A CRITIQUE OF AUTOMOBILE SAFETY RESEARCH METHODOLOGIES: IMPLICATONS FOR POLICYMAKERS AS WE ENTER THE AUTONOMOUS ERA

Vikram Maheshri
University of Houston

Clifford Winston*
Brookings Institution

Abstract. We argue in this paper that the established and varied research on highway safety has significant limitations that prevent it from either identifying causal influences on automobile accidents or from asking the most important policy-oriented questions about auto safety. We develop a theoretical framework to clarify the limitations of the three main approaches taken in automobile safety research: the use of controlled environments, disaggregated data, and aggregated data. We illustrate the limitations in the context of the vast empirical literature that has sought to assess the effectiveness of seatbelt use in reducing fatal accidents. We conclude by drawing important lessons for researchers and policymakers given the state of automobile safety research.

May 2025

***We are grateful to Jia Yan and Fred Mannering for helpful comments.**

1. Introduction

US policymakers and automakers have long prioritized improving automobile safety to reduce accidents and fatalities. Policymakers have spent hundreds of billions of dollars to improve the safety of roadways, expand and modernize traffic enforcement, and conduct public safety campaigns to discourage speeding and driving under the influence of alcohol and drugs. Automakers have strengthened vehicles' structures to provide greater resistance to crashes, improved braking and steering, installed occupant safety devices, and recently begun to make autonomous vehicle safety features available. As a result of the safety improvements in the driving environment and automobiles, the US experienced a sustained 3% annual decline, on average, in the rate of automobile fatalities from 1920 to 2010 (see figure 1). But since 2010, the US has not reduced the rate of automobile fatalities while experiencing a slight increase since COVID began in 2020.[1]

Although it is reasonable to expect that the US would not maintain the same rate of highway safety improvement indefinitely, the recent lack of improvement is cause for concern because the annual cost of automobile accidents, including fatalities, the loss of quality of life, vehicle damage, and disruption of road travel currently amounts to more than $1 trillion annually (TRIP, 2023). Such a large social cost should encourage the research community, including economists and transportation engineers, to improve our understanding of the causes of highway accidents and to provide evidence of the potential for new public policies to reduce this cost by addressing those causes. However, we argue in this paper that established and varied research approaches to study highway safety face significant limitations that prevent them from identifying causal influences on automobile accidents or from asking the most important questions about auto safety. Consequently, these flaws in automobile safety research are substantive obstacles to recommending new, constructive policies to policymakers.

The fundamental challenges in safety research arise from the fact that driving is a highly complex behavior that is comprised of many smaller yet interrelated decisions unfolding over time. Specific to safety research, observed outcomes (accidents) are caused by multiple decisions made by drivers, some of which are made even before they begin their trip, including where, when, under what road conditions, in which vehicle, and in what physical state they choose to drive. Other

[1] Estimates of the rate of change of automobile fatalities in the US from 1920 to 2010 and from 2010 to 2023 are based on annual fatalities data from the US Department of Transportation.

influential decisions on accidents are made by drivers during the trip, such as the speed and aggressiveness while driving. Those decisions in turn have countless determinants that are potentially highly correlated to one another and mostly unobservable, with unobserved risk preferences being the most important influence because they sort drivers into riskier or safer driving environments and behaviors.

Safety research consisting of analyses of data generated from controlled environments, disaggregated data obtained from police accident reports, and aggregated data from various administrative and private sources is unable to account fully for the various complex decisions and influences that lead to accidents. Each of these approaches faces distinct limitations that depend on the specifics of the research question, features of the data, and institutional details of the research settings. The controlled environment and disaggregate approaches have generally not been able to obtain credible causal explanations of the determinants of highway accidents and accurate estimates of the potential effects of government safety policies to reduce accidents. The aggregate approach can obtain credible causal explanations of the effects of certain policies, but it requires quasi-experimental variation that is not always available; hence, it can address only a limited set of policy questions.

In what follows, we propose a general, dynamic theoretical framework to analyze drivers' behavior and to identify the factors that contribute to the risk of driving and the probability of an accident. This framework encompasses essentially all of the automobile safety questions that are relevant to policymakers, so it can be used to compare the various empirical approaches that have been used by researchers and to clarify the assumptions that underly each of them. In many cases, those assumptions are implausible, or the empirical findings from the approaches do not coincide with the effects that researchers and policymakers are actually interested in estimating. As an illustrative exercise, we apply the framework to clarify the weaknesses of current safety research methodologies in the context of answering an important, long-studied, and policy-relevant issue: How effective are seatbelts in reducing fatal accidents?

Taking an optimistic look to the future, we stress that autonomous vehicles have the potential to make enormous improvements in safety that could eventually eliminate vehicle accidents by eliminating the risks posed by the nation's most dangerous drivers (Winston, Yan, and Associates, 2024). Importantly, we describe an aggregate approach that can accurately estimate the reduction in automobile accidents and fatalities that has been caused by even low levels of autonomous safety technologies (Maheshri, Winston, and Yu, 2025). In addition, as fully autonomous vehicle technologies are being adopted and if self-selection into autonomous vehicle technology adoption can be credibly analyzed, it will be possible to identify the causal effect of autonomous vehicles on accidents using disaggregated data because the technology makes the driving decisions that eliminate the influence of drivers' unobserved risk preferences on safety.

We conclude by drawing important lessons for the research community and policymakers. First, because the disaggregated approach using data from police accident reports has become the dominant approach for analyzing automobile safety, the broader research community must come to terms with the fact that it is unable to obtain causal estimates of the determinants of automobile accidents that are useful for safety policy. It will undoubtedly be difficult for the generations of safety researchers who have taken the disaggregated approach to accept its fatal flaws. Thus, it will take a new generation of transportation researchers to replace that approach with one that can produce causal and policy-relevant estimates of the determinants of automobile accidents. As Max Planck, the originator of the quantum theory in physics has said, "science makes progress funeral by funeral."

Second, as the United States adjusts to a new environment of autonomous vehicle technologies, policymakers should be aware of the tendency for flawed approaches to produce inflated estimates of their effectiveness when developing policy. Importantly, they should avoid past mistakes by not rushing to seize on a misguided opportunity to improve highway safety by prematurely requiring automakers to install autonomous safety features in their new vehicles on the basis of flawed or speculative estimates of the direct and external benefits of those features.

2. A Theoretical Framework to Study Driving Behavior

We develop a theoretical framework to understand driving behavior in order to frame and formalize our critique of empirical analyses of the determinants of an automobile accident. Driving is a complex dynamic activity because many agents (drivers) continuously update their positions on a road, which changes their exposure to other agents on the road and to roadway conditions. Previous work, for example, Tscharaktschiew (2020) and Yang et al. (2015), has modeled drivers' speed choices in a non-cooperative setting to obtain a traffic network equilibrium. We consider the simple case in which a fixed set of n drivers travel along the same highway route and choose their travel speed. The speed affects not only drivers' (future) position along the road, but it also affects the likelihood of a potentially significant cost; that is, the probability and severity of being involved in an automobile accident.

