FUEL ECONOMY AND SAFETY: THE INFLUENCES OF VEHICLE CLASS AND DRIVER BEHAVIOR

Mark R. Jacobsen*

September 2012

Abstract

Fuel economy standards change the composition of the vehicle fleet, influencing accident safety. The direction and size of the effect depend on interactions in the fleet. The model introduced here captures these interactions simultaneously with novel estimates of unobserved driving safety behavior and selection. I apply the model to the present structure of U.S. fuel economy standards, accounting for shifts in the composition of vehicle ownership, and estimate an adverse safety effect of 33 cents per gallon of gasoline saved. I show how two alternative regulatory provisions fully offset this effect, producing a near-zero change in accident fatalities.

* University of California at San Diego, Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093. Email: m3jacobsen@ucsd.edu. The University of California Energy Institute has generously provided funding in support of this work. This research is not the result of any consulting relationship. A short version appears in the 2011 proceedings of the American Economic Association annual meetings. I thank Kenneth Small and seminar participants at Harvard University, The University of Maryland, Columbia University, the UC Energy Institute, the American Economic Association, and the NBER Summer Institute for their helpful suggestions.

1. Introduction

Automobile policy in the U.S., particularly regulation to conserve gasoline, changes the composition of the vehicle fleet with the potential to influence accident safety. Competing interactions in the fleet leave both the direction and magnitude of the safety impacts as empirical questions.[1] Given the large annual cost of car accidents,[2] I show how changes in accident rates importantly alter the efficiency ranking of alternative fuel-economy policies. I build on a rich literature investigating the welfare impacts of fuel-economy standards and gasoline taxes (Goldberg [1998], Portney et al. [2003], Austin and Dinan [2005], Klier and Linn [2008], Bento et al. [2009], Busse, Knittel, and Zettelmeyer [2009], Anderson and Sallee [2011]).

While gasoline taxes are often argued to provide greater efficiency along a number of margins, including safety, U.S. policy instead focuses on fuel economy limits as the primary means to trim gasoline consumption: new corporate average fuel economy (CAFE) standards are set to make large improvements through 2016, with an even more ambitious limit for 2025 that nearly doubles fuel economy relative to today's fleet.[3] The economics literature considering CAFE and safety is relatively sparse, but work by Crandall and Graham (1989) and the National Research Council (2002) suggests significant adverse safety costs. Their estimates translate to more than $1.50 in safety cost per gallon of gasoline saved, rivaling the entire distortionary cost of CAFE appearing in recent studies.[4] A group of time-series studies considering both recent trends and the 1978 introduction of CAFE produce varied results, including the potential for safety benefits

[1] Reducing the number of large or heavy vehicles – substituting them evenly into the rest of the fleet – both conserves fuel and reduces the number of unevenly matched, risky accidents. At the same time, smaller or lighter weight vehicles tend to offer their own occupants less protection, operating the other direction on overall risk.

[2] There were 37,261 U.S. traffic fatalities and more than 2.3 million injured in 2008 (NHTSA, 2009).

[3] Environmental Protection Agency and Department of Transportation (2010), and The White House Office of the Press Secretary (2011).

[4] Jacobsen (2010) and Anderson and Sallee (2011) estimate the efficiency costs of CAFE at under $2.00 per gallon. To the extent much of vehicle safety is external (see Footnote 10) the safety implications of changing vehicle choice are not fully captured.

from the policy.[5] Finally, there is a long engineering and economics literature linking various vehicle attributes and safety: among the results a number of mechanisms offering CAFE the potential to improve safety have been identified.

The model in this paper overcomes two key challenges in the existing literature: i) a set of sometimes disparate implications for policy depending on the particular vehicle attribute and type of accident studied, and ii) a challenge in separating risks of the vehicle from risks due to driver behavior and selection on vehicle choice. My model also allows the novel ability to estimate accident rates in arbitrary counterfactual fleets after policy causes drivers to move across vehicles. I show that my application to CAFE has economically important implications for how policy is implemented. I believe the model can also be valuable in considering numerous other environmental and vehicle safety policies that change the composition of the fleet.

I address the first challenge by taking a semi-parametric approach based on interactions of vehicle classes rather than individual attributes, nesting prior results on accident risk: Crandall and Graham (1989) and others find very strong protective effects of vehicle weight, suggesting adverse effects of CAFE. Recent work by Anderson and Auffhammer (2011) instead focuses on the increased risks that weight imposes on other vehicles in accidents, demonstrating how they may be reduced using gasoline or weight-based taxes.[6] Another strand of the literature shows that the pickup truck and SUV classes, independent of weight, impose especially dramatic risk on others with only modest gains for their own occupants.[7] Still other studies focus on attributes like height, wheelbase, and rigidity, often finding dramatic changes to risk across interactions of these features.[8] In contrast to these approaches I group vehicles into a discrete set of ten classes that cut

---

[5] Khazzoom (1994), Noland (2004), and Ahmad and Greene (2005).
[6] The National Research Council (2002) and Kahane (2003) also provide broad summaries of the effects of weight.
[7] White (2004), Gayer (2004), Anderson (2008), and Li (2012) show evidence of the risks that pickups and SUV's impose on other classes.
[8] Kahane (2003) considers a list of attributes including height mismatch and frame rigidity.

flexibly across physical dimensions.[9] I estimate a separate risk coefficient for each of the 100 implied accident types, each class striking every other, allowing me to separate the protective and harmful effects of vehicles in each class without limiting the analysis to specific attributes.[10] The matrix of estimates from my approach can be mapped back to individual attributes *ex post*, allowing me to demonstrate how my results fit into findings in the prior literature.

The second key challenge is selection on unobservable driver attributes: not only might riskier drivers cluster in certain vehicles (biasing measures of how dangerous those vehicle models really are) but drivers will also move across models as they re-optimize according to the incentives placed by gasoline policy.[11] The movement of drivers requires direct estimates of driver risk by vehicle type in order to conduct policy counterfactuals. The approach I contribute addresses these questions by leveraging common factors attributable to drivers or geography that appear across a system of equations describing single-vehicle and two-vehicle accidents. I estimate unobserved driver risk while simultaneously considering the influence of physical features of vehicles on risk in accident interactions.

Among other factors the unobserved riskiness of drivers by vehicle type captures the safety of roads in the driver's geographical area, a tendency to drive drunk,[12] and Peltzman-type effects where the protective nature of a vehicle itself may affect driving behavior. The estimates address, for example, a puzzle in this literature related to

---

[9] The estimates are semi-parametric in the sense that no restrictions are placed on the combinations of physical characteristics available across vehicles. For example the classes below broadly span weight, volume, height, passenger capacity, frame type, and engine size.

[10] My estimates therefore distinguish between internal and external costs if we can assign changes in the protective effect across cars as internal and changes in damage to other vehicles as external. However, health, life, and disability insurance (not conditioned on vehicle choice) make part of the protective effect external, while automobile liability insurance, psychic costs, and the potential for civil and criminal liability internalize part of the damage a vehicle is expected to impose on others. I therefore focus on total accident cost in the fleet when comparing policy counterfactuals.

[11] A number of approaches to the selection portion of this issue appear in the literature. For example fatality risk can be measured conditional on an accident occurring (Anderson and Auffhammer [2011]), or using measures of induced exposure from police findings on fault.

[12] Levitt and Porter (2001) provide an innovative method to estimate drunk driving rates using innocent vehicles in accidents as control, but in most cases (including the present study) such personal characteristics are difficult to observe.

minivans: my model attributes their scarcity in fatal accidents largely to unobserved driver behavior rather than to the vehicles themselves. To my knowledge these estimates of selection on driver risk are novel to the literature. In the policy application this aspect also turns out to be pivotal to the welfare results.

Application of the model leads to a set of results investigating the safety effects of fuel economy policy. The estimated fleet-wide impact of a policy based on the historical CAFE rules is 149 additional annual fatalities per mile-per-gallon increment; a welfare cost of approximately 33 cents per gallon of gasoline saved.[13] I then consider a "unified" fuel economy policy that combines size reductions within the car and truck categories with broader switching across the two categories. When normalized to conserve the same amount of fuel this results in an increase in fatalities of only 8 per year, with a zero change included in the confidence band. In each of the two policies I demonstrate how accounting for driver behavior, and the movement of risky drivers through the fleet as policy changes car choices, influences the results.

Finally, I consider a "footprint" type rule similar to the provisions in fuel economy standards through 2016, and include alternative simulation approaches that address potential confounders. Among these is a set of simulations modeling Peltzman (1975) effects where driver risk behavior changes based on the vehicle selected. Further extensions to the simulation model could allow analysis in a variety of other settings. For example the U.S. "cash-for-clunkers" program as described in Knittel (2009) or incentives to switch among new and used vehicles in Busse, Knittel, and Zettelmeyer (2009) produce changes in the fleet that may importantly alter the efficiency of policy.

The rest of the paper is organized as follows: Section 2 describes U.S. fuel economy policy and the role of safety. Section 3 presents the model. Sections 4 and 5 respectively describe the data and empirical results. Section 6 presents the policy experiments, combining my empirical results with a model of fuel economy regulation. Section 7 considers alternative specifications and addresses robustness.

---

[13] Parry and Small (2005) estimate the external cost of gasoline consumption to be about $1.00 per gallon in the U.S.

2. Safety and Fuel Economy Regulation

The importance of automobile safety is evident simply from the scale of injuries and fatalities each year. In 2008 there were 37,261 fatalities in car accidents on U.S. roads and more than 2.3 million people injured.[14] The National Highway Traffic Safety Administration (NHTSA) is tasked with monitoring and mitigating these risks and oversees numerous federal regulations that include both automobiles and the design of roads and signals.

To motivate the concern about fuel economy standards with respect to safety consider the very rough estimate provided in NRC (2002): approximately 2,000 of the traffic fatalities each year are attributed to changes in the composition of the vehicle fleet due to the CAFE standards. If we further assume that the standards are binding by about 2 miles per gallon, this translates to a savings of 7.5 billion gallons of gasoline per year. When valuing the accident risks according to the Department of Transportation's methodology this implies a cost of $1.55 per gallon saved through increased fatalities alone.[15] This does not consider injuries, or any of the other distortions associated with fuel economy rules, yet by itself exceeds many estimates of the externalities arising from the consumption of gasoline.[16]

Conversely, a finding that accident risks improve with stricter fuel economy regulation would present an equally strong argument in favor of more stringent rules. The magnitude of the implicit costs involved in vehicle safety motivate the importance of a careful economic analysis, and mean that even small changes in the anticipated number of fatalities will carry great weight in determining the optimal level of policy.

---

[14] NHTSA (2009).
[15] The Department of Transportation currently incorporates a value of statistical life of $5.8 million in their estimates. This is conservative relative to the $6.9 million used by EPA.
[16] See Parry and Small (2005).

Current regulation

U.S. fuel economy regulation is in transition, with the rule through 2016 now complete (Environmental Protection Agency and Department of Transportation [2010]), while regulatory provisions beyond 2016 remain to be determined. I consider three possible regulatory regimes, each of which produces a unique effect on the composition of the fleet. The resulting impacts on the frequency of fatal accidents are similarly diverse:

1) The corporate average fuel economy (CAFE) rules: Light trucks and cars are separated into two fleets, which must individually meet average fuel economy targets. No direct incentive exists for manufacturers to produce more vehicles in one fleet than the other. Rather, the incentives to change composition occur inside each fleet: selling more small trucks and fewer large trucks improves the fuel economy and compliance of the truck fleet. The same is true inside the car fleet. This produces a distinctive pattern of shifts to smaller vehicles within each fleet, but without substitution between cars and trucks overall.

2) A unified standard: This type of standard was introduced in California as part of Assembly Bill 1493, and is under consideration federally.[17] It regulates all vehicles together based only on fuel economy. This includes the effects above while simultaneously encouraging more small vehicles, broadly shifting the fleet away from trucks and SUV's and into cars.

3) A "footprint" standard: This new type of rule is in place federally for the years 2012 - 2016 and is also expected for the years 2017 through 2020. It assigns target fuel economies to each size of vehicle (as determined by width and wheelbase), severely limiting the incentives for any change in fleet composition. As such it increases the technology costs of meeting a given target, but was required in the hopes of mitigating the costly safety consequences discussed above.[18]

---

[17] Strictly speaking the California bill preserves the fleet definition, but allows manufacturers to "trade" compliance obligations between fleets in order to achieve a single average target. The trading between fleets aligns incentives for all vehicles, making the rule act like a single standard.

[18] NHTSA (2010) discusses the safety rationale for the footprint rule. Technology costs are higher because most improvement must be achieved through technology; the earlier rules allow part of the improvement to come from technology and part via fleet composition.

3. A Model of Accident Counts

I model the count of fatal accidents in each vehicle class, normalizing by billion miles traveled. Vehicle classes will be a set of J categories covering the entire vehicle fleet; the physical characteristics of each vehicle class are interesting in policy examples but do not enter the general model specification.

Define Z_{ij} as the count of fatal accidents where vehicles of class i and j have collided and a fatality occurs in the vehicle of class i. The counts will be asymmetric, that is $Z_{ij} \neq Z_{ji}$, to the degree that some classes impose greater risk on others. If a fatality occurs in both vehicles in an accident then Z_{ij} and Z_{ji} are both incremented, though this will be relatively rare in the data.

The total count of fatal accidents in class i vehicles is then:

$$\text{(fatalities in class i)} = \sum_{j \in J} Z_{ij} \tag{3.1}$$

where J represents the set of all vehicle classes. By changing the order of subscripts we can similarly write the count of fatalities that are imposed on other vehicles by vehicles of class i:

$$\text{(fatalities imposed on others by class } i) = \sum_{j \in J} Z_{ji} \tag{3.2}$$

Total counts of fatal accidents reflect a combination of factors influencing risk and exposure. I divide the counts into three multiplicative components, of interest individually and for use in constructing policy counterfactuals: 1) The risk coming from the behavior of drivers in each vehicle class, 2) risk coming from physical vehicle characteristics alone – I will term this the "engineering" risk, and 3) the miles driven in each class. The combination of these three elements determines the number of fatal accidents of each type: Intuitively, the greater the driver recklessness, engineering risk, or miles driven, the more fatal accidents we should expect to see of type Z_{ij}.

