Gasoline Taxes and Fuel Economy: A Preference Heterogeneity Approach*

JOB MARKET PAPER

Mario Samano†

October, 2011

Abstract

This paper estimates an equilibrium model for the U.S. car market to measure the value of two policies aimed at reducing gasoline consumption. The first one is the Corporate Average Fuel Economy standard, and the second one is gasoline taxes. We use a structural model that allows for flexible substitution patterns across car models, measures preferences on cost per mile driven, accounts for the problem of endogeneity of prices, and jointly solves for the manufacturers' optimal responses. The data used include income and miles driven. Counterfactual results show that the welfare loss gross of externality costs from tightening the standard by 10 percent is about two times the cost of increasing net gasoline prices by 10 percent. When accounting for externalities, the two policies may be welfare increasing.

---

*I would like to thank Gautam Gowrisankaran, my main advisor, for all his encouragement, patience and feedback. I would also like to thank Kei Hirano, Derek Lemoine, Stanley Reynolds, Mo Xiao, and participants at University of Arizona seminars for all of their comments and suggestions. Adam Copeland generously shared data for this paper. Kathleen Nosal provided useful comments throughout the project. All errors are my own.

†University of Arizona: Department of Economics. email: msamano@email.arizona.edu, McClelland Hall 401 : 1130 E. Helen St. : P.O. Box 210108 : Tucson, Arizona 85721-0108

1 Introduction

Gasoline consumption is the second largest source of greenhouse gas (GHG) emissions in the U.S. In 2009, 27% of CO_2 emissions came from gas consumption[1]. This paper studies two policies that have been implemented to reduce gasoline consumption. First, in order to increase the efficiency of gasoline use, the U.S. introduced a vehicle efficiency standard. This efficiency measure is the Corporate Average Fuel Economy (CAFE) which is a weighted harmonic mean of the efficiency of vehicles− in miles per gallon− from a manufacturer. If a car maker's CAFE is under the standard, a fine proportional to the difference from the standard has to be paid.[2] Second, gasoline taxes have also become an instrument to induce lower gasoline consumption by increasing the cost per mile traveled[3]. To evaluate which one of the two policies is more effective we need to examine the social and environmental costs associated to each one in terms of the amount of gasoline consumption that can be avoided. This paper develops an empirical method to evaluate such costs.

In 2011, two main changes to the CAFE regulation are going into effect.[4] This is the first time that there has been a significant change in the policy since 1990. Both changes aim to increase the overall fuel efficiency to 35 mpg in the next five to nine years. This paper, although not directly modeling the new restrictions, sheds some light on what the effects are when the overall standard is mandated to increase in such a short time.

Evaluating the impact of CAFE standards and gasoline taxes requires knowledge on the substitution patterns across vehicle purchases and the elasticities with respect to the cost per mile. The CAFE policy mainly works on the extensive margin of purchases: the choice of the car model. This reallocates consumption to high fuel-economy cars. Therefore, the policy creates different incentives across manufacturers based on their fleet mix. CAFE acts on the intensive margin −the amount of miles driven− through the rebound effect: if consumers buy

[1] Just slightly below the 33% from power plants. Environmental Protection Agency (EPA), http://www.epa.gov/climatechange/emissions/usgginventory.html

[2] The value of the CAFE standard until 2010 was 27.5 mpg for car passengers, and it is to be increased to 35 mpg in 2016.

[3] Defined as price of gas divided by the rated miles per gallon of a vehicle.

[4] The first is an "attribute"-based policy where the standard is a decreasing function of the footprint of the vehicle (defined as the surface covered by the rectangle between the four wheels) embedded in the Energy Independence and Security Act (EISA) of 2007. The second is due to the National Fuel Efficiency Policy of 2009. Under this program, GHG emissions and fuel efficiency are jointly regulated. Its goal is to attain the new vehicle fuel efficiency average of 35 mpg by 2016, four years earlier than the goal in EISA.

a more fuel-efficient car, they might tend to drive it more since it is now cheaper to do so. On the other hand, gasoline taxes increase the cost per mile driven in the intensive margin. This in turn increases the value of those car models with high fuel-economy, an extensive margin effect. Taxes therefore act on both the intensive and the extensive margins, and on all the vehicles in the current fleet, both new and already existing in the fleet. The interaction of these forces is complex and may lead to deadweight loss relative to taxes.

The method in this paper consists in calculating what the prices of equilibrium would be under changes in each policy. We use a structural model that allows for flexible substitution patterns across car models, measures preferences on cost per mile driven, and jointly solves for the manufacturers' optimal responses. In the short run, adjustments to changes in the CAFE standard are made through car prices, but technology that can influence the cost per mile remains relatively fixed since it is difficult for manufacturers to change the characteristics from one year to the next. This change in prices would affect consumers decisions, potentially lowering consumer surplus since consumers might have to buy a different car had there been no change in car prices. When there is an increase in gas prices, the cost per mile increases by the same factor for all vehicles. However, those consumers who prefer vehicles with low cost per mile of driving might buy a different car model had there been no change in gas prices. This would change market shares and prices, lower consumer surplus and affect profits. This also causes a rebound effect, which consists of driving more given that it is less costly to drive, therefore diminishing the environmental value of the policy. By comparing the composition of the new cars fleet before and after these policies change, we can compute the change in gasoline consumption and the cost associated with it.

Understanding the relative efficiency of the CAFE regulation depends on consumers' preferences for fuel-efficient cars, the demand for miles traveled, the oligopolistic nature of the car industry, and a model that can capture any substitution pattern for car purchases. We believe the best way to incorporate the CAFE regulation into a general equilibrium model of the car industry is by using a demand model that is not subject to the independence of irrelevant alternatives (IIA) property[5] and that solves the problem of price endogeneity. This is complemented by a supply model that explicitly accounts for the fines incurred if there

[5]Models subject to the IIA property have relative market shares ratios that do not depend on price changes of other products. Even the nested-logit model is subject to the IIA property within the nests.

is a violation of the CAFE restrictions. To accomplish this I build the demand side of the model with a modification of the model for the automobile industry in Berry et al. [1995] (BLP). Then I add a supply model with the CAFE regulation which gives more accuracy to the recovered marginal costs. To our knowledge, this paper is the first study of the regulation that solves both the problem of price endogeneity without the IIA property and models the CAFE regulation in the supply side at the same time.

The model uses aggregated level data on market shares and individual level data on demographics conditional on car model. The absence of the IIA property is important because for example, it is unrealistic to think that the market shares ratio of two compact cars would not be affected by a change in the price of another compact car. Moreover, the elasticities obtained in this paper are sensitive to the cost per mile of all the other car models, appropriately weighted by the degree of similarity in the space of characteristics. Therefore, a counterfactual in gasoline price increases will more realistically reflect these substitution patterns.

The oligopolistic nature of the car industry is modeled assuming that observed prices correspond to the outcome of a Bertrand-Nash equilibrium among car manufacturers. Recovering the marginal costs is challenging since the CAFE standards affect the profit function through the fines imposed on the non-compliers and these fines are a non-linear function of prices. Also, doing comparative statics on the effects of increases in the CAFE standard is difficult, therefore it is not possible to measure its effects analytically. As shown later in the text with a simplified model, having accurate demand elasticities is crucial to the recovery of the true marginal costs, and therefore to the validity of the counterfactuals.

This paper builds on two literatures: empirical models of oligopolies, and structural estimation of discrete choice models. Goldberg [1998] proposed and calculated a model for the car industry that could be used to study the CAFE regulation. In order to make the model more tractable, she replaces the actual definition of the standard by using a linear form rather than the harmonic mean contained in the calculation of the CAFE for each manufacturer. We overcome this issue at the price of more complex computations. From the other branch of the literature, we follow Berry et al. [1995] (BLP) to estimate cross-price elasticities. This is done in a structural way as opposed to other studies that focus on reduced forms[6] in which case little can be said about the CAFE regulation since they do not model supply. Other studies have focused on the mechanics of CAFE without solving the problem of price endogeneity[7], or focusing on the demand side.[8]

There are three main limitations in this paper. First, this is a study of the short run impact of changes in these policies. This is the case because we are not modeling optimal decisions on changes in the car characteristics in the future. Therefore, it would be misleading to extrapolate the results of this paper to the long run. Second, this paper looks only at the market of new vehicles. Thus, the results in this paper constitute only one part of the overall impacts of such policies. However, it contributes to the literature in giving a precise account of the immediate welfare loss after a change in the policy so that it can be compared to the marginal cost of externalities. Third, to analyze the impact of the new choices under the counterfactuals on driving behavior, we take the elasticity of miles driven with respect to gasoline prices from the literature.

The model is applied to the new car market between the years 2000 and 2007, during which gas prices did not exhibit abrupt changes. We obtain data on car characteristics, prices and market shares from J.D. Power and Associates. Income and vehicle miles traveled distributions by manufacturers from the National Household Travel Survey 2009, and gasoline prices from the Energy Information Administration.

I find that increases in gasoline taxes of \$0.50 and \$1.00 lead to a welfare loss gross of externalities of \$1.13 and \$1.67 per gallon of gasoline saved. On the other hand, an increase of 10 percent in the CAFE standard leads to a welfare loss gross of externalities of \$2.11 per gallon of gasoline saved. When the savings from externalities avoided are included, these policies may even be welfare increasing as some authors have suggested (Parry et al. [2007], Portney et al. [2003], Austin and Dinan [2005]).

The remainder of the paper is divided as follows. Section 2 describes the CAFE regulation and reviews the literature. Section 3 describes the model of both demand and supply. Section 4 describes the data and the sources. In Section 5 I explain the details of the estimation

[6]Busse et al. [2009]. In their paper they find that, as expected, the largest effects from an increase in gas prices are on the most- and the least-fuel efficient cars

[7]Bento et al. [2005], Jacobsen [2009] use a Bayesian estimation method that does not explicitly account for price endogeneity. However, Yang et al. [2003] discuss how to include price endogeneity in the same kind of Bayesian estimation.

