DISCUSSION PAPER

December 2010 ■ RFF DP 10-50

New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard

Thomas Klier and Joshua Linn

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard

Thomas Klier and Joshua Linn

Abstract

By 2016, the Corporate Average Fuel Economy (CAFE) standard will increase by 40 percent from its current level, representing the first major increase in the standard since its creation in 1975. Previous analysis of the CAFE standard has focused on its short-run effects (1–2 years), in which vehicle characteristics are held fixed, or its long run effects (10 years or more), when firms can adopt new power train technology. This paper focuses on the medium run, when firms can choose characteristics such as weight and power, yet have only limited ability to modify current technology. We first document the historical importance of the medium run and then estimate consumers' willingness to pay for vehicle characteristics. We employ a novel empirical strategy that accounts for the vehicle characteristics' endogeneity by using variation in the set of engine models used in vehicle models. The results imply that consumers value an increase in power more than a proportional increase in fuel economy. Simulations of the medium-run effects of an increase in the CAFE standard suggest that regulatory costs are significantly smaller in the medium run than in the short run.

Key Words: new vehicles market, CAFE, medium run, demand estimation

JEL Classification Numbers: Q5, L5, L62

Contents

1. Introduction **1**

2. Data **5**

3. Fuel Economy Regulation and Engine Technology **6**

3.1 The CAFE Standard 7

3.2 CAFE and Market Outlook 7

3.3 The Medium Run 8

4. Response to the Initial CAFE Standard **9**

5. Medium-Run Model of the New Vehicle Market **12**

5.1 The Demand for New Vehicles 12

5.2 The Supply of New Vehicles 13

5.3 Demand Estimation 16

5.4 Supply Estimation 22

6. Short- and Medium-Run Effects of the CAFE Standard **23**

6.1 Short Run 23

6.2 Medium Run 25

7. Robustness of Demand Estimates **26**

8. Conclusion **27**

References **29**

Tables and Figures **31**

New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard

Thomas Klier and Joshua Linn*

1. Introduction

A vast literature has attempted to estimate the cost of a regulation or the welfare effects of a merger in differentiated product markets. In many such markets, firms choose the characteristics of the products they sell in response to consumer demand, regulation, and competition. Consequently, regulatory costs or welfare effects depend partly on how firms change their products' characteristics in response to regulatory or competitive shocks. For example, the welfare effects of a merger between two cable television providers depend partially on how firms subsequently change their programming menus. Yet most studies assume that product characteristics are exogenous—in other words, that they are not chosen by firms. This assumption greatly simplifies the modeling and empirical analysis, but risks misstating regulatory costs or the welfare effects of a merger.

The new vehicles market is a prominent example of these concerns. The Corporate Average Fuel Economy (CAFE) standard, which represents the minimum fuel economy that manufacturers of new vehicles must attain in the U.S. market, is the major regulation used to affect fuel economy and gasoline consumption in the United States. Past analysis of the market generally has taken most vehicle characteristics to be exogenous, both in modeling the market and in estimating consumers' willingness-to-pay for characteristics.[1] This paper makes two contributions, first by estimating the cost of CAFE to vehicle producers and consumers using a model in which firms choose characteristics endogenously, and second, by estimating

* We thank Michael Greenstone; James Sallee; conference participants at the National Bureau of Economic Research Summer Institute; and seminar participants at Carnegie Mellon, NYU, and Resources for the Future for helpful comments. Taft Foster, Vincent Liu, and Cristina Miller provided excellent research assistance. We are grateful to Eric Fedewa from CSM for providing data on engine specifications. Authors' email addresses: thomas.klier@chi.frb.org and linn@rff.org.

[1] Two exceptions are Shiau et al. (2009), who use a model with endogenous vehicle characteristics to analyze the CAFE standard but do not account for the endogeneity in estimating consumer willingness to pay for vehicle characteristics; and Gramlich (2009), who uses a model with endogenous fuel economy but focuses on the effect of gasoline prices on fuel economy. Some studies, such as Austin and Dinan (2005), use a model in which fuel economy is endogenous but other characteristics are exogenous.

consumers' willingness to pay for characteristics while accounting for the endogeneity of the characteristics. Estimates of both willingness to pay and the cost of CAFE are significantly different from previous estimates; in fact, we find the cost of compliance to be lower than in most previous work.

The Energy Independence and Security Act of 2007 increased the CAFE standard for new vehicles by about 40 percent, to be effective by the year 2020.[2] The legislation represents the first significant increase in the standard since it was created in 1975. It was passed following a period of vigorous public debate in which the law's proponents argued that it would reduce carbon dioxide emissions and oil imports without undermining the automobile industry. Opponents claimed that the costs to vehicle manufacturers and consumers would not justify the benefits, and that other policies would be more effective at reducing emissions and oil imports.

A number of studies have analyzed the costs to consumers and producers of using the CAFE standard to reduce gasoline consumption. These studies simulate the effect of an increase in the standard on market equilibrium and can be classified into two categories. Some, including Goldberg (1998), have used a short-run model, pertaining to one or two years after a change in the standard, in which vehicle characteristics and technology are held constant. Firms respond to an increase in the CAFE standard by adjusting vehicle prices. Other studies, such as Austin and Dinan (2005), use a long-run model to estimate costs; this pertains to 10 years or more after a change in the standard. In that analysis, firms choose vehicle prices as well as power train (the combination of a vehicle's engine and transmission) technology.

Yet casual observation of the new vehicles market suggests that both of these approaches are overly simplified: firms typically select vehicle prices every year, whereas major changes to a vehicle's power train technology are undertaken in much longer cycles, approximately every 10 years. But every four or five years, firms tend to redesign vehicles by changing their characteristics, such as interior cabin features. Of particular relevance to achieving compliance with the CAFE standard is the fact that during the redesign firms can change the fuel economy of a vehicle by altering its weight and power or by making relatively minor changes to the power train. For example, fuel economy can be increased by removing certain components, using lighter materials, or modifying the engine to reduce the number of cylinders that power the

[2] During the spring of 2009 the Obama administration moved up the deadline by which the new requirements have to be met from model-year 2020 to model-year 2016.

vehicle at low speeds. By contrast, in the long run, much larger changes to a vehicle's power train are possible, such as adopting hybrid power train technology. Because relatively minor changes are routinely introduced in the new vehicles market, they are also expected to occur in response to the new CAFE regulation. For example, in the spring of 2008, Honda introduced the 2009 version of the Acura TSX model. The new version had less power and greater fuel economy than the previous version. The vice president of corporate planning for Honda announced at the time of the introduction that "We feel comfortable there's plenty of horsepower already and wanted to focus on improving fuel efficiency and emissions. For us generally, you'll see more of that" (Ohnsman 2008). There is thus a medium-run response to CAFE that is distinct from previously modeled short-run price changes or long-run technology adoption.

The existing literature on CAFE has concluded that this regulation is far more costly than raising the gasoline tax to reduce gasoline consumption. However, because the previous analysis does not incorporate the medium-run margin of adjustment, total (discounted) costs to producers and consumers may be overstated. To the extent that reductions in weight and power or modifications to the power train are less costly than adjusting the sales mix, actual costs several years after a change in the regulatory standard could be much lower than the short-run analysis suggests; the argument is even stronger when a change in the standard is announced several years in advance. Medium-run changes in characteristics may also reduce the need to equip vehicle models with expensive advanced power train technologies in the long run, implying that the long run estimates may also be too high. Moreover, the distinction between the short run and the long run may overstate the time required for significant improvements in fuel economy to be realized. Although the medium-run margin represents a frequently used option for vehicle producers, it remains an open question as to whether it alters the cost of complying with CAFE.

Our paper is structured as follows. Sections 2–4 document the importance of changes in weight and power following the imposition of the initial CAFE standard in 1978. It turns out that reductions in weight and power explain much of the increase in fuel economy during the first years of the program. Afterwards, long-run technology adoption becomes increasingly important.

This pattern demonstrates the importance of medium-run decisions in complying with CAFE. It also suggests that the medium-run response to CAFE lasts about five years.[3]

The second part of the paper, in Sections 5–7, analyzes the medium-run effects of the CAFE standard on vehicle producers and consumers. We develop a model of the new vehicles market in which firms choose the characteristics of their vehicles, including fuel economy, weight and power. Implementation of the model poses two major challenges. First, to model the demand side, it is necessary to estimate the demand for vehicle characteristics while accounting for the fact that firms choose these very characteristics. The vast majority of the literature on consumer demand in the new vehicles market has assumed that vehicle characteristics are exogenous, in that characteristics observed by the econometrician are uncorrelated with unobserved characteristics. For example, Berry, Levinsohn and Pakes (1995; henceforth, BLP) construct a set of instrumental variables (IV) that is valid only if observed characteristics are uncorrelated with unobserved characteristics. That seems unlikely to be the case, however. For example, (observed) vehicle size may well be related to (unobserved) vehicle handling.

Several recent studies of other industries have in fact estimated consumer demand for endogenous characteristics (e.g., Ishii 2005). However, the new vehicles market poses the additional difficulty that observed vehicle characteristics are potentially correlated with unobserved characteristics. In this case, estimation requires an identifying assumption on the joint distribution of the observed and unobserved variables. For example, Sweeting (2007) assumes that changes in unobserved characteristics of radio stations occur after the firm has chosen the observed characteristics.[4] To address this issue we use an IV strategy that focuses on cost shocks across products and is broadly similar to Hausman *et al*. (1994). It exploits a particular feature of the new vehicles market: firms often sell vehicle models in different vehicle classes with the same engine. For example, the Ford F-Series (a pickup truck) and the Ford Excursion (a sport utility vehicle [SUV]) are available with the same engine. That allows us to instrument for a vehicle's endogenous characteristics (e.g., the characteristics of the F-Series) by

---

[3] Several studies have analyzed the changes in weight, power, and fuel economy after CAFE was adopted. Similarly to our study, Greene (1987, 1991) concludes that short-run changes in the sales mix explain a small share of the increase in fuel economy and that technology explains about half of the increase in fuel economy. Greene and Liu (1988) calculate the change in consumer surplus after CAFE was adopted using changes in these characteristics and willingness-to-pay estimates from other studies.

[4] In Sweeting (2007), unobserved station quality is exogenous, but is potentially correlated with observed characteristics. The timing assumption is used to construct a valid set of instruments using lagged variables.

using the engine characteristics of vehicles that have the same engine but are located in different vehicle classes (e.g., the Excursion). Our modeling results imply that consumers are willing to pay more for an increase in power than for a proportional increase in fuel economy.

The second challenge is that the supply side of the model must capture the technical trade-offs among vehicle characteristics, as well as the effect of these characteristics on production costs. For example, a given weight reduction raises fuel economy by a particular amount; it may also affect production costs. The technical relationship between fuel economy, weight, and power is estimated as part of the historical analysis presented in the first part of the paper. It is based on highly detailed engine data. The effect of vehicle characteristics on production costs is estimated from data provided in National Research Council (NRC 2008).

We use the empirical results to estimate the medium-run cost of the CAFE standard—that is, its effect on consumers and on vehicle producers' profits. Our simulations suggest that the medium-run costs are significantly lower than the short run costs, demonstrating the importance of accounting for vehicle characteristics' endogeneity when analyzing CAFE.

2. Data

This section describes the data sources used in the subsequent analytical work. We use a detailed data set of vehicle and engine characteristics and vehicle sales from 1975-2008 (see Klier and Linn [2010] for more detail). Vehicle sales are from the weekly publication Ward's Automotive Reports for the 1970s and from Ward's AutoInfoBank for subsequent years. Sales are matched to vehicle characteristics by vehicle model from 1975 to 2008.[5] The characteristics data, which are available in print in the annual Ward's Automotive Yearbooks (1975–2008), include horsepower, curb weight, length, displacement, fuel economy, and retail price. Note that the data do not include fuel economy from 1975 to 1977, as fuel economy was not reported prior to the CAFE program. We impute fuel economy from the other vehicle characteristics during these years, using the estimated relationship among characteristics for 1978 and 1979.

The data coverage for cars is far more extensive than for light trucks. For the 1970s, our sample includes all car models produced in the United States, but no light trucks. Consequently,

[5] The match is not straightforward because the two data sets are reported at different levels of aggregation. Vehicle characteristics are reported at the *trim level* to recognize differences in the manufacturer suggested retail price; for example, the data distinguish the two- and four-door versions of the Honda Accord sedan. We aggregate the characteristics data to match the model-based sales data and use means of the characteristics in the empirical work.

the historical analysis in this paper focuses on cars, which account for most of the market during the late 1970s and early 1980s. The share of light trucks in the new vehicles market was between 20 and 30 percent between the years 1975 and 1988 (U.S. Environmental Protection Agency [EPA], 2007).