Define the three components using:

α_i The riskiness of drivers selecting each vehicle class i (in estimation this will appear as a fixed effect on driver behavior for each class; in counterfactual simulations it will be allowed to vary as drivers switch across classes)

β_{ij} The risk per mile of a fatality in vehicle i when vehicles from class i and class j are driven by average drivers (β_{ij} will estimated for all possible combinations of vehicles)

n_i The number of miles driven in vehicles of class i (available as data below)

I normalize the measure of driver riskiness, α_i, to unity for the average driver so that it functions as an accelerator multiplying the overall risk per mile driven. For example, a value of $\alpha_i = 2$ corresponds to a driver who generates twice the average fatality risk for each mile they drive. High values of α_i come from a tendency of class i owners to live in locations with dangerous roads, travel at high risk times of day, drive recklessly, distracted or drunk, or have any other characteristic (observable or unobservable) that increases the risk per mile of fatal accidents. Notice that this means α_i can operate either through an increase in the number of collisions or through an increase in fatality risk after a collision has occurred; the distinction is not needed to consider total fatalities in the fleet.

Combining this definition of dangerous driving behavior with the engineering fatality risk results in:

$$\text{Probability of a fatal accident in vehicle } i \mid i, j \text{ driven 1 mile} = \alpha_i \alpha_j \beta_{ij} \tag{3.3}$$

The probability of a fatal accident for vehicle i, per mile traveled by i and j, is modeled as the product of the underlying engineering risk in a collision of that type, β_{ij}, and the parameters representing high risk coming from the drivers involved, α_i and α_j.

The multiplicative form contains an important implicit restriction: behaviors that increase risk are assumed to have the same influence in the presence of different classes and driver types. I argue that this is a reasonable approximation given that most fatal accidents result from inattention, drunk driving, and signal violations;[19] such accidents give drivers little time to alter behavior based on attributes of the other vehicle or driver.

Finally I include the effect of the number of miles traveled in each class, n_i, and further subdivide miles and accidents across time and location with the subscript s. If pickup trucks are less common on urban roads, or minivans tend to be parked at night, there should be differences in the number of accidents involving these vehicles across time and space. In the estimation below I divide the data into bins s according to time-of-day, average local income, and urban density – factors that appear to significantly influence both the composition of the fleet and the probability of fatal accidents.

The effect of miles driven in bin s on the number of fatalities again takes a natural multiplicative form: If twice as many miles are driven in a certain class then we expect twice as many cars of that class to be involved in an accident:

$$E(Z_{ijs}) = n_{is} n_{js} \alpha_{is} \alpha_{js} \beta_{ij} \tag{3.4}$$

I also add the bin s subscript to α since the risk multiplier may also differ across time and space. Broadly speaking, data will be available on Z and n leaving α and β to be estimated.[20]

The key challenge in this literature becomes clear in (3.4): Since the α_i terms include unobservable driving behavior, and the engineering risks β are also to be estimated, we need a way to separate the two. Is a vehicle class dangerous because of its engineering characteristics or do the drivers who select that class just happen to have high risk (from factors like the location where they live or poor driving habits)?

The method I propose here identifies driver risk via a second equation describing single-car fatalities, using the assumption that overall driver risk (in α_{is}) will influence both equations. I define the count of fatal single-car accidents in vehicle class i in location s as Y_{is}, such that:

---

[19] NHTSA (2008).

[20] Details are provided in Section 5, but I specifically will observe Z_{ijs} and n_i. The aggregation to n_i means differences in α_i across bins will not be observed, but recovery of an average α_i for each class is still possible.

$$E(Y_{is}) = n_{is}\alpha_{is}\lambda_s x_i \quad (3.5)$$

The four parameters are:

n_{is} (As above) The number of miles driven in class *i* and bin *s*

α_{is} (As above) The riskiness of drivers or the locations they live in

λ_s Controls the relative frequency of fatal single-car accidents separately for each bin

x_i The relative fatality risk to occupants of class *i* in a standardized collision (measured using crash test data, or in an alternative specification through additional restrictions on the β matrix)

The key identifying restriction across equations (3.4) and (3.5) is that dangerous locations or behaviors (in α) that differing across classes enter both the risk of single-car accidents and the risk of accidents with other vehicles. This may be a better assumption for some factors (geographical location, drunk driving, recklessness) than others (falling asleep) but I will argue below that the estimates closely match intuition on driver safety generally. λ_s allows flexibility in the relative frequency of single and two-car accidents (single car accidents are more frequent at night, for example), and importantly relaxes the stringency of the identifying restrictions; Section 5 on estimation provides more intuition on the exact nature of the restriction and the role of λ_s in the context of my data.

Comparison with other models of safety

Much of the work focusing on the influence of vehicle characteristics on safety (see Kahane [2003]) has taken a parametric approach in an attempt to isolate the effect of weight alone. By assigning a complete set of fixed effects for all possible interactions, β_{ij}, I can still recover information about vehicle weight while adding considerable flexibility in form and the ability to capture other attributes that vary by class. The cost to my approach comes in the degree of aggregation: I will consider 10 distinct classes, or 100 β_{ij} fixed effects. Since each class contains a variety of vehicles I must assume that changes caused by regulation inside a class are of relatively small importance compared with the changes across classes. The assumption will have the most influence at the extremes of the distribution, for example downsizing within the compact class and within the large pickup class. In the context of safety these biases will cancel out to some degree, though this remains an important caveat.

Wenzel and Ross (2005) describe overall risks using a similarly flexible class-based approach to vehicle interactions, but importantly do not model driving safety behavior and so are unable to separate it from underlying engineering risk. For purpose of comparison I provide estimates of a restricted version of my model where I set all the α_i's equal. The parameter estimates turn out to be quite different, so much so that the primary economic and policy implications are reversed in sign.

4. Data

I assemble data on each of the three variables needed to identify the parameters of (3.4) and (3.5):

- Comprehensive count of fatal accidents, Z_{ijs} and Y_{is}
- The number of miles driven in each class, n_i
- Crash test data to describe risks in single-car accidents, x_i

Fatal accident counts

The count data on fatal accidents is the core information needed to estimate my model. I rely on the comprehensive Fatal Accident Reporting System (FARS), which records each fatal automobile accident in the United States. The dataset is complete and of high quality, due in part to the importance of accurate reporting of fatal accidents for use in legal proceedings. If such complete data were available for accidents involving injuries or damage to vehicles it could be used in a framework similar to the one I propose, but reporting bias and a lack of redundancy checking in police reports for minor accidents make those data less reliable.

The FARS data include not only the vehicle class and information about where and when the accident took place (which I use to define bin *s* in the model), but a host of other factors like weather, and distance to the hospital. While the additional data is not needed in my main specification (which captures both observed and unobserved driver choices in fixed effects) I will make use of a number of these other values to investigate the robustness of my estimates.

I define the bins *s* using three times of day (day, evening, night), two levels of urban density, and three levels of income in the area of the accident. For the latter two items I use census data on the zip codes where the accidents take place. This creates 18 bins in my central specification that, together with time, produce the replicates on Z_{ij} used for estimation. The key parameters of interest are at the vehicle class, rather than bin, level and the selection of bin divisions turns out to have a relatively small impact empirically. An exploration of both more and less aggregate bin structures is provided in Appendix B.

For my main specification I pool data for the three years 2006-2008 and use weekly observations on fatal accident counts. I experiment with month-of-sample fixed effects and a non-overlapping sample of data from 1999-2001 and find no important differences in results. The selection of the time period is to match the timing of data on vehicle quantities and miles driven (see below), observed in 2001 and again in 2008. The pooling provides additional power in estimation.

Miles driven

I use total vehicle miles traveled (VMT) in each class as a measure of the quantity of vehicles of that class present on the road. This data is available from the National Household Transportation Survey (NHTS), which is a detailed survey of more than 20,000 U.S. households conducted in 2008. While I do have some information about the location of the VMT (for example the home state of the driver) I do not observe other important aspects like the time of day or type of road where the miles are driven. Fortunately, as shown in Section 5, it is possible to recover values for the parameters defining driver behavior using only the total VMT for each class: differences in bin *s* level VMT are absorbed in fixed effects.

While the NHTS enjoys wide use it remains subject to a number of important caveats: in my application sampling or reporting bias correlated with driver risk could bias the estimates. Many of the characteristics used in constructing the sample weights for the NHTS are also associated with safety (age, income, education level, and, of particular relevance here, location and geography). This offers some reassurance on the accuracy of aggregate VMT reported by class. Sampling bias at the level of individual models or localities is also of less concern in my application to the extent it remains uncorrelated with my aggregate measure of class. Finally, drawing from the NHTS enables me to make a direct match with the FARS data using common make and model codes assigned by NHTSA.

Crash test data

NHTSA has performed safety tests of vehicles using crash-test dummies since the 1970's, with recent tests involving thousands of sensors and computer-aided models to determine the extent of life-threatening injuries likely to be received. The head-injury criterion (HIC) is a summary index available from the crash tests and reflects the probability of a fatality in actual accidents very close to proportionally (Herman [2007]). This is important for my application since equation (3.5) requires a measure that reflects proportional risk across vehicle types: if the HIC for compact cars is twice that of full size cars I should expect to see twice the number of fatal accidents all else equal.

I have assembled the average HIC by vehicle class for high-speed frontal crash tests conducted by NHTSA over the period 1992-2008.[21] These tests are meant to simulate typical collisions with fixed objects (such as concrete barriers, posts, guardrails, and trees) that are common in many fatal single-car accidents. The values for each class are included in Table 1. Single-vehicle accidents in small pickup trucks, the most dangerous class, are

[21] Specifically, I include all NHTSA frontal crash tests involving fixed barriers (rigid, pole, and deformable) and a test speed of at least 50 miles per hour. This filter includes the results from 945 tests.

nearly twice as likely to result in a fatality as those occurring in large sedans, the safest class, all else equal.

The crash test data is more difficult to defend than my other sources since it relies on the ability of laboratory tests to reproduce typical crashes and measure injury risks. I therefore offer an alternative specification in Section 7 that abstracts altogether from crash-test data. It produces similar results but offers less precision since it places more burden on cross-equation restrictions.

Summary statistics

I define 10 vehicle classes spanning the U.S. passenger fleet, including various sizes of cars, trucks, SUV's, and minivans. Table 1 provides a list and summary of fatal accident counts, reflecting fatalities both in the vehicle and those of other drivers in accidents. The VMT data is summarized in column 3, displaying the total annual miles traveled in each class. Column 4 describes single-vehicle fatality rates per billion VMT while the final column displays the HIC data for each class. The different patterns in risks measured by the HIC and fatality rates observed in the data highlight the importance of controlling for driver location and behavior in the model.

Table 2 describes the data on fatal accidents, now divided according to bin s. The first three columns indicate total fatal accidents in my sample, summarizing only one and two-car accidents. Column 4 shows variance at the weekly level used in estimation. Columns 5 and 6 respectively display the fraction of accidents that involve one car and where the fatality is in a light truck. More than half of fatal accidents involve only one car. Finally, the last two columns show the accident types with the highest relative frequency. Pickups are involved in the most single-car accidents per mile everywhere except in the highest income cities. Two-car accidents are more varied, with luxury vehicles involved in the evening and at night, and compacts much more likely to have a fatality (the vehicle with the fatality is listed first). A summary of the accident rates in all 100 possible combinations of classes is provided in Table 3, and is discussed in detail in Section 5 below.

5. Estimation

Building on the model of accident counts outlined in Section 3, this section now turns to identification and estimation of the parameters. I will take x_i, n_i, Y_{is}, and Z_{ijs} as data and wish to estimate α_i, β_{ij}, and λ_s. The equations in Section 3 representing single and multi-car accidents are again:

$$E(Y_{is}) = n_{is}\alpha_{is}\lambda_s x_i \tag{5.1}$$
$$E(Z_{ijs}) = n_{is}n_{js}\alpha_{is}\alpha_{js}\beta_{ij} \tag{5.2}$$

Estimation first requires a reduction of the parameter space: since I do not observe miles driven, n_{is}, at the bin s level I also cannot estimate each α_{is} separately. Instead, I combine the effect of n_{is} and α_{is} into a single parameter for estimation: $\delta_{is} \equiv n_{is}\alpha_{is}$. This approach maintains flexibility across location and class in estimation, while still permitting calculation of average risks by class when applying data on n_i *ex post*.[22] δ_{is} is identified up to a constant so there are $(10 \cdot 18 - 1) = 179$ of these flexible bin-by-class effects. The remaining parameters in the model are the 100 β_{ij}'s and 18 λ_s's. I observe the HIC score by class, x_i, and weekly counts of Y_{is} and Z_{ijs}; pooling 3 years of data provides 2,808 observations on each of the 110 fatal accident types for a total of 308,880 counts.

The model for estimation is:

$$\begin{aligned} &Y_{is} \sim \text{Poisson}(\omega_{is}) \\ &E(Y_{is}) = \omega_{is} = \delta_{is}\lambda_s x_i \end{aligned} \tag{5.3}$$

$$\begin{aligned} &Z_{ijs} \sim \text{Poisson}(\mu_{ijs}) \\ &E(Z_{ijs}) = \mu_{ijs} = \delta_{is}\delta_{js}\beta_{ij} \end{aligned} \tag{5.4}$$

[22] In particular, define n_i as the aggregate quantity (miles) for class i such that $n_i = \sum_s n_{is}$. Then $\sum_s \delta_{is} \Big/ n_i = \sum_s n_{is}\alpha_i \Big/ n_i = \alpha_i$.