[8]Gillingham [2010], Busse et al. [2009]

and computation. Section 6 and Section 7 give the results and counterfactuals. Section 8 concludes.

2 CAFE Regulation

As a response to the Arab oil embargo in 1973-1974, it was proposed to the U.S. Congress a change to the Energy Policy Conservation Act to improve the fuel efficiency of vehicles. In 1975, the Corporate Average Fuel Economy (CAFE) standard was approved with a short-term goal of doubling the efficiency in miles per gallon of new cars by 1985. The CAFE is the weighted harmonic mean of miles per gallon of all the cars in a manufacturer's fleet for a given year. The weights come from the size of production of each model year [9]. Compliance occurs when the manufacurer's CAFE is above the standard for that year.

The National Highway Traffic Safety Administration (NHTSA)[10] establishes the CAFE standard to be proposed to Congress. The Environmental Protection Agency (EPA) calculates the CAFE for each manufacturer. The law specifies that the CAFE standards should be set at the "maximum feasible level" in consideration of four factors:[11] i) technological feasibility; ii) economic practicability; iii) effect of other standards on fuel economy, and iv) need of the nation to conserve energy. This paper addresses the interaction between the standard and ii), iii), and iv).

The level of the standard has changed over time. For passenger cars the standard was 18 mpg in 1978. In 1979 it was 19 mpg, in 1980 was 20 mpg and, for 1985 and thereafter the standard was established at 27.5 mpg.[12] In 2011, the standard became a function of the footprint of the vehicle, which is defined as the area between the four wheels of a car. Specifically, cars with a smaller footprint are subject now to higher standards than cars with a larger footprint. This discourages manufacturers from making smaller cars which has become a safety concern in recent years and it also discourages relabeling of vehicles to avoid

[9]For each manufacturer, two different CAFE are calculated, one for passenger cars and the other for light trucks. CAFE standards are established each year for both categories of vehicles. In this paper I abstract from this issue for computational reasons letting light trucks being always complying. This underestimates the impact of the standard on profits.

[10]A division of the Department of Transportation

[11]Energy Policy Conservation Act http://www.law.cornell.edu/uscode/42/ch77.html

[12]For light trucks Congress did not establish any standard at the beginning of the program in 1978. However, one year later, a standard of 17.2mpg was specified for vehicles with a gross vehicle weight rating (GVWR) of 6,000 pounds or less. The light truck standard increased to 22.2mpg in 2007.

stricter standards.

When a manufacturer's CAFE is below the standard, it incurs in a penalty of \$5.50 per tenth of a mile per gallon under the target value times the total volume of those vehicles in a given year. That is, \$55.00 are charged for each mpg under the standard multiplied by the total production volume of the manufacturer[13]. Since 1983, manufacturers have paid more than \$500 million in fines.[14] The non-compliers are typically European manufacturers of high-end cars.

The car industry represents 4.5% of total GDP, there were 8.9 million new cars in the U.S. in 2009 and 11.2 million in 2008, which together represent about 8% of total number of passenger cars for the U.S. Although the total number of cars has been mostly constant in the last ten years, the proportion of new cars in each of those years has been around 6%[15]. This is important because the CAFE standard policy, although it only affects new vehicles, is increasing the car fleet's efficiency at a rate of about 6% per year. It is expected that most of those cars produced in the last ten years are still functioning and therefore that would account for slightly more than half of the fleet in the best scenario.

There have been a number of previous studies on fuel efficiency, gas taxation, and CAFE standards (Goldberg [1998], West [2004], Kleit [2004], Austin and Dinan [2005], Jacobsen [2009], Busse et al. [2009]). The main difference of this paper with respect to others is the treatment of the interaction of cross-price elasticities and their relationship with the manufacturers' optimal decisions. Goldberg [1998] evaluates CAFE and computes the fuel consumption savings and profit loss in 1989. Then she calculates the tax increase in gas to get the same reduction in gas consumption. She finds that the tax increase would need to be several times higher at a point where, as she points out, it would be politically infeasible to implement it. Thus CAFE seems to be a more acceptable policy than taxation, from an average consumer perspective. Goldberg independently estimates the demand for gas and the profits of firms in a static context.[16] Studies that involve welfare effects include Kleit [2004],

[13]If the manufacturer is complying, it can earn CAFE credits to potentially offset future penalties. One improvement over this part of CAFE could be to create a market price for these credits in a cap-and-trade scheme as suggested in Portney et al. [2003]. In this paper we abstract from the carry-on of credits since our framework is static.

[14]http://www.nhtsa.gov/fuel-economy

[15]http://www.eia.gov/oog/info/gdu/gasdiesel.asp

[16]West (2004) uses a static nested logit model to find the effect of income and household characteristics on vehicle choice. She finds that tax on gas is regressive only across upper income groups.

Austin and Dinan [2005], and Jacobsen [2009]. Kleit [2004] considers the welfare cost of a 3 mpg increase in CAFE and its associated decrease in gas consumption. He finds that in order to get the same amount of fuel saved, an increase of 11 cents in the price of gas is needed, but the welfare cost in this case is 20 times less than the cost associated to the increase in standard. For his calculations he used the price elasticities supplied to him by General Motors and assumed the market is competitive which is a rather unrealistic assumption for this industry.

Austin and Dinan [2005] solve a similar exercise considering the effect of an increase in the CAFE standard of 3.8 mpg versus an increase in gas tax of 30 cents (in order to have the same gas reduction of 10%) over a 14-year period. They calculated present discounted values for welfare for these two scenarios and found that the cost of the second policy is only 71% of the cost of the first policy. Their model is static and the unobservable marginal costs are estimated from observed retail markups, however, cross-price elasticities are unspecified. Jacobsen [2009] studies increases in the CAFE standard that are gas reduction equivalent to gas tax increases. He finds that increasing the CAFE standard by 1 mpg produces a decrease of 3% in demand for gas. The same decrease in gas consumption can be achieved by increasing the tax on gas which causes a welfare cost of only one sixth of the welfare cost of increasing the CAFE standard. In Jacobsen's model, the estimates are obtained by calibrating the consumer and producer optimization problems so the two problems are solved independently from each other at each step. There are two potential problems with this approach. First, price endogeneity is not being taken into consideration, which biases the results. Second, the predicted demand of miles traveled would require longitudinal data to identify the elasticity of miles traveled with respect to gas prices; however the National Household Travel Survey does not contain enough information.[17]

3 A model for the car industry and CAFE

In this section I describe the utility function for consumers, how the regulation is incorporated into the supply model, and the distributional assumptions.

[17]This could be overcome with a data set that follows over time households on their consumption of vehicle miles traveled and the price of gasoline they paid for it.

3.1 Demand

I assume each individual has a utility function that depends on the choice of the car including the cost per mile associated at the time of purchasing. The decision making process consists of two simultaneous decisions: (i) a discrete choice of the car model[18], and (ii) a continuous choice of the amount of gas to be consumed through the price of miles traveled. The first decision can be modeled in the standard way using a random utility model. The second choice can be incorporated through the variable *dpm*[19] defined as price of gas divided by the miles per gallon of the car model. Appendix A describes in detail a model of these two simultaneous decisions and how the utility function described below can be thought of as a reduced form of that joint model. If we assume that the average car is driven on average 12,000 miles per year,[20] we get the distribution of market shares on prices and *dpm* of Figure 3 and Figure 4. On the left of those graphs we see the hybrids and other low-cost car models. In 2000, their market shares are very small. By 2007, the market shares of low-cost vehicles have considerably increased even if the price of those cars has not gone down. Simple correlations of market shares and price and *dpm* like the ones in Table 3 show that both car price and *dpm*, are negatively correlated with market shares. Therefore, it is crucial that the utility function capture the effects of *both* prices, as suggested by the model in Appendix A. Therefore, we assume consumers have an indirect utility function with the form

$$u_{ijt} = \underbrace{\alpha \log(dpm_{jt}) + x_{jt}\beta + \xi_{jt}}_{\delta_{jt}} + \underbrace{\gamma \log(y_i - p_{jt}) + \sigma^{dpm}\nu_i^{dpm} \log(dpm_{jt}) + \sum_{k=1}^{K} \sigma^k \nu_i^k x_{jt}^k}_{\mu_{ijt}} + \epsilon_{ijt}. \tag{1}$$

There are J_t car models and one outside good at year t; the latter represents the possibility that an individual chooses not to buy a car and use some other type of transportation. x_{jt} is a vector of car characteristics for model j, a constant term, a time trend, and manufacturer dummies. The time trend captures the degree to which utility from the same car model changes over time[21]. ξ_{jt} is a product time-specific demand shock which is unobserved by the

[18] I assume only one car can be owned by each household.
[19] Dollars per mile.
[20] Bureau of Transportation Statistics http://www.bts.gov/publications/national_transportation_statistics
[21] Most car models are sold over periods of two years or more. Their characteristics are exactly the same over time; however they are cheaper than in their first year they were on the market.

econometrician (Berry [1994]). This demand shock is often interpreted as a quality index that cannot be measured. The cost per mile *dpm* is price of gas divided by miles per gallon. The price of the car is p. The vector $(y_i, v_i, \epsilon_{ijt})$ contains individual-specific variable: y_i is household income, assumed independent of ν_i and ϵ_i[22]. ν_i is a $(K+1)$-vector of standard normal variables, independent and identically distributed (IID) across households, car models, and time. The last term, ϵ_{ij}, is a random draw from a type I extreme value distribution, IID across markets (years), across individuals, across car models, and independent of y_i and ν_i. The coefficients σ represent the degree of heterogeneity of the variable they are associated with. For example, the term $\sigma^{dpm}\nu_i^{dpm}$ gives the household variation around the mean taste for cost per mile α. Note that the utility function is separated into a mean utility component δ_{jt}, and an individual-specific deviation $\mu_{ijt} + \epsilon_{ijt}$. For computational reasons I restrict some of the σ^k to equal zero. I only allow for heterogeneity in the coefficients on price, cost per mile, acceleration[23], length, height, time trend, and the constant. The coefficients are assumed independent of income and other variables, and independent of each other within an individual. The utility from the outside good is defined as

$$u_{i0t} = \epsilon_{i0t}.$$

Because of the distributional assumptions made above, a standard result by McFadden [1981] is that predicted market shares are given by

$$\hat{s}_{jt}(x, p, dpm; \theta) = \int\int \frac{\exp(\delta_{jt} + \mu_{ijt}(x_j, p_{jt}, dpm_{jt}, \nu_i, y_i; \theta))}{1 + \sum_l \exp(\delta_{lt} + \mu_{ilt}(x_j, p_{jt}, dpm_{jt}, \nu_i, y_i; \theta))} dP_\nu(\nu_i) dP_y(y_i) \qquad (2)$$

where $\theta = (\gamma, \sigma)$, $P_v(\cdot)$ and $P_y(\cdot)$ are the distributions of taste shifters ν_i and income respectively.

