ENERGY INSTITUTE AT HAAS

EI @ Haas WP 214

Product Design Responses to Industrial Policy: Evaluating Fuel Economy Standards Using an Engineering Model of Endogenous Product Design

Kate Whitefoot, Meredith Fowlie, and Steven Skerlos

February 2011

Product Design Responses to Industrial Policy: Evaluating fuel economy standards using an engineering model of endogenous product design*

Kate Whitefoot† University of Michigan
Meredith Fowlie‡ University of California, Berkeley
Steven Skerlos University of Michigan

February 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 industrial policy impacts. As a case in point, we model automotive firms' medium-run compliance choices under the reformed Corporate Average Fuel Economy (CAFE) regulation. Physics-based simulations are used to characterize the potential for improving fuel efficiency through medium-run 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 perform counterfactual simulations of firms' pricing and medium-run design responses to the reformed CAFE regulation. Results indicate that compliant firms rely primarily on changes to vehicle designs to meet the CAFE standards, with a smaller contribution coming from pricing strategies designed to shift demand towards more fuel-efficient vehicles. Simulations that account for medium-run design responses yield estimates for compliance costs that are nine times lower than simulations that only account for price changes. Moreover, our simulations suggest that as compliant firms increase fuel efficiency, it becomes more profitable for non-compliant firms to adjust product designs to improve other performance attributes at the expense of fuel efficiency. This effect has the potential to significantly offset the fuel efficiency improvements achieved by compliant firms.

---

* 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 estimations, 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.

† Address:2350 Hayward St, Ann Arbor, MI 48109. Email: kwhitefoot@umich.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 durable goods. In 2007, Congress created efficiency standards for many household appliances, including dishwashers and furnaces. In 2009, the State of California adopted efficiency regulations for consumer products, such as battery chargers and televisions. The Department of Energy just recently announced new efficiency standards for refrigerators and clothes washers. One especially noteworthy effort to reduce energy consumption, enacted first by Congress and then by the Obama administration, raises fuel efficiency standards for new automobiles to 35 mpg by 2017. This mandates a more than 30% reduction in fuel consumption per mile. How firms respond to these types of policies can have significant implications for how efficiently energy-intensity reductions are achieved and who bears the costs.

When an energy-efficiency standard is introduced, firms typically have a variety of compliance strategies to choose from. They can adjust the relative prices of their products in order to shift demand toward their more efficient products. Alternatively, they can modify the designs of their products so as to increase the energy efficiency.[1] In much of the economics literature that investigates industry response to regulatory intervention, the abilities of firms to change product designs is underemphasized (e.g., Goldberg 1998; Nevo 2000; Jacobsen 2010). However, energy-efficiency regulations are typically announced a number of years before they become mandatory, providing firms the opportunity to respond to these policies through product design changes. Recent work on the automotive industry indicates that engineering design decisions have played a significant role in determining fleet fuel-efficiency trends, including gains under CAFE (Knittel 2009; Klier and Linn 2008).

In this paper, we develop an empirically tractable approach to incorporate design responses into economic analyses of industrial policies. To accomplish this, we draw from the engineering design literature. In many of the industries targeted by energy-efficiency policies, product design decisions are subject to a suite of engineering constraints and tradeoffs that play an important role in firm responses to policy interventions. These constraints and tradeoffs have

[1] These strategies are analogous to Grossman and Krueger's (1995) concepts of composition, technique, and scale as a basis for understanding links between changes in economic conditions and emissions at the country-level. We note that this framework can also be used to characterize firm- or industry-level responses to regulations.

been studied in detail and explicitly represented in engineering models of product design (e.g., Frischknecht and Papalambros 2008; Gholap and Khan 2007; Wright et al. 2002).

Our analysis focuses on the medium-run design modifications that firms can use to comply with energy-efficiency regulations. These design changes can be generally classified into two categories. "Technology implementation" involves choosing from a range of technology features (e.g., cylinder deactivation in vehicles), which may not be currently implemented in the market and improve energy efficiency at some additional cost. "Attribute tradeoffs" improve energy efficiency with some loss of another vehicle attribute (e.g., acceleration performance), but do not necessarily increase costs. Both classes of design changes have a potentially significant role in firms' responses to the energy-efficiency standards. To our knowledge, our analysis is the first to explicitly model both types of design changes.[2]

This approach is used to analyze automotive firms' medium-run options to respond to the reformed Corporate Average Fuel Economy (CAFE) regulation. Our analysis begins by highlighting some relevant features of the automotive design process. We first identify the aspects of vehicle design that must be chosen in the early stages of the design process. Conditioning on these longer run decisions, we define the production possibility frontiers (PPF) that capture the ability of firms to trade off fuel efficiency with acceleration performance in the later stages of the design process. To specify these relationships, we use physics-based vehicle simulation models that are used in the automotive industry to support the powertrain development process.

These PPFs are then integrated into a partial equilibrium, static oligopoly model of the automotive industry to analyze how firm design and pricing decisions are influenced by the reformed CAFE standards. The demand side of the partial-equilibrium model consists of a mixed-logit representation of consumer demand, estimated using disaggregate consumer data. We use the partial-equilibrium model to simulate firms' responses to replacing the unreformed CAFE standards with the model-year (MY) 2014 reformed CAFE standards under multiple sets of assumptions. First, we account for the full complement of medium-run policy responses and

[2] Purely-econometric studies can identify tradeoffs between product attributes and historical trends of technology implementation, but cannot incorporate technology features that could be implemented but are not observable in data. NHTSA (2009) uses engineering models to account for potential technology features that are not commonly implemented in the automotive market, but these analyses ignore any ability to trade off fuel efficiency with other product attributes.

then gradually restrict the ability of firms to exercise these responses, shutting down the ability to trade off vehicle attributes and then the ability to implement technology features.

Results highlight the importance of explicitly accounting not only for price responses but also for both types of design responses to fuel efficiency regulations. When we ignore the potential for tradeoffs between acceleration performance and fuel economy, our results suggest that the cost of the regulation to constrained firms is \$406 per short ton of CO_2 reduced.[3] When tradeoffs between acceleration performance and fuel economy are also considered, these costs fall to \$197. Furthermore, the simulations indicate that neither Chrysler nor Ford can meet the MY2014 reformed CAFE standards without adjusting vehicle designs.

Additionally, simulations that account for both price and design responses suggest that, among firms choosing to violate the standard and pay the fine, average fuel economy may *decrease* in response to the regulation. This leads to a form of leakage, where efficiency improvements from complying firms are offset by reductions in fuel efficiency from non-compliant firms. These findings not only highlight the adverse consequences of non-compliance but also further underscore the importance of explicitly accounting for design responses; simulations that consider only price changes imply much lower levels of leakage.

This paper is germane to the emerging literature on endogenous product design (e.g., Seim 2006; Sweeting 2007; Fan 2008). Much of the existing work on modeling endogenous attribute selection has focused on relatively straightforward design processes. The kinds of products that are typically targeted by energy efficiency standards, such as automobiles and household appliances, present additional challenges because the process of attribute selection includes many engineering constraints and tradeoffs. By integrating engineering simulations together with econometric models of consumer demand, we are able to explicitly account for the engineering relationships that define these more technologically complex product design processes. Our approach builds upon recent work by Klier and Linn (2008) who simulate the medium-run response to unreformed CAFE, coupling a demand-side estimation with an

---

[3] These costs assume a vehicle lifespan of 13 years, a baseline of 14,000 annual vehicle miles traveled (the average in 2006 as reported by the Department of Transportation) and a rebound effect of 10.3% (Small and Van Dender 2007).

econometric model of tradeoffs between endogenous vehicle attributes on the supply side.[4] Access to engineering vehicle simulations provides a relatively straightforward way to incorporate key features of the engineering design process into our endogenous attribute model.

The design relationships investigated in this paper are very similar to those analyzed by Knittel (2009). Knittel econometrically estimates the tradeoffs that automotive manufacturers face between the fuel economy, weight, and engine power of vehicles sold in the United States over the period 1980 and 2006. He documents both movements along and shifts in the PPF of these three vehicle attributes. This paper investigates very similar design relationships, but in lieu of using bundles of attributes observed in the market place, we use the outputs of physics-based engineering simulations to construct vehicle PPFs. For our purposes, this engineering-based approach confers two advantages. First, many combinations of product attributes are not observed in the market to date, but are technologically feasible and potentially optimal under the counterfactual policy instruments of interest. These unobserved but possible combinations of attributes are captured by the engineering simulations. Second, engineering tradeoffs between product attributes can be identified independent of unobserved product attributes. As we explain below, this latter advantage is instrumental to our identification strategy.

