Product design response to policy: Evaluating fuel economy standards using an engineering model of endogenous product design*

Kate S. Whitefoot†
University of Michigan

Meredith Fowlie‡
University of California, Berkeley
and NBER

Steven J. Skerlos
University of Michigan

December 2011

Abstract

Policies designed to improve industrial environmental performance are increasing in scope and stringency. These policies can significantly influence engineering design decisions as firms re-optimize their products and processes to meet compliance requirements at minimum cost. This paper demonstrates the importance of accounting for these design responses in the analysis of policy impacts. As a case in point, we model automotive firms' medium-run compliance choices under the reformed Corporate Average Fuel Economy (CAFE). Physics-based simulations are used to characterize the potential for improving fuel efficiency through design changes. These engineering simulation results are coupled with a partial-equilibrium, static oligopoly model in which firms choose prices and key vehicle design attributes. We simulate firms' pricing and medium-run design response to this regulation. Results indicate that firms rely primarily on changes to vehicle designs to meet the reformed CAFE standards, with a smaller contribution coming from pricing strategies designed to shift demand towards more fuel-efficient vehicles. The analysis also draws attention to two factors that could offset potential fuel efficiency improvements achieved under the reformed CAFE standards: an increase in the market share of light trucks and an increase in the market share of firms that choose to violate the standard.

* The authors would like to thank Severin Borenstein, Lawrence Goulder, Mark Jacobsen, Ryan Kellogg, Chris Knittel, James Sallee and seminar participants at the University of Michigan and the University of California, Berkeley for their helpful suggestions. We also thank Ashley Langer for valuable help with the demand estimation, Bart Frischknecht and Deokkyun Yoon for assistance with vehicle simulations, and Kevin Bolon for help managing the data sets. Financial support for this work is provided by the National Science Foundation and the University of Michigan Rackham Graduate School and is gratefully acknowledged.

† K. Whitefoot is now a Senior Program Officer at the National Academy of Engineering. Address:500 5th St, Washington, DC 20009. Email: kwhitefoot@nae.edu. Website: www.designscience.umich.edu/kwhitefoot.html.

‡ Address: 301 Giannini Hall, UC Berkeley, Berkeley CA 994720. Email: fowlie@berkeley.edu.

1 Introduction

In order to reduce greenhouse-gas emissions, local-air pollutants, and dependence on foreign energy sources, energy-efficiency standards and incentives are being established for many energy consuming durable goods. In 2007, Congress created efficiency standards for many household appliances, including dishwashers and furnaces. In 2010 the Department of Energy announced new efficiency standards for freezers, refrigerators, and clothes washers. One especially noteworthy policy intervention, enacted recently by Congress, raises fuel economy standards for new automobiles to at least 35 mpg by 2020.[1]

How firms respond to these types of policies can have significant implications for how efficiently mandated energy-intensity reductions are achieved and who bears the costs. In general, firms can comply with energy-efficiency standards through a combination of shifting production towards their more efficient products and modifying the designs of their products so that they are more energy efficient. This paper is particularly concerned with the latter "design response" to environmental regulation.

In much of the economics literature that investigates the response of a differentiated-products industry to a regulatory intervention, firms' ability to achieve regulatory compliance via product or process design changes is underemphasized or ignored (e.g., Goldberg 1998; Nevo 2000; Jacobsen 2010). The standard approach to modeling the response of firms to energy-efficiency standards allows firms to adjust the prices of their products but not product attributes or offerings. Recent work on the automotive industry indicates that changes in product designs have played a significant role in determining fleet fuel-efficiency trends, including gains achieved under the Corporate Average Fuel Economy (CAFE) standards (Knittel 2009; Klier and Linn 2008). This paper seeks to explicitly account for this design response into the modeling and analysis of energy efficiency standards.

The emerging literature on endogenous attribute selection in the context of differentiated products has carefully considered certain types of design choices, including decisions on where to locate a store, what broadcast frequency to choose for a radio station, and what content to cover in a newspaper (e.g., Seim 2006; Sweeting 2007; Fan 2008). However, the design

[1] This represents a more than 30% reduction in fuel consumption per mile.

decisions that apply to the technologically complex products that are typically targeted by energy-efficiency standards, such as automobiles and household appliances, present additional challenges. Design decisions for these types of products are subject to a suite of engineering constraints and tradeoffs. These constraints can play an important role in shaping firms' response to policy interventions.

In this paper, we develop an empirically tractable approach to modeling the design decisions of firms producing technologically complex products. To accomplish this, we draw from the engineering design literature, which examines the engineering constraints and tradeoffs that shape product design decisions. These constraints and tradeoffs have been explicitly represented in detailed engineering simulation models of product design (e.g., Frischknecht and Papalambros 2008; Gholap and Khan 2007; Wright et al. 2002). We demonstrate how these design models can be constructively integrated into economic models of strategic industry interactions.

We make use of a particularly rich engineering simulation model that is currently used by the automotive industry to support the powertrain development process. The model is too computationally intensive to be incorporated directly into our oligopolistic model of the automotive industry. Instead, we use the engineering model to simulate technologically feasible trade-offs between fuel efficiency, other vehicle performance attributes, and production costs. These vehicle design simulations are repeated many times using a range of feasible combinations of design parameters. The rich data generated by these simulations are used to estimate a flexible approximation to the vehicle design process. This more tractable representation of the vehicle design process can be nested within a standard oligopoly model with product differentiation. This facilitates a detailed analysis of both the price and the design response of automotive firms under the recently reformed CAFE standards.

Our approach builds upon recent work by Klier and Linn (2008) and Knittel (2009) who econometrically estimate the tradeoffs that automotive firms face between fuel economy, weight, and engine power. This paper investigates very similar design relationships, but in lieu of using bundles of attributes observed in the market place to econometrically estimate the vehicle production possibility frontiers, we use the outputs of physics-based engineering simulations.[2]

[2] .This is certainly not the first paper to make use of engineering estimates in a detailed economic analysis. For example, engineering estimates of costs have been used to benchmark electricity sector performance (see, for

For our purposes, this engineering-based approach confers two advantages. First, many combinations of product attributes are not observed in the data, but are technologically feasible and potentially optimal under counterfactual policy designs. Engineering simulation models are well suited to identify technologically possible combinations of attributes that have yet to manifest in existing vehicle designs. Second, correlations between unobserved attributes (such as luxury accessories) and attributes of interest (such as fuel economy) can make it difficult to identify design tradeoffs econometrically.[3] Physics-based engineering simulations can allow us to capture engineering tradeoffs independent of unobserved product attributes.

This paper also contributes to the literature that seeks to estimate demand parameters in the presence of correlations between observable and unobserved vehicle attributes. In previous work, researchers have used functions of non-price attributes of other vehicles (such as fuel efficiency or horsepower) as instruments (e.g., Berry et al. 1995; Train and Winston 2007). One criticism of this approach is that firms presumably choose some of these non-price attributes and prices simultaneously. We exploit the well-documented structure of the automotive design process to identify those vehicle attributes that are determined and fixed early in the design process, prior to the medium-run design decisions we analyze. These are the variables we use to instrument for vehicle attributes, including price and fuel efficiency, that are manipulable in the short- and medium-run. Our key identifying assumption is that powertrain architecture (e.g., hybrid), drive type (e.g., all-wheel-drive), and major vehicle dimensions are fixed prior to the powertrain tuning and selection of technology features that affect both fuel efficiency and acceleration performance. We present evidence in support of this assumption in Section 2.2.

In the second part of the paper, we use the modeling framework described above to analyze the impacts of the 2014 reformed CAFE standards on vehicle design decisions, producer profits, and consumer surplus. Notably, most automotive firms cannot meet the reformed standards through pricing adjustments alone. This highlights the importance of incorporating a model of firms' design responses into the analysis of this regulation. When the full complement of medium-run design responses are represented in our policy simulations, the sales-weighted

---

example, Wolfram (1999) and Borenstein (2002). Engineering models have also been used to simulate the response of electricity producers and automotive firms, respectively, to policies limiting NO_x emissions (Fowlie et al., 2012).

[3] Klier and Linn (2008) exploit an engine dataset to estimate tradeoffs between endogenous attributes using variation in observed attributes of vehicle models with the same engine program. One potential drawback of this approach is that unobserved attributes, such as luxury accessories that can further impact fuel economy, are often correlated with observed attributes such as horsepower and weight.

average fuel economy among compliant firms increases by 3.5 mpg under the reformed CAFE standards. A majority (88 percent) of these fuel economy improvements come from vehicle design changes; the remaining 12 percent derive from price changes. However, results also indicate that improvements in fuel economy achieved among complying firms are undermined by to two factors. First, firms have an incentive to adjust prices to increase the market share of light trucks because light trucks have lower fuel economy targets than passenger cars. Second, because the fines associated with violating the CAFE standards are relatively low, results suggest that the market share of European firms that typically violate the standards increases in response to the regulation, slightly offsetting improvements in fuel economy made by compliant firms

Vehicle fuel efficiency can be improved in the medium run, at the expense of acceleration performance, by tuning design parameters in the powertrain. This is often ignored in policy analyses of CAFE (e.g., Austin and Dinan 2005, NHTSA 2008). In a second set of policy simulations, we allow firms to improve vehicle fuel efficiency by implementing technology features, but shut down the ability to tradeoff acceleration performance and fuel efficiency. In these simulations, only 14 percent of gains in fuel economy are due to design changes. Importantly, estimated costs to producers of complying with the regulation are three times larger when we fail to account for tradeoffs between fuel economy and other vehicle attributes.

The remainder of the paper is organized as follows. Because it is important to account for the vehicle design process when constructing our model, we begin with a general overview of both the design process and our approach to modeling design decisions. Having laid down this foundation, Section 3 provides a more detailed discussion of how we integrate detailed, physics-based engineering simulations into our model of vehicle design. Section 4 describes the demand-side of the model. Section 5 describes the estimation results. Section 6 describes the reformed CAFE policy. Section 7 discusses the policy simulations. Section 8 concludes.

2 An overview of the vehicle design process

Modeling endogenous product design decisions for technologically complex products requires accurate representation of the engineering and economic tradeoffs associated with these decisions. We cannot directly observe all of the tradeoffs that firms make during the product development process. We can, however, generate detailed engineering estimates of the tradeoffs that play a significant role in determining the performance attributes of a product. We focus on two vehicle attributes that can be modified in the medium run: fuel efficiency (measured as gallons of fuel consumed per 100 miles) and acceleration (measured as the time in seconds to accelerate from 0-60 mph).[4]

Before delving into the details of the engineering model and the associated estimation, this section introduces the automotive development process in general terms. We then provide an overview of an automotive firm's medium-run design decisions. Finally, we contrast our approach with the econometric approaches taken elsewhere in the literature.

2.1 The Vehicle Development Process

Generally speaking, the automotive design process is a structured sequence of interrelated decisions, many of which constrain choices made at later stages (Braess 2005; Sörensen 2006; Weber 2009). We use this structure of the automotive design process to inform the development of our model.

The typical automotive design process begins with concept development, followed by a system-level design that defines the geometric layout of the vehicle (including target vehicle footprint), followed by a detailed design of all subsystems (Sörenson 2006; Weber 2009). For a newly designed vehicle model, the development process begins with targets for specific vehicle attributes, such as the vehicle segment (e.g., compact), powertrain architecture (e.g., hybrid), variations (e.g., four-door sedan), major dimensions, transmission types (e.g., automatic, torque classes) and engine versions (Braess 2005; Weber 2009). For a redesigned model, the development process begins with the determination of any changes to major properties of the vehicle and specifications for subsystems, such as how many drivetrain configurations or engine options will be available. In both new design and redesign contexts, there are certain earlier design decisions that must be finalized before the detailed engineering design of vehicle subsystems can begin (Braess 2005; Sörenson 2006; Weber 2009).

Figure 1 provides a stylized representation of this process. This figure is somewhat misleading insofar as it suggests that the design process proceeds in sequential, clearly defined stages. In fact, iteration loops and overlapping tasks may exist between the stages presented. This caveat notwithstanding, there is a point in the automotive development processes where vehicle

[4] Although there are some other performance attributes that can be modified in the medium run (such as vehicle handling), fuel consumption and acceleration are the most important in terms of determining the overall performance of (and thus demand for) the vehicle.

segment, powertrain architecture (e.g., conventional gasoline, hybrid, diesel), and major dimensions are finalized but changes to other aspects of the vehicle design (e.g., "tuning" the powertrain) are still possible. It is this latter stage of the vehicle design process, which involves what we term "medium-run" design decisions, that is the focus of our analysis.

For the purpose of our analysis, any design parameters that are determined in the earlier stages of the design process (Stage A in Figure 1) are assumed to be fixed. These include the segment of vehicle, key internal and external dimensions, and the powertrain architecture (e.g., conventional gasoline, hybrid, and diesel). Conditional on these features and attributes, we model manufacturers' choice of fuel economy and acceleration performance (Stage B) and vehicle pricing strategies (Stage C).[5] Importantly, we will use the variation in vehicle attributes determined in earlier stages of the design process (Stage A) to instrument for endogenous variables in our demand-side estimation such as fuel efficiency that are determined later in the design process (Stage B).

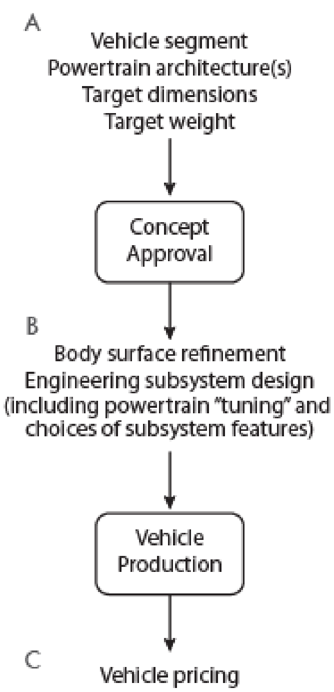

Figure 1: Simplified representation of an automotive development process illustrating short-run (Stage C), medium-run (Stage B) and longer-run (Stage A) design decisions

[5] Ideally, the supply side should be modeled as a two-stage game to represent the sequence of choosing product attributes before prices (or prices with smaller adjustments of product attributes). However, computational complexity prevents us from solving the second-stage using Newton-based methods and faster fixed-point methods accounting for the CAFE constraint are unknown.

2.2 Medium-run vehicle design decisions

At Stage B in the design process, the automotive manufacturer can adjust the fuel efficiency and acceleration performance by "tuning" design parameters in the powertrain (e.g., engine displacement and final drive ratio) and adding certain technology features (e.g., a high efficiency alternator).[6] Let j index a vehicle model and engine option (e.g., the non-hybrid Ford Escape). Let $\boldsymbol{x_j}$ denote the powertrain design parameters that are manipulable in the medium-run. Let $\boldsymbol{t_j}$ index the suite of technology features that can be added at Stage B to improve fuel efficiency and/or acceleration performance. Let $\widetilde{\boldsymbol{x}}_j$ represent the fixed design parameters determined earlier in the design process (Stage A).

