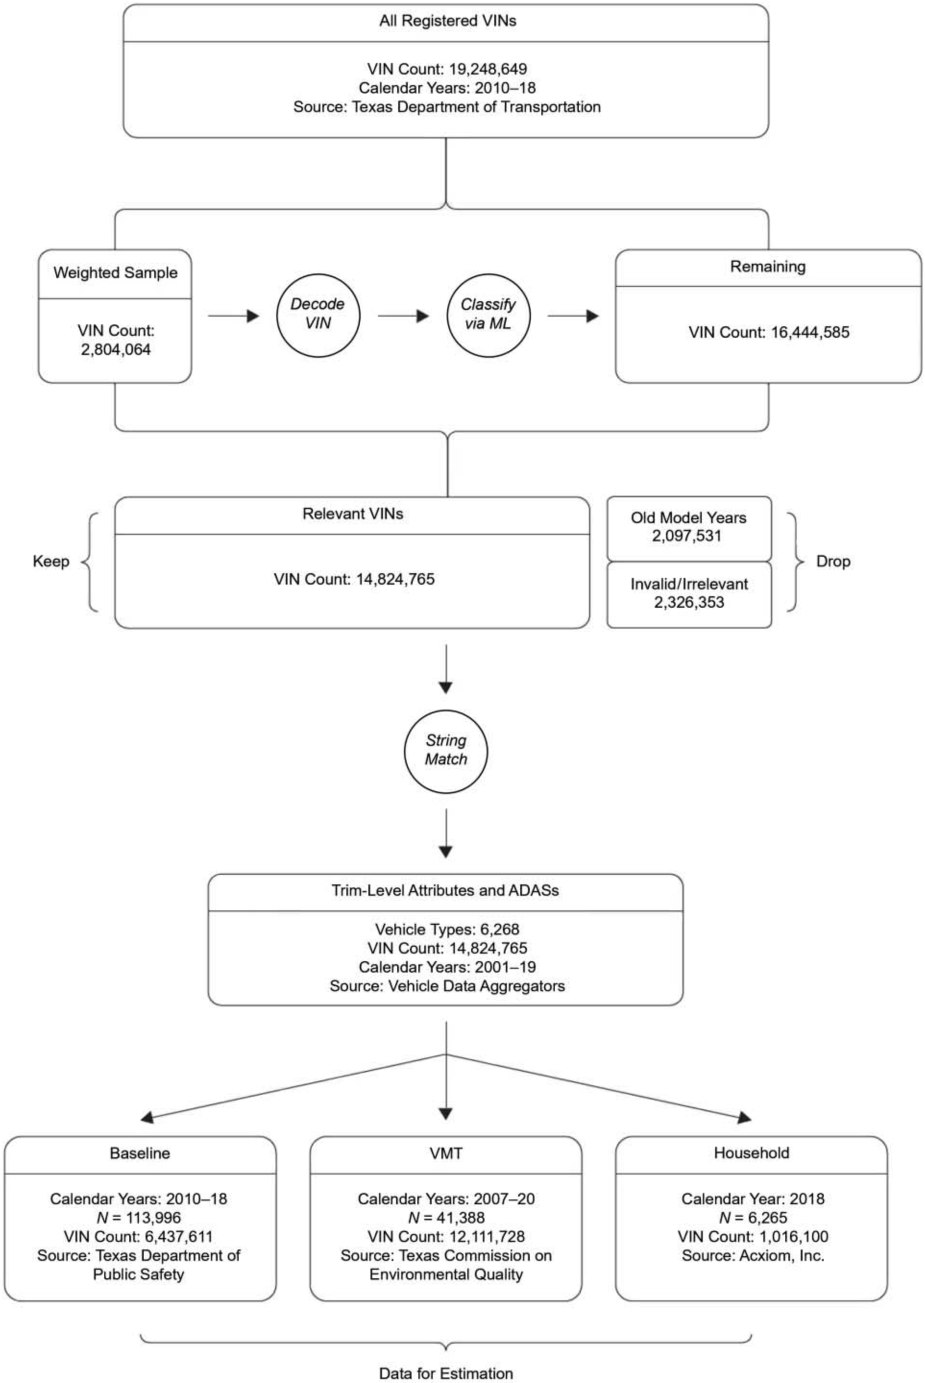

Figure 2. Overview of the dataset construction. ADAS = advanced driver-assistance system; ML = machine learning; VIN = vehicle identification number; VMT = vehicle miles traveled. Generated by the authors with indicated data sources.