This paper also extends the literature that seeks to estimate the effects of fuel economy standards on producer and consumer welfare (e.g., Anderson and Sallee 2009; Jacobsen 2010; Goldberg 1998). A defining challenge in much of this literature pertains to estimating demand parameters in the presence of correlation in observed prices and endogenous attributes with unobserved attributes. In previous work, some researchers have used functions of non-price attributes of other vehicles as instruments (e.g., Berry et al. 1995; Train and Winston 2007). One criticism of this approach is that firms presumably choose these non-price attributes and prices simultaneously. To address this issue, we exploit the well-documented structure of the automotive design process to identify vehicle attributes that are determined in earlier stages than the endogenous attributes of interest. Our key identifying assumption is that powertrain architecture (e.g., hybrid), drive type (e.g., all-wheel-drive), and major vehicle dimensions are

[4] 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 electric accessories that can further impact fuel economy, are often correlated with observed attributes such as horsepower or weight.

chosen earlier in the development process than detailed design variables in the powertrain that affect both fuel efficiency and acceleration performance. We present evidence in support of this assumption in Section 2.2.

The remainder of the paper is organized as follows. Section 2 provides an overview of the CAFE regulation and industry response, including a description of the vehicle development process. Section 3 details the simulations and estimation of engineering tradeoffs between endogenous attributes. Data used in the demand- and supply-side estimations are described in Section 4 and the basic model is presented in Section 5. Section 6 presents and discusses the empirical results and Section 7 discusses the counterfactual simulations. Conclusions are described in Section 8.

2 Overview of CAFE and Firm Response

2.1 Development of Fuel Economy Regulations

Since 1975, the CAFE policy has influenced automotive firms' decisions by setting a minimum standard for the average fuel economy of a manufacturer's fleet of vehicles sold in the United States. The principle motivation for Congress to create the CAFE regulation was to reduce dependence on oil consumption in the wake of the 1973-74 oil embargo. Since that time, interest in maintaining and strengthening the regulation has been driven by concerns about global climate change as well as dependence on foreign oil.

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). Using this particular formulation, a doubling of this average fuel economy corresponds with halving fuel consumption, assuming the same number of miles driven. In order to comply with the CAFE policy, this average must be greater than or equal to the CAFE standard, such that:

$$\frac{\sum_{j\in\mathfrak{J}_{f,c}} q_j(p_j)}{\sum_{j\in\mathfrak{J}_{f,c}} q_j(p_j)/mpg_j} \geq stand_c \tag{1}$$

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. 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).

The CAFE standards in place over the period 1975-2008 established a significantly lower standard for light trucks than for passenger cars. This distinction allowed minivans and SUVs, which composed a very small fraction of sales when the policy was introduced, to meet the lower light-truck standard despite their expanding role as a personal vehicle, giving rise to the so-called "SUV loophole". In 2007, Congress passed the Energy Independence and Security Act (EISA), phasing out this disparity by setting a target standard for both vehicle classes of 35 mpg by MY2020, later moved up to MY2016 by President Obama's administration.

In addition, Congress modified the design of the CAFE standard. The reformed CAFE establishes an individual fuel economy target, T_j, for each vehicle, based on vehicle footprint such that vehicles with larger footprints have lower standards.[5] 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:

$$stand_{f,c} = \frac{\sum_{j \in \Im_{f,c}} q_j(p_j)}{\sum_{j \in \Im_{f,c}} q_j(p_j)/T_j} \tag{2}$$

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

This change has a number of important implications. With the unreformed CAFE, 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

---

[5] 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.

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 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.

2.2 Vehicle Development Process

The design response of an automotive manufacturer to a fuel efficiency standard depends substantially on the structure of the vehicle development process. This process is a structured sequence of interrelated decisions, many of which constrain choices made at later stages (Sörensen 2006). The typical design process begins with a 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 a target catalog specifying 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 development 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, it is possible to identify a stage in current automotive development processes where vehicle segment, powertrain architecture (e.g., conventional gasoline, hybrid, diesel), and major dimensions are finalized but "tuning" of powertrain variables is still possible.

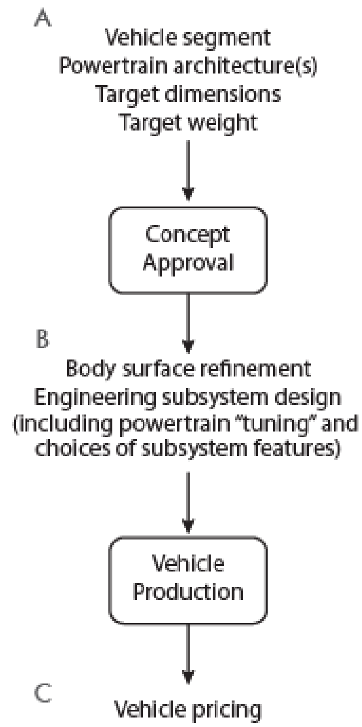

Figure 1: Simplified representation of an automotive development process illustrating short-run (Stage C), medium-run (Stage B) and longer-run (Stage A) design decisions

The structure of the automotive development process informs our research design in two important ways. First, we take as given the vehicle segment, powertrain architecture, and other vehicle dimensions that are determined in the earlier stages of the design process (Stage A in Figure 1). Conditional on these features and attributes, we model manufacturers' choice of fuel economy and acceleration performance (Stage B) and vehicle pricing strategies (Stage C).[6] Second, we use the variation in vehicle attributes determined in earlier stages of the design

[6] 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.

process (i.e., Stage A) to instrument for endogenous variables in our demand-side estimation, which are determined later in the design process (Stage B).

3 Endogenous Attribute Choice

Credible modeling of endogenous attribute selection in the design of technical products such as automobiles requires accurate representation of engineering and economic tradeoffs. We cannot directly observe all of the tradeoffs that firms make during different stages of the vehicle design process. We can, however, generate detailed engineering estimates of the tradeoffs that play an important role in determining vehicle fuel efficiency in medium-run design decisions.

Taking a "bottom up" approach, we construct the medium-run PPFs using detailed engineering simulations along with data of production costs and fuel-savings potential of various technology features. To explain this procedure, we first discuss the medium-run design decisions that are most relevant to our analysis. We then describe the vehicle simulation model, the simulations and cost data used in our analysis, and the estimation procedure to construct the PPFs. Finally, we conclude with a discussion of how our engineering approach to modeling this endogenous attribute choice contrasts with the econometric approaches that are more often used in the economics literature.

3.1 Medium-run design decisions affecting fuel efficiency

Our analysis of firms' response to the CAFE regulations focuses on medium-run vehicle design decisions and short-run vehicle pricing decisions (Stage B and C in Figure 1). At this point in the vehicle development process, many major parameters of the vehicle have been determined including the segment of vehicle, key internal and external dimensions, and the powertrain architecture (e.g., conventional gasoline, hybrid, and diesel). The automotive manufacturer can still adjust the fuel efficiency and acceleration performance at this point in the design process by "tuning" a number of variables in the powertrain (e.g., engine displacement and final drive ratio) and including technology features (e.g., a high efficiency alternator). 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.

Another means of improving fuel efficiency at this stage in the design process involves incorporating various extra “technology features” to the 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 increases the cost of vehicle production. Depending on the specific technology features chosen, acceleration performance may increase, decrease, or not be affected at all. For example, cylinder deactivation of the engine can improve fuel efficiency by effectively decreasing the size of the engine but, because it only is active during coasting, it will not affect acceleration. Reducing aerodynamic drag of the vehicle body can improve both fuel efficiency and acceleration performance, whereas early shifting logic can improve fuel efficiency but will reduce acceleration performance. Note that these technology features only affect demand through their influence on fuel efficiency and acceleration performance; they do not have intrinsic value to the consumer.

Although our goal is to identify the continuous iso-cost PPFs that define the relationship between fuel economy, acceleration performance, and production costs dependent on these medium-run design decisions, vehicle simulations and automotive data suggest that these PPFs are not in fact continuous. To illustrate this, again consider a given vehicle design such as the Honda Accord. The Accord could be described by its position on the two dimensional fuel-efficiency vs. acceleration-performance space. Honda can decrease the fuel consumption of the Accord without adding any additional technology features by trading off acceleration performance, which could be represented by the Accord moving along an “iso-technology curve” as in Figure 2. Considering that Honda could move along this curve in large increments by replacing the engine, smaller increments by decreasing the displacement size of the existing engine, or fine increments by adjusting the electronic control unit, approximating these possibilities as continuous is reasonable. However, incorporating technology features into the vehicle to increase fuel efficiency often causes discrete shifts in vehicle attributes. These discrete shifts of the iso-technology curves in the fuel-efficiency vs. acceleration-performance space cause discontinuities in the iso-cost PPFs as illustrated in Figure 2.