Our first task is to distill the complex engineering relationships between design parameters and vehicle attributes into a manageable number of estimable equations that can be nested within a standard economic model of the automotive industry. To accomplish this, we impose some additional structure and some simplifying assumptions. To fix ideas, consider the two vehicle attributes of primary interest: fuel consumption (*fuelcons*$_j$) and acceleration performance (*acc*$_j$)[7]:

$$fuelcons_j = \hbar_0^f(x_j, t_j; \widetilde{\boldsymbol{x}}_j)$$
$$acc_j = \hbar_0^a(x_j, t_j; \widetilde{\boldsymbol{x}}_j)$$

We assume that firms maximize profits and consumers value lower fuel consumption and higher acceleration performance. We also make use of the fact that, for any choice of acceleration performance and technology features, there is only one choice of $\boldsymbol{x}_{\mathrm{j}}$ that minimizes the production cost associated with a given level of fuel efficiency. Conditional on these assumptions the vehicle design decisions captured by the system of equations above can be reduced to:

$$fuelcons_j = \hbar_1\left(acc_j, t_j; \widetilde{\boldsymbol{x}}_j\right)$$

[6] For example, consider a given vehicle design such as the Honda Accord. If Honda wants to increase the fuel efficiency of the Accord, it could decrease the displacement size of the engine, or it could simply change the programming in the powertrain electronic control unit to favor fuel efficiency over acceleration performance. Each of these adjustments to improve fuel efficiency will cause some loss in acceleration performance.

[7] Although there are some other vehicle attributes that can be modified in the medium run (such as the styling of the vehicle body), we expect that fuel consumption and acceleration performance are the most important in terms of being affected by the CAFE standards.

This simplification reduces the set of choice variables and thus the computational complexity.[8]

These physics-based engineering relationships are nested within a standard differentiated product oligopoly model of the automotive industry. In the medium-run, multiproduct firms choose prices and medium run design parameters to maximize profits over the set of vehicles they produce:

$$\begin{array}{lll} \max\limits_{p_j, acc_j, t_j \, \forall j} & \pi = \sum\limits_j q_j \left(p_j - c_j\right) & \\ \mathit{subject\ to} & \boldsymbol{g}\left(\boldsymbol{x_j}, \boldsymbol{t_j}; \tilde{\boldsymbol{x}}_j\right) \leq \mathbf{0} & \\ & stand_l - CAFE_l \leq 0 \; \forall \, l & \quad (1) \\ \mathit{where} & q_j = f\left(p_j, fuelcons_j, acc_j; \tilde{\boldsymbol{x}}_j\right) & \\ & fuelcons_j = h_1\left(acc_j, t_j; \tilde{\boldsymbol{x}}_j\right) & \\ & c_j = h_2\left(acc_j, t_j; \tilde{\boldsymbol{x}}_j\right), & \end{array}$$

The variables q_j, p_j, and c_j are respectively the quantity demanded, price, and marginal cost associated with vehicle j. The constraint $\boldsymbol{g}\left(\boldsymbol{x_j}, \boldsymbol{t_j}; \tilde{\boldsymbol{x}}_j\right) \leq \mathbf{0}$ represents any applicable restrictions on powertrain design variables and technology features that are not feasible. Given these engineering constraints, $\boldsymbol{g}$, and the fixed design parameters $\tilde{\boldsymbol{x}}_j$, fuel consumption is determined jointly by the choice of technology features $\boldsymbol{t_j}$ and acceleration performance.

The reformed CAFE policy is represented in this formulation as a constraint for each vehicle class l (i.e. passenger cars and light trucks). We define $CAFE_l$ to be the sales-weighted average fuel economy of all vehicles in vehicle class l that the firm produces, according to the CAFE policy. This weighted average must equal or exceed the firm's standard for that vehicle class, $stand_l$. This policy is described in more detail in Section 4.

Conditional on the long-run design parameters $\tilde{\boldsymbol{x}}_j$, vehicle demand is assumed to be a function of vehicle price, fuel consumption, and acceleration performance. Note that the design parameters $\boldsymbol{x_j}$ and technology features $\boldsymbol{t_j}$ are not explicitly represented in these demand equations. Technology features and powertrain tuning parameters impact demand through their influence on fuel consumption and acceleration performance. They should not have any direct, intrinsic value to the consumer.

[8] We could instead have reformulated the optimization with $fuelcons$as a decision variable instead of acc. This convention is arbitrary and does not affect the formulation.

The engineering and economic tradeoffs associated with the medium-run design decisions are summarized in this formulation by the functions, $\hbar_1$ and $\hbar_2$. Taken together, $\hbar_1$ and $\hbar_2$ define a surface that we will subsequently refer to as the Production Possibility Frontier (PPF). Points along the frontier represent the set of values $\{fuelcons_j, acc_j, c_j\}$ that are attainable in the medium run conditional on $\tilde{\boldsymbol{x}}_j$.

The PPFs we estimate should not be interpreted as structural representations of the physical relationships between design parameters and vehicle performance. Instead, the functions $\hbar_1$ and $\hbar_2$ summarize the results of detailed engineering simulations and associated cost calculations. More precisely, we use an engineering vehicle simulation model called AVL Cruise AVL Cruise is widely used by major automotive manufacturers to inform powertrain design and development (Mayer 2008). Incorporating this engineering simulation model directly into our model of firms' design decisions is not possible due to large computational costs. Instead, we fit flexible functional forms, $\hbar_1$ and $\hbar_2$, to the results of a large set of detailed data generated from these engineering simulations. The engineering simulations and subsequent estimation are described in detail in Section 3.

2.3 Comparison of engineering versus econometric approaches

This is not the first paper to investigate the physical relationships that constrain vehicle design decisions. In the economics literature, recent studies have developed econometrically estimated models of endogenous product design for the automotive market (Gramlich 2008; Klier and Linn 2008). Although the econometric approach has its advantages, we argue that engineering estimates are more appropriate for our purposes.

Using physics-based simulations to identify the engineering tradeoffs between vehicle attributes allows us to model tradeoffs and attribute combinations that are technologically possible, but have not yet been implemented in existing vehicles. This is important because these new combinations of design attributes may become optimal under the policy regimes we are interested in analyzing. In fact, fuel economy standards historically forced the frontier of vehicle design (Klier and Linn, 2010). Moreover manufacturers have stated that they will rely on further advancing this frontier in order to meet the reformed CAFE standards.

Furthermore, our engineering approach allows us to isolate the tradeoffs between fuel efficiency, acceleration performance, and production costs without conflating changes in unobserved attributes that typically affect both fuel efficiency and consumer demand. In contrast, econometric approaches are more limited in their ability to account for correlation between endogenous attributes and unobserved attributes. For example, the 2010 Chrysler 300 Touring with a 2.7 L engine option has a combined fuel economy of 21.6 mpg and a 0-60 acceleration of 10.5 s, whereas the 3.5 L engine option has 19.7 mpg and 8.5 s. However, engine options are correlated with unobservable attributes; in addition to a larger engine, the 3.5 L Touring also contains a suite of electronic accessories including anti-lock brakes, electronic traction control, light-sensing headlamps, and an upgraded stereo system. The addition of these extra accessories could increase vehicle weight and consume additional energy, further reducing fuel economy. Moreover, these accessories typically increase demand, violating the *certeris paribus* assumption in counterfactual policy simulations.

3 Modeling medium-run vehicle design decisions

The previous section provided a high-level overview of our approach to modeling vehicle design decisions. In this section, we provide a more detailed explanation of how we implement the approach.

3.1 Conceptual framework

In the medium-run, there are two general classes of design changes that can affect the fuel efficiency performance of a vehicle. The first involves tuning the design of the powertrain. These powertrain design adjustments can generate approximately continuous changes in fuel efficiency. The second involves incorporating extra "technology features" into vehicle design. Examples include high efficiency alternators, low resistance tires, and improved aerodynamic drag of the vehicle body (NHTSA 2008). Adding one or more of these features can deliver discrete improvements in fuel efficiency. We make a conceptual distinction between continuous and discrete tradeoffs as they pertain to medium-run fuel efficiency improvements. The solid lines in Figure 2 represent what we call "iso-technology" curves. For each vehicle class, we define a set of viable combinations of technology features. We construct an iso-technology curve for each combination. The "baseline" curve assumes no additional technology features are installed. Movements along this curve capture the engineering tradeoffs between fuel consumption and 0-60 acceleration time for a vehicle with no extra technology features. The addition of technology features effectively shifts the baseline curve down (in the direction of fuel efficiency improvements).

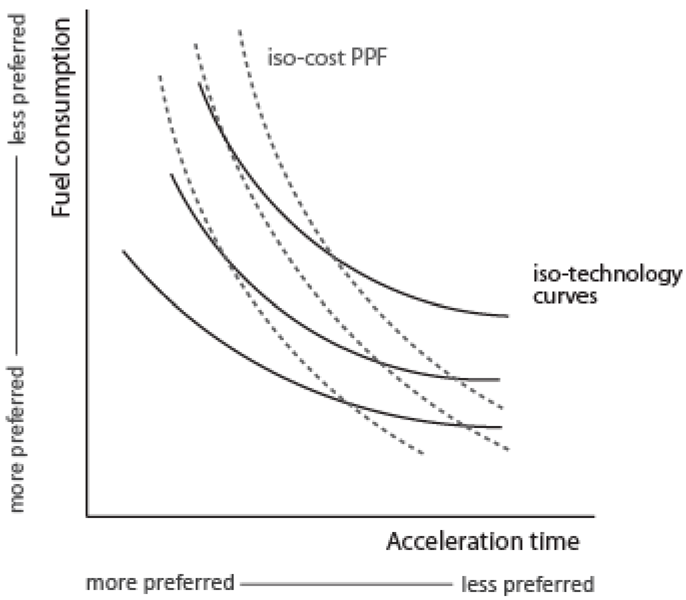

Figure 2: Illustration of medium-run Production Possibility Frontier (PPF) showing both iso-technology curves and iso-cost curves along the PPF

A simple example helps to build further intuition for Figure 2. Consider a particular vehicle design such as the Honda Accord. Honda can decrease the fuel consumption of the Accord without adding any additional technology features by moving along the baseline “iso-technology curve” represented by the highest curve in Figure 2, by adjusting the design of the powertrain (e.g., changing the engine displacement or final drive ratio). Alternatively, Honda can decrease fuel consumption via the addition of technology features. This is modeled as a movement to a lower iso-technology curve in Figure 2.

3.2 A tractable model of physical design constraints and tradeoffs

The first step in our simulation of vehicle design tradeoffs involves the construction of segment-specific “bundles” of design variables indexed $b=1\ldots B$. Each bundle is comprised of a set of values corresponding to the manipulable parameters in $\boldsymbol{x}_s$ (such as engine displacement size) and the fixed design parameters in $\tilde{\boldsymbol{x}}_s$ (such as curbweight). Parameter values are varied across segment-specific bundles at small intervals over the range of possible values.

Each segment-specific bundle is used to parameterize the AVL Cruise model. The fuel efficiency of a particular vehicle design, conditional on the assumed bundle of design parameters, is determined in AVL Cruise by simulating the EPA's fuel-economy test procedures. $Fuelcons_{sb}$ represents the simulated fuel consumption per mile for a vehicle in segment s with design variables b. Additional details of the vehicle simulations are discussed in Appendix A.

Following this process, we generate almost 30,000 sets of simulation results using the AVL Cruise model, each comprising the fuel consumption and acceleration performance of a simulated vehicle in a vehicle segment, s, and the corresponding bundle of design parameter inputs, b. These design variables and the corresponding simulated fuel consumption are used to estimate the baseline iso-technology curve. Several parametric specifications for this relationship were tested and the specification in equation (2) performed the best under the Akaike Information Criterion.

$$fuelcons_{sb} = \kappa_{1s} + \kappa_{2s}e^{-acc_{sb}} + \kappa_{3s}wt_{sb} + \kappa_{4s}wt_{sb} \cdot acc_{sb} + \varepsilon_{sb} \tag{2}$$

The parameter wt_{sb} in equation (2) is the vehicle curbweight. The ε_{sb} term represents the error associated with approximating the calculations performed in the vehicle simulations with this parametric function.

Figure 3 plots observed and estimated fuel consumption and 0-60 acceleration time for a particular vehicle type (a compact vehicle) with no extra technology features. Equation (2) predicts observed vehicle performance reasonably well (R^2=0.76). Appendix C discusses comparisons between estimated and observed performance attributes in more detail.

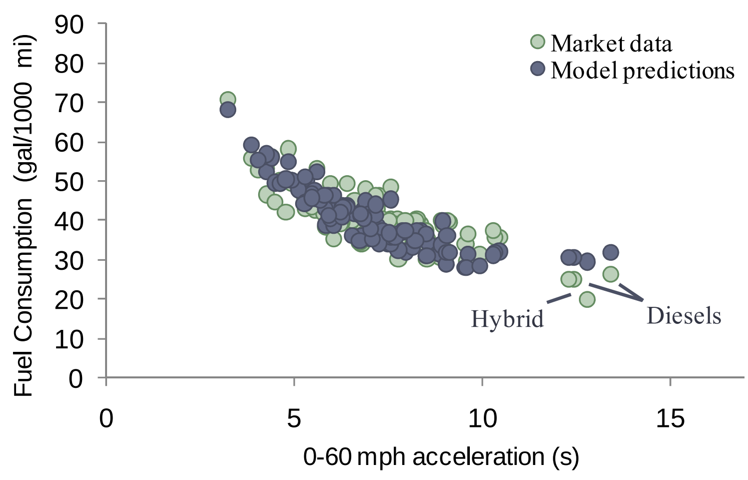

Figure 3: Comparison of compact-segment model to MY2006 vehicle data

Having estimated the baseline curves, we now turn to the technology features that can be added to improve fuel efficiency performance. NHTSA (2008) estimated the effect (in percentage terms) of each technology feature on the fuel economy vehicles in each segment. These estimates are based on values reported by automotive manufacturers, suppliers, and consultants. To determine how the addition of one or more technology features affects the position of the iso-technology curve relative to the baseline described above, we combine the AVL Cruise simulations and data from NHTSA (2008).[9] Further details of this process are described in Appendix D

We model the design implications of most, but not all, of the technology features that NHTSA considers in their analysis. The majority of technology features we omit from our analysis are only available in longer run planning stages. Some features are eliminated due to the challenges in simulating their effects (e.g., variable valve timing). Omitting these technology features will only make our estimated costs of CAFE regulations more conservative, representing an upper bound on costs, because we are failing to account for design features that could be cost effective.

We consider only those combinations of technology features that are cost effective—meaning that there is no lower cost combination that could achieve the same or better level of acceleration performance and fuel efficiency. This reduces the set of technology feature combinations under consideration to between 20 and 76, depending on vehicle segment. Although each technology feature we consider is represented explicitly in the engineering simulations, it is computationally infeasible to model the full set of discrete choices in our policy simulations. For the purpose of tractability, we introduce a simplification.

We augment equation (2) to accommodate the addition of one or more technology features as follows:

$$fuelcons_{sb} = \kappa_{1s} + \kappa_{2s}e^{-acc_{sb}} + \kappa_{3s}wt_{sb} + \kappa_{4s}wt_{sb} \cdot acc_{sb} + \kappa_{5s}\, tech + \kappa_{6s} tech \cdot acc^2 + \varepsilon_{sb} \quad (3)$$

[9] To determine how the baseline iso-technology curve shifts with technology features, we need to know the impact of each technology feature on 0-60 acceleration time, which is not reported by NHTSA. We accomplish this by using AVL Cruise to simulate vehicle performance effects from each technology feature, assuming fuel efficiency improvements that match those reported by NHTSA. For example, NHTSA reports a 0.5% improvement in fuel economy from using "low friction lubricants" in compact vehicles. We simulate this impact by reducing the friction losses in the engine of our representative compact vehicle model until we observe fuel economy improving by 0.5% and then observe the percentage improvement of 0-60 acceleration time.