The simple case that we analyze here raises a host of empirical issues that are not addressed in the empirical literature, such as treating route choice as an endogenous dynamic decision and endogenizing pre-trip decisions such as the choice of vehicle or time of day to drive. Adding further real-world complexity to our framework, such as motorists' choices of the extent and level of insurance coverage and how those choices may affect driving behavior, only exacerbates the problems with current empirical approaches to automobile safety.

We formulate driving behavior as a dynamic optimization problem where the state variable is the position of all drivers on the road. We denote the position of each driver i at time t as p_{it}, which we coalesce into an $n \times 1$ vector P_t. (Hereafter, upper-case variables correspond to the n vector of their lower case scalar counterparts.) The choice variable of each driver i is the travel speed at time t, which we denote as s_{it}. We characterize the driver's expected utility optimization problem by the following Bellman equation:

$$V_i(P_t) = \max_{s_{it}} \left\{ \begin{array}{c} \underbrace{u_i(p_{it} + s_{it}) - \sum_{j=1}^{J} E_{it}\left[\pi_{it}^{j}(P_t, S_t, X_t)\right] \times c_{ij}}_{\text{flow utility}} + \\ \underbrace{\left(1 - \sum_{j=1}^{J} E_{it}\left[\pi_{it}^{j}(P_t, S_t, X_t)\right]\right)}_{\text{Prob. of not getting in accident}} \cdot \underbrace{V_i(E_{it}[P_{t+1}(S_t)])}_{\text{Continuation Value}} \end{array} \right\}, \tag{1}$$

where u_i is driver i's utility function, which increases with their position along the road, and π_{it}^{j} is the probability that driver i gets into an accident of severity j at time t, with severity ranging from vehicle damage only to a driver fatality. This probability is a function of the positions p and speeds s of all the drivers on the road as well as the relevant characteristics of the drivers and vehicles on the road and the roadway conditions at time t. We coalesce those variables in the matrix X_t. Finally, c_{ij} is the cost to driver i of getting in an accident of severity j.

The instantaneous flow utility in the equation that accrues to a driver is given by the utility from travelling an additional distance of s_{it} between periods t and $t + 1$ net of the expected costs of getting into an accident during that time. Conditional on not getting into an accident between periods t and $t + 1$, driver i obtains a continuation value given their beliefs of where all other vehicles will be on the road in period $t + 1$.

This stylized formulation of a driver's dynamic optimization problem captures three important and plausible features of driving that have critical implications for analyses of automobile safety: (1) Drivers form expectations of their safety based on the driving environment and their speed choices (π_{it}^{j} is a function of s_{it} and X_t); (2) Drivers form expectations of where other drivers will be, implicitly taking into account that those drivers face their own optimization decisions (π_{it}^{j} is a function of P_t and S_t); and (3) Drivers understand the decisions they make at any point in time may influence the decisions of other drivers at future times (P_{t+1} is a function of S_t). Because drivers' expectations may affect their likelihood of getting into an accident, those expectations must be accounted for by researchers if they wish to explain the determinants of accidents empirically.

The complexity of modelling a driver's problem grows because each driver faces her own analogous optimization problem in each period *t*. Hence, driving can be thought of as a dynamic game of incomplete information. The perfect Bayesian equilibrium of this game consists of a series of strategies, or mappings from the state space to the action space, which we denote as $S_{it}^{*}(P_t)$, accompanied by a specification of driver beliefs that satisfy Bayes' Law. We therefore obtain the corresponding equilibrium accident probability functions $\pi_{it}^{j*}(P_t, X_t) = \pi_{it}^{j}(P_t, S_{it}^{*}(P_t), X_t)$.

3. Using the Framework to Analyze the Determinants of Automobile Safety

We have not seen previous empirical analyses of automobile safety framed by a dynamic model of drivers' behavior. In any case, the major questions of interest to automobile safety researchers can effectively be distilled into questions regarding the determinants of the accident probability, π_{it}^{j*}, in our framework. To connect this probability to several standard questions in the literature, we explicitly define the arguments in the matrix $X_t = (D_t, Z_t, R_t)$, where D_t is a matrix of driver socioeconomic characteristics, such as age and gender, Z_t is a matrix of vehicle characteristics, such as weight and horsepower, and R_t is a matrix of roadway characteristics, such as pavement condition and curvature.

Thus, the literature that seeks to inform policymakers and automakers by explaining how vehicle attributes affect auto safety, estimates $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$ for driver i and an attribute z_i of her vehicle. The literature that seeks to inform policymakers by identifying the characteristics of drivers that contribute to risk, estimates $\frac{\partial \pi_{it}^{j*}}{\partial d_i}$ for a driver characteristic d_i. Finally, the literature that seeks to inform highway engineers by identifying how roadway conditions affect accident risk, estimates $\frac{\partial \pi_{it}^{j*}}{\partial r_t}$ for a roadway condition r_t.

Following standard econometric practice, researchers would estimate an average of those effects on accident risk over multiple configurations of different drivers and vehicles on the road during different time periods, thereby obtaining a treatment effect that does not vary over time t and is not a function of drivers' positions on the road P_t. This choice of aggregation raises immediate concerns of how the estimates of interest can be identified because P_t is a driver's state variable in her dynamic optimization problem and it *does* affect accident risk. This relationship constitutes a potential source of endogeneity that must be addressed to obtain consistent estimates of important influences on safety. We discuss this problem and others in the context of the different empirical approaches that have been taken in the safety literature.

4. Using the Framework to Assess Empirical Approaches in the Safety Literature

Researchers have taken three different empirical approaches to estimate the determinants of an accident probability: (1) a controlled environment approach that generates empirical observations from simulated accidents; (2) a disaggregate approach based on accident data generated by individual drivers and included in police accident reports; and (3) an aggregate approach based on accident data generated by travelers and aggregated to a geographic level, such as a state. We clarify the identification problems that cause the controlled environment and disaggregate approaches to produce biased estimates. It does not appear that tractable methods exist at this time to circumvent the bias in those approaches. It is possible to circumvent the bias by taking an aggregate approach that limits the types of questions about accident safety that researchers can address.

We illustrate the limitations of each approach in the context of one of the most studied questions in the automobile safety literature: To what extent does wearing seatbelts reduce driving fatalities?

The Controlled Environment Approach

The controlled environment approach refers to a research design where researchers subject specific vehicles to simulated driving conditions and observe specific aspects of their safety performance. Crash tests and closed course observations are well-known examples of controlled environment approaches. Although policymakers are partial to this approach, in all likelihood because it resembles randomized controlled trials, which are broadly recognized as the gold standard of causal research (Kahane (2015)), the approach is far from being randomized.