As suggested by the model derived in Appendix A, the specification to be estimated is a reduced form of a model of the simultaneous decisions of miles driven and car choice. This study abstracts from any possibility of a highly sophisticated buyer who has expectations about the future behavior of gas prices and therefore incorporates this into his car purchase decision. Research on this issue shows mixed results on whether consumers fully account for future operating-costs savings but there is a mainstream opinion that they do not (Anderson

[22] A variant of this could be to have income correlated with ν_i by conditioning the distribution of ν_i on observed income.

[23] Measured as horsepower over weight.

et al. [2011], Allcott and Wozny [2010]). Anderson et al. [2011] found that most consumers construct beliefs of gas prices based on current values suggesting the validity of using current prices in research of durables associated with energy consumption. Thus, we assume that the consumer perceives the current value of *dpm* as the best estimate of future *dpm* values. Because data on purchases are highly aggregated, we assume the same gas price across the country.[24]

3.2 Supply and CAFE

As noted before, not complying with the regulation leads to fines, which negatively affect profits. Since the fines depend on the distribution of the manufacturer's market shares, they also depend on prices. The simplest case is the one with manufacturers that have never paid fines or have used the carry-on and credit system[25]. The profits for manufacturer f are

$$\Pi_f = \sum_r (p_r - mc_r) q_r(p, \theta) \tag{3}$$

where r is the index on the cars produced by the same manufacturer and q_r is the number of units produced of car model r. Producers whose CAFE is under the standard and pay the fines have a profit function of the form

$$\Pi_f = \sum_r (p_r - mc_r) q_r(p, \theta) - F(q) \tag{4}$$

where the last term represents the fines paid to the regulator. The CAFE regulation defines the way to compute the CAFE for a manufacturer. This number depends on the relative contribution of each of the manufacturer's car models to its aggregated level of fuel efficiency. The CAFE for each manufacturer is defined by the regulation as

$$c_f = \frac{\sum_{j \in f} q_j}{\sum_{j \in f} \frac{q_j}{mpg_j}}$$

where q_j is the number of cars of model j and mpg_j is the model's efficiency measure (miles per gallon). A few observations might be worth making. The harmonic mean is more appropriate here than the arithmetic mean because it is equivalent to the mpg needed to travel the same

[24] The author is currently working on alternative specifications to further analyze this issue.

[25] Domestic manufacturers have never paid fines even if some years they have not complied. Throughout I will assume that not having paid fines is equivalent to be a complier.

distance using the average of gallons needed when traveling the same distance in cars with different fuel efficiencies.[26] The CAFE is increasing in mpg, as expected. If the manufacturer produces the same number of units for each car model then c_f becomes just the harmonic mean of the mpg of the different models. This is why the CAFE depends on the market shares, so that cars with high fuel efficiency and large market shares are more represented in the CAFE.

The penalty fee $F(q)$ is a linear function of how far the manufacturer's CAFE is from the standard. For each 1 mpg away from the standard the fine is \$55 times the total number of units produced that year,

$$F(q) = \$55 \times \left(\sum_{j \in f} q_j(p, \theta) \right) \times (c_s - c_f(p, \theta))$$

where c_s is the value of the standard.

Each year, firms choose prices to maximize profits. I assume car makers compete in prices and we observe the outcome of a Bertrand-Nash equilibrium. First order conditions for the maximization problems for each firm characterize the equilibrium and define expressions for marginal costs by taking observed prices as the ones from equilibrium. The first order conditions also impose the structure of an oligopoly because a single manufacturer chooses prices for all the cars it produces. The first order conditions for compliers and for non-compliers in vector form are, respectively[27],

$$p_f - mc_f + {\Delta_f}^{-1} s_f(p) = 0 \tag{5}$$

and

$$p_f - mc_f + \Delta_f^{-1} \left(s_f + 55 \left(\sum_{r \in f} s_r \right) \frac{\partial c_f(p)}{\partial p} - 55(c_s - c_f(p, \theta)) \Delta_f 1_{(J,1)} \right) = 0 \tag{6}$$

where Δ_f is a matrix of cross price market shares derivatives, and $1_{(J,1)}$ is a column vector of ones.

To have an idea of the behavior of these conditions, consider the case of a non-complier

[26] For example, two cars with fuel efficiencies 20 mpg and 60 mpg. The average of gallons needed to travel 120 miles is 4. A vehicle with a fuel efficiency equal to the harmonic mean of the two can travel the same distance using 4 gallons of gas.

[27] They are derived in the appendix.

single-product manufacturer. In this case, the first order condition is

$$p - mc = 55\,(c_s - mpg) - \frac{s(p)}{\frac{\partial s(p)}{\partial p}}.$$

This shows that the mark-up is smaller for a higher fuel efficiency of that single car model, without accounting for the potential offsets from preferences, but it is larger than the mark-up of a complier that has the same market share and same own-price elasticity. So higher fuel economy has a negative effect on mark-up, but net effect is unknown if we do not know the effect on preferences. By comparative statics, we can see the effect of an increase in the standard on the optimal price. By assuming that the demand is convex, a complier with a large enough mark-up will decrease its price. A complier with relatively small mark-up will increase its price. These results are relative to the mark-up that a non-complier with same market share and same own-price elasticity would need to be for the same results to hold. That is, if it is optimal for the non-complier to increase the price, then the complier would need to have a larger mark-up for lowering its price to be optimal. In this case, the non-complier would attract more buyers but the complier needs to have a larger cushion in mark-up to be able to compete against that. Such neat results cannot be derived about a producer of two different car models. Under reasonable assumptions about elasticities, complying lowers mark-ups for both car models, and the lower the own-price elasticity is as compared to the cross-price elasticity the higher the mark-up will be.[28] These observations highlight the importance of knowing very precisely the own- and cross- price elasticities in the market. Without accurate elasticities, like the ones from models where the IIA assumption holds, the marginal costs recovered from the first order conditions can be largely biased.

4 Data

Table 2 summarizes the data on car characteristics, gas prices and miles driven. Data on car characteristics come from J.D. Power and Associates (JDPA) and from Ward's Communications[29]. The characteristics include length, height, width, horsepower/weight and mileage per gallon for each car model. The quantity horsepower/weight is a measure of acceleration.

[28]For this result I assumed the two models have the same market share.

[29]I thank Adam Copeland for generously sharing this data set for this paper. In Copeland et al. [2011] there is also a detailed description of this dataset.

This data set covers the period of 2000 through 2007. It includes weighted transaction prices for each car model in the year where the model appeared for the first time and the subsequent years when that model car was still being produced. This means the exact same car model is observed over more than one year with same characteristics but different price and different cost per mile, (see Figure 1). All prices are converted into 2007 dollars using the Consumer Price Index published by the Bureau of Labor Statistics[30].

Since the JDPA data do not have individual level information on characteristics of buyers, I matched this data with the National Household Travel Survey 2009 (U.S. Department of Transportation, Federal Highway Administration [2009]) to get income distributions by manufacturer. I also use these data to calculate in the counterfactuals the effect on miles traveled. These data are collected by a phone survey to households in the United States. For each household, the survey has information on car model, number of cars, miles traveled per year and demographics. Table 1 summarizes the empirical distributions of income and annual miles traveled.

Gasoline prices were obtained from the EIA[31] and converted into 2007 dollars using the Consumer Price Index. CAFEs for different manufacturers come from the Environmental Protection Agency (EPA) and the list of companies that have paid fines comes from the National Highway Traffic Safety Administration (NHTSA).[32] Figure 2 shows the CAFEs for some selected manufacturers. Asian companies have had a consistent trend in high fuel efficiency, domestic manufacturers oscillate around the standard, and European manufacturers are slightly below the standard.

Figure 3 and Figure 4 show market shares in 2000 and 2007 as distributed by price and annual average cost of driving.[33] The main difference is the big change in market shares of low-cost per mile vehicles. Another important change is that in 2000 the most expensive cars have costs per mile around the average. In 2007 this is not the case, the most expensive cars are also the ones with highest costs per mile. This means that during this period of time some consumers internalized the increase in gas prices by buying cheaper cars. Moreover,

[30] http://www.bls.gov/cpi/#tables

[31] http://www.eia.gov/petroleum/data.cfm#prices

[32] The NHTSA website http://www.nhtsa.dot.gov

[33] On average, a car is driven 12,000 miles per year in the U.S. according to EPA's MOBILE6 model http://www.epa.gov/oms/m6.htm and to the DOT's Federal Highway Administration http://www.fhwa.dot.gov/policyinformation/statistics/2009/

simple correlations in the data show that higher income households tend to travel less despite their high per mile cost. This is one factor that explains why high driving-cost vehicles can be high-priced. To further explore the data, Table 3 shows that under different specifications, *both* price and cost per mile are negatively correlated with market shares. Of course, this is not suggesting any causality, (the results from the full model confirm this relationship though). As suggested by Knittel [2010] and Gramlich [2009], there is a trade-off between fuel efficiency (mpg) and other car characteristics, like size and acceleration. Fuel efficiency only enters the utility function through the cost per mile.

5 Estimation and Computation

The calculations in this paper consist of several steps that I explain in this section. Some details are left for Appendix B. The goal of these computations is to characterize the utility function and the marginal costs so that we can find the prices in the new equilibrium when gasoline prices or the standard increase.