We obtained detailed data on engine specifications for the years 2000–2008 from CSM, a Michigan-based automotive sector consulting firm. The engine data distinguish two levels of aggregation: an engine program, which refers to a distinct power train technology, and a platform, which is a collection of related programs. For example, the Volkswagen Passat and Audi A4 are sold with power trains of the same program; the Volkswagen Jetta features power trains from a different engine program, but the power trains of both programs belong to the same platform. Firms may produce different versions of the same engine program that vary by power.[6]

For each vehicle model, we construct a list of engine programs that are sold with that model. For a given vehicle, there are three sources of variation over time in the engine technologies that are sold with it. First, the engine may be redesigned, in which case the program identifier changes. Second, firms may discontinue selling a particular engine with a specific vehicle model, as Honda did when it discontinued the hybrid version of its Accord sedan after model-year 2007.[7] Third, a firm can introduce a new version of the vehicle model that is sold with an engine that was previously sold only with other vehicle models. We have matched engine and vehicle model characteristics for 2000–2008.

We use annual Wards vehicle sales and characteristics data from 1975-2008 for the descriptive analysis in Section 4. Sales, characteristics and engine data for 2000-2008 are used for the demand estimation in Section 5.

3. Fuel Economy Regulation and Engine Technology

This section provides background on the CAFE standard and discusses decisions made by firms in the medium run.

[6] To further illustrate the distinction between programs and platforms, a single engine platform can include engine programs with different numbers of cylinders, whereas a single engine program features engines with the same number of cylinders.

[7] A model-year refers to the production cycle, which typically lasts one year and begins in August or September of the previous calendar year (Klier and Linn, 2010).

3.1 The CAFE Standard

Following the 1973 oil crisis, Congress passed the Energy Policy and Conservation Act in 1975 to reduce oil imports.[8] The Act established the CAFE program and required automobile manufacturers to increase the average fuel economy of passenger and non-passenger vehicles sold in the United States. Cars and light trucks are subject to separate standards, which have varied slightly over time; for model-year 2009, the standards were 27.5 miles per gallon (mpg) for cars and 23.1 mpg for light trucks.[9] Firms may also earn credits for over-compliance that can be applied towards compliance in future years. The standards are administered by the U.S. Department of Transportation (DOT) on the basis of the EPA's test procedure for measuring fuel economy.

The Energy Independence and Security Act of 2007 requires DOT to raise fuel economy standards, starting with model-year 2011, until they achieve a combined average fuel economy of at least 35 mpg. The initial law required this standard to be met by 2020, but the compliance date has since been changed to model-year 2016.[10] The CAFE standard continues to be extremely controversial; the 2007 law has been called "a victory for America" (Senator Carper, D-Del, quoted in Stoffer 2007), as well as "unnecessary at best and damaging at worst," (Ingrassia, 2008). Note that firms will be evaluated for compliance with the new standard using a different formula that is based on a vehicle's *footprint* (the product of length and width).

3.2 CAFE and Market Outlook

As Section 4 shows in more detail, when the original CAFE standard was introduced, automobile manufacturers rather quickly reduced horsepower and weight to raise fuel economy. Over time, power train technologies improved, which allowed firms to improve a vehicle's performance while continuing to meet the CAFE standard.

---

[8] This section draws extensively from National Research Council (2008).

[9] The 2009 standard for cars is identical to the standard from 1985, when the original CAFE requirements were phased in. The light truck standard for 1985 was 19.5 mpg.

[10] When moving up the compliance deadline to 2016, the Obama administration also instructed the EPA to regulate automobile greenhouse gas emissions. If it is assumed that the required reduction in greenhouse gas emissions will be met entirely through fuel economy improvements, the "GHG-equivalent" mileage requirement would be 35.5 mpg, which is slightly higher than in the initial legislation (Foster and Klier 2009). The final CAFE rule, which became effective in May 2010, states the target fleet fuel economy as 34.1 mpg by 2016. That number is less than 35.5 because some reductions of GHG emissions, such as making a vehicle's air conditioning system more efficient, will affect fuel economy (Yacobucci 2010).

Many industry analysts believe that because many of the “easy” improvements to power train technology were made in response to the initial CAFE standard, the pending increase in the standard may be much more costly to producers and consumers. Although new power train systems, such as those relying on hybrid electric and diesel technologies, have begun to penetrate the U.S. market, the vast majority of new vehicles sold continue to be powered by conventional gasoline-powered spark-ignition engines. Despite the fact that essentially every vehicle manufacturer is advertising its alternative power train research, as of late 2009, sales of hybrid vehicles represented only about 3 percent of total sales of cars and light trucks. Thus, the performance characteristics of the existing gasoline engine technology, as well as the related transmission technologies, are the focus of attention.

3.3 The Medium Run

We define the medium run as the period of time in which firms can adjust a vehicle’s weight, power, and fuel economy without adopting a new engine program. In the new vehicles market, the short, medium, and long runs arise from the timing of firms’ major decisions. Firms typically choose vehicle prices once each year, although firms can also offer price incentives during the year. Large changes in vehicle characteristics typically occur every 4–5 years during major model redesigns. Engine technologies change more slowly, as engines are redesigned roughly every 10 years.[11]

A firm can change vehicle characteristics in two ways in the medium run. First, the firm may improve fuel economy by reducing vehicle weight or power. Second, the firm can modify the power train in a way that does not require the firm to redesign the engine or transmission. Power trains are intentionally designed with this flexibility, allowing firms to respond to cost or

[11] The main reason for the lifespan of an engine program is the amount of time it takes for research and development and to source the parts—the *engine design cycle*, which tends to run just over four years. After the engine begins production, it typically takes about 10 years to amortize the development costs. According to Eric Fedewa, vice president of global power train forecasts at CSM, engine design cycles are getting somewhat shorter, but not significantly. To complicate matters, major automakers optimize the timing of production and research and development for their entire portfolio of vehicle and engine programs, suggesting that incremental, rather than major, changes to the power train would occur first in response to unanticipated shocks, such as an unforeseen increase in the CAFE standard (Hill *et al*. 2007). Thus, although it may be possible to shorten the engine design cycle, it would be very costly to do so for a large automaker that sells vehicles with many engine programs. It is unlikely that firms would significantly shorten the cycle for all programs, suggesting the medium run extends for several years.

demand shocks without having to redesign the power train. Table 1 provides examples of medium- and long-run changes to the engine or transmission (National Highway Traffic Safety Administration 2008). Compared to long-run changes, the medium-run changes are simple to implement and generally cost less, but result in smaller fuel economy gains. Thus, following an unexpected increase in the CAFE standard, firms may adjust prices in the short run; weight, power, and fuel economy in the medium run; and power train technology in the long run. In the medium run, the firm takes the engine platform and transmission technology as given, but can make changes to power, weight, and fuel economy without changing the engine platform or transmission technology.

4. Response to the Initial CAFE Standard

This section documents changes in fuel economy, weight, and power in the late 1970s and early 1980s. Changes in weight and power explain much of the increase in fuel economy during this period. This result motivates the use of a medium-run model to simulate the effect of CAFE in Sections 5 and 6. Section 4 also reports estimates of the technical trade-offs between fuel economy, weight, and power, which we use in the simulations in Section 6.

Figure 1 provides summary information on vehicle characteristics over time. The figure shows the CAFE standard and weight, power and fuel economy for all cars sold in the United States from 1975 to 2007, using data reported in EPA (2007). Average fuel economy increased dramatically in the late 1970s and early 1980s as the standard was phased in. During the same period, power and weight decreased significantly, before increasing again in subsequent years.

The increase in fuel economy in Figure 1 could be due to a number of factors: short-run changes in the prices and quantities; medium-run changes in power, weight, or fuel economy; or the long-run adoption of new power train technology. This section decomposes the total increase in fuel economy into these three effects. The analysis in this section focuses on cars sold by U.S.-based automobile manufacturers (AMC, Chrysler, Ford, and GM) for two reasons. First, as Jacobsen (2008) notes, during the 1970s and 1980s there were three categories of firms: firms that consistently exceeded the fuel economy standard by a large amount (e.g., Honda and Toyota); firms that were constrained by the standard and typically met it (e.g., U.S.-based manufacturers); and firms that consistently paid a fine for not meeting the standard (e.g., foreign-based luxury car producers). U.S. firms account for the vast majority of sales from the constrained category, so their response to CAFE is of particular interest. We focus on U.S. cars because data on light trucks are incomplete, as noted above.

For comparison with Figure 1, Figure 2 reports the fuel economy, weight, and power of cars sold by U.S. firms using data from Ward's instead of EPA. The figure shows that changes in the characteristics of U.S. firms' cars were similar to those of the overall market, which reflects the dominance of U.S. firms during this time period. Between 1975 and 1978, which was the first year the CAFE standard was in effect, fuel economy increased by about 2 mpg. Gasoline prices were fairly stable during this time period, suggesting that the increase was in anticipation of the standard. However, because we impute the values for fuel economy from 1975 to 1977, this result should be treated with caution. From 1978 until the early 1980s, fuel economy increased by an additional 4 mpg. From the mid-1980s until the end of the sample period, average fuel economy was slightly higher than the standard.

While fuel economy was increasing, weight and power were decreasing. Both power and weight decreased by about 25 percent between 1975 and 1982, after which they increased steadily. The increase in fuel economy following the imposition of the CAFE standard thus coincided with a large decrease in power and weight. Subsequently, weight and power increased while fuel economy did not change.

The remainder of this section compares the relative magnitudes of the short-, medium- and long-run responses to the initial CAFE standard. We first separate the short run from the medium and long run. The first data series in Figure 3 reproduces the sales-weighted fuel economy from Figure 2. The second series represents the short-run fuel economy. The series is set equal to actual fuel economy in 1975, and subsequent changes from one year to the next are equal to the sum over vehicles of the change in market share multiplied by the fuel economy from the previous year. Thus, an increase in the market shares of vehicle models that initially have high fuel economy would cause the sales-weighted average fuel economy to increase. The short-run changes explain a large fraction of the total change in fuel economy in 1978 and 1979, but a very small share in all other years. The limited duration of these changes is consistent with the definition of the short run.

After accounting for short-run changes in average fuel economy in Figure 3, the remaining changes in average fuel economy could be due to either medium- or long-run changes in vehicle characteristics and technology. Unfortunately, detailed technology data from this period are not available to distinguish clearly the medium run from the long run. However, we can estimate the effect of weight and power on fuel economy, which provides a lower bound to the full medium-run response.

We first estimate the within–engine technology trade-off between fuel economy, weight and power. To that effect we use data from 2000 to 2008 to estimate the following equation:

$$\ln M_{jey} = \delta_0 + \delta_1 \ln H_{jey} + \delta_2 \ln W_{jy} + \eta_e + \varepsilon_{jey}. \quad (1)$$

The dependent variable is the log of the fuel economy of vehicle j that has engine technology e (denoting either a platform or program) in year y, and the first two variables are the logs of power and weight. Equation (1) includes engine technology fixed effects. Thus, the coefficients on power and weight are the within–engine technology elasticity of fuel economy with respect to power and weight. The coefficients represent the technological trade-off between power, weight and fuel economy.

Table 2 reports the results of estimating equation (1). The two columns include engine program and engine platform fixed effects (recall that multiple engine programs may belong to the same platform). The reported coefficients are the within-program and within-platform effects of power and weight on fuel economy. The within-program elasticity of fuel economy with respect to power is –0.07 and for weight is –0.33; the coefficient on power is larger in column 2 with platform fixed effects. On the other hand, the effect of weight on fuel economy is the same across specifications, which is as expected because weight varies at the vehicle level and not the engine level.

Table 2 suggests that firms can increase fuel economy by decreasing power and weight in the medium run. Assuming that the elasticities have not changed over time, we can use the estimated parameters in equation (1) to obtain a lower bound of the medium-run response to CAFE. In particular, we use the actual weight and power each year from 1975 to 2007 for the sample in Figure 2, combined with the estimates in column 1 of Table 2, to predict the fuel economy of each vehicle. The predicted series captures the medium-run effect of weight and power on fuel economy. The difference between the actual and predicted mpg series can be interpreted as the effect on fuel economy of medium-run power train modifications (i.e., modifications that do not require a reduction in power) and long-run technology adoption. Figure 4 shows the actual and predicted fuel economy from 1975 to 2007. The figure demonstrates that decreases in power and weight explain about one-third of the increase in fuel economy in the late

1970s and early 1980s.[12] Given that this is probably a lower bound, we conclude that, historically, the medium run response to CAFE has been important.

5. Medium-Run Model of the New Vehicle Market

We first specify the demand and supply sides of the static medium-run model, followed by the estimation of the model parameters. Section 6 reports simulations of an increase in the CAFE standard.