The technology features are modeled as a continuous variable, *tech*, ranging from zero (the baseline case) to the maximum number of cost-effective combinations of technology features for each vehicle class. Note that a particular value for *tech* maps to a specific combination of technology features (e.g., low resistance tires and a high efficiency alternator) and does not represent the number of technology features. Technology feature combinations are ordered by decreasing fuel consumption for the same acceleration performance, which is also increasing in cost. Therefore, a higher value of *tech* corresponds to a lower fuel consumption and higher cost vehicle conditional on 0-60 acceleration time.

The impact of this continuous approximation on the engineering simulation results is relatively small with the average gap between discrete features less than 1 mpg. Furthermore, we provide evidence in Appendix B that the particular specification we use to estimate the relationships of the continuous *tech* variable to fuel consumption and cost preserves important properties of the discrete technology combinations.

3.3 Modeling the costs of medium-run design decisions

In addition to modeling the impact of medium-run design decisions on fuel efficiency and acceleration performance, we also need to account for the effect of these design decisions on vehicle production costs. The broken lines in Figure 2 are meant to represent iso-cost curves. Importantly, movements along iso-technology curves in the direction of improved fuel efficiency increase costs. It is possible to improve fuel efficiency without increasing production costs through the use of technology features, but this comes at a cost of reduced acceleration performance.

We use two separate sources of data to estimate these costs, one describing costs as a function of medium-run powertrain variables, $\boldsymbol{x}$, and another detailing the production costs associated with each technology feature, t_i.

The production costs of vehicle model j is composed of two parts:

$$c_j = engcost_j + \omega_j \tag{4}$$

The first cost component represents the portion of production costs that are dependent on the medium-run design decisions considered in this analysis. The second component represents the portion of costs that are independent of these medium-run design decisions. This second portion can be derived from the first order conditions of firms' profit maximization assuming that observed vehicle prices are in equilibrium. The procedure we use to construct these cost estimates is fairly standard. We follow Jacobsen's (2010) approach as described in Appendix C.

The first cost component in equation (4) is completely determined by medium-run design decisions. Production cost estimates associated with the baseline iso-technology curves are taken from Michalek et al. (2004). These authors collected detailed cost data from manufacturing, wholesale, and rebuilt engines of varying displacements and estimate a model to explain variation in production costs as a function of engine displacement. We make use of the fact that, conditional on all other design decisions, there is a one-to-one mapping between engine displacement and segment-specific combinations of acceleration performance and curbweight. This allows us to use the results of Michalek et al. (2004), together with our engineering simulation results, to construct the cost estimates.[10] Again, many parametric forms of this relationship were tested and the specification in equation (5) performed the best using the Akaike Information Criterion.

$$c_{sb} = \sigma_{1s} + \sigma_{2s} e^{-acc_b} + \sigma_{3s} wt_b + \sigma_{5s} 4wt_b \cdot acc_b \tag{5}$$

The additional production costs resulting from each technology feature are taken from NHTSA (2008). NHTSA generated these vehicle-segment specific estimates using reported values from automotive manufacturers, suppliers, and consultants. These data, which are shown in Table 1, are currently used to perform cost-benefit analyses of the CAFE regulations.

We treat the costs of technology features and the costs of adjusting powertrain tuning variables as additively separable. Engines are manufactured separately from other subsystems of the vehicle before assembly. Most of the technology features we consider do not require changes in engine design or affect the assembly of the engine with other vehicle subsystems. This is consistent with our assumption that costs are additively separable. There are only two exceptions. Two technology features—engine friction reduction and cylinder deactivation—do affect the engine subsystem. Even in these cases, it is reasonable to approximate technology costs as additively separable from the baseline production cost of the engine. For example,

[10] This assumes that firms minimize costs and that adjustments to the final drive ratio have negligible impact on production costs.

engine friction can be reduced by using lubricants, the costs of which are independent of all medium-run decisions considered.[11]

Table 1: Marginal Costs of Fuel-Saving Technology Costs Based on NHTSA's (2008) Analysis

Technology Costs	Two Seater	Compact	Midsize/ Minivan	Fullsize	SUV	Small Pickup	Large Pickup/Van
Low friction lubricants	$3	$3	$3	$3	$3	$3	$3
Engine friction reduction	$126	$84	$126	$126	$126	$126	$168
Aggressive shift logic	$38	$38	$38	$38	$38	$38	$38
Early torque converter lockup	$30	$30	$30	$30	$30	$30	$30
High efficiency alternator	$145	$145	$145	$145	$145	$145	$145
Aerodynamic drag reduction	$38	$38	$38	$38	$38	$38	$38
Low rolling resistance tires	$6	$6	$6	$6	$6	$6	$0
Cylinder deactivation	n/a	n/a	$203	$203	$203	$203	$229

Notes: *cost* represents the unit production cost in $/vehicle produced, % *mpg* is the percentage increase in combined highway-city fuel economy, and % *acc* is the percentage reduction in 0-60 mph acceleration time in seconds. Cost and fuel economy figures are taken from NHTSA (2008). The change in acceleration is calculated in the engineering vehicle simulation model, AVL Cruise.

Combining these two sets of cost estimates, equation (6) models the portion of production costs that are dependent on the medium-run design decisions considered:

This portion of production costs is a function of curbweight, acceleration time, and the continuous measure of technology features.

4 Modeling vehicle demand

In this section, we introduce an empirical model of vehicle demand. The specification follows the seminal work by Berry et al. (1995), and subsequent work by Berry et al. (2004), Train and Winston (2007), Langer (2011), and others. One distinguishing feature of our identification approach is our choice of instruments which is informed by structure the vehicle design process (as described in Section 2).

[11] The case of cylinder deactivation poses a larger challenge for treating technology costs as additively separable from engine costs. Given large changes in engine displacement achieved by switching the engine architecture (e.g., replacing a V-8 engine with a V-6) would slightly reduce the costs of cylinder deactivation due to a smaller number of cylinders. However, even with this cost reduction, cylinder deactivation is the highest-cost technology feature considered and therefore would not significantly affect counterfactual results.

4.1 Demand side data

We use two sources of data to estimate the model of vehicle demand: a household-level survey conducted by Maritz Research, and vehicle characteristic data available from Chrome Systems Inc.. The Maritz Research U.S. New Vehicle Customer Study (NVCS) is a monthly survey of households that purchased or leased new vehicles. This survey provides information on socio-demographic data, household characteristics, and the vehicle identification number (VIN) for the purchased vehicle. The survey also asks respondents to list up to three other vehicles considered during the purchase decision. We use data from this survey during the twelve-month 2006 model-year. Approximately one-third of respondents in this dataset listed at least one considered vehicle. Because the survey oversamples households that purchase vehicles with low market shares, we take a choice-based sample from this data such that the shares of vehicles purchased by the sampled households matches the observed 2006 model-year market shares.

We supplement the survey data with information on vehicle characteristics using Chrome System Inc.'s New Vehicle Database and VINMatch tool. Vehicle alternatives are identified using the reported VIN, distinguishing vehicles by their make, model, and engine option, with a few modifications. We eliminate vehicles priced over $100,000, which represent a small portion of market sales, and remove seven vehicle alternatives that were not chosen or considered by any survey respondent. We further reduce the data set by consolidating pickup truck and full-size van models with gross vehicle weight ratings over 8,000 lb to only two engine options each. Summary vehicle data are described in Table 2.

Table 2: Summary of Vehicle Characteristic Data

Attribute	Unit	Mean	Std. Dev	Min	Max
MSRP	*1,000 2006$*	32.67	16.73	11.93	97.49
Fuel Economy	*mpg*	21.46	5.14	10.98	56.55
Horsepower	*hp*	241	78	65	520
Curb weight	*1,000 lb*	3.87	0.85	1.98	6.40
Footprint	*1,000* in^2	13.92	2.00	9.52	20.05
Make Grps.	38				
Obs.	473				

4.2 Demand side model

Following Berry et al. (1994), we model vehicle demand at the household level and then aggregate up to obtain product level demands. New automobiles are described as bundles of attributes. Consumers are assumed to choose the vehicle that maximizes the utility derived from these attributes. As noted above, the design parameters $\boldsymbol{x_i}$ and technology features $\boldsymbol{t_i}$ affect utility only indirectly through their effect on vehicle attributes such as fuel efficiency and acceleration.

The utility that consumer *n* derives from purchasing vehicle model *j* is defined as:

$$u_{nj} = \delta_j + \sum_{\boldsymbol{kr}} \boldsymbol{a_{jk} z_{nr}}\, \beta_{kr} + \sum_{\boldsymbol{k}} \boldsymbol{a_{jk} \nu_{nk} \mu_k} + \epsilon_{nj} \tag{7}$$

where $j=0\ldots J$ indexes all of the vehicles competing in the market. The utility obtained when no new vehicle is purchased (i.e. the outside option) is u_{i0}. The $\boldsymbol{a}_j$ are observable vehicle attributes such as the ratio of vehicle price to income; "gallons per mile", *gpm*, (the inverse of fuel economy); the inverse of 0-60 acceleration time; and vehicle footprint.

Conceptually, this utility function can be decomposed into four components. The vehicle model-specific fixed effect, δ_j, represents the portion of utility that is the same across all consumers:

$$\delta_j = \sum \quad a_{jk}\overline{\beta_k} + \xi_j \tag{8}$$

The attribute-specific coefficients $\bar{\beta}_k$ are common across all consumers. The ξ_j are vehicle attributes that are valued by the consumer, but not observed by the econometrician. Examples include vehicle handling. One might expect the ξ_j to be correlated with the vehicle attributes of primary interest: vehicle price, fuel efficiency, and acceleration performance. The key insight, illustrated by equations (7) and (8), is that this structural error only enters the mean utility level. Furthermore, the mean utility level is a linear function of ξ_j. Following Berry (1994), this allows us to move the endogeneity problem out of the non-linear equation (7) and into a linear regression framework in equation (8) where the problem is more easily dealt with using two-stage least squares.

The second component in equation (7) captures the component of utility that varies systematically with observable consumer characteristics $\boldsymbol{z_n}$. Interactions between consumer characteristics and vehicle attributes play an important role in determining substitution patterns.

Our specification allows preferences for pickup trucks to vary across rural and urban areas, and includes interactions between children in the household and SUV and minivan segments.

The third component captures the effects of interactions between vehicle attributes $\boldsymbol{a}_j$ and household characteristics we cannot observe. This allows for some random variation in consumer preferences for specific vehicle attributes. The random coefficients $\boldsymbol{\mu}$ are assumed to be distributed normally in the population according to the distribution $f(\boldsymbol{\mu}|\boldsymbol{\sigma})$.

Finally, the disturbance term ϵ_{nj} captures idiosyncratic individual preferences which are assumed to be independent of the product attributes and of each other the effects of unobserved determinants of utility that vary randomly across consumers. We assume that these idiosyncratic errors have an i.i.d. Type I extreme value distribution. This assumption yields the familiar logit functional form for the vehicle choice-share probabilities, P_{ni}, conditional on z , v, and the parameters to be estimated, θ, as shown in equation (9).

$$P_{ni} = \Pr(y_n = i|z_n, v_n, \theta) = \frac{\exp\left(\delta_i + \sum_{kr} \boldsymbol{a}_{ik}\boldsymbol{z}_{nr}\,\beta_{kr} + \sum_k \boldsymbol{a}_{ik}\boldsymbol{v}_{nk}\boldsymbol{\mu}_k\right)}{1 + \sum_j \exp\left(\delta_j + \sum_{kr} \boldsymbol{a}_{jk}\boldsymbol{z}_{nr}\,\beta_{kr} + \sum_k \boldsymbol{a}_{jk}\boldsymbol{v}_{nk}\boldsymbol{\mu}_k\right)} \quad (9)$$
$$\equiv \frac{\exp(u_{nj})}{1 + \sum_j \exp(u_{nj})}.$$

The predicted market share of a vehicle, j, is $\sum_n P_{jn}$.

Following Train and Winston (2007), the utility formulation is extended to include information about ranked choices when these data are available for a respondent. The ranking is specified as $u_{ni} > u_{nh_1} \ldots u_{nh_m} > u_{nj}$ for all $j \neq i, h_1, \ldots, h_m$ where i is the chosen vehicle; h_1 is the second ranked vehicle (the vehicle that would have been chosen if vehicle i was not available) and h_m is the m ranked vehicle. Therefore, the probability that respondent n purchased vehicle i and ranked vehicle h_1 through h_m is defined as:

$$L_{nih_1 \ldots h_m} = \left(\frac{e^{U_{ni}}}{1 + \sum_j e^{U_{nj}}}\right)\left(\frac{e^{U_{nh_1}}}{\sum_{k \neq i} e^{U_{nk}}}\right) \ldots \left(\frac{e^{U_{nh_m}}}{\sum_{l \neq i, h_1, \ldots, h_{m-1}} e^{U_{nl}}}\right) \quad (10)$$

The first two terms of this formulation correspond to the probability that the consumer purchased vehicle i, given all available vehicle models and the outside good, and the probability that they would have purchased vehicle h_1 if vehicle i and the outside good were not available.[12] The

[12] The outside good is removed from the ranked choice set (all but the first term in equation 7) because respondents indicated that they considered the ranked vehicles during their purchasing decision, but it is not clear if

outside good is excluded from the denominator of every term but the first because we do not observe whether the respondents would have chosen not to purchase a vehicle if their first choice was not available. When no ranking data are available for a respondent, the likelihood consists of only the first term in equation (10).

Recently, significant concerns have been raised about the sensitivity of parameter estimates using similar random-coefficient discrete choice demand models (Knittel and Metaxoglou 2008). We estimate the model using a series of randomly selected initial values to test the robustness of our estimates. Specifically, ten initial values are randomly selected from a uniform distribution from -15–15. Initial values outside of this range were also tested but these initial points often produced values of the log-likelihood that were near negative infinity, indicating a very poor fit to the model, which prevents the algorithm from solving the estimation problem.

4.3 Identification strategy

One distinguishing feature of our demand estimation is our choice of instruments. It has become standard in the literature to use functions of non-price attributes, including vehicle dimensions, horsepower and fuel economy, as instruments for price (e.g., Berry et al. 1995; Berry et al., 2004; Train and Winston 2007). This approach has been criticized on two accounts: 1) firms presumably choose these non-price attributes and prices simultaneously, and 2) decisions regarding unobserved attributes may depend on previously determined non-price attributes, rendering them invalid as instruments.[13] As Heckman (2007) notes, to obtain valid instruments in this context requires a model of the determinants of product attributes. The engineering design literature provides this kind of model.

Our choice of instruments is informed by the engineering simulations and related information regarding how design decisions are made. We use only those attributes determined from longer run product-planning schedules to instrument for variation in price, fuel efficiency,

---

they would have chosen the 1st ranked vehicle, for instance, if the vehicle they purchased was not available or if they would instead have chosen to not purchase a new vehicle.

[13] Berry et al. (1995), and Train and Winston (2007) both focused on short-run pricing decisions and therefore the assumption that many vehicle attributes are exogenous to their analysis is justified. However, anecdotal evidence suggests that automotive manufacturers routinely adjust the electronic control unit of vehicle engines, which affects fuel economy and acceleration performance, in the same time frame as setting suggested retail prices and thus fuel economy may not be exogenous to pricing decisions.

and acceleration performance.[14] More precisely, we use the moments of vehicle dimensions of same-manufacturer vehicles (*din, dsqin*) and different-manufacturer vehicles (*dout, dsqout*), powertrain architecture (i.e., hybrid, turbocharged, and diesel), and drive type (i.e., all wheel drive or 4-wheel drive). These instruments can be considered fixed in the medium run.