If we knew $\hat{\delta}_{jt}$ then we could recover parameters (α, β) by OLS. The rest of the terms in the utility function can be put together into $\mu_{ijt} = \gamma \log(y_i - p_{jt}) + \sigma^{dpm}\nu_i^{dpm}\log(dpm_{jt}) + \sum_{k=1}^{K}\sigma^k\nu_i^k x_{jt}^k$. This term captures individual deviations from the mean valuation. Thus, we can write

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt}.$$

We do not have individual level data on car purchases, but only market shares. A utility maximizer consumer i will choose car model j if and only if $u_{ijt} > u_{iht}$ for all h. The distribution on ϵ_{ijt} implies that

$$\Pr(u_{ijt} > u_{iht}) = \int\int \frac{\exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_l \exp(\delta_{lt} + \mu_{ilt})} dP(\nu) dP_y(y_i)$$

where the integration is over the distribution of the ν_i, which are assumed normally distributed and independent from characteristics and from each other, and over the distribution of income by manufacturer (empirical distribution from the data). We can interpret market shares as

$$s_{jt} = \Pr(u_{ijt} > u_{iht} \text{ for all } h).$$

therefore, individual data on car purchases are not needed. Estimation will "match" model predictions with observed market shares by choosing appropriate parameters.

The estimation relies on the minimization of a set of moments by a standard general method of moments (GMM). This procedure has been widely used in empirical industrial organization problems (Berry et al. [1995], Petrin [2002], Nevo [2001]). The predicted probabilities $\hat{s}_{jt}(\delta;\theta) = \Pr(u_{ijt} > u_{iht}$ for all $h)$ are taken to the data by solving the system of equations

$$\hat{s}_{jt}(\delta;\theta) = s_{jt}$$

where s_{jt} are the actual market shares and δ is the column vector obtained by stacking all the mean valuations δ_{jt}. BLP show that it is possible to solve for the vector δ using a contraction mapping. This inversion will give the vector of mean valuations as a function of parameters, $\delta = \delta(\theta)$. By the definition of δ_{jt} in (1), we have

$$\xi_{jt} = \delta_{jt} - (\alpha \log(dpm_{jt}) + x_{jt}\beta). \tag{7}$$

This means that there is a closed form expression for the demand unobservables ξ_{jt}. These unobservables might contain covariates that are correlated with price; therefore there is an endogeneity problem. As shown in previous literature (Berry et al. [1995], Berry [1994]), this problem can be solved by using the moment condition $E(\xi|Z) = 0$ where Z is a set of instruments that are functions of characteristics of other products of the same manufacturer and of characteristics of products of competitors. These instruments are the same used in BLP. They are aggregations of other manufacturer's products and this is what solves the problem of price endogeneity. Specifically, the instruments are differences between overall annual average characteristics and the annual average of the manufacturer's characteristics, differences between model characteristics and annual averages of the manufacturer competitors' characteristics, and the characteristics themselves.

A second set of moments is given by the first order conditions from the supply side of the model. This is done by solving for the marginal costs in (5) and (6) and assuming they are a function of car characteristics, quantity produced, a time trend, and other controls (W_{jt}). Thus, the form of this moment condition is assumed to be $\omega_{jt} = \log(mc_{jt}) - W_{jt}\gamma$. Once it is determined from the data which manufacturers were compliers and which ones non-compliers, it is asummed that that was their optimal response. This means that during the GMM estimation, if in certain year a manufacturer was a non-complier, then its FOC will always be during the estimation equation (6). By the definition of c_f,

$$c_f = M \times \frac{\sum_{j \in f} s_j(p, \theta)}{\sum_{j \in f} \frac{s_j(p,\theta)}{mpg_j}}$$

where M is the market size, we can see that its value depends on the parameters through the market shares (where we have ommited the dependency on δ). But the contraction mapping in BLP will always give a δ such that the system of equations $\hat{s}(\delta, p, \theta) = s$ is always satisfied for *any* vector of parameters θ. Therefore, the value of c_f keeps unchanged during the estimation as well as the manufacturer's nature of non-compliance. There is in fact no need to compute c_f during the estimation, but the value of its derivative. So for any vector of parameters, each of the moment conditions is continuous since there is no switching from one group to the other. Observe also that by construction, car prices observed in the data satisfy the FOCs by calculating marginal costs using equations (5) and (6).

The variation of prices for the same car model over time helps to identify some of the parameters. From equation (2) for the same car model-year over time, since there is variation in actual market shares and cost per mile, the mean utility values δ solve for different values. When those values are taken to equation (7) the vector of car characteristics is the same since it is the same car model-year at different points in time, but again there is variation in the cost per mile. Thus, this case helps to identify α. For the case of two different car models in the same year, (2) shows that there is still variation in μ_{ijt} through the cost per mile, car price, and income, leading to potentially different mean utility values. When taking those values to equation (7), there is variation in the characteristics, this helps to identify the vector β. Supply parameters are more difficult to tract. However, the author has initialized the optimization at different values of the parameters without finding any significant change in the optimal solution.

6 Results

The estimates of the structural parameters are reported in Table 4 and Table 5. For the demand parameters, I report the mean and the standard deviation of the variables with random coefficients, each of them with their respective standard error. The quantities that were not considered to have a random coefficient − or variables with individual variation like disposable income − only have an estimated mean reported. The coefficient on disposable income is positive and significant, which reflects the presence of income effects on car purchases. The coefficient of the mean on the price per mile traveled is significant and negative. This gives evidence for the consumers' sensitivity to changes in the cost of gasoline through the cost of mile driven[34]. The standard deviation for this random coefficient is the largest among all the car characteristics, reflecting the variety in tastes for this attribute: some consumers have a strong preference for low fuel economy vehicles while others may even have a positive value on this coefficient. The mean coefficients on the other car characteristics are all positive except for acceleration. In general there is a trade-off between accelaration and other characteristics because of technological constraints. Therefore, buying a car with relatively high value on this characteristic is a commitment on low levels for other attributes. Because this quantity was assumed to have a random coefficient, I am allowing for many different types of preferences. Since the standard deviation is larger than the absolute value of the mean coefficient, there are some consumers who have a positive preference on accelaration as expected. The negative coefficient on the time trend reflects that most consumers perceive a decrease in the value of the same car model over time, controlled for the change in price. The standard deviation for the trend is large; therefore there is a lot of variety on the preferences for the same model over time. Usually the same model is in the new cars market for three or four years, and the price drops considerably, as the market share does too. This time trend quantity captures the decline in how popular the car model is, which attempts to explain the decline in market share despite the decrease in its price.

There were also included manufacturer dummies in both the supply and demand specifications. These are constants over time and over all the car models from the same manufacturer. The manufacturer VW is omitted in the estimation to avoid multicollinearity. The coefficients on these constants can also be thought of as indicators of overall quality for the manufacturer that could not be controlled for by the other covariates. BMW, Porsche, and Chrysler have the largest values. A bit surprising is the small coefficient on MB, which is generally thought

[34]In Appendix B we show some robustness checks on this coefficient.

as a producer of high-end vehicles, but it is not significant.

The coefficients on cost parameters are positive on all the characteristics except for length. This suggests that smaller vehicles are more expensive to manufacture, which is directly related to the trade-off between fuel economy and other characteristics. The negative coefficient on quantity produced is evidence of economies of scale in the car industry. There is also a negative coefficient on the time trend covariate since it is expected that as technology moves forward, it is less costly to produce the good. As a way to summarize other estimates of the model it is reported in Table 7 the mean own-price elasticity for cars, and own-price and own-*dpm* elasticities by quartile of fuel efficiency for model year 2007. We divided the sample of cars in quartiles according to their fuel efficiency. The least fuel-efficient cars have a more elastic demand than the most fuel efficient cars. These might be due for example to the fact that some of the least fuel efficient cars are luxuries, and that the most fuel efficient cars may be considered as a necessity in order to make savings in gas consumption. More difficult to capture is the elasticity with respect to the price per mile traveled at the moment of purchase. In this case, there does not seem to be a lot of variation according to the level of fuel efficiency. This is because our model does not take into account the variation in individual miles traveled. However, it is important to see that there is an economically significant level of response to this price even from the moment of purchase. The mean Lerner index is 79%. Other studies on the industry have found average Lerner indexes that go from 10% up to 41%; however none of them have used the whole specification of the CAFE regulation like in this paper. The difference might be attributed to higher market power exercised by those manufacturers that are more fuel efficient.

7 Counterfactuals

One of the main advantages of a structural model is that we can use the estimates to analyze counterfactuals. In the first one we increase the gasoline taxes. In the second one we increase the CAFE standard to 10 percent. For both cases we calculate the welfare loss gross of externalities per gallon of gasoline saved, the reduction in miles traveled using individual data from the NHTS, and we quantify the losses in terms of avoided gasoline consumption.

7.1 Gasoline tax

For each of the increases in gasoline taxes, we recompute the cost per mile driven for each of the car models in 2007. Since the coefficient on this price is negative, the mean valuation of each car model will decrease but less so for those ones that have high fuel economy. Thus the utility levels change, predicted market shares change and this makes the first order conditions fail at the original prices. Because in the short-run manufacturers will not be able to change the technology, marginal costs of production will remain the same under the counterfactual. Finally, we solved for the new equilibrium prices in the system of equations of first order conditions for all car models in 2007. The vector of new equilibrium prices is used to compute new market shares.

We matched old and new market shares to annual miles traveled as reported in the NHTS by car model. In order to calculate the change in gallons of gas used after the tax increase, We calculate what the gas consumption would be in the counterfactual and the gas consumption in the initial situation. The general form of this calculation is

$$\Delta \text{gallons} = \sum_{i,j} M \frac{1}{I_j \text{mpg}_j} \left(\text{miles}_i^c \times s_j^c - \text{miles}_i \times s_j \right), \tag{8}$$

where M is the size of the market, I_j is an appropriate factor to average the miles driven when there is more than one observation in the NHTS for the same car model, s_j^c are the market shares obtained by using the new equilibrium prices, miles_i^c is a hypothetical case in which there is a consumer response in miles traveled to the increase in gasoline prices (this term has three different specification that we explain below), and miles_i is the observed number of miles traveled in the NHTS.