5.1 The Demand for New Vehicles

Consumers decide whether to purchase a new vehicle and, if so, which vehicle. Consumer i derives utility from vehicle j according to the utility function:

$$U_{ij} = \alpha \ln p_j + \beta_D D_j + \beta_H HW_j + \beta_W W_j + \xi_j + \varepsilon_{ij}, \tag{2}$$

where p_j is the manufacturer's suggested retail price (MSRP) of the vehicle; D_j is dollars-per-mile, equal to the ratio of the price of gasoline to the fuel economy (M_j) of the vehicle; HW_j is the ratio of the vehicle's horsepower (H_j) to weight (W_j); ξ_j represents the characteristics of the vehicle that are unobserved by the econometrician; ε_{ij} is an error term that varies by consumer and vehicle; and α, β_D, β_H, and β_W are parameters to be estimated. The price of gasoline is constant in the static framework, so dollars-per-mile is proportional to the expected fuel costs of the vehicle. Thus, the coefficients α and β_D represent the disutility of the reduction in income from purchasing the vehicle, including the up-front price of the vehicle and future fuel costs (maintenance costs are included in ξ_j). Equation (2) allows power-to-weight and weight to enter the utility function separately, whereas many other studies, such as Petrin (2002), omit weight. The parameter ξ_j includes all unobserved characteristics of the vehicle. The error term represents the consumer-specific shock to the consumer's utility from purchasing the vehicle.

Demand follows a nesting structure, with the outside good defined as the purchase of a used vehicle.[13] We define eight classes based on the vehicle classification system available in the

[12] This finding is consistent with Greene (1987), who concludes that about half of the increase in fuel economy between 1978 and 1985 was due to technology.

[13] Defining the outside good as a used vehicle is a simplifying assumption that does not affect the estimation.

Ward's database (McManus 2005). Under this nesting structure, the error term includes a common shock for all vehicles within a class, and an idiosyncratic term:

$$\varepsilon_{ij} = \omega_{ic} + (1-\sigma)\eta_{ij},$$

where ω_{ic} is the shock for vehicles in class c; σ is the similarity coefficient, which represents the extent to which consumers receive similar shocks to individual vehicles within a class; and η_{ij} is the idiosyncratic shock for consumer i and vehicle j. The class shock for consumer i is the same for all vehicles within nest c. Intuitively, consumers first decide whether to purchase a new vehicle, and then select a class, and finally, a vehicle model. Following Berry (1994), the market share of each vehicle model can be expressed as:

$$ln\, s_j - ln\, s_0 = \alpha\, ln\, p_j + \beta_D D_j + \beta_H HW_j + \beta_W W_j + \sigma\, ln\, s_{j|c} + \xi_j. \quad (3)$$

The left-hand side of equation (3) is the difference between the log market share of vehicle model j and the log market share of the outside good; the denominator in the market shares includes total new and used vehicle purchases.

There are trade-offs with using the nested logit framework. On the one hand, equation (3) can be estimated by linear regression models, and the nesting structure partially relaxes the independence of irrelevant alternatives (IIA) assumption in the basic logit model. Furthermore, the structure is consistent with the empirical strategy discussed below. On the other hand, equation (3) imposes restrictions on cross-price demand elasticities, and it does maintain the IIA assumption for models within a nest. Moreover, we have assumed that the coefficients in the utility function do not vary across consumers.[14]

5.2 The Supply of New Vehicles

The supply side of the model is static, following BLP. The model contains one period, in which a set of multiproduct firms select vehicle prices and the characteristics of the vehicles they sell. The end of this section discusses the implications of the main assumptions.

[14] The assumption is made for practical reasons. In our case it is not possible to estimate the larger number of coefficients in a model with heterogeneity because of the limited number of available instruments.

Firms compete in a Bertrand–Nash manner, choosing optimal prices and characteristics while taking as given the prices and characteristics of vehicles sold by other firms. Each firm is subject to the CAFE constraint, which requires that the harmonic means of its car and truck fleets exceed particular thresholds. If the firm does not satisfy the constraint it must pay a fine. We assume that, in equilibrium, the constraint is satisfied exactly, but we relax this assumption in the simulations.

As a benchmark against which the medium-run model can be compared, we first specify the firm's optimization problem in a short-run model (designated SR). This represents the standard case, which is considered by BLP and many others. Vehicle characteristics are exogenous and the firm chooses the vector of prices of its set of vehicles J_f:

$$\max_{\{p_j\}} \sum_{j \in J_f} (p_j - c_j) q_j [p_j, H_j, W_j, M_j, \xi_j] \qquad \text{(SR)}$$

$$s.t. \sum_{j \in J_f} q_j \left[p_j, H_j, W_j, M_j, \xi_j\right] / C_j \geq \sum_{j \in J_f} q_j \left[p_j, H_j, W_j, M_j, \xi_j\right] / M_j$$

where q_j is vehicle sales, H_j is horsepower, M_j is fuel economy, and c_j is the marginal cost of the vehicle. The parameter C_j is the CAFE standard that applies to vehicle model j. Quantity is a function of price, observed characteristics (horsepower, weight and fuel economy), and unobserved characteristics.

In the medium-run model (designated MR) firms choose prices and characteristics:

$$\max_{\{p_j, H_j, W_j, T_j\}} \sum_{j \in J_f} (p_j - c_j\,[H_j, W_j, T_j]) q_j [p_j, H_j, W_j, M_j, \xi_j] \qquad \text{(MR)}$$

$$s.t. \sum_{j \in J_f} q_j \left[p_j, H_j, W_j, M_j, \xi_j\right] / C_j \geq \sum_{j \in J_f} q_j \left[p_j, H_j, W_j, M_j, \xi_j\right] / M_j \qquad \text{(a)}$$

$$\ln M_j = \delta_0 + \delta_1 \ln H_j + \delta_2 \ln W_j + \ln T_j \qquad \text{(b)}$$

$$\ln c_j = \gamma_{0j} + \gamma_1 \ln H_j + \gamma_2 \ln W_j + \gamma_3 \ln T_j \qquad \text{(c)}$$

Equation (a) is the same CAFE constraint as in the short-run model. Equation (b) is a technological equation that represents the medium-run trade-off between fuel economy, power, and weight. According to the equation, the fuel economy of vehicle model j depends on the engine's horsepower, the vehicle's weight, and the level of the engine technology.[15] The level of technology, T, is a scalar that summarizes all of the components that increase the vehicle's fuel economy without affecting weight or power. Real world examples of power train components that affect T include engine calibration, piston compression rate, maximum speed, and variable valve timing. To simplify the numerical simulations, we approximate the effect of changes in these components on fuel economy using the linear function in (b). The technology variable is scaled so that a 1 percent increase in T causes a 1 percent increase in fuel economy. For example, suppose that adding variable valve timing raises the fuel economy of a particular model by 2 percent. The addition would be represented by an increase in T of 0.02.

Adding variable valve timing would also increase the marginal cost of producing the vehicle. Equation (c) describes the medium-run relationship between the marginal cost of production and the characteristics of the vehicle. In equation (c), it is assumed that any proportional increase in T has the same effect on marginal costs, as represented by γ_3. Note that in this setting, fuel economy indirectly affects marginal costs. If a firm wants to increase fuel economy, it has to decrease power or weight or increase T, all of which affect fuel economy. In other words, firms choose power, weight, and T, which determine fuel economy and marginal costs, which in turn enter the profit function and CAFE constraint. We could eliminate fuel economy and marginal costs in the profit function and CAFE constraint using (b) and (c), but we show the equations separately for clarity.

The trade-offs associated with increasing fuel economy can be understood by examining the demand equation (3) and the constraints in (MR). Decreasing power would increase fuel economy according to (b), relax the CAFE constraint, and decrease marginal costs according to (c). The net effect on demand is ambiguous because the decrease in power reduces demand, whereas the increase in fuel economy increases demand. Decreasing weight would increase fuel economy and marginal costs and would relax the CAFE constraint, but the effect on demand is similarly ambiguous and depends on demand for power-to-weight and weight. Increasing T would raise fuel economy according to (b), which increases demand and relaxes the CAFE

[15] Equation (b) is similar to equation (1) above, with the subscript jey in equation (1) replaced by j.

constraint, but also increases marginal costs. The unobserved characteristics, ξ_j, are exogenous to the firm and thus include variables that the firm cannot choose directly, as well as variables that can only be chosen in the long run.

In equilibrium, firms choose the profit-maximizing vectors of prices and observed vehicle characteristics. Consumers choose vehicles based on the prices and on the observed and unobserved characteristics. The first-order conditions to (MR) imply that the prices and observed characteristics are correlated with the unobserved characteristics in equilibrium.

Before turning to estimation, we discuss the implications of the main assumptions. First, the use of a static model is common in the literature and greatly simplifies the simulations. We do not include entry decisions, which would require an analysis of entry costs and production line decisions, which are beyond the scope of the paper. The exogeneity of ξ_j is necessary for the simulations in Section 6, although the assumption is relaxed partially in the demand estimation below. We model the variables that directly affect fuel economy, which suggests that this assumption does not have a large effect on the simulation results. However, in principle, the firm could change ξ_j, which would affect the demand for the vehicle, and indirectly affect the costs of CAFE.

5.3 Demand Estimation

The data used to estimate demand span the model-years 2000–2008 (see Section 2). Consequently, we introduce time subscripts for the discussion of demand estimation, where y denotes the model-year. Thus, the demand equation becomes:

$$\ln s_{jy} - \ln s_{0y} = \alpha \ln p_{jy} + \beta_D D_{jy} + \beta_H HW_{jy} + \beta_W W_{jy} + \sigma \ln s_{jy|c} + \xi_{jy}. \tag{4}$$

Market shares are calculated by model and model-year. The Ward's data contain the MSRP, fuel economy, power, and weight for each model and model-year. Dollars-per-mile is the ratio of the national average real retail price of gasoline over the model-year to the vehicle's fuel economy. The ratio is proportional to expected fuel costs under the assumption that the price of gasoline follows a random walk.

We allow for the possibility that ξ_{jy} includes characteristics chosen by the firm. The central challenge to estimating the parameters in equation (4) is that ξ_{jy} is potentially correlated with the vehicle price and other characteristics. For example, the firm may choose a higher price and greater horsepower for a vehicle that consumers perceive as being "sporty," or of higher "quality." More formally, in the medium run supply model, the first-order conditions for price and characteristics include ξ_{jy}, so that the observed characteristics and price are functions of ξ_{jy}.

5.3.1 Instrumental Variables Strategy for Estimating Demand Parameters

Because of the correlation between the observed characteristics and ξ_{jy}, estimating equation (4) by ordinary least squares (OLS) would yield biased estimates of all coefficients. Similarly, three standard approaches would also yield biased estimates. First, including vehicle fixed effects would only address the problem if one assumes that unobserved characteristics do not change over time (i.e., $\xi_{jy} = \xi_j$). In that case, the parameters would be identified by within-model changes in prices, power, and weight. We suggest that this assumption is not appropriate because many of the unobserved characteristics, such as reliability, can change over time.

The second approach would be to follow many previous studies of automobile demand, such as BLP, and use moments of vehicle characteristics of other vehicles in the same class or other vehicles sold by the same firm to instrument for the price and within-class market share. Such instruments are valid if the measured characteristics are exogenous. In that case, the instruments would be correlated with vehicle prices (via first-order conditions in model SR), but not with the unobserved characteristics. That condition, however, does not apply in the medium-run model. A similar problem applies to the third standard approach, the application of a hedonic analysis (e.g., McManus 2005).

Given these issues, we use an estimation strategy that is similar in spirit to Hausman *et al*. (1994), in that it takes advantage of common cost shocks across subsets of the market. Methodologically, the difference between this study and Hausman *et al*. is that we exploit the technological relationships across vehicle models sold by the same firm, which allows us to use the characteristics of other vehicle models to instrument for a model's characteristics and prices; by contrast, Hausman *et al*. instrument solely for prices using prices of the same product observed in other markets.

In practice, vehicle models from different product classes often share the same engine platform. This is common for SUVs and pickup trucks, but is not confined to those classes; Section 5.3.2 documents the prevalence of this relationship across the entire market. As a result, when vehicles in different classes share the same engine platform, they have very similar engine characteristics. For example, pickup trucks of the Ford F-Series come with the same engine platform as the Ford Excursion, an SUV. Therefore these two vehicle models have very similar fuel economy and power characteristics.

To illustrate, consider two vehicle models, j and j', which are in different vehicle classes but have engines e and e' that belong to the same engine platform. The profit-maximizing power of vehicle j depends on the engine platform technology, consumer demand, and characteristics of other vehicles in the same class. Therefore, the power of vehicle j will be a function of the power of vehicle j', plus a constant:

$$H_{jec} = f(H_{j'e'c'}) + \eta_c \tag{5}$$

The power of the two vehicles is correlated because they have the same engine. The class intercept, η_c, is unobserved and is an arbitrary, potentially nonlinear, function of the characteristics of other vehicles in the same class as vehicle j, as well as non-engine characteristics of the same vehicle. It allows for class-specific demand and supply shocks, so that the power of the two vehicles j and j' will differ because of variation across classes in consumer preferences and the characteristics of the other vehicles in the respective classes.