5 Empirical Results

In this section, we first present the estimates of the design model parameters for each vehicle class as described in Section 3. We then discuss the demand-side estimation.

5.1 Endogenous design model estimation

Parameters defining the tradeoffs between vehicle fuel efficiency, acceleration performance, and production costs are estimated using the data generated from the vehicle simulations described in Section 3. Estimated parameters for equation (3), summarizing the relationship between vehicle attributes dependent on technology features and powertrain parameters for each vehicle class, are reported in Table 3. The estimated relationships fit the vehicle data in each class reasonably well (R^2>0.89) except for the two-seater class (R^2=0.44). However, the two-seater class comprises less than 1% of vehicle sales in MY2006 so the poorer fit of this class should not significantly affect the policy simulation results.

All parameter estimates have expected signs. The negative sign on the variable representing technology features and positive sign on the interaction between the technology variable and acceleration indicate that implementing more fuel-efficient combinations of technology features reduces fuel consumption with decreasing returns, as illustrated in Appendix B. The positive sign on the weight parameter and negative sign on the weight-acceleration interaction imply that the iso-technology curves in Fig. 2 shift up and rotate clockwise with vehicle weight. This indicates that heavier vehicles will have worse fuel consumption given the same 0-60 mph acceleration time, as expected, but that this effect increases for vehicles with faster acceleration. Validation tests of these results, comparing model predictions to observed market data, are described in Appendix C.

[14] Literature detailing the automotive design process allows us to address the first criticism of instrument choice. A remaining assumption in our approach is that these longer run attributes do not affect choices of unobserved attributes in the medium run.

Table 3: Estimation Results for Fuel Consumption in Technology and Design Model

	Two seater			Compact			Midsize / Minivan			Fullsize		
	param		**std. err.**	**param**		**std. err.**	**param**		**std. err.**	**param**		**std. err.**
constant	20.8484	***	0.9414	10.7920	***	0.2158	11.5531	***	0.5674	10.8515	***	0.4034
$\exp(-acc_j)$	89.6806	***	31.1595	69.5244	***	5.3977	733.706	***	55.3714	452.552	***	23.7850
$tech_j$	-0.2049	***	0.0245	-0.2605	***	0.0129	-0.0829	***	0.0018	-0.0794	***	0.0042
$acc_j^2 \cdot tech_j$	0.0016	*	0.0005	0.0013	**	0.0003	0.0002	***	2.3E-05	0.0002	***	2.7E-05
wt_j	2.9159	***	0.3667	12.9897	***	0.3047	8.6680	***	0.1119	8.3583	***	0.1208
$wt_j \cdot acc_j$	0.0280		0.0476	-0.5593	***	0.0456	-0.2954	***	0.0223	-0.2539	***	0.0079
R^2	0.443			0.941			0.896			0.976		
Obs.	4473			5117			3542			3542		
	SUVs			**Small Pickup**			**Large Pickup / Van**					
	param		**std. err.**	**param**		**std. err.**	**param**		**std. err.**			
constant	14.0535	***	0.4985	10.5032	***	0.4049	10.5185	***	0.4083			
$\exp(-acc_j)$	1329.854	***	115.804	13387.530	***	1836.191	2979.096	***	411.241			
$tech_j$	-0.1540	***	0.0082	-0.0933	***	0.0071	-0.0890	***	0.0065			
$acc_j^2 \cdot tech_j$	0.0006	***	0.0001	0.0003	**	0.0001	0.0003	**	0.0001			
wt_j	9.0653	***	0.1069	8.9932	***	0.3207	8.6604	***	0.2813			
$wt_j \cdot acc_j$	-0.2934	***	0.0148	-0.2224	***	0.0295	-0.2225	***	0.0299			
R^2	0.965			0.952			0.952					
Obs.	16863			9450			9450					

*p<0.1, **p<0.05, ***p<0.01, all others p>0.1, standard errors are clustered by vehicle curb weight

The estimates describing the relationship between production costs and choices of acceleration performance and technology implementation, described by equation (4), are reported in Table 4. These estimates fit all vehicle classes reasonably well (R^2>0.83). Production costs increase with the level of technology implementation by design. As expected, production costs decrease as 0-60 mph acceleration times get slower. The positive sign on the weight term and negative sign on the weight-acceleration interaction term indicate that incrementally improving acceleration is more costly in heavier vehicles, and this effect is magnified for vehicles with relatively better acceleration performance. All parameter estimates in both equations are significant to the 90% level or better.

Table 4: Estimation Results for Cost of Technology and Powertrain Design

	Two seater		Compact		Midsize / Minivan		Fullsize	
	param	**std. err.**	**param**	**std. err.**	**param**	**std. err.**	**param**	**std. err.**
constant	0.3669 *	0.0865	0.7800 ***	0.0091	0.5540 ***	0.0355	0.4029 ***	0.0450
$\exp(-acc_j)$	10.6686 *	2.2311	1.9716 ***	0.1631	24.3842 ***	1.3842	24.0527 ***	2.7273
$tech_j$	0.0175 ***	0.0001	0.0016 ***	0.0002	0.0054 ***	0.0001	0.0057 ***	3.0E-05
wt_j	0.2579 ***	0.0132	0.2250 ***	0.0051	0.1963 ***	0.0103	0.2339 ***	0.0052
$wt_j \cdot acc_j$	-0.0082 ***	0.0013	-0.0123 ***	0.0005	-0.0071 ***	0.0012	-0.0069 ***	0.0004
R^2	0.890		0.898		0.898		0.931	
Obs.	4473		5117		3542		3542	
	SUVs		**Small Pickup**		**Large Pickup / Van**			
	param	**std. err.**	**param**	**std. err.**	**param**	**std. err.**		
constant	0.0200 **	0.1337	0.3025 *	0.0607	0.3025 *	0.0607		
$\exp(-acc_j)$	92.3965 ***	16.4768	719.579 *	162.643	160.560 *	36.291		
$tech_j$	0.0038 ***	0.0003	0.0066 ***	0.0001	0.0066 ***	0.0001		
wt_j	0.3470 ***	0.0143	0.2621 ***	0.0137	0.2538 ***	0.0117		
$wt_j \cdot acc_j$	-0.0108 ***	0.0016	-0.0055 *	0.0014	-0.0055 *	0.0014		
R^2	0.887		0.831		0.831			
Obs.	16863		9450		9450			

*p<0.1, **p<0.05, *** p<0.01, standard errors are clustered by vehicle curb weight

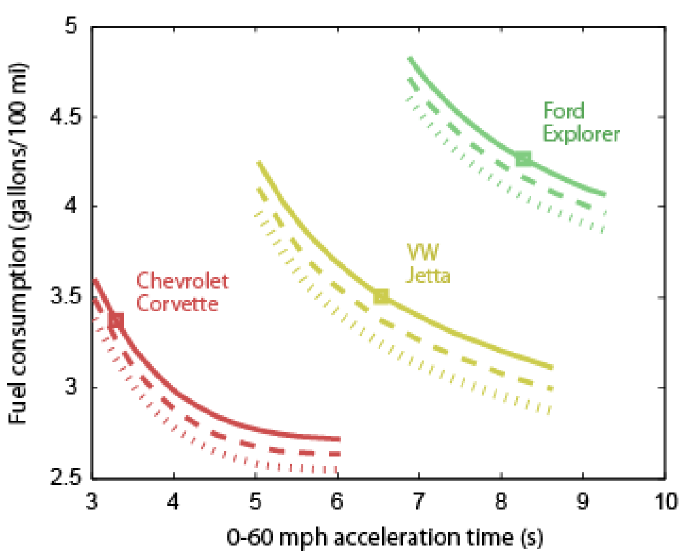

Figure 4: Estimated PPFs illustrated as iso-cost curves for selected vehicles (■ current location, – baseline frontier,‒ ‒ $100 design changes,···· $200 design changes)

5.2 Demand-side estimation

Table 5 summarizes the demand model parameter estimates. All coefficient estimates have the expected signs. Recall that the σ parameters represent the standard deviations of the demand parameters in equation (5) that are allowed to vary randomly in the population. Only the standard deviation of the fuel consumption coefficient is found to be statistically significantly different from zero. Several of the parameters that capture the effects of interactions between vehicle attributes and observable consumer attributes are found to be statistically significant, including the ratio of price to income (*p/inc*), and the interactions between minivans and children, SUVs and children, and pickup trucks and living in a rural location.

Table 5: Heterogeneous Demand Parameter Results

	σ	
	param	**st. err.**
p	0.0366	0.0237
gpm	0.0215	0.0131
accinv	0.0335	0.0533
ftp	0.0390	0.0534
	μ	
	param	**st. err.**
p/inc	-0.1721	0.0246
minivan-child	7.4388	0.5603
suv-child	0.9447	0.1468
truck-rural	1.8365	0.2422

Notes: The μ's are the estimates of the demand parameters for attribute-demographic interactions in equation 5, and the σ's are the estimates of the standard deviations of the normally-distributed random-variable parameters on vehicle attributes.

Results of initial-value tests of these estimates are reported in Table 6. All ten initial values resulted in the same estimate solutions within 1e-4. Infinite-norms of the gradients for each solution were on the order of 1e-3 to 1e-4, and the hessians at these solutions were all verified to be positive definite. Table 7 presents the OLS estimations of the first-stage regressions of endogenous decisions (price, fuel consumption, and inverse 0-60 mph acceleration time), with F-tests of 20.68, 19.62, and 21.38, respectively.

Table 6: Initial-value tests of heterogeneous demand parameter estimates

test	p	gpm	accinv	ftp	p/inc	minivan-child	suv-child	truck-rural	Log-Lik.	grad. norm
1	0.036643	0.021473	0.033539	0.039035	-0.172102	7.438812	0.944718	1.836548	7053.39256	0.0005
2	0.036643	0.021473	0.033540	0.039033	-0.172102	7.438795	0.944720	1.836550	7053.39256	0.0006
3	0.036643	0.021473	0.033539	0.039034	-0.172102	7.438796	0.944720	1.836558	7053.39256	0.0014
4	0.036643	0.021473	0.033539	0.039035	-0.172102	7.438812	0.944718	1.836548	7053.39256	0.0005
5	0.036643	0.021473	0.033539	0.039034	-0.172102	7.438774	0.944718	1.836565	7053.39257	0.0013
6	0.036643	0.021473	0.033540	0.039033	-0.172102	7.438796	0.944719	1.836554	7053.39256	0.0005
7	0.036644	0.021474	0.033538	0.039035	-0.172101	7.438782	0.944723	1.836571	7053.39256	0.0018
8	0.036643	0.021473	0.033539	0.039034	-0.172103	7.438800	0.944717	1.836572	7053.39256	0.0005
9	0.036643	0.021473	0.033539	0.039032	-0.172102	7.438778	0.944720	1.836556	7053.39256	0.0005
10	0.036643	0.021473	0.033539	0.039277	-0.172100	7.438800	0.944720	1.836600	7053.39257	0.0009

Notes: This table presents the initial value tests of the estimates presented in Table 5**,** representing the heterogeneous demand parameters. The value of the log-likelihood, and the infinity norm of the gradient is reported. Initial values were randomly selected from uniform distributions from -15.0–15.0.

Table 7: First Stage Instrumental Variable Results

	p	gpm	accinv
din	0.0007	-0.0001	0.0005**
dout	0.1368***	-0.0216*	0.0448***
dsqin	-0.0054***	-0.0012**	0.0002
dsqout	-0.0468	0.0086	-0.0112
awd	0.7622***	0.4325***	0.0512
turbo	-0.1561	-0.1118	0.0896
diesel	0.6513	-0.9733***	-0.2685**
hybrid	0.0554	-1.6513***	-0.3875***
ftp	11.4430***	3.0836***	2.6178***
sport	2.885***	0.8498***	0.6702***
truck	-1.4788***	0.065436	-0.089059
suv	0.6541***	0.5843***	0.0843
minivan	-0.9603**	-0.4317***	-0.2751***
constant	-12.5413***	0.3783	-2.1626***

* p<0.1, **p<0.01, ***p<0.001

Notes: *din* and *dout* are the distances of vehicle dimensions (length x width x height) to the average dimensions of same and different manufacturers, respectively; *dsqin* and *dsqout* are these values squared. The remaining variables represent powertrain architectures— *turbo* (turbocharged), *hybrid, diesel*— and the type of drive—all wheel or four wheel drive, *awd*.

Table 8 reports the second-stage IV estimates of the parameters in equation (8). The SUV indicator variable is positive and significant suggesting a preference for SUVs over the omitted category (sedans) when the observable attributes are controlled for. The minivan indicator is negative and significant. The parameter estimate for two-seater sports cars is negative and the parameter for pickup trucks is slightly positive, but neither is significant.

Table 8: Homogeneous Demand Parameter Results

	α	
	param	**st. err.**
p	-0.4591	0.0998
gpm	-0.3677	0.1679
accinv	1.1262	0.3956
ftp	2.4541	0.7504
sport	-0.4654	0.3129
truck	0.0440	0.2412
suv	0.6669	0.2116
minivan	-5.1827	0.3032
constant	-8.0903	0.7262

Notes: This table presents the 2nd stage IV estimators of the demand parameters of vehicle attributes following equation 8.

The average price elasticity implied by these demand parameter estimates is -1.9, (95% CI: -2.0, -1.8), and the sales-weighted average elasticity is -1.7, (95% CI: -1.8,-1.7). These estimates are somewhat lower than those found in previous studies, which range from -2.0 to -3.1 (Jacobsen 2010; Klier and Linn 2008, Train and Winston 2007; Goldberg 1998). Similar to prior studies (Berry et al. 1995; Goldberg 1998; Beresteanu and Li 2008) we find that, in general, demand is more elastic for cheaper "economy" vehicles and less elastic for higher priced vehicles, although this relationship is not monotonic.

With the demand model parameter estimates in hand, we can calculate the implied willingness-to-pay for an increase in fuel economy and contrast this with our estimates of technology costs. Figure 5 illustrates how willingness-to-pay for fuel economy improvements varies across vehicle models. One intuitive reason for this variation is that the value (in terms of fuel cost savings conditioning on driving distance) of a 1 mpg improvement in fuel efficiency varies with the fuel economy of the vehicle. Consequently, all else equal, consumers should be willing to pay more for a 1 mpg improvement in a less fuel efficient car. We also find that consumers who are more likely to purchase cheaper vehicles are, on average, less willing to pay for improvements in any vehicle attributes (including fuel economy).

Figure 5 also shows that our estimates of the technology costs to incrementally increase fuel economy vary considerably between vehicles depending on the vehicle class and characteristics such as fuel economy and weight. For many vehicle models, the technology cost to increase fuel economy by 1 mpg is lower than $200 but these costs are substantially more for larger vehicles.

This variation notwithstanding, across all vehicles we find that the willingness-to-pay for fuel economy improvements is generally lower than the technology costs associated with increasing fuel economy in that vehicle, and considerably less than the value of the associated (discounted) fuel savings over the vehicle's lifetime.[15] This discrepancy between willingness-to-pay for fuel economy with the net-present-value of fuel savings is documented in other studies (Helfand and Wolverton 2009; Alcott and Wozny 2009).