The Poisson also dictates the variance of the observed counts, coming from the underlying binomial.[23] A generalization allowing an additional source of error is discussed below and produces very similar estimates in this setting.

Identification

Equations (5.3) and (5.4) are estimated in combination since neither of the two is identified in isolation: λ_s and δ_{is} cannot be separated in the first equation and δ_{is} and β_{ij} cannot be separated in the second. This reflects the key identification challenge: δ_{is} contains unobserved location and driving safety behavior which we wish to separate from risks due to the vehicles themselves (assuming an average driver and location) represented by β_{ij}.

Algebraically, separate identification of the parameters is possible via the presence of δ_{is} in both equations and the implied cross-equation restrictions. More intuitively, the assumption I need is that factors causing δ_{is} to differ (for example a tendency to drive recklessly or in dangerous locations) simultaneously influence risk of fatal single car accidents and fatal accidents with other cars. The λ_s parameters in (5.3) allow me to importantly weaken the strength of this assumption: factors contributing to single-car accidents in bin s that are common across classes are absorbed by λ_s.[24]

As an example, consider the role of dangerous rural highways: to the extent the number of single-car fatalities across classes in the rural bins is higher than would be predicted by the HIC scores, x_i, this will be captured in a large λ_s for those bins. If after taking out λ_s there remains a particular excess of fatal accidents among pickup trucks (which is the case in the data), then the δ_{is} parameters on pickup trucks will be increased. The assumption across equations is that this part of the variation, the risk multiplier specific to pickup trucks in the rural bins, also multiplies the risks they impose in two-car

---

[23] If $Y_{is} \sim \text{Poisson}(\omega_{is})$ then $\text{Var}(Y_{is}) = E(Y_{is}) = \omega_{is}$.

[24] Separate identification of an increase in λ_s from an increase in each of the δ_{is}'s for that bin comes via the overall frequency of single- vs. multi-car fatal accidents in that bin and the fact the β_{ij} is defined independent of bin: if a particular bin experiences an unusually large number of single-car accidents but an average number of multi-car accidents then a large value of λ_s and average values for the δ_{is}'s will fit the data best.

accidents. Since λ_s is common to classes within a bin my assumption is violated, for example, if the connection between single- and two-car accidents is stronger for some classes than others. A variety of alternative bin structures, influencing the degree and type of flexibility allowed by the λ_s parameters, are explored in Appendix B. The policy results are robust.

As a final note on identification, it may be useful to consider a simplified version of (5.3) that abstracts from the λ_s parameters: We would then have $E(Y_{is}) = \delta_{is} x_i$. In this setting δ_{is} would just be the count of single-car fatalities divided by the number we would expect to have based on crash-test results. In this sense the δ_{is} parameters are a measure of residual riskiness, allowing them to include unobserved variation in geography and driving at the bin-class level.

Overdispersion and Error

The Poisson specification above assumes that the only source of deviation in the count of fatal accidents across observations comes from an underlying low-probability binomial event, here the binomial occurrence of a fatality for each vehicle mile driven. However, additional sources of error can create overdispersion where the counts will differ by more than the underlying binomial would imply. The negative binomial generalization of the Poisson allows overdispersion to be modeled explicitly, adding an error component and associated variance parameter. I follow the specification of the negative binomial model given in Cameron and Trivedi (1986); the full model appears in Appendix A along with further discussion of error.

In my application, estimation of the negative binomial model produces parameter estimates that are nearly unchanged relative to the simple Poisson (a comparison appears in the appendix). However, a likelihood-ratio test does reject the Poisson and so I report results from the more general negative binomial throughout.

Results from a restricted model

It provides a useful comparison to first consider a restricted model where driving behavior and underlying engineering safety are combined into a single parameter. The next subsection will display the full model, where the effects are separated.

In the restricted model I retain the full set of fixed effects on bins *s* and vehicle interactions β_{ij}, but simply drop the terms for driver behavior:

$$
\begin{aligned}
&Z_{ijs} \sim \text{Poisson}(\tilde{\mu}_{ijs}) \\
&E(Z_{ijs}) = \tilde{\mu}_{ijs} = \tilde{n}_{is}\tilde{n}_{js}\tilde{\beta}_{ij}
\end{aligned}
\qquad (5.5)
$$

n and β are defined as before; the ~ modifier indicates the restricted model.

Table 3 presents the restricted estimates of $\tilde{\beta}_{ij}$. The parameters have a simple interpretation: they are the total fatality rates in interactions between each pair of classes. The most dangerous interaction in the table occurs when a compact car collides with a large pickup truck, resulting in 38.1 fatalities in the compact car per billion miles that the two vehicles are driven. The chance of a fatality in the compact in this case is about 3 times greater than if it had collided with another compact, and twice as large as if it collided with a full-size sedan. However, this table cannot address the possibility that some classes contain more fatalities due to dangerous driving or locations, as opposed to any inherent risk in the engineering.

Biases of this sort are particularly evident when examining minivans in Table 3. Minivans are much larger and heavier than the average car yet appear to impose very few fatalities on any other vehicle type, even compacts. This is noted as a puzzle in the engineering literature (Kahane [2003]) since simple physics suggests minivans will cause considerable damage in collisions. I find below that this is resolved by allowing flexibility in driving behavior; minivans tend to be driven much more safely which accounts for the low rate of fatalities.

Results from the full model

By combining (5.3) and (5.4) my full model is able to separate the accident rates shown in Table 3 into two pieces: The portion attributable to driver location and behavior, and the portion that comes from the physical characteristics of the vehicles themselves. The semi-parametric form allows me to be agnostic about which physical attributes of the vehicles cause the changes in underlying safety; the influence of any one characteristic of interest (for example vehicle weight, or category definition as a light truck) can be easily calculated *ex post* from my full matrix of estimates.

My central estimates appear in Table 4. The first row displays estimates of α_i, or the average across bins of the driving safety risks among people who select vehicles in each of the ten classes. Average safety is normalized to unity and standard errors appear in parentheses. For easier comparison, I also display 95% confidence intervals graphically in Figure 1. I find that minivan drivers are the safest among all classes, with accident risks that are approximately 1/3 of the average. This is due both to driving behavior and the locations and times of day that minivan owners tend to be on the road. Small SUV drivers also have very low risk for fatal accidents, about half of the average. Small SUV's tend to be driven in urban areas (which are much safer than rural areas in terms of fatal accidents) and are among the more expensive vehicles. Pickup trucks are driven significantly more dangerously than SUV's of similar sizes, also intuitive given their younger drivers and prevalence in rural areas. Among passenger cars, large sedans are driven somewhat more dangerously than other car types. Again the urban-rural divide may explain some of this (there are more compacts in cities) as well as the higher average age of large sedan drivers.

The next ten rows of Table 4 are my estimates of the underlying safety across all vehicle interactions. The fatality rates shown are per billion miles traveled and represent a situation where driving behavior is fixed at the average in both vehicles. The miles-driven weighted sum of the β parameters is scaled to match the predicted number of fatalities overall, allowing comparison with the restricted model in Table 3. The change when moving from the restricted to full model is determined by the interaction of the α terms in the row and column: if both classes have unusually high α parameters, for example, the β

coefficient in the full model will be much smaller. I also plot the estimates of β_{ij} in Tables 3 and 4 against one other to demonstrate the overall pattern of changes: this appears in Figure 2 with the changes in pickups and minivans highlighted.[25]

A number of key differences in β_{ij} appear relative to the summary of accident rates shown in Table 3: without including differences in driving behavior large pickup trucks appear much more dangerous to other drivers than large SUV's (compare columns 7 and 9 of Table 3). After correcting for driving safety, the two classes of vehicles now appear similar (columns 7 and 9 of Table 4). This is an intuitive result in terms of physical attributes: large SUV's and large pickups have similar weight and size, often being built on an identical light truck platform. Minivans now also look like the light trucks that they are based on (in fact becoming statistically indistinguishable from them in most accident combinations). This validates engineering predictions based on weight and size, resolving the puzzle of why they appear in so few fatal accidents.

β_{ij} and the effects of vehicle weight

While this paper focuses on the policy implications of driver behavior combined with engineering safety, an examination of the engineering coefficients in isolation can also be useful as a check relative to the existing literature: much of the related work in engineering and economics has focused on carefully measuring the physical effect of vehicle weight on accident fatalities, controlling away driver behavior. In my model these effects should appear within the β matrix, though will be only a rough measure due to aggregation.

Changes in β_{ij} across the columns in Table 4 can be interpreted to reflect the external effect of a class; that is, the average number of fatalities that each class imposes on the other vehicle involved in an accident after driver behavior has been removed. Similarly, changes across the rows of Table 4 may be interpreted as the internal effect of

[25]Regressing the β parameters from the full model on those from the restricted model yields an R-squared of 0.47, suggesting the degree of variation missed by the restricted model.

that class on safety. To reduce these effects to the dimension of average weight in each class I fit the following simple equation by least squares:

$$\ln\left(\beta_{ij}\right) = a + b \cdot weight_i + c \cdot weight_j \tag{5.6}$$

where $weight_i$ measured in thousands of pounds for the class with the fatality (so $-b$ is the protective effect) and $weight_j$ is the average weight of the class without the fatality (making c the increase in risk to others). The estimate of c is 0.46 (standard error 0.065), suggesting that 1000 pounds of weight increases the number of fatalities in other vehicles by about 46%. The average protective effect given in b suggests each 1000 pounds of vehicle weight reduces own risk by 54% (standard error 6.5%). For context, the average weight in my sample is about 3500 pounds with a standard deviation of 800. Among my ten classes, large pickups are on average 2000 pounds heavier than compacts.

Evans (2001) estimates both the external and internal effects of vehicle weight using differences in the number of occupants in the striking and struck car. This strategy helps avoid a host of selection issues, since it allows weight to vary holding all other attributes of the vehicle fixed. He finds that 1000 pounds increases external risk by 42% and decreases own risk by 40%.[26] Kahane (2003) focuses on own safety risk: for passenger cars the central estimate of the protective effect is 44% per 1000 pounds of weight.[27] Kahane's estimates for light trucks, in contrast, are not robust and vary between -30% and +70% depending on accident type and vehicle size. Kahane speculates in his report that the difficulty in getting consistent estimates for light trucks may be due to selection by driver type. I now have evidence to support this: the selection effects I find among different types of light trucks are much stronger than those among passenger cars.

Anderson and Auffhammer (2011) isolate the effect of weight by conditioning on the occurrence of an accident (either fatal or not) and controlling for observable

[26] In particular, they estimate that each adult occupant adds 190 pounds on average and that striking vehicles with an extra adult occupant increase the fatality risk in the other car by 8.1%.

[27] The report includes a very large number of estimation strategies; the central statistic I quote for cars is taken from the conclusion to Chapter 3 and the results for trucks from Chapter 4.

characteristics.[28] They find that 1000 pounds of weight increases external risk by 47%. The rough estimate of the weight externality contained in my β_{ij} parameters is very similar, suggesting that at least along the dimension of vehicle weight the structure I impose in equations (3.4) and (3.5) has not restricted the underlying pattern in the data.

Anderson and Auffhammer use their findings to investigate the ability of gasoline taxes and weight-based mileage taxes to correct the weight externality in the fleet. In contrast, my approach allows me to consider accident risk in counterfactual fleets where the composition of vehicles and distribution of drivers across those vehicles have changed. This is ideal for analysis of the U.S. CAFE standard and will be the focus of the policy simulations below. The two papers also take quite distinct approaches on empirical identification: here it comes from the relation between single- and multi- vehicle accidents, permitting considerable flexibility in the correlation between unobserved driver characteristics and class.

6. Policy Simulations

An economic analysis of safety, fuel economy, and fleet composition turns on three factors: The underlying engineering causes of fatal accidents, the driving risk of the individuals who choose different vehicle types, and the re-optimization of vehicle choices that occurs due to the regulation. I recover the first two of these as empirical estimates in my framework above. The third, modeling which individuals change their car choice as a result of the standard, is included as the first stage of the simulation here.

Simulating vehicle choice begins with a measure of the shadow costs that various types of fuel economy policy will impose: implicitly, existing CAFE policy increases the purchases of small cars and decreases the purchases of large cars in order to meet an

---

[28] Anderson and Auffhammer argue that conditioning on accident occurrence controls for most of driver selection such that the remaining fatality risk can be attributed to the vehicle. Notice that this can remain consistent with the large differences I find in α_i: since I condition on miles driven α_i will include a tendency to get in more accidents per mile (Anderson and Auffhammer suggest this is the dominant component) and also allow a tendency toward increased severity once an accident has occurred.

average target. Policy also creates an incentive for technological change that I am assuming does not alter safety in itself; I instead focus on the changes in fleet composition. All of my empirical measures are per-mile driven, and that continues to hold in simulation. The vehicle choice model assumes constant own and cross- price elasticities of demand taken from the literature, and that consumers re-optimize based on the shadow costs present under different types of fuel economy standard.

The behavior of drivers, a key focus of this paper, also enters the simulation. I first assume that drivers carry their residual term with them as they switch vehicles. For example if a minivan driver switches to a large sedan, that will lower (all else equal) the fatality rate per mile in sedans. On the other hand, if a pickup truck driver switches to the same sedan that would increase the fatality rate per mile in sedans. Simulating a movement of the residual with the driver assumes that exogenous characteristics of drivers make up most of the safety residual (safety of nearby roads, geography, age, income, alcohol use, children in the vehicle, etc.).

However, Peltzman (1975) points out that larger, safer vehicles should induce more risk-taking behavior. Gayer (2004) also makes the case that light trucks and SUV's are more difficult to drive, working in the same direction as the Peltzman effect.[29] In my context the Peltzman effect means that a portion of the safety residual should stay with the vehicle class even as drivers re-optimize. I compute an upper bound on these effects below: intuitively, Peltzman-type effects make all fuel economy standards look better on safety since we are now arguing that movement to smaller vehicles causes an improvement in driving behavior on average. Importantly my main policy conclusions, including the adverse effect of the current standard and the improvement offered by a unified standard, will remain fully robust to this alternative model.