Table 2
Effects of ADASs on Accidents and Fatalities

	All Accidents		Fatal Accidents	
Dependent Variable	Single Vehicle (1)	Multivehicle (2)	Single Vehicle (3)	Multivehicle (4)
ADAS dummy	.86**	.89*	.68*	.99
	[−1.31]**	[−8.74]**	[−.12]*	[−.00]
	(.04)	(.04)	(.12)	(.12)
Pseudo-R^2	.77	.87	.21	.24
N	4,776	4,983	1,643	2,315

Note. Incidence risk ratios are presented from Poisson maximum likelihood regressions. Marginal effects are in brackets. Heteroskedasticity-robust standard errors for the model parameters clustered by model year-make-model are in parentheses. Vehicle trims that are never equipped with advanced driver-assistance systems (ADASs) are excluded. All regressions include make-model-trim-calendar year and make-model-model year-calendar year fixed effects.
* 95 percent significance.
** 99 percent significance.

4. Estimation Results

Table 2 presents the effects of the availability of ADASs 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 percent. An IRR of less than 1 corresponds to a negative effect on vehicle accidents, and an IRR of .80 means the expected count decreases by 20 percent. We also present marginal effects in Table 2 as if we were presenting standard regression 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 ADASs is likely to simultaneously influence VMT and accidents and fatalities. Even if one of those influences were small, it would still prevent us from determining the distinct effects of ADASs on accidents and fatalities. However, we directly estimate the effect of the adoption of ADASs on VMT to provide further context for our results.

In each regression, we restrict our sample to model year-make-model-trim combinations with at least 10 registered vehicles to ensure that the results are not affected by rare vehicles, such as Ferraris.[15] We also restrict the sample to vehicle types that are equipped with ADASs at some point during the sample period to ensure that the results are not affected by variation among never-treated vehicles; when we relax this assumption and include never-treated vehicle types, our findings are unchanged.

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

4.1. Basic Findings

According to the estimated IRRs in Table 2, the availability of ADASs reduces the total number of accidents of a given vehicle and trim type by 14 percent in single-vehicle accidents (column 1) and by 11 percent 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.[16]

Bear in mind that only a small share of the entire vehicle fleet is equipped with autonomous vehicle safety features. It therefore is plausible that ADASs are less effective in reducing multivehicle accidents compared with reducing single-vehicle accidents because the other vehicles involved in a multivehicle accident may not be equipped with autonomous vehicle safety features. As more of the nation's vehicle fleet is equipped with autonomous vehicle safety features, ADASs may become equally effective in reducing single-vehicle and multivehicle accidents.[17]

These safety technologies are even more effective at reducing single-vehicle fatal accidents than they are at reducing nonfatal 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 in which the 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 from occurring.

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

Notwithstanding this uncertainty, the upper bound of ADASs' effectiveness in reducing multivehicle fatal accidents is notably less than their mean effect on reducing single-vehicle fatalities. This is plausible because ADASs' features are designed to be more effective in reducing single-vehicle accidents. That is, they can

[16] 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 as a robustness test of our findings and obtained slightly larger but less precise average treatment effects of ADASs. The loss in statistical precision arises because the estimator proposed by Wooldridge is less efficient than the simpler estimator that we primarily used. In any case, our quantitative estimates of the effects of ADASs on automobile safety appear to be consistent with those obtained by using the more sophisticated estimator.

[17] 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.

Table 3

Effects of ADASs on Driving

Dependent Variable	Vehicle Miles Traveled (1)	log(Vehicle Miles Traveled) (2)
ADAS dummy	−31.56	.04
	(138.07)	(.04)
Pseudo-R^2	.74	.77
N	6,565	6,464

Note. Heteroskedasticity-robust standard errors clustered by model year-make-model are in parentheses. Vehicle trims that are never equipped with advanced driver-assistance systems (ADASs) are excluded. All regressions include make-model-trim-calendar year and make-model-model year-calendar year fixed effects.