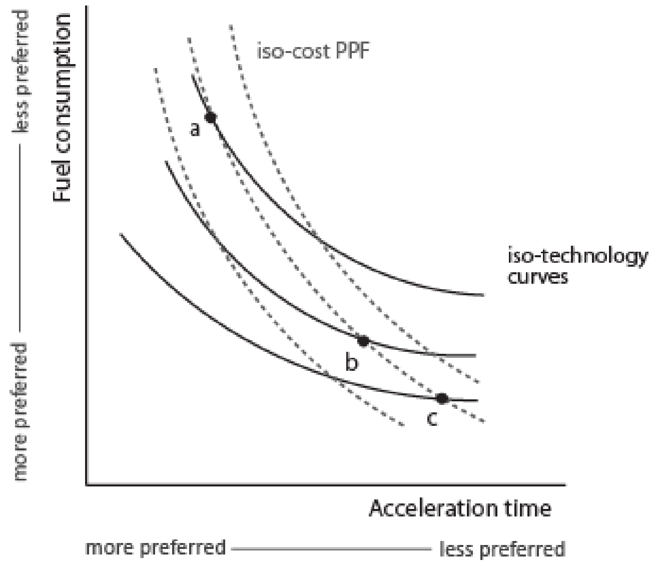

Figure 2: Illustration of discontinuities in iso-cost PPFs (a to b to c) due to the discrete effect technology features have on possible attribute combinations

Ideally, we may like to model the discontinuities in the iso-cost PPFs caused by discrete technology options. However, because of the large number of discrete combinations of technology features, further described in Section 3.4, this is computationally infeasible. We address this challenge by approximating the effect of the technology features as a continuous variable. To do this, we first construct the iso-technology curves for each combination of technology features, then order the technology-feature combinations by the position of their corresponding iso-technology curves. We then approximate the technology features as continuous in the counterfactual simulations to construct the iso-cost PPFs.

3.2 Engineering modeling of medium-run design

We use detailed engineering simulations to construct a "baseline" iso-technology curve. This represents the engineering tradeoffs between fuel consumption and 0-60 acceleration time for a vehicle with no extra technology features. To determine the baseline iso-technology curve, we employ detailed simulation models that have been developed in the private sector to assist the engine and drivetrain industry with product design and development. Specifically, we use "AVL Cruise", which is an integrated simulation platform that is widely used by major automotive manufacturers in powertrain development (Mayer 2008), to characterize the engineering tradeoffs between fuel efficiency and 0-60 acceleration time for each vehicle class. The fuel efficiency of a particular vehicle design is determined in AVL Cruise by simulating the EPA's fuel-economy test procedures. Acceleration performance is determined by simulating a shifting program of the vehicle from standstill to 60 mph. Additional details of the vehicle simulations are discussed in Appendix A.

Our next step is to determine how the addition of one or more technology features affects the position of the iso-technology curve relative to the baseline. To accomplish this, we combine the AVL Cruise simulations and data from NHTSA (2008). We consider only a subset of the types of technology features identified by NHTSA in our analysis. The majority of technology features we omit from our analysis are only available in longer run planning stages, but 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.

NHTSA (2008) estimated the effect of each technology feature listed in Table 1 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 endogenous attribute model. The results of these simulations are reported in Table 1.

3.3 Cost of medium-run design decisions

In addition to representing the impact of medium-run design decisions on vehicle attribute performance, we also need to account for the effect of these decisions on vehicle production costs. We use two separate sources of data to estimate these costs, one describing costs dependent on the powertrain tuning variables, which we use to determine costs along the baseline iso-technology curve, and another dataset detailing production costs for each technology feature. The production cost of the baseline iso-technology curve—representing the costs dependent on choices of engine size and final drive ratio without any extra technology features—is taken from Michalek et al. (2004). The authors collected cost data from manufacturing, wholesale, and rebuilt engines of varying displacements. The additional production costs resulting from each technology feature is taken from NHTSA (2008), which are shown in Table 1. These cost data were estimated by NHTSA based on reported values from automotive manufacturers, suppliers, and consultants, and 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. The specific technology features we consider do not require changes in engine design or affect the assembly of the engine with other vehicle subsystems, consistent with our assumption that costs are additively separable, with 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, engine friction can be reduced by using lubricants, the costs of which are independent of all medium-run decisions considered.[7]

3.4 Estimating a tractable model of engineering tradeoffs and costs

Ideally, all of the detailed information about design tradeoffs that are captured by the AVL Cruise model would be incorporated directly in our model of supply-side design and pricing decisions. However, because of the computational time required to execute the vehicle simulations, and the large number of discrete combinations of technology features, this is

[7] 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.

computationally infeasible. Instead, we approximate these relationships with a flexible parametric form.

Taking all possible combinations of technology features gives automotive firms, depending on the vehicle segment, 128 or 256 options to choose from for each vehicle. From this set, we consider only those combinations of 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. Although this reduces the set of technology feature combinations to between 20 and76, depending on vehicle segment, it is still computationally infeasible to model this number of choices per vehicle per manufacturer in our counterfactual simulations, so further simplifications are necessary.

We approximate the discrete choices of technology features 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. The set of cost-effective technology feature combinations is ordered by increasing fuel efficiency (decreasing fuel consumption) for the same acceleration performance, which is also increasing in cost. Therefore, a higher value of *tech* corresponds to a higher fuel efficiency and higher cost vehicle conditional on 0-60 acceleration time. The impact of the continuous approximation on the results is relatively small with the average gap between discrete features less than 1 mpg. Furthermore, we provide some evidence in Appendix B that the particular specification we use to estimate the relationships of the continuous *tech* variable to fuel consumption and cost preserve important properties of the discrete technology combinations.

We use the results of the engineering simulations together with data on the technology features and technology costs to estimate Eq. 3 and 4, which together define the PPFs between vehicle fuel consumption, acceleration performance, and production costs.

$$fuelcons_j = \kappa_{1s} + \kappa_{2s} e^{-acc_j} + \kappa_{3s}\, tech_j + \kappa_{4s} tech_j \cdot acc_j^2 + \kappa_{5s} wt_j + \kappa_{6s} wt_j \cdot acc_j + \varepsilon_j \quad (3)$$

$$w_j = \sigma_{1s} + \sigma_{2s} e^{-acc_j} + \sigma_{3s} tech_j + \sigma_{4s} wt_j + \sigma_{5s} wt_j \cdot acc_j + \nu_j \quad (4)$$

The subscript *j* in Eq. 3 and 4 denotes the vehicle design in the vehicle class *s*.[8]

The dependent variable in Eq. 3 is the fuel consumption of a vehicle, *fuelcons*. The 0-60 acceleration time is denoted *acc* , *wt* is the curb weight of the vehicle, and *tech* is the scalar measure of technological features incorporated. Equation 4 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 curb weight, acceleration time, and the continuous measure of technology features. The terms ε_j and ν_j represent the error associated with approximating the calculations performed in the vehicle simulations with the simplified relationships. Curb weight is included in Eq. 3 and 4 although it is considered fixed in our medium-run analysis. Including curb weight, and estimating the parameters for each vehicle class, conditions the fuel consumption and cost relationships on vehicle parameters that are exogenous in the medium run and greatly improves estimation fit. Several specifications for each equation were tested; the equations above performed the best under the Akaike Information Criterion.

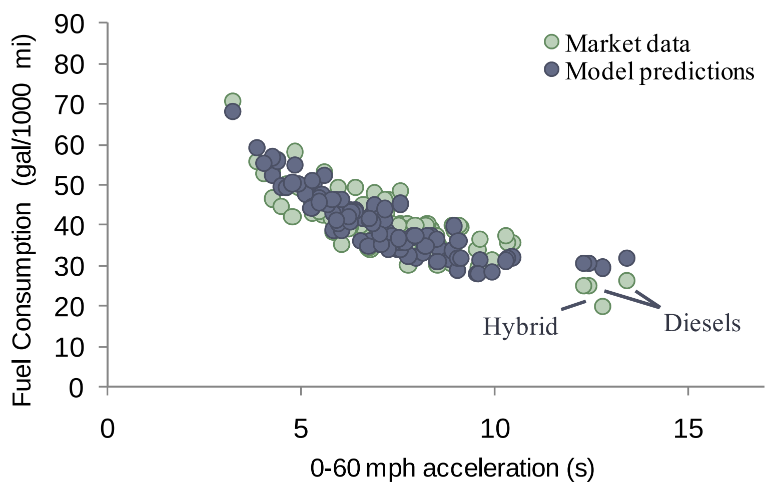

Figure 3: Comparison of compact-segment model to MY2006 vehicle data

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. Eq. 3 predicts observed vehicle performance reasonably well (R^2=0.76). Appendix C discusses comparisons between estimated and observed performance attributes in more detail.

[8] The vehicle design refers to a combination of design parameters input into the simulation, described in Appendix A, and the vehicle class (e.g., minivan).

3.5 An engineering approach to endogenous attribute choice modeling

This is not the first paper to investigate the physical relationships between vehicle attributes that constrain vehicle design decisions. In the economics literature, a few recent studies have developed econometrically estimated endogenous attribute models for the automotive market (Gramlich 2008; Klier and Linn 2008). Although the econometric approach has its advantages, 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 possible, but are not yet observed in the data. This is important because these unobserved vehicle designs may be optimal under the policy regimes we are interested in analyzing. In fact, not only have fuel economy standards historically forced the frontier of vehicle design but manufacturers have also 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 tradeoff between fuel efficiency, acceleration performance, and technology features 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, which are common in the automotive industry. 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 counterfactuals.