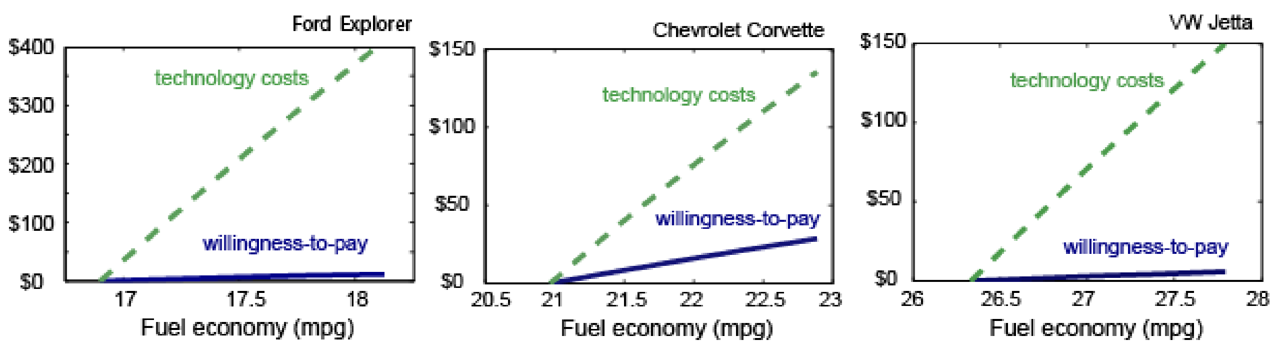

Figure 5: Select vehicles' technology costs and willingness-to-pay to increase fuel economy

Interestingly, our estimates also imply that, in general, consumers' willingness-to-pay for a 1 mpg improvement in fuel economy is lower than their willingness-to-pay for the increase in acceleration performance that would correspond to a loss in fuel economy of 1 mpg. As expected, we find that the consumers who are more likely to purchase luxury vehicles or opt for

[15] Back of the envelope calculations, assuming s a discount rate of 4.5%, a vehicle lifespan of 13 years, constant gas prices at $2.60 (the average in MY2006) and 14,000 annual vehicle miles traveled (the average in 2006 as reported by the Department of Transportation) give a net present value fuel savings of $1,100 for increasing the fuel economy of a vehicle with 21 mpg by 1 mpg.

the higher horsepower vehicle options are willing to pay relatively more for acceleration performance relative to other consumers, and relatively less for fuel economy improvements.

6 Reformed CAFE standards

In this section, we provide an overview of the recently reformed CAFE standards and explain how these standards are represented in our policy simulations.

6.1 Overview of reformed CAFE standards

First enacted by Congress in 1975, the Corporate Average Fuel Economy (CAFE) standards are designed to reduce energy consumption by increasing the fuel economy of cars and light trucks. The CAFE regulation sets a minimum standard for the average fuel economy of a manufacturer's fleet of vehicles sold in the United States. Under this regulation, the average fuel economy for each manufacturer is calculated as a sales-weighted harmonic mean fuel economy across the manufacturer's fleet of vehicles in a particular class (i.e., passenger cars or light trucks). In order to comply with the CAFE policy, this average must be greater than or equal to the CAFE standard, such that:

$$\frac{\sum_{j \in \mathfrak{I}_{f,c}} q_j(p_j)}{\sum_{j \in \mathfrak{I}_{f,c}} q_j(p_j)/mpg_j} \geq stand_c$$

where q_j and mpg_j are the number of sales and fuel economy of vehicle j, and $stand_c$ is the fuel economy standard for vehicle j's class.

If a firm violates this standard, they must pay a fine of $5.50 per 0.1 mpg below the standard for each vehicle produced.[16] Historically, there have been three categories of firm responses to the CAFE standard: all domestic manufacturers (GM, Ford, and Chrysler) have met the standard within an allowable deviation, certain Asian manufacturers (e.g., Toyota and Honda) have consistently exceeded the standard, and many European manufacturers have violated the standard and paid the fine (Jacobsen 2010).

In 2007, Congress passed the Energy Independence and Security Act (EISA), which both raised the requirement for the average fuel economy of new vehicles and modified how the

[16] In addition, a Gas Guzzler Tax is levied on individual passenger car models (but not trucks, vans, minivans, or SUVs) that get less than 22.5 miles per US gallon (10.5 l/100 km).

standards would be set. This reformed CAFE regulation establishes an individual fuel economy target, T_j, for each vehicle, based on vehicle footprint (one measure of vehicle size) such that vehicles with larger footprints have lower standards.[17] Specifically, the fuel economy standard for firm f and vehicle class c is a harmonic average of the fuel economy targets of the firm's vehicles in class c:

$$stand_{f,c} = \frac{\sum_{j \in \Im_{f,c}} q_j(p_j)}{\sum_{j \in \Im_{f,c}} q_j(p_j)/T_j}$$

Unlike the unreformed CAFE standards, the reformed standards vary across manufacturers.

CAFE reform has a number of important implications. Under the previous CAFE regulation, sales of any vehicle that had a higher fuel economy than its class standard (27.5 mpg for passenger cars and 21.6 mpg for light trucks in MY2006) helped a firm comply with the regulation, and any vehicle under the standard hindered a firm's ability to comply. Under the reformed CAFE, any vehicle that has a fuel economy higher than its individual footprint-based target will help a firm comply with the regulation. For example, a firm may prefer to produce a larger vehicle that can exceed its target versus a smaller vehicle that has higher fuel economy but does not exceed its target. Also, because domestic manufacturers tend to have larger vehicles than their competitors, the footprint-based standards allow domestic manufactures to meet a lower standard than European or Asian manufacturers.

Both the unreformed and the reformed CAFE regulations provide some flexibility to meet the fuel economy standards. Specifically, both regulations allow firms to bank and borrow fuel economy credits. This allows a firm to meet the standard in a given year by applying any available banked credits earned from exceeding the standard in previous years or by borrowing credits, which will have to be repaid in future years. In addition to this, the reformed CAFE allows a trading program of credits within each firm, between the fuel economy and light truck

[17] This decision was based on a National Association of Science report which raised concerns that the CAFE regulation encouraged production of smaller vehicles, and that smaller vehicles were more unsafe for the public (NAS 2002, 24; and dissent to this opinion, app. A). NHTSA responded to these concerns by defining the reformed CAFE standards as a function of the footprint (track width multiplied by wheelbase) of the vehicles in a manufacturer's fleet.

standard, as well as among firms. An upper limit of credit trading was set at 1 mpg through 2013 and 1.5 mpg through 2017.

While we do account for trading of fuel economy credits within firms in our policy simulations, we do not permit trading between firms or banking and borrowing credits in our simulations. Therefore, our results should be interpreted as upper bounds on the producer surplus losses resulting from the regulation. This is consistent with our assumptions in the endogenous attribute model, which is constructed to be a conservative representation of possible producer options to respond to the policy. For a discussion of banking and borrowing of fuel economy credits, interested readers should refer to Jacobsen (2010).

6.2 The constraints imposed by reformed CAFE

Consistent with historical behavior, auto firms are characterized into two groups: those that operate at the CAFE standard (constrained), and those that can violate the standard and pay the corresponding fine. Similar to Jacobsen (2010) and Klier and Linn (2008), our policy simulations account for heterogeneity in the compliance behavior of firms, distinguishing between firms that are constrained to meet the CAFE standards and those that can violate the standards and instead pay a fine.

Although both the unreformed and reformed CAFE regulations allow manufacturers to violate the standards and pay corresponding fines that are proportional to the number of miles per gallon under the standard, there is evidence that domestic manufacturers should be treated as though they are constrained to meet the standards . Historically, domestic firms have always met the CAFE standards within allowable levels (Jacobsen 2010). These firms have stated that they view CAFE as binding, believing that they would be liable for civil damages in stockholder suits were they to violate the standards (Kleit 2004). In contrast, many European firms, such as BMW and Audi, have chosen to violate the standards and pay the fines many times, so we do not model them as constrained in our simulations. It is difficult to know whether other foreign firms that have historically met the standards, such as Toyota and Honda, would choose to violate the higher reformed CAFE standards if it were more profitable.

We account for this heterogenous treatment of the CAFE standards in our policy simulations. BMW, Jaguar, Mercedes, Porsche, and VW are allowed to violate the standards if it is more profitable to pay the corresponding fines than to comply with the regulations. All other firms are constrained to meet the CAFE standards.

The optimization problem solved by a constrained firm is to maximize profit subject to meeting the CAFE standards (*stand*$_{PC}$ and *stand*$_{LT}$) for their fleet of cars, $\mathfrak{J}_c$, and their fleet of light trucks, $\mathfrak{J}_T$. The regulatory constraint in equation (1) can be more accurately formulated as:

$$\max_{acc_j, tech_j, p_j \forall j} \sum_j q_j(p_j)\left(p_j - c_j\right)$$

$$\begin{aligned} subject\ to\!: \quad & \textstyle\sum_{j \in \mathfrak{J}_C} q_j(p_j)\, \mathrm{r_{PCj}} \geq 0 \\ & \textstyle\sum_{j \in \mathfrak{J}_T} q_j(p_j)\, \mathrm{r_{LTj}} \geq 0 \end{aligned}$$

where $\mathrm{r_C}$ is $1 - stand_{PC}/mpg_j$ if $j \in \mathfrak{J}_{PC}$ and zero otherwise; and similarly $\mathrm{r_T}$ is $1 - stand_{LT}/mpg_j$ if $j \in \mathfrak{J}_{LT}$ and zero otherwise.

For firms able to violate the CAFE penalty, the profit maximization problem is given by:

$$\max_{acc_j, tech_j, p_j \forall j} \sum_j q_j(p_j)\left(p_j - c_j\right) - F_{PC} - F_{LT}$$

where F_{PC} and F_{LT} are the respective fines if the firm violates either the passenger car or light truck standard:

$$F_{PC} = 55 \sum_{j \in \mathfrak{J}_{f,PC}} q_j(p_j) \left(stand_{PC} - \frac{\sum_{j \in \mathfrak{J}_{f,C}} q_j(p_j)}{\sum_{j \in \mathfrak{J}_{f,C}} q_j(p_j)/mpg_j} \right)$$

$$F_{LT} = 55 \sum_{j \in \mathfrak{J}_{f,LT}} q_j(p_j) \left(stand_{LT} - \frac{\sum_{j \in \mathfrak{J}_{f,T}} q_j(p_j)}{\sum_{j \in \mathfrak{J}_{f,T}} q_j(p_j)/mpg_j} \right)$$

7 Policy Simulations

Having estimated the parameters of the model we introduced in Section 5, we can use the model to simulate firms' response to the reformed CAFE standards that will be applied to the MY2014 fleet of vehicles.[18]

[18] An underlying assumption of these, or any, counterfactuals is that the structure of decision-making is unaffected by the policy change. Because our supply model is constructed from physics-based simulations and we

7.1 Policy scenarios

Twenty firms are represented in these simulations, producing a total of 473 vehicle model and engine options. This represents all vehicle models and engine options in MY2006, which is a considerably larger scale than previous studies (e.g., Goldberg 1998; Jacobsen 2010).

We simulate the effect of replacing the unreformed CAFE standards with the reformed standards under two sets of assumptions. First, we account for the full set of medium-run responses in our model: technology implementation, tradeoffs between fuel efficiency and acceleration performance, pricing adjustments, and trading fuel economy credits between the passenger-car and light-truck standards within a firm. We will refer to this subsequently as the "medium-run full design response" scenario. Second, we shut down the ability of firms to trade off acceleration performance with fuel efficiency, treating acceleration performance as exogenous, but allowing technology implementation, pricing adjustments, and trading of fuel-economy credits. We do this to assess the relative importance of the tradeoffs between acceleration performance and fuel efficiency, which are not included in many analyses of CAFE (e.g., Austin and Dinan 2005, NHTSA 2008). We refer to this as the "medium-run partial design response" scenario.

One important distinction between these two scenarios is that, in the case where tradeoffs between acceleration performance and fuel efficiency are not considered, firms can only improve a vehicle's fuel efficiency implementing technology features that raise the marginal cost of the vehicle. In the case where acceleration tradeoffs are considered, firms can implement these technology features but can also improve a vehicle's fuel efficiency by reducing its acceleration performance, which *reduces* marginal costs (smaller engines are generally cheaper than larger ones). This distinction proves important in the policy simulations, as discussed in the sections below.

A natural extension of this policy analysis would involve a comparison between the aforementioned sets of results and a scenario in which all design responses are shut off. More precisely, we could follow the convention of some earlier studies and consider only the price response to the CAFE policy, treating all aspects of vehicle design exogenous. However, for

---

have no indication that demand would be directly impacted by the change in CAFE, this assumption is justifiable. One possible caveat, however, is that firms may have an incentive to allow adjustments of vehicle footprint later in the development process because the regulation allows manufacturers of larger vehicles to meet lower standards.

many firms, none of the vehicle models produced exceeds the corresponding 2014 fuel economy target. In other words, most firms cannot meet the 2014 CAFE standards through pricing adjustments alone. Consequently, simulations that accommodate only price responses are not informative from an applied policy perspective. This highlights the importance of explicitly accounting for a design response in an analysis of the reformed CAFE standards.

Our policy simulations predict the partial-equilibrium price, marginal production cost, fuel economy, acceleration performance, level of technology implementation, and demand for every vehicle. Taken together, these simulation results can be used to calculate profits and consumer surplus. Note that this analysis does not account for any indirect benefits associated with reduced fuel consumption, such as reduction in environmental damages.

We first discuss the results of the first specification in Section 7.2, which represents the full complement of medium-run design and pricing decisions, and then contrast these results to the analysis that ignores tradeoffs between fuel economy and acceleration performance in Section 7.3.

7.2 Medium-run full design and pricing response

When the full complement of medium-run design responses are accommodated, simulation results indicate that the sales-weighted average[19] fuel economy among compliant firms increases by 3.5 mpg under the reformed CAFE standards. This increase in fuel economy is achieved to a large extent by compromising acceleration performance and implementing technology features although firms also use mix-shifting to comply with CAFE. Results indicate that the sales-weighted average 0-60 acceleration time under the reformed CAFE standards is 2.7 s slower than under the unreformed CAFE standards. These results are consistent with Knittel's (2009) finding that meeting the reformed CAFE standards will require a non-trivial "downsizing" of vehicle performance attributes, such as acceleration, but is clearly attainable. Technology implementation is also significantly affected by the new CAFE standards; cylinder deactivation is included in 10% of vehicle sales and all other technology features considered are included in over 50% of vehicle sales.

[19] All fuel economy averages in section 7 are sales-weighted harmonic means; all other averages are sales-weighted arithmetic means.

Conceptually, we can decompose the aggregate improvement in fuel economy into a design response-induced component and a price response-induced component. To do this, we simulate vehicle demand using the vehicle attributes and technology features obtained in our "full medium-run" simulations and the vehicle prices associated with the pre-policy-change baseline. This effectively holds vehicle prices constant and allows us to isolate the effect of the design response on fuel economy. We find that 88% of fuel economy improvements derive from changes to vehicle designs; the remaining 12% derive from price changes. These results highlight the importance of explicitly accounting for design responses as well as price responses.

Although mix-shifting accounts for a smaller percentage of fuel-economy improvements, it causes a number of important changes. The simulations indicate that firms have an incentive to adjust prices such that production shifts toward light trucks, as opposed to passenger cars, because the fuel-economy targets for these vehicles are lower. Results show that the average price of a light truck reduces by $1,500 and the average price of a passenger car increases by $1,000. Accordingly, the share of light trucks increases 11%.

