ABSTRACT

Title of dissertation: ESSAYS ON THE ECONOMICS OF AUTOMOBILE FUEL ECONOMY

Takahiko Kiso, Doctor of Philosophy, 2013

Dissertation directed by: Professor Kenneth McConnell
Department of Agricultural and Resource Economics

This dissertation consists of three chapters that analyze issues relating to automobile fuel economy. Chapter 1 discusses automobile fuel economy regulations in the United States. The new U.S. Corporate Average Fuel Economy (CAFE) standards not only tighten the target fuel economy to be achieved by automakers, but also make significant changes to the design/structure of CAFE standards by introducing three policy instruments (footprint-based targets, intra-firm transferring of fuel efficiency credits between passenger cars and light trucks, and inter-firm trading of fuel efficiency credits). While there are a number of previous studies on the impact of tightening CAFE standards, economists have paid little attention to the design of CAFE standards. I use policy simulation to evaluate these policy instruments relating to the design of CAFE standards. First, I model and estimate the demand- and supply-sides of the U.S. vehicle market using various data sets. Then, based on the estimation results, I simulate the vehicle market and the demand for driving under four counterfactual CAFE policies with different designs, and examine the impacts of

the three policy instruments. Simulation results suggest: (1) footprint-based targets have little impact on market shares, producer profits, consumer surplus, and gasoline use; (2) inter-firm credit trading lowers overall compliance costs by about $110-$140 million, and thus increases social welfare; and (3) allowing intra-firm credit transferring (but not inter-firm credit trading) reduces aggregate gasoline consumption by 0.1-0.25%.

Chapter 2 proposes a new approach to analyzing how automobile fuel economy is valued in the market, using a hedonic regression framework. A distinctive feature of my approach is the use of each vehicle's miles traveled: a consumer's marginal willingness to pay (MWTP) for fuel economy is inferred with her vehicle's miles traveled. With the inferred MWTP, we apply the steps of the standard hedonic method backward and estimate each vehicle's marginal and total price of fuel economy, and consumers' discount rate for future fuel cost savings. We find that the standard hedonic method may not provide a stable and reasonable estimate of the value of fuel economy, likely due to the omitted variable bias from vehicle attributes such as safety features, interior equipment and reliability. This method makes it possible to separate the portion of vehicle price that is attributable to fuel economy, and significantly alleviates the omitted variable bias. Applying the procedure to model year 2001 vehicles in the U.S. market, we estimate that consumers discount future fuel cost savings at the annual rate of 26-43%, that for the middle case of the discount rate of 34%, the price of a 0.1 gallon per 100 miles improvement in fuel efficiency is on average $75 (in 2000 U.S. dollars), and that for the same case, the average total price of fuel economy is $1,950. We also find that larger, less fuel efficient vehicles tend to have higher marginal and total prices of fuel economy.

Chapter 3 examines whether Japanese fuel economy regulations established in the 1990s induced technological progress in Japanese automakers' technology for providing fuel economy. By observing how fuel economy of automobiles has improved after controlling for changes in vehicle characteristics such as weight and power, I find that fuel economy improvement accelerated after regulations were introduced, implying induced innovation in fuel efficiency technology.

ESSAYS ON THE ECONOMICS OF AUTOMOBILE FUEL ECONOMY

by

Takahiko Kiso

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:
Professor Kenneth McConnell, Chair
Professor Anna Alberini
Professor Maureen Cropper
Professor Erik Lichtenberg
Professor Roberton Williams III

Acknowledgments

To my advisor, Ted McConnell, for his guidance and encouragement throughout my studies at the Ph.D. program. He not only gave me excellent suggestions on my research, but also valuable advices on many other aspects of an academic researcher including teaching and presentation.

To my committee members—Maureen Cropper, Rob Williams, Anna Alberini, and Erik Lichtenberg—for always being approachable and helpful.

To Ginny McConnell and Shanjun Li for having discussions with me and providing comments and suggestions.

To my friends at the Department of Agricultural and Resource Economics for helping me develop research ideas through many informal discussions.

To the University of Maryland, the Economic Club of Washington, D.C., and Nishihara Cultural Foundation for their generous financial support.

Last but not the least, to my family—Clara and Mia—for their love, support, and patience all the time.

Table of Contents

List of Figures

Chapter 1: Evaluating New Policy Instruments of the Corporate Average Fuel Economy Standards: Footprint, Credit Transferring, and Credit Trading

1.1 Introduction

Reducing automotive fuel consumption has been an important policy issue in the United States. Vehicle fuel combustion is a large source of greenhouse gas emissions and local air pollutants.[1] In addition, large demand for gasoline increases imports of foreign oil and makes the U.S. economy more sensitive to oil supply disruptions and price shocks.[2] Different types of policies have been implemented to cut automotive fuel use, including gasoline taxes, Corporate Average Fuel Economy (CAFE) standards, and the gas guzzler tax. Among these policies, this paper focuses on CAFE standards, a set of federal regulations that are believed to have played a key role in curbing automotive fuel consumption. CAFE standards have been administered by the National Highway Traffic Safety Administration (NHTSA) and, roughly speaking, require an automaker to achieve a certain level of average fuel economy in its

[1] For example, in 2010 carbon dioxide emissions from vehicle fuel combustion accounted for 25.6% of total U.S. greenhouse gas emissions (Environmental Protection Agency, 2012, Table 2-12).

[2] Imported crude oil accounts for about 50% of U.S. consumption.

fleet.