4 Data

We employ a combination of household-level data conducted by Maritz Research and vehicle characteristic data available from Chrome Systems Inc. to construct our demand-side estimations. The Maritz Research U.S. New Vehicle Customer Study (NVCS) collects data monthly from 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. Approximately one-third of respondents 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.

Additionally, data on dealer transactions purchased from JD Power and Associates are used to estimate dealer markups in the supply-side model. 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.

5 Model

In this section, we introduce an econometric model of vehicle demand and a supply-side representation of decisions in response to CAFE. Specification and estimation of the demand-side model draws from the seminal work by BLP, and subsequent work by Train and Winston (2007), and others. The distinguishing feature of the demand estimation has to do with our choice of instruments, which is informed by our understanding of the vehicle design process as described in Section 2. Our model of the supply side departs significantly from much of the previous literature insofar as critical medium run design decisions are endogenous to the model.

5.1 Demand side

Demand for new automobiles is modeled using a mixed-logit representation of consumer preferences. The indirect utility V_{nj} that consumer n derives from purchasing vehicle model j is defined as:

$$V_{nj} = U_{nj} + \epsilon_{nj} = \delta_j + \boldsymbol{\beta_n}'\boldsymbol{x_{nj}} + \boldsymbol{\mu}'\boldsymbol{w_{nj}} + \epsilon_{nj} \tag{5}$$

where the model-specific fixed effect, δ_j, represents the portion of utility that is the same across all consumers; $\mathbf{w_{nj}}$ are interactions between vehicle attributes and consumer characteristics that affect utility homogenously across the population; $\mathbf{x_{nj}}$ are vehicle attributes and attribute-demographic interactions that heterogeneously affect utility, entering the equation through draws , $\boldsymbol{\beta_n}$, of normal distributions with standard deviations $\boldsymbol{\sigma}$. Attributes in the vector $\mathbf{w_{nj}}$, which are assumed to have homogeneous effects in the utility specification, include the interaction between living in a rural area and the pickup truck segment, and interactions between children in the household and SUV and minivan segments. Attributes in the vector $\mathbf{x_{nj}}$, which have random effects in the utility specification, are the ratio of vehicle price to income; "gallons per mile", *gpm*, or the inverse of fuel economy; the inverse of 0-60 acceleration time; and vehicle footprint. The disturbance term ϵ_{nj} is the unobserved utility that varies randomly across consumers. Assuming that the disturbance term in Eq. 5, ϵ_{nj}, is independent and identically distributed (iid) Type I extreme value, the probability that consumer n chooses vehicle i over all other vehicle choices j≠i or the outside option[9] takes the form:

$$P_{in} = \frac{e^{U_{ni}}}{1 + \sum_j e^{U_{nj}}} \tag{6}$$

where U_{nj} is the non-stochastic portion of utility of vehicle j for consumer n from Eq. 5. The predicted market share of vehicle i is $\sum_n P_{in}$.

The model specific fixed effects, δ_j, capture the average utility associated with the observed vehicle attributes denoted $\mathbf{z_j}$ and unobserved attributes denoted ξ_j. Vehicle attributes in the vector $\mathbf{z_j}$ include price, fuel consumption, the inverse of 0-60 acceleration time, vehicle footprint, and vehicle segment. The segments are based on the EPA's classes (e.g., minivans). The observable variables of primary interest (namely price, fuel consumption, and acceleration

[9] The utility of the outside option of not purchasing a new vehicle is assumed to be $U_{nO} = \delta_O + \epsilon_{nO}$, where ϵ_{nO} is a draw from an extreme value Type 1 distribution, and δ_O is normalized to zero.

time) are likely to be correlated with unobserved attributes that are simultaneously determined and captured in the error term. Consequently, estimating Eq. 6 directly will yield inconsistent parameter estimates. Following Berry (1994), we move this endogeneity problem out of the non-linear Eq. 6 and into a linear regression framework. This allows us to address the endogeneity problem using well developed two-stage least squares. More precisely, we define the model specific fixed effects to be a function of observed vehicle attributes, $\mathbf{z_j}$, including all endogenous attributes, and the average utility of unobserved attributes, ξ_j:

$$\delta_j = \boldsymbol{\alpha}'\boldsymbol{z_j} + \xi_j \tag{7}$$

Berry (1994) showed that, given a set of values for $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, a unique $\boldsymbol{\delta}$ exists such that predicted market shares match observed market shares. Given a set of exogenous attributes and instrumental variables for endogenous attributes, $\mathbf{y}$, the condition that $\mathbb{E}(y_{j,}\xi_j) = 0$ for all j is sufficient for the instrumental variable estimator of $\boldsymbol{\alpha}$ to be consistent and asymptotically normal conditional on $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.

In related automotive studies (e.g., Berry et al. 1995; Train and Winston 2007), researchers use functions of non-price attributes as instruments, including vehicle dimensions, horsepower and fuel economy. This approach has been criticized due to two concerns: 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.[10] As Heckman (2007) notes, to obtain valid instruments in this context requires a model of the determinants of product attributes. Access to the engineering design literature provides a description of this model.

Availability of literature detailing the automotive development process allows us to limit the choice of instruments to only those attributes determined from longer run product-planning schedules than the endogenous variables, increasing the credibility of these instruments for

[10] 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.

medium-run analyses.[11] Instruments are selected from attributes that can be considered fixed in the medium run as supported by evidence in Section 2.2: 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).

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 $\mathrm{V_{ni}} > \mathrm{V_{nh_1}} \dots \mathrm{V_{nh_m}} > \mathrm{V_{nj}}$ for all $\mathrm{j} \neq \mathrm{i}, \mathrm{h_1}, \dots, \mathrm{h_m}$ 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:

$$L_{nih_1 \dots h_m} = \left(\frac{e^{U_{ni}}}{1+\sum_j e^{U_{nj}}}\right)\left(\frac{e^{U_{nh_1}}}{\sum_{k\neq i} e^{U_{nk}}}\right)\cdots\left(\frac{e^{U_{nh_m}}}{\sum_{l\neq i,h_1,\dots,h_{m-1}} e^{U_{nl}}}\right) \tag{8}$$

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 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 Eq. 8.

Recently, significant concerns have been raised about the sensitivity of parameter estimates using similar random-coefficient discrete choice demand models (Knittel and Metaxoglou 2008). While more work is needed to demonstrate the robustness of our estimation results, initial robustness tests have been reassuring. Specifically, we have estimated the model using a series of randomly selected starting values and found that the estimates converge to the

---

[11] 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.

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

same values within small tolerances. The resulting maximum satisfies both the first and second order conditions, with a zero-gradient within a tolerance of 1e-4 and a negative definite hessian.

5.2 Supply side

The automotive industry is modeled as an oligopoly of multiproduct firms that maximize the expected value of profits with respect to vehicle attributes and prices. 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 counterfactual 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, which 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 the standards instead. First, domestic firms have historically always met the CAFE standards within allowable levels but have never significantly exceeded them (Jacobsen 2010). Second, 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 (Kliet 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 model these firms as choosing whether to meet the standards based on the most profitable option and discuss the effect of this assumption on our simulation results.

In our counterfactual simulations, the domestic Big 3 manufacturers (Chrysler, Ford, and General Motors) are constrained to the standard following their historic behavior, and expected future behavior, of consistently meeting the CAFE standards within allowable banking and borrowing credits (Jacobsen 2010). The remaining firms are allowed to violate the standards if it is more profitable to pay the corresponding fines than to comply with the regulations.

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, $\mathfrak{J}_T$. Rearranging Eq. 1, this formulation can be written as:

$$\max_{acc_j, tech_j, p_j \forall j} \sum_j q_j(p_j)\left(p_j - c_j\right) \tag{9}$$

$$\begin{aligned} subject\ to: \quad & \textstyle\sum_{j \in \mathfrak{J}_C} q_j(p_j)\, \mathrm{r}_{\mathrm{Cj}} \geq 0 \\ & \textstyle\sum_{j \in \mathfrak{J}_T} q_j(p_j)\, \mathrm{r}_{\mathrm{Tj}} \geq 0 \end{aligned} \tag{10}$$

where 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.

