

Pounds That Kill: The External Costs of Vehicle Weight

MICHAEL L. ANDERSON

University of California, Berkeley and NBER

and

MAXIMILIAN AUFFHAMMER

University of California, Berkeley and NBER

First version received January 2012; final version accepted May 2013 (Eds.)

Heavier vehicles are safer for their own occupants but more hazardous for other vehicles. Simple theory thus suggests that an unregulated vehicle fleet is inefficiently heavy. Using three separate identification strategies we show that, controlling for own-vehicle weight, being hit by a vehicle that is 1000 pounds heavier generates a 40–50% increase in fatality risk. These results imply a total accident-related externality that exceeds the estimated social cost of US carbon emissions and is equivalent to a gas tax of \$0.97 per gallon (\$136 billion annually). We consider two policies for internalizing this external cost, a weight-varying mileage tax and a gas tax, and find that they are similar for most vehicles. The findings suggest that European gas taxes may be much closer to optimal levels than the US gas tax.

Key words: Traffic accidents, Highway safety, Gasoline tax, Energy Policy, Arms race.

JEL Codes: H23, I18, Q48, Q58, R41

1. INTRODUCTION

The average weight of light vehicles sold in the United States has fluctuated substantially over the past 35 years. From 1975 to 1980, average weight dropped almost 1000 pounds (from 4060 pounds to 3228 pounds), likely in response to rising gasoline prices and the passage of the Corporate Average Fuel Efficiency (CAFE) standard. As gasoline prices fell in the late-1980s, however, average vehicle weight began to rise, and by 2005 it had attained 1975 levels (US Environmental Protection Agency, 2009). A rich body of research examines the effects of CAFE and gasoline prices on consumers' vehicle choices (Portney *et al.*, 2003; Austin and Dinan, 2005; Bento *et al.*, 2009; Li *et al.*, 2009; Klier and Linn, 2012; Busse *et al.*, 2013).

One area of intense research interest is how the choices consumers make in response to gasoline prices and fuel economy standards affect traffic fatalities. Traffic accidents are the leading cause of death for persons under the age of 40, and they are a major source of life-years lost.¹ Intuitively, heavier cars are safer than lighter cars, and previous research has argued

1. Lung cancer, a disease that is generally the result of smoking, kills approximately four times as many Americans each year as traffic accidents. However, the average lung cancer decedent is 71 years old while the average traffic accident decedent is only 39 years old. The number of life-years lost to traffic accidents is thus similar in magnitude to the number of life-years lost to lung cancer.

that a heavier vehicle fleet is a safer vehicle fleet (Crandall and Graham, 1989). Much of the subsequent transportation safety literature has focused on the effects of average vehicle weight on safety, reaching varying conclusions.

From an economic standpoint, however, an unregulated vehicle fleet must be inefficiently heavy. A heavier vehicle is safer for its own occupants but more hazardous for the occupants of other vehicles. The safety benefits of vehicle weight are therefore internal, while the safety costs of vehicle weight are external. Consumers' vehicle choices thus have the important features of an "arms race". To date no comprehensive attempt has been made to quantify the *external* safety costs of vehicle weight. This figure is essential for determining the socially optimal weight of the vehicle fleet, and it cannot be inferred from the net effects on traffic safety of average vehicle weight or fuel economy regulations.

We quantify the external costs of vehicle weight using a census of police-reported crashes across 8 heterogeneous states. Unlike data sets employed in the existing transportation literature or Jacobsen (2013), our data include both fatal and non-fatal accidents. Using unique vehicle identifiers (VINs), we determine the curb weight of each vehicle in an accident, thus minimizing concerns about attenuation bias due to measurement error. The rich set of vehicle, person, and accident observables in the data set allow us to minimize concerns about omitted variables bias. Using these data, we estimate the external effects of vehicle weight on fatalities and serious injuries conditional on a collision occurring.

Two key results emerge from our estimates. First, we show that vehicle weight is a critical determinant of fatalities in other vehicles in the event of a multivehicle collision; our preferred estimate implies that a 1000 pound increase in striking vehicle weight raises the probability of a fatality in the struck vehicle by 47%. When we translate this higher probability of a fatality into external costs (relative to a small baseline vehicle), the total external costs of vehicle weight from fatalities alone are estimated at \$136 billion per year. Second, by separately controlling for vehicle weight and whether the striking vehicle is a light truck (*i.e.* a pickup truck, sport utility vehicle, or minivan), we show that light trucks significantly raise the probability of a fatality in the struck car—in addition to the effect of their already higher vehicle weight.

Our unique data set allows us to condition on a collision occurring and thus ensures that our results cannot be generated by differences in collision rates between drivers of lighter and heavier vehicles. Nevertheless, driver selection could bias our results if drivers of heavy vehicles have a tendency towards severe accidents. We rule out this possibility through falsification tests and two alternative sources of identification. First, we show that our estimates persist even when controlling for specific vehicle type via make and model fixed effects. Second, we estimate the effect of striking vehicle weight using variation in the number of occupants in the striking vehicle and find estimates that are close to our main estimates. Finally, we show that vehicle weight does not predict fatalities when two vehicles of equal weight collide. This suggests that drivers of heavy vehicles are not predisposed towards severe accidents. All three tests suggest that we successfully identify the causal effect of vehicle weight on the probability of fatalities in two-car collisions.

One way to internalize the identified externality is through a weight varying mileage tax. However, such a tax could be logistically difficult to implement. We apply our estimates to consider whether a simple gasoline tax could be an alternative to internalize most of the external costs and conclude that it could. Our calculations suggest that the level of the optimal gasoline tax is much higher than previously estimated (*e.g.* Parry and Small, 2005) and that the external traffic fatality costs of vehicle weight eclipse the sum of all other vehicle-related externalities (Portney *et al.*, 2003). Total accident-related external costs exceed central estimates of the annual social cost of US carbon emissions (US EIA, 2009; Greenstone

et al., 2011). We estimate that a weight or gasoline tax would not substantially increase total traffic fatalities.

This article is organized as follows. Section 2 conducts a literature review. Section 3 presents the analytic and empirical framework. Section 4 details the data, and Section 5 presents the main results. Section 6 presents two alternative sources of identification and falsification tests to confirm that selection bias does not drive our results. Section 7 links the results to energy policy implications, focusing on the gasoline tax, and estimates the potential effects of such a tax on total fatalities. Section 8 concludes.

2. LITERATURE REVIEW

A large traffic safety literature examines the relationship between average vehicle weight and traffic fatality rates. Most of this literature estimates aggregate time series correlations (Robertson, 1991; Khazzoom, 1994; Noland, 2004; Ahmad and Greene, 2005). Two exceptions are Kahane (2003) and Van Auken and Zellner (2005), which use micro data on fatal accidents only.² These studies come to varying conclusions regarding the sign of the relationship between average vehicle weight and overall fatality rates, but all conclude that the magnitude of this relationship is relatively modest.

In recent years economists have studied the “arms race” nature of vehicle choice. This work focuses on the internal and external risks posed by the largest vehicles—pickup trucks and sport utility vehicles (SUVs)—relative to the typical passenger car. White (2004), Gayer (2004), Anderson (2008), and Li (2012) all conclude that light trucks (pickups and SUVs) impose significant risks relative to passenger cars.

This study builds upon the existing literature by considering the fundamental role that vehicle weight plays in determining external risk. We recognize that *any* vehicle that is heavier than the smallest feasible vehicle poses some external risk to other roadway users. We quantify that risk and find total external safety costs of vehicle weight that are at least 11.6 times larger than Li’s estimates and 7.4 times larger than Anderson’s estimates.³ Our comprehensive results span the entire range of the vehicle fleet and allow us to consider the broader implications of vehicle weight for energy policy. We also develop a theoretical model that captures the central role that weight disparity plays in determining traffic fatalities. The model, in combination with the empirical results, establishes several results with significant policy implications. First, we show that the fatality externality does not result from a failure to coordinate on a single vehicle weight. Second, we show that “downsizing” the vehicle fleet does not significantly increase traffic fatalities unless it dramatically increases fleet heterogeneity. Finally, we show that decreasing the weight of lighter vehicles can be welfare enhancing even if traffic fatalities increase. Our estimates of the function relating weight disparity and traffic fatalities also illuminate the mechanisms underlying existing findings in the traffic safety literature.

Concurrent work by Jacobsen (2013) is particularly relevant to this article. Jacobsen uses data on fatal accidents to explore the traffic safety implications of different fuel economy regulatory

2. They supplement the fatal accident data with data on police-reported accidents from several states to estimate the rate at which different types of vehicles enter into collisions. However, unlike this study, they do not use police-reported accident micro data to estimate their econometric specifications.

3. Li (2012) and Anderson (2008) calculate lifetime accident externalities of \$2,444 and \$3,850, respectively for each light truck sold. In the steady state, annual truck sales must average 5.8 million to maintain the current fleet of 108 million light trucks (California Air Resources Board, 2004; US Department of Transportation, 2011). The annual safety externality for light trucks is thus between \$14.2 billion (Li) and \$22.3 billion (Anderson). We find an annual safety externality of \$135.8 billion from vehicle weight alone, and an additional \$28.7 billion from light truck bodies.

schemes across 10 broad vehicle classes. While both Jacobsen and this article explore issues related to energy policy and traffic fatalities, the two papers address different questions. We estimate a parameter, the external cost of vehicle weight, which is crucial for determining the appropriate level of vehicle or gasoline taxation. Jacobsen estimates the net effect of changes to an existing policy, CAFE, on net traffic fatalities. Loosely speaking, Jacobsen concludes that scheduled changes to CAFE may not increase traffic fatalities, while we demonstrate that increased energy taxes or stricter CAFE standards may internalize an important externality. Empirically, the two papers use different data and identification strategies. Jacobsen uses fatal accident data and adjusts for differences in driver collision rates by controlling for fatality rates in single-vehicle accidents. We use data on all police reported accidents, which automatically adjusts for differences in driver collision rates.

3. ANALYTIC AND EMPIRICAL FRAMEWORK

The wider impacts of consumers' vehicle choices represent a classic example of an externality. Purchasing a heavier vehicle enhances safety for the individual, but it also increases the risk to other drivers in the case of an accident. The net benefit of vehicle weight on traffic fatalities is thus smaller than the private benefit of vehicle weight on traffic fatalities, and consumers are incentivized to purchase heavier vehicles than is socially optimal. The following stylized model outlines how the socially optimal vehicle weight differs from the privately chosen weight.

Consider a population of N consumers. Consumer i spends income y_i on x_i units of the composite good x , which is normalized to have price 1, and on a vehicle weighing w_i pounds, which costs p dollars per pound. The price per pound, p , is positive because heavier vehicles cost more to build and fuel. For simplicity we assume an additive utility function:

$$U_i(x_i, w_1, \dots, w_N) = x_i + E_{j \neq i} [f(w_i - w_j)] + g_i(w_i) \quad (1)$$

Consumer i gets utility from the composite good x , the safety benefits of driving a heavier vehicle than other consumers, captured by the function f , and the larger capacity of a heavier vehicle, captured by the function g_i . One way to view f is as $-\alpha \cdot h(w_i - w_j)$, where α is positive and $h(w_i - w_j)$ is the probability that driver i dies if a collision occurs between vehicles i and j . In this formulation, increases in the overall fatality rate due to factors other than the relative weight distribution (e.g. an increase in the number of drunk drivers) will increase α and thus increase f 's impact in the utility function. The function g_i also captures private benefits from other features related to weight, including engine size and even safety features like reinforced bars. We assume that $f', g' > 0$, $f'' \leq 0$, and $g'' < 0$. We allow heterogeneity in vehicle preferences by indexing g by i . For tractability we assume that safety is a function of relative weight; *i.e.*, extra weight does not independently make vehicle i safer except insofar as it increases its weight relative to other vehicles. We relax this assumption in our regressions, but there are two reasons to believe it holds empirically. First, adding vehicle i 's weight to a regression of fatalities on the difference in vehicle weights does not increase the regression's explanatory power (see Section 5). Second, vehicle weight is uncorrelated with performance in controlled crash tests (see Section 6.3). We assume consumer i takes the weight of other vehicles as given and that the safety benefits and costs do not vary with the size of the vehicle fleet.⁴ We substitute the consumer's budget constraint

4. In practical terms this implies that the average probability that consumer i experiences a collision does not increase with the size of the vehicle fleet. The effect of congestion on accident rates is the focus of Edlin and Karaca-Mandic (2006), and it is a different externality than the one we estimate in this article.

$y_i = x_i + pw_i$ into the utility function and take the derivative with respect to w_i . The first-order condition (FOC) is:

$$g'_i(w_i^*) + E_{j \neq i}[f'(w_i^* - w_j^*)] = p \quad (2)$$

Consumer i chooses a vehicle weighing w_i^* pounds, where at w_i^* the expected marginal safety benefits, $E_{j \neq i}[f'(w_i^* - w_j^*)]$, plus the marginal benefits of larger capacity, $g'_i(w_i^*)$, equal p .⁵ The marginal safety benefits of vehicle weight are averaged over the entire fleet of vehicles that the consumer may collide with. The consumer's choice may thus depend on the distribution of vehicle weight in the existing fleet.

The social planner maximizes the sum of the individual indirect utility functions:

$$\max_{w_1, \dots, w_N} \sum_{i=1}^N [y_i - pw_i + E_{j \neq i}[f(w_i - w_j)] + g_i(w_i)] \quad (3)$$

For simplicity we assume the social planner has a sufficiently long time horizon to view every vehicle as eligible for replacement.⁶ Taking the derivative with respect to w_i for each consumer and solving all of the first-order conditions gives an optimal weight of w_i^{**} for consumer i , which is the weight at which the sum of the marginal benefits from larger capacity and the net safety benefits (marginal private safety benefits minus the marginal safety costs to other vehicles) equals p , or:

$$g'_i(w_i^{**}) + E_{j \neq i}[f'(w_i^{**} - w_j^{**}) - f'(w_j^{**} - w_i^{**})] = p \quad (4)$$

Since $f' > 0$ and g is strictly concave, the social planner chooses weight $w_i^{**} < w_i^*$, $g'_i(w_i^{**})$ must be greater than $g'_i(w_i^*)$ to satisfy both equations (2) and (4). Complicating the comparison between the two FOCs is the fact that $w_j^{**} < w_j^*$, but since f is concave this, if anything, attenuates $f'(w_i^{**} - w_j^{**})$ relative to $f'(w_i^* - w_j^*)$. The social planner chooses a lower weight than consumer i 's choice of weight in the private market equilibrium for two reasons. First, the planner accounts for the negative effect that vehicle i 's weight has on other vehicles' safety, captured by $-f'(w_j^{**} - w_i^{**})$. Second, if f is strictly concave, the marginal safety benefit of vehicle i 's weight, $f'(w_i^{**} - w_j^{**})$, decreases as the weight of other vehicles in the fleet, w_j^{**} , decreases. This is the "arms race" dynamic at work.

In two special cases the social planner's FOC reduces to $g'_i(w_i^{**}) \approx p$, and she chooses weight based only on cargo and passenger capacity benefits. One case is if f is approximately linear over the support of $w_i^{**} - w_j^{**}$. Another case is if g_i does not vary across consumers, so that $w_i^{**} = w_j^{**}$ for all consumers. In either case the marginal internal safety benefit and external safety cost of weight are identical for all consumers, and adding weight for safety becomes a zero-sum game. Beyond establishing the externality, the model yields five results with significant policy implications:

Result 1. *The externality is not caused by consumers failing to coordinate on the same weight. Suppose there is no heterogeneity in g_i , so that all consumers purchase vehicles of identical weight w^* in the private market equilibrium. In this equilibrium consumers choose w^* such*

5. For simplicity this stylized model ignores heterogeneity in safety preferences among consumers. We could account for this heterogeneity by indexing $f(\cdot)$ by i . However, in practice we find that heterogeneity in the value of a statistical life (which implies heterogeneity in f) has minimal effect on the distribution of corrective taxes.