To illustrate the lack of randomness, consider that a series of carefully conducted crash tests showed that the use of seatbelts reduced the risk of a fatal accident by 45% (e.g., Lave and Weber (1970)). Assuming that seatbelts are the vehicle attribute z_i of interest and that their effectiveness is determined by drivers' behavior to wear them, this finding corresponds to $\frac{\partial \pi_{it}^{j}}{\partial z_i}$ evaluated at a particular value of vehicles' positions, drivers' and vehicles' characteristics, and roadway conditions; that is, (P, X) corresponding to the details of the tests.[2] But this calculation differs from the calculation of interest in actual driving environments, $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$, which is determined as the equilibrium choices of driver i that are evaluated at the equilibrium levels of (P^*, X^*). So, for instance, if safer drivers were more likely to wear seatbelts, then the calculation of $\frac{\partial \pi_{it}^{j}}{\partial z_i}$ based on the crash test would overestimate $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$.[3] Alternatively, if wearing seatbelts make drivers feel safer and thus more willing to drive in potentially hazardous road conditions, such as during a snowstorm, then $\frac{\partial \pi_{it}^{j}}{\partial z_i}$ would underestimate $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$.[4]

Of course, the potential selection bias from individuals' choice of seatbelt use is generally well known and some researchers may feel that it merits only a qualification. But we have shown that the bias is likely to be quite serious by showing how it arises in a plausible dynamic model of drivers' behavior and by clarifying that drivers' safety outcomes are based on a series of decisions that they make prior to and while driving. As pointed out, the decision they make while driving is their choice of speed s_{it} at time *t*, which determines their position p_{it} on the road at time *t*. The

---

[2] The 45% risk reduction technically corresponds to $\frac{\partial \log \pi_{it}^{j}}{\partial \log z_i}$. This elasticity can be recovered from the marginal effect $\frac{\partial \pi_{it}^{j}}{\partial z_i}$.

[3] Descriptive correlational evidence indicates that drivers who use seatbelts are less likely to engage in risky behaviors like speeding or impaired driving, based on observational data and crash statistics. See the "National Occupant Protection Use Survey" published as "Seat Belt Use in 2019—Overall Results" (Report No. DOT HS 812 821).

[4] In this case, drivers' behavior would be consistent with Peltzman's (1975) and Wilde's (1982) risk compensation hypothesis. Winston, Maheshri, and Mannering (2006) found that motorists' increase in risky driving behavior appeared to offset the technological effectiveness of airbags.

decisions they make prior to driving include the type of vehicle to buy, which determines vehicle characteristics Z_t, the kinds of behaviors to engage in, which determines driver characteristics D_t, and the roads they will traverse and when they will travel, which determines roadway characteristics R_t.

By incorporating those considerations in a model of driver behavior that may result in accidents, it becomes clear that estimates of the effect of seatbelts on driving fatalities based on the results of a crash test would be applicable to actual highway driving conditions only if it could be assumed that seatbelt use was determined independently of all the decision variables in our framework. Such an assumption is implausible and cannot be ignored by a qualification.

The Disaggregate Approach

Researchers have attempted to circumvent one shortcoming of the controlled environment approach by using disaggregated data obtained from accidents involving actual drivers; thus, the disaggregate approach refers to a research design where researchers use observational, driver-level data to estimate the effects of various highway and vehicle characteristics and safety policies on accidents.

Our theoretical framework characterizes the behavior of all drivers regardless of whether they are involved in an accident. In addition, we do not assume that drivers who are involved in accidents do not differ from drivers who are not involved in accidents in terms of observed and unobserved influences on accidents. The determinants of the decisions that drivers make on the road are also correlated to observed influences on accidents, including speed choice, vehicle characteristics, some driver characteristics, and roadway conditions for their trip. Important examples of determinants of accidents that are not proxied, measured, or observed by the researcher are unobserved characteristics of the driver, such as their temperament and judgment.[5]

The primary weakness of the disaggregate approach is that because researchers obtain data from police accident reports, they are forced to make the implausible assumption that drivers who are involved in accidents do *not* differ from drivers who are not involved in accidents in order to attach external validity to their results. This assumption merits more than a qualification and is much stronger than researchers realize. That is, it is assumed that the decisions drivers make prior to and while driving and their unobserved characteristics do not have different effects on drivers

[5] It could be argued that researchers use the number of motorists' speeding tickets as a proxy for risk preferences. For example, see Vertlib et al. (2023).

who get into accidents and on drivers who do not get into accidents. However, drivers effectively self-select to be included in accident reports by being involved in an accident; otherwise, they are not included in those reports.

Similar to researchers who take a controlled environment approach, researchers who take a disaggregate approach do not obtain findings based on a random sample. Researchers rarely acknowledge this problem and implicitly attempt to deal with it by effectively comparing drivers who get into accidents of different severities (i.e., fatality, serious injury, minor injury, or property damage only). But to avoid biasing parameter estimates, researchers must estimate an effect that captures both the marginal effect of getting into an accident and the conditional effect of the severity of that accident. Even if researchers use sophisticated econometric methodologies to, for example, control for motorists' heterogenous behavior, estimating the determinants of the severity of an accident conditional on an accident occurring is simply unable to address the fundamental identification problems that we have stressed here.

Researchers have taken different empirical approaches to estimate disaggregate models, but they all fundamentally estimate the probability of getting in an accident of severity $j > 1$ conditional on getting in an accident of any severity (denoted as π_{it}^{*}), where j = 1, if no accident occurred. We can write this probability as:

$$\pi_{it}^{j**} = \frac{\pi_{it}^{j*}}{\pi_{it}^{*}} \tag{2}$$

It is clear that $\pi_{it}^{j**} \neq \pi_{it}^{j*}$. Importantly, this fact dramatically diminishes the value of the entire empirical exercise if its ultimate purpose is to inform policymakers how highway safety could be improved.

Consider, for example, the policy question of the causal effect of seatbelt use on automotive fatalities. Denote seatbelt use with the binary variable z_i (where 1 indicates the use of a seatbelt; 0 otherwise). The effect that is identified in a disaggregate analysis can then be written as:

$$\text{Disaggregate Effect} = \left(\pi_{it}^{j*} | z_i = 1, A_i = 1\right) - \left(\pi_{it}^{*} | z_i = 0, A_i = 1\right), \tag{3}$$

where the binary variable A_i is equal to 1 if the vehicle got in an accident. Policymakers are interested in what we call the true causal effect (TCE) of seatbelt use on fatalities, which can be expressed as:

$$\begin{aligned} \text{TCE} &= (\pi_{it}^{*} | z_i = 1) - (\pi_{it}^{*} | z_i = 0) \\ &= \left(\pi_{it}^{j*} | z_i = 1, A_i = 1\right) \times (\pi_{it}^{*} | z_i = 1) - \left(\pi_{it}^{j*} | z_i = 0, A_i = 1\right) \times (\pi_{it}^{*} | z_i = 0) \end{aligned} \tag{4}$$

Even if we could perfectly estimate the probability of getting in an accident (π_{it}^{*}), we would be unable to use the estimates of the disaggregate causal effect in equation (3) to obtain the TCE unless we made the strong additional assumption that $(\pi_{it}^{*}|z_i = 1) = (\pi_{it}^{*}|z_i = 0)$. This assumption is highly implausible because it states that the likelihood that a person who wears a seatbelt gets in an accident is the same as the likelihood that a person who does not wear a seatbelt gets in an accident, which ignores that a person's propensity to wear a seatbelt is correlated to their attitude toward risk and, in turn, to their driving behavior. The assumption is further weakened because some drivers may adjust their behavior if they are wearing a seatbelt.