Finally, the farther out of sample I wish to look in simulation (i.e. very extreme changes to the fleet) the more strain is placed on the empirical estimates. Fortunately, there is a substantial amount of variation in the fleet already included in the data: For

[29] The recent widespread adoption of unibody SUV designs and electronic traction and stability control may reduce this effect.

example the fraction of the fleet that are large pickup trucks varies by more than factor of two across bins *s*.[30] The changes as the result of fuel economy rules span only a small piece of this variation.

Simulation Model

I begin with a set of estimates for own and cross-price elasticities of demand among the 10 vehicle classes. The central-case elasticities I use are shown in Table 5 and come from Bento et al (2009); alternatives will be explored in Section 7 and Appendix D. I will analyze an improvement in fuel economy of 1.0 miles-per-gallon (MPG) overall. My model of fleet composition and safety needs to include just the portion of the improvement we expect to come via composition: to remain conservative, I will assume that only 0.1 MPG comes through composition and allow the remaining 0.9 MPG to come via other changes, for example improved engine technologies. Alternative assumptions on this division can be easily accommodated by scaling the results in the tables below.[31]

The matrix of elasticities, combined with the shadow tax implicit in fuel economy regulation, uniquely determine the pattern of vehicle choices that will create this 0.1 MPG improvement in the fleet.[32] I assume that the gain in MPG is realized throughout the new and used fleets, meaning the results below should be taken as long run. Table 6 displays the shadow taxes under each of the three policies I consider.

1) Extension of the current CAFE rule

The shadow tax in this case is proportional to fuel economy within the light truck fleet and within the car fleet. This means that large pickups receive a shadow tax while small pickups receive a shadow subsidy. Similarly large cars receive a shadow tax while compacts receive a shadow subsidy. There is no incentive to switch from trucks and

---

[30] It ranges from 10% (high-income, urban, daytime) to 22% (low-income, rural, night).
[31] If we define C as the desired alternative role of fleet composition the scaling factor to apply in the tables is $10 \cdot C$. If half of the gain is expected via composition, for example, the safety impacts per MPG would be 5 times larger.
[32] Average fuel economy regulation places a shadow tax on vehicles that fall below the average requirement and a shadow subsidy on vehicles that are more efficient than the requirement.

SUV's into cars with this policy, since they are regulated by separate average requirements.

2) Single standard

Here the shadow tax is very simple: The least efficient vehicles receive the highest tax and the most efficient ones the highest subsidy. All are in proportion to fuel economy. In general trucks receive a shadow tax (the worse their fuel economy the more so) and cars receive a shadow subsidy.

3) Footprint-based CAFE standard

This more complicated policy targets fuel economy for vehicles based on their wheelbase and width. Large footprint vehicles are given a more lenient target, leaving little or no incentive for manufacturers to change the composition of vehicle types they produce. The only residual effect on fleet composition will be for classes that are either particularly efficient relative to their footprint (non-luxury cars) or particularly inefficient relative to their footprint (SUV's). This implies relatively little switching across vehicle types and therefore only small changes in safety. The aggregation up to class level in my model presents a caveat that is important here: if the correlation between weight and fuel economy is low for vehicles within the same class, then the footprint standard may cause more finely detailed compositional changes that I cannot observe.

The main simulation uses the elasticities, shadow costs, and estimates from the safety model above to calculate the final composition of the fleet under each policy alternative and also track types of drivers as they switch across vehicles. Depending on which types of drivers are switching into the smaller vehicles their accident rates per mile can either rise or fall. For example: If the policy causes a lot of large-pickup drivers to now buy small SUV's instead, I would predict that the average driving safety behavior in small SUV's worsens. The small SUV class will now contain the relatively safe, urban drivers it originally included, and now also add some drivers from the more dangerous category that formerly owned large pickups.

More formally, I compute the updated driver behavior, $\hat{\alpha}_i$, by taking a quantity-weighted average of the safety characteristics of drivers from all the classes who have switched into class *i* as a result of policy. This is combined with those who choose class *i* both before and after the regulation. The predicted number of fatalities under the new policy scenarios is given by:

$$\hat{Z}_{ijs} = \hat{n}_{is}\hat{n}_{js}\hat{\alpha}_i\hat{\alpha}_j\beta_{ij} \tag{6.1}$$
$$\hat{Y}_{is} = \hat{n}_{is}\hat{\alpha}_i\lambda_s x_i \tag{6.2}$$

where $\hat{\alpha}_i$ is the new driver safety residual and $\hat{n}_i$ reflects the new fleet composition induced by the policy. In constructing the counterfactual it is also important to note that α_i is an average of underlying α_{is} parameters that can vary by bin. I implicitly assume here that the cross-price elasticities in Table 5 apply in all bins, so that the average switcher from each class is still accurately described by the average α_i.

Simplifying assumptions

In order to keep the analysis tractable I abstract from issues of scale and accidents outside the passenger fleet as follows:

i) Commercial vehicles: I assume that the fleet of commercial vehicles (mainly heavy trucks for which a commercial driver's license is required) remains fixed since they are not covered by CAFE regulation. Fatalities occurring in passenger vehicles colliding with these commercial vehicles make up about 8.4% of fatalities (NHTSA, 2009) and I scale these using the same risk factors I estimate for single-car accidents.[33] If the relation between class and fatality risk is less strong for accidents with commercial vehicles, in the extreme keeping fatalities constant for those accidents, the magnitude of the changes I

[33] This approximation relies on the much larger mass of commercial trucks meaning collisions with them resemble collisions with fixed objects.

estimate below would be reduced. The difference this assumption makes is small, however, when compared with the estimates overall.[34]

ii) The scale of the fleet and miles driven: It may be that fuel economy rules will change the total number of cars sold (likely decreasing it) or the number of miles driven (likely increasing that in a "rebound" effect).[35] I focus here on fatalities per mile driven in order to keep the simulation transparent: to the extent that either the increase in overall miles or decrease in fleet size is important it will scale total fatalities either up or down. The comparison in policy provisions that I focus on is unaffected by changes in overall scale.[36]

iii) Pedestrians and bicyclists: About 14% of fatalities involving passenger vehicles are pedestrians and bicyclists. These fatality rates are nearly identical among cars and light trucks, consistent with the observation that the mass of the passenger vehicle is many times larger regardless of its class.[37] I therefore assume a constant rate of fatal accidents involving pedestrians. To the extent that smaller vehicles can reduce pedestrian fatalities – for example because of better visibility when reversing – it will serve to accentuate the benefits of the uniform policy that I identify below.

Results of policy simulations

The results of the three main policy simulations are contained in Tables 7 through 9. The standard errors reflect the estimates of the safety parameters made in this paper; the

[34] The largest policy effect I simulate is roughly 0.5% of the base fatality rate; interacted with the 8.4% of accidents involving commercial vehicles this only reduces the results below by about 6%. The adjustment is smaller for the other cases.

[35] A decrease in quantity might come from cost increases as fuel-saving technologies are introduced. An increase in miles is known as the rebound-effect; better fuel economy results in cheaper miles at the margin.

[36] Differential changes in driving across vehicle types will have more complicated effects and an extension to the paper could involve a richer simulation model to account for this. These effects would not change the estimation strategy or empirical results.

[37] Pedestrian and cyclist fatalities in my data are 2.82 per billion miles for cars and 2.81 per billion miles for light trucks. Within trucks, fatality rates are somewhat higher for larger vehicles. Surprisingly, the opposite effect holds within cars: larger vehicles have lower pedestrian fatality rates.

hypothetical changes in fleet composition are treated as deterministic. Appendix D and the sensitivity analysis explore the effect of error in the elasticities and an alternative source for the elasticities data.

1) Increment of 1.0 MPG to the current CAFE rules:

The left panel of Table 7 displays the change in total traffic deaths that are predicted using the restricted model, where driving behavior is not estimated. The restricted model suggests that CAFE offers an improvement in safety: 135 lives would be saved.

A very different picture emerges when I use the full model, including the selection on driving behavior at the class level. The central estimate is that the increment to CAFE will result in 149 additional traffic-related fatalities per year. The final row applies the value of statistical life figure used in EPA benefit-cost analyses to convert this change in risk to dollar cost, exceeding one billion annually. If about 3.1 billion gallons of gasoline are saved this translates to 33 cents per gallon. Placing this in context, an external cost of 25 dollars per ton CO_2 amounts to 22 cents per gallon of gasoline. Parry and Small (2005) include damages from local air pollution emissions of about 16 cents per gallon.

It is straightforward to see the intuition behind the reversal in sign: large SUV's and pickups (and large sedans) cause and experience a lot of fatal accidents in the data. The naive restricted model assumes that when you take away these large (and seemingly dangerous) vehicles an improvement in safety results. Unfortunately I must argue that the picture is not so favorable: much of the danger in the larger vehicle classes appears to be due to their drivers, not the cars themselves. When we move those people into smaller vehicles it does not diminish the risk, and in some cases can even magnify it since smaller vehicles do more poorly in most single-car accidents.

It is important to point out that the driver effects here are not all habits that we would fault the drivers themselves for (like running through traffic signals). A significant portion is simply geography and the urban-rural divide: drivers who currently choose large vehicles tend to live in rural areas, where accident fatality rates are already very high. As rural drivers change to smaller vehicles the dangers of accidents on rural highways remain. These are very often single-car accidents, as reflected in the composition of additional fatalities I predict. Finally, an important limitation in these results comes from the aggregation I make across classes: substitution within a class to different engine sizes or technologies may not have a large impact on safety, but there is also considerable variation in characteristics like weight and volume within a class that do influence safety. More detail on class definitions and aggregation appears in Appendix E.

2) Unified standard achieving a 1.0 MPG improvement

Table 8 presents results under a unified standard, which has a strikingly different effect from an increment to current CAFE rules. My full model shows an increase of only 8 fatalities per year under a unified standard. A zero change lies within the confidence bounds. This represents a highly statistically significant improvement over current CAFE rules and comes as the result of two effects canceling each other out in the fleet:

The first effect reiterates the undesirable outcome I find in the first experiment, that is, changes within the car fleet and within the truck fleet lead to smaller and lighter vehicles and increase the number of fatalities.

Recall though that the unified standard adds a second incentive: It encourages switching away from light trucks and SUV’s and into cars. This second effect improves overall safety substantially. There are aspects of light trucks (for example the height of their center of mass) that make them more dangerous vehicles than cars, even after controlling for their drivers. My model is able to measure the importance of the difference between light trucks and cars, and then compare it with the deterioration of safety within the car and truck fleets also resulting from the CAFE standard.

3) Footprint-based standard

Table 9 presents results under the footprint-based standard that is currently coming into effect. The standard discourages most types of composition changes by design, with the regulatory brief stating: "With the footprint-based standard approach, EPA and

NHTSA believe there should be no significant effect on the relative distribution of different vehicle sizes in the fleet." (NHTSA 2010)

The most significant compositional changes likely to remain are a modest movement away from SUV's and into pickup trucks and cars; this is due to the relatively small footprint of SUV's relative to their fuel consumption.[38] My full model shows a very small deterioration in safety from the footprint standard, with an increase of only 6 fatalities per year.

It is important to point out that these small safety effects come paired with large efficiency costs: Fuel savings under the footprint standard must be accomplished almost exclusively through engine technology. Movement to a smaller and lighter fleet is likely to be a much cheaper way to save gasoline and that channel is shut down by the new rules.

My results on the unified standard are encouraging in this regard: I show that savings in gasoline from movement to a smaller fleet can come with the same minimal effect on safety that appears under the footprint standard. As the U.S. presses toward even more fuel efficiency in coming years, changes in fleet composition will prove valuable (even necessary) and I show here that these changes can be made without severe safety consequences.

Comparison with a Gasoline Tax

While increases in the U.S. gasoline tax are typically met with strong political opposition, they do provide an efficient benchmark for reduction of gasoline use. CAFE rules typically compare very unfavorably to a gasoline tax and consideration of safety outcomes tends to further that conclusion:

Consider a gasoline tax that achieves half of its gasoline savings through fleet composition and half through a reduction in miles driven.[39] The portion saved via fleet

---

[38] The steepness of the slope set for the footprint standard determines the extent to which it shuts down switching. A slope too low to shut down switching altogether will still involve some change into smaller vehicles, though always less than under the unified rule as long as the slope is greater than zero.

[39] The literature differs on this fraction. Jacobsen (2010), for example, suggests a much greater portion may come from miles driven. Here the fraction will simply scale the results up or down: the larger the fraction from miles driven, the larger the safety gains from a gasoline tax.

composition will look just like the unified standard above: the shadow tax proportional to fuel economy is now an actual tax proportional to fuel economy. Following the results in Table 8, this type of movement in fleet composition is predicted to produce only small changes in fatalities. The remaining reduction in gasoline use from the tax would come from miles driven; the 1.0 MPG improvement considered above reduces gasoline use by about 3.8%, so half of that is a 1.9% reduction in miles driven from the equivalent gasoline tax. All else equal, this reduction in miles will create a reduction of about 500 fatalities per year, representing a dramatic improvement relative to any version of CAFE considered above.[40]

7. Alternative Models

Driver-vehicle interactions correlated with size

Peltzman (1975) argues that safer vehicles (in particular those with seatbelts installed) will be driven more aggressively as a result of the driver's tradeoff in utility.[41] Gayer (2004) presents evidence of a similar effect, where drivers in light trucks appear to take more risks or have less control when driving.