6. One potential benefit we ignore is the possibility that the social planner may factor in domestic automaker profits, which in the United States have historically been increasing in light truck share.

that $g'(w^*) + f'(0) = p$. However, the socially optimal weight is w^{**} such that $g'(w^{**}) = p$. Thus w^* remains higher than the optimal weight w^{**} since $f'(0) > 0$; identical weights alone are not sufficient to guarantee a socially optimal outcome.

Result 2. Coordinated reductions in fleet weight need not increase total fatalities. If all consumers reduce weight by some constant w_c , then $E_{j \neq i}[f(w_i - w_j)]$ remains unchanged for everyone. This result hinges on the assumption that, at the margin, extra weight does not independently increase safety (this appears to hold empirically). However, heterogeneous reductions in fleet weight may affect total fatalities.

Result 3. Fleet heterogeneity can affect total fatalities iff f is strictly concave. Suppose consumers move from a world of equal weight vehicles to one in which they choose heavy vehicles weighing w_h or light vehicles weighing w_l , with $w_h > w_l$. In any collision between two vehicles of different weights, the heavy vehicle's safety gain over colliding with an equal weight vehicle, $f(w_h - w_l) - f(0)$, is offset by the light vehicle's safety loss over colliding with an equal weight vehicle, $f(0) - f(w_l - w_h)$. If f is strictly concave, then $f(0) - f(w_l - w_h) > f(w_h - w_l) - f(0)$, and the safety loss to the light vehicle exceeds the safety gain to the heavy vehicle. Fleet heterogeneity can thus increase total fatalities.

Result 4. Even if f is strictly concave, reducing the weight of light vehicles can be welfare enhancing. Consider a consumer with low demand for cargo capacity that privately chooses a vehicle of weight $w_i^* < \bar{w}$. If the social planner forces consumer i to further reduce w_i , the net change in social welfare is $p - g'_i(w_i^*) + E_{j \neq i}[f'(w_j^* - w_i^*) - f'(w_i^* - w_j^*)]$. Since $p - g'_i(w_i^*) - E_{j \neq i}[f'(w_i^* - w_j^*)] = 0$ by equation (2), the net change in social welfare is positive. Due to the externality, reducing the weight of a light vehicle can increase social welfare even if total fatalities rise.

Result 5. Whether the externality is a function of the fleet's weight distribution depends on the concavity of f . The marginal externality of consumer i 's vehicle weight is $E_{j \neq i}[-f'(w_j - w_i)]$. If f is approximately linear over the support of the weight distribution, then the marginal externality reduces to β , where $f(u) \approx \beta u$. If f is strictly concave, then $E_{j \neq i}[-f'(w_j - w_i)]$ may change as the distribution of fleet weight changes. This implies that the marginal externality under the current fleet weight distribution may differ from the marginal externality under the optimal fleet weight distribution.

Empirically, we estimate the effects of vehicle weight on fatalities in two-vehicle collisions. Our regressions estimate $h(w_i - w_j) = -\alpha \cdot f(w_j - w_i)$. We switch the order of w_i and w_j because we focus on the external costs of weight, and we reverse the sign on f because our outcome—a fatality—represents negative utility. Concavity in f thus corresponds to convexity in our regression function h . In our main specification we relax the assumption that safety is only a function of relative weight and allow a more general functional form of $h_0(w_i) - h_1(w_j)$. However, we cannot reject the hypothesis that only relative weight matters. To test for convexity in h (concavity in f) we also estimate $h(w_i - w_j)$ flexibly over the support of weight differences (see Section 5.2). We apply our estimates to calculate the total external safety costs accruing from the weight of the current vehicle fleet and to forecast the effects on total fatalities of coordinated and uncoordinated reductions in fleet weight.

Note that the primary costs of this externality do not accrue in the form of traffic fatalities, which on net may change little with a reduction in fleet weight (Results 2 and 3). Rather they accrue in the form of purchases of larger vehicles that are more expensive to operate. In this

sense it is similar to an arms race, which may not increase the probability of conflict even as both countries spend more on new weapons. In the notation of our model, the welfare loss arises from choosing a weight w that does not maximize $g_i(w) - pw$.

In principle, liability rules and insurance regulations could internalize many of the external costs due to vehicle weight. If drivers of heavy vehicles know that they will be held liable for deaths in other vehicles, then they should take these risks into account when purchasing their own vehicles. If insurance companies understand that heavier vehicles pose more danger to other roadway users, then they should charge higher liability premiums to drivers of heavy vehicles. In practice, however, liability rules and insurance regulations fail to internalize the fatality risks generated by heavy vehicles.

Tort liability rules are inadequate to internalize fatality risks for two reasons. First, liability only applies in cases in which a driver behaves negligently (White, 2004). This implies that the driver of any given vehicle may not be liable in the event of a multivehicle accident. Second, even if found liable, few drivers possess assets that are sufficient to cover the cost of a fatality. The value of a statistical life (VSL) used by the United States Department of Transportation in cost-benefit analyses is \$5.8 million (2008 dollars), but only 7% of families in the United States had a net worth exceeding \$1 million in 2001 (Kennickell, 2003).

Though few drivers can cover the cost of a fatality, liability insurance regulations could force most drivers to pay the expected liability costs of operating their vehicles. Again, however, the mandated levels of liability insurance are inadequate to cover the costs of a fatality. Two states (Florida and New Hampshire) do not require drivers to carry any liability coverage at all for injuries, and 44 states require drivers to carry \$25,000 or less in liability coverage for each person injured (Insurance Information Institute, 2010). Many drivers remain uninsured despite the regulations, and few drivers have policies that exceed several hundred thousand dollars of coverage.

While liability rules and insurance regulations cannot internalize the majority of fatality costs, they may internalize a significant fraction of incapacitating injury costs. Estimates of the value of an incapacitating injury are far lower than the value of a statistical life, and it is plausible that insurance policies carried by many drivers could cover the costs of an incapacitating injury.⁷ For this reason, our policy analysis focuses on external fatality costs (*i.e.* costs from fatalities that occur outside of the driver's own vehicle) and ignores external incapacitating injury costs. Accounting for injury costs increases the magnitude of our results, though we cannot accurately estimate what fraction of injury costs are already internalized. We calculate an upper bound on external injury costs in Section 7.

To measure the effect of vehicle weight on external fatalities under ideal conditions, we would randomly assign vehicles of differing weights to drivers and observe external fatality rates by vehicle type. Such an experiment is infeasible in practice, and even an analogous study using observational data is impractical due to substantial measurement error in vehicle stocks and model-level vehicle miles travelled in most states. Instead, we focus on the risk of a fatality conditional on a collision occurring. A key assumption when we interpret our estimates in a policy context is that vehicle weight has no causal effect on the probability of a collision. We discuss this assumption below and conclude that, if it is violated, then the effect of vehicle weight on the probability of a collision is likely positive. Our estimates thus represent a lower bound on the effect of weight on external fatalities. In Section 7 we explore how much our estimates might increase if we relax the assumption.

7. The National Safety Council (2010) estimates the comprehensive cost of an incapacitating injury at \$214,000 (2008 dollars). In comparison, the council estimates the comprehensive cost of a fatality at \$4.2 million.

Consider the expected external fatalities for a vehicle of type i during time interval t . For simplicity, assume that t is short enough that the probability of multiple collisions during t is effectively zero.

$$E[\text{fatalities}_{it}] = E[E[\text{fatalities}_{it} | \text{collision}_{it}]] = E[\text{fatalities}_{it} | \text{collision}_{it} = 1] \cdot P(\text{collision}_{it} = 1) \quad (5)$$

Equation (5) must hold via the law of iterated expectations. If weight has no causal effect on the probability of a collision, then the total effect of weight on external fatalities is proportional to the effect of weight on external fatalities conditional on a collision occurring. Weight may affect the probability of a collision in two ways, however. First, from an engineering perspective, heavier vehicles are less manoeuvrable and have longer braking distances.⁸ Even if driver behaviour is unchanged, heavier vehicles may therefore get into more accidents. Second, heavier vehicles may also affect driver behaviour. On the margin, drivers may respond to the internal safety benefits of heavy vehicles by increasing their optimal collision rate (Peltzman, 1975). Both the physical characteristics of heavier vehicles and the potential driver response to heavier vehicles could therefore generate a *positive* effect of vehicle weight on collision rates.⁹

Empirical evidence also suggests that, if anything, heavier vehicles have higher collision rates than lighter vehicles. Evans (1984) examines the relationship between accident rates and vehicle weight using accident data and vehicle registration data from North Carolina, New York, and Michigan. He finds that, after conditioning on driver age, 4000 pound vehicles have accident rates that are 39% higher than 2000 pound vehicles. More recently, White (2004) and Anderson (2008) estimate that light trucks are 13% to 45% more likely to experience multivehicle collisions than passenger cars. Of course, some of the observed differences in crash rates may be due to driver selection; careless drivers may choose heavier vehicles. Nevertheless, both theory and empirical evidence suggest that weight may directly increase the probability of experiencing a collision. We thus interpret our estimates—which are conditional on a collision occurring—as lower bounds on the causal effect of weight on external fatalities.¹⁰

4. DATA

The data set consists of the population of police-reported accidents for eight states: Florida, Kansas, Kentucky, Maryland, Missouri, Ohio, Washington, and Wyoming. These data come from the State Data System, maintained by the National Highway Traffic Safety Administration (NHTSA). We obtained permission from each state's police force to use the data. The SDS data include information on injuries and fatalities, geographic location, weather conditions, use of

8. We confirm this relationship using braking and manoeuvrability data from *Consumer Reports*. Analysing data from 58 *Consumer Reports*-tested vehicles of varying weights, we find that an extra 1,000 pounds of curb weight is associated with 4.5% worse performance in an emergency handling test ($t=8.9$), 2.2% worse performance in a dry braking distance test ($t=3.8$), and 3.8% worse performance in a wet braking distance test ($t=5.0$).

9. To the best of our knowledge, the only factor that might reduce the probability of a collision for heavier vehicles is visibility. Larger vehicles provide their drivers with a better view of the road ahead, which may decrease the probability of an accident. However, they also make it more difficult for drivers behind them to see ahead, which may increase the probability of an accident. The net impact of these two effects is unclear, but the resulting dynamic is again an example of an arms race; the visibility benefits are internal while the visibility costs are external. Visibility would thus be another reason to tax larger vehicles more than smaller vehicles.

10. Note that the concern here is whether weight has a *causal* effect on collision probabilities. This concern arises because we consider the policy implications of inducing some drivers to switch to lighter vehicles via a tax. This exogenous manipulation of vehicle choice will affect collision probabilities only if vehicle weight has a causal effect on collision probabilities. Weight may also be correlated with the type of driver, which could generate selection bias in our regressions. We consider this issue separately in Section 6.

safety equipment, and driver and occupant characteristics. We selected these eight states out of the 32 states currently participating in the SDS as they report the vehicle identification number (VIN) for the majority of vehicles in the data set. We purchased data tables from DataOne Software to match the first 9 digits of the VIN to curb weight data for each vehicle (a vehicle's curb weight is its weight with standard equipment and a full tank of fuel, but not loaded with any passengers or cargo). We therefore observe curb vehicle weight for approximately 64% of the vehicles in our data set (we confirm in Section 5 that the missing weight data do not appear to bias our estimates). For analytic purposes, we decompose the data set into three sub-samples, two-vehicle crashes, three-vehicle crashes, and single-vehicle crashes. The two-vehicle crash data set is the focus of most of our analyses. It contains 4.8 million vehicles in collisions in which both vehicles have complete curb weight data.¹¹

One important feature of the SDS data is that accidents only appear in the data set if the police take an accident report. According to NHTSA documentation, various estimates suggest that only half of all motor vehicle accidents are police reported. While many of the unreported accidents are single vehicle accidents, some no doubt involve two vehicles as well. This sampling frame could affect our estimates if vehicle weight affects the probability of a police report, all other factors held constant. Serious multivehicle accidents are always reported to the police regardless of vehicle weight, but vehicle weight could affect the probability that a minor accident is reported to the police. Unlike the probability of a collision, there is no *a priori* reason to believe that vehicle weight must have a positive effect on the probability of a police report. On the one hand, collisions involving heavier vehicles cause more property damage, all other factors held constant, because more kinetic energy must be dissipated through deformation of materials. On the other hand, some heavier vehicles, such as pickup trucks, are more likely to be involved in rugged work. These trucks may have accumulated more dents, reducing the likelihood that the owners will report property damage from a minor accident.

If vehicle weight positively affects the reporting probability of minor accidents, then our estimates will represent a lower bound on the effect of weight on external fatalities. If vehicle weight negatively affects the reporting probability of minor accidents, however, then our estimates of the effect of weight on external fatalities could be upwardly biased. To test whether the "ruggedness" hypothesis affects our results, we estimate our regressions while limiting the sample to collisions that do not involve any light trucks. This sample restriction does not reduce the coefficient estimates.¹² We also conduct a series of robustness tests in Section 6 that imply that the sampling frame does not bias our results.

Table 1 presents summary statistics from our two-vehicle collision data set. This data set contains all collisions involving two light vehicles built after 1980. We define a light vehicle as any car, pickup truck, SUV, or minivan that weighs between 1500 and 6000 pounds. We exclude collisions involving heavy trucks. The first two columns report statistics for the entire two-vehicle collision data set. The mean vehicle weight in this data set is 3076 pounds, and approximately 24.5% of vehicles are light trucks (pickups, SUVs, or minivans). The average model year is 1992, and the average number of occupants per vehicle is 1.41. The probability of a fatality in each vehicle is 0.19% (*i.e.* 0.0019), and the probability of a serious injury in each vehicle is 2.7%. Alcohol is involved in 8.3% of collisions.

11. The data set contains the population of police-reported accidents for Florida (1989–2005), Kansas (2001–05), Kentucky (1998–2005), Maryland (1989–99), Missouri (1989–2005), Ohio (1991–2005), Washington (2002–05), and Wyoming (1998–2005).

12. In the sample that excludes all collisions involving light trucks, the estimated effects are of similar magnitude to the analogous estimates from the main sample, reported in Table 2. This implies that the "ruggedness" hypothesis is not upwardly biasing our main results (see Supplementary Appendix Table A1).

TABLE 1
Summary statistics for two-vehicle collision data set

	Base sample		Complete covariates sample	
	Mean (Std dev)	Sample size	Mean (Std dev)	Sample size
Weight	3,076 lbs (685)	4,849,575	3,113 lbs (694)	2,829,768
Light truck	24.5% (43.0)	4,849,575	25.8% (43.8)	2,829,768
Model year	1992 (5.6)	4,849,575	1993 (5.7)	2,829,768
Accident year	1998 (4.4)	4,849,575	1999 (4.3)	2,829,768
Occupants	1.41 (0.84)	2,608,821	1.45 (0.87)	1,476,441
Fatality	0.19% (4.36)	4,849,575	0.23% (4.83)	2,829,768
Serious injury	2.7% (16.1)	4,849,575	3.4% (18.0)	2,829,768
Alcohol involved	8.3% (27.6)	2,753,533	10.0% (30.1)	1,723,694

Notes: Both samples are limited to collisions involving two light vehicles built post-1980. The complete covariates sample is further limited to collisions in which all covariates in our preferred specification are non-missing

The last two columns of Table 1 report summary statistics for the estimation sample with complete covariates. This sample is smaller than the overall two-vehicle collision sample because we drop collisions in which any of the covariates from our preferred specification are missing. This restriction reduces the sample from 4.8 million observations to 2.8 million. Nevertheless, the two samples appear similar along most observable measures.