Although we have shown how the identification problem prevents one from determining the true causal effect on safety when the treatment variable is discrete, the same identification problem extends to the case when the treatment variable is continuous. For example, the problem arises in Anderson and Auffhammer (2014), where the treatment variable z_i corresponds to the curb weight of the vehicle.

Possible Responses to Justify the Disaggregate Approach. There are a number of possible responses to justify a disaggregate estimation approach. First, it could be argued that the identification problem is mitigated if estimates of the effect of seatbelt use, for example, on injury severity, $\frac{\partial \pi_{it}^{j**}}{\partial z_i}$, could be interpreted as proxy estimates of $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$ because they were obtained from empirical models that are insensitive to the inclusion of additional control variables or because they are based on plausibly exogenous instrumental variables for seatbelt use z_i.

However, this argument obscures but does not address the fundamental identification issue. Note that differentiation of equation (2) yields:

$$\frac{\partial \pi_{it}^{j**}}{\partial z_i} = \frac{1}{\pi_{it}^{*}}\left(\frac{\partial \pi_{it}^{j*}}{\partial z_i} - \pi_{it}^{j**}\,\frac{\partial \pi_{it}^{*}}{\partial z_i}\right). \tag{5}$$

Even though π_{it}^{*} is observable, equation (5) implies that $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$ can be recovered from an estimate of $\frac{\partial \pi_{it}^{j**}}{\partial z_i}$ only if $\frac{\partial \pi_{it}^{j*}}{\partial z_i} = 0$, which is the same identifying assumption indicated above. This relationship is implausible because it is difficult to believe that *any* vehicle attribute or driver/roadway attribute would affect the unconditional probability of getting in a severe accident without affecting the probability of getting in any accident. Indeed, our theoretical framework shows that, in general, drivers' decisions will influence the (unconditional) probabilities of getting in accidents of all types of severity.

Second, our framework, which stresses that drivers make many endogenous decisions before and during their trip, reveals that disaggregate approaches cannot be used to obtain consistent estimates of the determinants of the *marginal* probability π_{it}^{*}. This is a critical limitation for two reasons. First, explaining the probability of getting in any kind of accident is one of the most important objects of interest to policymakers. Second, it is not possible to use the marginal probability of getting in an accident as a selection equation to obtain consistent estimates of the determinants of accident severity, which do not suffer from selectivity bias.[6]

Finally, it could be argued that an estimate of the *conditional* effect of any determinant of safety on reducing severe or fatal accidents (conditional on any accident occurring) is interesting in its own right. But the estimate will still suffer from endogeneity bias that cannot be addressed using disaggregated data. To see this in the case of seatbelts, suppose there was some confounding cause of accidents that was unobservable and correlated with seatbelt use, such as whether the driver was extremely distracted and delayed fastening her seatbelt. Then researchers would need to block the pathways from this confounding variable to both π_{it}^{j**} and π_{it}^{*}. That is, by not being able to study π_{it}^{j*} directly, researchers would need to make an additional identifying assumption. In this example, the assumption would be that being extremely distracted does not affect the likelihood of getting in a fatal accident, even if it affected the likelihood of getting in any type of accident. This is not only an implausible assumption, but it appears that researchers taking a disaggregate approach to estimate conditional probabilities may not even be aware that they are making it.

Why Have Transportation Researchers Continued to Use the Disaggregated Approach? Notwithstanding the significant shortcomings of the disaggregated approach, a large body of transportation science researchers continue to estimate disaggregated models of the determinants of highway safety and publish well-cited papers in academic journals, including but not limited to *Accident Analysis and Prevention* and *Analytic Methods in Accident Research* (Haghani and

---

[6] Eluru and Bhat (2010) effectively take a selectivity approach by jointly modelling seatbelt use and accident severity. In this approach, the endogenous treatment of seatbelt use takes the role of a selection equation to reduce the biased parameter estimate of seatbelt use in the accident severity equation. But the authors do not have clean variation in seatbelt use that is uncorrelated to the determinants of accident severity. Thus, identification is achieved by the choice of functional form, which does not address the fundamental endogeneity problem in their model.

Bliemer, 2022, Zoe, Vu, and Huang, 2020). We argue that transportation researchers would be more aware that the disaggregated approach to analyzing automobile safety is not useful and would be more likely to abandon it if the ultimate goal of this research program were to produce scholarly research that could potentially inform policymakers' efforts to improve highway safety. As the research program currently stands, it effectively consists of demonstration papers that use different econometric methods and data sets to obtain parameter estimates, but that do not reach any substantive conclusions that have accumulated and can guide new safety policies or policy reforms.

As economists, we are certainly aware of researchers' fascination with methodologies at the expense of any concerns with substantive issues. We also are aware that policy issues have provided "reality checks" for some methodologies and have led to the death of certain research programs. Hopefully, non-economists will find the following examples to be informative.

Winston (2021) discusses a body of research in industrial organization during the 1960s and 1970s that estimated regressions of the effect of industry concentration on profits with the goal of justifying an active deconcentration policy by the antitrust authorities to reduce excess profits and consumer welfare losses. Much of the research attempted to improve econometric specifications, use better data, and provide a stronger theoretical underpinning for the regressions that were estimated. The academic debate eventually crystalized between University of Chicago economists, who wanted a less aggressive antitrust policy and Harvard University economists, who wanted a more aggressive antitrust policy. Importantly, the academic policy debate also attracted the attention of policymakers.

Demsetz (1973) provided the decisive contribution to the debate by arguing that the source of high concentration in an industry's output could derive only from a few firms' superiority in producing and marketing products. Methodologically, Demsetz argued that concentration-profits regressions suffered from an identification problem that obscured the fact that some firms are simply much more efficient than other firms and either have lower costs or are able to offer better products or both. Of course, as Demsetz noted, it is not easy to ascertain just why General Motors and IBM (times change!) perform better than their competitors, but their profits do not arise from reducing output or colluding. Demsetz's idea, reinforced by the work of other Chicago scholars, contributed to a profound shift to a less aggressive antitrust policy that persists today. In addition, the research program of estimating concentration-profits regressions was abandoned and new generations of researchers developed new empirical approaches toward studying industry competition and concentration.

During the 1960s, macroeconomists estimated Keynesian econometric models of the US economy with the goal of identifying fiscal and monetary policies that could improve the US economy's performance. The research coincided with methodological and computational developments in econometrics that enabled simultaneous equation models of the macroeconomy to be estimated by three-stage-least squares and full information maximum likelihood methods. The policy debate crystalized between freshwater schools, including the University of Chicago and other schools located near the Great Lakes, which supported a less active monetary and fiscal policy, and saltwater schools, including Harvard and MIT located along the east coast, which supported a more active macro policy. Again, much of the research focused on improving the econometric specifications, data, and methodologies.