Recently, CAFE standards have been radically reformed in terms of stringency and structure.[3] With the Energy Independence and Security Act of 2007 and the initiative of the Obama Administration, CAFE standards have been aggressively tightened. The current goal is to improve the average fuel economy of new passenger cars and light-duty trucks to 34.1 miles per gallon (mpg) in model year (MY) 2016, and to 40.3-41.0 mpg in MY 2021, as compared to 29.3 mpg achieved in MY 2010. NHTSA expects large economic benefits of the reformed fuel economy regulations: The net benefits over the useful lives of MY 2012-2016 vehicles is estimated to be \$130.7 billion (at a 3% discount rate) or \$94.5 billion (at a 7% discount rate), in 2007 dollars (Environmental Protection Agency and Department of Transportation, 2010, Table I.C.1-1).

In addition to increased stringency, CAFE standards have also experienced significant structural/design changes. From the 1970s until recently, the standards had required each automaker's (sales-weighted) average fuel economy in each vehicle category to exceed the category's target value, which is common to all automakers.[4] The new CAFE standards have completely reformed this structure. The three most important instruments introduced in the standards are footprint-based targets, intra-firm credit transferring and inter-firm credit trading.

Footprint-based targets make an automaker's target value dependent on the size

[3]Beginning in model year 2012, automobile fuel economy is additionally regulated by the Environmental Protection Agency's national greenhouse gas (GHG) emissions standards as well. These two sets of federal regulations form a consistent, harmonized national program for improving automobile fuel economy.

[4]The three vehicle categories are domestic passenger cars, import passenger cars, and light-duty trucks.

of vehicles it sells. In each vehicle category, a firm that on average produces smaller vehicles needs to attain better average fuel economy than another firm that on average produces larger vehicles. As compared to conventional standards, footprint-based standards are intended to reduce the incentive for automakers to downsize vehicles in order to improve fuel economy and comply with the standards. Downsizing vehicles is the easiest way to improve fuel economy, but smaller vehicles are considered to have higher fatality risks than larger vehicles if involved in a traffic accident.

The other two instruments provide flexibility in how automakers meet the targets, and they are expected to reduce the overall costs for automakers to comply with the standards. Intra-firm credit transferring allows an automaker to over-comply in a category and earn credits, then use the credits to offset under-compliance in another category. With inter-firm credit trading, an automaker can sell extra credits from over-compliance to other automakers; conversely, it can buy credits from other automakers to offset under-compliance.[5] These new instruments have structurally changed CAFE standards, and may have various impacts on the vehicle market and gasoline consumption.

A number of previous studies have analyzed how CAFE standards and other policies affect consumers, producers and gasoline consumption. For this purpose, these studies model consumer demand for automobiles and producer behavior in an imperfectly competitive auto market. Then, the demand and supply models predict how consumers and producers react to a given policy, allowing them to compare the

[5] For a complete description of these instruments, see Environmental Protection Agency and Department of Transportation (2010).

impacts of different policies through counterfactual policy simulations. For example, Goldberg (1998), Kleit (2004), Klier and Linn (2012), Jacobsen (forthcoming) and Whitefoot et al. (2012) analyze a tightening of CAFE standards (i.e., an increase in target fuel economy that automakers need to achieve). Bento et al. (2009) examine an increase in U.S. gasoline taxes. Austin and Dinan (2005) compare the effects of tightening CAFE standards and increasing gasoline taxes.

I also analyze CAFE standards based on a similar framework, but I focus on a different aspect of CAFE standards than most of the above-mentioned studies. Specifically, this paper aims to examine the structure/design of CAFE standards, rather than the stringency of the standards as in most previous studies. I do this by simulating the impact of counterfactual policies that differ in structure, but not in stringency. These counterfactual policies can feature one or more of the three instruments above.

Among studies that use counterfactual simulations to analyze CAFE standards, very few have looked at these instruments. Coleman and Harrington (2010) analyze footprint-based targets and intra-firm credit transferring between categories, but not inter-firm credit trading. Austin and Dinan (2005) look at inter-firm credit trading without considering footprint-based targets and intra-firm credit transferring.[6]

This paper differs from these previous studies on the design of CAFE standards in several ways. First, I consider all of the three instruments introduced in the reformed CAFE standards. As discussed above, the literature has not simulated

[6]Although Whitefoot et al. (2012) also consider footprint-based targets and intra-firm credit transferring, their focus is on tightening of CAFE standards, but not on footprint-based targets and credit transferring. As an extension to his analysis, Jacobsen (forthcoming) simulates a simplified version of footprint-based targets.

a case where the three instruments are in effect simultaneously. Since the actual reformed CAFE standards include all of these, it is important to examine how they altogether affect consumers, producers and gasoline consumption.

Second, I make use of recent developments in the economic modeling and estimation of the vehicle market, which would lead to a more accurate evaluation of the instruments of interest. In particular, following the approach of Bento et al. (2009), I simultaneously estimate consumers' vehicle purchase and use decisions.[7] This approach can analyze the two connected decisions in a unified framework that is consistent with economic theory. For example, it can account for the tendency that households with high driving demand consider fuel economy an important attribute in making a vehicle purchase decision, as compared to households with low driving demand. Previous studies on the structure of CAFE standards model vehicle choice decisions only. Since the value of fuel economy to consumers crucially depends on how much the vehicle is driven, incorporating a model of VMT demand is essential in analyzing policies relating to fuel economy and gasoline consumption. Indeed, my simulations suggest the fact that light-duty trucks are typically driven longer distance than passenger cars greatly influences the impacts of policy instruments on aggregate gasoline consumption.

Third, I carefully set parameters of counterfactual standards, which determine target fuel economy levels of each automaker, so that the market average fuel efficiency at the simulated equilibrium is almost the same across all policies.[8] This

[7]Vehicle use is measured by vehicle miles traveled, or VMT.