5. SPECIFICATION AND RESULTS

Consider a collision involving two vehicles, Vehicle 1 and Vehicle 2. Suppose that we label Vehicle 1 as the “striking vehicle” and Vehicle 2 as the “struck vehicle”. These labels are for expositional purposes only—they do not signify which vehicle may be at fault in the collision.¹³ The external effects of vehicle weight are given by the effect of striking vehicle weight on the probability of fatalities in the struck vehicle. The internal effects of vehicle weight are given by the effect of struck vehicle weight on the probability of fatalities in the struck vehicle. The former is the quantity of policy interest, but we report results for the latter as well to calculate the effect of changes in fleet weight on fatalities.

We estimate the conditional expectation of a fatality in the struck vehicle as a function of striking vehicle weight, struck vehicle weight, and a rich set of covariates. We estimate the conditional expectation function (CEF) using either a linear probability model (LPM) or a probit. For robustness, we report estimates for both models.

We specify the linear probability model as follows:

$$E[\text{struck veh fatality}_i | \text{striking veh weight}_i, \text{struck veh weight}_i, X_{1i}, X_{2i}, W_i] \quad (6)$$

$$= \beta_1 \text{striking veh weight}_i + \beta_2 \text{struck veh weight}_i + X_{1i}\delta_1 + X_{2i}\delta_2 + W_i\delta_3$$

13. The labels are symmetric in that each vehicle enters our data set twice, once as the striking vehicle and once as the struck vehicle.

In equation (6), β_1 is the coefficient of interest, X_{1i} represents a set of characteristics pertaining to the striking vehicle in collision i , X_{2i} represents a set of characteristics pertaining to the struck vehicle in collision i , and W_i represents a set of collision-specific characteristics. The probit model modifies equation (6) to include the link function Φ , the normal CDF. The marginal effect of striking vehicle weight then varies with striking vehicle weight. For comparability with the LPM results, for each probit regression we report the average marginal effect across all observations included in that regression.¹⁴

5.1. *Effects of vehicle weight on fatalities and serious injuries*

Table 2 presents results from estimating the LPM and probit on the two-vehicle collision data set. The sample includes all accidents for which there is complete vehicle weight data for both vehicles; analyses restricted to states with low rates of missing weight data suggest that this constraint does not bias our results.¹⁵ Each vehicle appears in the two-vehicle collision data set twice, once as the struck vehicle and once as the striking vehicle. We therefore cluster the standard errors at the collision level to account for correlation between observations pertaining to the same collision. Alternatively, clustering at the vehicle model level for both the striking and struck vehicles does not affect our conclusions.

The first and second columns in Table 2 include the following covariates: vehicle curb weight, light truck indicators, and year fixed effects. A striking vehicle and struck vehicle version of each of the first two variables is included. The first column implies that a 1000 pound increase in weight in the striking vehicle is associated with a statistically significant 0.09 percentage point increase in the probability of a fatality in the struck vehicle ($t = 22.0$). This coefficient represents a 46% increase over the average probability of a fatality in a struck vehicle in this sample (0.19%). In comparison, a 1000 pound increase in weight in the struck vehicle is associated with a smaller 0.05 percentage point decrease in the probability of a fatality in the struck vehicle ($t = -11.8$). Striking light trucks increase the probability of a fatality in the struck vehicle by 0.12 percentage points (62% of the sample mean), even after controlling for striking vehicle weight ($t = 19.5$). The results from the probit model in column (2) display z -statistics that are similar to the t -statistics in column (1), and the average marginal effect generated by the probit model is of similar magnitude to the LPM coefficient (0.08 percentage points versus 0.09 percentage points).

Subsequent columns in Table 2 add additional covariates to the regressions. Columns (3) and (4) add controls for rain, darkness, day of week (weekday versus weekend), interstate highway, a quadratic in model year for each vehicle, and year, hour, and county fixed effects. The estimated effect of striking vehicle weight changes little in both the LPM and probit models. Columns (5) and (6) add controls for any seat belt usage, a quadratic in driver age, indicators for drivers under 21 or over 60, and indicators for male drivers or young male drivers. A striking vehicle and struck

14. Some of our probit regressions include fixed effects, raising the possibility of inconsistency due to the incidental parameters problem. However, in most cases we have many observations for each fixed effect, and as shown in Fernandez-Val (2009), the incidental parameters problem generates a trivial degree of bias in the probit model when estimating marginal effects (which are our quantities of interest).

15. Weight data are missing for vehicles for which we do not have VINs. The percentage of vehicles with missing weight data ranges from 17.4% (Ohio) to 54.5% (Maryland). When estimating our main statistical models on the four states with the lowest rates of missing weight data (Kentucky, Ohio, Washington, and Wyoming), we find that an additional 1,000 pounds of striking vehicle weight increases the probability of a fatality in the struck vehicle by 46–51%. When estimating the same models on the four states with highest rates of missing weight data (Florida, Kansas, Maryland, and Missouri), we find that an additional 1,000 pounds of striking vehicle weight increases the probability of a fatality in the struck vehicle by 44%. The rate of missing weight data thus appears to have little impact on our estimates (see Supplementary Appendix Table A2).

TABLE 2
Effect of vehicle weight on fatalities

Dependent variable: presence of fatalities in struck vehicle	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Weight of striking vehicle (1000s of lbs)	0.00088 (0.00004)	0.12685 (0.00531)	0.00093 (0.00005)	0.12797 (0.00616)	0.00101 (0.00005)	0.13440 (0.00685)	0.00110 (0.00006)	0.00104 (0.00004)	0.00065 (0.00006)	0.00064 (0.00006)
Effect of 1000 lb increase in striking vehicle weight	0.00088	0.00077	0.00093	0.00085	0.00101	0.00086	0.00110	0.00104	0.00065	0.00064
percent increase over sample mean	46%	41%	42%	38%	44%	37%	47%	44%	49%	48%
Weight of struck vehicle (1000s of lbs)	-0.00047 (0.00004)	-0.08196 (0.00644)	-0.00053 (0.00004)	-0.08548 (0.00727)	-0.00101 (0.00005)	-0.15988 (0.00815)	-0.00097 (0.00005)	-0.00104 (0.00004)	-0.00060 (0.00006)	-0.00065 (0.00006)
Striking vehicle is light truck	0.00117 (0.00006)	0.15977 (0.00861)	0.00105 (0.00008)	0.13016 (0.00958)	0.00088 (0.00009)	0.10113 (0.01075)	0.00093 (0.00009)	0.00098 (0.00008)	0.00054 (0.00010)	0.00054 (0.00010)
Struck vehicle is light truck	-0.00014 (0.00006)	-0.02541 (0.00987)	-0.00036 (0.00007)	-0.06192 (0.01081)	-0.00001 (0.00007)	-0.03605 (0.01154)	0.00021 (0.00008)	0.00026 (0.00007)	-0.00015 (0.00009)	-0.00012 (0.00009)
Specification	OLS	Probit	OLS	Probit	OLS	Probit	OLS	OLS	OLS	OLS
Weather, time, and county fixed effects		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Driver characteristics					Yes	Yes	Yes	Yes	Yes	Yes
City fixed effects							Yes	Yes	Yes	Yes
Weight coefficients set to be equal magnitude								Yes		
Occupants and seat belt usage										
Sample size	4,849,575	4,849,575	3,572,439	3,536,684	3,223,746	3,197,882	2,829,768	2,829,768	1,470,596	1,470,596

Notes: Each column represents a separate regression. The estimation sample is limited to collisions involving two vehicles. Parentheses contain standard errors clustered at the collision level. Effects of a 1000 lb increase in striking vehicle weight are computed as the average effect of a 1000 lb increase in weight across all observations included in the regression. All regressions include as right-hand-side variables the weight of each vehicle, indicators for whether each vehicle is a light truck, and year fixed effects. Weather, time, and county fixed effects controls include rain, darkness, day of week (weekday versus weekend), Interstate highway, a quadratic in model year for each vehicle, and year, hour, and county fixed effects. Driver characteristic controls include quadratics in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, and indicators for any seat belt usage in the vehicle.

vehicle version of each of these variables is included. The inclusion of these driver characteristics has minimal impact on the primary coefficient of interest (striking vehicle weight). They do, however, increase the magnitude of the struck vehicle weight coefficient to -0.10 percentage points ($t = -20.2$). It is now identical in magnitude to the striking vehicle weight coefficient.

Column (7) of Table 2 adds city fixed effects and is our preferred specification. City fixed effects should absorb any geographic heterogeneity in fatality rates that could be correlated with average vehicle weight. This issue would arise if, for example, heavy vehicles clustered in rural areas and these areas had deadlier accidents due to a prevalence of undivided highways or a sparseness of hospitals. At this point there are too many regressors to reliably estimate a probit model, and for many cities the city fixed effect perfectly predicts the fatality indicator, forcing the city to be dropped. We thus estimate only linear probability models in columns (7) through (10). The addition of city fixed effects has little impact on the coefficient on striking vehicle weight, changing it from 0.10 percentage points to 0.11 percentage points ($t = 18.3$). This coefficient represents a 47% increase over the average probability of a fatality in a struck vehicle in this sample.¹⁶

Comparing the striking vehicle weight coefficient (β_1) and struck vehicle weight coefficient (β_2) in columns (5) or (7) suggests that they may be of opposite sign but equal magnitude. Indeed, we cannot reject the hypothesis that $\beta_2 = -\beta_1$ in either column (5) or column (7). Column (8) estimates the same specification as column (7) but restricts β_2 to equal $-\beta_1$. This specification corresponds to our theoretical model; it imposes the assumption that weight only matters in a relative sense. Imposing this restriction has little impact on the estimates; a 1000 pound increase in striking vehicle weight now increases the probability of a fatality in the struck vehicle by 0.10 percentage points ($t = 26.0$).

Column (9) estimates the same specification as column (7) but limits the sample to observations for which we have data on the number of occupants per vehicle and the seat belt usage of each occupant (two controls we add in the next column). This restriction shrinks the sample in half and reduces the coefficient on striking vehicle weight to 0.07 percentage points ($t = 10.8$). However, the ratio of the coefficient to the average fatality rate in the sample remains stable (49%). The change in the coefficient simply reflects the fact that the restricted sample contains states with a lower threshold for reporting accidents, and thus a lower fatality rate per reported accident. Column (10) adds controls for the number of occupants per vehicle and seat belt usage rate of these occupants. The coefficient on striking vehicle weight is unchanged from column (9).¹⁷

The results in Table 2 suggest that selection bias has little impact on the striking vehicle weight coefficient but may affect the struck vehicle weight coefficient. In particular, the addition of driver characteristics in columns (5) and (6) notably impacts the struck vehicle weight coefficient but has little impact on the striking vehicle weight coefficient. When adding covariates one at a time, we find that almost all of the change in the struck vehicle weight coefficient between columns (4) and (6) can be attributed to the addition of the controls for driver age. The patterns suggest that older drivers drive heavier vehicles and that older drivers are more susceptible to dying in crashes. Since there is little correlation between the age of the struck vehicle's driver and the

16. An alternative question is whether the effect of weight in rural areas differs from its effect in urban areas. When we separate the Florida sample (which has a rural/urban indicator) into rural and urban areas, we do not find a statistically significant difference in the proportionate effect of weight across the two areas.

17. We experimented with flexibly controlling for manufacturer's suggested retail price (MSRP) as a proxy for driver wealth. This does not affect our coefficient estimates. In the Florida sample we also experimented with adding controls for drivers' insurance status, alcohol involvement, and negligent driving. These controls do not affect our coefficient estimates.

weight of the striking vehicle, however, the addition of driver age controls has no impact on the striking vehicle weight coefficient.¹⁸ Stated simply, heavy vehicles do not “seek out” elderly drivers to crash into.¹⁹

The results in Table 2 also suggest that the external risk posed by light trucks is not due solely to their heavy weight. The coefficient on the indicator for whether the striking vehicle is a light truck is positive and statistically significant in every column. In our preferred specification, column (7), the coefficient implies that being struck by a light truck increases the probability of a fatality by 0.09 percentage points ($t = 10.3$), even after conditioning on striking vehicle weight. This represents a 40% increase over the average fatality rate in the sample. In comparison, if we do not control for vehicle weight, then the light truck coefficient doubles to 0.18 percentage points (*i.e.* 0.0018). This effect is roughly similar in magnitude to the external effects of light trucks in two-vehicle collisions that White (2004) and Anderson (2008) estimate. The additional risk posed by light trucks may be due to the stiffness of their frames or their height incompatibility with other vehicles (Hakim, 2003). However, the robustness tests that we perform in Section 6 for the vehicle weight coefficient do not apply to the light truck coefficient. Thus we cannot rule out the possibility that some of the light truck coefficient may represent driver selection effects; *i.e.*, consumers that purchase light trucks may drive in an aggressive manner that generates particularly severe collisions. For this reason we do not incorporate the light truck coefficient when calculating the total externality across all vehicles in Section 7. If we were to incorporate the light truck coefficient, the total externality would be even larger. In the context of CAFE standards, however, we do consider the potential risks that light trucks pose.

While 90% of multivehicle collisions involve two vehicles, 9% involve three vehicles, and 1% involve four or more vehicles (a three or four vehicle collision is one in which three or four vehicles are damaged in the same collision, though each vehicle need not have collided with every other vehicle in the collision). Adding 1000 pounds to a vehicle in a three-vehicle collision should increase the risk of a fatality in the other two vehicles by less than 47% each (our preferred estimate from the two-vehicle collision data set). This attenuation occurs because the extra mass of the first vehicle is, in expectation, now distributed across two other vehicles rather than one other vehicle. We estimate the relationship between vehicle weight and fatalities in three-vehicle collisions in Table 3. For expositional purposes, assume that Vehicle 1 is the struck vehicle and that Vehicles 2 and 3 are the striking vehicles. In Table 3, the striking vehicle weight coefficient represents the average effect of a 1000 pound increase in the weight of either Vehicle 2 or 3 (but not both) on the probability of a fatality in Vehicle 1. The striking vehicle weight coefficient is positive and statistically significant in all specifications, and the magnitude of the coefficient ranges from 28% to 42% of the average probability of a fatality. Our preferred estimate, column (7), implies that a 1000 pound increase in one vehicle raises the probability of a fatality in either of the other two vehicles by 35%.

Table 4 presents results from estimating versions of the LPM and probit in which the dependent variable is the presence of serious injuries in the struck vehicle. The regressions are analogous to those in Table 2, but the dependent variable has changed from any fatalities to any serious injuries. The striking vehicle weight coefficients (or marginal effects, for the probit regressions) in Table 4 are approximately 6 times larger than the corresponding coefficients in Table 2. This difference

18. A related question is whether dangerous driver characteristics modify the striking vehicle weight effect. We tested whether the effect of striking vehicle weight differs when the striking vehicle is driven by a young male or someone over the age of 60. We did not find statistically significant differences.

19. A related concern is that fatalities may consist disproportionately of elderly drivers with below-average VSLs. However, in a subsample with the most detailed data, accounting for age-specific VSLs changes our estimates by only 5.5%. Data limitations prevent us from incorporating age-specific VSLs into our main regressions.