Lucas (1976) provided a major contribution to the debate with his famous critique, which argued that Keynesian macroeconometric models were fundamentally flawed because the estimated parameters that were used in policy simulations of the effects of changes in macro policy would themselves change in the new policy regime. Thus, estimates of the effects of stabilization policies were biased and econometric models of the US economy were not a reliable tool to help guide macroeconomic policy reform. Of course, debates about appropriate macroeconomic policy continue today, but estimation of Keynesian simultaneous equations models has been abandoned and replaced by new empirical and theoretical methodologies.

These episodes in economics of evaluating the usefulness of and subsequently abandoning certain methodologies suggest that transportation science researchers have been able to avoid careful scrutiny of their work using the disaggregated approach because their specific empirical findings have not been used to guide policy. Importantly, this research has not led to any intense policy debates with potentially significant stakes.

If the empirical findings were used to guide policy, both policymakers and researchers would have considerable incentives to assess the credibility of the research on substantive grounds and would undoubtedly abandon disaggregate models because they could not produce credible causal findings that were useful for policy. Because that denouement has yet to occur, researchers have continued to use the disaggregated approach to study the severity of accidents drawing on econometric methodologies to analyze discrete data and on non-econometric techniques.

Most of the research uses econometric methodologies to analyze discrete data to estimate a severity model given by:

$$S_{kn} = \beta'_k x_{kn} + \varepsilon_{kn}\ , \tag{6}$$

where S_{kn} is an injury-severity function determining the probability of injury severity category k for vehicle occupant n, x_{kn} is a vector of explanatory variables that affect the occupant's injury severity level k, β_k is a vector of estimable parameters, and ε_{kn} is an error term. Because the severity outcomes are clearly described in police accident reports, ranging from a fatality to vehicle damage only, researchers can approach the problem using methodologies to analyze discrete data that account for preference heterogeneity, such as mixed-logit (Eluru and Bhat, 2007).

In this model, the effectiveness of seatbelt use, for example, in reducing a fatality is estimated by including in the specification whether a seatbelt was used when an accident occurred (e.g., Eluru and Bhat, 2007). The police officer investigating an accident will report this variable in the police accident report after inspecting the accident. But, as discussed, the use of a non-random sample will still cause the estimate of the effect seatbelt use to be biased and as noted, a selection equation cannot be used that will be uncorrelated with all the omitted influences caused by selectivity bias that influence the occurrence of an automobile accident.

A minority of researchers also have estimated the determinants of severity without taking an econometric approach by performing simple data comparisons. For example, Evans (1986) compares the severity outcomes of pairs of passengers in the same car involved in an accident, with one passenger wearing and the other passenger not wearing a seatbelt.

Although this approach explicitly controls for differences in vehicle occupants who are involved in different accidents, it will still yield biased estimates of the effectiveness of seatbelts in reducing the probability of a fatality because it is based on a non-random sample of automobile travelers. That is, it consists of only those travelers who travel with companions who have distinctly different habits of wearing a seatbelt than they do. A sample designed to include automobile travelers' distinct seatbelt wearing habits, which are correlated with the travelers' attitudes toward safety, will yield biased estimates because seatbelt use will necessarily be correlated with the driver's attitude toward safety. A finding that motorists who wear seatbelts are less likely to be involved in a fatal accident may simply reflect that safer drivers, who are less likely to be involved in a fatal accident than are other drivers, are more likely to wear seatbelts.

The Aggregate Approach

Independent of research based on the controlled environment and disaggregate approaches, economists were conducting safety studies that relied on observational data and allowed for the analysis of unconditional accident probabilities by collecting aggregated data that included drivers who did and did not get involved in accidents. The aggregate approach attempted to explain the fatality rate per vehicle mile traveled (VMT) at the national, state, or regional level for a given time period as a function of safety policy variables, such as speed limits and seatbelt laws, and other influences, such as alcohol and drug consumption. Because the approach aggregated individual drivers' VMT and accidents, it included many drivers who never got into a fatal accident.

In the context of our model of driver behavior, if data were collected on the universe of all vehicles on the road for a given time period along with the severity outcomes in accident reports for those vehicles involved in an accident, measures of an accident or fatality rate could be constructed that are analogous to π_{it}^{j*}. Thus, for example, the effect of seatbelt use on the fatal accident rate, $\frac{\partial \pi_{it}^{j*}}{\partial z_i}$, could be identified provided we had variation in z_i that is orthogonal to other elements of X_{jt}.

The limitation of the aggregate approach is that it restricts the questions that can be asked about how to improve automobile safety. For example, aggregated data may not be available for particular socioeconomic groups of drivers, such as teenagers or less-affluent motorists. Thus, it may not be possible to estimate the effect of the introduction of states' seatbelt laws on the fatality rate of teenagers and less-affluent households. So, researchers may find that seatbelt laws reduce fatality rates, but they can only speculate about the primary sources of the safety improvement.

Researchers have made effective use of the aggregate approach to estimate broad impacts of policy changes. For example, Dee (1998) and Cohen and Einav (2003) leveraged the staggered rollout of state-level mandatory seat belt use laws to identify the effects of seatbelt use on the rate of overall driving fatalities. In those studies, the change in seat belt use is constant with aggregation (state-year combination), and the entire universe of fatal accidents is reported for each state. Thus, these studies provide consistent estimates of the effects of seatbelt use on the rate of overall driving fatalities.

5. Comparing Some of the Findings of the Approaches and a Caution About Policy

Table 1 summarizes the findings from a selection of studies taking different methodological approaches to estimate the extent that wearing seatbelts reduces driving fatalities. Interestingly, the studies that take the controlled environment and disaggregate approaches, which we argued are particularly subject to bias that could inflate the safety effects of seatbelts, find that seatbelts produce sizable reductions, roughly 40% to 60%, in auto fatalities. In contrast, studies that take the aggregate approach, which we argued are not subject to the same bias that affects estimates obtained from the controlled environment and disaggregate approaches, find that wearing seatbelts produce notably smaller reductions in auto fatalities on the order of 10%.

Circumstantial evidence on seatbelt use and automobile fatalities in the United States in recent decades suggests that the smaller estimates obtained from the aggregate approach are more plausible than the larger estimates obtained from the alternative approaches. As shown in figure 1, highway fatalities have declined more slowly during the 2000s than in previous decades, roughly 2% from 2000 to 2023. During the same period, because of stronger and more comprehensive seatbelt laws at the state level, greater enforcement of those laws, and public awareness campaigns, seatbelt use in the United States has increased from roughly 70% in 2000 to roughly 92% in 2023.[7] Thus, the 30% increase in seat belt use during the period is associated with a 2% decrease in auto fatalities, or an elasticity of roughly 7%, which is much closer to the estimates obtained from the aggregate studies than it is to the estimates obtained from the controlled environment and the disaggregate studies. Of course, this comparison does not hold everything else constant. But it is difficult to identify other changes in drivers and the driving environment during that period that could have significantly reduced the effect of the increase in seatbelt use on fatalities.