I am able to investigate this in the context of my model by further decomposing α_i into two pieces. The first portion is the part of α_i predicted by the own-safety of the vehicle.[42] In that sense it is an upper limit on the size of the Peltzman effect.[43] The second portion is whatever idiosyncratic variation remains in α_i and I will assume that continues to move with the driver. Table 10 presents the results of these policy experiments. The third column is my upper bound on the Peltzman effect over all driving safety residuals. The

[40] A portion of these gains is expected to be external; see Footnote 10.
[41] Subsequent empirical research has shown this effect may be small, see Cohen and Einav (2003).
[42] Using least squares regression of α_i on own-safety, calculated as the average of β_{ij} in a row.
[43] Unobserved countervailing selection in initial vehicle choice could potentially make the Peltzman effect even larger; these more extreme cases could still be modeled in simulation, possibly using estimates from other studies.

fourth column controls first for census region (there are more light trucks and dangerous roads in the west) and then applies the same method to divide the residual into two pieces.

As expected, the outcomes in Table 10 show that all fuel economy standards are improved if smaller vehicles indeed cause safer driving. However, even at the limit defined above I show that the existing fuel economy standard continues to have adverse effects on safety. Controlling for census region seems reasonable (as driver residence is unlikely to change with fuel economy standards), and the result becomes even closer to my central case. Adding further support to the importance of location, data on accident fault in Appendix C suggests that much of α_i is coming from factors like location or time of day rather than from behaviors associated with fault.

Finally, the improvement that can be offered by unifying the standard appears fully robust to the case where I allow a Peltzman-type effect. This is shown in the final row of the table. Because the difference in policies is maintained and overall safety is improved, we see that the unified standard even begins to offer substantial improvements in overall safety in the final column of the table.

Estimating driver behavior without using crash test data

It is possible to identify my empirical model (including the measurement of driver behavior by class) without the use of crash test data, relying instead on the physical properties of accidents. Accidents between two vehicles of similar mass and speed closely resemble accidents with fixed objects since both crashes result in rapid deceleration to a stationary position.[44] When vehicles of different mass collide, the heavier vehicle will decelerate more slowly (pushing the smaller vehicle back) which creates asymmetry in the degree of injuries.

My alternative identification strategy makes use of this property, setting risk in single car accidents proportional to the risk in accidents between cars of the same class, β_{ii}. The model described in Section 5 becomes:

---

[44] See Greene (2009). Each vehicle's change in velocity raised to the 4th power closely predicts injury severity.

$$E(Y_{is}) = n_{is}\alpha_{is}\lambda_s\beta_{ii} \quad (7.1)$$
$$E(Z_{ijs}) = n_{is}n_{js}\alpha_{is}\alpha_{js}\beta_{ij} \quad (7.2)$$

The restriction on the diagonal elements of β is sufficient for identification.

The first two columns of Table 11 provide a summary of results from my preferred specification in Section 5. The third column shows the results from estimating (7.1) and (7.2) above; the standard errors are much larger in this specification, reflecting the reduction in data available to the model. The results on existing CAFE and the unified standard confirm those in the main specification: the unified standard continues to offer a statistically significant improvement. The final row, for the footprint standard, now displays an improvement in safety in contrast to the near zero result in the central case. This effect stems from relatively high fatality rates in SUV-SUV collisions, which in this specification translate to larger estimates of their engineering risk and a gain when these classes are discouraged by the footprint standard.[45]

Alternative demand elasticities

The general pattern in the simulation, that fewer large vehicles and more small ones will be sold, is fundamental to a reduction in fuel economy. However, my simulation also embeds more subtle changes in substitution across classes. For example: Is a driver giving up a large SUV more likely to buy a small SUV or switch to a small pickup truck?

I first investigate the robustness of my simulation results by introducing a separate set of substitution elasticities, shown in Table 12. These are reported in Kleit (2004) and are also employed by Austin and Dinan in their 2007 study. The elasticities derive mainly from survey data on second-choices of new car owners, providing a very different view than the cross-sectional variation used to generate the elasticities in my main simulation.

[45] The fatality rate in matched SUV collisions is large relative to that expected given SUV crash tests into fixed objects, perhaps due to increased rollover risk in this type of collision. The resulting beta coefficients on SUV's are 13% larger here.

The fourth column of Table 11 summarizes the results under the alternative elasticities. My main findings remain intact, though the effectiveness of a single fuel economy standard at mitigating safety consequences is somewhat muted relative to my preferred model. For further analysis along these lines, Appendix D also returns to the Bento et al (2009) source for the elasticities and takes 50 different draws from the posterior density. The model in that paper is estimated relatively precisely, such that the results here remain robust to even the extreme draws.

Additional robustness checks

I also investigate the robustness of my findings in a number of subsamples of the data. Columns 3 through 5 of Table 13 summarize my main results in various subsamples, with total fatalities scaled by the number of observations used so that the columns are comparable.

1998 and newer model years

1998 was the first model year where both passenger and driver airbags were required in all new vehicles. Airbags dramatically alter safety risks, and if their presence also influences driving behavior or changes relative risks across classes we might expect a different set of results to emerge. My estimates, however, appear robust in this dimension.

Drivers under 55

There is evidence that elderly drivers may more often be the subjects of fatal traffic accidents due to their relative frailty.[46] This introduces a potential asymmetry in my model: Older drivers may place themselves at greater risk but don't necessarily impose this risk on those around them. I restrict my sample to driver fatalities among those less than 55 years old and find similar results in the aggregate outcomes for fuel economy rules.

[46] Loughran and Seabury (2007) investigate this issue in detail.

Clear weather

My simulations assume that the locational or behavioral factors influencing driver safety remain with the driver after the change in composition. A potentially important caveat has to do with weather. If a driver switches away from an SUV, for example, they may be less likely to drive in the rain or snow. I therefore experiment with a sample limited to fatalities that occur in clear weather (any weather condition, even fog or mist, is excluded). Notably, this only removes 10% of observations; 90% of fatal accidents occur in clear conditions. My results are again unchanged, suggesting that even if there is substantial behavioral response to weather conditions it would not be relevant to most accident fatalities.

8. Conclusions

I introduce a new empirical model of vehicle accidents that provides estimates of both the behavior of drivers and the underlying risk associated with engineering characteristics in a single framework. To my knowledge this is the first study to capture unobserved driver behavior and the impact of unobserved physical vehicle characteristics both within and across vehicle categories. The framework has application to fuel economy policy (the simulations performed here) and also to a much broader set of policy initiatives. I show that in the case of fuel economy, correctly accounting for driver behavior significantly alters conclusions about fleet composition and safety.

Two main effects appear in the empirical estimates. First, there is considerable diversity in driving behavior across vehicle classes: the most dangerous drivers (pickup truck owners) are nearly four times as likely to be involved in fatal accidents as the safest drivers (minivan owners) after controlling for the physical safety attributes of their vehicles. Second, controlling for driver safety produces estimates of the physical safety of interactions between vehicles that closely mirrors theoretical engineering results. Large and heavy vehicles are the safest to be inside during an accident but also cause the most damage to others. When reduced to the single dimension of vehicle weight, my estimates of the own and external effects of heavier vehicles match those in the literature closely.

I use these results to address the motivating question relating safety and fuel economy regulation. I find that the provision in existing CAFE regulation to separate light trucks and SUVs from passenger cars is harmful to safety: incrementing the standards by 1.0 mile per gallon causes an additional 149 fatalities per year in expectation. The increase in statistical risk would be valued at 33 cents per gallon of gasoline saved, with any additional injuries or property damage (assuming they are correlated with fatalities) further increasing the cost of this type of regulation.[47] Intuitively, my estimates measure the degree to which greater diversity in the vehicle fleet leads to more fatal accidents. Current CAFE standards, by encouraging light trucks while at the same time making passenger cars smaller and lighter, increase the diversity of the fleet.

In contrast, I find that a unified fuel economy standard has almost no harmful effect on safety. Two effects are operating in opposing directions: weight reductions increase risk while substitution away from light trucks makes the fleet more homogeneous. In contrast to the literature, my model can compare the relative importance of these two effects and I find they offset almost exactly under the shadow costs implied by a uniform fuel economy standard.

Extensions of the model here could address some of the remaining limitations and potentially uncover additional effects of interest: a more detailed disaggregation of car classes, for example by manufacturer, fuel economy, or footprint, could identify changes within the current class definitions and lead to additional insight on fuel economy rules. More detailed forecasts for the evolution of the fleet over time could reveal important short run impacts on safety, and a longer time series with finer detail on miles driven could similarly enhance the identification. Combining this with improved resolution on the location of miles driven by class might also allow further relaxation of the restrictions I impose here across equations, influencing the part of risk that enters both single- and multi-vehicle accidents.

[47] The gasoline savings here reflect only fleet composition changes, holding miles driven fixed. To the extent that a "rebound effect" increases miles driven, the safety cost per gallon saved would be even larger.

Table 1: Summary Statistics by Class

	Count of Accident Fatalities[1]		Total Miles	Single Vehicle	Crash Test
Class	Own Vehicle	Other Vehicle	Driven[2]	Fatality Rate[3]	HIC[4]
Compact	2812	1068	247.7	14.3	528.7
Midsize	2155	1280	249.7	11.3	491.4
Fullsize	733	507	83.2	10.2	353.9
Small Luxury	317	236	54.5	13.5	424.3
Large Luxury	364	307	50.8	11.9	469.3
Small SUV	719	1129	216.0	9.4	626.3
Large SUV	477	1379	148.9	12.8	531.2
Small Pickup	594	624	87.1	15.9	666.2
Large Pickup	716	2293	159.5	18.2	585.9
Minivan	469	532	126.7	4.9	577.9

[1] Two-vehicle accidents, annual average 2006-2008.
[2] In billions of miles per year (2008 National Household Transportation Survey).
[3] Fatal single car accidents per billion miles traveled, annual average 2006-2008.
[4] Results from NHTSA testing 1992-2008. The head-injury criterion (HIC) score has been shown to be closely and linearly related to fatality rates (when controlling for driver behavior, a doubling in the score should correspond to a doubling of fatality rates).

Table 2: Summary Statistics by Bin *s*

Density[1]	Income[1]	Time of Day	Fatalities (1 and 2 Car Accidents) 2006	2007	2008	Variance (Weekly)	Fraction 1-Car	Fraction Light Trucks	Greatest Relative Frequency[2] 1-Car Accidents	2-Car Accidents
Rural	Low	Night	705	673	582	3.92	0.882	0.557	Lg Pickup	Fullsize/Fullsize
		Evening	374	373	320	2.77	0.718	0.485	Lg Pickup	Fullsize/Fullsize
		Day	1574	1475	1310	6.66	0.643	0.519	Lg Pickup	Sm Pickup/Lg Pickup
	Medium	Night	501	518	414	3.47	0.883	0.537	Lg Pickup	Compact/Lg Lux
		Evening	254	257	210	2.16	0.756	0.535	Sm Pickup	Fullsize/Fullsize
		Day	1022	1003	897	4.97	0.585	0.498	Sm Pickup	Sm Pickup/Lg Pickup
	High	Night	341	308	266	2.71	0.897	0.460	Lg Pickup	Sm Lux/Sm Pickup
		Evening	150	133	144	1.65	0.728	0.478	Sm Pickup	Lg Lux/Lg Lux
		Day	639	645	540	4.02	0.550	0.459	Lg Pickup	Compact/Lg Pickup
Urban	Low	Night	587	570	532	3.53	0.827	0.528	Lg Pickup	Compact/Lg Pickup
		Evening	283	265	222	2.37	0.655	0.491	Lg Pickup	Sm Lux/Sm Lux
		Day	1133	1062	953	4.91	0.609	0.528	Lg Pickup	Sm Pickup/Lg Pickup
	Medium	Night	1038	995	946	4.83	0.822	0.491	Lg Pickup	Lg Lux/Lg Lux
		Evening	478	437	368	2.95	0.652	0.471	Lg Pickup	Lg Lux/Lg Pickup
		Day	1850	1671	1569	6.72	0.571	0.473	Lg Pickup	Compact/Lg Pickup
	High	Night	4234	4085	3565	11.96	0.766	0.380	Sm Lux	Compact/Sm Lux
		Evening	1490	1404	1229	5.96	0.599	0.385	Sm Lux	Compact/Lg Pickup
		Day	5786	5525	4801	14.16	0.511	0.386	Compact	Compact/Lg Pickup
All			22439	21399	18868	41.85	0.650	0.441	Lg Pickup	Compact/Lg Pickup

[1] Based on zip-code level classifications from the U.S. Census.
[2] Relative frequencies are calculated as accident counts within group divided by total miles traveled. A combination of vehicle popularity and driver behavior within group determines the accident with greatest relative frequency.