TABLE 3
Effect of vehicle weight on fatalities in three-vehicle accident

Dependent variable: presence of fatalities in struck vehicle	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Weight of striking vehicle (1000s of lbs)	0.00088 (0.00010)	0.11785 (0.01203)	0.00086 (0.00012)	0.10934 (0.01391)	0.00081 (0.00013)	0.10541 (0.01583)	0.00086 (0.00014)	0.00065 (0.00019)	0.00064 (0.00019)
Effect of 1000 lb increase in striking vehicle weight	42%	37%	36%	32%	32%	28%	35%	42%	41%
Weight of struck vehicle (1000s of lbs)	-0.00062 (0.00011)	-0.09830 (0.01865)	-0.00078 (0.00014)	-0.11463 (0.02104)	-0.00127 (0.00016)	-0.19317 (0.02415)	-0.00109 (0.00017)	-0.00079 (0.00023)	-0.00084 (0.00024)
Striking vehicle is light truck	0.00064 (0.00014)	0.08021 (0.01863)	0.00055 (0.00017)	0.06840 (0.02044)	0.00056 (0.00019)	0.06530 (0.02345)	0.00061 (0.00020)	0.00039 (0.00030)	0.00039 (0.00030)
Struck vehicle is light truck	-0.00022 (0.00018)	-0.03293 (0.02969)	-0.00031 (0.00022)	-0.05346 (0.03178)	0.00031 (0.00024)	0.00197 (0.03484)	0.00051 (0.00026)	-0.00014 (0.00036)	-0.00012 (0.00036)
Specification	OLS	Probit	OLS	Probit	OLS	Probit	OLS	OLS	OLS
Weather, time, and county fixed effects			Yes	Yes	Yes	Yes	Yes	Yes	Yes
Driver characteristics					Yes	Yes	Yes	Yes	Yes
City fixed effects						Yes	Yes	Yes	Yes
Occupants and seat belt usage									
Sample size	518,378	518,378	391,456	356,970	348,543	306,684	317,769	110,541	110,541

Notes: Each column represents a separate regression. The estimation sample is limited to collisions involving three vehicles. Striking vehicle weight coefficients represent the average effect of increasing the weight of one striking vehicle by 1000 pound; they are the average of the coefficients on the first and second striking vehicles. Parentheses contain standard errors clustered at the collision level. Effects of a 1000 lb increase in striking vehicle weight are computed as the average effect of a 1000 lb increase in the weight of one striking vehicle across all observations included in the regression. All regressions include as right-hand-side variables the weight of each vehicle, indicators for whether each vehicle is a light truck, and year fixed effects. Weather, time, and county fixed effects controls include rain, darkness, day of week (weekday versus weekend), Interstate highway, a quadratic in model year for each vehicle, and year, hour, and county fixed effects. Driver characteristic controls include quadratics in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, and indicators for any seat belt usage in the vehicle.

TABLE 4
Effect of vehicle weight on serious injuries

Dependent variable: presence of serious injuries in struck vehicle	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Weight of striking vehicle (1000s of lbs)	0.00484 (0.00014)	0.07616 (0.00207)	0.00573 (0.00017)	0.08332 (0.00242)	0.00615 (0.00019)	0.08494 (0.00260)	0.00687 (0.00021)	0.00324 (0.00020)	0.00316 (0.00020)
Effect of 1000 lb increase in striking vehicle weight	0.00484	0.00470	0.00573	0.00543	0.00615	0.00563	0.00687	0.00324	0.00316
percent increase over sample mean	18%	18%	19%	18%	19%	17%	20%	20%	20%
Weight of struck vehicle (1000s of lbs)	-0.00720 (0.00013)	-0.12392 (0.00226)	-0.00797 (0.00016)	-0.12874 (0.00261)	-0.00921 (0.00018)	-0.14280 (0.00282)	-0.00891 (0.00020)	-0.00479 (0.00019)	-0.00514 (0.00019)
Striking vehicle is light truck	0.00567 (0.00022)	0.08524 (0.00322)	0.00456 (0.00027)	0.06400 (0.00367)	0.00412 (0.00029)	0.05399 (0.00392)	0.00448 (0.00032)	0.00215 (0.00030)	0.00215 (0.00030)
Struck vehicle is light truck	-0.00033 (0.00020)	-0.00894 (0.00343)	-0.00212 (0.00024)	-0.03734 (0.00389)	-0.00075 (0.00027)	-0.02072 (0.00414)	0.00005 (0.00029)	-0.00069 (0.00028)	-0.00048 (0.00028)
Specification	OLS	Probit	OLS	Probit	OLS	Probit	OLS	OLS	OLS
Weather, time, and county fixed effects			Yes	Yes	Yes	Yes	Yes	Yes	Yes
Driver characteristics					Yes	Yes	Yes	Yes	Yes
City fixed effects							Yes	Yes	Yes
Occupants and seat belt usage									
Sample size	4,849,575	4,849,575	3,572,439	3,571,255	3,223,746	3,223,344	2,829,768	1,470,596	1,470,596

Notes: Each column represents a separate regression. The estimation sample is limited to collisions involving two vehicles. Parentheses contain standard errors clustered at the collision level. Effects of a 1000 lb increase in striking vehicle weight are computed as the average effect of a 1000 lb increase in weight across all observations included in the regression. All regressions include as right-hand-side variables the weight of each vehicle, indicators for whether each vehicle is a light truck, and year fixed effects. Weather, time, and county fixed effects controls include rain, darkness, day of week (weekday versus weekend), Interstate highway, a quadratic in model year for each vehicle, and year, hour, and county fixed effects. Driver characteristic controls include quadratics in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, and indicators for any seat belt usage in the vehicle.

arises because the probability of a serious injury in this sample is roughly 15 times higher than the probability of a fatality. In the preferred specification, column (7), a 1000 pound increase in striking vehicle weight raises the probability of serious injuries in the struck vehicle by 0.7 percentage points ($t=32.7$). This figure represents 20% of the average probability of a serious injury in this sample.²⁰ Drawing on our discussion of liability insurance above, it is not clear what proportion of serious injuries that represent external costs is internalized through existing insurance contracts. We therefore focus on fatalities for the remainder of the article, which is a conservative approach.

5.2. *Non-linear effects of vehicle weight*

The linear specifications in Table 2 may obscure significant nonlinearity in the relationship between vehicle weight and fatalities. As our model in Section 3 establishes, nonlinearity in the relationship between weight and fatalities can have important implications. First, the marginal externality may vary with a vehicle's weight. Second, if weight has nonlinear effects, then the marginal externality may change if the fleet's weight distribution changes. Finally, the effect of the fleet's weight distribution on total fatalities depends on nonlinearity in the relationship between weight and fatalities. We focus on the first implication in this section and discuss the latter two implications in Section 7.

Figure 1 presents the relationship between striking vehicle weight and struck vehicle fatalities for three specifications. The first specification is linear in striking vehicle weight and corresponds to column (7) of Table 2. The solid line in Figure 1 plots the predictions from this specification; they are linear by construction.²¹ The second specification relaxes the linearity assumption. In this specification, striking vehicle weight and struck vehicle weight each appear as 5th order polynomials in the estimating equation, and we include an interaction between striking vehicle and struck vehicle weight as well. The dashed line in Figure 1 plots the predictions from this specification; they closely track the predictions from the linear specification but diverge above 4000 pounds. However, the confidence intervals become large at high weights, and we cannot reject the hypothesis that the two specifications yield similar predictions above 4000 pounds. The third specification models fatalities in the struck vehicle as a 5th order polynomial of the difference in vehicle weights:

$$\text{struck veh fatality}_i = \sum_{j=1}^5 \beta_j (\text{striking weight}_i - \text{struck weight}_i)^j + X_{1i}\delta_1 + X_{2i}\delta_2 + W_i\delta_3 + \varepsilon_i \quad (7)$$

The dotted line in Figure 1 plots the predictions from this specification; they closely track the predictions from the linear specification.²² We find no evidence that struck vehicle weight has a

20. It is notable that weight's effect on serious injuries (20%) is smaller than weight's effect on fatalities (47%). A major finding in the crash-safety literature is that Δv —the change in velocity—is the most important determinant of crash fatalities (Joks, 1993). This implies that relative weight has a protective effect in multi-vehicle collisions because the heavier vehicle experiences lower Δv ; if both vehicles are traveling at similar speeds, the heavier vehicle continues to travel forward after the initial collision, while the lighter vehicle actually reverses direction. Δv is also an important determinant of injuries. However, Δv 's effect on injuries is less than its effect on fatalities. Elvik *et al.* (2004) conduct a meta-analysis of studies relating crash velocities, fatalities, and injuries. They find that, on average, fatalities are proportional to $(\Delta v)^{4.0}$ while serious injuries are proportional to $(\Delta v)^{2.4}$. The result that the fatality effect is larger than the injury effect is thus broadly consistent with the crash-safety literature.

21. In Figures 1 and 2 we set regressors that are not functions of striking vehicle weight to their sample means.

22. At each striking vehicle weight w , we form predictions from equation (7) by computing the average difference in vehicle weights for collisions in which the striking vehicle weighs w pounds. We do the same for the average square, cube, quartic, and quintic of difference in vehicle weights.




FIGURE 1

Relationship between striking vehicle weight and struck vehicle fatalities

protective effect independent of its effect through relative weight; if we include struck vehicle weight as a separate regressor in equation (7), it has a coefficient of -0.003 percentage points ($t = -0.4$). Overall the results in Figure 1 suggest that the relationship is roughly linear over the support of striking vehicle weight, though it appears somewhat convex for very light striking vehicles (sub-2500 pounds).²³

Figure 2 presents the relationship between struck vehicle fatalities and the difference in vehicle weights. The solid line plots the predictions from estimating equation (7) over the support of the difference in vehicle weights. The marginal effect of striking vehicle weight increases from -2000 pounds to -500 pounds (*i.e.* when the struck vehicle outweighs the striking vehicle by 500 to 2000 pounds) and is roughly constant thereafter. The convexity in Figure 2 implies concavity in f , the private safety benefit of relative weight, from our theoretical model. The relationship does not flatten at high weight differences (*e.g.* $+2000$ pounds), implying that adding weight to the heavier vehicle continues to increase the risk in the lighter vehicle even in cases of severe mismatch. Nevertheless, the convexity at weight differences below -500 pounds has a modest impact on the marginal externality for lighter vehicles (see Figure 1) because, even among sub-2500 pound striking vehicles, the majority of collisions involve a weight difference between -700 and $+2000$ pounds.

6. ALTERNATIVE SOURCES OF IDENTIFICATION AND FALSIFICATION TESTS

The results in Section 5 demonstrate a strong relationship between striking vehicle weight and struck vehicle fatalities. The robustness of this relationship to the inclusion of a rich set of

23. In a previous version of this article we also estimated more flexible versions of the probit model. A probit in which weight enters the normal CDF linearly fits the data poorly, as it is an inherently non-linear model that forces the marginal effect of weight to steeply increase in accidents involving heavy striking vehicles. A probit in which weight enters the normal CDF as a higher order polynomial approximates the LPM specifications we present in this figure.

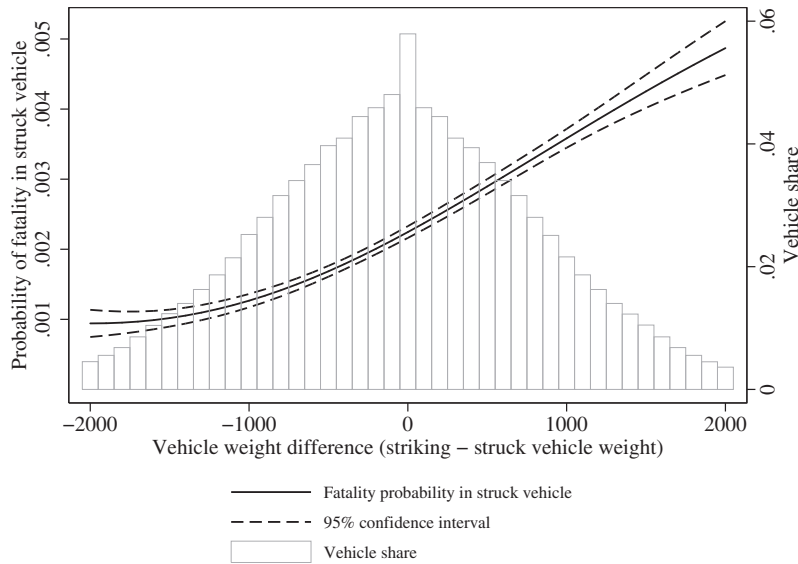


FIGURE 2

Relationship between difference in vehicle weights and struck vehicle fatalities

accident and driver characteristics, as well as very fine geographic fixed effects, suggests that the striking vehicle weight coefficients represent causal effects of weight on fatality risk. However, two potential sources of upward bias seem particularly plausible. First, driver selection may bias the coefficient estimates if heavier vehicles attract aggressive drivers who get into deadlier accidents. Note, however, that only selection of drivers who get into deadlier accidents, rather than drivers who get into more accidents, could bias our estimates.²⁴ Second, the sampling frame might bias the coefficient estimates if minor collisions involving heavier vehicles are less likely to be reported to the police, all other factors held constant.²⁵ To test whether either of these factors could bias our results, we conduct three exercises. First, we estimate the effect of striking vehicle weight on fatalities using within-model changes in vehicle weight that occur when models are refreshed. Second, we estimate the effect of striking vehicle weight on fatalities using striking vehicle occupants as an additional source of variation in weight. Finally, we implement a series of falsification tests that we benchmark against engineering safety estimates.

24. Because our estimates are conditional on a collision occurring, only specific types of driver selection can generate bias. Selection of “careless” drivers who simply get into more accidents of the same expected severity would not bias our results. It would increase the number of times we observe these drivers in the sample, but it would not increase the probability that someone dies in a collision conditional on the collision occurring. Selection of “aggressive” drivers who get into more severe accidents could bias our results, however. These drivers could increase the probability that someone dies in a collision conditional on the collision occurring.

25. Note that, unlike the struck vehicle weight coefficients, striking vehicle weight coefficients are unlikely to be biased by any correlation between vehicle weight and vehicle safety features. It is plausible that heavier vehicles may be more or less likely to have safety features such as airbags, side impact protection beams, and unibody construction. However, these safety features are much more helpful to the striking vehicle’s own occupants than they are to the occupants of other vehicles that the striking vehicle hits.

6.1. *Vehicle model fixed effects results*

To establish the robustness of our results, we explore two alternative sources of identification. Our first alternative leverages within-model changes in vehicle weight to estimate the effect of striking vehicle weight on fatalities. To implement this design, we include vehicle model fixed effects for the striking vehicle in our preferred specification. The effect of striking vehicle weight on fatalities is thus identified on the basis of changes in vehicle weight that occur when a vehicle model is refreshed. This design minimizes the impact of driver selection as long as the composition of customers for a particular vehicle model remains relatively stable when the model is refreshed.

A key concern with including vehicle model fixed effects is that they may absorb almost all of the variation in striking vehicle weight, leaving little variation remaining to identify the effect of interest. However, summary statistics imply that there is sufficient within-model weight variation to identify an effect. For example, the overall standard deviation in vehicle weight is 520 pounds, while the within-model standard deviation in vehicle weight is 280 pounds.²⁶ We observe within-model deviations in vehicle weight as large as 1025 pounds in the data (this is the 99th percentile of within-model deviations in vehicle weight). Substantial variation in vehicle weight thus remains even after including the vehicle model fixed effects.

Table 5 reports estimates from models that include vehicle model fixed effects (as well as year fixed effects and all other controls from column (5) of Table 2). Column (1) presents results from our preferred specification estimated on the sample for which we have complete vehicle model data. The sample size is substantially smaller than our main analytic sample because only four states—Kentucky, Maryland, Ohio, and Wyoming—report detailed vehicle model data. In this subsample, a 1000 pound increase in striking vehicle weight is associated with a 47% increase in the probability of a fatality in the struck vehicle (0.06 percentage points, $t=7.3$). This effect is consistent with the estimates from Section 5. Column (2) presents results from the same specification with vehicle model fixed effects added. A 1000 pound increase in striking vehicle weight is now associated with a 58% increase in the probability of a fatality in the struck vehicle (0.07 percentage points, $t=4.5$). The correspondence between the two coefficient estimates suggests that driver selection does not seriously bias our results, and we cannot reject the hypothesis that both coefficients converge to the same value.