In order to calculate the loss in consumer surplus, we use the concept of compensating variation. This means that we calculate the income amount that it would be needed to keep the individual at the same utility level had prices changed.[35] The calculation follows Small and Rosen [1981] and for this paper it takes the form

$$CV = \int -\frac{1}{\lambda_i} \left(\log \sum_j \exp u_{ij}^c - \log \sum_j \exp u_{ij} \right) dF(\nu_i), \tag{9}$$

[35] Jacobsen [2009] uses equivalent variation instead.

where λ_i is the marginal utility of income, u^c_{ij} is the utility level under the counterfactual, u_{ij} is the utility level in the initial situation, and ν_i is a vector that contains the random shocks to the preference parameters. Finally, producer surplus is calculated by taking the difference in profits using the new and old market shares in (3) and (4).

Results for two of the tax increases are presented in Table 8. The additional gasoline taxes were \$0.50 and \$1.00. Annual change in gas consumption is calculated under three cases: a perfectly inelastic demand for miles traveled, a demand for miles traveled with elasticity of -0.40[36], and using the individual heterogeneity in miles traveled from the NHTS together with elastic demand for miles traveled. Using the gas savings from the very last specification, the welfare loss per gallon of gasoline saved gross of externalities is relatively small under both gas taxes, \$1.13 and \$1.67 respectively, and the loss is increasing with the size of the tax as shown in Table 10. The annual percentage changes in gas consumption go from -0.639% in the inelastic case and \$0.50 of additional tax up to -0.967% using heterogeneity in driving habits and \$1 of additional tax. These numbers may seem small but the demand for gas in the U.S. is very high. There is a consumption of 47 million gallons/day of gas in the U.S. according to the EIA. This is equivalent to 17 billion gallons per year. There are 137 million passenger cars total, but only 11 million new cars each year. So we are talking about 8% of the total fleet. This means that these calculations on savings are only for new cars without taking into consideration the savings that a tax on gasoline would induce in households that owned a car already. The market share of the outside option increased by 2.12% and 2.13% for each of the two scenarios. Although small, these changes show the effect of those who are indecisive about buying a new car or not, and that they opt not to buy it if net gas prices increase.

The optimality of the tax is beyond the scope of this paper.[37] The results suggest at least that consumers are sensitive to this tax and that there is a welfare loss from an increase in it, but it is small. The relationship between the tax and the externalities from gas consumption is discussed below.

[36] As in Parry and Small [2005]

[37] In Parry and Small [2005] it is claimed that gas taxes in the U.S. are below their optimal value.

7.2 Increase in CAFE standard

When the CAFE standard increases, some manufacturers that were compliers now become non-compliers. Therefore, they change the kind of first order condition that characterizes their behavior, moving from equation (5) to (6). Keeping the same marginal costs from the estimation, we solved for the prices in the new system of equations. These new equilibrium prices change the market shares, utility levels, and profits. Compensating variation is computed using (9), and profits are calculated using equations (3) and (4) taking into account that some manufacturers may have switched from one group to the other. We expect a negative effect on profits from the fines that some of the manufacturers now have to pay. But either positive or negative effects from the new market shares according to consumer preferences.

The results are presented in Table 9. Since we assumed there is no increase in gas taxes in this case, equation (8) is used only with an inelastic demand on miles driven and then with individual heterogeneity in miles driven. The increase of 10 percent in the standard is motivated by two factors: first, with a 10 percent increase some manufacturers become non-compliers; and second, this increase is well within the range of potential values of the new CAFE standard that were contemplated in 2007. The welfare loss gross of externalities per gallon of gasoline saved is \$2.11, which is higher than the two proposed taxes above. The fraction of total welfare loss that corresponds to consumer surplus is about 70% when gasoline taxes increase. On the other hand, consumer surplus represents only about half of the total welfare loss in the cases analyzed for CAFE standard increases. For the 5% column in Table 11, there was a decrease in consumption and a net decrease in welfare, but producer surplus actually increased: manufacturers with very high fuel-economy cars increased profits and this offset the losses from those manufacturers that were affected by the new standard since the amount of fees is relatively small compared to profits.

The 27% increase in the standard corresponds to the 35 mpg goal for 2016. The model does not seem to capture this large increase to give a large reduction in gasoline consumption. The effect measured only comes from the change in market shares, not from changes in driving habits, which are assumed to remain the same. That is, these changes do not include a rebound effect. If we accounted for it, since consumers would drive more, the reduction in gas consumption would be even smaller or even positive. Therefore, these results are an upper bound of such effects. One conclusion from these results is that the welfare loss gross of environmental costs from tightening the standard by 10 percent is about two times the cost of increasing net gasoline prices taxes by 10 percent.

7.3 Taxes, CAFE, and externalities

Previous studies have tried to answer the question of whether CAFE standards are welfare increasing (Parry et al. [2007], Portney et al. [2003], Austin and Dinan [2005]). The results in this paper are gross of externalities. We can take from the literature estimates of the social costs of these externalities and subtract them from the welfare losses per gallon of gasoline saved. Table 10 and Table 11 show the welfare losses at different tax levels and CAFE standard increases. Also shown is the value of this welfare loss gross of externalities per ton of CO_2 of emissions saved. This number is easily obtained by dividing the loss per gallon by a factor that converts gallons of gas into average emissions.[38] We see that the costs are increasing with the amount of tax and with the amount of CAFE standard increase. However, at 5% of CAFE increase, there is gain from the policy. Although not directly comparable, the short-run losses from the tax increases are almost always less than the losses from a change in CAFE. When the savings from the externalities avoided are subtracted we obtain graphs Figure 6 and Figure 7. These graphs show evidence that at least in the short run, these policies can even be welfare neutral.

If the current gasoline tax were already equal to the cost of the externalities, then by having CAFE standards we would be adding extra costs. In this case, CAFE would increase consumption of fuel economy, increasing the mileage-related externalities such as local pollution, congestion, and accidents due to the rebound effect. But this is probably not the case since Parry and Small [2005] conclude that the U.S. tax is below its optimal value.

8 Conclusions

This paper presents an application of estimation techniques in industrial organization to an environmental policy problem. The estimated model gives a more complete picture of

[38]EPA http://www.epa.gov/otaq/climate/420f05001.htm

the interaction of demand for cars and the CAFE policy by estimating consumer preferences with random coefficients, accounting for price endogeneity, including income effects, and a model of supply with the specifics of the regulation. The model uses information on market shares, prices, miles traveled, and household income by car model owner to obtain structural demand and supply parameters. It is shown that it is crucial to know the cross- and own-price elasticities in order to understand the impacts of increases in gasoline taxes and in the CAFE standard. This paper is also an example of how to estimate models in which there are not individual data available for all the different pieces in the model. The results show that both policies, gasoline taxes and the CAFE regulation, diminish gasoline consumption. Using data on new cars, we found that the welfare loss gross of environmental costs from tightening the standard by 10 percent is about two times the cost of increasing net gasoline prices by 10 percent.

There are two main characteristics of the regulation that should be added in future work. The first is the used car market. The continuous renovation of the car fleet into one with more fuel-efficient cars is what makes this regulation a success. The second characteristic is the credit system. In a year where a manufacturer's CAFE is above the standard, the NHTSA gives credits to that manufacturer which can be used in the next three years to offset potential fines. A model where the discounted expected profits are maximized could capture these characteristics of the industry. Each period, firms would play a static Bertrand competition game like the one described in this paper and they could make decisions on the technology for fuel efficiency as well.

The scarcity of non-renewable energy sources and the effects of CO_2 emissions to the atmosphere make the study of this regulation an important issue in terms of national security and environmental impacts. The analysis of policies on the regulation of GHG has become of interest due to the ongoing policy changes on these issues. This paper uses techniques from industrial organization to evaluate two policies that are under current discussion.

Appendix A Joint Model of Miles Driven and Car Choice

Consumer has a utility function over the amount of driving m miles and the cost of driving such number of miles at a price $dpm = \frac{p_g+T}{mpg}$, where p_g is the price of gasoline, T is the tax, and mpg is the rated miles per gallon of the car,

$$u_2(m, dpm) = f(m) - dpm \times m,$$

f is a concave function. The amount of miles driven that maximizes this utility function is $m^* = f'^{-1}(dpm)$. At the moment of purchase, consumer maximizes a Cobb-Douglas utility function on the discounted expected utility from driving and the utility from the specific car model she buys,

$$U_1(E(u_2), x, p, y) = (\delta E(u_2))^{\tilde{\alpha}}(g(x, p, y))^{\eta} \exp(\epsilon)$$

where x is a vector of car characteristics, p is the price of the car, y is consumer's income, ϵ is a random shock, and the expected value is over the distribution of the gasoline prices in the future. By taking the log of the utility function and substituting in the amount of miles that maximizes u_2 we get

$$u_1(E(u_2), x, p, y) = \alpha \log(E(f(f'^{-1}(dpm)) - dpm \times f'^{-1}(dpm))) + \eta \log(g(x, p, y)) + \epsilon.$$

Assume $f(m) = -am^2$, which gives increasing returns to utility in some range of miles driven and decreasing returns in utility above certain threshold. Then,

$$u_1(E(u_2), x, p, y) = \alpha \log(dpm) + \eta \log(g(x, p, y)) + \epsilon$$

where

$$g(x, p, y) = \exp\left(x\beta + \sum_{k=1}^{K} \sigma^k \nu^k x^k + \xi\right)(y - p)^{\gamma}.$$

as in (1). Following previous studies on gasoline prices and their effects on car purchases, we assume that the consumer perceives the current value of dpm as the best estimate of future dpm, so we drop the expectation operator from the utility.