[8]Roughly speaking, by adjusting the stringency of the standards, I (as the policymaker) can control the fuel economy level achieved by automakers that are marginally complying with the standards. Thus, I can also affect the market average fuel economy realized in the equilibrium.

ensures that different policies are essentially at the same level of stringency as measured by the market average fuel efficiency. The difference is less than 0.1% between any counterfactual policies in this paper, while in Coleman and Harrington (2010) it is in some cases larger than 1%. Since I want to evaluate the effectiveness of different policy instruments, it is important that counterfactual policies featuring these instruments are at the same level of stringency. Otherwise, it is unclear whether differences in simulation outcomes are attributable to differences in structure or in stringency. Setting counterfactual policies at the same level of stringency makes it possible to directly compare the effectiveness of different policy instruments.

My simulations overall provide the following policy implications. First, footprint-based targets do not have significant impacts on producer profits, consumer surplus, gasoline use, and market shares of different vehicle models. How targets are set does not matter so much compared to other policy instruments.

Second, allowing intra-firm credit transferring but not inter-firm credit trading is effective in reducing aggregate gasoline consumption. This policy induces automakers to improve light-duty trucks' fuel economy and worsen passenger cars'. Because light-duty trucks are typically driven longer distance, this leads to a reduction in aggregate gasoline use (by about 0.1-0.25%). Under the circumstances in which gasoline is expensive and/or externalities from gasoline use are taken seriously, this policy option would become more attractive for the regulator and the society.

Third, introducing inter-firm credit trading lowers compliance costs to achieve a given level of market average fuel economy by about $110-140 million. This is because inter-firm trading shifts the role of improving fuel economy at additional production costs to automakers that can do so more cheaply. Due to this production cost reduction, inter-firm credit trading increases social welfare. Indeed, I find that the policy with inter-firm trading gives the highest social welfare among the counterfactual policies considered. However, allowing inter-firm credit trading in addition to intra-firm credit transferring nullifies the effect of the latter on cutting aggregate gasoline consumption.

The rest of the paper is organized as follows. Section 2.6 describes the datasets used in the study. Section 1.3 focuses on the demand side of the new vehicle market, and estimates a model of consumers' vehicle choice and VMT demand. The estimates from Section 1.3 are called on repeatedly in later sections. Section 1.4 considers automakers' profit maximization in an imperfectly competitive U.S. new vehicle market and under the actual CAFE standards. From first order conditions, I derive estimates of how improving fuel efficiency changes the production costs of vehicles. Based on the results of Sections 1.3 and 1.4, Section 1.5 simulates four counterfactual CAFE standards to evaluate and compare the effects of footprint-based targets, credit transferring and credit trading. Section 1.6 concludes.

1.2 Data

I use household and vehicle data from 2001 to evaluate the impacts of different mechanisms of CAFE standards. I estimate households' vehicle and VMT choices by using 2001 National Household Travel Survey, and data from various sources on gasoline prices and vehicle sales/prices/attributes in 2001. Then, I analyze the supply side of the year 2001 vehicle market, with taking into account actual CAFE standards and each automaker's compliance status at that time. Counterfactual policies are simulated based on year 2001 data and estimates from the demand- and supply-side models. Below I first describe the datasets used in this study, and then discuss why I choose to focus on year 2001.

1.2.1 Description of the Datasets

The data used for this study come from various sources. First, household data are from 2001 National Household Travel Survey (NHTS), a national survey conducted by the Department of Transportation. This survey contains information on the vehicle(s) each surveyed household owns (such as make, model and model year) and the estimated annual VMT of the vehicle(s), as well as household characteristics (such as income and household size). Additionally, gasoline prices that households face come from state level data from the Energy Information Administration.

Second, the data on vehicle models (such as sales, price, fuel economy, horse power, weight and footprint) are obtained from several sources. They are *Wards Automotive Yearbook*, the Environmental Protection Agency's (EPA) "Fuel Economy Test Car List Data", MSN Autos (`http://home.autos.msn.com/`) and manufacturers' automotive fuel economy reports submitted to the National Highway Traffic Safety Administration.

Among these datasets, EPA's "Fuel Economy Test Car List Data" and manufacturers' automotive fuel economy reports contain detailed, disaggregate information on vehicle attributes and market sales.[9] In particular, manufacturers' automotive fuel economy reports have not been used in previous studies, and while other sources include U.S. market sales data only up to the nameplate level (e.g., Chevrolet Malibu, Ford Explorer, Toyota Camry), manufacturers' reports provide it at a more disaggregate level. Specifically, they distinguish vehicle configurations that are under the same nameplate but with different specifications (in terms of such attributes as engine size, weight class, horsepower and transmission), and have information on each configuration's model year sales in the entire U.S. market. This detailed information allows me to distinguish different models (e.g., Toyota Camry CE and Toyota Camry LE) within the same nameplate in the following estimation and simulation, while other studies at the U.S. market level generally cannot do so, and they analyze only up to the nameplate level. In practice, different models under the same nameplate are often equipped with very different attributes (e.g., engine size, weight class, horsepower and transmission), and consumers choose to buy, and automakers set the price and attributes of, each model, but not each nameplate. Thus, using these datasets makes estimation and simulation closer to the choices faced by actual consumers and producers.