Our IV strategy is based on equation (5), in which we instrument for a vehicle's price, dollars-per-mile, power-to-weight, weight, and within-class market share in equation (4). The instruments are the means of seven characteristics of vehicle models that share the same engine platform, but are located in other classes.[16] The variables are fuel economy, power, weight, torque, the number of valves, the number of cylinders, and displacement.

The rationale for the IV strategy is that, for technological reasons, engine characteristics can vary within a platform, but only to a limited extent. Therefore, the observed engine characteristics of one vehicle are likely to be highly correlated with the observed engine characteristics of another vehicle that shares the same platform. They are not identical, however, and we argue that the variation derives from unobserved class-specific cost and demand shocks. For example, consider a hypothetical SUV and pickup truck that share the same platform. Owners of the pickup truck may have a greater demand for power than owners of the SUV. This

[16] We prefer to construct the instruments using engine platforms rather than engine programs because the sample size is much larger and the instruments for a particular vehicle are constructed from a wider range of vehicles. Note that the results are sensitive to this distinction, however, as we find the demand for power to be small and not statistically significant when using instruments based on engine programs. See Section 7 for a more extensive discussion of the robustness analysis.

will result in the pickup having greater power. Below, we document the extent of the within-platform variation in characteristics.

More precisely, the exclusion restriction is that the engine characteristics of one vehicle do not affect the demand for another vehicle in a different class that has the same engine, independent of the common technology. For example, the characteristics of the Ford Fusion (a midsize car) do not affect the demand for the Ford Ranger (a pickup truck that shares the same engine as the Fusion). Note that this assumption is considerably weaker than the standard assumption that observed and unobserved characteristics are uncorrelated.

Although this approach relaxes the assumption that vehicle characteristics are exogenous, several potential sources of bias remain. First, the assumption that the engine characteristics of one vehicle do not affect demand for vehicles in another nest is not strictly true in the nested logit framework. Demand is correlated across nests, although the correlation is second-order in magnitude compared to the within-nest correlation. For example, the within-nest cross-price elasticity of demand is typically several orders of magnitude greater than the cross-nest, cross-price elasticity of demand. Because of this potential source of bias, Section 7 reports a wide range of robustness checks.

Second, there may be unobserved company-specific fixed effects or trends, such as perceived quality. To address this concern, the specification we estimate includes company–year interactions; for example, such an approach would be robust if all Honda models share common unobserved characteristics.

Third, the estimates would be biased if equation (4) omitted some engine characteristics. However, we believe that the variables included in equation (4) capture the main features that consumers use to differentiate engines. Section 7 reports specifications with additional vehicle characteristics.

A final concern is that a firm may choose the unobserved characteristics of the vehicle in response to a demand or cost shock. For example, a firm may change its marketing strategy based on the successful entry of another firm’s vehicle model. Note that our demand estimation is robust to this possibility as long as the underlying shocks are uncorrelated across nests and within a platform. The same argument applies to the consideration that in many cases, firms link certain unobserved characteristics (e.g., a sunroof) to the presence of a certain power train. That is, the firm chooses which platform is offered with which vehicle. In that case, the instruments are valid even if ξ_{jy} is chosen by the firm, as long as the underlying shocks that affect ξ_{jy} are uncorrelated across nests; offering the sunroof depends on the demand for vehicles in the same nest, and not on the engine platform.

Thus, the instruments would be invalid in the presence of unobserved, cross-nest, platform-level shocks. We address this possibility by using lagged values of the instruments in Section 7. Although engine platforms are redesigned roughly every ten years, cost or demand shocks may cause the firm to change the engine characteristics between redesigns. Using lagged values of the instruments would reduce bias if unobserved platform-level cost or demand shocks are not highly persistent.

5.3.2 Within–Engine Platform Variation and First-Stage Estimates

Before reporting the main estimates, we summarize the engine variation across vehicle models that underlies the first stage and discuss the first-stage estimates for equation (4). Each row in Table 3 includes a different vehicle class. Column 1 shows the total number of vehicle models in 2008, and column 2 shows the number of vehicle models in the sample. The sample includes only those vehicles that have an engine platform that is also found in a vehicle from a different class—in other words, those for which the instruments can be constructed. Only about two-thirds of the vehicle models are in the sample, but column 4 shows that the sample includes 89 percent of total sales. Furthermore, except for small cars, the sample includes nearly all of the sales for each class. Note that it would be possible to increase the sample size by defining narrower vehicle classes. However, there is a trade-off between sample size and bias as demand shocks are more likely to be correlated across classes when classes are narrower, possibly invalidating the IV approach.

Table 4 reports summary statistics for the dependent variable and five endogenous right-hand-side variables in equation (4). Appendix Table 1 reports the first-stage estimates, in which the endogenous variables are regressed on the instruments. All specifications include company-year interactions and the reported engine-based instruments. The instruments are jointly strong predictors of the endogenous variables, as indicated by the high F statistics at the bottom of the table. Note that some of the coefficients are counterintuitive. For example, the coefficient on power in the power-to-weight regression is negative, but this reflects the high correlation between power and torque.

Figure 5 provides information about the extent of variation of the instruments. To construct the figure, we determined the engine platform most commonly sold with each model. Each vehicle model and model-year represents a unique observation. For each observation, the figure plots the vehicle characteristic on the horizontal axis and, on the vertical axis, the average value of that characteristic across other vehicles that have the same platform but belong to different classes. That is, the vertical axis shows the variables used to construct the instruments. As a benchmark, in the absence of within-platform variation, the data points would all lie along the 45 degree line—that is, for each model and platform, the value of the characteristic would exactly equal the average value for models in other classes that share the platform. Note that the figures use the most common engine platform, as opposed to the instruments, which are computed as the average across platforms sold with the model. The figure shows that, although the characteristics are highly correlated within platforms, there is considerable within-platform variation for all characteristics, even for the number of cylinders.

5.3.3 Estimation of Demand Parameters

Table 5 reports the demand estimates from equation (4). The dependent variable is the log of the vehicle model's sales, and the independent variables are the log price of the vehicle, dollars-per-mile, power-to-weight, weight, the within-class market share, and a set of company–year interactions.[17]

Column 1 reports the OLS estimates of equation (4) for comparison with the IV estimates. The coefficient on the log price of the vehicle is statistically significant but is small in magnitude, as the implied average own-price elasticity of demand is –0.30. The coefficient on dollars-per-mile is not statistically significant and the sign is not consistent with theory. The coefficient on power-to-weight is positive, as expected, but is not statistically significant. The price coefficient is probably biased toward zero because the price is expected to be positively correlated with unobserved characteristics, but the direction of the bias for the observed characteristics is ambiguous as they may be positively or negatively correlated with unobserved characteristics.

Column 2 of Table 5 reports a specification that follows the previous literature (e.g., BLP) and uses other characteristics as instruments—in particular, the sum of the characteristics

[17] Because there is no cross-sectional variation of used vehicle purchases, the year dummies absorb the denominator in the left-hand side variables in equation (4). Consequently, the dependent variable is the log sales of the new model.The regression model used in this paper differs from that in a previous version of the paper in several ways. First, we use the log vehicle price instead of the price level. Second, we previously estimated demand in two stages. In the first stage, we used monthly sales and gasoline price data to estimate the coefficient on dollars-per-mile. The first stage included model by model-year interactions, which we used to estimate the remaining demand parameters in the second stage, as suggested by Nevo (2000). The parameter estimates presented here are similar to the previous version.

of other vehicles in the same class and the sum of characteristics of other vehicles sold by the same firm. The coefficient on the vehicle's price is larger in magnitude than the OLS estimate and implies an average elasticity of demand of –5.4. Both the coefficient on dollars-per-mile and the coefficient on power-to-weight have the expected sign and are statistically significant at the five percent level. The estimates imply that consumers value similarly an increase in fuel economy and a proportional increase in power.

Column 3 reports the baseline specification using our engine-based instruments. The estimated own-price elasticity of demand, –1.40, is smaller than the BLP specification. The coefficient on dollars-per-mile is quite similar to the estimate in column 2, although it is not statistically significant (p-value of 0.18). The coefficient on power-to-weight is much larger than in column 2, and is significant at the one percent level. The point estimates imply that the willingness to pay for horsepower is greater than the willingness to pay for a proportional increase in fuel economy. This result is consistent with Figures 2 and 4, which show that, as power train technology has improved over time, firms have increased power and weight while keeping fuel economy constant.

Figure 6 summarizes the magnitudes of the engine-based IV estimates in Table 5. If the willingness to pay for power and weight is sufficiently large relative to fuel economy, the decrease in power and weight in the late 1970s and 1980s for U.S. cars would have reduced consumers' willingness to pay for these vehicles. Figure 6 plots the change in willingness to pay for the average car sold by U.S. firms from 1975 to 2007, using the characteristics in Figure 2 and the estimates from column 3 of Table 5, and holding vehicle prices fixed. The figure shows that willingness to pay was flat soon after CAFE was implemented, but increased steadily beginning around 1980.[18] Note that Figure 6 does not allow for an inference about the causal effect of CAFE on willingness to pay, yet it is useful for summarizing the magnitudes in Table 5.

5.4 Supply Estimation

Finally, the parameters in the technology and cost equations in (MR) must be estimated. Equation (b) represents the trade-off between fuel economy, weight, and power. The coefficients on weight and power are estimated in Section 4, with the results reported in Table 2.

[18] Greene and Liu (1988) perform a similar analysis and reach a similar conclusion using estimates of willingness-to-pay for characteristics from other studies conducted in the 1970s and 1980s.

In the marginal cost equation, we estimate the coefficients on weight and T from NRC (2008). This report estimates that, for the average vehicle in the sample, reducing weight by 1 percent raises production costs by 0.25 percent. The report also contains estimates of the costs and fuel economy improvements of specific engine and transmission technologies, some of which are listed in Table 1. Taking each technology in the report as an independent observation, we estimate a linear regression in which the percentage change in production costs is the dependent variable and the percentage change in fuel economy is the independent variable.[19] The coefficient estimate is 0.12 (standard error 0.04), which implies that an increase in T_j that increases fuel economy by 1 percent would increase marginal costs by 0.12 percent.[20] In addition, we calculate the constant term in equation (c) from the first-order conditions for vehicle prices. This ensures that the first-order conditions hold at the initial equilibrium.

We estimate the cost of increasing horsepower from proprietary data on engine production costs provide by CSM. We estimate that increasing horsepower by 1 percent raises marginal costs by 0.2 percent.

6. Short- and Medium-Run Effects of the CAFE Standard

This section uses the model and parameter estimates from Section 5 to compare the short- and medium-run costs to vehicle producers and consumers of the CAFE standard. We simulate the equilibrium under a 1-mpg increase in the CAFE standard for cars and light trucks. This section reports the main results and the next section reports additional robustness results.

6.1 Short Run

For comparison with the previous literature and with our medium-run analysis, we first simulate the short-run effects of the CAFE standard. Section 5.2 summarizes the model, in which

[19] Treating each technology as being independent may not be appropriate if the ability to use one technology depends on the presence or absence of another technology. We introduce this simplification to maintain consistency with the simulation model.

[20] We note that the underlying cost estimates for vehicle technology, weight and horsepower are engineering-based estimates, and may not represent the true marginal cost.

firms choose a vector of prices to maximize profits.[21] Following Jacobsen (2008), we separate firms into three categories: unconstrained firms that exceed the standard, constrained firms that meet the standard, and firms that pay the fine for not meeting the standard. Assignments to the three categories are based on firms' past behavior. Honda, Toyota, and several smaller Asian firms have consistently exceeded the standard by a wide margin and are unconstrained; Chrysler, Ford, GM, and a few other firms have generally been close to the standard and are constrained; and all other firms have been well below the standard. Constrained firms solve problem (SR) (see Section 5.1), whereas the other firms do not have a constraint. In performing the simulations, we assume that firms do not change categories as a result of the increase in the standard, which we believe is a reasonable assumption for the modest increase in the CAFE standard considered here.

Table 6 shows the estimated effects of a 1-mpg increase in the CAFE standard.[22] The compensating variation of the new standard is -$6.51 billion, which represents a significant cost to consumers. Total profits decrease by $9.07 billion, which is comparable to other recent estimates (e.g., Jacobsen, 2010). Many previous studies find that consumers pay most of the costs of the policy, which makes the results found here somewhat surprising; we return to this issue below when discussing the medium-run results.