These policy simulations suggest that there is a heterogeneous response to CAFE across firms. Among the firms that, by assumption, comply with the CAFE standards average fuel economy increases by 3-5 mpg. In contrast, the average fuel economy across fine-paying firms *decreases* by 0.3 mpg on average. This behavior is also noted by Jacobsen (2010). Intuitively, when the CAFE regulation is introduced and the compliant firms improve the fuel efficiency of their vehicles, the residual demand curve for lower fuel economy vehicles, which are larger or have better acceleration, shifts out. Therefore, non-compliant firms have an incentive to improve the acceleration performance of their vehicles by compromising fuel economy. This leads to a form of leakage, where efficiency improvements among compliant firms are offset somewhat by increased demand of vehicles produced by non-compliant firms. Our results indicate that the market-share of compliant firms decreases by 4% so that the sales-weighted average fuel economy across all new vehicle sales is only 2 mpg. We also find that 2% of consumers choose the outside good instead of purchasing a new vehicle. This could also have implications for fleet fuel efficiency because used vehicles typically have lower fuel efficiency than newer vehicles.

It is important to emphasize that, unlike leakage problems that have been characterized elsewhere in the literature (e.g. Fowlie, 2010), the leakage we observe in our simulations can be readily mitigated through policy design. Specifically, the leakage effect would decrease if the fines that noncompliant firms are required to pay were increased. It should be noted that the level of this fine, \$5.5 per vehicle sold per 0.1 mpg below the standard, has not been changed since the original CAFE standards were created in 1975. As Shiau et al. (2009) concluded, if it is desirable to encourage firms to meet a high fuel-economy standard, the fine for violating this standard must also be increased.

Based on the design and pricing responses of firms to the reformed CAFE standards, point estimates of the effects of the policy on economic surplus (i.e., the sum of producer and consumer surplus) are shown in Table 9. All values are measured relative to a baseline equilibrium with respect to vehicle prices, acceleration performance, and technology implementation in the presence of the unreformed CAFE standards. Details of this baseline can be found in Appendix D.

Table 9: Impact of MY2014 CAFE on Economic Surplus Using 1st Specification

Full Medium Run Design & Pricing Analysis

Change in Consumer Surplus	
Total (billions)	-\$266
Constrained firms	-\$377
Fine-paying firms	\$111
Change in Producer Welfare	
Total (billions)	-\$4
Constrained firms	-\$22
Fine-paying firms	\$18

Notes: This table reports point estimates of changes in producer and consumer surplus resulting from replacing the unreformed CAFE standards with the reformed CAFE standards accounting for price changes, technology implementation, and tradeoffs between fuel economy and acceleration performance.

The simulation results indicate that profits of constrained firms decrease by \$22 billion as a result of the 2014 CAFE standards. This represents approximately \$1,400 per vehicle sold. Results also indicate that consumer surplus decreases by \$266 billion. The loss in consumer surplus is largely due to decreases in acceleration performance as well as increases in the prices of passenger cars. These estimates are upper bounds of the impact of CAFE on firm profits and consumer surplus due to conservative assumptions in our endogenous design model as described in Section 3. Partial design response and pricing response

In the second counterfactual policy simulation, we shut down the ability of firms to trade off acceleration performance with fuel economy, equivalent to treating the baseline acceleration performance as exogenous. Figure 6 illustrates a comparison of these results compared to the full medium-run specification. As the figure illustrates, technology implementation is much higher in this specification. This is intuitive considering that adding fuel-saving technology features is the only design strategy that firms can use to increase fuel efficiency. Decomposing the aggregate improvement in fuel economy into a design response-induced component and a price response-induced component indicates that using this specification only 14% of gains in fuel economy are due to design changes (namely technology implementation) with the remaining 86% due to price changes.

Figure 6: Changes in Sales-Weighted Average Vehicle Attributes Using Two Specifications

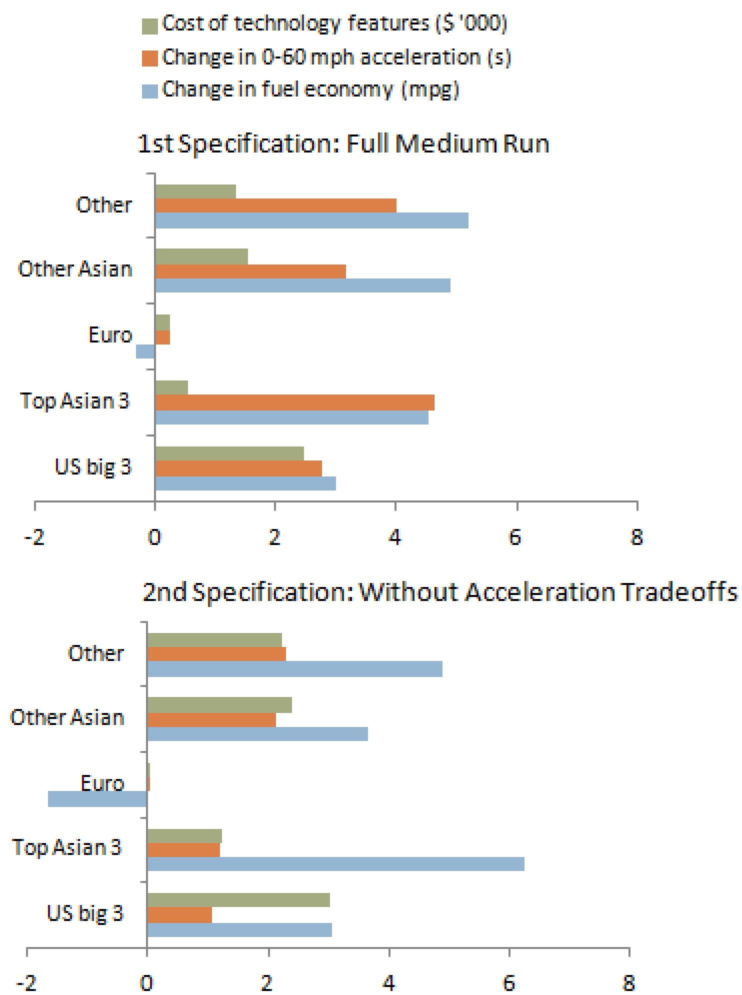

Compared to the first policy scenario that accounts for the full medium-run design response, the leakage effect is larger in this second scenario. The market share of compliant firms decreases 9%, as opposed to 4% in the first specification. As a result, the sales-weighted average fuel economy across all new vehicles is only 1.5 mpg. In addition, 7% of consumers choose the outside good instead of purchasing a new vehicle.

Table 10 reports the point-estimates of the impact of CAFE on producer profits and consumer surplus under this second scenario that considers only the partial design response. Comparing these results to the first specification indicates that the costs of the regulation to compliant firms, in terms of profits lost, are three times larger than estimates produced from the analysis that accounts for tradeoffs between acceleration performance and fuel economy. However, the results suggest that these tradeoffs have the opposite effect on estimates of consumer surplus. Estimates of consumer surplus losses are 20% lower when tradeoffs between acceleration and fuel economy are ignored. These results suggest that analyses of CAFE that do not account for tradeoffs between fuel economy and other vehicle attributes may substantially overestimate the costs of the regulation to firms but may also underestimate the impact of CAFE on consumer surplus.

The smaller profit losses in the first specification compared to the second are intuitive. In the first specification, firms have more options available to them to maximize profits subject to the CAFE standards. Therefore, we would expect that the costs of compliance would be lower as compared to the second scenario which shuts down the ability to trade off acceleration performance with fuel economy. The differences in consumer surplus between the two specifications are somewhat less intuitive. There are a number of competing forces that influence these consumer surplus results. In the second specification, compliant firms rely on technology implementation to a greater extent to meet the CAFE standards. As a result, prices for vehicles produced by these firms are higher, which reduces consumer surplus. However, in the second specification, firms are not allowed to compromise acceleration performance, which increases consumer surplus results compared to the first specification where acceleration performance decreases. Results suggest that this second factor is larger than the first; consumer surplus losses are lower once tradeoffs with acceleration performance are considered. However, due to the sensitivity of these results to the relative impact of acceleration performance, fuel economy, and vehicle price on consumer utility, future work should assess the sensitivity of this particular result to alternative demand-model specifications.

Table 10: Impact of MY2014 CAFE on Economic Surplus Using 2nd Specification

Medium Run Without Acceleration Tradeoffs

Change in Consumer Surplus	
Total (billions)	-$206
Constrained firms	-$322
Fine-paying firms	$116
Change in Producer Welfare	
Total (billions)	-$29
Constrained firms	-$54
Fine-paying firms	$25

Notes: This table reports point estimates of changes in producer and consumer surplus resulting from replacing the unreformed CAFE standards with the reformed CAFE standards accounting for price changes and technology implementation. Unlike the 1st specification, these results ignore tradeoffs between fuel economy and acceleration performance, treating acceleration performance as exogenous.

8 Conclusions

This paper demonstrates the importance of accounting for design responses in the analysis of industrial policy impacts using a case study of medium-run firm responses to the reformed CAFE standards. In addition to accounting for fuel-saving technologies that firms can implement in response to CAFE, our model explicitly accounts for engineering tradeoffs between fuel efficiency and acceleration performance, which provide another mechanism for firms to adjust product designs to respond to energy-efficiency regulations.

Our modeling approach extends the literature on endogenous product attributes in several ways. We use physics-based engineering simulations of vehicles to construct the production possibility frontiers of vehicle fuel efficiency, acceleration performance, and production costs. This allows us to 1) represent the effects of fuel-saving technologies that are not observable in the data but that are potentially optimal given policy counterfactuals, and 2) account for engineering tradeoffs between vehicle attributes, without conflating effects of unobserved attributes. Additionally, we contribute to the endogeneity problem in the demand-side estimation by informing the choice of instruments with documentation of the automotive development process. This literature allows us to identify product attributes that are fixed earlier in the design process than prices and the endogenous attributes of interest, which we use as instruments. The methods we develop here are generalizable to other technical products such as household appliances and consumer electronics.

We use this model to estimate the effects of the model-year 2014 CAFE regulation on producer and consumer surplus and fuel economy. Results indicate that almost 90% of the improvements in fuel-economy are due to changes in product designs. In addition to design changes, compliant firms also adjust product prices to comply with the fuel economy standards, but this has a smaller effect on fuel-efficiency improvements than design changes.

Results highlight the substantial sensitivity of estimates of profit losses and consumer surplus to the product design strategies that firms use to comply with the regulation. When we ignore the potential for tradeoffs between acceleration performance and fuel economy, our results suggest that the profit losses of constrained firms are three times as high as when these tradeoffs are considered. However, consumer surplus losses considering these constraints are 20% higher than when the tradeoffs are ignored. These results suggest that welfare analyses of CAFE or similar policy instruments that ignore the potential for changes in product design decisions could significantly overestimate the impact on firm profits but may also underestimate the impact on consumer surplus losses.

Furthermore, our results illustrate two important factors that can offset improvements in fuel economy under the reformed CAFE standards. Without consumer incentives for higher fuel economy, a number of leakage effects are possible as consumers shift to purchasing vehicles that have lower fuel efficiency. The first factor manifests when some firms choose to violate the CAFE standards and instead pay corresponding fines. Results indicate that the market share of these non-compliant firms increases 4%, which slightly offsets fuel economy gains from compliant firms. The second factor occurs because fuel economy targets for light trucks are lower than passenger cars. This creates an incentive for firms to adjust prices so as to increase the market share of light trucks. In our policy simulations, the market share of light trucks increases 11%. This offsets the fuel economy improvements that would have been realized had market shares been unchanged. These effects could be reduced, respectively, by increasing the fine that noncompliant firms are required to pay and by narrowing the difference between fuel economy targets for passenger cars and light trucks.

The results of our policy analysis are sensitive to our modeling assumptions and the parameter values we use to calibrate the policy simulation modeling. Future work will test the robustness of our results to alternative estimates of engine and technology-feature costs and to alternative specifications of our consumer utility function.

9 References

Alcott, H. and N. Wozny. 2009. Gasoline Prices and the Fuel Economy Discount Puzzle. Princeton University.

Anderson, S. and J. Sallee. 2009. Using loopholes to reveal the marginal cost of regulation: the case of fuel-economy standards. Michigan State University. Working paper.

Beresteanu, A. and S. Li. 2008. Gasoline prices, government support, and the demand for hybrid vehicles in the U.S. Duke University. Working paper.

Borenstein, Severin, James B. Bushnell, and Frank A. Wolak. 2002. "Measuring Market Ineffiencies in California's Restructured Wholesale Electricity Market." *American Economic Review*, 92(5): pp. 1376-1405.

Berry, S. 1994. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics* 25(2): 242-262.

Berry, S., J. Levinsohn, and A. Pakes. 1995. Automobile Prices in Market Equilibrium. *Econometrica* 63: 841–890.

Braess, H.-H. and U. Seiffert, ed. 2005 *Handbook of automotive engineering.* Warrendale, PA: SAE International.

Bresnahan,T. and P. Reiss. 1989 .Dealer and manufacturer margins. *Rand Journal of Economics.*16(2): 253-268.

Fan, Y. 2008. Market structure and product quality in the U.S. daily newspaper industry. Yale University. Working paper.

Fowlie, M., C.R. Knittel and C. Wolfram. 2012. "Sacred Cars? Cost-effective Regulation of Stationary and Non-stationary Pollution Sources." *American Economic Journal: Economic Policy.* Forthcoming.

Frischknecht, B. and P. Papalambros. 2008. A Pareto Approach to Aligning Public and Private Objectives in Vehicle Design. ASME IDETC, DETC2008-49143.

Gholap, A.K. and J.A. Khan. 2007. Design and multi-objective optimization of heat exchangers for refrigerators. *Applied Energy*. 84(12):1226–1239.

Goldberg, P. K. 1998. The effects of the corporate average fuel efficiency standards in the US. *The Journal of Industrial Economics* 46(1): 1–33.

Gramlich, J. 2008. Gas Prices and Endogenous Product Selection in the U.S. Automobile Industry. Yale University. Job Market Paper.

Grossman, G. M. and A. B. Krueger. 1995. Economic Growth and the Environment. *Quarterly Journal of Economics*. 110(2): 353-77.

Helfand, G. and A. Wolverton. Evaluating the Consumer Response to Fuel Economy: A Review of the Literature. National Center for Environmental Economics. Working Paper #09-04.

Jacobsen, M. 2010. Evaluating U.S. Fuel Economy Standards In a Model with Producer Household Heterogeneity. Stanford University. Working Paper.

Klier, T. and J. Linn. 2008. New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard. Federal Reserve Bank of Chicago. Working Paper No. 2008-13.

Kleit, Andrew. 1990. The Effect of Annual Changes in Automobile Fuel Economy Standards. *Journal of Regulatory Economics*, 2(2):151-172.

Kleit, Andrew. 2004. Impacts of Long-Range Increases in the Fuel Economy (CAFE) Standards. *Economic Inquiry.* 42(2): 279-294.

Knittel, Chris. Automobiles on Steroids: Product Attribute Trade-offs and Technological Progress in the Automobile Sector. *American Economic Review*. Forthcoming.

Knittel, C. and K. Metaxoglou. 2008. Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings. Working Paper.

Levinson, A. 2009. "Technology, International Trade, and Pollution from U.S. Manufacturing," *American Economic Review* 99(5) pp. 2177-92, December 2009.

Mayer, B. 2008. Fuel saving under development. *Automotive News Europe.* International Supplier Fair, Oct. 27, 2008. Crain Communications GmbH.