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 event—likely during our sample period—that the second vehicle is not equipped with ADAS-related features.[18]

Finally, as a check that our results do not simply reflect changes in driving intensity, we present the effects in Table 3 of ADASs' availability on VMT. We are unable to obtain statistically significant effects, but using the estimates in specification (2), we can rule out that the availability of ADAS-related technology will not reduce VMT by more than 4 percent, which implies that changes in VMT cannot explain the sizeable effects that ADAS availability has on both all accidents and fatal accidents.[19]

4.2. 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

[18] Importantly, we acknowledge that it is not clear why the estimate of ADASs' effect on multivehicle fatal accidents is less precise than the estimate of ADASs' effect on multivehicle nonfatal 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 accidents are systematically different from other types of accidents. A more reliable estimate of ADASs' effect on reducing fatalities in multivehicle accidents is likely to be possible in future research as the adoption of ADAS-equipped vehicles increases.

[19] 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 percent of all the vehicles.

effects of other important influences on auto fatalities that are available in the automobile safety literature. The estimates summarized in Table OA1 in the Online Appendix indicate that, all else being constant, the observed pattern of ADAS adoption would have resulted in a 1.89 percent 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 percent increase in fatalities, and the observed decrease in alcohol use while driving (as proxied by arrests) would have resulted in a 6.55 percent 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.[20]

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

4.3. Heterogeneity

In Figure 3, we show the extent that the effects of ADASs on accidents vary in any systematic way with vehicle characteristics, which could have implications for consumers' selective purchase behavior of vehicles when an ADAS is available.[21] Generally, we find little heterogeneity in the effects of ADASs in reducing accidents on the basis of vehicle size, price, and manufacturer nationality, with the puzzling exception that its availability appears to be more effective in lighter than heavier vehicles.[22] 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 ADASs on fatal accidents.[23]

The effectiveness of ADASs on accidents also may vary by driving behavior and conditions when an accident occurs. These findings also could reflect selec-

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

[21] In Figure 3, parameter estimates and 95 percent confidence intervals (formed from heteroskedasticity-robust standard errors clustered by model year-make-model) are truncated at 2 for clarity.

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

[23] We are unable to estimate precise heterogeneous effects of ADASs on fatalities for many of the vehicle type/weight/price/automaker categories, in all likelihood because of the infrequency of fatal accidents.

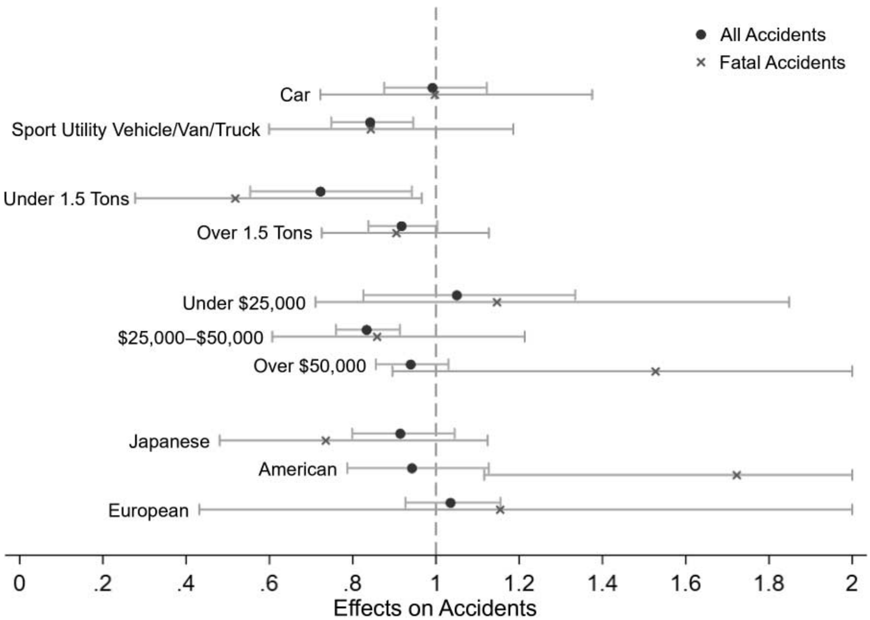

Figure 3. Heterogeneous effects of advanced driver-assistance systems on accident rates by vehicle type.

tivity to the extent that more risk-taking drivers tend to purchase vehicles with ADAS-related features because they are more likely than more risk-averse drivers to drive more dangerously and to drive in more dangerous conditions. An ADAS could therefore possibly offset risk-taking drivers' choices of how, when, and where to drive. Figure 4 shows, however, that there is no evidence of heterogeneity in the effects of ADASs on all accidents by speed of crash, roadway type, roadway conditions, weather conditions, and day or week or time of day. If anything, ADASs appear 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

We have assessed important potential sources of bias to our estimates of the effects of ADASs 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 ADASs have significant effects on reducing all accidents and fatal accidents: selection bias, offsetting behavior, and contamination of the control group.