One issue with the vehicle model fixed effects specification is that our sample contains vehicles constructed from 1981 to 2006. In an extreme case, the fixed effects regression could compare a 1981 Honda Accord to a 2006 Honda Accord. These two model years differ in weight by over 1000 pounds, but it is likely that the owners of 1981 Honda Accords are very different from the owners of 2006 Honda Accords. To ensure that we only compare vehicles of roughly similar vintage, we interact the vehicle model fixed effects with four model-year group fixed effects: 1986–90, 1991–95, 1996–2000, and 2000–06 model years (the omitted category is 1981–85). We include the interacted set of fixed effects as controls.

Column (3) reports results from a specification that controls for the full set of interactions between the vehicle model fixed effects and the model-year group fixed effects. In this specification a 1,000 pound increase in vehicle weight is associated with a 57% increase in the probability of a fatality in the struck vehicle (0.07 percentage points, $t=3.9$). This estimate is almost identical to the estimate in column (2), and the standard error is only 13% larger. The drop

26. These figures represent standard deviations of the portion of vehicle weight that is not explained by the other controls in our regressions. They are thus smaller than the raw standard deviation of vehicle weight reported in Table 1. It is important to remove the portion of vehicle weight that is explained by the other controls in our regressions because this variation is not used to identify the effect of vehicle weight regardless of whether we include model fixed effects.

TABLE 5
Effect of vehicle weight on fatalities using alternative sources of identification

Dependent variable: presence of fatalities in struck vehicle	(1)	(2)	(3)	(4)	(5)
Occupant weight in striking vehicle (1000s of lbs)					0.00062 (0.00026)
Curb weight of striking vehicle (1000s of lbs)	0.00058 (0.00008)	0.00072 (0.00016)	0.00071 (0.00018)	0.00064 (0.00006)	0.00063 (0.00006)
Percentage effect of 1000 lb increase in striking vehicle weight	47%	58%	57%	48%	46%
Specification	OLS	OLS w/model FEs	OLS w/model by model-year-group FEs	OLS	OLS
Sample size	1,011,982	1,011,982	1,011,982	1,475,762	1,475,762

Notes: Each column represents a separate regression. The estimation sample is limited to collisions involving two vehicles. Parentheses contain standard errors clustered at the collision level. All regressions include the following right-hand-side variables: weight of each vehicle, a quadratic in model year for each vehicle, indicators for whether each vehicle is a light truck, rain, darkness, day of week (weekday versus weekend), Interstate highway, quadratics in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, indicators for any seat belt usage in the vehicle, and year, hour, and county fixed effects. OLS regressions with model fixed effects contain fixed effects for each vehicle model. OLS regressions with model by model-year-group fixed effects contain fixed effects for each vehicle model interacted with model year indicators for the 1986–90, 1991–95, 1996–2000, and 2001–09 model years. Regressions in columns (4) and (5) contain city fixed effects. Occupant weight in the striking vehicle is calculated as the number of occupants in the striking vehicle times 164 lbs per occupant.

in statistical power from adding the model–year group interactions is modest for two reasons. First, almost 60% of vehicles in our sample are built between 1989 and 2000. Second, vehicle model identifiers typically change over long periods of time, so a single vehicle model fixed effect often does not span two decades. As a result, the average difference in model year between two randomly selected vehicles is only 3.8 years after controlling for the vehicle model fixed effects. The specification in column (2) thus primarily compares vehicles of similar vintage even without including model-year group interactions.

6.2. Occupant weight results

Our second alternate source of identification leverages the number of occupants in the striking vehicle as an additional source of variation in striking vehicle weight. The number of occupants in the striking vehicle directly affects the striking vehicle’s weight, so we estimate the regression:

$$\text{struck veh fatality}_i = \alpha_1 \text{striking veh occupant weight}_i + \alpha_2 \text{striking veh curb weight}_i + X_{1i}\gamma_1 + X_{2i}\gamma_2 + W_i\gamma_3 + \varepsilon_i \quad (8)$$

In this regression the striking vehicle’s occupant weight equals the number of occupants in the striking vehicle multiplied by 164 pounds (the average weight of an additional occupant circa 2000).²⁷ The coefficient α_1 represents the effect of additional striking vehicle occupant weight on

27. We calculate this figure as follows. First, for the subset of accidents for which we have detailed occupant characteristics, we tabulate the share of additional occupants that are male adults, female adults, male children, and female children. We find that 21.6% of additional occupants are male adults, 39.2% are female adults, 19.2% are male children, and 20.0% are female children. Using national statistics on body weight by gender and age we then compute the average weight of an additional occupant as $0.216 \times 190 \text{ lbs} + 0.392 \times 163 \text{ lbs} + 0.192 \times 110 \text{ lbs} + 0.200 \times 114 \text{ lbs} = 149 \text{ lbs}$ (Ogden *et al.*, 2004). Finally, we add 15 lbs per occupant to account for clothing, outerwear, and personal belongings ($149 \text{ lbs} + 15 \text{ lbs} = 164 \text{ lbs}$).

struck vehicle fatalities. It is worth emphasizing that the regression controls for the curb weight of each vehicle (*i.e.* each vehicle's weight absent any passengers or cargo). Identification of α_1 thus comes from variation in the number of occupants in the striking vehicle *after* controlling for the curb weight of the striking vehicle. This means that the variation used to identify α_1 is, by the properties of linear regression, orthogonal to the variation in curb weight (*i.e.* the variation that we use in Section 5).

Nevertheless, it is not obvious that the number of occupants in the striking vehicle is uncorrelated with any other factors that affect fatalities in the struck vehicle. It is possible that, even after controlling for vehicle curb weight and other characteristics, drivers who carry additional occupants in their vehicles drive more aggressively than drivers who do not carry additional occupants. If this were true, then our estimates of α_1 would be biased upward. We thus do not interpret our estimates of α_1 as being more robust than our estimates of α_2 (which correspond to the estimates reported in Section 5). Instead, we recognize that the identifying variation for α_1 is orthogonal to the identifying variation for α_2 (a fact guaranteed by the inclusion of curb weight as a control in the regression). If the regression produces similar estimates of α_1 and α_2 , this suggests that both coefficients are estimating causal effects. If the regression produces very different estimates of α_1 and α_2 , this suggests that one (or both) estimates may be biased. This comparison is similar in spirit to general overidentification tests that test whether different instruments generate similar coefficient estimates.

The last two columns of Table 5 report coefficients from the sample with occupant data. The occupant data sample is approximately half the size of our main analytic sample because data on the number of occupants is not available in every state. Column (4) presents results from estimating the preferred OLS specification (column (7) of Table 2) on the occupant data sample. A 1000 pound increase in striking vehicle weight is associated with a statistically significant 0.064 percentage point increase in the probability of a fatality in the struck vehicle ($t=10.7$). This coefficient represents a 48% increase over the average probability of a fatality in a struck vehicle, which is consistent with the results in Section 5. Column (5) presents results from estimating equation (8). The first reported coefficient is α_1 , the coefficient on occupant weight in the striking vehicle. An additional 1000 pounds of occupant weight in the striking vehicle is associated with a statistically significant 0.062 percentage point increase in the probability of a fatality in the struck vehicle ($t=2.4$). This coefficient represents a 46% increase over the average probability of a fatality in the struck vehicle. It is almost identical to the striking vehicle curb weight coefficient in column (4). The second reported coefficient is α_2 , the coefficient on striking vehicle curb weight. This coefficient is 0.063 percentage points ($t=10.5$), which is virtually identical to α_1 . The correspondence between the two coefficients increases our confidence in both sources of identifying variation.

6.3. *Falsification tests*

Suppose that heavier vehicles pose no additional risk to other vehicles, and that the estimates reported in Section 5 simply reflect the possibility that drivers of heavier vehicles are more aggressive (regardless of vehicle weight) or that heavier vehicles are less likely to generate police reports. In that case, there should be a strong positive correlation between vehicle weight and fatalities or injuries when analysing two-vehicle collisions between vehicles of the same weight. These accidents therefore provide an opportunity to test whether driver selection bias or sampling frame bias are generating our results.

It is possible, however, that heavier vehicles are safer than lighter vehicles. In that case, a positive driver selection effect might be mitigated by a negative weight effect. Put simply, even if drivers of heavier vehicles drive aggressively, our falsification test might generate a small

TABLE 6
Relationship between vehicle weight and NHTSA crash test performance

Dependent variable	HIC (1)	HIC>700 (2)	HIC (3)	HIC>700 (4)
Weight of vehicle	17.7 (16.1)	0.024 (0.018)	38.2 (43.5)	0.018 (0.040)
Percentage effect of 1000 lb increase	3.0%	8.7%	6.7%	7.2%
Sales share weighted			Yes	Yes
Sample size	4788	4788	2847	2847

Notes: Each column represents a separate regression. The estimation sample in the first two columns contains all NHTSA vehicle-to-barrier frontal crash test results. The estimation sample in the last two columns contains only crash tests involving vehicles for which we have sales share data. Parentheses contain standard errors clustered by vehicle make. All regressions include the following right-hand-side variables: weight of tested vehicle, a quadratic in model year, a light truck indicator, and a quadratic in collision speed. Sales share weighted regressions are weighted by the tested vehicle's sales share for a given year.

coefficient because the heavier vehicles are fundamentally safer. We therefore benchmark the results of our falsification tests against the results of NHTSA crash tests. NHTSA crash tests entail colliding a vehicle with a concrete barrier; they are meant to simulate the results of a collision with a stationary object or a head-on collision with another vehicle of similar weight.²⁸ The primary outcome in the NHTSA crash test is the Head Injury Criterion (HIC). This variable is derived from an accelerometer mounted on the crash test dummy's head and measures the forces that the head is exposed to. A higher HIC value corresponds to a higher probability of severe or fatal head injury.

Table 6 presents results from regressions of HIC scores on vehicle weight using the NHTSA crash test data. All regressions include as controls a light truck indicator, a quadratic in vehicle model year, and a quadratic in collision speed. The estimation sample in the first two columns contains all NHTSA vehicle-to-barrier frontal crash tests conducted from 1980 to 2009 (the mean year is 1997). Column (1) reports regression results when the dependent variable is HIC. The results indicate that an additional 1000 pounds of vehicle weight is associated with a statistically insignificant 3% increase in HIC (17.7 points). Column (2) reports regression results when the dependent variable is an indicator for whether HIC exceeds 700. This threshold is of interest because it represents the point at which there is a significant (5%) chance of severe brain injury (Mertz *et al.*, 1997). The results indicate that an additional 1000 pounds of vehicle weight is associated with a statistically insignificant 8.7% increase in the probability that HIC exceeds 700 (2.4 percentage points). The composition of vehicles that NHTSA tests, however, is not identical to the composition of vehicles on the roadways. To account for this, we estimate regressions in which each test result is weighted by the sales share of the tested vehicle (Ward's Reports Incorporated, 1994). Columns (3) and (4) report results from these regressions. The sample size falls because we do not have sales share data for every tested vehicle, but the results are qualitatively unchanged. An additional 1000 pounds of vehicle weight is associated with small, statistically insignificant increases in HIC or the probability that HIC exceeds 700. Overall, there is a weak positive relationship between vehicle weight and HIC values. The point estimates suggest that an additional 1000 pounds of vehicle weight could raise the fatality rate by 3% to

28. Many real world collisions involve side or rear impacts. In the Florida SDS data, we find that rear impacts are the safest, with a fatality rate that is 30% of the average, frontal impacts are the next safest, with a fatality rate that is 77% of the average, and side impacts are the most dangerous, with a fatality rate that is 363% of the average. However, in all collisions types the effect of adding 1,000 pounds of striking vehicle weight ranges from 37% to 45% of the average fatality rate for that collision type.

TABLE 7
Effect of vehicle weight in collisions between two equal weight vehicles

Dependent variable: presence of fatalities in struck vehicle	(1)	(2)	(3)	(4)
Average vehicle weight in collision (1000s of lbs)	−0.00004 (0.00017)	−0.00110 (0.00108)	0.00003 (0.00023)	−0.00181 (0.00137)
Percentage effect of 1000 lb increase in average weight	−2%	−19%	2%	−35%
Max weight difference between vehicles	200 lbs	200 lbs	100 lbs	100 lbs
Frontal collisions only		Yes		Yes
Sample size	539,350	39,242	288,988	20,488

Notes: Each column represents a separate regression. The estimation sample is limited to collisions in which the difference in weight between the two vehicles is less than 200 lbs (first two columns) or 100 lbs (last two columns). Parentheses contain standard errors clustered at the collision level. Effects of a 1000 lb increase in striking vehicle weight are computed as the average effect of a 1000 lb increase in weight across all observations included in the regression. All regressions include the following right-hand-side variables: weight of each vehicle, a quadratic in model year for each vehicle, indicators for whether each vehicle is a light truck, rain, darkness, day of week (weekday versus weekend), Interstate highway, quadratics in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, indicators for any seat belt usage in the vehicle, and year, hour, and city fixed effects.

9%, but none of the coefficients are statistically significant. We thus expect a weak relationship between vehicle weight and fatalities in collisions between two equal weight vehicles if our research design is sound.

Table 7 presents results from regressions in which the estimation sample consists of collisions involving two vehicles of similar weight—the difference in vehicle weight cannot exceed 200 pounds. In each regression, an indicator for fatalities in the struck vehicle is regressed on the average weight of the two vehicles and the set of controls from our preferred specification. Column (1) indicates that an increase of 1000 pounds in average vehicle weight predicts a statistically insignificant 2% decrease in the probability of a fatality (0.00 percentage points). Column (2) restricts the sample to head-on collisions between two vehicles of the same weight, the type of collision simulated by NHTSA. In this sample, an increase of 1000 pounds in average vehicle weight predicts a statistically insignificant 19% decrease in the probability of a fatality (0.11 percentage points).²⁹ Columns (3) and (4) replicate columns (1) and (2) but restrict the sample so that the difference in vehicle weight cannot exceed 100 pounds. The estimates remain small or negative and statistically insignificant, but are less precisely estimated.

Overall, the estimates in Table 7 indicate that there is a weak relationship between vehicle weight and fatalities in collisions between two vehicles of equal weight, and we cannot reject the hypothesis that this relationship is zero. This finding is consistent with NHTSA crash test results (Table 6) and inconsistent with the hypothesis that driver selection bias or sampling frame bias is generating the results in Section 5. The most precise estimate in Table 7—column (1)—suggests that increasing average vehicle weight by 1000 pounds decreases the fatality rate by 2%. This figure is not statistically different from the coefficients implied by the NHTSA crash test data. In contrast, if the relationship between striking vehicle weight and struck vehicle fatalities were generated by driver selection bias or sampling frame bias, then we would expect a large positive coefficient on average vehicle weight when two vehicles of equal weight collide. The preferred estimate from Section 5 indicates that a 1000 pound increase in striking vehicle weight raises the probability of a fatality in the struck vehicle by 47%. If this coefficient represented driver

29. The average probability of a fatality is much higher in column (2) than in column (1) because head-on collisions are more dangerous than the average collision.