On the other hand, the NHTS identifies vehicles only up to the nameplate level, so the model level information on households' vehicle choice is unavailable. That is,

[9]Fuel economy values contained in these sources are unadjusted values used for CAFE standards. Fuel economy ratings that consumers see are adjusted by the EPA to account for actual driving conditions. Roughly, adjusted values are 15% less fuel efficient. Throughout this paper, I take account of the difference between unadjusted and adjusted fuel economy: I use adjusted fuel economy when dealing with the consumer side (e.g., the price of driving, gasoline consumption), and unadjusted fuel economy when calculating the (corporate or market) average fuel economy, and calculating credits.

in the data, we observe only the nameplate purchased by each household, but not the model under the nameplate. Below I construct an estimation method that enables analysis of vehicle choice at the model level, even when choice data at that level is unavailable.

Table 1.1 reports summary statistics of the sample of households and vehicles used in this study. The sample consists of 5884 households who purchased new model year (MY) 2001 vehicles, and 457 MY 2001 vehicle models under 185 nameplates.

1.2.2 Reasons for Using Year 2001 Data

The National Household Travel Survey (NHTS), which contains information on households' vehicle and VMT choices, is conducted between mid-2001 and mid-2002 (2001 NHTS), and between mid-2008 and mid-2009 (2009 NHTS). Since I analyze households' purchase and use of new vehicles, these survey periods imply that focusing on model year (MY) 2001 vehicles or MY 2008 vehicles is most appropriate for my analysis.[10]

I choose to concentrate on 2001 rather than 2008 because gasoline prices were relatively stable since the 1990s until around 2002, but highly volatile around 2008. The model used in my estimation and simulation (as well as in many previous studies discussed above) is an equilibrium model, in which, roughly speaking, manufacturers

[10]The NHTS does not tell us whether a vehicle in the sample is purchased new or used. Yet, MY 2001 vehicles in the 2001 NHTS and MY 2008 vehicles in the 2009 NHTS are mostly purchased new. A vehicle from a prior model year (i.e., MY 2000 and earlier for the 2001 NHTS, and MY 2007 and earlier for the 2009 NHTS) is more likely to be purchased used, and in such a case we cannot obtain the desired information about the household who purchased it new. The 2001 NHTS also includes MY 2002 vehicles, and the 2009 NHTS includes MY 2009 vehicles, but the number of observations is much smaller than that for MY 2001 or 2008 vehicles.

determine fuel economy of vehicles to respond to consumers' demand (willingness to pay) for it. The price of gasoline is a crucial factor affecting consumers' demand for fuel economy: The higher the price, the more consumers are willing to pay for better fuel economy. In the actual market, consumers' demand for fuel economy can respond quickly to gasoline price changes, while automakers cannot adjust fuel economy very frequently to respond to the demand changes. Thus, in the actual vehicle market, gasoline price changes disturb the "market" for fuel economy in a way that demand and supply do not balance. Therefore, applying the equilibrium model to analyze a situation in which gasoline prices are changing significantly may lead to inaccurate results. Gasoline prices in the U.S. had been very stable since the 1990s until around 2002, and then started to increase sharply and became volatile. Gasoline prices were extremely volatile around 2008.[11] For this reason, I analyze the vehicle market in year 2001, whose conditions are more appropriate for the structure of my model than year 2008's.

Simulations based on the 2001 data can provide useful insights for understanding and predicting manufacturer behavior under the current CAFE standards as well. This is because automakers in 2001 were overall under the circumstances that are qualitatively similar with respect to fuel economy and CAFE standards to those faced by automakers in more recent years, and my simulations reproduce these circumstances. Table 1.2 shows selected automakers' standard (target) to be achieved, average fuel economy (in mpg) and sales volume in 2001 and 2012 (or 2011).[12] The

[11]In 2008, the highest price is over $4.10, and the lowest is around $1.60 (weekly U.S. regular gasoline prices per gallon).

[12]The data are taken from National Highway Traffic Safety Administration (2012). Model year (MY) 2012 data are projected values. Sales data for 2012 are still unavailable, so 2011 data are

firms shown in the table are the six largest automakers in the U.S. market in 2001, and continue to be so in the current market as well. Both in 2001 and 2012, the U.S. Big Three usually achieve lower fuel economy, and seem more constrained by CAFE standards than Japanese automakers.[13] My simulations reproduce similar situations, and thus can provide helpful policy implications for analyzing the current CAFE standards.

1.3 Demand Side of the New Vehicle Market

This section considers the demand side of of the new vehicle market. I set up and estimate a model of household vehicle and VMT choice with data from model year 2001, and discuss estimation results. Fundamentally, the following model resembles the discrete-continuous model of consumer demand pioneered by Dubin and McFadden (1984) and also used in the vehicle literature to analyze the discrete choice of vehicle purchase and the continuous choice of VMT (e.g., Mannering and Winston, 1985; West, 2004). As in other papers based on Dubin and McFadden (1984), my model is a structural model of consumer demand for vehicle purchase and utilization, and the discrete and continuous choices are connected by Roy's identity. Unlike the sequential estimation approach of Dubin and McFadden (1984), I employ a variant of the more recent approach developed by Bento et al. (2009) that simultaneously es-

shown instead. Honda's MY 2012 standards are strangely low in all vehicle categories; they are most likely wrong and obtained because targets were calculated with a formula for MY 2011 by mistake.

[13]Footprint-based targeting makes MY 2012 targets differ across automakers even within the same category. In 2012, intra-firm credit transferring is available as well as inter-firm credit trading. With intra-firm transferring, Ford's under-compliance import passenger cars can be offset by its over-compliance in other categories, and so is General Motors, Toyota and Nissan's under-compliance light duty trucks.

timates the two choices in a full information maximum likelihood framework. Other papers using a similar approach include Gillingham (2011) and Spiller and Stephens (2012).