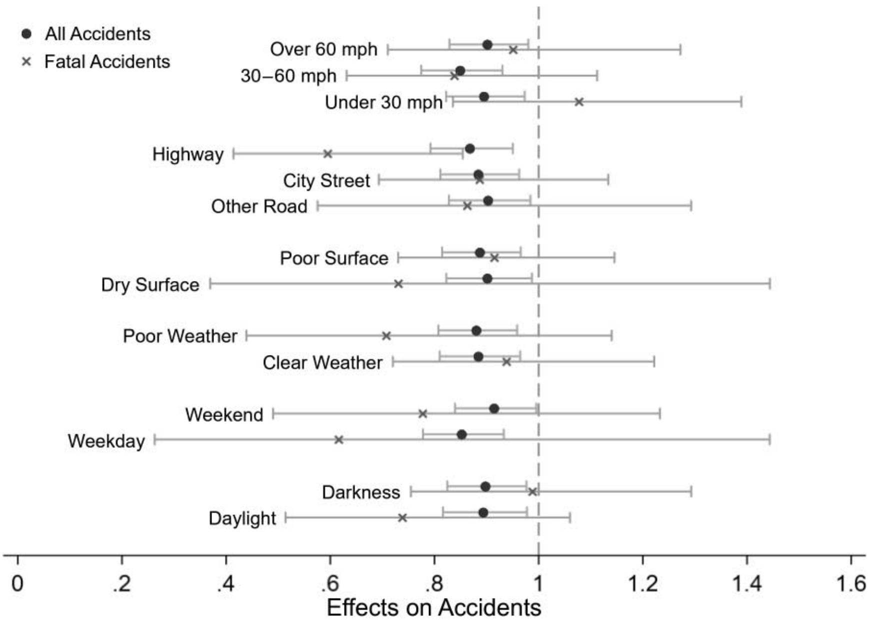

Figure 4. Heterogeneous effects of advanced driver-assistance systems on accident rates by accident conditions.

5.1. Selection Bias

We have stressed that drivers' decisions to self-select into treatment—that is, to drive a vehicle with an ADAS—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 vehicles with ADASs than were risk-taking drivers, then our estimates of the effects of ADASs on automobile safety would be biased upward. Conversely, our estimates of the effects of ADASs on automobile safety would be biased downward if risk-taking drivers were more likely to adopt ADASs than safer drivers. The latter behavior would be more relevant in the case of a safety feature like an ADAS that can compensate for a driver's risk-taking, instead of a safety feature like airbags that does not compensate for a driver's risk-taking 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

$$A_{yijt} = \exp\left(\sum_{\tau=-3,\ldots,4,\tau\neq 0} \beta^{\tau} \times 1(\tilde{y}_{ij} - y = \tau - 1) + \text{Controls} + \lambda_{ijt} + \lambda_{iyt} + \varepsilon_{yijt} \right), \quad (2)$$