The third and fourth columns show that the increase in the standard significantly reduces the profits of constrained firms (mostly U.S.-based automakers), and has no effect on the profits of unconstrained firms (i.e., firms that exceed the standard and do not pay the fine). The omitted category is the constrained firms that pay the fine, whose profits increase. This reflects the substitution of consumers away from the models sold by constrained firms and towards the larger

[21] The new CAFE standard, effective in model-year 2011, will be computed for each firm based on the footprint of the vehicles sold by the firm. The footprint-based standard was designed to raise the fuel economy of each vehicle model. To the extent that larger vehicles will face a less stringent CAFE standard under the new rules, reductions in weight will be less beneficial for meeting the constraint. In our paper, we use the approach of the initial CAFE standard, which applies one fuel economy standard for cars and a separate one for light trucks. We do not model the trading of CAFE credits; although not yet permitted, trading could reduce the total cost of the program.

[22] The simulations do not include vehicles that are not in the estimation sample, and thus, they do not include a significant share of small cars. This does not seem to significantly affect the magnitudes of the welfare estimates, however. The results are similar if we include all vehicles in the market and assume that the demand parameters for the vehicles in the sample are the same as for vehicles that are not in the sample. Also, recall that the price elasticity of demand is smaller than many other estimates in the literature; as expected, the estimated cost of CAFE to consumers is significantly lower if we use a larger price elasticity in the simulations.

vehicles sold by firms that pay the fine. This consumer-substitution effect outweighs the fact that the latter firms pay a larger fine under the new standard.

The next two columns report the change in overall fuel economy and the fuel economy of unconstrained firms (both are weighted by sales). Overall fuel economy increases by less than 1 mpg for two reasons. The first is that only a subset of the firms is constrained by the standard. The second is that the unconstrained firms decrease their average fuel economy (i.e., decrease prices of vehicles with low fuel economy) in an effort to attract consumers that previously purchased the low-fuel economy vehicles of the constrained firms.

The last two columns show that the changes in the sales mix significantly reduces the sales-weighted average weight and power of vehicles sold. These changes are consistent with the large reduction in consumer welfare.

6.2 Medium Run

We also simulate the medium-run equilibrium under a 1-mpg increase in the CAFE standard. All firms choose prices and vehicle characteristics to maximize profits. In particular, constrained firms can increase fuel economy by modifying power trains as well as by reducing weight or power. We impose the constraint that constrained firms can increase fuel economy of each vehicle by no more than 0.25 mpg.

We assume that the unobserved product characteristics, ξ_j, are exogenous to the firm, and do not change in response to the change in the CAFE standard. Implications of this assumption were discussed above.

The second row of Table 6 reports summary statistics from a simulation of the medium-run effects of the standard. The simulation results demonstrate the importance of allowing for endogenous characteristics in estimating willingness to pay and in modeling the new vehicles market. The medium-run responses significantly reduce the negative effect of the CAFE standard on profits; on the other hand consumers are slightly worse off in the medium run than in the short run. The distribution of costs across firms and consumers is closer to previous studies that tend to find that consumers pay the majority of the costs.

We also find that the constrained firms respond to the higher standard by increasing T (power train technology). This response significantly affects the distribution of costs across firms, as the reduction in profits for constrained firms is much smaller in the medium run than in the short run because raising T is a less costly option than changing vehicle prices. (Note that the effect would be even more pronounced if we relax the constraint on the maximum fuel economy increase achieved by raising T.) By comparison, unconstrained firms and firms that pay the fine are worse off in the medium run than the short run because they are unable to attract as many consumers who bought vehicles from constrained firms before the standard increased. These results can be seen in the columns showing the profits of constrained firms, the profits of unconstrained firms, and the fuel economy of unconstrained firms.

The decrease in weight and power in the medium run arises from changes in market shares, and not because constrained firms reduce the weight or power of individual vehicle models. Given the parameters related to the demand, cost and technological tradeoffs, reducing weight and power does not result in a sufficient increase in fuel economy to justify the lower demand that would result. However, the reduction in sales-weighted power and weight is smaller in the medium run than in the short run because the constrained firms increase T rather than relying exclusively on changes in the sales mix. We note that that this result may not hold for larger increases in the standard, which could induce firms to reduce power or weight.

The results suggest that the medium run cost of CAFE is significantly lower, but they do not necessarily overturn the conclusions of the previous literature that the gasoline tax is much less costly than the CAFE standard. The drawbacks associated with CAFE—including the *rebound effect* and the fact that CAFE gradually affects fuel consumption via changes in the composition of the set of vehicles in use—are still present in the medium run. A full comparison of the two policies is beyond the scope of this paper, however.[23]

7. Robustness of Demand Estimates

Table 7 reports a number of robustness checks for equation (4). We focus on the coefficient estimates for log vehicle price, dollars-per-mile, and power-to-weight. Column 1 shows that the coefficients are larger than the baseline estimates (column 3 of Table 5), although they are similar using brand–year interactions in place of company–year interactions (for example, Honda includes the Honda and Acura brands).

[23] The ability to improve fuel economy in the medium run may reduce the cost of the gas tax as well, but this effect is likely to be small in magnitude because fuel economy of new vehicles is a relatively small component of the total effect of the gasoline tax on gasoline consumption.

Column 2 shows that the standard errors increase if they are clustered by vehicle model. Column 3 includes the lagged dependent variable and shows evidence of serial correlation; nonetheless, the coefficient estimates are broadly similar to the estimates in Table 5.

Columns 4–6 show specifications in which the characteristics that enter the utility function are different. Column 4 separates power and weight, column 5 adds vehicle length to column 4, and column 6 adds engine displacement to column 4. Note that it is not possible to separately identify the demand for displacement, which probably arises from the high correlation between power and displacement; similar results obtain for torque (not reported).

Column 7 reports a specification that addresses the possibility that unobserved characteristics varying by platform are correlated with the observed characteristics. The specification uses the three-year lags of the instruments. The coefficient on dollars-per-mile is smaller than in Table 5, but the other coefficients are quite similar. This specification thus provides some evidence against the presence of bias due to unobserved platform characteristics that vary by platform.

Overall, the coefficient estimate for power-to-weight appears to be somewhat more robust than dollars-per-mile. For comparison, Appendix Table 2 shows the same specifications with the BLP instruments in place of the engine-based instruments (omitting the lagged instrument specification). The coefficients vary across specifications, although perhaps less than with our engine-based instruments. The BLP instruments appear to yield more stable coefficient estimates, but there is a fair bit of variation across specifications in either case, something we believe has not been properly emphasized in the vehicle demand literature.

8. Conclusion

The upcoming increase in the CAFE standard will significantly affect the new vehicles market. We demonstrate the importance of accounting for the effect of the regulation on vehicle characteristics, both in modeling firms' responses to the regulation and in estimating consumer willingness to pay. The focus is on the medium-run effect of CAFE, a time frame of 4–5 years, during which firms can choose vehicle characteristics and modify power trains. This paper first shows that, in response to the initial standard, firms significantly reduced the power and weight of vehicles sold in the late 1970s and early 1980s to increase fuel economy, but technological progress caused weight and power to recover in the long run.

We then estimate consumers' willingness to pay for fuel economy, power, and weight to analyze quantitatively the medium-run effects of the CAFE standard. Estimating willingness to pay is complicated by the fact that firms select vehicle characteristics endogenously—an issue that previous empirical work has not addressed. We propose an IV strategy that accounts for time-varying unobserved characteristics and the endogeneity of observed characteristics. The estimates suggest that consumers value an increase in power more than a proportional increase in fuel economy. The parameter estimates are significantly different from those obtained using standard estimation approaches.

The simulation results demonstrate the importance of allowing for endogenous characteristics in estimating the cost of CAFE, as we find that the medium-run costs are substantially lower than the short-run costs. The results also suggest that firms can attain larger improvements in fuel economy in a shorter amount of time than is implied by a long-run analysis.

A few limitations of the analysis should be noted. First, it is assumed that unobserved characteristics do not change in response to the increase in the CAFE standard. Second, the policy scenario discussed above considers the medium-run effect of the CAFE standard, without entry of new vehicles. Third, there may be other dynamic considerations if, for example, medium-run decisions affect long-run decisions. Addressing these limitations would require the use of a dynamic model, the estimation of entry costs, and the modeling of production line decisions; these could be subjects for future research. Finally, as Section 3 notes, the upcoming CAFE standard will be footprint-based, meaning that the fuel economy standard will depend on the size of the vehicle. Also, the new framework will allow CAFE credit trading across firms. Incentives for improving the fuel economy of individual vehicle models may be different under the new framework, a possibility that we may also explore in future work.

References

Atkinson, S. 1981. Rising Gasoline Prices and Federal Automotive Efficiency Standards: Their Impact on Consumer Choice. Research Study #23. Washington, DC: American Petroleum Institute.

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

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. 2006. Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes. American Economic Review 99.

Berry, Steven. 1994. Estimating Discrete Choice Models of Product Differentiation. RAND Journal of Economics 25: 242–262.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. Automobile Prices in Market Equilibrium. Econometrica 63(4): 841–890.

Foster, Taft and Thomas Klier. 2009. Raising Automotive Fuel Efficiency. Chicago Fed Letter September 2009, Number 266. Chicago, IL: The Federal Reserve Bank of Chicago.

Goldberg, Penelope Koujianou. 1998. The Effects of the Corporate Average Fuel Efficiency Standards in the U.S. The Journal of Industrial Economics 46(1): 1–33.

Gramlich, Jacob. 2009. Gas Prices, Fuel Efficiency, and Endogenous Product Choice in the U.S. Automobile Industry. Washington, DC: Georgetown University.

Greene, David L. 1987. Advances in Automobile Technology and the Market for Fuel Efficiency, 1978–1985. Transportation Research Record 1155: 18–27.

Greene, David L. 1991. Short-Run Pricing Strategies To Increase Corporate Average Fuel Economy. Economic Inquiry 29: 101–114.

Greene, David L. and Jin-Tan Liu. 1988. Automotive Fuel Economy Improvements and Consumers' Surplus. Transportation Research A 22-A: 203–218.

Hausman, Jerry, Gregory Leonard, and J. Douglas Zona. 1994. Competitive Analysis with Differentiated Products. Annales D'Economique et de Statistique 34: 159–180.

Hill, Kim, Morgan Edwards, and Steven Szakaly. 2007. How Automakers Plan Their Products—A Primer. July. Ann Arbor, MI: Center for Automotive Research.

Ingrassia, Paul. 2008. Detroit's (Long) Quest for Fuel Efficiency. The Wall Street Journal, February 19, A19.

Ishii, Joy. 2005. Compatibility, Competition and Investment in Network Industries: ATM Networks in the Banking Industry. Stanford, CA: Stanford University.

Jacobsen, Mark. 2008. Evaluating Fuel Economy Standards in a Model with Producer and Household Heterogeneity. La Jolla, CA: University of California, San Diego.

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

Klier, Thomas, and Joshua Linn. 2010. The Price of Gasoline and the Demand for Fuel Economy: Evidence from Monthly New Vehicles Sales Data. American Economic Journal: Economic Policy August: 120–141.

Kranz, Rick. 2008. GM Cars To Get Smaller Engines with Turbos. Automotive News. http://www.autonews.com/apps/pbcs.dll/article?AID=/20080128/OEM01/301289931 (accessed October 22, 2010).

McManus, Walter S. 2005. The Effects of Higher Gasoline Prices on U.S. Light Vehicle Sales, Prices, and Variable Profit by Segment and Manufacturer Group, 2001 and 2004. Ann Arbor, MI: University of Michigan Transportation Research Institute.

National Highway Traffic Safety Administration. 2008. Preliminary Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2011–2015 Passenger Cars and Light Trucks. Washington, DC: U.S. Department of Transportation.

Nevo, Aviv. 2000. A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. Journal of Economics and Management Strategy 9: 513–548.

NRC (National Research Council). 2008. Interim Technology Assessment. Letter to National Highway Traffic Safety Administration from the National Research Council Committee on Assessment of Technologies for Improving Light-Duty Vehicle Fuel Efficiency. February 14.

Ohnsman, Alan. 2008. Honda Bucks Industry with TSX Favoring Fuel Efficiency over Power. Bloomberg. February 15.

Petrin, Amil. 2002. Quantifying the Benefits of New Products: The Case of the Minivan. Journal of Political Economy 110: 705–729.

Shiau, Ching-Shin Norman, Jeremy J. Michalek, and Chris T. Hendrickson. 2009. A Structural Analysis of Vehicle Design Response to Corporate Average Fuel Economy Policy. Transportation Research Part A 43: 814–828.

Stoffer, Harry. 2007. Senate OKs CAFE Hike. *Automotive News*, December 14. http://www.autonews.com/apps/pbcs.dll/article?AID=/20071214/ANA02/71213013.

Sweeting, Andrew. 2007. Dynamic Product Repositioning in Differentiated Product Markets: The Case of Format Switching in the Commercial Radio Industry. NBER Working Paper #13522. Cambridge, MA: National Bureau of Economic Research.

U.S. EPA (Environmental Protection Agency). 2007. Light-Duty Automotive Technology and Fuel Efficiency Trends: 1975 through 2007. EPA 420-R-07-008. Washington, DC: U.S. EPA.