1.3.1 Econometric Model

Suppose that consumer i is choosing a new vehicle to purchase. The household's decision making process can be modeled as follows. Given a utility function and vehicle models available in the market, the household predicts the utility level that will be attained from owning each vehicle. Comparing all vehicles, it decides to buy the one that is predicted to give the highest utility. In calculating the utility from owning vehicle model jk (model k of nameplate j), household i considers the following utility maximization problem.[14]

$$V(r_{ijk}, w_{ijk}; \mathbf{\Omega}_{ijk}) = \max_{\tilde{q},\tilde{m}} U(\tilde{q}, \tilde{m}; \mathbf{\Omega}_{ijk}) \quad \text{s.t.} \quad \tilde{q} + r_{ijk}\tilde{m} \leq w_{ijk}, \tag{1.1}$$

where $\tilde{m}$ is vehicle miles traveled (VMT) measured in 100 miles; $\tilde{q}$ is the amount of the numéraire good (the composite of all goods/services other than VMT) consumed; r_{ijk} is the price of household i driving model jk for 100 miles (per-gallon price for model jk's fuel at household i's location times model jk's fuel consumption rate in gallons per 100 miles); w_{ijk} is household i's budget available after purchasing model

[14] A "nameplate" refers to, for example, Toyota Camry, and a "model" refers to Toyota Camry CE or Toyota Camry LE. A nameplate may include several models, which typically differ in powertrain specifications.

jk; and $\boldsymbol{\Omega}_{ijk}$ is a vector of variables that are predetermined for this problem, but shift the utility function, such as household characteristics and vehicle attributes. $V(r_{ijk}, w_{ijk}; \boldsymbol{\Omega}_{ijk})$ is the indirect utility function resulting from the maximization problem, which is interpreted as the maximum achievable utility if household i purchases model jk. For the sake of econometric tractability, this problem explicitly considers only one period, while vehicles are likely used over years. I will discuss below (in subsection 1.3.2) how to incorporate the multi-period nature of vehicle ownership into this framework.

For the purpose of estimation, I parameterize the indirect utility function $V(r_{ijk}, w_{ijk}; \boldsymbol{\Omega}_{ijk})$ as follows.

$$V_{ijk} \equiv V(r_{ijk}, w_{ijk}; \boldsymbol{\Omega}_{ijk}) = \tilde{V}_{ijk} + \epsilon_{ijk}, \tag{1.2}$$

$$\tilde{V}_{ijk} = \alpha_{1i} w_{ijk} - \alpha_{2ijk} r_{ijk}^{\alpha_3} + \sum_h \eta_{hi} d_{hjk} + \xi_{jk}, \tag{1.3}$$

$$\alpha_{1i} = \exp(\boldsymbol{\beta}_1 \boldsymbol{x}_i + \delta_{1i}), \tag{1.4}$$

$$\alpha_{2ijk} = \exp(\boldsymbol{\beta}_2 \boldsymbol{x}_i + \boldsymbol{\gamma}_2 \boldsymbol{z}_{jk} + \delta_{2i}), \tag{1.5}$$

$$0 < \alpha_3 < 1, \tag{1.6}$$

where ϵ_{ijk} is an iid error drawn from the standard type I extreme value distribution for each i-jk combination; d_{hjk} is a dummy variable that takes 1 if model jk is of vehicle type h and 0 otherwise;[15] ξ_{jk} is model jk's fixed effect; $\boldsymbol{x}_i$ is a vector of household i's (observed) characteristics; $\boldsymbol{z}_{jk}$ is a vector of model jk's attributes; δ_{1i},

[15] I consider four categories: $a = 1$ for passenger cars, 2 for vans, 3 for SUVs, and 4 for pickup trucks.

δ_{2i} and η_{hi}'s represent household i's unobserved characteristics and are randomly and independently drawn from different normal distributions.[16]

Note that V_{ijk} is the (maximum) utility level household i can achieve from purchasing model jk, and that ϵ_{ijk} is drawn independently from the standard type I extreme value distribution. Therefore, conditional on household i's unobserved heterogeneity $\boldsymbol{\Delta}_i \equiv (\delta_{1i}, \delta_{2i}, \eta_{1i}, \eta_{2i}, \eta_{3i}, \eta_{4i})$ (as well as other household characteristics, and vehicle prices/attributes), the probability that household i purchases model jk is

$$\Pr_i(jk|\boldsymbol{\Delta_i}) = \frac{\exp\left[\tilde{V}_{ijk}\right]}{\sum_{j'}\sum_{k'}\exp\left[\tilde{V}_{ij'k'}\right]}, \tag{1.8}$$

where $\tilde{V}_{ijk}$ is as defined in (1.3)-(1.6). I explicitly write the left hand side as the conditional probability given $\boldsymbol{\Delta_i}$ in order to explain the method of maximum simulated likelihood below.