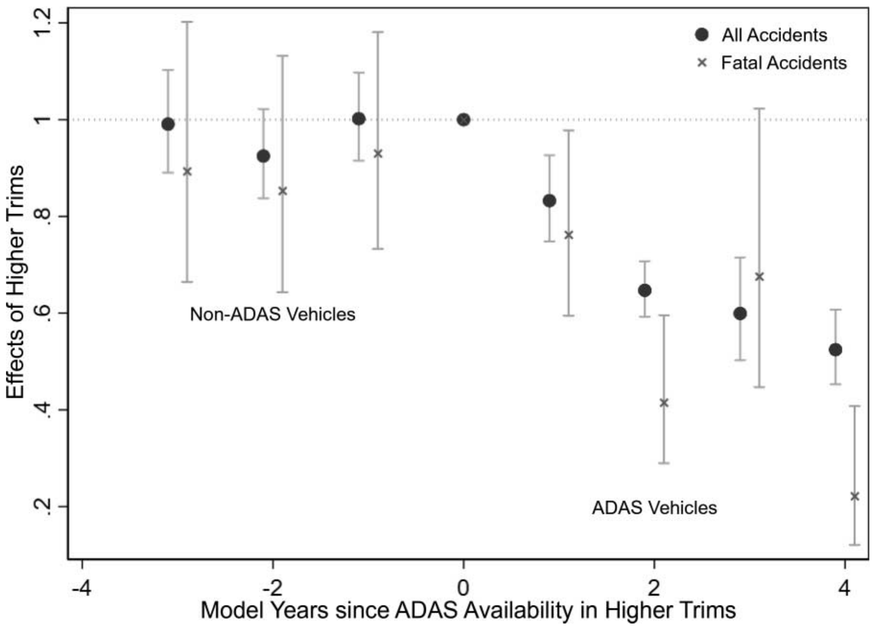

Figure 5. Event-study-style plot of the effects of advanced driver-assistance systems (ADASs) on accidents.

where $\tilde{y}_{ij}$ denotes the model year in which vehicle ij is first equipped with an ADAS and $1(\cdot)$ represents the indicator function. The coefficient β^{τ} corresponds to the effect of an ADAS in the τ vehicles equipped with an 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 (for example, 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 ADASs 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, that is, 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, which shows that the effectiveness of ADASs at reducing accidents and especially at reducing fatalities increases with the model years that ADAS-related technologies are available on higher trims.[24] There are two potential explanations for this pattern: ADASs reduce the prevalence of accidents by an amount that is quan-

[24] We refer to the plot as "event-study-style" because our data are 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 five observations to the estimation of the point with -1 model year because an ADAS was available in the higher trim calendar years 2014–18, but the 2016 Acura MDX contributes only three observations to the estimation of the point with 1 model year because an ADAS was available in higher trim calendar years 2016–18.

titatively consistent with the parameter estimates in column 2 of Table 2, which are based on our original specification in equation (1), or drivers systematically switch to ADAS-equipped trims only when they are made available, and they avoid higher level trims in earlier model years when ADASs were 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 ADASs. Those dimensions include non-ADAS vehicle safety features, such as side curtain and seat-mounted side impact airbags, as well as nonsafety features, such as a premium leather collection. In Section OA3 in the Online Appendix, we report a complete list of the non-ADAS safety features and non-safety-related trim features that were available for vehicles with high trims but not available for vehicles with low trims. That trim choice is influenced by more than just the availability of an ADAS lends credence to our first explanation that the pattern of results is credibly aligned with the estimates of the effect of ADASs on all accidents and fatal accidents.

We find that the effects of ADASs are smaller in the first model year when automakers introduce ADASs than after subsequent periods because over time automakers may improve the integration of ADAS-related technologies in a vehicle's safety performance or drivers might use the system's features more effectively or both. We cannot offer a more precise explanation because the estimates 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

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

where X_{yij} is a demographic characteristic of owners of vehicle yij.[25] Recall that we obtained this variable from a survey by Acxiom, which was conducted for a given calendar year and did not 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 an 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 ADASs 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.[26] Drivers who switch to ADAS-equipped trims immediately

[25] We continue to estimate the specification using Poisson regression because the demographic variables are nonnegative with a small number of values. Using Poisson regression in such instances has become common practice.

[26] 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 ADASs became available, but the fit of the prediction was poor. Independently, we