The upward bias in the controlled environment and disaggregate approaches could have contributed to a costly introduction of mandatory seatbelt laws if the studies based on those approaches helped to influence policymakers to enact those laws prematurely. Thaler and Rosen (1976) and Mannering and Winston (1987) found that although federal law in 1968 required seat belts to be installed in all vehicles except buses, many motorists eschewed their safety benefits based on a rational cost-benefit assessment of the time and bother costs to fasten seat belts and their effect on reducing the probability of a fatal accident.

---

[7] These figures are from the National Highway Traffic Safety Administration's (NHTSA) National Occupant Protection Use Survey (NOPUS).

Even by 1985, when New York was the first state to introduce a mandatory seatbelt law, seatbelt use in the nation was only 19%. By 1995, when 49 states had introduced some type of mandatory seatbelt law (New Hampshire has yet to introduce one), seatbelt use in the nation was roughly 68%. Thus, roughly one-third of US motorists found that the cost of using a seatbelt was sufficiently onerous that they were willing to disobey the law and eschew the safety benefits of wearing a seatbelt. Undoubtedly, during the period when seatbelt laws were being introduced by the states, a large share of motorists could have concluded that they incurred costs from being forced to use seatbelts that exceeded the benefits that they perceived from wearing one.

Of course, seatbelt use is much higher today and there is little evidence that a notable share of motorists is, on net, incurring costs from using them. But between 1985 and 1995, well-intentioned policymakers, who believed that seatbelt use would reduce the probability of a fatality by 40% to 60%, could have been influenced to prematurely introduce mandatory seatbelt laws, which were opposed by nearly two-thirds of the public in a 1984 Gallup Poll[8], produced fewer benefits than expected, and were exceeded in many cases by motorists' costs from being required by law to wear them. It also appears that policymakers prematurely mandated in 1998 that automakers install airbags in all new cars and light trucks despite consumers adopting them and automakers installing them in a manner that was consistent with cost-benefit analysis.[9]

6. Toward Efficient Adoption of Autonomous Safety Features in Automobiles

Given the limitations we have identified in automobile safety research, it is tempting to suggest that the recent stagnation in automobile safety since 2010 is attributable to the lack of solid policy guidance provided by the research community. However, an alternative and more plausible explanation in our view is that since highway safety has continually improved for roughly a century and safety performance is determined by a complex set of drivers' decisions that interact with a dynamic driving environment, it is unlikely that any purely regulatory policy could further improve automobile safety in the United States in a quantitatively meaningful way.

[8] https://tpmblegal.com/how-seatbelt-use-has-changed/

[9] Mannering and Winston (1995) found that, on average, motorists were willing to pay the average cost of installing air bags in their vehicles and that automakers were steadily installing airbags on those vehicles for which motorists were willing to pay the average cost of air bag installation. Nonetheless, in 1998, federal law required that all cars and light trucks sold in the United States have air bags on both sides of the front seat without policymakers carefully assessing whether such a requirement was justified on cost-benefit grounds, accounting for the welfare loss to motorists who valued air bags at less than the cost that was passed on to them through higher vehicle prices.

However, automobile safety could be potentially improved by either changes in motorists' driving behavior that reduce the risk of an accident or advances in automobile technology that reduce the severity and likelihood of an accident. Widespread constructive changes in driver behavior that have the potential to improve safety could include psychological and emotional changes in how drivers value their time, their relationships with other drivers and road users, their relationships with their vehicles, and the like. Indeed, such behavioral variation may help explain the substantial variation in automobile fatality rates across high-income countries with well-developed infrastructure and modern vehicle fleets. Unfortunately, it is extremely difficult to change the culture of driving in any country; a rare but notable exception occurred during the 1980s when Mothers Against Drunk Driving conducted their effective campaign. Moreover, government policymakers are not well-suited to changing the current culture of driving in the United States.

Instead, the most promising source of a significant safety improvement in the future is the major technological advance represented by the widespread adoption of autonomous vehicles (AVs). AVs would replace the drivers' optimization problem that we have discussed here, along with their choice of speed s_{it} and the various influences on safety outcomes, with the network optimization problem of determining vehicles' speeds and routings without the threats to safety created by drivers' heterogeneous preferences for risky behavior (Winston and Karpliow (2020), Winston, Yan, Associates (2024)).

Although AVs' technology is currently being perfected and tested, their widespread adoption is still decades away. But beginning in the late 2000s, automakers have taken an important step to introduce AV technology to the public 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). In contrast to other vehicle safety features, such as seatbelts and airbags, which enhance safety by reducing the severity of an injury if an accident occurs, ADAS substitutes for a driver's attention by making its own decisions to prevent an accident from occurring; for example, it may brake automatically to avoid a collision.

We apply the findings we have presented here to gain a better understanding of ADAS performance to show the usefulness of an aggregate approach for analyzing the effect of the latest automobile technological innovation on safety, illustrate a framework to guide appropriate policy toward the widespread adoption of a new innovation, and set the stage for analyzing the nation's transition to full automobile autonomy.

Theoretical arguments and empirical evidence indicate that the controlled environment and disaggregate approaches to assess the effects of a vehicle safety attribute are seriously flawed but that a well-executed aggregate approach can identify some informative effects, especially when a natural experiment is available such as the rollout of state seatbelt laws. Maheshri, Winston, and Wu (2025) take an aggregate approach and exploit the rollout of advanced driver assistance systems by model year, make and model, and vehicle trim from 2010 to 2018 to assess the effect of ADAS on reducing drivers' accidents and fatalities.

Their data consists of all registered *vehicles* in Texas from 2010 to 2018, which have and have not been involved in accidents. The data are linked to a dataset of all accidents for which a police report was filed in Texas during the period. Finally, data were obtained that identify the availability of ADAS-related safety features on each trim of every vehicle that was registered during the sample period. The dataset is used to compare the aggregate safety performance of vehicles with and without ADAS.

As in any econometric study, an identifying assumption must be made. Given that ADAS became available at different times for different trim levels—notably within vehicles of the same make and model—the causal effect of ADAS on accidents is assumed to be identified under the plausible assumption that drivers did not systematically opt for higher trim level vehicles solely because of the availability of ADAS. That is, new vehicles with a low trim level were never equipped with ADAS during the sample period; new vehicles with a medium trim level were equipped with ADAS only in 2018, the last year of the sample; and vehicles with a high trim level were equipped with ADAS beginning in 2015 but not before that calendar year.

Of course, drivers of higher trim vehicles may 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 the identifying assumption, which relies on a combination of the choice of higher trim versus lower trim *and* the timing of ADAS availability.