Applying Roy's identity to the indirect utility function (1.2) and taking the logarithm of both sides yields

$$\begin{aligned}
\log(\tilde{m}_{ijk}) &= \log(-\frac{\partial V_{ijk}/\partial r_{ijk}}{\partial V_{ijk}/\partial w_{ijk}}) \\
&= -\log(\alpha_{1i}) + \log(\alpha_{2ijk}) + \log(\alpha_3) + (\alpha_3 - 1)\log(r_{ijk}) \\
&= -\delta_{1i} + \delta_{2i} + \log(\alpha_3) + (-\boldsymbol{\beta}_1 + \boldsymbol{\beta}_2)\boldsymbol{x}_i + \boldsymbol{\gamma}_2\boldsymbol{z}_{jk} + (\alpha_3 - 1)\log(r_{ijk}), \quad (1.9)
\end{aligned}$$

[16]It can be shown that this indirect utility function is associated with the following quasi-linear utility function:

$$U(\tilde{q}, \tilde{m}; \boldsymbol{\Omega}_{ijk}) = \alpha_{1i}\tilde{q} - (1-\alpha_3)\left(\frac{\alpha_{1i}}{\alpha_3}\right)^{\frac{\alpha_3}{\alpha_3-1}} \alpha_{2ijk}^{\frac{1}{1-\alpha_3}} \tilde{m}^{\frac{\alpha_3}{\alpha_3-1}} + \sum_h \eta_{hi} d_{hjk} + \xi_{jk} + \epsilon_{ijk}. \tag{1.7}$$

Note that if $0 < \alpha_3 < 1$, $U(\tilde{q}, \tilde{m}; \boldsymbol{\Omega}_{ijk})$ is concave in $\tilde{m}$.

where $\tilde{m}_{ijk}$ is VMT planned by household i at the time of purchase given that household i decides on model jk. The logarithm of planned VMT is expressed as a linear function of household characteristics, vehicle attributes and (the logarithm of) the price of driving r_{ijk}. In particular, note that $\alpha_3 - 1$ is interpreted as the elasticity of VMT with respect to the price of driving, conditional on vehicle choice.

In practice, the household's planned VMT is unobservable to the econometrician. What he observes is VMT reported in the NHTS. Therefore, in order to work with observed VMT, I add to (1.9) an (independently, identically and normally distributed) error term μ_{ijk}, which corresponds to the difference between observed and planned VMT:

$$\begin{aligned}\log(m_{ijk}) &= \log(\tilde{m}_{ijk}) + \mu_{ijk} \\ &= -\delta_{1i} + \delta_{2i} + \log(\alpha_3) + (-\boldsymbol{\beta}_1 + \boldsymbol{\beta}_2)\boldsymbol{x}_i + \boldsymbol{\gamma}_2 \boldsymbol{z}_{jk} + (\alpha_3 - 1)\log(r_{ijk}) + \mu_{ijk},\end{aligned} \tag{1.10}$$

where m_{ijk} is observed VMT (in 100 miles) reported in the NHTS for household i who purchased model jk.[17]

Conditional on household i purchasing model jk and planning to drive it $\tilde{m}_{ijk}(\times 100$ miles), (1.10) implies that the probability that we observe household i

[17] For the sample of vehicles analyzed in this study (model year 2001 vehicles), the NHTS reports annual VMT estimates for the first year after vehicle purchase, while at the time of purchase consumers might consider VMT and fuel spending over a period longer than a year. Later in this section I will consider how to account for this difference.

drive the vehicle $m_{ijk}(\times 100)$ miles is

$$\Pr_i(m_{ijk}|jk, \boldsymbol{\Delta_i}) = f\left[\log(m_{ijk}) - \log(\tilde{m}_{ijk})\right], \tag{1.11}$$

where $f\left[\cdot\right]$ is the probability density function for the error μ_{ijk}. As in (1.8), I explicitly write the left hand side conditional on $\boldsymbol{\Delta_i}$ in order to explain the method of maximum simulated likelihood below.

Then, the probability that household i purchases model jk and we observe the vehicle driven for $m_{ijk}(\times 100)$ miles is

$$\Pr_i(jk, m_{ijk}|\boldsymbol{\Delta_i}) = \Pr_i(jk|\boldsymbol{\Delta_i}) \times \Pr_i(m_{ijk}|jk, \boldsymbol{\Delta_i}), \tag{1.12}$$

where $\Pr_i(jk|\boldsymbol{\Delta_i})$ and $\Pr_i(m_{ijk}|jk, \boldsymbol{\Delta_i})$ are as defined in (1.8) and (1.11), respectively.

One change is required in the indirect utility function (1.2)-(1.6) to make estimation possible. In (1.3), w_{ijk} is household i's budget after purchasing model jk. That is, $w_{ijk} = y_i - p_{jk}$, where y_i is household i's pre-purchase budget and p_{jk} is model jk's price. The problem is that y_i is unobservable in the data set. However, plugging $w_{ijk} = y_i - p_{jk}$ into (1.3), and noting that any term not varying over alternatives (jk) does not affect the probability in (1.8), we can drop the term $\alpha_{1i}y_i$, which involves unobservable pre-purchase budget, from (1.3). Making this change, I re-define the indirect utility function as follows, with a slight abuse of notation:

$$V_{ijk} = \tilde{V}_{ijk} + \epsilon_{ijk}, \tag{1.13}$$

$$\tilde{V}_{ijk} = -\alpha_{1i} p_{jk} - \alpha_{2ijk} r_{ijk}^{\alpha_3} + \sum_h \eta_{hi} d_{hjk} + \xi_{jk}, \tag{1.14}$$

$$\alpha_{1i} = \exp(\boldsymbol{\beta}_1 \boldsymbol{x}_i + \delta_{1i}), \tag{1.15}$$

$$\alpha_{2ijk} = \exp(\boldsymbol{\beta}_2 \boldsymbol{x}_i + \boldsymbol{\gamma}_2 \boldsymbol{z}_{jk} + \delta_{2i}), \tag{1.16}$$

$$0 < \alpha_3 < 1. \tag{1.17}$$

The only change from (1.2)-(1.6) is that (1.3) is replaced by (1.14). Demand estimation that follows will be based on (1.13)-(1.17) for the part of discrete vehicle choice. That is, vehicle choice probabilities defined in (1.8) and used in (1.12) are calculated by (1.14)-(1.17). Since w_{ijk} does not appear in the VMT demand equation (1.9), estimation can be based on (1.10) for the part of continuous VMT choice.