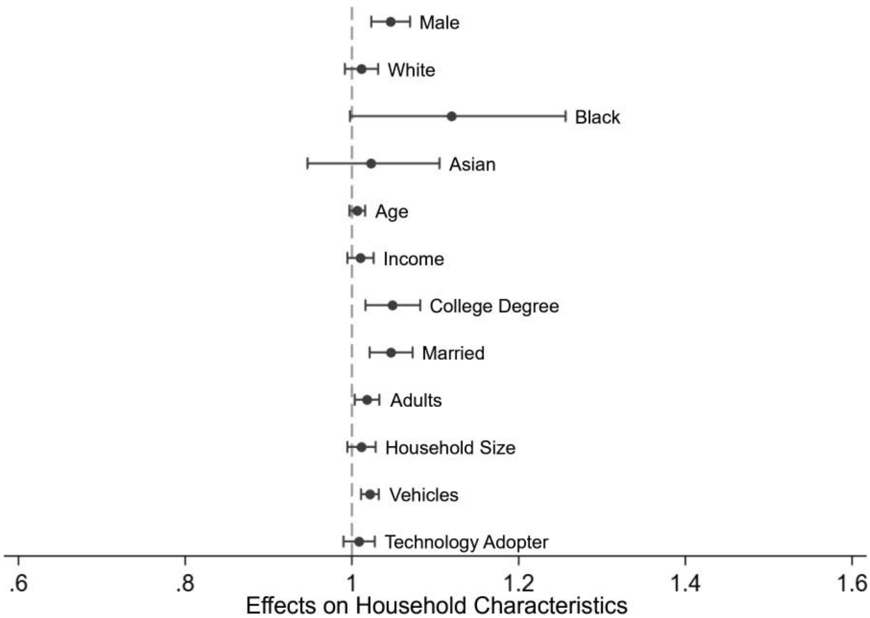

Figure 6. Effects of advanced driver-assistance systems by household characteristic

after an ADAS becomes available 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 an ADAS becomes available, the effects for virtually any characteristic are very small (less than a 5 percent change).[27] In sum, the evidence bolsters our claim that the findings of ADASs' 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 ADASs 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 are at similar levels. As we would expect, given the effectiveness of ADASs, the crash rate of drivers who never switched into ADAS-equipped vehicles is somewhat

did not find any increase in the share of higher-trim-level vehicles within a model year-make-model when an ADAS was made available.

[27] The only exception is that Black drivers adopt ADAS-equipped trims when they are first made available at close to a 20 percent change. But 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.

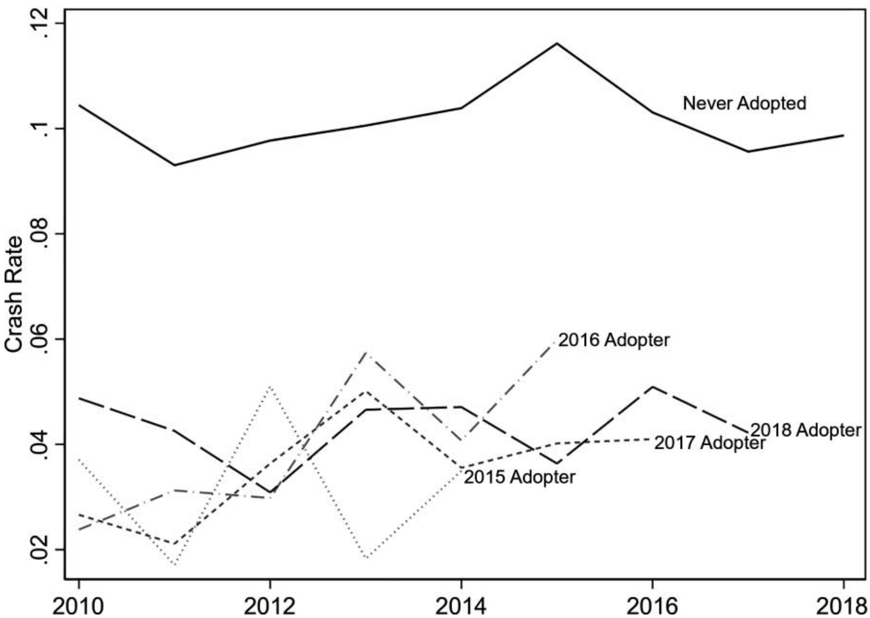

Figure 7. Crash rate by household advanced driver-assistance system adoption year

higher than the crash rate of drivers who switched into ADAS-equipped vehicles at some point during our sample period.

5.2. *Offsetting Behavior*

A second potential source of bias to our estimates is that the adoption of an ADAS might affect a driver's behavior on the road. For example, a driver with an ADAS-equipped vehicle might take more risks while driving, like texting and paying less attention to traffic conditions, which would offset the safety benefits of an ADAS.[28] Alternatively, because ADAS-related features include auditory and visual warnings to drivers when other vehicles are approaching, ADASs may induce drivers to make safety-augmenting responses. In any case, given that our interest is to estimate the effect of ADASs 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 ADASs 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 ADASs by vehicle characteristics did not

[28] Peltzman (1975) and Winston, Maheshri, and Mannering (2006) find evidence that drivers' offsetting behavior reduces the overall safety benefits of seatbelts and of airbags and antilock brakes, respectively.

suggest that drivers' risk preferences led them to systematically change their behavior in response to adopting ADASs.