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

$$\max_{acc_j, tech_j, p_j \forall j} \sum_j q_j(p_j)\left(p_j - c_j\right) - F_C - F_T \tag{12}$$

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

$$\begin{aligned} F_C &= 55 \sum_{j \in \mathfrak{J}_{f,C}} q_j(p_j)\left(stand_C - \frac{\sum_{j \in \mathfrak{J}_{f,C}} q_j(p_j)}{\sum_{j \in \mathfrak{J}_{f,C}} q_j(p_j)/mpg_j}\right) \\ F_T &= 55 \sum_{j \in \mathfrak{J}_{f,T}} q_j(p_j)\left(stand_T - \frac{\sum_{j \in \mathfrak{J}_{f,T}} q_j(p_j)}{\sum_{j \in \mathfrak{J}_{f,T}} q_j(p_j)/mpg_j}\right) \end{aligned} \tag{13}$$

Note that, for all types of firms, fuel consumption is determined through decisions on acceleration performance and technology features from Eq. 3 and therefore is not listed explicitly as a decision variable. We could instead have listed *mpg* and *tech* as the decision variables and implicitly determined acceleration performance; this convention is arbitrary and does not affect the formulation.

For all firms, the marginal cost of producing automobile j is represented as Eq. 14.

$$c_j = engcost_j + \omega_j \tag{14}$$

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:

$$\textit{fee-paying:}\quad q_j(p_j) + \sum_i \frac{\partial q_i}{\partial p_j}(p_i - c_i - \lambda_C r_{Ci} - \lambda_T r_{Ti}) - \frac{\partial F_C}{\partial p_j} - \frac{\partial F_T}{\partial p_j} = 0 \qquad (15a)$$

$$\textit{constrained:}\quad q_j(p_j) + \sum_i \frac{\partial q_i}{\partial p_j}(p_i - c_i - \lambda_C r_{Ci} - \lambda_T r_{Ti}) = 0 \qquad (15b)$$

For the firms able to violate the standard, the costs can be directly determined from equations 12a and 12b. 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 will be negative, represent the effect on firm profits of incrementally increasing the unreformed CAFE constraints in Eq. 11 holding vehicle design fixed. Notice that if we assumed that the Lagrange multipliers were 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 are a fixed percentage of manufacturer markups (Bresnahan and Reiss 1989):

$$b_j = \gamma(p_j - c_j + \varepsilon_j) \qquad (16)$$

Substituting in Eq. 15b, we can obtain estimates for the Lagrange multipliers:

$$b_j = \gamma(p_j - c_j + \lambda_C r_{Cj} + \lambda_T r_{Tj} + \varepsilon_j) \qquad (17)$$

Using these recovered parameters, we can then solve for the equilibrium marginal vehicle costs for the constrained firms. This estimation does have 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_{Cj} and r_{Tj} such that are estimates for marginal costs are not biased. We therefore expect any errors in the estimation of λ_C and λ_T will only have a small effect on the results of our counterfactual simulations. Future work will include performing sensitivity tests on counterfactual results to various values of λ_C and λ_T.

6 Empirical Results

In this section, we first present estimates of the medium-run PPF parameters for each vehicle class as described in Section 3. We then introduce the demand-side parameter estimates. Finally, we present the estimations of the Lagrange multipliers, described in Section 5, which are used to recover marginal production costs used in the counterfactuals.

6.1 Endogenous attribute 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 Eq. 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 counterfactual 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.

The estimates describing the relationship between production costs and choices of acceleration performance and technology implementation, described by Eq. 4, are reported in Table 4. These estimates fit all vehicle classes reasonably well (R^2>0.83). As expected, these results indicate that production costs increase with the level of technology implementation and decrease with longer 0-60 mph acceleration times. 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.

Taken together, these estimates of Eq. 3 and 4 define the iso-cost PPFs conditional on longer run vehicle-design decisions. These results suggest that significant improvements in fuel efficiency are possible in the medium run but require either substantial tradeoffs with other aspects of vehicle performance or an increase in production costs. To illustrate this, Figure 4 plots the estimated iso-cost PPFs for selected vehicles, indicating that a 10% reduction in fuel consumption can be achieved in many vehicles without increasing production costs by reducing acceleration performance by 1 s or less.

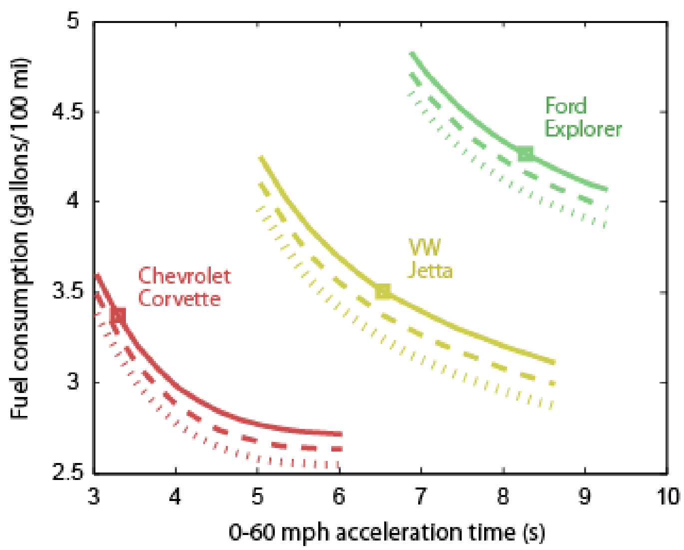

Figure 4: Estimated iso-cost production possibility frontiers for selected vehicles (■ current location,— baseline frontier,- -$100 design changes, ····$200 design changes)

On the other hand, fuel efficiency can be increased without affecting acceleration performance by implementing technology features, which instead raise production costs. Results indicate that fuel economy can be increased by 1 mpg in the majority of vehicle models by implementing $400 or less worth of technology features. Figure 5 in Section 6.2 plots the relationship between technology costs and fuel economy for select vehicle models. This figure illustrates that the required 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.

6.2 Demand-side estimation

Table 5 reports results from estimating Eq. 7. All estimates have the expected signs. Recall that the σ parameters represent the standard deviations of the demand parameters in Eq. 5 that are allowed to vary randomly in the population and assumed normally distributed. 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 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 6 reports the second-stage IV estimates of the parameters in Eq. 5. The OLS estimations of the first-stage regressions of endogenous decisions (price, fuel consumption, and inverse 0-60 mph acceleration time) are presented in Table 7, with F-tests of 20.68, 19.62, and 21.38, respectively. The SUV indicator variable in Eq. 7 is positive and significant, implying that they are preferred more than sedans; and the minivan indicator is negative and significant, implying that they are preferred less than sedans. The parameter estimate for two-seater sports cars is negative and the parameter for pickup trucks is slightly positive, but neither is significant.

Based on these estimates, the average price elasticity is -1.9, (95% CI: -2.0, -1.8), and the sales-weighted average 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 other models of the automotive industry (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.

Using the demand-side estimates, we calculate the expected willingness-to-pay for an increase in fuel economy as illustrated in Figure 5. These results indicate that the willingness-to-pay for fuel economy varies substantially across vehicle models—both because the reduction in fuel consumption due to an increase in 1 mpg varies with the fuel economy of the vehicle and because the consumers that are likely to buy cheaper vehicles are less willing to pay for any improvements to vehicle attributes. However, the willingness-to-pay for fuel economy improvements in any given vehicle 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.[13] This discrepancy of willingness-to-pay for fuel economy with the net-present-value of fuel savings is well documented in other studies (Helfand and Wolverton 2009; Alcott and Wozny 2009).

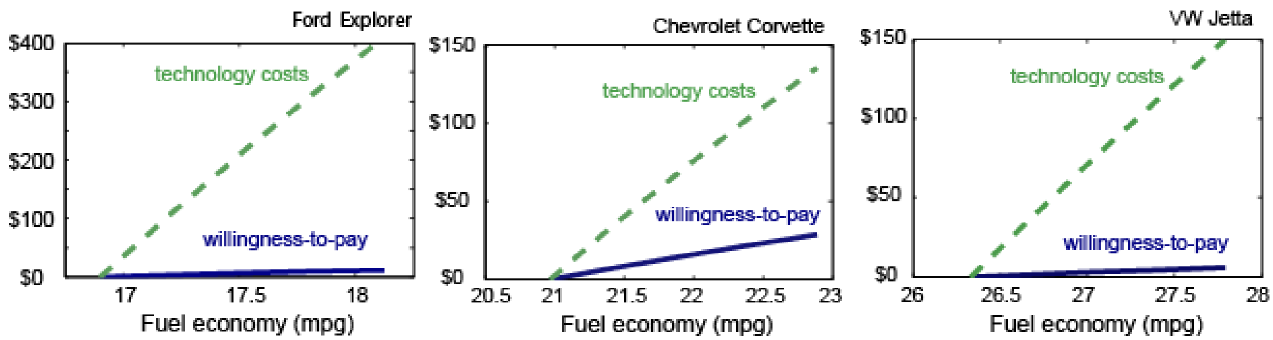

Figure 5: Select vehicles' technology costs and willingness-to-pay to increase fuel economy

Our estimates imply that, in general, consumers' willingness-to-pay for a 1 mpg improvement in fuel economy is lower than their willingness-to-pay for an 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 the higher horsepower vehicle options are willing to pay more for acceleration performance relative to other consumers, and much less for fuel economy improvements.