TABLE 8
Effect of vehicle weight on fatalities in single-vehicle collision

Dependent variable: presence of fatalities in struck vehicle	(1)	(2)	(3)	(4)
Weight of vehicle (1000s of lbs)	0.00044 (0.00026)	0.00030 (0.00048)	0.00153 (0.00655)	0.00434 (0.01330)
Effect of 1000 lb increase in vehicle weight percent increase over sample mean	0.00044 3%	0.00030 2%	0.00005 0%	0.00014 1%
Collision type	1 vehicle	1 vehicle, frontal	1 vehicle	1 vehicle, frontal
Specification	OLS	OLS	Probit	Probit
Sample size	774,790	224,696	916,766	223,236

Notes: Each column represents a separate regression. The estimation sample is limited to collisions involving a single vehicle. Parentheses contain robust standard errors. Effects of a 1000 lb increase in vehicle weight are computed as the average effect of a 1000 lb increase in weight across all observations included in the regression. All regressions include the following right-hand-side variables: weight of vehicle, a quadratic in model year, indicators for whether a vehicle is a light truck, rain, darkness, day of week (weekday versus weekend), Interstate highway, quadratic in driver age, indicators for drivers under 21 or over 60, indicators for male drivers and young male drivers, indicators for any seat belt usage in the vehicle, and year, hour, and either city fixed effects (OLS) or county fixed effects (probit).

selection bias, and if two aggressive drivers were twice as dangerous as one aggressive driver, then we might expect a 1000 pound increase in both vehicles weights to raise the probability of a fatality by 94% ($2 \times 47 = 94$). However, no coefficient in Table 7 is above 2%, and we can reject any effect above 12% in either column (1) or (2) at the 5% significance level.

As an additional set of falsification tests, we examine the relationship between vehicle weight and fatalities in collisions involving a single vehicle. If drivers of heavier vehicles are more aggressive, then we expect a strong positive relationship between vehicle weight and fatalities in these collisions. Table 8 presents results for single-vehicle collisions. In these collisions, we regress a fatality indicator on vehicle weight and other controls. The results in column (1) pertain to all single-vehicle collisions; a 1000 pound increase in vehicle weight is associated with a 3% increase in the probability of a fatality (0.04 percentage points).³⁰ Column (2) pertains to single-vehicle frontal collisions, the type of collision simulated by NHTSA. A 1000 pound increase in vehicle weight is associated with a 2% increase in the probability of a fatality (0.03 percentage points). Columns (3) and (4) present results that are analogous to columns (1) and (2) but are estimated using a probit specification instead of a linear probability model.³¹ In both columns, a 1000 pound increase in vehicle weight is associated with an increase of less than 1% in the probability of a fatality. In all columns, the percentage effects fall close to the range implied by the NHTSA crash test data, further suggesting no substantial bias due to driver selection.

7. POLICY IMPLICATIONS

The econometric evidence demonstrates that the impact of heavier striking vehicles on fatalities in struck vehicles is statistically significant and robust to the inclusion of an extensive set of vehicle,

30. The raw magnitude of the coefficients is much larger in Table 8 than in Table 7 because the fatality rate in single-vehicle collisions is approximately 7 times higher than the fatality rate in two-vehicle collisions. This occurs because observed single-vehicle collisions tend to be more severe; drivers have no incentive to report minor single-vehicle collisions to their insurers or the police.

31. The sample size increases in column (3) relative to column (1) because (3) includes county fixed effects while (1) includes city fixed effects (which are missing for some observations). The probit estimator does not reliably converge with city fixed effects due to the large number of incidental parameters. However, the inclusion of city versus county fixed effects has little impact on the linear probability model estimates.

driver and accident covariates, estimation methods and identification strategies. Our estimates also scale to the national level.³²

7.1. *External costs and corrective taxes*

We now explore whether the estimated causal effect of vehicle weight on fatalities is economically significant. Consistent with the model from Section 3, we calculate the traffic fatality-related external cost of adding curb weight to a vehicle weighing w_i pounds and aggregate this external cost across all vehicles on the road. In any given collision this cost is proportional to $h'(w_i - w_j) = -\alpha \cdot f'(w_j - w_i)$. This quantity varies depending on w_j , the weight of the vehicle that i collides with. At each striking vehicle weight w_i , the slope of the dotted line in Figure 1 represents the average value of $h'(w_i - w_j)$. This average, $\bar{h}'(w_i)$, is computed across all accidents involving vehicles of weight w_i . The dotted line in Figure 1 plots $\bar{h}'(w_i)$. Figure 1 indicates that $\bar{h}'(w_i)$ is linear over most striking vehicle weights, but it diverges somewhat for vehicles weighing less than 2400 pounds. Therefore $\bar{h}'(w_i)$ for vehicles below 2400 pounds is less than $\bar{h}'(w_i)$ for vehicles weighing over 2400 pounds. In our calculations we use a piecewise linear function for $\bar{h}'(w_i)$, with a slope change at 2400 pounds. This piecewise linear function fits the dotted line in Figure 1 very well, generating an R^2 of 0.997. Note that using a linear $\bar{h}'(w_i)$ or the dashed line would result in a modestly higher external cost estimate (11% higher). In order to aggregate these external costs across vehicles, we need to make an assumption about what share of the vehicle fleet weighs less than 2400 pounds. However, this share is very low in recent years (only 2.5% of 2005 model year vehicles in our data weigh below 2400 pounds).

To calculate the total external cost of curb weight for a vehicle weighing w_i pounds, we compare its external cost, $\bar{h}(w_i)$, to the external cost of the lightest available vehicle, $\bar{h}(w_{cf})$. This calculation requires an assumption about w_{cf} , the weight of the smallest “counterfactual” vehicle that is available on the market. One could argue that any curb weight over zero pounds increases the probability of a fatality and that the appropriate value of w_{cf} is zero. However, a “zero pound” vehicle lies far outside the support of our data. We thus consider two more reasonable counterfactual vehicles below, and we assume that the only fatality-related externalities imposed by our smaller counterfactual vehicle (a car weighing 1850 pounds) involve pedestrians and motorcyclists. We experiment with a “zero pound” counterfactual in Table 9.

The average individual in our baseline scenario chooses a vehicle weighing the same as the average 2005 model year vehicle in our data (3616 pounds). We calculate the individual external costs relative to two counterfactual vehicles that the individual could have bought—a slightly lighter vehicle and the lightest possible vehicle. The slightly lighter counterfactual vehicle is a proxy for the average 1989 model year vehicle in our data, which weighs 2953 pounds. The lightest possible counterfactual vehicle is the smallest drivable car in mass production in 2005, which weighs 1850 pounds.

The external cost of an individual buying vehicle model i weighing w_i pounds over a lighter counterfactual vehicle weighing w_{cf} pounds is given by:

$$\text{External Cost}_i = [\bar{h}(w_i) - \bar{h}(w_{cf})] \cdot P(\text{accident}) \cdot VSL \quad (9)$$

32. In a previous working paper, we estimated the same models using data from the NHTSA General Estimates System (GES). The GES is a random subsample of police reported accidents in all states. It thus has fewer observations, but greater geographic coverage, than our merged state data sets. If we estimate our preferred specification using GES data, we find that 1,000 pounds of additional vehicle weight increases the probability of a fatality by 40% in the other vehicle. This estimate is statistically significant ($t = 4.8$) and similar in magnitude to our preferred estimate of 47% from the state data sets. We cannot reject the hypothesis that both estimates converge to the same value.

TABLE 9
Accident-related external costs under different assumptions

	Total external costs (billions)	External costs per gallon
Fatalities only	\$135.8	\$0.97
Fatalities + injuries	\$175.5	\$1.25
Fatalities + injuries + increased collision rates	\$235.9	\$1.69
Fatalities + injuries + increased collision rates + light trucks	\$264.6	\$1.89
Fatalities + injuries + increased collision rates + light trucks + 0 pound counterfactual	\$304.4	\$2.17

Notes: This table reports total accident-related external costs under a set of assumptions. For injuries, we apply the vehicle weight coefficient from Table 3 (column 7) and assume that the value of a statistical injury is 214,000 (as reported by the National Safety Council). We also assume that serious injury costs are not internalized through existing liability insurance. For collision rates, we assume that, *ceteris paribus*, an average weight vehicle (3616 lbs.) is 40% more likely to have an accident than an 1850 lb. baseline vehicle. This figure is an approximation of the observed relationships between accident rates and vehicle weight in Evans (1984) and White (2004). For the effects of light truck frames, we apply the light truck coefficient from Table 2 (column 7) and a light truck market share of 50% (Anderson, 2008). The zero pound counterfactual assumes that even an 1850 pound vehicle imposes external costs on other vehicles, at a rate of 0.00076 fatalities per collision (this rate comes from our data).

We employ the estimate of the causal effect of curb weight on the probability of a fatality in an accident as shown by the dotted line in Figure 1 (more precisely, a piecewise linear function approximating that line). For weights above 2400 pounds we apply the estimated 0.109 percentage points of risk for each additional 1000 lbs in striking weight. For weights below 2400 pounds we apply 0.058 percentage points. Both values are the averages of $\hat{h}'(w_i)$ over the respective curb weight ranges. We set w_i to 3616 pounds in all simulations, as discussed above. As the fatality externality is not a simple linear function of vehicle weight, we need to determine the share of vehicles in each range. We employ the shares observed in our accident data, which is 2.5% for 2005 model year vehicles and 15% for the 1989 model year vehicles. We calculate the probability of a vehicle being involved in a police-reported multivehicle collision at 3.65% per year (NHTSA, 2007),³³ and we apply the DOT value of a statistical life of \$5.8 million.³⁴

Critical in interpreting the expression in equation (9) as an external cost is the assumption that consumers recognize and value the safety-related internal benefits of weight, $f(w_i - w_j)$. If not, then they will not invest additional resources in buying heavier cars (except to gain the non-safety-related benefits), and no distortion of preferences will occur. However, we are confident that consumers do recognize and value these internal benefits. An abundance of VSL studies in the auto safety context demonstrate that consumers value automotive safety (De Blaeij *et al.*, 2003). Furthermore, a large-scale survey in Hellinga *et al.* (2007) reveals that 84% of parents believe a midsize to large vehicle is safer than a small one, with 93% of these specifically citing the larger size as the feature that enhances safety.³⁵

33. We estimate the probability of being involved in an accident by dividing the total number of vehicles involved in reported multivehicle collisions by the total number of registered vehicles in 2005 (US DOT BTS, 2010 Table 1–11).

34. This value is consistent with an extensive body of research on consumer valuation of auto safety. For example, de Blaeij *et al.* (2003) review a large number of studies analysing willingness to pay for automotive safety. These include seven studies conducted in the United States post-1980. Across these seven studies the median (mean) VSL is \$5.5 million (\$8.8 million) in 2008 dollars.

35. Even among the 16% of parents that believe a smaller vehicle is safer, 88% cited the easier handling of the smaller vehicle as the reason it is safer, suggesting that they too understand that small size itself is not helpful in surviving a crash except insofar as it allows one to avoid the crash to begin with.

In the first counterfactual scenario our simulated individual chooses an average vehicle weighing 3616 lbs (w_i) instead of one weighing 2953 lbs (w_{cf}). The total external costs aggregate to \$33.9 billion per year. One can think of this estimate as the external costs arising from the fleet weight gain since 1989.

Our second counterfactual scenario assumes the individual purchases an average vehicle weighing 3616 lb (w_i) instead of one weighing 1850 lbs (w_{cf}), which represents the lightest automobile in mass production that can transport two adult passengers. This is the approximate weight of Toyota's iQ, Mercedes Benz's Smart Car, or the first generation Honda Insight. The intuition behind calculating the total external cost using this baseline vehicle is that individuals privately choose the size of the externality by choosing a heavier vehicle than required to provide baseline transportation services. This calculation recognizes that a driver of a Smart Car poses minimal risk to other roadway users (except bicyclists or motorcyclists). The total external costs from this scenario sum to \$78.8 billion per year.

Both scenarios ignore the external fatality risks that vehicles pose to pedestrians and motorcyclists. In 2005, there were 2659 motorcycle crash fatalities involving light vehicles and 5864 non-motorist fatalities involving light vehicles (NHTSA, 2010). This is equivalent to an external "baseline" fatality cost of \$49.4 billion. Our interpretation of this as a "baseline" cost that applies across all vehicles assumes that the fatality risk to pedestrians and motorcyclists is independent of weight and that it can be eliminated by not driving. This appears to be true empirically.³⁶ The total external cost of "excess" vehicle weight (relative to the 1850 lb. counterfactual vehicle) and baseline fatality risk from collisions involving pedestrians and motorcyclists is \$128.2 billion ($\$78.8 + \$49.4 = \128.2).

The above calculations also ignore the impact of higher striking vehicle weight in multivehicle collisions with more than two vehicles. Almost all of these accidents involve three vehicles. We repeat the simulation above but add the external costs in three-vehicle collisions. We assume that striking vehicle weight has half the causal effect (per vehicle struck) in three-vehicle accidents as compared to its effect in two-vehicle collisions. This assumption is conservative in comparison to our three-vehicle collision estimates in Table 3. These calculations raise external costs in the "weight gain since 1989" scenario to \$37.2 billion and external costs in the "lightest possible vehicle" scenario to \$86.4 billion. Total external costs rise from \$128.2 to \$135.8 billion.³⁷ Notably, this figure exceeds central estimates of the social cost of US carbon emissions, which total \$123.7 billion per year.³⁸

The calculations above reveal the fatality-related externality for an average driver. However, collision rates vary across drivers, and a dangerous driver generates a larger externality than a safe driver, vehicle weight held constant. In the context of our model a dangerous driver has a higher value of α in the expression $h(w_i - w_j) = -\alpha \cdot f(w_j - w_i)$. From a policy perspective, it would be attractive to rescale existing liability insurance rates by vehicle weight, as existing insurance rates already account for driver heterogeneity and provide discounts to low-mileage drivers. Liability insurance would then be an increasing function of a driver's record, miles

36. Not all states reliably report pedestrian data, but we have a sample of 23,280 crashes involving pedestrians and 6,831 crashes involving motorcyclists. Regressing a pedestrian fatality indicator on vehicle weight (plus control variables, including a light truck indicator), we find that an additional 1,000 pounds of vehicle weight is associated with a statistically insignificant 6.5% increase in the probability of a fatality ($t = 0.9$). Regressing a motorcyclist fatality indicator on vehicle weight (plus control variables, including a light truck indicator), we find that an additional 1,000 pounds of vehicle weight is associated with a statistically insignificant 0.03% decrease in the probability of a fatality ($t = -0.0$).

37. A spreadsheet detailing these calculations is available from the authors.

38. In 2008 the US generated 5.89 billion metric tons of CO₂ (US EIA, 2009). Using a central estimate of the social cost of CO₂ at \$21 per metric ton, the annual social cost of US carbon emissions is \$123.7 billion (Greenstone *et al.*, 2011).

driven, and the weight of the vehicle that he chooses to buy, and the three factors could interact with each other. If this approach is not feasible, we discuss below two simpler ways to distribute the fatality-related external costs across vehicles. These mechanisms are less precise at distributing external costs in contexts with substantial driver heterogeneity in accident rates, as they do not allow for an interaction between a driver's record and his vehicle's weight. Note however that allowing for driver heterogeneity in the VSL has little impact on any of the policy instruments.³⁹

On the one hand, one could incorporate the fatality-related external costs as a per mile charge, in the spirit of "pay as you drive" (PAYD) insurance proposals. In contrast to existing proposals for PAYD insurance (*e.g.* Parry, 2005; Bordoff and Noel, 2008), our results demonstrate that the per mile charge should vary sharply by weight—a heavier car generates greater expected external costs per mile than a lighter car. An appropriately set per mile charge should be similar in effect to a weight-based excise tax on automobiles; the primary difference is whether the tax is collected at the time of sale or over the life of the vehicle.⁴⁰ However, to assess a charge that varies per pound and per mile, one needs accurate information on vehicle miles travelled (VMT) for each vehicle, which given today's monitoring technology creates significant but not insurmountable technical challenges.