Michalek, J., Papalambros, P. Y., and Skerlos, S. 2004. A Study of Fuel Efficiency and Emission Policy Impact on Optimal Vehicle Design Decisions. Journal of Mechanical Design. 126(6): 1062-1070.

Nevo, A. Mergers with differentiated products: the case of the ready-to-eat cereal industry. *Rand Journal of Economics.* 31(3):395-421.

NHTSA. 2008. Preliminary Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2011-2015 Passenger Cars and Light Trucks. U.S. Department of Transportation.

Seim, K. 2006. An Empirical Model of Firm Entry with Endogenous Product-Type Choices. Stanford University. Working Paper.

Shiau,C-S., J. Michalek and C. Hendrickson. 2009. A structural analysis of vehicle design responses to Corporate Average Fuel Economy policy. *Transportation Research Part A*.43:814-28.

Small, K. and K. Van Dender. 2007. Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect. Energy Journal. 28(1): 25-51.

Sörenson, D. 2006. *The automotive development process:A real options analysis.*Frankfurt: Deutscher Universitäts-Verlag.

Stavins, R. N. 2008. A meaningful U.S. cap and trade system to address climate change. *Oxford Review of Economic Policy*. 24(2): 298–321.

Sweeting, A. 2007. Dynamic product repositioning in differentiated product markets: The case of format switching in the commercial radio industry. NBER Working Paper #13522

Train, K. and C. Winston. 2007. Vehicle choice behavior and the declining market share of U.S. automakers. *International Economics Review.* 48(4): 1469–1496 .

U.S. Environmental Protection Agency, 2004, "EPA's Fuel Economy and Emissions Programs" http://www.epa.gov/fueleconomy/420f04053.htm.

Weber, J. 2009. *Automotive development processes: Processes for successful customer oriented vehicle development*. New York: Springer.

Wolfram, Catherine D. 1999. "Measuring Duopoly Power in the British Electricity Spot Market." *American Economic Review*, 89(4): 805{826.

Wright, J.A., H.A. Loosemore, and R. Farmani. 2002. Optimization of building thermal design and control by multi-criterion genetic algorithm. *Energy and Buildings.* 34:959–972.

Appendix A Engineering vehicle simulations using AVL Cruise

AVL Powertrain Engineering, Inc. (AVL) is an independent company, founded in 1948 and headquartered in Austria, specializing in the development of powertrain systems, simulation methods, and engine instrumentation and test systems. The vehicle simulation software Cruise, developed by AVL, is commonly used by automotive original equipment manufacturers to aid in powertrain development (Mayer, 2008). Cruise simulates vehicle-driving performance, fuel consumption, and emissions based on kinematic calculations.

Cruise models the physical dynamics that occur between subsystems in a vehicle, which translate inputs from a driver into motion of the vehicle. For example, as Figure A1 shows, the Engine module is physically connected to the modules making up the transmission, which include the Torque Converter, Gear Box, Final Drive, and Differential modules. The Combustion Engine module calculates the fuel consumption, speed, and torque of the engine based on user inputs, such as fuel consumption maps, and input information from other vehicle subsystems, including the load on the acceleration pedal from the Cockpit (driver) module and the external temperature from the Vehicle module. It then transmits information about the torque and speed to the transmission modules.

Table A1: Screen shot of the AVL Cruise simulation interface

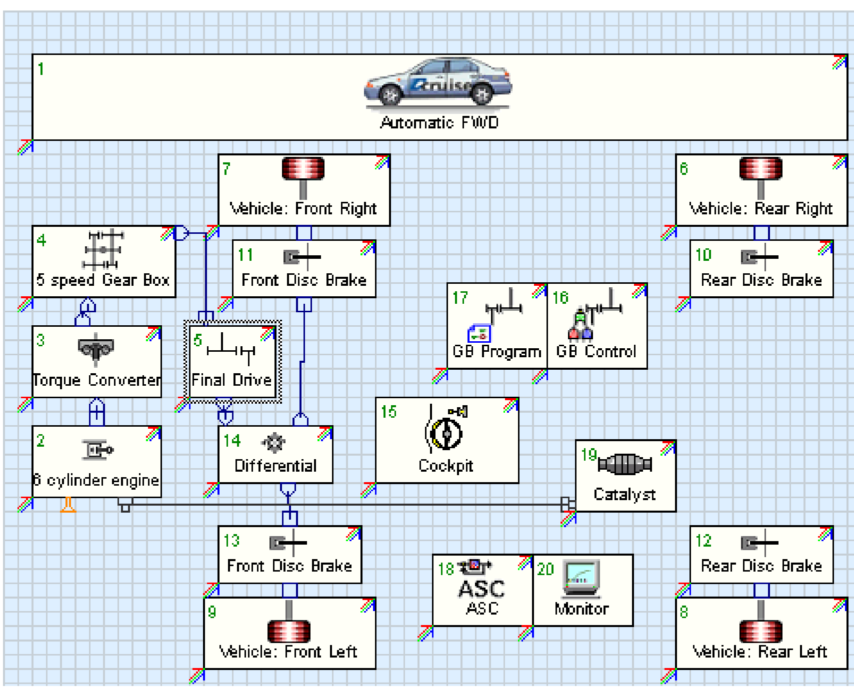

The modular structure of Cruise allows researchers to simulate multiple vehicle architectures by customizing the subsystem modules (e.g., front or rear wheel drive, automatic or manual transmissions), and modifying various input parameters. For example, with the Vehicle module, a user can adjust the aerodynamic drag coefficient of the vehicle body and the curbweight of the vehicle.

Using Cruise, a total of 29,575 vehicle simulations were conducted. Design input parameters are varied at small intervals so that we can observe the influence of each of these parameters and their interactions on attributes of interest (i.e. acceleration performance and fuel efficiency). Table A2 summarizes the range of parameter values we consider in our analysis. These include the powertrain variables that can be changed in the medium run (i.e., engine displacement and final drive ratio) as well as longer-run design decisions that are continuous (i.e., curbweight), which we condition on in the supply-side model.

Table A2: Ranges and intervals of vehicle simulation parameters

Vehicle Class	Base Vehicle	Displacement			Curbweight			Final Drive		
		Min	Max	Int.	Min	Max	Int.	Min	Max	Int.
2seater/Mini	Ford Mustang	1,000	8,200	400	1,900	4,300	800	1.0	6.0	0.4
Sub/Compact	Honda Civic	1,000	4,200	400	2,200	4,800	800	1.0	6.0	0.4
Midsize	Toyota Camry	1,000	4,200	400	2,400	4,800	800	1.0	6.0	0.4
Fullsize	Ford Taurus	1,600	6,800	400	3,200	5,400	400	1.0	6.0	0.4
SUV	Ford Explorer	2,000	8,400	400	3,200	8,000	800	1.0	6.0	0.4
Small pickup	Toyota Tacoma	1,600	8,400	400	2,800	5,200	800	1.0	6.0	0.4
Stand. pickup	Ford F150	2,000	8,400	400	3,200	6,600	400	1.0	6.0	0.4

Notes: This table lists the min, max, and interval of input parameters used in the "AVL Cruise" vehicle simulations. Engine *Displacement* is in cm^3, *Curbweight* is in lb, and *Final Drive* is the final drive gear ratio. All other input parameters for the simulations (e.g., front-wheel drive) were taken using data for the "base vehicle".

All other vehicle parameters are determined from a representative base vehicle for each class.[20] Many of these parameters (e.g., front-wheel drive) are determined prior to the medium-run decisions we are interested in, but for some parameters (e.g., transmission gear ratios), it is possible that they could be modified in the same time period. In these cases, any potential bias in our counterfactual results caused by holding these design parameters fixed will be toward overestimating negative impacts of the CAFE regulations on producers and consumers.

[20] The classes are based on the EPA segment classifications, with some grouping of segments based on similar ranges of engine displacement, final drive ratios, and curbweight as well as similar predicted outputs from AVL Cruise.

NHTSA (2008) estimated the effect of each technology feature listed in Table A3 on fuel economy, in terms of the percentage improvement, based on values reported by automotive manufacturers, suppliers, and consultants. We use these estimates to determine how the baseline iso-technology curve changes with the addition of one or more technology features. To do this we also need to know the impact of each technology feature on 0-60 acceleration time, which is not reported by NHTSA. We determine these impacts by simulating each technology feature in AVL Cruise to a level that matches the improvement in fuel economy reported by NHTSA. For example, NHTSA reports a 0.5% improvement in fuel economy from using "low friction lubricants" in compact vehicles. We simulate this impact by reducing the friction losses in the engine of our representative compact vehicle model until we observe fuel economy improving by 0.5% and then observe the percentage improvement of 0-60 acceleration time. When NHTSA provided a range of fuel economy improvement for a technology feature, the lower bound of this range is used, consistent with our other assumptions in creating a conservative engineering design model. The results of these simulations are reported in Table A3.

Table A3: Impact of technology features on fuel economy and 0-60 mph acceleration time

	Two Seater		Compact		Mid/Minivan		Fullsize		Fullsize		Fullsize		Fullsize	
Technology	% mpg	% acc	% mpg	% acc	% mpg	% acc	% mpg	% acc	% mpg	% acc	% mpg	% acc	% mpg	% acc
Low friction lubricants	0.5	0.3	0.5	4.1	0.5	0.6	0.5	0.6	0.5	1.6	0.5	0.4	0.5	0.7
Engine friction reduction	1	0.3	1	5.6	1	1.5	1	1.2	1	3.1	1	0.7	1	1.5
Aggressive shift logic	1	-0.2	1	-5.0	1	-0.2	1	-0.3	1	0.0	1	-0.2	1	-2.8
Early torque converter lockup	0.5	-	0.5	-	0.5	-	0.5	-	0.5	-	0.5	-	0.5	-
High efficiency alternator	1	0.3	1	5.6	1	1.5	1	1.2	1	3.1	1	0.7	1	1.5
Aerodynamic drag reduction	3	0.3	3	5.1	3	0.5	3	0.3	3	1.4	2	0.5	2	0.4
Low rolling resistance tires	1	0.1	1	2.5	1	0.2	1	0.1	1	0.5	1	0.2	1	0.1
Cylinder deactivation	n/a		n/a		4.5	-	4.5	-	4.5	-	4.5	-	4.5	-

Appendix B Details about the technology combinations and the *tech* variable

The specifications for equations 3 and 4 were chosen after examining the relationship between the discrete technology feature combinations with cost and fuel economy. For example, Figure B1 below plots production cost against the *tech* variable assigned to each cost-effective combination of technology features conditional on 0-60 mph acceleration time. Each point on the plot represents a potential vehicle design with an engine size, final drive ratio, and set of discrete technology features. The gaps between vehicle designs achieving the same acceleration time is an artifact of the ranges of input variables used in the AVL Cruise vehicle simulations. We would expect that as the intervals of these input variables approached zero, the gaps would disappear.

Figure B1: Relationship of ordered technology feature combinations to production cost conditional on 0-60 mph acceleration time

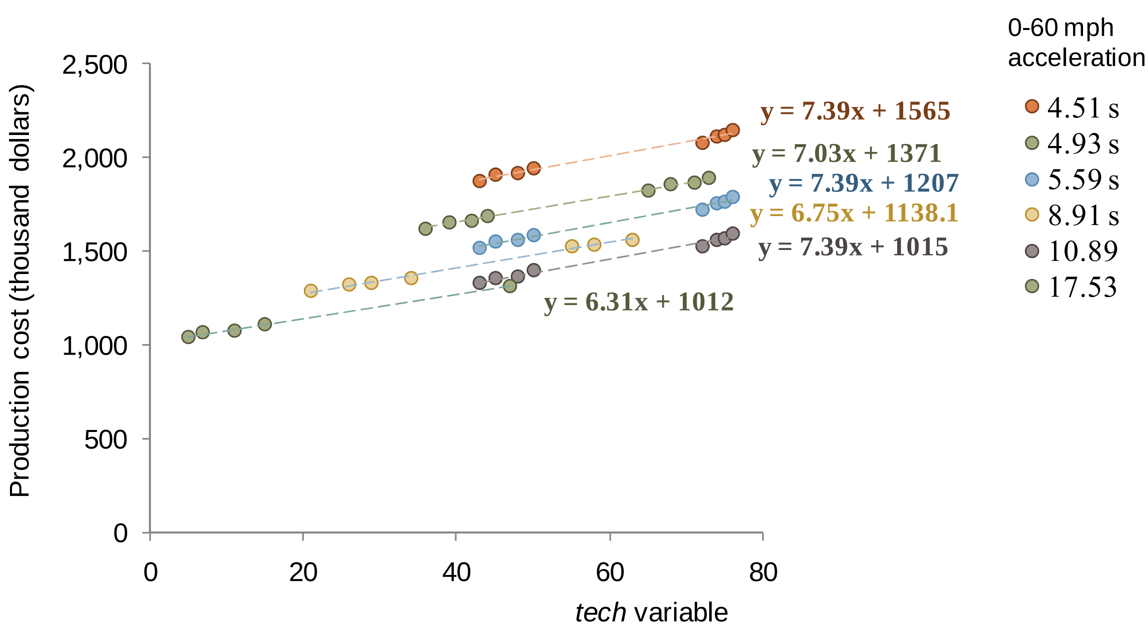

Figure B1 is generated for a specific vehicle segment (an SUV) and a specific curbweight (3,200 lb). Similar trends were found for other segments and other curbweights. The figure indicates that conditional on vehicle segment, curbweight, and 0-60 acceleration time, moving "up the line" of combinations of technology features increases cost linearly. It also indicates that the incremental change in cost of changing technology features is roughly constant across the various levels of acceleration performance. This structure is preserved in the specification of equation 4 where cost is linear in technology conditional on vehicle segment, curbweight, and acceleration time.

Figure B2 below plots fuel consumption against combinations of technology features for the same vehicle segment (SUV) and curbweight (3,200 lb). This figure indicates that conditional on vehicle segment, curbweight, and 0-60 acceleration time, the set of combinations of technology features linearly decrease fuel consumption. However, unlike cost, the incremental change in cost of changing technology features varies across the various levels of 0-60 mph acceleration time. Figure B2 shows that the incremental decrease in fuel consumption from moving to a higher ordered combination of technology features becomes larger as acceleration time becomes faster. Also, the rate of this change increases as acceleration time gets faster, with the slopes in Figure B2 roughly constant when acceleration time is relatively large but more negative for relatively faster acceleration times. Similar trends were found for other segments and other curbweights. These properties are represented in the specification of equation 3 by including a linear *tech* term as well as an interaction term multiplying *tech* by *acc* squared.