Table 3: Estimates of $\tilde{\beta}_{ij}$ in Restricted Model (No class-level driver safety effects)[1]

Vehicle *i*: / Vehicle *j*:	Compact	Midsize	Fullsize	Small Luxury	Large Luxury	Small SUV	Large SUV	Small Pickup	Large Pickup	Minivan
Compact	12.4 (0.5)	14.9 (0.5)	17.7 (1.0)	12.6 (1.0)	17.2 (1.2)	16.2 (0.6)	26.4 (0.9)	20.2 (1.0)	38.1 (1.1)	12.1 (0.6)
Midsize	8.8 (0.4)	11.8 (0.5)	12.9 (0.8)	9.2 (0.8)	12.8 (1.0)	11.2 (0.5)	20.4 (0.8)	16.5 (0.9)	30.5 (1.0)	8.9 (0.5)
Fullsize	8.7 (0.7)	11.9 (0.8)	16.0 (1.5)	8.8 (1.4)	14.9 (1.9)	11.6 (0.8)	19.0 (1.3)	17.4 (1.6)	30.6 (1.6)	9.8 (1.0)
Small Luxury	8.5 (0.8)	6.5 (0.7)	11.2 (1.6)	11.8 (2.0)	10.8 (2.0)	9.6 (0.9)	12.1 (1.2)	6.9 (1.2)	16.6 (1.4)	5.1 (0.9)
Large Luxury	6.6 (0.7)	8.7 (0.8)	11.6 (1.7)	6.1 (1.5)	11.2 (2.1)	10.3 (1.0)	20.4 (1.7)	13.3 (1.7)	22.9 (1.7)	8.2 (1.1)
Small SUV	3.6 (0.3)	4.2 (0.3)	4.6 (0.5)	4.2 (0.6)	6.8 (0.8)	4.3 (0.3)	7.9 (0.5)	4.9 (0.5)	12.2 (0.6)	3.4 (0.4)
Large SUV	4.2 (0.3)	4.2 (0.3)	3.8 (0.6)	3.7 (0.7)	5.2 (0.8)	3.5 (0.3)	7.9 (0.6)	5.4 (0.6)	11.1 (0.7)	3.7 (0.4)
Small Pickup	8.2 (0.6)	8.4 (0.6)	10.1 (1.2)	4.6 (1.0)	6.6 (1.2)	7.4 (0.6)	14.0 (1.1)	13.0 (1.3)	29.1 (1.5)	7.7 (0.8)
Large Pickup	4.8 (0.4)	5.2 (0.4)	5.9 (0.7)	4.5 (0.7)	6.3 (0.9)	4.4 (0.4)	10.1 (0.7)	7.4 (0.7)	21.5 (1.0)	3.6 (0.4)
Minivan	3.5 (0.3)	3.8 (0.3)	6.1 (0.8)	3.5 (0.7)	3.9 (0.8)	5.0 (0.4)	8.9 (0.7)	7.7 (0.8)	14.4 (0.9)	4.7 (0.5)

[1] Estimates are from the multi-car accident equation alone, with all class-level safety effects restricted to unity. The parameters and standard errors (shown in parentheses) are computed by maximum likelihood estimation of the negative binomial version of the model. Without single-car accidents or variation over bins there are 15,600 observations and the log likelihood is -19637. These coefficients provide a summary of fatal accident rates without controlling for driver behavior.

Table 4: Central Estimation Results[1]

	Compact	Midsize	Fullsize	Small Luxury	Large Luxury	Small SUV	Large SUV	Small Pickup	Large Pickup	Minivan
α_i: Driver Safety Behavior	1.14 (0.06)	0.98 (0.06)	1.25 (0.08)	1.19 (0.08)	1.05 (0.07)	0.65 (0.04)	1.06 (0.06)	1.09 (0.07)	1.45 (0.08)	0.39 (0.02)
β_{ij}: Fatality rate in vehicle *i*										
Compact	8.6 (1.0)	12.1 (1.4)	11.5 (1.5)	7.6 (1.1)	12.4 (1.7)	19.8 (2.4)	19.9 (2.4)	16.3 (2.0)	24.3 (2.9)	24.9 (3.2)
Midsize	7.1 (0.9)	11.0 (1.4)	9.7 (1.3)	6.6 (1.0)	10.7 (1.5)	15.9 (2.0)	17.7 (2.2)	15.1 (1.9)	22.2 (2.7)	21.0 (2.8)
Fullsize	5.6 (0.8)	8.9 (1.2)	9.5 (1.5)	5.2 (1.0)	10.1 (1.8)	13.0 (1.8)	13.0 (1.8)	12.5 (1.8)	17.5 (2.2)	18.2 (2.8)
Small Luxury	5.1 (0.8)	4.6 (0.7)	6.6 (1.2)	5.6 (1.2)	6.7 (1.5)	10.5 (1.6)	8.3 (1.3)	5.3 (1.1)	10.2 (1.5)	9.5 (2.0)
Large Luxury	4.8 (0.8)	7.3 (1.1)	7.8 (1.5)	3.8 (1.0)	8.3 (1.9)	13.1 (2.0)	15.9 (2.3)	11.2 (2.0)	15.5 (2.2)	17.5 (3.2)
Small SUV	4.4 (0.6)	6.0 (0.8)	5.1 (0.8)	4.7 (0.9)	8.7 (1.4)	9.1 (1.3)	10.2 (1.4)	6.7 (1.1)	13.3 (1.7)	11.8 (1.9)
Large SUV	3.1 (0.4)	3.7 (0.5)	2.6 (0.5)	2.5 (0.6)	4.0 (0.8)	4.6 (0.7)	6.2 (0.9)	4.5 (0.8)	7.3 (1.0)	7.9 (1.3)
Small Pickup	6.6 (0.9)	7.7 (1.1)	7.2 (1.2)	3.5 (0.9)	5.5 (1.2)	10.1 (1.5)	11.6 (1.6)	11.0 (1.7)	19.4 (2.5)	17.4 (2.8)
Large Pickup	3.1 (0.4)	3.8 (0.5)	3.4 (0.5)	2.8 (0.5)	4.2 (0.8)	4.8 (0.7)	6.6 (0.9)	4.9 (0.8)	11.1 (1.4)	6.4 (1.1)
Minivan	7.3 (1.1)	8.9 (1.3)	11.3 (2.0)	6.5 (1.6)	8.3 (1.9)	17.7 (2.6)	19.0 (2.7)	17.4 (2.8)	25.9 (3.4)	27.1 (4.7)

Negative binomial regression
Number of obs: 308880
Log likelihood: -89321
Wald chi2(297): 233212

[1] Estimates of α_i reflect driver safety risks by class. These are identified up to a constant and normalized here such that a value of one represents the average driver. Values larger than one reflect increased risk. β_{ij} are estimated rates of fatalities in car *i* (row) when colliding with car *j* (column) per billion miles traveled by average drivers. β_{ij} are scaled such that their VMT weighted sum equals the total predicted number of fatalities, making them comparable to the values in Table 3. The parameters and standard errors (shown in parentheses) are computed by maximum likelihood estimation of the negative binomial version of the model.

Table 5: Matrix of Own and Cross-Price Demand Elasticities by Class[1]

	Compact	Midsize	Fullsize	Small Luxury	Large Luxury	Small SUV	Large SUV	Small Pickup	Large Pickup	Minivan
Compact	-3.51	0.97	0.42	0.32	0.21	0.67	0.49	0.41	0.51	0.52
Midsize	0.80	-3.01	0.31	0.16	0.15	0.41	0.31	0.32	0.32	0.29
Fullsize	0.79	0.73	-4.94	0.14	0.21	0.31	0.44	0.30	0.45	0.30
Small Luxury	0.59	0.35	0.14	-5.15	0.15	0.46	0.16	0.13	0.24	0.16
Large Luxury	0.42	0.36	0.22	0.16	-4.18	0.24	0.22	0.10	0.21	0.12
Small SUV	0.76	0.54	0.19	0.28	0.14	-2.39	0.25	0.19	0.30	0.29
Large SUV	0.62	0.48	0.31	0.11	0.15	0.27	-2.95	0.19	0.37	0.21
Small Pickup	0.68	0.66	0.26	0.12	0.08	0.29	0.24	-3.96	0.23	0.18
Large Pickup	0.92	0.68	0.44	0.24	0.19	0.48	0.51	0.25	-2.81	0.43
Minivan	0.69	0.47	0.23	0.12	0.08	0.34	0.23	0.15	0.32	-3.31

[1] These elasticities are derived from Bento et al (2009) and are used in the central case of the policy simulations. I investigate the robustness of the results to alternative elasticities (see Table 12) and variance in the Bento et al estimates (see Appendix D).

Table 6: Average Fuel Economies and Shadow Taxes by Class

		Shadow Tax of Policy Increment[1]		
Class	Fuel Economy (MPG)	Increase current CAFE	Unified standard	Footprint CAFE
Compact	30.2	0.28	0.22	0.06
Midsize	27.0	-0.09	0.12	0.05
Fullsize	25.4	-0.31	0.06	0.06
Small Luxury	26.0	-0.22	0.08	-0.02
Large Luxury	23.8	-0.56	-0.01	0.00
Small SUV	24.1	0.37	0.01	-0.11
Large SUV	19.0	-0.44	-0.28	-0.14
Small Pickup	22.5	0.16	-0.07	0.02
Large Pickup	19.1	-0.41	-0.27	0.01
Minivan	23.4	0.29	-0.02	0.06

[1] The shadow taxes and shadow subsidies are placed by the fuel economy policy and differ according to the type of standard in place. They are proportional to the distance of each vehicle (in gallons-per-mile) from the applicable fuel economy target. The units are in thousands of dollars per vehicle, though only the resulting changes in composition of the fleet are relevant to the safety outcomes modeled here.

Table 7: Effect of an Increase in Current CAFE Rules on Total Traffic Deaths

	No driver effects[1]			*Full model*[2]		
	One car	Two car	Total	One car	Two car	Total
Compact	226.3	142.4	368.6	236.1	177.6	413.6
Midsize	-60.1	-75.4	-135.5	-51.3	-50.6	-101.9
Fullsize	-55.0	-57.0	-112.0	-55.1	-51.0	-106.1
Small Luxury	-30.8	-16.1	-46.8	-30.9	-13.4	-44.2
Large Luxury	-34.6	-25.6	-60.2	-34.6	-22.3	-57.0
Small SUV	78.4	16.4	94.8	142.4	45.3	187.7
Large SUV	-85.9	-27.1	-113.0	-85.8	-23.2	-109.0
Small Pickup	47.8	11.9	59.7	50.9	18.4	69.3
Large Pickup	-168.7	-54.6	-223.2	-171.4	-50.8	-222.3
Minivan	22.4	10.2	32.6	69.1	50.2	119.3
Total	-60.0	-75.0	**-135.0**	69.3	80.2	**149.5**
Standard error[3]			(6.1)			(9.4)
Cost of risk change (millions)[4]			-932			1031

[1] This case reflects the restricted model, where driving safety behavior is assumed constant across all classes; only the quantity of cars in each class changes. The change is expressed in fatalities per year with a one-MPG increment to fuel economy rules.

[2] Here the full model is used to predict changes in safety, including the parameters that account for differences in driving safety behavior across classes.

[3] The variance of the simulation results G is approximated by $(\partial G(\hat{\gamma})/\partial\hat{\gamma})' Var(\hat{\gamma})(\partial G(\hat{\gamma})/\partial\hat{\gamma})$ where $Var(\hat{\gamma})$ is the variance-covariance matrix from estimation of the negative binomial model and γ includes the λ, β and δ parameters.

[4] Annual costs in millions of dollars, calculated as the change in total fatality risk multiplied by a value of statistical life of $6.9 million, following EPA benefit-cost analysis.

Table 8: Effect of a Unified Fuel Economy Standard on Total Traffic Deaths[1]

	No driver effects			*Full model*		
	One car	Two car	Total	One car	Two car	Total
Compact	167.8	105.7	273.5	153.3	97.7	251.0
Midsize	39.4	7.5	47.0	44.7	13.9	58.6
Fullsize	6.7	-1.5	5.2	5.6	-1.6	4.0
Small Luxury	5.7	0.8	6.5	4.9	0.7	5.6
Large Luxury	-2.6	-5.6	-8.1	-2.1	-4.8	-6.9
Small SUV	-12.5	-11.8	-24.3	-0.3	-6.7	-7.0
Large SUV	-62.1	-19.6	-81.7	-62.1	-19.1	-81.2
Small Pickup	-32.6	-20.4	-53.0	-32.3	-19.7	-52.0
Large Pickup	-122.4	-39.2	-161.6	-122.9	-38.9	-161.8
Minivan	-5.6	-10.0	-15.6	2.0	-3.8	-1.8
Total	-18.0	5.9	**-12.1**	-9.3	17.8	**8.5**
Standard error			(3.8)			(4.3)
Cost of risk change (millions)[2]			-84			59

[1] The unified standard induces two kinds of changes in the fleet: i) Small vehicles replace large ones within the car and light truck divisions. ii) Light trucks overall (the second set of five classes) replace cars overall (the first group). The change is expressed in fatalities per year for a one-MPG increment to fuel economy rules and standard errors are calculated as in Table 7.
[2] In millions of dollars annually.

Table 9: Effect of a Footprint Fuel Economy Standard on Total Traffic Deaths[1]

	No driver effects			*Full model*		
	One car	Two car	Total	One car	Two car	Total
Compact	45.6	31.4	77.0	38.0	24.4	62.4
Midsize	15.9	8.5	24.4	15.0	6.9	21.9
Fullsize	8.9	6.7	15.6	7.3	5.0	12.3
Small Luxury	-3.4	-1.9	-5.3	-3.9	-2.3	-6.2
Large Luxury	-0.5	-1.2	-1.7	-0.8	-1.5	-2.2
Small SUV	-31.6	-12.5	-44.1	-31.3	-12.7	-44.0
Large SUV	-32.6	-8.7	-41.3	-32.6	-8.9	-41.5
Small Pickup	1.8	0.3	2.1	0.9	-0.4	0.5
Large Pickup	-4.1	-2.0	-6.2	-10.0	-4.0	-14.0
Minivan	4.1	2.2	6.4	10.3	6.8	17.1
Total	4.2	22.7	**26.9**	-7.1	13.4	**6.3**
Standard error			(1.3)			(1.5)
Cost of risk change (millions)[2]			185			43

[1] A footprint standard (by design) involves much smaller changes in the composition of the fleet than either of the first two policies simulated. The changes in accident fatalities are similarly small. The change is again expressed in fatalities per year for a one-MPG improvement and standard errors are as above.
[2] In millions of dollars annually.

Table 10: Peltzman Effects and the Influence of a Driver-Vehicle Specific Residual[1]

	No driver effects	Full model (central)	Peltzman effect (upper limit)	Peltzman within census divisions (upper limit)
Current CAFE within fleet	-135.02 (6.15)	149.47 (9.36)	69.80 (9.36)	101.72 (9.36)
Unified standard	-12.14 (3.81)	8.50 (4.35)	-57.00 (4.35)	-64.43 (4.35)
Footprint-based standard	26.88 (1.28)	6.27 (1.52)	-18.94 (1.52)	-4.49 (1.52)
Improvement offered by unified standard	**-122.9**	**141.0**	**126.8**	**166.1**

[1] The values in the right two columns allow driving behavior to improve as drivers switch to smaller vehicle classes. They are upper limits in the sense that all of the correlation between estimated driver behavior and size is attributed to the vehicle (e.g. large vehicles are driven more aggressively or are more difficult to control). As expected, all safety outcomes from CAFE improve in these columns. The sign of the effect on the current CAFE standard is preserved and the improvement offered by a unified standard is robust. Standard errors in parentheses are calculated as above.