A practical alternative is to distribute the total external costs by raising the gasoline tax assessed per gallon. Taxing gasoline is appealing because it is simple and because gasoline usage is positively related to both miles driven and vehicle weight. The United States consumed 140 billion gallons of gasoline in 2005 (US EIA, 2010). If we distribute the total external costs calculated above across 140 billion gallons of gasoline, this translates into 26 cents per gallon in the "weight gain since 1989" scenario (\$37.2 billion/140 billion gallons = 26 cents/gallon). The total externality due to vehicle fatalities when the baseline vehicle is 1850 pounds translates into a tax of 61 cents per gallon (\$86.4 billion/140 billion gallons = 61 cents/gallon). Including pedestrian and motorcycle fatalities translates into a tax of \$0.97 per gallon (\$135.8 billion/140 billion gallons = 97 cents/gallon).

While the per gallon tax does not differ by the weight of the vehicle, it results in a higher per mile charge for heavier vehicles as these have worse fuel economy. Figure 3 plots a Lowess smoother of miles per gallon (mpg) against vehicle weight, estimated for model year 2005 cars using data from Knittel (2011).⁴¹ There is a strong negative and slightly nonlinear relationship between the two variables. A linear regression indicates that an additional 1000 pounds in vehicle weight decreases fuel economy by 4.5 mpg. A gas tax thus results in heavier vehicles indirectly paying a higher per mile tax because they get fewer mpg. In this sense, the gas tax approximates a weight varying mileage charge.

A natural question is how close the gasoline tax comes to achieving the desired weight varying mileage charge. We perform a back of the envelope calculation using a large set of vehicles for which we have vehicle weight and fuel efficiency ratings (Knittel, 2011). We examine 8201 model-year combinations built from 1997 to 2006, which includes most cars and light trucks sold in the United States during this period.⁴²

39. For example, even in an extreme case in which drivers sort into vehicles such that a driver's VSL is proportional to her vehicle's weight, the appropriate value of a weight-based tax or a gas tax changes only 11%.

40. One difference between the two taxes is that under the excise tax owners would have an incentive to drive their vehicles a few more miles before retiring them, as this would amortize the tax over a larger mileage base.

41. For this comparison we require vehicle weight and EPA fuel economy ratings. The latter are not in our VIN decoder database, but Chris Knittel graciously shared his model level data on weight and fuel economy ratings.

42. For the analysis we remove boutique vehicles, which have zero market share (*e.g.* Ferrari, Bentley, etc.), flex fuel vehicles, which have inflated mpg ratings for accounting reasons, and a few miscoded observations.

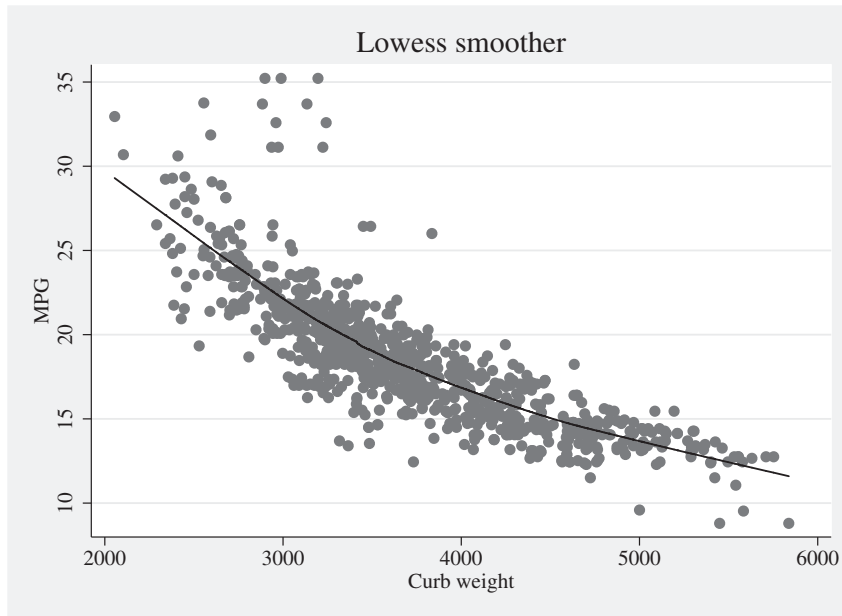


FIGURE 3
Fuel economy vs. weight for 2005 model year light vehicles

The per mile weight based charge for a given vehicle is equal to c_i^e , which is the per mile weight based external cost for vehicle type i and given by

$$c_i^e = \frac{[\bar{h}(w_i) - \bar{h}(w_{cf})] \cdot P(\text{accident}) \cdot \text{VSL}}{\text{VMT}_i} + \frac{c_{\text{ped-mot}}^e}{\text{VMT}_{\text{total}}} \quad (10)$$

where the numerator of the first term is the external cost given in equation (9). Now w_i is the curb weight for each 2005 model-year vehicle in Knittel's database, w_{cf} is the baseline vehicle's weight (in our case 1850 lbs), VSL is the value of a statistical life, and $P(\text{accident})$ is the probability of being involved in a multivehicle collision. VMT_i are set at 11,000 miles per year for each model. Note that VMT may fall after the tax is implemented.⁴³ In fact, this is desirable, as reducing VMT is one way to reduce the total externality. As long as prices are set correctly, consumers should be free to choose whichever vehicles and driving patterns they wish. The parameter $c_{\text{ped-mot}}^e$ is the total number of fatalities in collisions between vehicles and pedestrians, bicyclists and motorcycles, multiplied by the VSL; it sums to \$49.4 billion per year, or 1.65 cents per VMT.⁴⁴ We calculate c_i^e for each model in our database using the piecewise linear function discussed above. The average value of c_i^e across all models is 4.8 cents per mile.

43. VMT may fall less than expected if smaller vehicles require drivers to take more trips. Descriptively, however, we find a positive relationship between VMT and weight in the National Household Travel Survey.

44. As noted in footnote 36, there is no evidence that heavy cars pose greater risks to pedestrians and motorcyclists than light cars. The only vehicle characteristic we are aware of that influences pedestrian and motorcyclist fatality risk is whether a vehicle is a light truck. Anderson (2008) finds that a pedestrian's or motorcyclist's fatality risk when struck by a car is approximately 40% lower than his risk when struck by a light truck. In theory one could charge a lower value of $c_{\text{ped-mot}}^e$ to cars and a higher value to light trucks. In practice, however, this modification makes little difference when comparing a gasoline tax to the weight varying mileage charge. Without this modification, the average absolute difference between the weight-based tax and the gas tax is 0.88 cents per mile. With this modification, the average absolute difference between the weight-based tax and the gas tax becomes 1.00 cents per mile.

Denote the alternative instrument, the per gallon gas tax, as \bar{c}_e , which we calculate at \$0.97 above. We translate this per gallon tax into a per mile tax for vehicle i as follows:

$$c_i^g = \frac{\bar{c}_e}{mpg_i} \quad (11)$$

where c_i^g is the gasoline tax per mile for vehicle i , \bar{c}_e is the \$0.97 per gallon external cost and mpg_i is miles per gallon for vehicle i . Here we use the standard 45/55 weighting of the EPA city and highway fuel economy ratings.⁴⁵ The gas tax per mile therefore only varies across models through differences in fuel economy.

Figure 4 presents a scatterplot of the gas tax versus the weight tax for all models from 1997 to 2006 in the cleaned Knittel (2011) data. The difference between the two taxes is small for most models, but it can be significant at the extremes, ranging between -3.7 cents to 4.7 cents per mile. A one-cent difference per mile equates to \$110 on an annual basis. For 63% of the models in our data, the absolute value of the difference between the two taxes is less than one cent per mile, and for 96% of the models the absolute value of the difference is less than 2 cents per mile. The average difference between the two taxes is 0.88 cents per mile, which represents 18.3% of the average value of the per mile weight tax. In either case, the revenues could be redistributed to make the taxes revenue neutral.

Recall that our estimates of accident-related externalities are conservative along several dimensions. First, we assumed that the effect of weight on serious injuries is internalized by liability insurance. Second, we assumed that weight has no causal effect on the probability of a collision. Third, we ignored the external risk that light truck frames appear to pose independent of weight. Finally, we assumed that the lightest production vehicle (1850 pounds) poses no risk to other vehicles. Table 9 presents estimates of the total external costs after including these factors. Accounting for potential injury costs increases the total externality from \$135.8 billion to \$175.5 billion.⁴⁶ Accounting for the possibility that weight may have a causal effect on collision rates increases the externality another \$60.4 billion to \$235.9 billion. Accounting for the additional risk from light truck frames increases the total externality to \$264.4 billion. Allowing for a “zero pound” counterfactual vehicle increases our estimate of the total externality to \$304.4 billion.⁴⁷ This is equivalent to a \$2.17 per gallon gas tax.

It is worth noting factors held constant in our policy analysis. Chief among these is the accident rate. Even if weight does not affect the accident rate, it may change over time based on congestion

45. Pre-2008 EPA fuel economy ratings are widely recognized to overstate the actual mileage achieved by the average driver (Edmunds, 2006). This affects our subsequent analysis because the \$0.97 gas tax was derived from actual fuel economy rather than the EPA’s forecast fuel economy. We thus rescale the EPA ratings so that the average fuel economy in this sample matches the average fuel economy observed nationwide (17.8 mpg), after adjusting for weight differences between the two samples. The rescaling factor that achieves this equivalence is 0.73. Our conclusions in the subsequent analysis are unchanged if we instead leave the EPA ratings untouched and recalculate the gas tax using EPA mileage ratings—in both cases the per mile gas tax closely tracks the weight based mileage tax.

46. Even injuries within a driver’s own vehicle may represent external costs if medical treatment is paid for by government or group insurance. Accounting for this potential positive externality (since own vehicle weight is protective) would decrease total external costs related to fatalities and injuries by 2.6% (Parry, 2004). Other potential injury-related externalities include emergency response costs and traffic jams at accident sites. Our calculations imply that accounting for these externalities would change the total externality by less than 1%.

47. In 62,057 collisions in which the striking vehicle weighed less than 2,000 lbs, there was a fatality in the struck vehicle in only 47 cases (0.076% of cases). This is 65% lower than the average fatality rate. If we assume that a zero pound car imposes no fatality externality, then an 1,850 pound car would impose a 0.076 percentage point fatality externality (relative to a zero pound car). Including this risk in our calculations would increase the total fatality-related externality by \$39.8 billion.




FIGURE 4

Sunflower scatterplot of gas tax vs. weight tax for cars and trucks

Notes: The graph above displays the joint distribution of the weight tax and gas tax per mile for the sample of cars and trucks with model years 1997–2006 from the database provided by Knittel (2011). We remove boutique cars, flex fuel vehicles, and a few outliers with incorrectly recorded fuel ratings. The sunflower plot bunches multiple observations into single flowers, where the number of petals indicates the total number of observations represented by the flower. The petals of light flowers represent one observation each and the petals of darker flowers represent 13 observations each.

levels, road design, driver behaviour, or other influences. A policymaker implementing a weight-based tax would thus do well to recalibrate the tax periodically to reflect changes in accident rates or overall fatality rates.

The imposition of a high gasoline tax or a weight-based mileage tax should ultimately change the weight distribution of the vehicle fleet. Our regressions estimate the external effects of vehicle weight given the current weight distribution of the vehicle fleet, but these effects could change as the fleet downsizes (Result 5). To simulate how this shift might affect our regression estimates, we consider three downsizing scenarios. In the first scenario, all vehicles reduce their weight by 20%. This results in larger absolute weight reductions for heavier vehicles. In the second scenario, vehicles above the average weight reduce their weight by 10%, and vehicles below the average weight reduce their weight by 20%. In the third scenario, vehicles above the average weight remain unchanged, and vehicles below the average weight reduce their weight by 20%. In all three scenarios $\bar{h}(w_i)$, the function relating striking vehicle weight and struck vehicle fatalities, remains approximately linear and very close to its original values. This suggests that the appropriate marginal tax on vehicle weight would not change as the vehicle fleet became lighter. It is also possible that in the long run the relationship between weight and gasoline usage may break down, particularly if hybrid and electric vehicles become a significant fraction of the vehicle fleet. In that case the correlation between a gas tax and a weight-based tax would diminish, and a tax that varies directly with weight could become preferable.

While many countries encourage fuel efficiency through high gasoline taxes, the United States encourages fuel efficiency through CAFE standards. In principle, fuel economy standards could

achieve the same downsizing of the vehicle fleet as a gasoline tax—a properly specified fuel economy standard should act as a de facto tax on heavier vehicles. A primary difference between the two instruments is that the fuel economy standard “tax” would be collected when purchasing the vehicle and would be amortized over the vehicle’s lifetime VMT, while the gas tax would be collected in small increments throughout the life of the vehicle. This difference could be important if consumers exhibit high discount rates or if salience is important (Finkelstein, 2009). Furthermore, the gas tax directly incentivizes a reduction in travel, while CAFE standards have an ambiguous effect on total VMT (on the one hand they increase the cost of purchasing a vehicle, but on the other hand they reduce operating costs). The gas tax is thus more likely than CAFE standards to reduce other VMT-related externalities (*e.g.* congestion). Calculating the exact fuel economy standards that achieve equivalent weight distribution effects to a \$0.97 per gallon gas tax is beyond the scope of this article, as it requires a variety of supply and demand elasticities. Nevertheless, we note two important points in the context of CAFE standards.

First, current CAFE standards are insufficient to internalize the externality presented in this article. Goldberg (1998) estimates that CAFE increases the price of pickup trucks by 0.6% and reduces the price of subcompacts by 0.5%. This equates to a tax on pickup trucks (relative to subcompacts) of approximately \$200. The gasoline tax discussed above, however, equates to a tax on pickup trucks (relative to subcompacts) of over \$4000 over the life of the vehicle. Second, the light truck coefficient in Table 2 suggests that removing the historical split in CAFE standards between cars and light trucks would improve welfare. The results in Table 2 imply that light truck frames impose significant external risks upon other roadway users but provide little or no safety benefit to their own occupants. This suggests that light truck purchases should be discouraged, but historical CAFE standards encourage light truck production by imposing a much lower mileage standard on trucks than on cars.

7.2. *Fleet weight and total fatalities*

Our results demonstrate that a tax on vehicle weight, or a high gasoline tax, could be welfare enhancing. In some cases, however, policymakers may focus only on a policy’s effect on total traffic fatalities. Considering the net effect of weight on both external and “internal” fatalities links our results back to the existing safety literature.

Our model implies that total traffic fatalities depend on the distribution of vehicle weight (Result 3). Estimating $h(w_i - w_j) = -\alpha \cdot f(w_j - w_i)$ flexibly over the support of weight differences allows us to simulate total traffic fatalities for different vehicle weight distributions. Figure 2 presents estimates of $h(w_i - w_j)$. There is an inherent symmetry in the distribution of $w_i - w_j$ because a $+X$ pound weight difference for one vehicle in a collision implies a $-X$ pound weight difference for the other vehicle. The total expected fatalities in any given collision can thus be calculated by adding $h(w_i - w_j)$ and $h(w_j - w_i)$. The convexity in Figure 2 indicates that collisions between vehicles of different weights are on net deadlier than collisions between vehicles of similar weights; the decreased risk of moving from -500 to -1000 pounds, for example, is smaller than the increased risk of moving from $+500$ to $+1000$ pounds. We apply the relationship in Figure 2 to simulate changes in total fatalities for three downsizing scenarios.