Formally, the model is specified in a Poisson regression framework, because the dependent variables (vehicle accidents and fatal accidents) take on small, discrete, non-negative values, as:

$$A_{yijt} = exp(\beta\, S_{yijt} + \lambda_{ijt} + \lambda_{yjt} + \epsilon_{yijt}) \,, \tag{7}$$

where A_{yijt} indicates vehicle or fatal accidents for model year y of make and model i with trim level j in year t. 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_{yjt} are make-model-model year-calendar year fixed effects; and ϵ_{yijt} is an error term.

The parameter β can be interpreted as the causal effect of the availability of ADAS on selected vehicles on the total number of accidents or fatalities if 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. Several pieces of empirical evidence were provided to support the assumption. For example, the authors showed that over time the safest drivers did not disproportionately switch to vehicles equipped with ADAS when those safety features were first made available.

Estimates of β indicated that ADAS is highly effective at improving automobile safety even after accounting for drivers' behavioral responses to its availability and installation. Specifically, ADAS technologies reduced the risk of a motorist getting in a single vehicle accident by 13%, reduced the risk of a motorist getting in a multivehicle accident by 10%, and reduced 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, which is likely to involve a vehicle not equipped with ADAS. Thus, ADAS is likely to reduce the fatality risk of multivehicle accidents as a greater share of the nation's vehicle fleet is equipped with autonomous vehicle safety features.

The current small share of ADAS-equipped vehicles in the nation's entire vehicle fleet also makes it difficult to perform some useful reality checks. However, 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-equipped vehicles on the *nation's* automobile fatalities and insurance costs. The latter 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.

For purposes of comparison, a disaggregate approach would compare the safety performance of *drivers* who got into accidents using vehicles that were and were not equipped with ADAS. Thus, the effect of ADAS would be identified only under the implausible assumption that a driver's propensity to purchase an ADAS-equipped vehicle was uncorrelated with her attitudes toward safety and her driving abilities. As noted, the bias that results from this assumption is likely to inflate the estimates of β and could play a role in influencing policymakers to require automakers to install ADAS in all their new vehicles. The upward bias from a controlled environment approach that assesses the effect of ADAS also could play a role in justifying policymakers' mandate to automakers.

Once fully autonomous *vehicles* start to be adopted, it will be possible to assess those vehicles' effects on safety with renewed expectations about the feasibility of using disaggregated data because the technology will be making the driving decisions, which eliminates the feedback effects from drivers' risk preferences and their behavior. That is, fully autonomous vehicles eliminates the heterogeneity in drivers' riskiness. Of course, the availability of autonomous vehicles does not eliminate the self-selection issue of the type of drivers who choose to adopt AVs as AVs are gradually adopted in the US. But the self-selection issue is well-defined and could potentially be addressed with a conventional selection equation that specifies exogenous socioeconomic characteristics, such income, education, family size, occupation, and the like, as proxies for travelers' risk preferences.

7. Final Comments

This paper offers two important conclusions about automobile safety research and safety performance and policy, both of which will take some time to be realized in practice. First, the flaws associated with a disaggregate or a controlled environment approach prevent this line research from obtaining causal estimates of the determinants of automobile accidents that are identified and useful for safety policy. Causal estimates of the determinants of automobile accidents obtained from an aggregate approach can be identified in limited contexts.

Unfortunately, the disaggregate approach has become the dominant approach in safety research and as have we have discussed, researchers have continued to use it because the method and the findings have not been assessed on the grounds of whether they are useful for policy. The fact that the disaggregated approach is not useful for policy in a research environment where policy issues are largely ignored suggests that the method will be abandoned only when new researchers come along with the ambition of conducting safety research to eventually help policymakers reduce the nation's trillion dollars cost of automobile accidents by making methodological advances that can produce causal estimates of the determinants of automobile accidents that are useful for policy. The transition to a new generation of transportation researchers that makes this advance is likely to take a number of decades.

Second, notwithstanding the limited contributions of safety research, advances in automobile technology and public investments in infrastructure have enabled automobile safety to steadily improve for roughly a century. However, government policymakers have periodically overreacted to occupant safety improvements by prematurely mandating that motorists use them and that automakers install them on all their new vehicles. Those mandates have imposed costs on consumers whose value of the increased safety is less than the costs of time and bother to use the new occupant safety features and the increase in prices they must pay to cover their installation cost.

The recent response by policymakers to the introduction of autonomous vehicle safety technologies in the form of advanced driver-assistance systems (ADAS) appears to be aligned with the preceding history. An aggregate research approach has provided credible evidence of the positive effects of ADAS, which is identified. At the same time, a study taking a controlled environment approach (Haus, Sherony, and Gabler, 2019) estimated that a component of ADAS, autonomous emergency braking (AEB), could reduce pedestrian fatality risk by roughly 85%, which is likely to be significantly upward biased. In any case, federal policymakers appear to be prematurely mandating that automakers install AEB on all their new model year 2030 passenger cars and light trucks by 2029. Thus, similar to consumers who, on net, incurred costs from policymakers' premature decisions to enact seatbelt laws and to mandate automakers' adoption of airbags in new vehicles, consumers who do not value the increased safety attributable to AEB by as much as the $500 to $2000 cost to install it (depending on the vehicle) will, on net, incur costs.

Of course, we do not know the factors that influenced policymakers to mandate that automakers install AEB in new cars and light trucks by 2029. But if they were influenced to any extent by inflated estimates of the safety effects of AEB from, for example, a study taking a controlled environment approach, then it would be clear that the importance of the concerns that we raise about current automobile safety methodologies extend well beyond the academic community.

Given the roughly $1 trillion annual cost of automobile accidents, is understandable that policymakers are eager to enact a potentially effective safety policy. But the research community has not produced credible evidence of AEB's benefits that supports mandating its installation in all new cars at some future date, while the technology itself is costly and still evolving.

Policymakers and researchers must proceed with humility as both autonomous vehicle technologies and their value to consumers evolve. The transition to an autonomous vehicle fleet also is likely to take at least three decades. In the meantime, claims by researchers that they have obtained causal findings that address the most important safety questions should be regarded with skepticism until they make methodological advances that enable their empirics to be clearly identified. Hopefully, transitions in the new generation of researchers and the transition to an autonomous vehicle fleet will converge such that transportation safety research can still be useful in the evolving autonomous environment.

Finally, policymakers should heed the concerns that we have raised here and should be careful attempting to use researchers' findings to justify safety policies without understanding their flaws. Of course, we live in a political world that has recently placed little value on reliable scholarly research findings. In the case of automobile safety, as in other areas, public policies based on erroneous empirical claims will significantly harm consumers.

References

Cohen, Alma, and Liran Einav. 2003. "The Effects of Mandatory Seat Belt Laws on Driving Behavior and Traffic Fatalities." *Review of Economics and Statistics*, vol. 85, no. 4, pp. 828–843.

Dee, Thomas S. 1998. "Reconsidering the Effects of Seat Belt Laws and Their Enforcement Status." *Accident Analysis & Prevention*, vol. 30, no. 1, pp. 1–10.

Demsetz, Harold. 1973. "Industry Structure, Market Rivalry, and Public Policy." *Journal of Law and Economics*, vol. 16, no. 1, pp. 1–9.