Parameters and variables that appear in both (1.13)-(1.17) and (1.10) affect both vehicle and VMT choices in a way that is consistent with Roy's identity. In other words, they introduce correlation between the two choices. For example, if there is a household characteristic (e.g., large household size, or commuting by car) that, other things equal, tends to increase the household's driving demand, then conditional on purchasing a particular model, households with this characteristic likely show higher VMT than households that are without the characteristic, but are otherwise equal. In addition, since higher driving demand means larger benefits from fuel efficiency, households with this characteristic tend to purchase more fuel efficient vehicles. Therefore, there is correlation between vehicle and VMT choices that households with high driving demand likely buy a fuel efficient vehicle and drive it long distance. This kind of correlation between the two choices is taken care of in the economic model by parameters and variables appearing in both (1.13)-(1.17) and (1.10). Among them, $\boldsymbol{x}_i$ and $\boldsymbol{z}_{jk}$ are, respectively, nrobserved household characteristics and vehicle attributes that affect both choices. In the following estimations, $\boldsymbol{x}_i$ includes i's household size, population density at i's location, the number of vehicles owned by i, and i's annual income, while $\boldsymbol{z}_{jk}$ includes vehicle footprint, acceleration capacity, and vehicle class dummies. On the other hand, δ_{1i} and δ_{1i} are unobserved (to the econometrician) household tastes affecting both choices. For example, δ_{1i} and δ_{1i} may include whether or not household i plans to use the vehicle for commuting.

Terms ξ_{jk}, ϵ_{ijk} and $\sum_h \eta_{hi} d_{hjk}$ all appear in (1.13)-(1.14), but not in the VMT equation (1.10). That is, they are assumed to affect discrete vehicle choice, but not continuous VMT choice. Fixed effect ξ_{jk} represents the portion of utility of jk that does not vary across households. As explained below, it is set to equate jk's share predicted by the econometric model to its observed market share.

The sum of unobserved terms $\epsilon_{ijk} + \sum_h \eta_{hi} d_{hjk}$ captures the random (unobserved) portion of utility from jk that varies over i, and results from factors not explained in the econometric model (Train, 2009). As explained above, ϵ_{ijk} is the random portion of utility that varies over both i and jk. Random parameter η_{hi}, which varies over i but not over jk, represents i's unobserved preference on vehicle type h that is common to all models in type h. Thus, $\sum_h \eta_{hi} d_{hjk}$ induces correlation among different models in each type. That is, the household's unobserved heterogeneous preference ($\epsilon_{ijk} + \sum_h \eta_{hi} d_{hjk}$) is correlated among two models of the same vehicle type, but uncorrelated among two alternatives in different categories. With these terms, this econometric model is analogous to a nested logit model. The magnitude of the correlation within type h is controlled by the variance of the distribution of η_{hi}.

The error term in (1.10), μ_{ijk}, consists of unobserved (to the econometrician) factors that affect reported VMT, but not the vehicle choice. For example, μ_{ijk} may include various factors affecting realized VMT but not considered by the consumer at the time of vehicle purchase, such as unpredicted changes in the household's driving demand that occur after the purchase.

1.3.2 Automobiles as Durable Goods

In practice, automobiles are a durable good and consumers most likely plan to use them over years. In the utility maximization problem (1.1) and the following argument so far, I have considered only one period in order to make the model econometrically tractable. However, I can incorporate the multi-period nature of vehicle ownership into this modeling framework as follows. First, $\tilde{q}$ in (1.1) is interpreted as the present value sum of spending on the numéraire good (the composite of all goods/services other than VMT) over multiple years, and w_{ijk} is interpreted similarly. If we consider multiple periods, the present value sum of fuel spending is calculated as

$$\sum_{t=1}^{T} d_f^{(t-1)} r_{ijkt} \tilde{m}_t, \tag{1.18}$$

where T is the length (in years) of the household's planning horizon, r_{ijkt} and $\tilde{m}_t$ are, respectively, the anticipated price of driving and planned VMT in year t, and d_f is a discount factor for future fuel spending, which may be lower than one for other spending or income, so that future fuel spending is discounted faster than other spending or income (the energy paradox).

Assuming that the household expects the price of driving to change over years based on $r_{ijkt} = d_r^{(t-1)} r_{ijk1}$ (due to fuel price changes or fuel efficiency deterioration over time), and also that VMT changes based on $\tilde{m}_t = d_m^{(t-1)} \tilde{m}_1$, we can write the present value sum of fuel spending anticipated at the time of vehicle purchase as

$$\sum_{t=1}^{T} d_f^{(t-1)} r_{ijkt} \tilde{m}_t = r_{ijk1} \frac{1-(d_f d_r d_m)^T}{1-d_f d_r d_m} \tilde{m}_1. \tag{1.19}$$

Therefore, we may take account of the multi-period nature of vehicle use by interpreting r_{ijk} and $\tilde{m}$ in (1.1) as

$$r_{ijk} = r_{ijk1}, \tag{1.20}$$

$$\tilde{m} = \frac{1-(d_f d_r d_m)^T}{1-d_f d_r d_m} \tilde{m}_1. \tag{1.21}$$

The 2001 NHTS includes an estimate of realized annual VMT of each surveyed vehicle. For a model year 2001 vehicle, we can interpret that the survey shows an estimate of the first year's actual VMT, which is the sum of the first year's planned VMT, $\tilde{m}_1$ in (1.21), and an error. Therefore, to account for durability of vehicle use, we need a constant equal to the average household's $(1 - (d_f d_r d_m)^T)/(1 -$

$d_f d_r d_m)(\equiv K)$ that scales up the first year's observed VMT.[18][19] Unfortunately, this K is unobservable, so I instead choose a few reasonable values (e.g., 1, 4 and 7) for K and report results for each case. The larger K, the more valuable fuel economy is for the household because improved fuel economy is expected to provide greater savings in household fuel spending.

