Online appendix for "Fuel Economy and Safety:
The Influences of Vehicle Class and Driver Behavior"
Mark Jacobsen

A. Negative Binomial Specification

Begin by stacking the model in (7) and (8) to write the estimation as a single equation. Combining the count data and dividing out the HIC score, define the vector:

$$\mathbf{f} \equiv \begin{bmatrix} F_{i0st}/h_i \\ F_{ijst} \end{bmatrix}$$

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

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

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

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

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

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

$$E(f_l \mid \mathbf{x}_l) = \exp(\mathbf{x}_l'\gamma)$$
$$\text{with } \operatorname{Var}(f_l \mid \mathbf{x}_l) = \exp(\mathbf{x}_l'\gamma) \qquad \text{(A.1)}$$

When applying the definition of the indicators in $\mathbf{x}_l$ and taking logs (A.1) is equivalent to (7) and (8) in the main text.

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

$$E(f_l \mid \mathbf{x}_l, \varepsilon_l) = \exp(\mathbf{x}_l'\gamma + \varepsilon_l) \qquad \text{(A.2)}$$

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

$$E(f_l \mid \mathbf{x}_l) = \exp(\mathbf{x}_l'\gamma)$$
$$\operatorname{Var}(f_l \mid \mathbf{x}_l) = \exp(\mathbf{x}_l'\gamma) \cdot \left(1 + \theta \exp(\mathbf{x}_l'\gamma)\right) \qquad \text{(A.3)}$$

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

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

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

B. Alternative Definitions of the Bin Structure

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

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

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

C. Accident Fault Data

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

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

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

D. Influence of Peltzman Effects in the Simulation

In order to consider the potential size of Peltzman-type effects in the unobserved portion of driving behavior I further decompose α_i into two pieces. The first portion is the part of α_i predicted by the own-safety of the vehicle. (Constructed using least squares regression of α_i on own-safety, calculated as the average of β_{ij} in a row.) In that sense it is an upper limit on the size of the Peltzman effect. The second portion is whatever idiosyncratic variation remains in α_i and I will assume that continues to move with the driver. Table A4 presents the results of the policy experiment where α_i changes with the own-safety of the vehicle. The third column is my upper bound on the Peltzman effect over all driving safety residuals. The fourth column controls first for census region (there are more light trucks and dangerous roads in the west) and then applies the same method to divide the residual into two pieces.

As expected, the outcomes in Table A4 show that all fuel economy standards are improved if smaller vehicles indeed cause safer driving. However, even at the limit defined above I show that the existing fuel economy standard continues to have adverse effects on safety. Controlling for census region seems reasonable (as driver residence is unlikely to change with fuel economy standards), and the result becomes even closer to my central case. The safety improvement that can be offered by unifying the standard appears fully robust to the case where I allow a Peltzman-type effect. This is shown in the final row of the table. Because the difference in policies is maintained and overall safety is improved, we see that the unified standard even begins to offer substantial improvements in overall safety in the final column of the table.

E. Alternative Demand Elasticities

Table A5 presents the demand elasticities used in the simulation model, taken from Bento et al (2009). These control the pattern of substitution across vehicles when the shadow prices implied by the fuel economy rules are imposed. The elasticities come originally from a Bayesian model of vehicle choice, making it possible to draw a variety of alternatives from the estimated posterior densities to further explore the sensitivity of my results. Table A6 displays the minimum and maximum change in fatalities for each policy simulation, taken over 50 different draws on the elasticities. The extremes remain quite close to the central case, likely reflecting the large dataset and relatively high precision in the source paper.

Structural uncertainty in the elasticity estimates may be a larger source of variation, attested to by the wide range of estimates produced in the literature. The paper therefore also presents simulation results based instead on elasticities from Kleit (2004). The simulation results appear in Table 9 of the main text and the elasticities themselves are shown in Table A7 below.

F. Vehicle Class Aggregates

Table A8 below provides details on the mean and standard deviation of the fuel economy and weight of each class. Fuel economies across classes range from 19 to 30 MPG, capturing much of the variation among popular vehicles; the aggregation, however, means I do not capture the extremes as well. Hybrid compacts on the high end, for example, or luxury SUV's on the low end will be missed, though demand for these vehicles is also relatively inelastic meaning they play less of a role in the compositional changes expected under regulation.

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

G. Comparison with a Gasoline Tax

Increases in the gasoline tax, while typically eschewed in the U.S. for political reasons, provide an efficient benchmark for reduction of gasoline use. Gasoline taxes tend to compare very favorably to CAFE rules in the existing literature and my estimates on safety impacts further that argument.

To examine safety impacts I divide the effect of the tax into two components: a change in composition and an overall reduction in miles driven. The change in composition is similar to the unified standard presented in the main text: the shadow tax proportional to fuel economy becomes an actual tax proportional to fuel economy. I show that this type of movement in fleet composition produces only a small change in fatalities. The remaining reduction in gasoline use from the tax would come from miles driven.

Consider for example a case where the tax achieves half of gasoline savings through fleet composition and half through a reduction in miles driven. The 1.0 MPG improvement considered in the main policy examples reduces gasoline use by about 3.8%, so half of that is a 1.9% reduction in miles driven from the equivalent gasoline tax. All else equal, this reduction in miles will create a reduction of about 500 fatalities per year, representing a dramatic improvement relative to any version of CAFE I consider in the main text.

Table A1: Comparison of Poisson and Negative Binomial Estimation Results

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

Table A2: Effects of Alternative Bin Structures

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

Table A3: Fault by Class[a]

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

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

Table A4: Peltzman Effects and the Influence of a Driver-Vehicle Specific Residual[a]

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

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

Table A5: Matrix of Own and Cross-Price Demand Elasticities by Class[a]

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

[a] These elasticities are derived from Bento et al (2009) and are used in the central case of the policy simulations.

Table A6: Simulation Results over a Sample of Varying Demand Elasticities[a]

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

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

Table A7: Alternative Demand Elasticities by Class[a]

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

[a] Elasticities from Kleit (2004) aggregated to match the ten class definitions in my model. In order to isolate the effects of fleet composition I also proportionally adjust the cross-price elasticities such that fleet size is exactly maintained. The results from using these elasticities appear in Table 9 of the main text.

Table A8: Fuel Economy and Weight by Class

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