Yacobucci, Brent. 2010. EPA and NHTSA: The New Auto Greenhouse Gas and CAFE Standards. Presentation at automotive industry conference. May 11, Detroit, MI. http://www.chicagofed.org/digital_assets/others/events/2010/automotive_perfect_storm/2010_Detroit_autoconference_Yacobucci.pdf (accessed September 29, 2010).

Tables and Figures

See following pages.

Table 1

Examples of Medium and Long Run Engine and Transmission Changes

Medium run			Long run		
Technology	Cost ($)	Percent increase in mpg	Technology	Cost ($)	Percent increase in mpg
Low friction lubricants	3	0.5	Turbocharge/ downsize	120	5-7.5
Variable valve timing	59-209	1-3	Continuously variable trans	139	3.5
5-speed automatic transmission	76-167	0.5-2.5	Automatic manual transmission	141	4.5-7.5
Cylinder deactivation	203	4.5-6	PHEV	6750	28

Source: NHTSA (2008). All figures represent estimates for a mid-size car.

Table 2

Tradeoff Between Fuel Efficiency, Weight and Power for Cars

	Dependent variable: log fuel economy	
	(1)	(2)
Log horsepower	-0.06 (0.03)	-0.15 (0.03)
Log weight	-0.33 (0.07)	-0.33 (0.09)
R^2	0.90	0.84
Number of observations	1989	1989
Fixed effects	Engine program	Engine platform

Notes: Standard errors in parentheses, clustered by engine. Observations are by engine and year for 2000-2007. All specifications are estimated by Ordinary Least Squares. The dependent variable is the log of the fuel economy of the corresponding vehicle model. All columns include the log of the engine's power and the log of the vehicle model's weight. Column 1 includes engine program dummies and column 2 includes engine platform dummies.

Table 3

Sample Coverage by Vehicle Class, 2007

Vehicle class	(1) Number of vehicle models	(2) Number of vehicle models with instruments	(3) Fraction sales	(4) Fraction sales with instruments
Small cars	32	18	0.14	0.12
Mid-size cars	31	19	0.18	0.17
Large cars	14	8	0.05	0.03
Luxury and specialty cars	79	48	0.11	0.08
Small SUVs	53	39	0.16	0.14
Large SUVs	47	35	0.13	0.12
Vans	21	13	0.07	0.07
Pickup trucks	21	18	0.17	0.17
Total	298	198	1.00	0.89

Notes: Vehicles are assigned to the vehicle classes, which are defined in the Wards database. The number of vehicle models is the number of unique models in each class in the 2007 model-year. The number of vehicle models with instruments is the number of models for which there is another model that belongs to a different class and has the same engine. Fraction sales is the share of sales of vehicle models in the class in total sales in the 2007 model-year. Fraction sales with instruments is the fraction of sales in total sales for the vehicle models with instruments.

Table 4

Summary Statistics		
Variable name	Mean	Standard deviation
Log sales	10.446	1.467
Vehicle price	36.848	19.984
Dollars-per-mile	0.111	0.038
Power-to-weight	0.059	0.014
Weight	1.904	0.424
Log within-class share	-4.057	1.478

Notes: The table reports the mean and standard deviation of log sales, vehicle price (thousands of dollars), dollars-per-mile, power-to-weight (horsepower per pound), weight (tons) and the log of the within-class market share. The sample contains 1819 observations.

Table 5

Willingness-to-Pay for Power and Weight

	Dependent variable: log sales		
	(1)	(2)	(3)
Log vehicle price	-0.28 (0.05)	-1.86 (0.33)	-1.28 (0.31)
Dollars-per-mile	1.11 (0.88)	-12.96 (2.69)	-11.05 (8.40)
Power-to-weight	0.06 (1.03)	9.53 (4.74)	38.75 (9.51)
Weight	0.11 (0.06)	1.62 (0.28)	1.14 (0.53)
Log within-class share	0.92 (0.01)	0.34 (0.07)	0.91 (0.05)
R^2	0.94	0.74	0.89
Number of observations	1819	1819	1819
Estimation model	OLS	IV, BLP Instruments	IV, Engine Instruments

Notes: The table reports the results from estimating equation (4). Standard errors are in parentheses, and are robust to heteroskedasticity. The dependent variable is the log sales of the vehicle model. The independent variables are the price of the vehicle, in thousands of dollars; dollars-per-mile; power-to-weight, in horsepower divided by weight, in pounds; weight, in tons; the log of the within class share of sales; and a full set of company-year interactions. Column 1 is estimated by Ordinary Least Squares and columns 2 and 3 are estimated by Instrumental Variables. Column 2 instruments for vehicle price using the sum of characteristics of vehicle models in the same category produced by other firms and the sum of characteristics of other models produced by the firm. Column 3 uses as instruments the independent variables in Appendix Table 1.

Table 6

Effects of a 1 MPG Increase in the CAFE Standard

	Compensating variation (billion $)	Change in total profits (billion $)	Change in constrained firms' profits (billion $)	Change in unconstrained firms' profits (billion $)	Change in fuel economy (mpg)	Change in unconstrained firms' fuel economy (mpg)	Change in horsepower	Change in weight (pounds)
Short run	-6.51	-9.07	-12.51	-0.04	0.54	-0.09	-6.76	-87.13
Medium run	-7.81	-5.58	-7.01	-1.39	0.52	0.01	-5.42	-70.87

Notes: The table reports the compensating variation of a 1 mpg increase in the CAFE standard and the effects on total profits; profits of constrained firms; profits of unconstrained firms that do not pay the fine (all in billions of 2007 dollars); average fuel economy for the entire market (mpg); average fuel economy for unconstrained firms (mpg); power (HP); and weight (pounds). All vehicle characteristics are weighted by sales. The two rows report the results of different simulations. In the first row, weight, power and fuel economy of each vehicle model are held constant, while in the second row these characteristics are chosen by the firm. See text for details on the simulations.

Table 7

Alternative Specifications

	Dependent variable: log sales						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log vehicle price	-1.79 (0.51)	-1.28 (0.63)	-1.34 (0.55)	-0.99 (0.25)	-0.60 (0.30)	-1.06 (0.26)	-1.49 (0.42)
Dollars-per-mile	-13.24 (9.94)	-11.05 (20.64)	-22.94 (11.36)	-3.95 (6.26)	-0.98 (6.68)	0.43 (7.36)	-3.29 (13.70)
Power-to-weight	47.20 (10.70)	38.75 (19.41)	40.74 (15.35)				42.18 (11.84)
Power				0.008 (0.002)	0.005 (0.003)	0.008 (0.002)	
Weight	1.47 (0.68)	1.14 (1.20)	1.82 (0.81)	-0.25 (0.27)	-0.99 (0.33)	-0.18 (0.29)	0.82 (0.80)
Log within-class share	0.89 (0.09)	0.91 (0.10)	0.39 (0.20)	0.96 (0.04)	0.84 (0.05)	0.93 (0.04)	0.87 (0.06)
Lag dep var			0.45 (0.19)				
Length					0.42 (0.12)		
Displacement x 1000						-0.14 (0.12)	
R^2	0.89	0.89	0.87	0.91	0.60	0.92	0.89
Number of observations	1819	1819	1530	1819	1819	1802	1335
Spec	Brand-year interactions	Cluster std errors by model	Add lag dep var	Separate power, weight	Add length	Add displacement	Lagged instrument

Notes: The table reports the specifications indicated in the bottom row, using column 3 of Table 5 as the baseline. Standard errors are robust to heteroskedasticity, except in column 2 where standard errors are clustered by vehicle model. Column 1 includes brand-year interactions instead of company-year interactions. Column 3 includes the lag of the dependent variable. Column 4 includes weight and power separately. Column 5 adds vehicle length to column 4 and column 6 adds displacement to column 4. The displacement coefficient is multiplied by 1000. Column 7 uses the three-year lags of the instruments from the corresponding engine platform.

Appendix Table 1

First Stage Estimates

	Dependent Variable:				
	Vehicle price (thousand $)	Dollars-per-mile	Power-to-weight (horsepower / pound)	Weight (tons)	Log within-class share
Fuel economy	-12.11 (2.45)	-1.97 (1.46)	-6.21 (1.32)	0.63 (0.28)	-2.71 (1.26)
Power	-2.28 (0.63)	-1.19 (0.31)	-1.05 (0.36)	0.03 (0.06)	1.75 (0.29)
Weight	-174.61 (39.15)	6.82 (20.95)	-21.19 (21.07)	8.22 (4.15)	132.66 (17.48)
Torque	5.19 (0.55)	1.42 (0.38)	0.59 (0.21)	0.27 (0.06)	-1.66 (0.27)
Number of valves	16.40 (2.91)	3.05 (1.75)	9.36 (1.19)	-0.19 (0.32)	-8.49 (1.34)
Number of cylinders	42.48 (23.86)	-39.27 (15.19)	-7.88 (8.72)	-8.82 (2.61)	-27.93 (11.17)
Displacement	-0.08 (0.04)	0.17 (0.03)	0.02 (0.02)	0.02 (0.01)	0.03 (0.02)
F-Statistic	289.29	184.19	51.08	167.10	37.91
R^2	0.69	0.84	0.31	0.58	0.29
Number of observations	1819	1819	1819	1819	1819

Notes: Instruments for vehicle price, dollars-per-mile, power-to-weight, weight, and within-class market share are constructed from the matched engine model-vehicle model data set. The instruments are the means across vehicles belonging to other classes that have the same engine. The sample includes all models for which the instruments can be calculated, and spans 2000-2008. The table reports coefficient estimates with standard errors in parentheses. All regressions include brand-year interactions. Standard errors are robust to heteroskedasticity. For readability, log vehicle price is multiplied by 1,000, dollars-per-mile is multiplied by 10,000, power-to-weight is multiplied by 10,000, weight is multiplied by 100, and log within-class market share is multiplied by 100.

Appendix Table 2

Alternative Specifications with BLP Instruments

	Dependent variable: log sales					
	(1)	(2)	(3)	(4)	(5)	(6)
Log vehicle price	-1.59 (0.39)	-1.86 (0.60)	-0.43 (0.27)	-1.77 (0.33)	-1.77 (0.34)	-1.65 (0.31)
Dollars-per-mile	-14.11 (2.59)	-12.96 (4.32)	-13.08 (2.29)	-12.06 (2.53)	-12.52 (2.63)	-12.44 (2.44)
Power-to-weight	8.25 (5.36)	9.53 (9.13)	3.19 (3.90)			
Power				0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
Weight	1.47 (0.29)	1.62 (0.49)	0.90 (0.23)	1.27 (0.17)	1.28 (0.18)	1.28 (0.17)
Log within-class share	0.41 (0.07)	0.34 (0.12)	0.11 (0.06)	0.35 (0.07)	0.33 (0.07)	0.40 (0.06)
Lag dep var			0.70 (0.08)	0.35 (0.07)		
Length					0.018 (0.008)	
Displacement						-0.075 (0.032)
R^2	0.82	0.74	0.84	0.75	0.73	0.78
Number of observations	1819	1819	1530	1819	1819	1802
Spec	Brand-year interactions	Cluster std errors by model	Add lag dep var	Separate power, weight	Add length	Add displacement

Notes: The table reports the specifications indicated in the bottom row. Each column reports the same specifcation as the corresponding column in Table 7, except that the BLP instruments from column 2 of Table 5 are used in place of the engine-based instruments.