5.3. Contamination and Externalities

A final potential source of bias could be 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 ADASs, also may improve the safety of untreated vehicles and cause an estimate of the effectiveness of ADAS-related 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 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 an ADAS available as an option at the time of manufacture.[29] Second, nearly 50 percent of the fatal accidents in our sample were single-vehicle accidents.

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 ADASs 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 multi-vehicle 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 because drivers' behavior in response to the safety feature must be taken into account. We have addressed this issue empirically in the context of automakers' introduction of ADAS-related safety features. We have presented causal evidence that ADASs have 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 vehicles. We also have tested for the possibility that our findings 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 ADASs have measurably improved automobile safety, for two reasons it is

[29] Slightly more than 25 percent of all the vehicle models in our sample have ADASs, while the share of the total number of vehicles on the road that have ADASs is much smaller.

premature to conclude that policymakers should require automakers to equip all new passenger cars and light trucks with ADASs or even just AEB systems, which is required by fall 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 ADASs or even AEB systems is small but increasing. Thus, our evidence of some 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 ADASs in all vehicles. Other potentially informative evidence from insurance markets is not yet available.[30] Looking further ahead, we hope 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 transportation technologies that use artificial intelligence (Winston and Karpilow 2020; Winston, Yan, and Associates 2024).

References

af Wåhlberg, 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* 88:217–29.

Anderson, Michael L., and Maximilian Auffhammer. 2014. Pounds That Kill: The External Costs of Vehicle Weight. *Review of Economic Studies* 81:535–71.

Blanco, Myra, Jon Atwood, Sheldon Russell, Timothy Trimble, Julie McClafferty, and Miguel Perez. 2016. *Automated Vehicle Crash Rate Comparison Using Naturalistic Data.* Virginia Tech Transportation Institute, Blacksburg.

Cameron, A. Colin, and Pravin K. Trivedi. 1998. *Regression Analysis of Count Data*. Cambridge University Press.

Edlin, Aaron S., and Pinar Karaca-Mandic. 2006. The Accident Externality from Driving. *Journal of Political Economy* 114:931–55.

Gourieroux, C., A. Monfort, and A. Trognon. 1984. Pseudo Maximum Likelihood Methods: Applications to Poisson Models. *Econometrica* 52:701–20.

Haghani, Milad, and Michiel C. J. Bliemer. 2023. Emerging Trends and Influential Outsiders of Transportation Science. *Transportation Letters* 15:386–422.

Maheshri, Vikram, and Clifford Winston. 2024. The Effect of Reductions in Vehicle Miles Traveled on Highway Fatalities and Congestion with Heterogeneous Motorists. Pp. 11–25 in *Revitalizing a Nation: Competition and Innovation in the US Transportation System* by Clifford Winston, Jia Yan, and Associates. Brookings Institution Press.

———. 2025. Methodological Challenges in Automobile Safety Research and Opportunities as We Enter the Autonomous Era. Working paper. University of Houston, Houston.

[30] We conducted some interviews with insurance industry personnel to see if the adoption of ADASs 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 composed of a large share of ADAS-equipped vehicles, it would be useful to estimate the effect of the staggered adoption of ADAS-equipped vehicles on the nation's insurance costs, which will reflect a tradeoff between the potentially lower claims caused by ADASs' reduction of accidents and the potentially higher claims caused by ADASs' increase in the cost of a car and repairs.

Mosquet, Xavier, Michelle Andersen, and Aakash Arora. 2016. A Roadmap to Safer Driving through Advanced Driver Assistance Systems. *Auto Tech Review* 5:20–25.

Peltzman, Sam. 1975. The Effects of Automobile Safety Regulation. *Journal of Political Economy* 83:677–725.

Slovick, Murray. 2020. How Technology Is Driving the Democratization of ADAS. Electronic Design, March 17. https://www.electronicdesign.com/markets/automotive/article/21126132/how-technology-is-driving-the-democratization-of-adas.

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

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* 32:83–99.

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

Wooldridge, Jeffrey M. 1999. Distribution-Free Estimation of Some Nonlinear Panel Data Models. *Journal of Econometrics* 90:77–97.

———. 2023. Simple Approaches to Nonlinear Difference-in-Differences with Panel Data. *Econometrics Journal* 26:C31–C66.