Appendix B Estimation and Computation

Demand for Cars

We observe a vector s of size J of market shares. These shares add to one. Since we also want the share for the outside good we multiply these shares by the probability of purchasing a car in that year. This number is obtained from the CEX[39] (share of consumers who bought

[39]Consumer Expenditure Survey http://www.bls.gov/cex/

a car). The share for the outside good is just the difference of 1 and the sum of the car shares, that is, the fraction of the population that does not buy a new car. The idea is to find the parameters such that the predicted market shares and the actual ones are as close as possible. In order to accomplish this, we use a BLP type estimation. Three main contributions in BLP are relevant to this paper: (i) they allow for interaction of individual characteristics and product attributes using random coefficients, giving price elasticities that are not subject to the IIA, (ii) the estimation relies on the generalized method of moments and it can accomodate extra moment conditions, and (iii) their method uses optimal instrumental variables to solve the problem of price endogeneity. The only theorem in Berry et al. [1995] asserts that for any guess of the parameters (α, β) (the "mean" parameters) we can find a vector of unobserved demand shocks ξ that solves the system of market shares using a contraction mapping. Remember that by (1), if we have a guess for (α, β) then we can find the value of the vector ξ. This defines a unique value for the vector δ; therefore, the moment condition becomes $E(\xi_j|Z_j) = E(\delta_{jt} - (\alpha \log(dpm_{jt}) + x_{jt}\beta)|Z) = 0$. The sample analog of this expression suggests a function to minimize with respect to the parameters (α, β, γ).

The algorithm to find the desired moment condition is as follows:

1. Initialize values for δ, (α, β) (we will call all these parameters simply θ) and also initialize ξ.

2. Get the n_s random draws ν_i from normal distribution and from the empirical income distribution for each manufacturer.

3. Simulate the integral by

$$\hat{s}_j(p, \theta) = \frac{1}{n_s} \sum_{i=1}^{n_s} \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{k=1}^{J} \exp(\delta_k + \mu_{ik})}.$$

   where p includes the prices of all the cars.

4. For $h = 1, 2, \ldots$

$$\delta^{(h+1)} = \delta^{(h)} + \log(s) - \log(\hat{s}(p, \theta))$$

   until convergence to get the new δ. This is the contraction mapping.

5. Obtain the unobservables

$$\xi_{jt} = \delta_{jt} - (\alpha \log(dpm_{jt}) + x_{jt}\beta)$$

Supply and CAFE

For the supply side assume the functional form

$$\log(mc_j) = W_j\tau + \omega_j$$

where W_j is a subset of the characteristics X_j (includes mpg_j but it does not include dpm) and some cost shifters like steal price, coal price, etc.

As discussed above, there are two types of manufacturers:

1. Compliers case. The problem for such a firm is

$$\max_p \sum_r (p_r - mc_r)\hat{s}_r(p, \theta)$$

   where r is the index on the cars produced by the same manufacturer, p is the vector of prices of that manufacturer, and instead of writing the problem in terms of quantities, we can write it terms of the shares $\hat{s}_r(p, \theta)$ since they only differ by a constant. The first order conditions in vector form are

$$mc_f = p_f + \Delta_f(p, \theta)^{-1}\hat{s}_f(p, \theta),$$

   the ij-th entry of the matrix Δ is $\partial \hat{s}_j / \partial p_i$ for car i made by firm j and 0 otherwise. Note that because a manufacturer produces more than one car model, the matrix Δ is block diagonal (and therefore the inverse is as well). Each block contains the partial derivatives of market shares of that manufacturer with respect to their own prices, because it does not have control over other manufacturer's prices.

2. Non-compliers case. Here firms prefer to pay the fines rather than comply with the standard.

$$\max_p \sum_r (p_r - mc_r)\hat{s}_r(p, \theta) - 55 \left(\sum_r \hat{s}_r(p, \theta) \right) (c_s - c_f(p, \theta))$$

where c_s is the CAFE standard and c_f is the manufacturer's CAFE. First, we calculate

$$\frac{\partial c_f(p)}{\partial p_j} = \frac{1}{\left(\sum_r \frac{s_r}{mpg_r}\right)^2}\left(\left(\sum_r \frac{s_r}{mpg_r}\right)\left(\sum_r \frac{\partial s_r}{\partial p_r}\right) - \left(\sum_r s_r\right)\left(\sum_r \frac{1}{mpg_r}\frac{\partial s_r}{\partial p_j}\right)\right)$$

where the sums are over the cars produced by this manufacturer. The first order condition is

$$\begin{aligned}&\textstyle\sum_r (p_r - mc_r)\frac{\partial s_r(p,\theta)}{\partial p_j} + s_j(p,\theta)\\ &+55\left(\textstyle\sum_r s_r(p,\theta)\right)\frac{\partial c_f(p,\theta)}{\partial p_j} - 55(c_s - c_f(p,\theta))\textstyle\sum_r \frac{\partial s_r(p,\theta)}{\partial p_j} = 0.\end{aligned}$$

which by stacking those equations we get in vector form

$$\Delta_j(p - mc) + s + 55\left(\sum_r s_r\right)\frac{\partial c_f(p)}{\partial p} - 55(c_s - c_f(p,\theta))\Delta_j * 1_{(J,1)} = 0$$

which gives the marginal cost expression

$$mc = p + \Delta_j^{-1}\left(s + 55\left(\sum_r s_r\right)\frac{\partial c_f(p)}{\partial p} - 55(c_s - c_f(p,\theta)) * \Delta_j * 1_{(J,1)}\right)$$

where Δ_j is the submatrix of Δ of partial derivatives of market shares of firm j with respect to the prices of the cars manufactured by this firm, and $1_{(J,1)}$ is a column vector of ones.

Minimization of the Moment Conditions

The two sets of moments described above cannot be estimated independently because there are parameters that are shared in two or more set of moments and because for each guess of the parameters we get new predicted market shares that enter the profit maximization problem.

Define

$$\psi(\theta) = \begin{bmatrix}\xi(\theta)\\ \omega(\theta)\end{bmatrix} = \begin{bmatrix}\delta_{jt} - (\alpha \log(dpm_{jt}) + x_{jt}\beta)\\ \log(mc_j) - W_j\tau\end{bmatrix}$$

where it should be understood that there are two types of expressions for calculating mc depending on which type of manufacturer is. As explained in the main text, if it is observed in the data that a manufacturer was a non-complier, that is taken as the optimal behavior, and therefore it will be assigned equation (6) as its first order condition and that is going to be fixed througout the entire estimation. Construct the moment condition

$$\begin{aligned} G(\theta) &= \psi' Z \Omega^{-1} Z' \psi \\ \Omega^{-1} &= \left(E(Z'\psi\psi' Z)\right)^{-1} \end{aligned}$$

where Z is the matrix of instruments and it is block diagonal. These instruments are the same used in BLP; they are aggregations of other manufacturers' products. The linearity of some of the parameters allows to write down an analytical solution for the optimal parameters without having to use a search method for the minimization.

The outer loop of the overall estimation can be summarized as follows,

1. For a guess of the parameters get the value of the function $G(\theta)$.
2. Get the linear parameters $\hat{\theta}_{GMM}$.
3. Recompute $G(\cdot)$ with these new values of the linear parameters so that now it will be a non-linear function of remaining parameters. Minimize this function with respect to these non-linear parameters.

Robustness checks

Table 6 shows some robustness checks on the main model. The first column repeats the main results of the paper. Model 2 replaces the time trend with a variable 'age' that is equal to the number of years in the market since the first time that car model was sold for the first time. This quantity is replaced in the consumer utility function but not in the cost function. The coefficient on log(dpm) is still negative but the value of the moment function at the optimal solution is higher than in Model 1, keeping the same instruments. When the variable 'age' is also included in the cost function, we get Model 3, with still a negative coefficient but higher moment value. When no manufacturer dummies are added to either the utility or the cost function, nor any time variable, nor log(q), the coefficient on log(dpm) is positive as shown in Model 4. The sign seems to be very sensitive to the introduction of any time variable as suggested by Model 5. Models 2 and 3 have lower mean Lerner index than Model 1, but they have higher moment value and the coefficients on the other car characteristics are negative. Models 4 and 5 have very high Lerner indexes, which would mean firms have large market power but this contradicts the general perception that the car industry is more competitive than that. This also indicates that quality in cars is largely perceived through the brand name, manufacturer dummies capture this.

References

Allcott, H. and Wozny, N. (2010). Gasoline prices, fuel economy, and the energy paradox. Working paper.

Anderson, S., Kellogg, R., Sallee, J., and Curtin, R. (2011). Forecasting gasoline prices using consumer surveys. *American Economic Review Papers and Proceedings*, 101(3):110–114.

Austin, D. and Dinan, T. (2005). Clearing the air: The costs and consequences of higher CAFE standards and increased gasoline taxes. *Journal of Environmental Economics and Management*, 50:562–582.

Bento, A., Goulder, L., Henry, E., Jacobsen, M., and Haffen, R. V. (2005). Distributional and efficiency impacts of gasoline taxes: An econometrically based multi-market study. *American Economic Review*, 95:2.

Berry, S. (1994). Estimating discrete choice models of product diferentiation. *Journal of Economics*, 25:242–262.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 60(4):889–917.

Busse, M., Knittel, C., and Zettelmeyer, F. (2009). Pain at the pump: The differential effect of gasoline prices on new and used automobile markets. *American Economic Review.*

Copeland, A., Dunn, W., and Hall, G. (2011). Inventories and the automobile market. *RAND Journal of Economics*, 42(1):121–149.

Gillingham, K. (2010). How do consumers respond to gasoline price shocks? heterogeneity in vehicle choice and driving behavior. Working paper.

Goldberg, P. K. (1998). The effects of the Corporate Average Fuel Efficiency Standards in the U.S. *Journal of Industrial Economics*, 46:1–33.

Gramlich, J. (2009). Gas prices, fuel efficiency, and endogeneous product choice in the U.S. automobile industry. Working paper.

Jacobsen, M. (2009). Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. Working paper.

Kleit, A. (2004). Impacts of long-range increases in the Corporate Average Fuel Economy Standard. *Economic Inquiry*, 42:279–294.