Figure 1a: Fuel Economy and the CAFE Standard for Cars, 1975-2007 (MPG)

Figure 1b: Power and Weight of Cars, 1975-2007

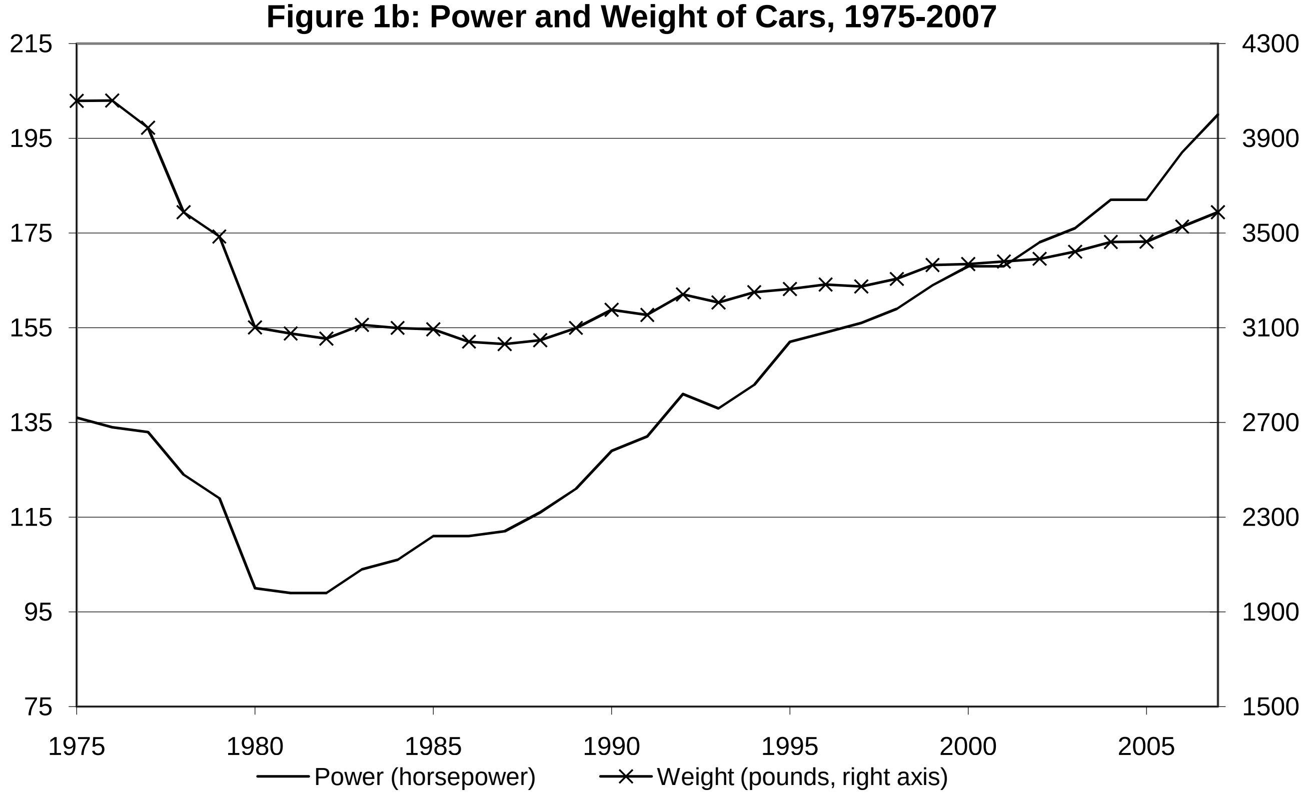

Notes: Figures are constructed using data reported in U.S. EPA (2007).

Figure 2a: Fuel Economy, Weight and Displacement for Cars of U.S. Manufacturers, 1975-2008

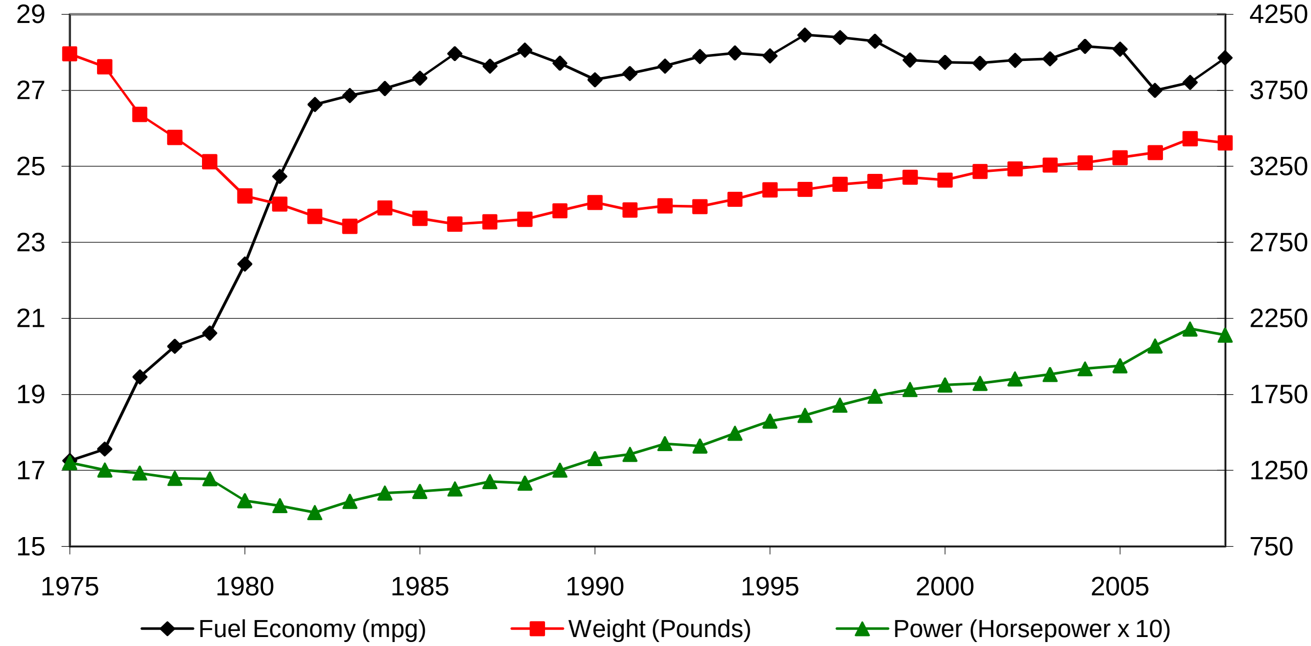

Figure 2b: Fraction Change in Fuel Economy, Weight and Power, 1975-2008

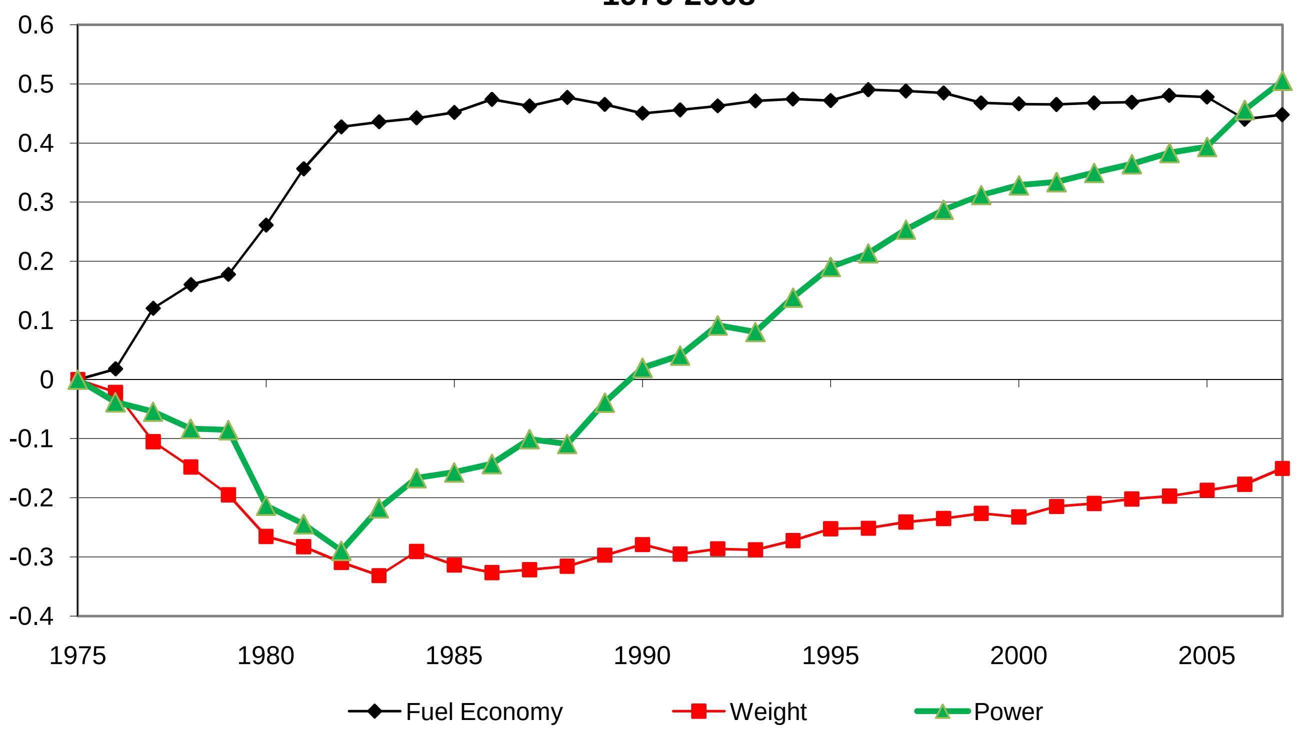

Notes: Figure 2a reports the sales-weighted mean fuel economy (in MPG), weight (in pounds) and horsepower (multiplied by 10) of all cars sold by U.S. companies for each year. Figure 2b reports the percent change in each variable, relative to 1975.

Figure 3: The Effect of Changes in Sales on Average Fuel Economy of U.S. Cars, 1975-1984

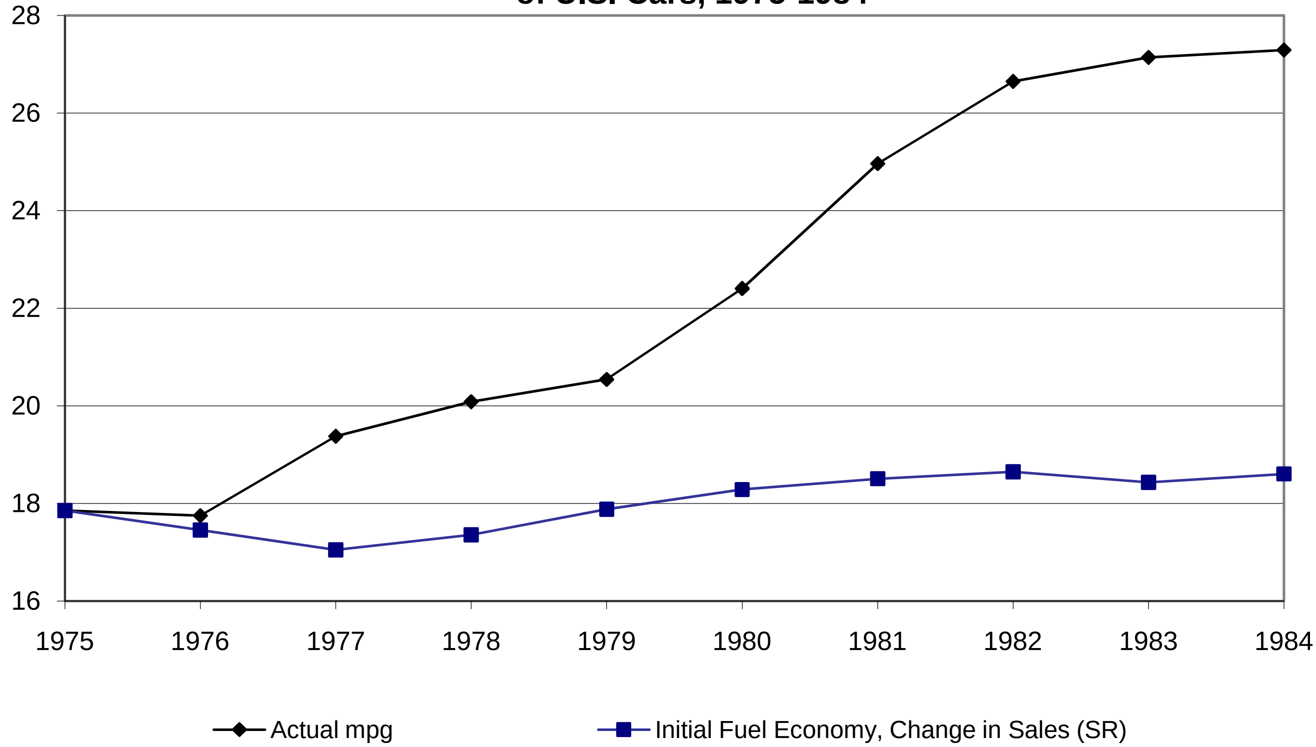

Notes: Actual Fuel Economy is the sales-weighted mean MPG of all cars sold by U.S. firms as in Figure 2a. The change in Initial Fuel Economy is the sum over vehicles of the change in market share multiplied by the fuel economy from the previous year. See text for details. Starting in 1975, the Initial Fuel Economy series plots the cumulative sum of previous changes.