Table 11: Alternative Identification Strategy and Simulation Elasticities[1]

	No driver effects	Full model (central)	Alternative identification	Alternative elasticities
Current CAFE within fleet	-135.02 (6.15)	149.47 (9.36)	222.00 (53.97)	156.15 (10.38)
Unified standard	-12.14 (3.81)	8.50 (4.35)	7.31 (21.11)	32.97 (2.85)
Footprint-based standard	26.88 (1.28)	6.27 (1.52)	-47.55 (5.72)	8.18 (1.27)

[1] The alternative identification strategy removes the need for crash test data. The standard errors are calculated as above and are much higher given the additional cross-equation restrictions. The final column includes results from an alternative source for substitution elasticities in the choice model.

Table 12: Alternative Demand Elasticities by Class[1]

	Compact	Midsize	Fullsize	Small Luxury	Large Luxury	Small SUV	Large SUV	Small Pickup	Large Pickup	Minivan
Compact	-3.12	0.94	0.06	0.10	0.00	0.10	0.01	0.12	0.03	0.03
Midsize	1.64	-3.92	1.10	0.15	0.06	0.39	0.07	0.06	0.02	0.19
Fullsize	0.65	4.28	-5.00	0.15	0.75	0.20	0.09	0.03	0.07	0.19
Small Luxury	1.32	0.94	0.32	-2.50	0.03	0.49	0.12	0.31	0.25	0.06
Large Luxury	0.11	0.90	1.06	0.05	-1.93	0.49	0.23	0.00	0.03	0.25
Small SUV	0.52	0.62	0.10	0.15	0.03	-4.05	0.96	0.31	0.44	0.38
Large SUV	0.24	0.45	0.14	0.09	0.05	3.73	-2.29	0.16	0.40	0.93
Small Pickup	0.39	0.22	0.00	0.05	0.00	0.49	0.08	-3.32	0.88	0.03
Large Pickup	0.15	0.16	0.02	0.05	0.00	0.30	0.16	0.81	-1.72	0.06
Minivan	0.19	0.38	0.06	0.00	0.03	0.30	0.46	0.03	0.06	-2.54

[1] Elasticities from Kleit (2004) aggregated to match the ten class definitions in my model. In order to isolate the effects of fleet composition I also proportionally adjust the cross-price elasticities such that fleet size is exactly maintained.

Table 13: Additional Robustness Checks

	No driver effects	Full model (central)	1998 and newer	Drivers under 55	Clear weather
Current CAFE within fleet	-135.02	149.47	142.15	132.82	148.52
Unified standard	-12.14	8.50	6.27	-2.47	8.26
Footprint-based standard	26.88	6.27	0.56	3.36	6.99
Fraction of accidents		1.00	0.52	0.77	0.90

[1] Changes in overall safety through time (perhaps most importantly the airbag requirement in 1998) do not affect the relative safety performance of classes enough to alter my conclusions on fuel economy rules. The potential frailty of older drivers and selection of vehicle type by weather conditions similarly have very small impacts on the results.

Figure 1: Estimates of α_i in Full Model[1]

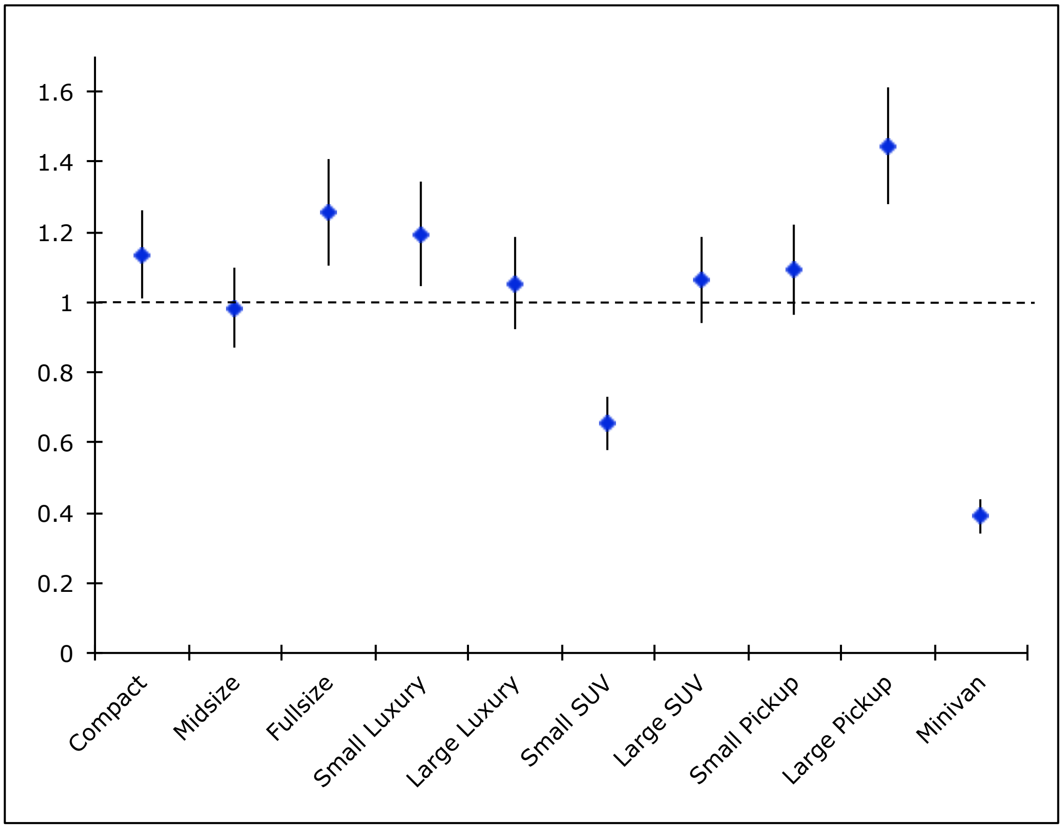

[1] Values are taken from the first row of Table 4 and bars indicate 95% confidence intervals. The average driving safety is normalized to 1 and larger values indicate more risk.

Figure 2: β_{ij} in the Restricted and Full Models[1]

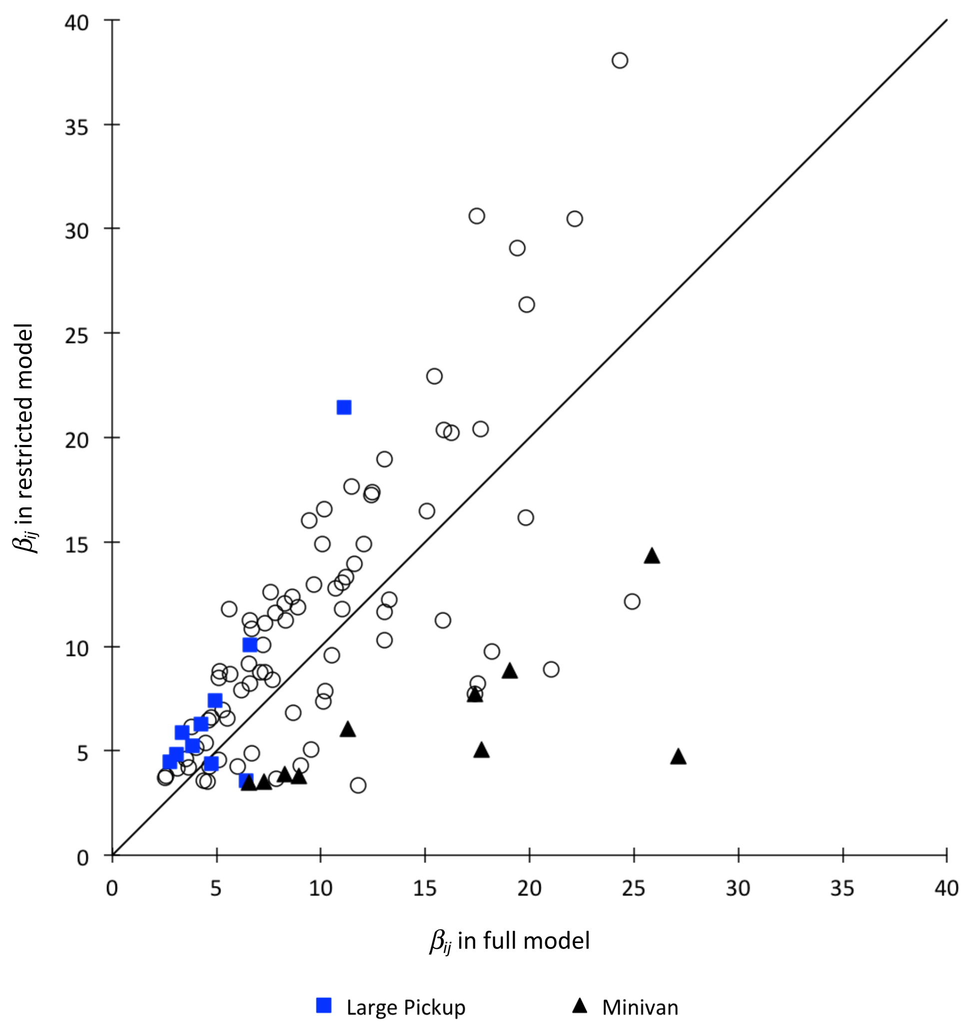

[1] The 100 β_{ij} parameters from Tables 3 and 4 are plotted relative to one another. The 45-degree line represents no change across specifications and markers for large pickups and minivans (for parameters in rows *i*) are shown to highlight the pattern of changes. The miles-driven weighted change on both sides of 45 degrees is equal, reflecting the fact that predicted risk in the two models matches the data overall.

References

Ahmad, Sanjana and David L. Greene. 2005. "Effect of Fuel Economy on Automobile Safety: A Reexamination." *Transportation Research Record* 1941: 1-7.

Anderson, Michael. 2008. "Safety for whom? The effects of light trucks on traffic fatalities." *Journal of Health Economics* 27(4): 973–989.

Anderson, Michael and Maximilian Auffhammer. 2011. "Pounds That Kill: The External Costs of Vehicle Weight." National Bureau of Economic Research Working Paper 17170.

Anderson, Soren and James Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." *American Economic Review* 101 (4): 1375-1409.

Austin, David and Terry Dinan. 2005. "Clearing the Air: The Costs and Consequences of Higher CAFE Standards and increased Gasoline Taxes." *Journal of Environmental Economics and Management* 50: 562-582.

Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes." *American Economic Review* 99 (3): 667:699.

Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. 2009. "Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing." National Bureau of Economic Research Working Paper 15590.

Cohen, Alma and Liran Einav. 2003. "The Effect of Mandatory Seat Belt Laws on Driving Behavior and Traffic Fatalities." *Review of Economics and Statistics* 85 (4): 828-843.

Cameron, A. Colin, and Pravin Trivedi. 1986. "Econometric Models Based on Count Data: Comparisons and Applications of Some Estimators and Tests." *Journal of Applied Econometrics* 1: 29–54.

Crandall, Robert W. and John D. Graham. 1989. "The Effect of Fuel Economy Standards on Automobile Safety." *Journal of Law and Economics* 32 (1): 97-118.

Environmental Protection Agency and Department of Transportation. 2010. *Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards; Final Rule*. Federal Register 75 (88), May 2010.

Evans, Leonard. 2001. "Causal Influence of Car Mass and Size on Driver Fatality Risk." *American Journal of Public Health* 91: 1076-81.

Gayer, Ted. 2004. "The Fatality Risks of Sport-Utility Vehicles, Vans, and Pickups Relative to Cars." *The Journal of Risk and Uncertainty* 28 (2): 103-133.

Goldberg, Pinelopi K. 1998. “The Effects of the Corporate Average Fuel Efficiency Standards in the U.S.” *Journal of Industrial Economics* 46: 1-33.

Greene, David L. 2009. “Feebates, Footprints and Highway Safety.” *Transportation Research Part D* 14: 375–384.

Herman, Irving P. 2007. *The Physics of the Human Body*. New York: Springer.

Jacobsen, Mark R. 2010. “Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity.” http://econ.ucsd.edu/~m3jacobs/Jacobsen_CAFE.pdf.

Kahane, Charles. 2003. *Vehicle Weight, Fatality Risk and Crash Compatibility of Model Year 1991-99 Passenger Cars and Light Trucks*. NHTSA Technical Report, DOT HS 809 662.

Khazzoom, J. Daniel. 1994. “Fuel Efficiency and Automobile Safety: Single-Vehicle Highway Fatalities for Passenger Cars.” *The Energy Journal* 15 (4): 49–102.

Kleit, Andrew N. 2004. “Impacts of Long-Range Increases in the Corporate Average Fuel Economy (CAFE) Standard.” *Economic Inquiry* 42: 279-94.

Klier, Thomas and Joshua Linn. 2008. “New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard.” Federal Reserve Bank of Chicago Working Paper 2008-13.

Knittel, Christopher. 2009. “The Implied Cost of Carbon Dioxide under the Cash for Clunkers Program.” Center for the Study of Energy Markets Working Paper 189, August.

Levitt, Steven D. and Jack Porter. 2001. “How Dangerous are Drinking Drivers.” *Journal of Political Economy* 109 (6): 1198-1237.

Li, Shanjun. 2012. “Traffic Safety and Vehicle Choice: Quantifying the Effects of the ‘Arms Race’ on American Roads.” *Journal of Applied Econometrics* 27: 34–62.

Loughran, David S. and Seth A. Seabury. 2007. “Estimating the Accident Risk of Older Drivers.” RAND Technical Reports TR-450-ICJ.