In our first scenario, all vehicles reduce their weight by 20% (600 pounds on average).⁴⁸ This results in larger absolute weight reductions for heavier vehicles, reducing the variance of the weight distribution. Total traffic fatalities fall by 0.8% (367 fatalities). In our second scenario,

48. To simulate each scenario, we subtract the specified weight reduction from each vehicle in our collision data set. We then tabulate the counterfactual distribution of weight differences ($w_j - w_i$), and use this distribution to calculate the average value of $h(w_j - w_i)$ over the set of counterfactual collisions.

vehicles above the average weight reduce their weight by 10%, and vehicles below the average weight reduce their weight by 20%. This slightly increases the variance of the weight distribution, and traffic fatalities increase by 0.3% (134 fatalities). In our third scenario, vehicles above the average weight remain unchanged, and vehicles below the average weight reduce their weight by 20%. This significantly increases the variance of the weight distribution, and traffic fatalities increase by 1.6% (690 fatalities). However, this scenario is somewhat extreme and is inconsistent with recent evidence in Busse *et al.* (2013).⁴⁹ In summary, it takes a large and highly uneven downsizing of the vehicle fleet to increase total traffic fatalities by 1% or more, and a uniformly proportionate downsizing of the fleet slightly decreases traffic fatalities. Higher gasoline prices are thus unlikely to significantly increase traffic fatalities.

Our estimates of $h(w_i - w_j)$ illuminate the mechanisms underlying other results in the traffic safety literature. The finding that changes in fleet weight have modest impacts on fatalities is consistent with the literature discussed in Section 2; our estimates of $h(w_i - w_j)$ and our simulations reveal it occurs because changes in fleet weight do not significantly increase fatalities unless they dramatically increase fleet heterogeneity. Our results are also consistent with recent work by Jacobsen (2013). Jacobsen finds that a 1 mpg increase in current CAFE standards increases traffic fatalities by 149 per year, but that tighter standards do not increase fatalities as long as they are “footprint based” or unified across cars and trucks. Our results clarify the mechanisms underlying these results. Current CAFE standards are detrimental to safety because they encourage light truck ownership, and light trucks are unambiguously dangerous. A unified or footprint-based standard, however, removes the preference for light trucks and encourages a more uniform lightening of the fleet across different size vehicles. As our simulations reveal, a uniform lightening of the fleet need not affect fatalities because the distribution of weight differences in crashes remains unchanged.

8. CONCLUSION

The US vehicle fleet has become significantly heavier over the past two decades. The average car on the road in 2008 was roughly 530 pounds heavier than the average car on the road in 1988, representing a 20% increase. This trend and its potential traffic safety implications have been widely discussed by policymakers when contemplating more stringent fuel economy standards or greenhouse gas emissions standards. However, it is less widely recognized that an unregulated vehicle fleet is inefficiently heavy due to the “arms race” nature of vehicle choice. In this article, we estimate the external effects of choosing a heavier vehicle on fatalities in two-vehicle collisions. We present robust evidence that increasing striking vehicle weight by 1000 pounds increases the probability of a fatality in the struck vehicle by 40–50%. This finding is unchanged across different specifications, estimation methods, and different subsets of the sample. We show that there are also significant impacts on serious injuries.

The external costs of fatalities are currently not internalized in the form of a first- or second-best policy. We calculate that a simple gasoline tax that internalizes the fleet weight gain since 1989 is \$0.26 per gallon. We further calculate that internalizing the total cost of external fatalities and injuries due to vehicle weight and operation, including crashes with motorcycles and pedestrians,

49. Busse *et al.* (2013) estimate the reduced form effect of an increase in gasoline prices on short-run equilibrium prices and quantities of cars of different fuel economies. They find that the market share of vehicles in the lowest mpg quartile (*i.e.* the heaviest vehicles) decreases by 27% with a \$1 increase in gas prices. In comparison, the market shares of the second and third mpg quartiles change by -7% and 0% , respectively. These results suggest that buyers of heavy vehicles do respond to increases in gasoline prices, making it unlikely that a gasoline or weight tax would have no effect on purchases of heavy vehicles.

requires a tax of at least \$0.97 per gallon, and as much as \$2.17 per gallon. Parry and Small (2005), applying a lower VSL to monetize other external costs and not accounting for the vehicle weight externality, calculate an optimal value of \$1.01 per gallon for the US gas tax (approximately \$0.60 above its current level) and \$1.34 per gallon for the UK gas tax (approximately \$2 below its current level). Internalizing the vehicle weight externality could increase this optimal value by approximately \$2, implying that European gas taxes may be much closer to optimal levels than the US gas tax.

The primary social costs of the vehicle weight “arms race” accrue in the form of higher operating costs rather than changes in total fatalities. While our calculation of external fatality costs provides the information that a policymaker needs for setting the correct prices, a calculation of the potential cost savings is also useful. We find that the 2005 model year fleet consumes \$92.8 billion more gasoline annually than the lightest possible fleet, which exceeds the comparable external fatality cost of \$86.4 billion.⁵⁰ This is consistent with our model and empirical estimates, which imply that consumers are willing to pay approximately \$86 billion for the internal safety benefits of added weight (recall that internal safety benefits and external safety costs are approximately equal), plus an additional amount for the other private benefits of weight. The \$92.8 billion figure should not be interpreted as a welfare calculation, as we do not know the valuation of the other private benefits or the costs of manufacturing heavier cars.⁵¹ Nevertheless, it indicates the potential magnitudes at stake.

Acknowledgments. We gratefully acknowledge support from the Robert Wood Johnson Foundation and the University of California Energy Institute. We thank Larry Goulder, Michael Greenstone, Ryan Kellogg, Ian Parry, Ken Small, and Florian Zettelmeyer for valuable feedback. Seminar participants at CESifo München, Duke University, UC Berkeley, the UC Energy Institute, UC Irvine, University of Michigan, Cornell University, Tinbergen Institute, Erasmus University, RAND, NBER, and the Occasional Workshop on Environmental and Resource Economics have provided helpful comments. Grace Chan, Erick Gong, Tammie Vu, and Angela Yu provided excellent research assistance. Any errors in the paper are the authors’.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- AHMAD, S. and GREENE, D.L. (2005), “Effect of Fuel Economy on Automobile Safety: A Reexamination”, *Journal of the Transportation Research Board*, (1941), 1–7.
- ANDERSON, M. (2008), “Safety for Whom? The Effects of Light Trucks on Traffic Fatalities”, *Journal of Health Economics*, **27**, 973–989.
- AUSTIN, D. and DINAN, T. (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, A.M., GOULDER, L.H., JACOBSEN, M.R. and VON HAEFEN, R.H. (2009), “Distributional and Efficiency Impacts of Increased US Gasoline Taxes”, *American Economic Review*, **99**, 667–699.
- DE BLAEIJ, A., FLORAX, R.J.G.M., RIETVELD, P. and VERHOEF, E. (2003), “The Value of Statistical Life in Road Safety: a Meta-Analysis”, *Accident Analysis and Prevention*, **35**, 973–986.

50. The savings from reduced gasoline consumption depend upon four parameters: The effect of weight on mpg, the number of passenger vehicles in 2005, the gasoline price in 2005, and the number of miles driven per vehicle. In Section 7 we find that that a 1,000 pound lighter vehicle achieves 4.53 additional mpg. The US DOT (2010) estimates that there were 232 million light duty vehicles registered in 2005. The EIA reports an average 2005 gasoline price of \$2.31 per gallon, and we assume that each vehicle drives 11,000 miles per year.

51. A regression of MSRP on weight finds a clear positive correlation. However, the relationship is only weakly monotonic. A key challenge in interpreting the relationship is that one can make vehicles lighter by making them smaller, which would create savings, or by using different materials (*e.g.* high-strength aluminum or carbon composites), which could actually make them more expensive. We thus emphasize the operating costs of heavier vehicles—which are unambiguously higher—rather than the manufacturing costs.

- BORDOFF, J. and NOEL, P. (2008), "Pay-As-You-Drive Auto Insurance: A Simple Way to Reduce Driving-Related Harms and Increase Equity" (Discussion Paper No. 2008–09, The Hamilton Project).
- BUSSE, M.R., KNITTEL, C.R. and ZETTELMEYER, F. (2013), "Are Consumers Myopic? Evidence from New and Used Car Purchases", *American Economic Review*, **103**, 220–256.
- CALIFORNIA AIR RESOURCES BOARD. (2004), *Estimation of Average Lifetime Vehicle Miles of Travel* (Sacramento: California Air Resources Board).
- CRANDALL, R.W. and GRAHAM, J.D. (1989), "The Effect of Fuel Economy Standards on Automobile Safety", *The Journal of Law and Economics*, **32**, 97–118.
- EDLIN, A.S. and KARACA-MANDIC, P. (2006), "The Accident Externality from Driving", *Journal of Political Economy*, **114**, 931–955.
- EDMUNDS, D. (2006), "EPA Overhauls Fuel Economy Estimates for 2008". Available at: <http://www.edmunds.com/fuel-economy/epa-overhauls-fuel-economy-estimates-for-2008.html>.
- ELVIK, R., CHRISTENSEN, P. and AMUNDSEN, A. (2004), *Speed and Road Accidents: an Evaluation of the Power Model* (Oslo: Institute of Transportation Economics).
- EVANS, L. (1984), "Accident Involvement Rate and Car Size", *Accident Analysis and Prevention*, **16**, 387–406.
- FERNÁNDEZ-VAL, I. (2009), "Fixed Effects Estimation of Structural Parameters and Marginal Effects in Panel Probit Models", *Journal of Econometrics*, **150**, 71–85.
- FINKELSTEIN, A. (2009), "E-ZTAX: Tax Salience and Tax Rates", *Quarterly Journal of Economics*, **124**, 969–1010.
- GAYER, T. (2004), "The Fatality Risks of Sport-Utility Vehicles, Vans, and Pickups Relative to Cars", *Journal of Risk and Uncertainty*, **28**, 103–133.
- GOLDBERG, P.K. (1998), "The Effects of the Corporate Average Fuel Efficiency Standards in the US", *The Journal of Industrial Economics*, **46**, 1–33.
- GREENSTONE, M., KOPITS, E. and WOLVERTON, M. (2011), "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation" (Working Paper No. 16913, National Bureau of Economic Research).
- HAKIM, D. (2003), "S.U.V.'s to Be Redesigned to Reduce Risk to Cars", *The New York Times*. Available at: <http://www.nytimes.com/2003/12/04/business/suv-s-to-be-redesigned-to-reduce-risk-to-cars.html?pagewanted=all>.
- HELLINGA, L.A., MCCARTT, A.T. and HAIRE, E.R. (2007), "Choice of Teenagers' Vehicles and Views on Vehicle Safety: Survey of Parents of Novice Teenage Drivers", *Journal of Safety Research*, **38**(6), 707–713.
- INSURANCE INFORMATION INSTITUTE. (2010), "Compulsory Auto/Uninsured Motorists". Available at: <http://www.iii.org/IU/Compulsory-Auto-Uninsured-Motorists/>.
- JACOBSEN, M. (2013), "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior", *American Economic Journal: Applied Economics*, **5**, 1–26.
- JOKSCH, H.C. (1993), "Velocity Change and Fatality Risk in a Crash – A Rule of Thumb", *Accident Analysis and Prevention*, **25**.
- KAHANE, C. (2003), "Vehicle Weight, Fatality Risk and Crash Compatibility of Model Year 1991–99 Passenger Cars and Light Trucks" (Technical Report No. DOT HS 809 662, National Highway Traffic Safety Administration).
- KENNICKELL, A.B. (2003), "A Rolling Tide: Changes in the Distribution of Wealth in the US" (Working Paper, Federal Reserve Board of Governors).
- KHAZZOOM, J.D. (1994), "Fuel Efficiency and Automobile Safety: Single-Vehicle Highway Fatalities for Passenger Cars", *The Energy Journal*, **15**, 49–102.
- KLIER, T. and LINN, J. (2012), "New-vehicle Characteristics and the Cost of the Corporate Average Fuel Economy standard", *The RAND Journal of Economics*, **43**, 186–213.
- KNITTEL, C.R. (2011), "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector", *American Economic Review*, **101**, 3368–3399.
- LI, S. (2012), "Traffic Safety and Vehicle Choice: Quantifying the Effects of the 'Arms Race' on American Roads", *Journal of Applied Econometrics*, **27**, 34–62.
- LI, S., TIMMINS, C. and VON HAEFEN, R.H. (2009), "How do Gasoline Prices Affect Fleet Fuel Economy?", *American Economic Journal: Economic Policy*, **1**, 113–137.
- MERTZ, H.J., PRASAD, P. and IRWIN, A.L. (1997), "Injury Risk Curves for Children and Adults in Frontal and Rear Collisions", (SAE 973318). In *Proceedings of the 41st Stapp Car Crash Conference*. 13–30.
- NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION. (2010), FARS Encyclopedia. Available at: <http://www.fars.nhtsa.dot.gov/Main/index.aspx>.
- NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION. (2007), "Traffic Safety Facts 2005".
- NATIONAL SAFETY COUNCIL. (2010), "Estimating the Costs of Unintentional Injuries". Available at: http://www.nsc.org/news_resources/injury_and_death_statistics/Pages/EstimatingtheCostsofUnintentionalinjuries.aspx.
- NOLAND, R.B. (2004), "Motor Vehicle Fuel Efficiency and Traffic Fatalities", *The Energy Journal*, **25**, 1–22.
- OGDEN, C.L., FRYAR, C.D., CARROLL, M.D. and FLEGAL, K.M. (2004), "Mean Body Weight, Height, and Body Mass Index, United States 1960–2002", *Advance Data*, **347**, 1–17.
- PARRY, I.W.H. (2004), "Comparing Alternative Policies to Reduce Traffic Accidents", *Journal of Urban Economics*, **56**, 346–368.
- PARRY, I.W.H. (2005), "Is Pay-as-You-Drive Insurance a Better Way to Reduce Gasoline than Gasoline Taxes?", *The American Economic Review*, **95**, 288–293.
- PARRY, I.W.H. and SMALL, K.A. (2005), "Does Britain or the United States Have the Right Gasoline Tax?", *The American Economic Review*, **95**, 1276–1289.

- PELTZMAN, S. (1975), "The Effects of Automobile Safety Regulation", *The Journal of Political Economy*, **83**, 677–725.
- PORTNEY, P.R., PARRY, I.W.H., GRUENSPECHT, H.K. and HARRINGTON, W. (2003), "Policy Watch: The Economics of Fuel Economy Standards", *The Journal of Economic Perspectives*, **17**, 203–217.
- ROBERTSON, L.S. (1991), "How to Save Fuel and Reduce Injuries in Automobiles", *The Journal of Trauma*, **31**, p.107.
- US ENVIRONMENTAL PROTECTION AGENCY. (2009), "Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 Through 2009" (Report No. EPA420-R-09-014, US Environmental Protection Agency).
- US DEPARTMENT OF TRANSPORTATION. (2011), *Highway Statistics 2010* (Washington, DC: US Department of Transportation).
- US ENERGY INFORMATION ADMINISTRATION. (2010), "Use of Gasoline - Energy Explained, Your Guide To Understanding Energy". Available at: http://tonto.eia.doe.gov/energyexplained/index.cfm?page=gasoline_use.
- VAN AUKEN, R.M. and ZELLNER, J.W. (2005), "An Assessment of the Effects of Vehicle Weight and Size on Fatality Risk in 1985 to 1998 Model Year Passenger Cars and 1985 to 1997 Model Year Light Trucks and Vans", *SAE Transactions*, **114**, 1607–1622.
- WARD'S REPORTS INCORPORATED. (1994), *Ward's Automotive Yearbook* (Southfield: Ward's Communications).
- WHITE, M.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.