Knittel (2010). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review*.

McFadden, D. (1981). Econometric models of probabilistic choice. in Structural Analysis of Discrete Data with Economic Applications.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Parry, W. H. and Small, K. A. (2005). Does Britain or the United States have the right gasoline tax? *American Economic Review*, 95:1276–1289.

Parry, W. H., Walls, M., and Harrington, W. (2007). Automobile externalities and policies. *Journal of Economic Literature*, 45:373–399.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.

Portney, P. R., Parry, W. H., Gruenspecht, H. K., and Harrington, W. (2003). Policy watch: the economics of fuel economy standards. *Journal of Economic Perspectives*, 17(4):203–217.

Small, K. A. and Rosen, H. (1981). Applied welfare economics with discrete choice models. *Econometrica*, 49(1):105–131.

U.S. Department of Transportation, Federal Highway Administration (2009). 2009 National Household Travel Survey. http://nhts.ornl.gov.

West, S. E. (2004). Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics*, 88:735–757.

Yang, S., Chen, Y., and Allenby, G. (2003). Bayesian analysis of simultaneous demand and supply. *Quantitative Marketing and Economics*, 1:251–275.

Tables and Figures

Table 1: Empirical distributions of miles traveled and annual income by manufacturer

manufacturer	**annual miles traveled ('000s)**		**annual income ('000s 2007$)**	
	mean	std.dev.	mean	std.dev.
BMW	10.0	6.2	92.6	21.9
Chrysler	11.7	8.8	64.9	30.4
Ford	12.2	11.0	69.3	29.7
Fuji	11.7	7.3	76.5	29.1
GM	12.3	9.5	70.4	29.6
Honda	11.9	8.7	78.1	28.4
Hyundai	11.7	7.8	64.6	30.6
Kia	11.7	8.9	56.4	28.7
Mazda	11.8	9.2	74.3	30.4
MB	10.8	7.7	91.2	24.3
Mitsubishi	12.8	11.1	63.4	31.3
Nissan	12.6	10.3	73.4	29.5
Porsche	6.2	4.7	98.9	13.7
Toyota	12.3	8.3	74.6	29.6
VW	11.5	8.9	80.9	26.9

Table 2: Summary statistics. 2000-2007

covariate	**mean**	**s.d.**	**min**	**max**
dollars per mile (dpm)	0.09	0.03	0.02	0.21
miles traveled per year	13,765	9,001	199	79,998
acceleration (HP/ton)	214.25	75.29	66.00	550.00
mpg	20.99	5.96	11.00	61.00
width (in)	71.27	3.33	55.50	85.85
length (in)	183.77	13.28	109.40	218.40
height (in)	57.78	5.47	44.30	77.60
market share (no outside good)	0.0043	0.0093	2.9e-7	0.0896
price gallon gas (2007$)	1.91	0.56	1.34	2.80

Table 3: OLS Dependent variable: market share

	1	**2**	**3**	**4**	**5**
dpm	-.0160301**	-.0171999**	-.0173562**	-.0292226**	-.0522166**
	(.0023918)	(.0023839)	(.0024412)	(.0038933)	(.0080592)
price ('000 $)	-.0000259**	-.0000154**	-.0000293**	-.0000336**	-.0001018**
	(4.76e-06)	(4.95e-06)	(5.49e-06)	(5.67e-06)	(.0000113)
acceleration	2.61e-06	3.66e-06*	6.11e-06**	.0000105**	.0000219**
	(1.72e-06)	(1.75e-06)	(1.83e-06)	(1.88e-06)	(3.83e-06)
length	.0000381**	.000029**	.0000323**	.0000323**	.0000872**
	(4.33e-06)	(5.46e-06)	(5.69e-06)	(5.74e-06)	(.0000119)
width	.000026	.0000673**	.0000548**	.0000552**	.0001345**
	(.0000183)	(.0000188)	(.0000195)	(.0000193)	(.0000375)
height	.0000283**	.0000174	.0000121	.0000287+	2.32e-06
	(8.34e-06)	(.0000153)	(.0000156)	(.0000156)	(.0000321)
intercept	-.0068643**	-.0079926**	-.0063878**	-.0051111**	-.0135438**
	(.0009322)	(.0011525)	(.0012376)	(.0012397)	(.0025316)
segment dummies	N	Y	Y	Y	Y
manufacturer dummies	N	N	Y	Y	Y
year dummies	N	N	N	Y	Y
N	5010	5010	5010	5010	1841
R^2	.06	.08	.10	.12	.33

Note: Model 5 uses only the observations corresponding to the first time the car model appeared in the market.

Table 4: Demand Parameters

variable	**mean**	**std. dev.**
non-linear parameters		
log (dpm)	-1.5922^{**} (0.0057)	0.4167 (1.8183)
length	0.0115^{**} (0.0001)	0.0037 (0.0042)
height	0.0297^{**} (0.0001)	0.0253 (0.0175)
acceleration	-0.0027^{**} (0.0001)	0.0069^{**} (0.0033)
time trend	-1.5107^{**} (0.0005)	1.0823^{**} (0.0912)
constant	-25.5123^{**} (0.0236)	2.0600^{**} (0.1260)
linear parameters		
log (income-price)	1.3016^{**} (0.3568)	
BMW	1.2425^{**} (0.0060)	
Chrysler	1.1527^{**} (0.0043)	
Ford	0.4917^{**} (0.0040)	
Fuji	0.5895^{**} (0.0069)	
GM	0.5644^{**} (0.0039)	
Honda	0.7130^{**} (0.0048)	
Hyundai	0.8513^{**} (0.0059)	
Kia	1.1567^{**} (0.0063)	
Mazda	0.5430^{**} (0.0057)	
MB	0.0038 (0.0054)	
Mitsubishi	0.1435^{**} (0.0058)	
Nissan	0.6301^{**} (0.0047)	
Porsche	1.0980^{**} (0.0083)	
Toyota	0.8871^{**} (0.0043)	

Note: Standard errors in parentheses.

Note: Omitted manufacturer in dummy variables: VW.

Table 5: Supply Parameters

variable	
log(length)	-1.3179^{**} (0.0038)
log(height)	0.2559^{**} (0.0027)
log(acceleration)	2.4929^{**} (0.0016)
log(mpg)	0.3681^{**} (0.0024)
time trend	-0.0704^{**} (0.0001)
log (q)	-0.0214^{**} (0.0001)
constant	-7.2571^{**} (0.0267)
BMW	0.3583^{**} (0.0018)
Chrysler	-0.3095^{**} (0.0013)
Ford	-0.0557^{**} (0.0012)
Fuji	-0.6566^{**} (0.0021)
GM	-0.3357^{**} (0.0012)
Honda	-0.2356^{**} (0.0015)
Hyundai	-0.1910^{**} (0.0024)
Kia	0.3403^{**} (0.0017)
Mazda	-0.5186^{**} (0.0057)
MB	0.1468^{**} (0.0016)
Mitsubishi	-0.2449^{**} (0.0017)
Nissan	-0.3590^{**} (0.0014)
Porsche	1.1991^{**} (0.0025)
Toyota	0.0553^{**} (0.0013)

Note: Standard errors in parentheses.

Note: Omitted manufacturer in dummy variables: VW.

Table 6: Coefficient on log(dpm)

	1	**2**	**3**	**4**	**5**
mean	−1.59(0.0057)	−0.12(0.010)	−0.10(0.010)	1.16(0.015)	−1.54(0.006)
standard dev.	0.42(1.8183)	0.50(3.529)	0.51(3.351)	2.28(0.652)	0.71(1.827)
Lerner index	79	62	61	94	96
utility					
mfr. dumm.	Y	Y	Y	N	N
‘age’	N	Y	Y	N	N
trend	Y	N	N	N	Y
mc					
mfr. dumm.	Y	Y	Y	N	N
log(q)	Y	Y	Y	N	N
‘age’	N	N	Y	N	N
trend	Y	Y	N	N	N
moment	99.125	362.544	471.747	303.345	116.075

Table 7: Elasticities

average own-price elasticity of cars (2007)	
1st quartile mpg	−2.03
2nd quartile mpg	−1.95
3rd quartile mpg	−1.57
4th quartile mpg	−0.87
average own-*dpm* elasticity (2007)[a]	
1st quartile mpg	−1.642
2nd quartile mpg	−1.657
3rd quartile mpg	−1.654
4th quartile mpg	−1.647
mean Lerner index (%)	79

Note: Standard errors in parentheses.

[a]Note: At moment of purchase.

Table 8: Effects of tax increase. The average federal plus local tax on gasoline in 2007 was $0.41.

additional tax (2007$)	0.50	1.0
welfare loss ($/gallon saved)[a]	1.13	1.67
perfectly inelastic demand		
annual change in gas cons.(%)	−0.639	−0.641
gallons (millions per year)	−108.5	−108.93
elastic demand[b]		
annual change in gas cons.(%)	−0.801	−0.951
gallons (millions per year)	−136.17	−164.11
heterogeneity in VMT[c] and elast. demand		
annual change in gas cons.(%)	−0.803	−0.967
gallons (millions per year)	−136.47	−164.39
change in outside option share (%)	2.12	2.13

[a]Note: In 2007$

[b]Note: Elasticity of −0.4 as suggested in Parry and Small [2005]

[c]Note: Vehicle-miles-traveled using NHTS data.

Table 9: Effects of 10% increase in CAFE standard

welfare loss ($/gallon saved)[a]	2.11
no heterogeneity in VMT[b]	
annual change in gas cons.(%)	−0.00034
gallons (millions per year)	−0.058
heterogeneity in VMT[b]	
annual change in gas cons.(%)	−0.00026
gallons (millions per year)	−0.044
change in outside option share (%)	0.011

[a]Note: In 2007$

[b]Note: Vehicle-miles-traveled using NHTS data.