National Highway Traffic Safety Administration. 2008. *National Motor Vehicle Crash Causation Survey: Report to Congress*. DOT HS 811 059, July.

National Highway Traffic Safety Administration. 2009. *Traffic Safety Facts: A Brief Statistical Summary*. DOT HS 811172, June.

National Highway Traffic Safety Administration. 2010. “Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards; Final Rule.” *Federal Register* 75 (88): 24351-487.

National Research Council. 2002. *Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards*. Washington, DC: National Academy of Sciences.

Noland, Robert B. 2004. "Motor Vehicle Fuel Efficiency and Traffic Fatalities." *The Energy Journal* 25 (4): 1–22.

Parry, Ian W. H. and Kenneth A. Small. 2005. "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review* 95 (4): 1276–89.

Peltzman, Sam. 1975. "The Effects of Automobile Safety Regulation." *The Journal of Political Economy* 83 (4): 677-725.

Portney, Paul R., Ian W.H. Parry, Howard K. Gruenspecht and Winston Harrington. 2003. "The Economics of Fuel Economy Standards." *Journal of Economic Perspectives* 17: 203–217.

The White House Office of the Press Secretary. 2011. *President Obama Announces Historic 54.5 MPG Fuel Efficiency Standard*. Washington DC, July.

Wenzel, Thomas P. and Marc Ross. 2005. "The effects of vehicle model and driver behavior on risk." *Accident Analysis and Prevention* 37: 479-494.

White, Michelle J. 2004. "The 'Arms Race' on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety." *The Journal of Law and Economics* 47: 333-355.

Appendix A: Negative Binomial Specification

Begin by stacking the model in (5.3) and (5.4) to write the estimation as a single equation. Combining the count data and dividing out the HIC score, define the vector:

$$\mathbf{q} \equiv \begin{bmatrix} Y_{ist} / x_i \\ Z_{ijst} \end{bmatrix}$$

With 10 classes i and j, 18 bins s, and 156 weeks t, $\mathbf{q}$ contains the 308,880 rows of count data. Individual observations will be q_l. Similarly, write the parameters as a single vector (taking logs for convenience in the expressions below):

$$\gamma \equiv \begin{bmatrix} \ln(\delta_{is}) & \ln(\lambda_s) & \ln(\beta_{ij}) \end{bmatrix}$$

The right hand side of the model will contain only indicator variables, determining the set of observations to which individual δ, λ, and β parameters apply. For example, an observation of the count of multi-car accidents between classes 2 and 3 occurring in location bin 5 should receive indicators turning on $\delta_{2,5}$, $\delta_{3,5}$, and $\beta_{2,3}$. The corresponding vector of indicators for each observation l is:

$$\mathbf{m}_l \equiv \begin{bmatrix} \mathbf{d}_{is} & \mathbf{d}_s \cdot \mathbf{I}_{\text{single}} & \mathbf{d}_{ij} \cdot \mathbf{I}_{\text{multi}} \end{bmatrix}$$

where $\mathbf{d}_{is}$ has i times s elements, set to 1 for the vehicle(s) and bin involved and zero otherwise. Similarly, $\mathbf{d}_s$ is a vector containing indicators for bin and $\mathbf{d}_{ij}$ a vector containing indicators for all accident combinations. $\mathbf{I}_{\text{single}}$ is an indicator for single car accidents and $\mathbf{I}_{\text{multi}}$ an indicator for multi-car accidents. These allow λ_s to enter the first set of counts in $\mathbf{q}$ (corresponding to equation [5.3]) and β_{ij} to enter the second set (corresponding to [5.4]).

In the combined notation the Poisson version of the model becomes:

$$\begin{aligned} &E(q_l \mid \mathbf{m}_l) = \exp(\mathbf{m}_l{}'\gamma) \\ &\text{with } Var(q_l \mid \mathbf{m}_l) = \exp(\mathbf{m}_l{}'\gamma) \end{aligned} \qquad \text{(A.1)}$$

When applying the definition of the indicators in $\mathbf{m}_l$ and taking logs (A.1) is equivalent to (5.3) and (5.4) in the main text.

We now wish to generalize by adding an error term to the observed counts in $\mathbf{q}$:

$$E(q_l \mid \mathbf{m}_l, \varepsilon_l) = \exp(\mathbf{m}_l{}'\gamma + \varepsilon_l) \tag{A.2}$$

If we take $\exp(\varepsilon_l)$ to be distributed gamma with mean 1, variance θ, and independent of $\mathbf{m}$, then, following Cameron and Trivedi (1986), $\mathbf{q}_l$ will be distributed negative binomial with the following properties:

$$\begin{aligned} E(q_l \mid \mathbf{m}_l) &= \exp(\mathbf{m}_l{}'\gamma) \\ Var(q_l \mid \mathbf{m}_l) &= \exp(\mathbf{m}_l{}'\gamma)\cdot\left(1+\theta\exp(\mathbf{m}_l{}'\gamma)\right) \end{aligned} \tag{A.3}$$

The assumptions on ε provide an expected value of $\mathbf{q}_l$ that remains the same as above. The variance, however, is now allowed to exceed that of the Poisson model as the variance of the additional error in (A.2) grows away from 0. The model reduces back to Poisson as θ goes to 0.

The variation allowed in the ε_l error would include, for example, randomness in weather or driving patterns across time that is independent of the variables in $\mathbf{m}_l$. The requirement of independence with $\mathbf{m}$ is softened by noting that the δ_{is} and λ_s fixed effects already flexibly capture much of the unobserved variation we would expect at the bin-class level. Unobserved factors in the error that violate the cross-equation restrictions, for example a temporary traffic pattern that makes accidents between two particular classes more frequent but has no influence on other pairs or single car accidents, could bias the estimates.

The bias from violations on the error assumptions, though, may be quite limited considering that the influence of overdispersion more generally appears very small. Table A1 compares the estimates of α_i and the results of the main policy simulations using the estimates from the Poisson (left column) and the negative binomial (right column). The estimates are nearly indistinguishable and the standard errors increase only slightly, a function of the small magnitude of the estimated θ. Nevertheless, the estimate for θ is statistically significant and the more restrictive Poisson is rejected by a likelihood ratio test, so I have reported results from the negative binomial version throughout.

Appendix B: Alternative Definitions of the Bin Structure

As discussed in Section 5, the bins in the model relax the strictness of the cross-equation restrictions. Unobserved factors that vary across bin, but not classes within that bin, are allowed to enter the single- and multi- car equations differently. More flexible bins along this dimension should improve the model, though since the number of bins also rapidly increases the computation time required I must consider the influence of different possibilities separately.

Table A2 first reproduces the model without driver effects. Six possibilities for the bin structure in the full model follow, with the first four building up to the central case. The next row refines the central case bins even further, subdividing each into three road types (interstate, rural highway, and local roads) for a total of 54 bins. Finally, a version with bins by U.S. state is shown.

The results in the first column, applying to the usual CAFE standard, are largely robust. The second column, showing the unified standard, reveals somewhat larger differences: the no bins and state-level bins rows show modest improvements in safety. However, when we instead bin the data on factors more directly related to the relative frequencies of single- and multi-car accidents, these improvements become statistically indistinguishable from zero or even a slight deterioration of safety. Time-of-day and urban density will have a natural influence on the divide between single- and multi- car accidents since they change the density of cars on the road. Income has a substantial effect as well, perhaps proxying for commuting patterns or local road quality, making these the three I choose for the central case. The two key qualitative findings, understatement of fatalities in the naive model and the dramatic improvements available under a unified standard, remain robust to all of these bin structures.

Appendix C: Accident Fault Data

Data on accident fault can provide additional insight on the composition of factors within the α_i parameters, though the analysis here remains only suggestive given the degree of noise and potential biases in fault assignment. The table below suggests that behaviors for which drivers are faulted, like traffic law violations or distraction and inattention, are a relatively small component of α. Factors related to geography and the timing of trips may therefore explain a greater portion of the risk, helping to further reduce concern of Peltzman-type effects influencing the simulation.

I consider fatal two-car accidents where the vehicles involved are from unmatched classes, taking a measure of fault from the FARS data. I assign fault to a vehicle if the driver is either charged with a traffic violation in conjunction with the accident or if the FARS notes a "driver contributing factor" (for example, sleep, drug use, or distractions). I remove accidents where fault appears on both sides and also exclude counts where both vehicles are of the same class, since these counts only reflect the safety of the class itself. Table A3 counts the number of times fault was with the listed class and the number of times fault was with the opposing class. The ratio, removing the overall tendency of individual classes to appear in fatal accidents, provides a measure of differences in fault: unity indicates that the listed class is exactly as likely to be faulted as any opposing vehicle.

A key result from this exercise is that fault appears more evenly distributed than α_i: this suggests that location, time, and other factors not associated with fault are also important determinants. Some classes, like luxury cars in Table A3, appear to receive a lot of blame in the accidents they end up in but in fact appear in relatively few fatal accidents overall (as seen in α). In these cases, effects other than fault are then the dominant determinants of α. To the extent these other factors (for example geographical location) are also more likely to remain fixed when a driver switches cars this adds a piece of suggestive evidence in support of the central case simulation assumptions.

Appendix D: Alternative Demand Elasticities

The demand elasticities I use in the central case simulations come originally from a Bayesian model of vehicle choice, making it possible to draw a variety of alternatives from the estimated posterior densities to further explore the sensitivity of my results. Table A4 displays the minimum and maximum change in fatalities for each policy simulation, taken over 50 different draws on the elasticities. The extremes remain quite close to the central case, likely reflecting the large dataset and relatively high precision in the source paper.

However, structural uncertainty in the elasticity estimates likely remains an important issue, attested to by the wide variety of estimates produced in the literature. This highlights the importance of the sensitivity analysis across different sources for the elasticity parameters, shown in Tables 5 and 12 of the main text.

Appendix E: Vehicle Class Aggregates

Table A5 below provides details on the mean and standard deviation of the fuel economy and weight of each class. Fuel economies across classes range from 19 to 30 MPG, capturing much of the variation among popular vehicles; the aggregation, however, means I do not capture the extremes as well. Hybrid compacts on the high end, for example, or luxury SUV's on the low end will be missed, though demand for these vehicles is also relatively inelastic meaning they play less of a role in the compositional changes expected under regulation.

The largest variation within class for fuel economy comes in compacts, where a growing fraction of hybrids and ultra-compacts enter alongside more typical compacts like the Ford Focus or Honda Civic. The largest variation in weight is in the large SUV class, likely coming from the presence of so-called "premium" large SUV's that feature weights near the maximum permissible without a commercial driving license.

Table A1: Comparison of Poisson and Negative Binomial Estimation Results

	Poisson	Negative binomial
Estimates of α_i		
Compact	1.137 (0.0641)	1.136 (0.0644)
Midsize	0.983 (0.0572)	0.983 (0.0576)
Fullsize	1.254 (0.0769)	1.254 (0.0774)
Small Luxury	1.193 (0.0749)	1.193 (0.0754)
Large Luxury	1.054 (0.0669)	1.053 (0.0673)
Small SUV	0.653 (0.0384)	0.654 (0.0387)
Large SUV	1.062 (0.0626)	1.063 (0.0631)
Small Pickup	1.094 (0.0652)	1.094 (0.0657)
Large Pickup	1.445 (0.0838)	1.446 (0.0844)
Minivan	0.389 (0.0245)	0.389 (0.0247)
Central policy results		
Current CAFE within fleet	149.06 (9.27)	149.47 (9.36)
Unified standard	8.19 (4.29)	8.50 (4.35)
Footprint-based standard	6.22 (1.50)	6.27 (1.52)

Table A2: Effects of Alternative Bin Structures

	Current CAFE within fleet	Unified standard	Footprint-based standard
No driver effects	-135.02	-12.14	26.88
Full model			
Bins:			
None	142.17	-25.04	3.77
Time-of-day	136.79	-14.58	4.13
Time-of-day, income	143.22	-4.94	4.99
Time-of-day, income, urban *(central case)*	149.47	8.5	6.27
Time-of-day, income, urban, road type	136.65	10.79	10.43
Fifty states	125.34	-29.86	2.67

Table A3: Fault by Class[1]

All accidents involving	Own fault	Others fault	Ratio
Compact	4262	3404	1.25
Midsize	3748	3039	1.23
Fullsize	1208	1218	0.99
Small Luxury	660	453	1.46
Large Luxury	702	602	1.17
Small SUV	1673	1959	0.85
Large SUV	1540	2091	0.74
Small Pickup	1187	1218	0.97
Large Pickup	2654	3344	0.79
Minivan	817	1123	0.73

[1] Fault is assigned here if the FARS data either indicate a moving violation charged or include a driver contributing factor.

Table A4: Simulation Results over a Sample of Varying Demand Elasticities[1]

	No driver effects			Full model		
	Min	Central case	Max	Min	Central case	Max
Current CAFE within fleet	-147.2	-135.0	-127.1	144.2	149.5	156.6
Unified standard	-21.0	-12.1	-1.9	5.6	8.5	11.8
Footprint-based standard	18.3	26.9	31.8	5.3	6.3	8.1

[1] The minimum and maximum simulation outcomes are shown over 50 draws from the posterior density of the parameters controlling demand elasticity.

Table A5: Fuel Economy and Weight by Class

	Fuel Economy (MPG)		Weight (pounds)	
Class	Mean	Standard deviation	Mean	Standard deviation
Compact	30.2	3.50	2680	415.5
Midsize	27.0	2.39	3150	312.6
Fullsize	25.4	2.05	3598	345.3
Small Luxury	26.0	2.95	3332	472.4
Large Luxury	23.8	1.42	3801	285.9
Small SUV	24.1	3.28	3506	465.3
Large SUV	19.0	2.53	4652	489.9
Small Pickup	22.5	2.85	3236	325.2
Large Pickup	19.1	2.49	4718	435.6
Minivan	23.4	1.46	3688	301.8