Figure 4: Effect of Power and Weight on Fuel Economy for U.S. Manufacturers, 1975-2008

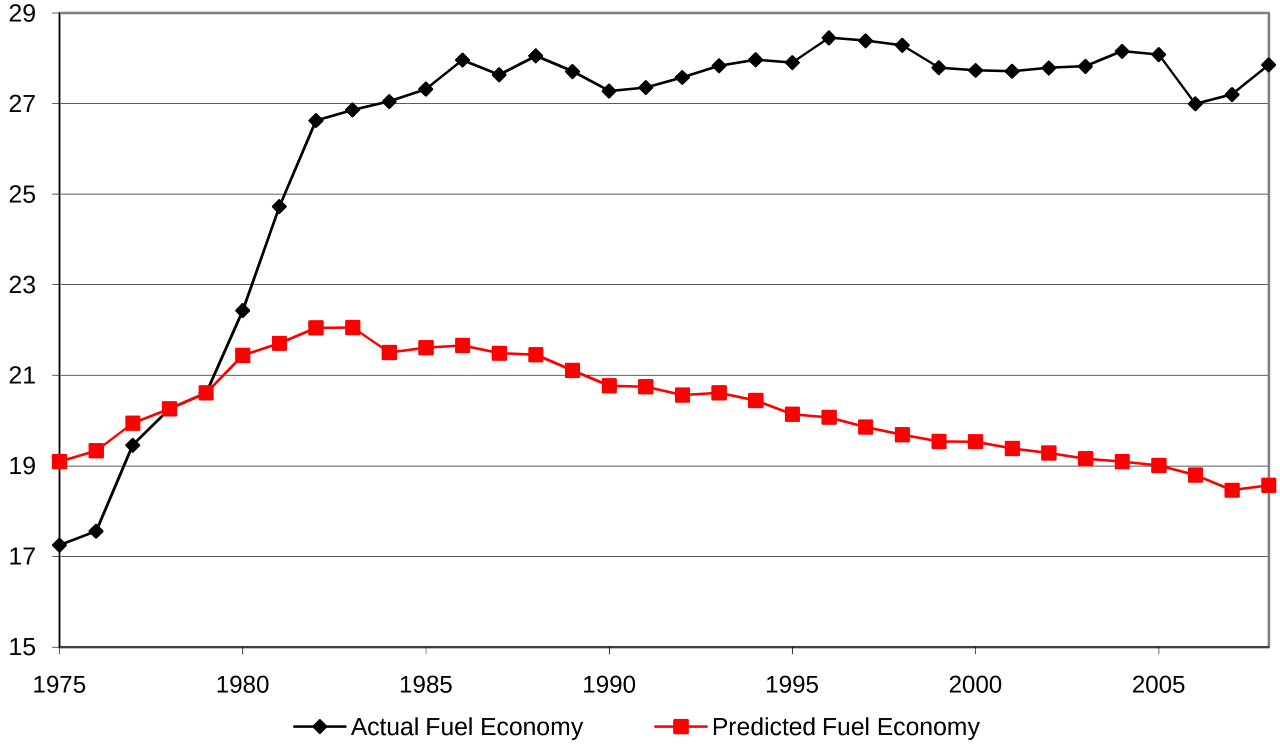

Notes: The Actual Fuel Economy series is the same series as reported in Figure 2. The change in predicted fuel economy is calculated using equation (1), the estimated coefficients reported in column 1 of Table 2. For each model the change in predicted fuel economy is calculated using the model's change in weight and power and the corresponding coefficients. The overall change in predicted fuel economy reported in the figure is the sales-weighted mean of the model-specific chnages in predicted fuel economy.

Figure 5:
Instruments Versus Engine Characteristics

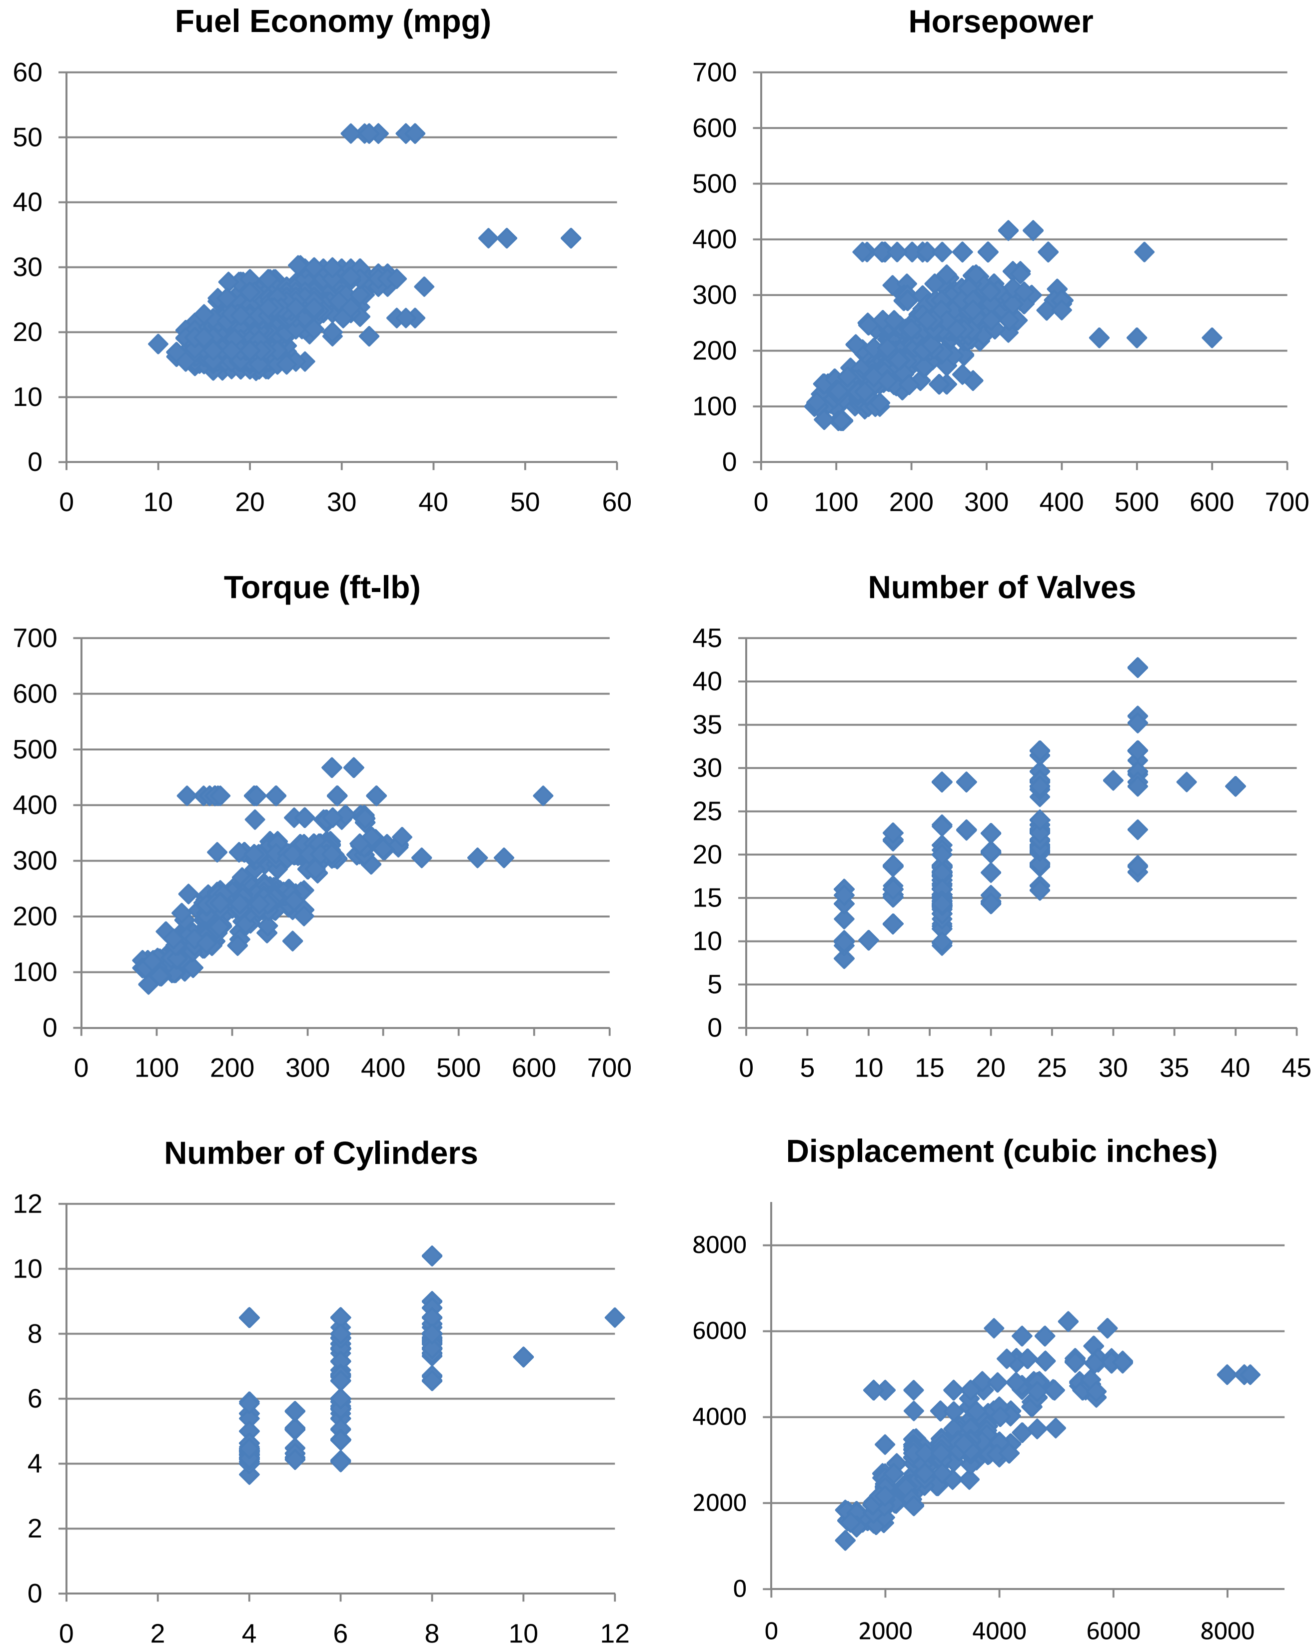

Figure 6: Change in Willingness-to-Pay Due to Changing Vehicle Characteristics for U.S. Firms, 1975-2008

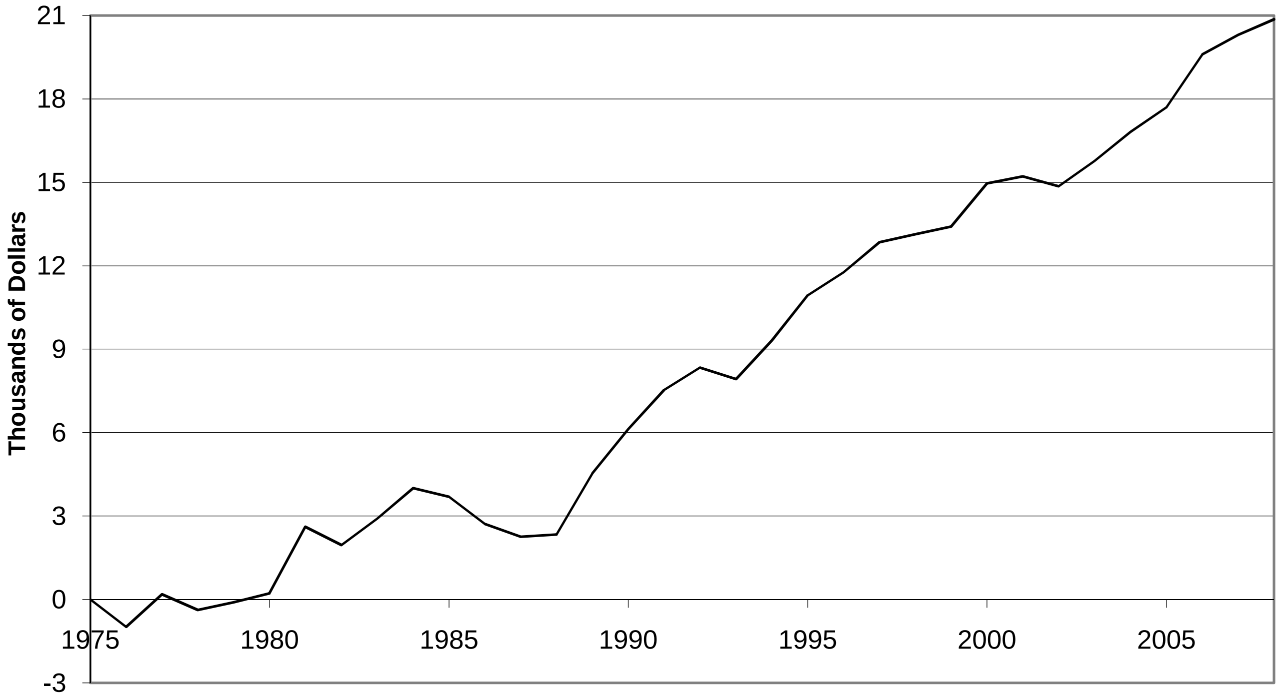

Notes: The figure plots the change in willingess-to-pay for U.S. cars, using 1975 as the baseline year, in 2007 dollars. Change in willingness-to-pay is calculated using the change in sales-weighted power and weight from Figure 2 and the estimates from column 3 of Table 5.