6.3 Supply-side estimation

Table 8 shows the estimates of the effect of incrementally increasing the regulatory constraints in Eq. 10 (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 9. 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,

[13] 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.

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.

7 Counterfactual Experiments

Having estimated the parameters of the model we introduced in Section 5, we can use the model to perform counterfactual simulations of the reformed CAFE standards that will be applied to the MY2014 fleet of vehicles.[14] We simulate the effect of replacing the unreformed CAFE standards with the reformed standards under four sets of assumptions with increasing constraints. 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. Second, we disallow within-firm trading. Third, we shut down the ability of firms to trade off acceleration performance with fuel efficiency, treating acceleration performance as exogenous. Fourth, we simulate short-run responses, only allowing for price adjustments.

Twenty firms are represented in these simulations, producing a total of 473 vehicle models, described by Table 2. 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). The choice of firms represented in these simulations is determined following the EPA's classification of manufacturers as listed on the MY2006 fuel-economy test data. For example, Saab is considered as part of GM but Land Rover is considered a separate firm from Ford.

Before summarizing the simulation results, it is instructive to highlight some of the differences between the unreformed CAFE and the MY2014 reformed CAFE policy regimes that

---

[14] 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 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.

are captured in the simulations. First, the reformed standards set individual fuel economy targets for each vehicle based on its footprint. Although these footprint-based standards allow larger vehicles to meet lower standards, even the lowest standard is 1.5 mpg higher than the unreformed CAFE standards. Second, the reformed CAFE standards allow trading within a firm between the passenger car standard and the light truck standard although they set a floor on the minimum passenger car standard of 32.4 mpg.

We simulate the partial-equilibrium price, marginal production cost, fuel economy, acceleration performance, and amount of technology implementation for every vehicle. Taken together, these simulation results can be used to calculate profits and consumer surplus. Point estimates of the effects of the reformed CAFE on economic surplus (i.e., the sum of producer and consumer surplus) are shown in Table 10. Note that this welfare analysis does not account for any indirect benefits associated with reduced fuel consumption, such as reduction in environmental damages. All values are measured relative to a baseline of partial equilibrium with respect to prices, acceleration performance, and technology implementation in the presence of the unreformed CAFE standards. Details of this baseline can be found in Appendix D.

While we do account for trading of fuel economy credits within firms in our counterfactuals, 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).

7.1 Impact of firm heterogeneity on fuel economy

The first set of counterfactual simulations allows firms to use the full complement of medium-run responses to the CAFE regulations including price changes and design changes (i.e., technology implementation and attribute tradeoffs). Firms constrained to the CAFE standards use a combination of these strategies to increase fuel economy, but the majority of improvements are due to changes in product design. In order to test the relative impact of design changes and price changes on fuel economy, we assign the prices from our pre-policy-change baseline to the counterfactual vehicle designs from the reformed CAFE simulations. This indicates that 62% of fuel economy improvements derive from changes to vehicle designs and 38% derive from price changes. Results indicate that the domestic Big 3 manufacturers improve the fuel economy of their vehicle fleets by an average[15] of 4.3 mpg, compromising acceleration performance such that the average 0-60 acceleration time is 2.7 s slower. This result is 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.

However, despite the rise in fuel economy by constrained firms, we find that the extent of fuel-consumption leakage to non-compliant firms is large. Manufacturers that violate the standard offset the increase in fuel efficiency by compliant firms such that both the average fuel economy and acceleration performance across the market remain approximately the same—fuel economy increases by only 0.1 mpg and acceleration performance changes by less than 0.3 s. This behavior, also noted by Jacobsen (2010), can be explained by the residual demand curve for lower fuel economy vehicles, which are larger or have better acceleration, shifting to the right in response to other firms producing higher fuel economy vehicles. Because of this effect, our simulations suggest that it is more profitable for even Toyota and Honda to violate the reformed CAFE standards and pay the corresponding fines due to the substantial increase in fuel economy required by constrained firms under the reformed CAFE standards.

Results of our short-run simulations, where vehicle designs are taken as exogenous and firms can only adjust prices to respond to the reformed CAFE standards, highlight the important effect that vehicle-design strategies have on firm responses to fuel economy regulations. Neither Ford nor Chrysler has passenger cars that exceed their individual MY2014 footprint-based standards. Therefore, they cannot meet the reformed CAFE standard for passenger cars through price adjustments alone. Simulation results imply that neither of these firms can meet the minimum passenger-car standard in the short run, even allowing for trading of fuel-economy credits between each firm's passenger-car and light-truck standards. Ford violates the passenger car standard by 4.0 mpg and Chrysler violates it by 4.2 mpg.

Furthermore, the leakage effect resulting from foreign firms decreasing the fuel economy of their vehicles is much smaller in the short run. Foreign manufacturers only decrease fuel economy by an average of 0.3 mpg, resulting in an overall increase in average new-vehicle fuel

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

economy by 2.3 mpg. Similar to the medium-run simulations, production of more fuel-efficient vehicles by compliant firms increases the residual demand curve for less fuel-efficient vehicles. But, because firms cannot adjust the fuel economy of their vehicle fleet as easily in the short run, the leakage effect is much smaller as compared to the medium-run simulations. Although non-compliant firms have an incentive to decrease fuel economy in both the short run and medium run, they are more able to respond in the medium run due to the ability to adjust vehicle designs.

7.2 Impact of attribute tradeoffs on firm behavior and costs

In the third specification, shown in Table 10, we exclude the ability of firms to trade off acceleration performance with fuel economy, considering the determination of acceleration performance as exogenous. We do this to investigate the effect of ignoring the tradeoffs between these attributes on the resulting welfare estimations. Results suggest that the costs of the regulation to compliant firms, in terms of profits lost, are twice as large as when these tradeoffs 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. Compared to the short-run specification, which treats all aspects of vehicle design as exogenous, costs to compliant firms are almost nine times lower in the specification that accounts for both attribute tradeoffs and technology implementation. Including within-firm credit trading further reduces these costs by 23%.

Contrasting the results of the "full medium run" and "medium-run without tradeoffs" specifications highlights the complex relationship between firm compliance strategies and production decisions. In the "medium-run without tradeoffs" specification, firms can only increase the fuel economy of their vehicles by implementing technology features, which increase vehicle production costs. Consequently, domestic firms (which are constrained to the standards) increase the prices of their vehicles. This behavior leads to a 13% decrease in the market share of domestic firms. These firms also shift production substantially from passenger cars to light trucks, in order to take advantage of the lower fuel-economy standards for light trucks. As a result, the share of light trucks across the domestic firms' fleets increases by 25%.

This behavior does not occur in the simulations that account for tradeoffs between acceleration performance and fuel economy. In the "full medium run" specification, constrained firms choose to increase fuel economy primarily by compromising acceleration performance, which decreases marginal costs but also decreases consumer utility. As a result, these firms lower their vehicle prices considerably, by an average of \$1,300. This behavior leads to a much smaller loss in the market share of domestic firms (5.5%) and a much smaller shift to light truck production (4.5%). Including within-firm credit trading tempers these responses further: vehicle prices decrease by an average of \$750, the market share of domestic firms declines by only 4.2% and the shift to light trucks is less than 1%.

7.3 Implications for climate and fuel consumption policies

Our results suggest that the cost-effectiveness of the reformed CAFE, in terms of profit loss per reduction in fuel consumption, is significantly better than cost-effectiveness estimates of the unreformed CAFE (e.g., Goldberg 1998; Jacobsen 2010). For example, Goldberg (1998) estimated the costs of the CAFE standards in 1998 as approximately \$1,140 per short ton CO_2 reduced, or \$11.06 per gallon reduced. Assuming vehicles have a lifetime of 13 years and are driven 14,000 miles per year with a rebound effect of 10.3% (Small and Van Dender 2007), our counterfactual results indicate that compliant firms reduce fuel consumption of their vehicles by 14% at a cost of \$197 per short ton of CO_2 reduced, or \$1.91 per gallon reduced. Considering CAFE solely as climate policy, this estimate is notably larger than cost estimates for other potential climate policy instruments such as a comprehensive cap and trade system (e.g., Stavins 2008). However, the welfare benefits from reducing fuel consumption are not limited to only the benefit of reducing CO_2 emissions. Reduced local air pollution and dependency on foreign oil also contribute to welfare benefits from CAFE.[16]

Despite significant gains in fuel economy by compliant firms, our results underscore that the effectiveness of CAFE—in terms of both fuel consumption reductions and costs—is highly dependent on the behavior of Asian and European manufacturers. Given the current fine of \$55 per vehicle per mpg under the standard, our results indicate that it is more profitable for each of these firms (including Toyota and Honda) to violate the standard and pay the corresponding fines. Furthermore, we find that these firms have an incentive in the medium run to substantially

---

[16] Estimates of the value of these welfare benefits are notably hard to obtain. For a rough comparison, estimates of welfare benefits from Parry et al. (2004) attributed to local air pollution and oil dependency sum to 22 cents per gallon of gasoline reduced. On the other hand, if the rebound effect leads to substantially higher driving during congested periods, CAFE may lead to welfare losses from increased traffic congestion.

decrease fuel economy such that the fuel consumption across all new vehicles is approximately the same. These results imply negligible reductions in CO_2 at exceptionally high costs.