Figure B2: Relationship of ordered technology feature combinations to fuel consumption conditional on 0-60 mph acceleration time

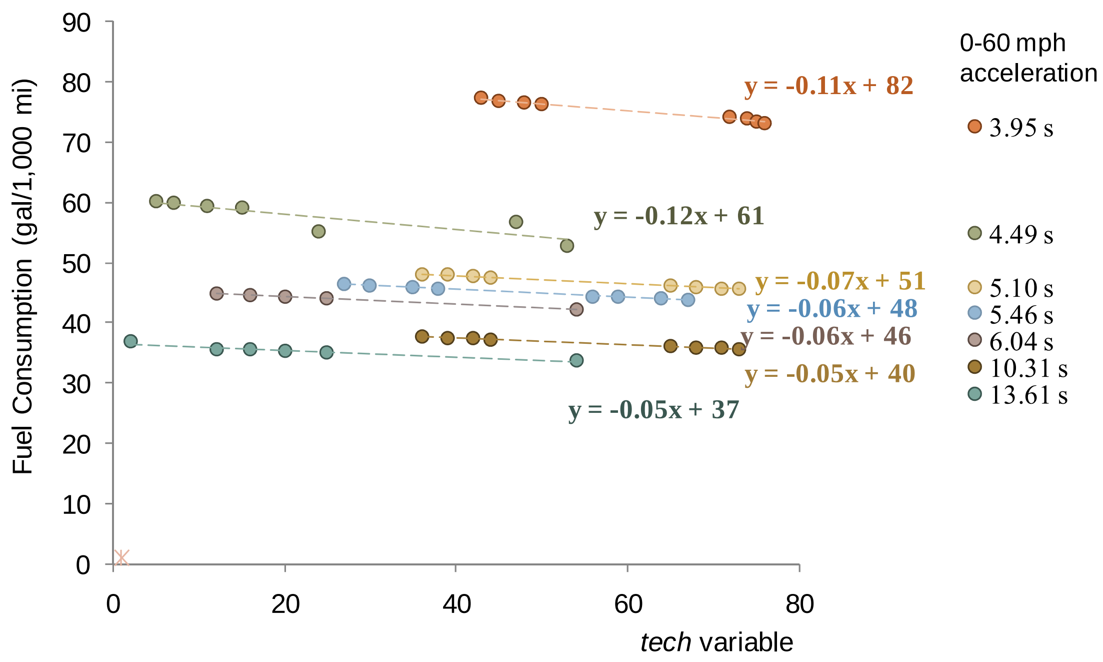

Appendix C Estimation of Production Cost Parameters

For all firms, the marginal cost of producing automobile j is represented as:

$$c_j = engcost_j + \omega_j \tag{C1}$$

The variable $engcost_j$ represents the portion of marginal cost dependent on the endogenous selection of attributes as described in Section 3.3. The remaining portion of marginal cost, ω_j, is determined from the first order conditions of Bertrand equilibrium following Jacobsen's (2010) procedure. Firms' profit maximization problems follow the standard Bertrand equilibrium but are also subject to the CAFE regulations. Similar to Jacobsen, we distinguish between American firms, who behave as though they are constrained to the CAFE standards, and European firms who often violate the CAFE standards and pay corresponding penalty fines. Asian firms are treated similarly to European firms but, because the Asian firms exceed the CAFE standards over the time period in the data, the constrained and unconstrained formulations are equivalent.

The optimization problem solved by a constrained firm is to maximize profit subject to meeting the CAFE standards ($stand_C$ and $stand_T$) for their fleet of cars, $\mathfrak{J}_C$, and their fleet of light trucks, $\mathfrak{J}_T$, as defined in equation C2. In this equation, q_j, and p_j are respectively the quantity sold and price of vehicle j, r_C is $1 - stand_C/mpg_j$ if $j \in \mathfrak{J}_C$ and zero otherwise; and similarly r_T is $1 - stand_T/mpg_j$ if $j \in \mathfrak{J}_T$ and zero otherwise.

$$\begin{aligned} &\max_{p_j \forall j} \sum_j q_j(p_j)\,(p_j - c_j) \\ &subject\ to: \quad \sum_{j \in \mathfrak{J}_C} q_j(p_j)\,\mathrm{r_{Cj}} \geq 0 \\ &\qquad\qquad\qquad \sum_{j \in \mathfrak{J}_T} q_j(p_j)\,\mathrm{r_{Tj}} \geq 0 \end{aligned} \tag{C2}$$

For firms able to violate the CAFE standards, the profit maximization problem is given by:

$$\max_{,p_j \forall j} \sum_j q_j(p_j)\,(p_j - c_j) - F_C - F_T \tag{C3}$$

where F_C and F_T are the respective fines if the firm violates either the passenger car or light truck standard:

$$F_C = 55 \sum_{j \in \Im_{f,C}} q_j(p_j) \left(stand_C - \frac{\sum_{j \in \Im_{f,C}} q_j(p_j)}{\sum_{j \in \Im_{f,C}} q_j(p_j)/mpg_j} \right)$$
$$F_T = 55 \sum_{j \in \Im_{f,T}} q_j(p_j) \left(stand_T - \frac{\sum_{j \in \Im_{f,T}} q_j(p_j)}{\sum_{j \in \Im_{f,T}} q_j(p_j)/mpg_j} \right) \quad \text{C4}$$

Therefore, the first order conditions of these two optimization problems can be written in vector notation as in equations C5 and C6:

$$\textit{fine-paying:} \qquad \boldsymbol{q}(\boldsymbol{p}) + \nabla_{\mathbf{p}} \boldsymbol{q}^{\boldsymbol{T}} (\boldsymbol{p} - \boldsymbol{c}) - \nabla_{\mathbf{p}} \boldsymbol{F_C} - \nabla_{\mathbf{p}} \boldsymbol{F_T} = \mathbf{0} \quad \text{C5}$$

$$\textit{constrained:} \qquad \boldsymbol{q}(\boldsymbol{p}) + \nabla_{\mathbf{p}} \boldsymbol{q}^{\boldsymbol{T}} (\boldsymbol{p} - \boldsymbol{c} - \lambda_C \boldsymbol{r}_C - \lambda_T \boldsymbol{r}_T) = \mathbf{0} \quad \text{C6}$$

Given data on vehicle sales and prices, and estimates of the cross-price elasticities, $\nabla_{\mathbf{p}} \boldsymbol{q}^{\boldsymbol{T}}$, from the demand model, the vehicle costs for fine-paying firms can be directly determined from equation C5. However because the Lagrange multipliers, λ_C and λ_T, are unknown and $\boldsymbol{r}_C$ and $\boldsymbol{r}_T$ depend on fuel economy, which is correlated with marginal cost, we cannot directly solve for marginal cost for the firms constrained to the CAFE standard. The Lagrange multipliers, which are negative, represent the effect on firm profits of incrementally increasing the constraints in equation C2 holding vehicle design fixed. If we assume that the Lagrange multipliers are zero then we would overestimate the marginal costs of vehicles with fuel economies below the standard and underestimate the costs of vehicles that exceed the standard.

Following Jacobsen (2010), we estimate these multipliers using the relationship of dealer markups to manufacturer markups. Specifically, there is evidence that dealer markups, b_j, for each vehicle j are a fixed percentage of manufacturer markups (Bresnahan and Reiss 1989):

$$\boldsymbol{b} = \gamma(\boldsymbol{p} - \boldsymbol{c} + \boldsymbol{\varepsilon}) \quad \text{C7}$$

Substituting in equation C6, we can relate dealer markups to the Lagrange multipliers:

$$\boldsymbol{b} = \gamma \left(-\left(\nabla_{\mathbf{p}} \boldsymbol{q}^{\boldsymbol{T}} \right)^{-1} \boldsymbol{q} + \lambda_C \boldsymbol{r}_C + \lambda_T \boldsymbol{r}_T + \varepsilon_j \right) \quad \text{C8}$$

Using data on dealer markups and parameters estimated in the demand model, we can obtain estimates for the Lagrange multipliers and then solve for the equilibrium vehicle costs for constrained firms from equation C6. This estimation has the disadvantage of relying on the imposed form of the relationship between dealer and manufacturer markups. However, our interest in the estimates of λ_C and λ_T is limited to their role in controlling for the correlation of marginal vehicle costs with r_C and r_T. We conduct sensitivity analyses of marginal cost estimates to the estimates of λ_C and λ_T and find that this sensitivity is low.

Data on dealer transactions purchased from JD Power and Associates are used to estimate the Lagrange multipliers according to equation C8. These data were collected from approximately 6,000 dealers from the proprietary Power Information Network data, aggregated to quarterly invoice costs and transaction prices for each vehicle model.

Table C1 shows the estimates of the effect of incrementally increasing the regulatory constraints in equation C2 (i.e., λ_C and λ_T) on domestic firm profits. Recall that these constraints are represented as a nonlinear function of the CAFE standards, and therefore these estimates do not directly correspond to an incremental increase in the CAFE standards. Using point estimates of λ_C and λ_T, we calculate the corresponding impact of incrementally increasing the passenger car and light truck standards of the unreformed CAFE regulation on firm profits, shown in Table C2. These profit losses, or shadow costs, of increasing the CAFE standards are in the range of those estimated by Anderson and Sallee (2009).

The estimates indicate that, for passenger cars, Chrysler faces a higher cost of compliance than Ford or GM. This result is intuitive given that in MY2006, Chrysler neither offered many small vehicles nor had any passenger cars with a fuel economy higher than 26 mpg. For light trucks, the estimates indicate that Ford has the lowest cost of compliance and GM has the highest. This result can be explained by the fact that, while Ford produces fewer models of light trucks than GM, Ford produces a number of high-efficiency light trucks including the Escape Hybrid.

Table C1: Lagrange Multiplier Estimators

		I	II	III
Chrysler	λ_C	250.5	863.76***	861.06***
	λ_T	340.98**	355.70***	351.29***
Ford	λ_C	690.95**	733.56***	715.59***
	λ_T	520.84***	55.33	58.92
GM	λ_C	1042.98***	768.31***	760.75***
	λ_T	762.03***	739.42***	736.28***
Fixed Effects		None	Manufacturer and class	Manufacturer class, and quarter
Obs.		708	708	708
R^2		0.278	0.689	0.9973

legend: * p<.05; ** p<.01; *** p<.001

Notes: This table presents the impacts of incrementally increasing the constraints in equation C2 on firm profits. Because the constraints are nonlinear functions of the CAFE standard, these values are not the shadow costs of the regulation, but the shadow costs can be derived from these estimates as shown in Table C1.

Table C2: Shadow Costs Estimates of Unreformed CAFE Regulation

		profit losses (millions)	losses per vehicle
Chrysler	passenger cars	\$19.849	\$41
	light trucks	\$30.654	\$19
Ford	passenger cars	\$27.905	\$31
	light trucks	\$5.649	\$4
GM	passenger cars	\$52.122	\$31
	light trucks	\$92.015	\$41

Notes: This table presents the impacts of incrementally increasing the unreformed CAFE passenger car and light truck standards on firm profits. These values were determined from point estimates of the third specification of the Lagrange multipliers presented in Table C1.

The sensitivity of production cost estimates to the estimates of λ_C and $\lambda_T$was assessed by increasing the value of these Lagrange multipliers for each firm by 10% and observing the change in production costs. Table C3 reports the results of these tests, indicating that the absolute value of changes in production costs are between less than \$0.01 to approximately \$63, with a mean absolute value change of \$1–\$4 depending on the firm. The maximum absolute change in these estimates due to increasing the Lagrange multipliers is less than 1%. We further tested the effect of completely ignoring the effect of the shadow costs of CAFE on production cost estimates, setting each value of λ_C and λ_T to zero. This test indicated that the absolute value changes in production costs are between less than \$0.01 and \$170, with a mean absolute value change of \$5.60. These results suggest that sensitivity of cost estimates to estimates of Lagrange multipliers are low and so we do not expect any errors in the Lagrange multiplier estimates to significantly affect any counterfactual results using this cost model.

Table C3 Sensitivity of production costs to estimates of Lagrange multipliers

	mean absolute change	min absolute change	max absolute change
Chrysler	\$1.05	< \$0.01	\$15.57
Ford	\$4.06	\$0.19	\$25.36
GM	\$3.98	< \$0.01	\$62.68

Appendix D Endogenous Attribute Model

We compare the results of our approximated endogenous attribute model, fit using simulation data, to market vehicle data as a validation test. Figure D1 and D2 show this comparison for the compact and SUV segments, respectively. Similar comparisons were done for all other segments. The midsize, fullsize, small pickup, and large pickup/van segment comparisons are comparable to those shown in the figures. The two-seater model, however, does not predict the market data as well as the other segments, but because this segment represents less than 1% of sales, this is unlikely to affect our counterfactual results.

Figure D1: Comparison of compact-segment model to MY2006 compact vehicle data

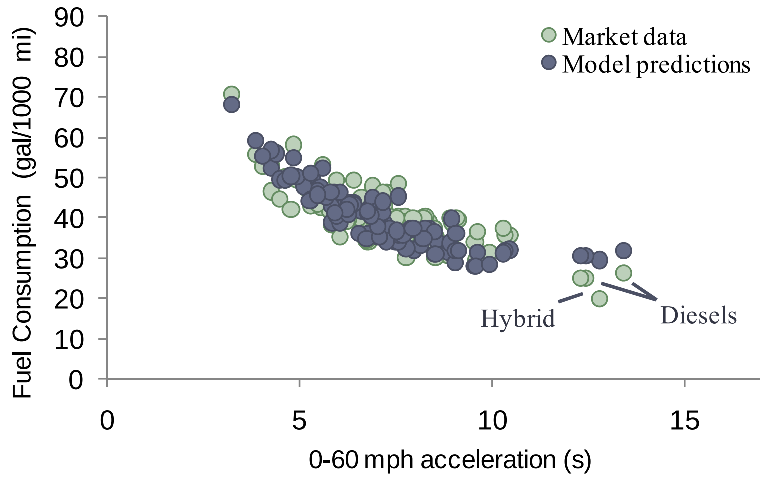

Figure D2: Comparison of SUV-segment model to MY2006 SUV data

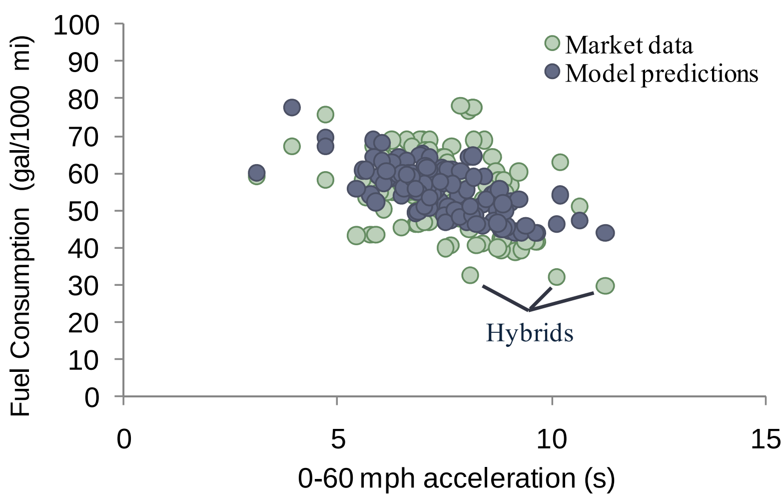

Appendix E Counterfactual baseline

Because our endogenous attribute model is derived from engineering simulations and cost data, observed attributes are not necessarily restricted to be in equilibrium. We therefore perform simulations of the CAFE regulations that were applied to the MY2006 automotive market. The simulation results are in Nash equilibrium with respect to firm decisions on 0-60 acceleration time, and technology implementation—which implicitly determines fuel economy—as well as price for each of their vehicles. All other counterfactual results are measured from this baseline. Although the technology features we consider weren't necessarily available to the same extent in 2006, we account for them in this "baseline" equilibrium so that our policy simulations isolate the effect of the 2014 CAFE standards from additional effects caused by the availability of additional technology features. Additionally, this baseline scenario does not allow for the banking and borrowing of fuel-economy credits, as described in Section 7.

Table E1: Comparison of observed attributes and baseline simulation

	Sales-Weighted Average Fuel Economy (mpg)	Sales-Weighted Average Acceleration (s)	Sales-Weighted Average *tech*
Observed	21.5	7.56	-
Baseline	24.5	5.30	5.14