Table 10: Effects from Taxes. All values in 2007$.

tax increase	0.25	0.50	0.75	1.00
welfare loss per gallon of gas. cons. saved	1.09	1.13	1.17	1.68
welfare loss per tCO_2 saved	124.19	128.44	132.55	190.32
annual change in gas. cons. (%)	−0.442	−0.803	−1.099	−0.967
CS loss as a % of total welfare loss	67.97	69.42	70.75	72.01

Table 11: Effects from tightening the standard. All values in 2007$.

CAFE standard increase	5%	7.5%	10%	27%
welfare loss per gallon of gas. cons. saved	−0.57	2.22	2.11	2.126
welfare loss per tCO_2 saved	−64.24	251.81	239.47	241.59
annual change in gas. cons. (%)	−0.00014	−0.00022	−0.00026	−0.00056
CS loss as a % of total welfare loss		45.61	47.83	54.87

Figure 1: Car prices

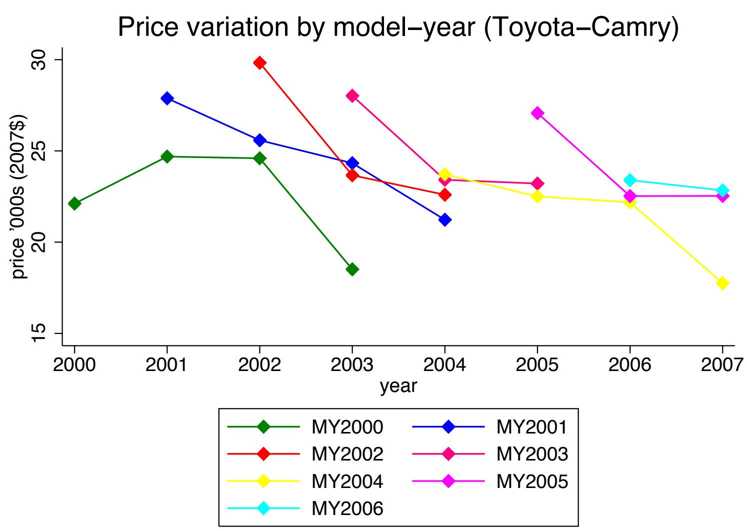

Figure 2: CAFE for selected manufacturers

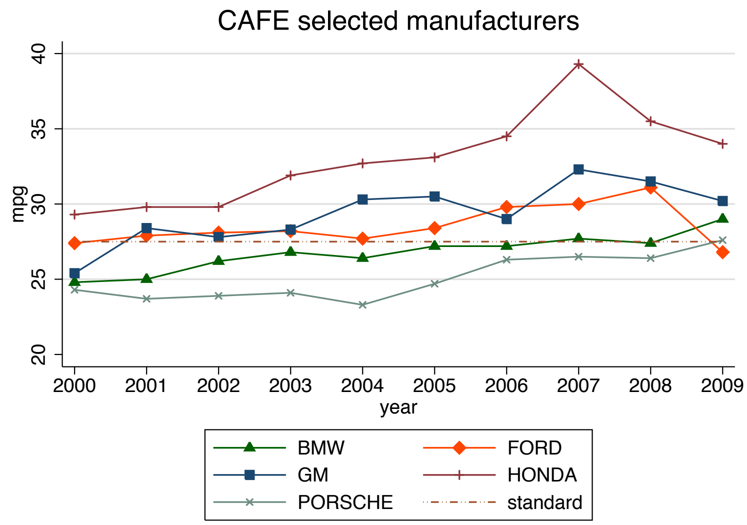

Figure 3: Circles proportional to size of market shares. Assuming 12,000 miles traveled per year in average for 2000.

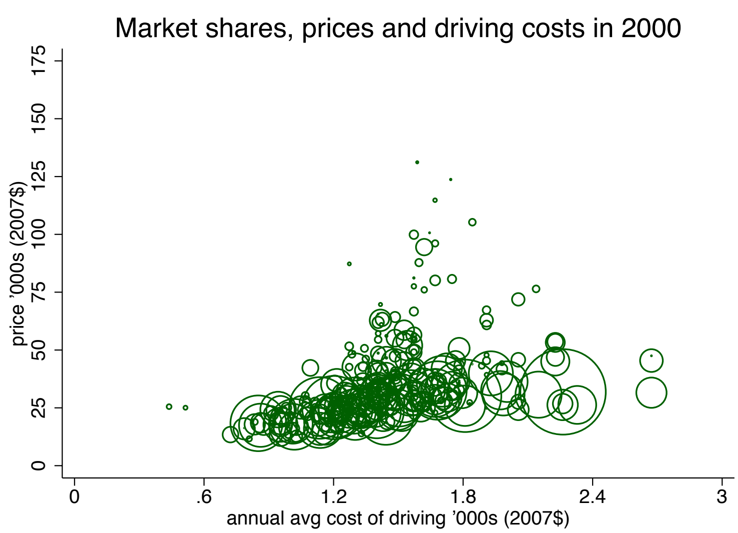

Figure 4: Circles proportional to size of market shares. Assuming 12,000 miles traveled per year in average for 2007.

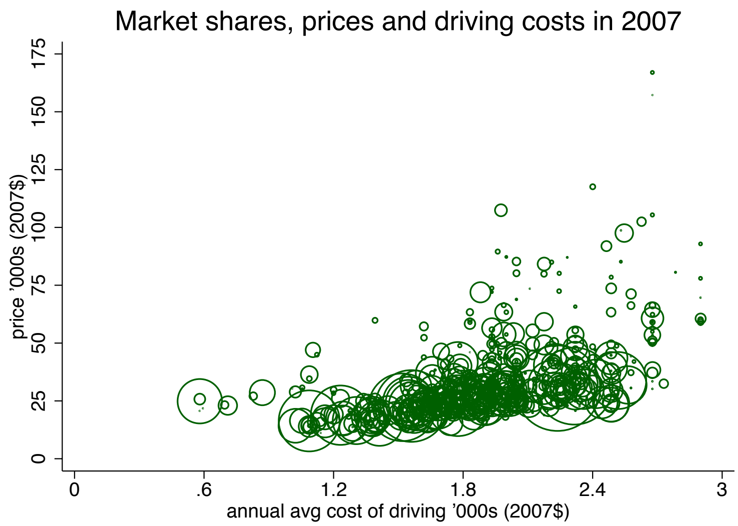

Figure 5: Each circle represents the market share in 2009 for major car models. The annual miles driven and income are averages by car model from reported households in the NHTS

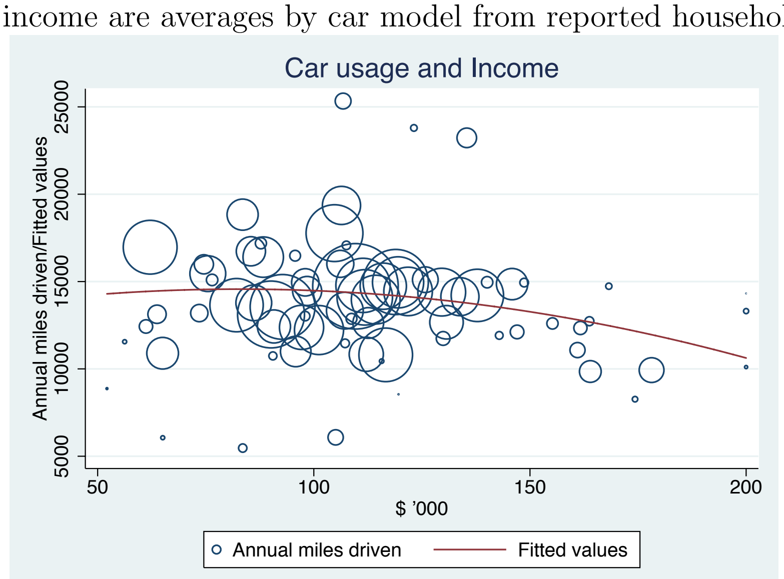

Figure 6: Counterfactual results for increases in gasoline taxes (over the $0.41/gal of average total tax in 2007). Graph shows the marginal welfare loss per gallon of gasoline saved due to the change in the policy. The dashed lines show the welfare loss net of the marginal cost of externalities of gasoline consumption taken from Parry et al. [2007]. The fuel-related costs include greenhouse warming ($0.06/gal), and oil dependency ($0.12/gal). The mileage-related costs include local pollution ($0.42/gal), congestion ($1.05/gal), and accidents ($0.63/gal). All values in 2007$.

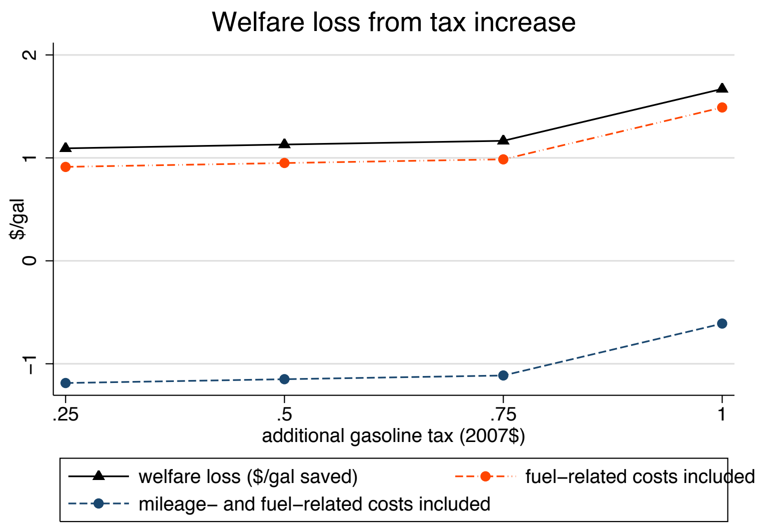

Figure 7: Counterfactual results for changes in the CAFE standard. Graph shows the marginal welfare loss per gallon of gasoline saved due to the change in the policy. The dashed lines show the welfare loss net of the marginal cost of externalities of gasoline consumption taken from Parry et al. [2007]. The fuel-related costs include greenhouse warming ($0.06/gal), and oil dependency ($0.12/gal). The mileage-related costs include local pollution ($0.42/gal), congestion ($1.05/gal), and accidents ($0.63/gal). All values in 2007$.