It is important to emphasize that, unlike leakage problems that have been characterized elsewhere in the literature, the leakage we observe in our simulations can be readily mitigated through policy design—namely, by increasing the level of the fine. 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. These findings have direct implications for the recently announced greenhouse-gas emission standards for light-duty vehicles, which give the EPA the authority to set the penalties for noncompliance on a case-by-case basis. Ideally, the fines should be set equal to the marginal damages that would result from increasing fuel consumption above the level necessary to exactly meet the standard. Because fuel consumption is the inverse of fuel economy, this relationship varies with the standard and with the number of miles per gallon that a firm's average fuel economy is below the standard.

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 the majority of fuel-efficiency improvements of compliant firms are from changes in product designs. Compliant firms also adjust product prices to shift demand to more efficient vehicles in response to fuel economy standards, but this has a smaller effect on fuel-efficiency improvements than design changes.

Results highlight the substantial sensitivity of profit losses and fuel efficiency gains 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 twice as high as when these tradeoffs are considered. When product designs are considered exogenous, the profit losses to constrained firms are over nine times greater. 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 policy costs.

Furthermore, our results highlight the notable difficulty of significantly increasing fuel economy across the U.S. market of new vehicles when firms choose to violate the CAFE standards and instead pay a fine. Without consumer incentives for higher fuel economy, results indicate that increases in fuel efficiency from one set of firms could be almost entirely offset in the medium run by decreases in fuel efficiency from other firms. These results are sensitive to the assumption that Asian manufacturers will choose to violate the standard and pay the corresponding fines when it is more profitable, but underscore the large effect their behavior has on the outcomes of the CAFE standards.

A number of limitations to the current analysis exist and may affect these results. Sensitivity analyses are needed to test the robustness of our results to the assumptions used in our model and the standard errors on parameter estimates. Most notably, future work will include sensitivity tests to estimates of engine and technology-feature costs, the shadow costs of the unreformed standards, and the specification 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.

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

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.

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.

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.

10 Tables

Table 1: Technology Costs and Percentage Improvements in Fuel Economy and 0-60 Acceleration Time

	Two Seater			Compact			Midsize / Minivan			Fullsize		
Technology	cost	% mpg	% acc	cost	% mpg	% acc	cost	% mpg	% acc	cost	% mpg	% acc
Low friction lubricants	3	0.5	0.3	3	0.5	4.1	3	0.5	0.6	3	0.5	0.6
Engine friction reduction	126	1	0.3	84	1	5.6	126	1	1.5	126	1	1.2
Aggressive shift logic	38	1	-0.2	38	1	-5.0	38	1	-0.2	38	1	-0.3
Early torque converter lockup	30	0.5	-	30	0.5	-	30	0.5	-	30	0.5	-
High efficiency alternator	145	1	0.3	145	1	5.6	145	1	1.5	145	1	1.2
Aerodynamic drag reduction	38	3	0.3	38	3	5.1	38	3	0.5	38	3	0.3
Low rolling resistance tires	6	1	0.1	6	1	2.5	6	1	0.2	6	1	0.1
Cylinder deactivation		n/a			n/a		203	4.5	-	203	4.5	-

	SUV			Sm Pickup			Lrg Pickup / Van		
Technology	cost	% mpg	% acc	cost	% mpg	% acc	cost	% mpg	% acc
Low friction lubricants	3	0.5	1.6	3	0.5	0.4	3	0.5	0.7
Engine friction reduction	126	1	3.1	126	1	0.7	168	1	1.5
Aggressive shift logic	38	1	0.0	38	1	-0.2	38	1	-2.8
Early torque converter lockup	30	0.5	-	30	0.5	-	30	0.5	-
High efficiency alternator	145	1	3.1	145	1	0.7	145	1	1.5
Aerodynamic drag reduction	38	3	1.4	38	2	0.5	38	2	0.4
Low rolling resistance tires	6	1	0.5	6	1	0.2	0	1	0.1
Cylinder deactivation	203	4.5	-	203	4.5	-	229	4.5	-

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".

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				

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

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

Table 5: Heterogeneous Demand Parameter Results

σ		
	param	**st. err.**
p	0.0377	0.0239
gpm	0.0220	0.0131
accinv	0.0321	0.0531
ftp	0.0393	0.0529

μ		
	param	**st. err.**
p/inc	-0.1756	0.0248
minivan-child	1.2806	0.2456
suv-child	1.0925	0.1594
truck-rural	1.8500	0.2444

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

Table 6: 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 Eq. 7.

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

*p<0.1, **p<0.01, ***p<0.001, all others p>0.1

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: 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 Eq. 10 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 9.

Table 9: 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 10.

Table 10: Impact of MY2014 CAFE on Social Surplus

	Change in Consumer Surplus			
	Full Medium Run with within-firm Trading	*Full Medium Run*	*Medium Run Without Tradeoffs*	*Short Run with within-firm Trading*
Total (billions)	-\$59	\$8	\$11	-\$113
	Change in Producer Welfare			
	Full Medium Run with within-firm Trading	*Full Medium Run*	*Medium Run Without Tradeoffs*	*Short Run with within-firm Trading*
Total (billions)	-\$144	-\$142	-\$163	-\$182
Constrained firms	-\$17	-\$21	-\$44	-\$149 †
Fine-paying firms	-\$127	-\$121	-\$119	-\$33

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. The *Short-Run with within-firm Trading* specification only allows firms to adjust the prices of their vehicles and trade fuel-economy credits between the passenger car and light truck standards; all vehicle designs are considered exogenous. The *Medium-Run Without Tradeoffs* specification accounts for price changes and the ability of firms to implement fuel-saving technologies but does not allow for any tradeoffs between fuel economy and acceleration performance. The *Full Medium-Run* specification accounts for price changes, technology implementation, and tradeoffs between fuel economy and acceleration performance. The *Full Medium-Run with within-firm Trading* specification further adds the ability of firms to trade fuel-economy credits between their passenger car and light truck standards.
† In the short run, neither Chrysler nor Ford could meet the reformed CAFE minimum requirement for passenger cars; Chrysler violated it by 4.2 mpg and Ford violated it by 4.0 mpg.

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

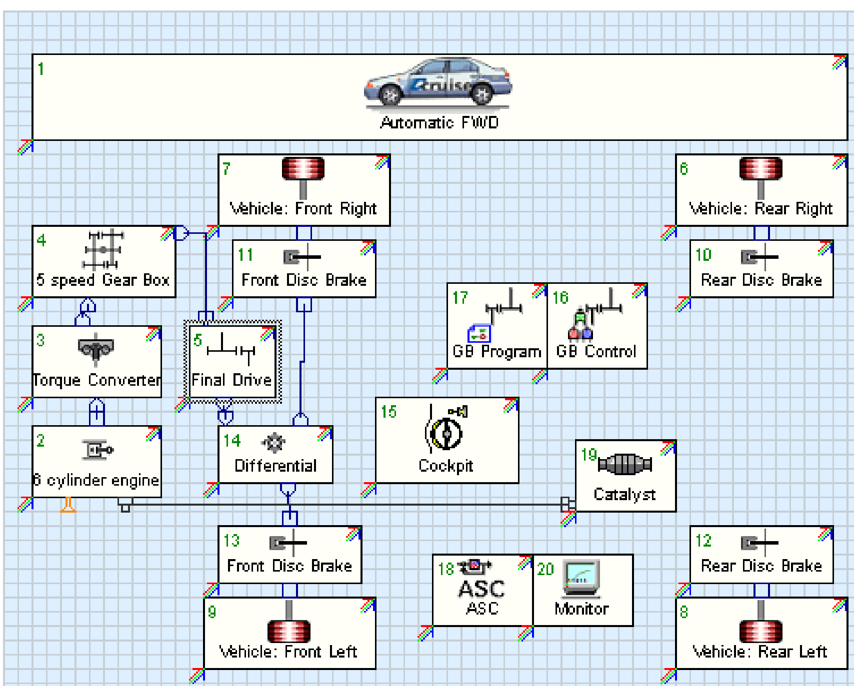

The modular structure of Cruise allows researchers to simulate multiple vehicle architectures by customizing the subsystem modules (e.g., front or rear wheel drive, automatic or manual transmissions), and modifying various input parameters. For example, with the Vehicle module, a user can adjust the aerodynamic drag coefficient of the vehicle body and the curb weight 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., curb weight), 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.[17] 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.