Eluru, Naveen, and Chandra R. Bhat. 2007. "A Joint Econometric Analysis of Seat Belt Use and Crash-Related Injury Severity," *Accident Analysis and Prevention*, vol. 39, no. 5, pp. 1037–1049.

Evans, Leonard. "The Effectiveness of Safety Belts in Preventing Fatalities." *Accident Analysis and Prevention*, vol. 18, no. 3, 1986, pp. 229–241.

Haus, Samantha H., and Rini Sherony, and Hampton C. Gabler. 2019. "Estimated Benefit of Automated Emergency Braking Systems for Vehicle–Pedestrian Crashes in the United States," *Traffic Injury Prevention,* vol. 20, No. S1, pp. S171-S176.

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.

Kahane, C. J. 2015. Lives saved by vehicle safety technologies and associated Federal Motor Vehicle Safety Standards, 1960 to 2012 – Passenger cars and LTVs – With reviews of 26 FMVSS and the effectiveness of their associated safety technologies in reducing fatalities, injuries, and crashes. (Report No. DOT HS 812 069). Washington, DC: National Highway Traffic Safety Administration, January.

Lucas, Robert E. 1976. "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, volume 1, pp. 19-46.

Maheshri, Vikram, and Clifford Winston. 2016. "Did the Great Recession Keep Bad Drivers Off the Road?," *Journal of Risk and Uncertainty,* vol. 52, no. 3, pp. 255-280.

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.

Maheshri, Vikram, and Clifford Winston, and Yidi Wu. 2025. " AI at the Wheel: The Effectiveness of Advanced Driver-Assistance Systems and its Implications for Policy," working paper.

Mannering, Fred, and Clifford Winston. 1987. "Recent Automobile Occupant Safety Proposals," in Clifford Winston and Associates, *Blind Intersection? Policy and the Automobile Industry,* Brookings Institution, pp. 68-88.

Mannering, Fred, and Clifford Winston. 1995. "Automobile Air Bags in the 1990s: Market Failure or Market Efficiency?," *Journal of Law and Economics,* volume 38, October, pp. 265-79.

Peltzman, Sam. 1975. "The Effects of Automobile Safety Regulation," *Journal of Political Economy*, vol. 83, no. 4, pp. 677–725.

Thaler, Richard, and Sherwin Rosen. 1976. "The Value of Saving a Life: Evidence from the Labor Market," in Nestor E. Terlecky, editor, *Household Production and Consumption,* National Bureau of Economic Research.

TRIP. 2023. *Addressing America's Traffic Safety Crisis: Examining the Causes of Increasing US Traffic Fatalities and Identifying Solutions to Improve Road User Safety*, Washington, DC, June.

Tscharaktschiew, Stefan. 2020. "Why are Highway Speed Limits Really Justified? An Equilibrium Speed Choice Analysis," *Transportation Research Part B: Methodological,* vol. 138, August, pp. 317-351.

Vertlib, Shani R., Stav Rosenzweig, Ofir D. Rubin, and Aviv Steren. 2023. "Are Car Safety Systems Associated with More Speeding Violations? Evidence from Police Records in Israel," PLoS ONE 18(8): e0286622. https://doi.org/10.1371/journal. pone.0286622

Wilde, Gerald J.S. 1982. "The Theory of Risk Homeostasis: Implications for Safety and Health," *Risk Analysis,* vol.2, issue 4, pp. 209-225.

Winston, Clifford. 2021. "Back To The Good—or Were They the Bad—Old Days of Antitrust? A Review Essay of Jonathan B. Baker's The Antitrust Paradigm: Restoring a Competitive Economy," *Journal of Economic Literature,* vol. 59, issue 1, pp. 265-284.

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,* 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.

Yang, Hai, Hongbo Ye, Xinwei Li, and Bingqing Zhao. 2015. "Speed Limits, Speed Selection and Network Equilibrium," *Transportation Research Part C: Emerging Technologies,* vol. 51, February, pp. 260-273.

Zoua, Xin, Hai L. Vua, and Helai Huang. 2020. “Fifty Years of Accident Analysis & Prevention: A Bibliometric and Scientometric Overview,” *Accident Analysis and Prevention,* vol. 144, September, pp. 1-26.

Figure 1. US Automotive Fatality Rate Over Time (in Logs)[a]

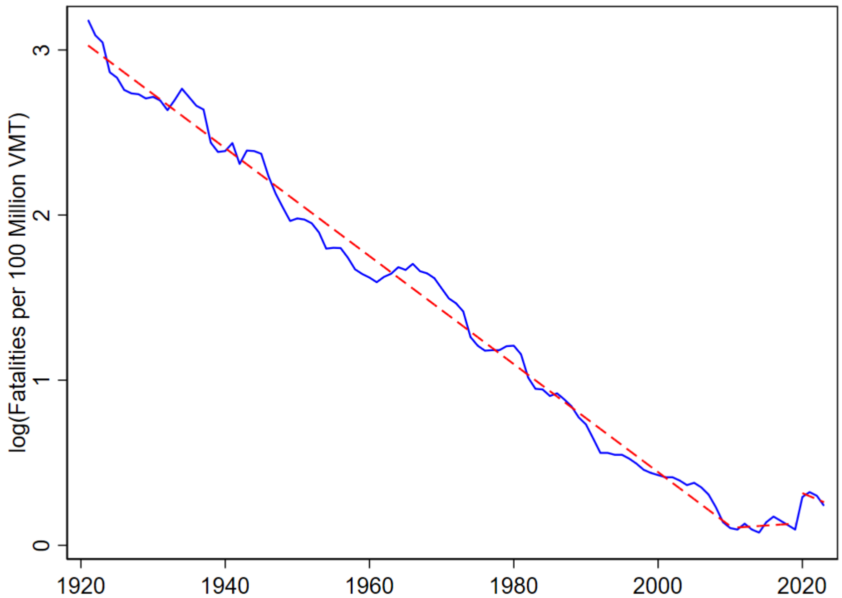

[a] Fatalities data from the US Department of Transportation.

Table 1. Selected Empirical Studies of the Effect of Seatbelts on Automobile Fatalities.

Year	Authors	Findings	Notes
Controlled Environment			
1970	Lave and Weber	40-50% reduction	Use biomechanical evidence from government crash tests.
2015	Kahane (NHTSA)	25-69% reduction	Use biomechanical evidence from government crash tests.
Disaggregate			
1986	Evans	42% reduction	Compares pairs of passengers in the same car, one belted, one unbelted.
2007	Eluru and Bhat	64% reduction	Joint model of seat belt use and accident severity conditional on a collision.
Aggregate			
1975	Peltzman	0% overall, accounting for pedestrian deaths.	Compares trends before and after 1968 federal safety regulations.
1998	Dee	5-6% reduction	Exploits staggered rollout of mandatory seat belt laws by states in diff-in-diff estimation.
2003	Cohen & Einav	13% elasticity of fatalities to usage	Exploits staggered rollout of mandatory seat belt laws by states as IV for reported usage.