[17] 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 curb weight as well as similar predicted outputs from AVL Cruise.

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

The specifications for Eq. 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

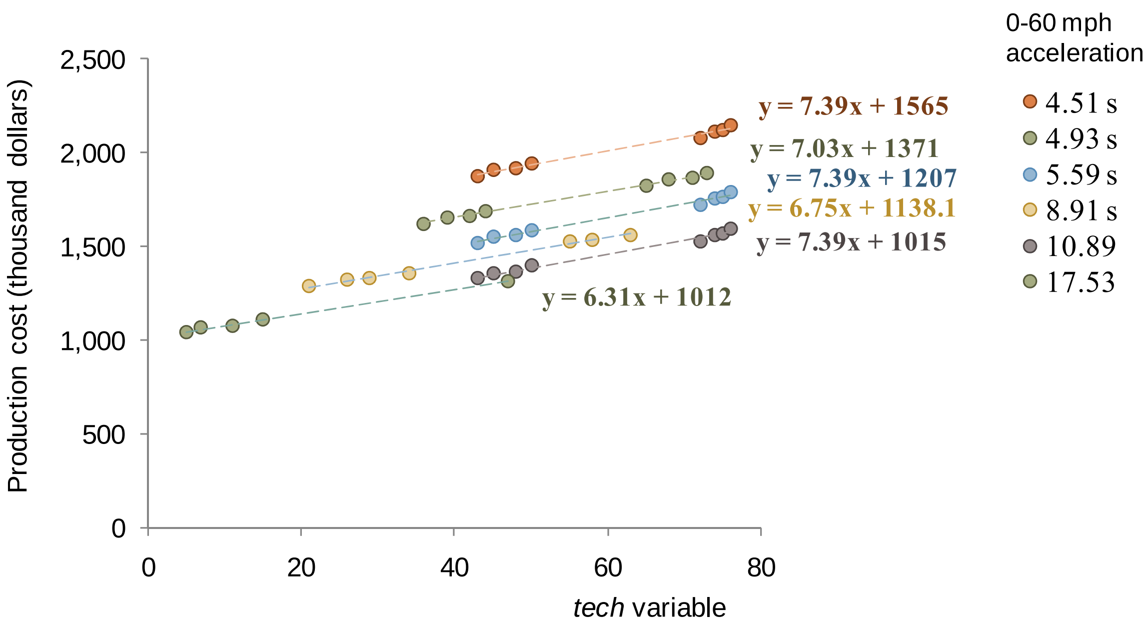

Figure B1 is generated for a specific vehicle segment (an SUV) and a specific curb weight (3,200 lb). Similar trends were found for other segments and other curb weights. The figure indicates that conditional on vehicle segment, curb weight, and 0-60 acceleration time, moving "up the line" of combinations 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 Eq. 4 where cost is linear in technology conditional on vehicle segment, curb weight, and acceleration time.

Figure B2 below plots fuel consumption against combinations of technology features for the same vehicle segment (SUV) and curb weight (3,200 lb). This figure indicates that conditional on vehicle segment, curb weight, 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 Eq. 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

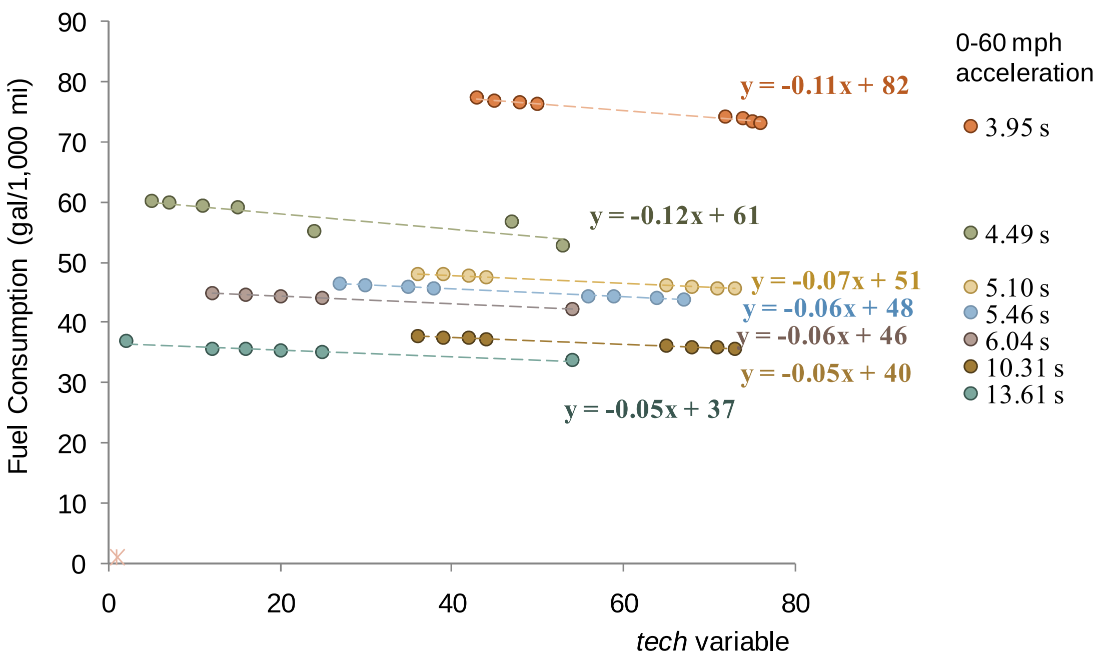

Appendix C Endogenous Attribute Model

We compare the results of our approximated endogenous attribute model, fit using simulation data, to market vehicle data as a validation test. Figure A1 and A2 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 C1: Comparison of compact-segment model to MY2006 compact vehicle data

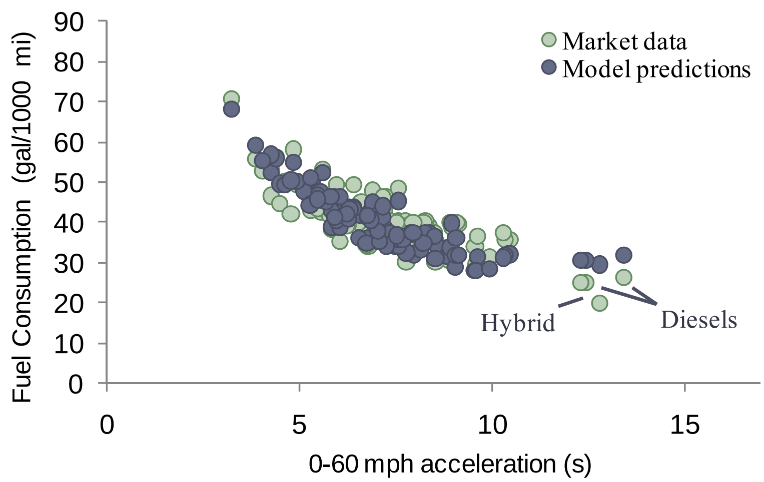

Figure C2: Comparison of SUV-segment model to MY2006 SUV data

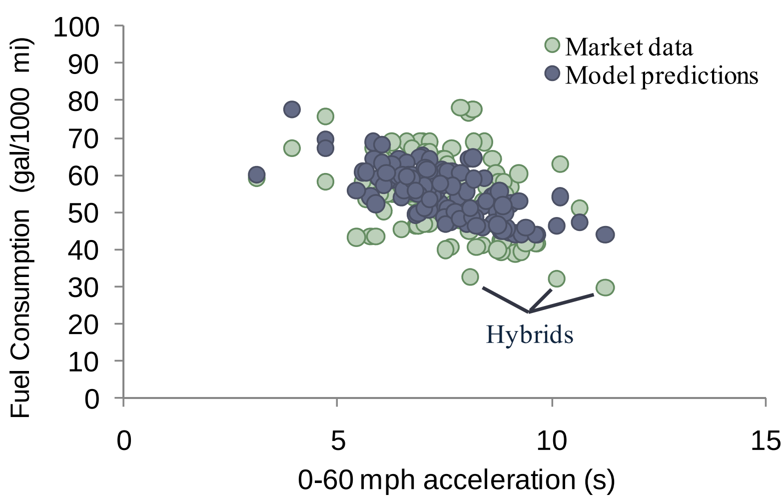

Appendix D Counterfactual baseline

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. Our results shown in Table E1 indicate that fuel economy and 0-60 acceleration time for MY2006 vehicles are very close to equilibrium.

Due to a lack of data on adoption of each technology feature, we cannot compare the results of adopted technology features to those in the MY2006 vehicles. But, the simulations results appear to match general information about technology. For instance, very few vehicles in MY2006 included cylinder deactivation, which is supported by the simulations.

Table E1: Comparison of observed attributes and baseline simulation

	Average Fuel economy (mpg)	Average 0-60 acceleration (s)	Average *tech*
Observed	21.5	7.11	-
Baseline Simulations	20.7	7.42	0.19