1.3.3 Estimation Procedure

In the National Household Travel Survey, vehicles are identified only up to the nameplate (i.e., j) level, so the model (i.e., k) level information is unavailable. That is, in the data, we observe only the nameplate purchased by household i and its VMT (denoted respectively by j_i and m_{ij_i}).

With this structure of the data set as well as (1.8), (1.11) and (1.12) considered, conditional on household i's unobserved heterogeneity, the probability that we observe household i purchase any model under nameplate j_i and drive it for $m_{ij_i}(\times 100)$ miles is

$$\Pr_i(j_i, m_{ij_i}|\mathbf{\Delta}_i) = \sum_k \Big\{ \Pr_i(j_i k|\mathbf{\Delta}_i) \times \Pr_i(m_{ij_i}|j_i k, \mathbf{\Delta}_i) \Big\}, \tag{1.22}$$

where $\Pr_i(j_i k|\mathbf{\Delta}_i)$ and $\Pr_i(m_{ij_i k}|j_i k_i, \mathbf{\Delta}_i)$ are defined by (1.8) and (1.11), respectively.

Integrating (1.22) over $\mathbf{\Delta}_i$ gives the unconditional (on $\mathbf{\Delta}_i$) probability:

$$\Pr_i(j_i, m_{ij_i}) = \int \Pr_i(j_i, m_{ij_i}|\mathbf{\Delta}_i) \mathrm{d}\Phi(\mathbf{\Delta}_i), \tag{1.23}$$

[18] As an illustration, if $T = 2$, $d_f = 0.8$, $d_r = 1$ and $d_m = 0.95$, then $K = 1.76$. If $T = 8$, $d_f = 0.95$, $d_r = 1$ and $d_m = 0.98$, then $K = 6.313$.

[19] The constant to scale up the first year's VMT is likely to be different across households. This heterogeneity is absorbed into household-specific random parameters, so we can assume that K is common to all households.

where $\Phi(\cdot)$ is the joint distribution function of $\boldsymbol{\Delta}_i$. The above integral does not have a closed form solution, so needs to be approximated by simulation:

$$\widetilde{\Pr}_i(j_i, m_{ij_i}) = \frac{1}{A}\sum_{a=1}^{A} \Pr_i(j_i, m_{ij_i}|\boldsymbol{\Delta}_i^a), \tag{1.24}$$

where $\boldsymbol{\Delta}_i^a$ is the ath set of simulated independent Halton draws from $\Phi(\cdot)$, and A is the number of sets ($A = 200$ in the following estimation). Since choice probability clearly depends on parameters to be estimated, in the following I will denote the probability on the left hand side of (1.24) by $\widetilde{\Pr}(j_i, m_{ij_i}|\boldsymbol{\xi}, \boldsymbol{\theta})$ to make this dependence explicit, where $\boldsymbol{\xi}$ is a vector of ξ_{jk}'s, and $\boldsymbol{\theta}$ is a vector of other parameters to be estimated.[20]

I estimate parameters by the maximum simulated likelihood (MSL) estimator.[21] The MSL estimator maximizes the following simulated likelihood with respect to parameters.

$$\ln \widetilde{\mathrm{L}}(\boldsymbol{\xi}, \boldsymbol{\theta}) = \sum_i W_i \ln \widetilde{\Pr}(j_i, m_{ij_i}|\boldsymbol{\xi}, \boldsymbol{\theta}), \tag{1.25}$$

where W_i is the sample weight of household i, calculated from the NHTS data set.

Vehicle fixed effects $\boldsymbol{\xi}$ are estimated with the contraction mapping method used in Berry et al. (1995). With this method, ξ_{jk} is set to equate the predicted market share of jk from the econometric model with the observed market share of jk. jk's predicted market share is calculated with each household's probability of choosing

[20] $\boldsymbol{\theta} = [\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\gamma}_2, \alpha_3, \bar{\delta}_1, \sigma_{\delta_1}, \bar{\delta}_2, \sigma_{\delta_2}, \sigma_{\eta_1}, \sigma_{\eta_2}, \sigma_{\eta_3}, \sigma_{\eta_4}]$, where $\mathcal{N}(\bar{\delta}_g, \sigma_{\delta_g})$ is the distribution of δ_{gi} ($g = 1, 2$), and $\mathcal{N}(0, \sigma_{\eta_h})$ is the distribution of η_{hi} ($h = 1, 2, 3, 4$).

[21] Bento et al. (2009) uses a Markov Chain Monte Carlo framework rather than maximum likelihood.

jk and sample weight. Note that though each household's choice information at the vehicle model (k) level is unavailable in the NHTS as mentioned above, the aggregate sales information in the U.S. market can be obtained at this level, making it possible to set ξ_{jk} for each model.

Specifically, given parameter values $\check{\boldsymbol{\theta}}$ obtained at each iteration of the numerical maximization of (1.25) and some initial vector $\boldsymbol{\xi}^0$, I use the following contraction mapping to find $\boldsymbol{\xi}(\check{\boldsymbol{\theta}}, \mathbf{s})$ that equates the predicted and observed market shares for all jk's.

$$\xi^t_{jk}(\check{\boldsymbol{\theta}}, \mathbf{s}) = \xi^{t-1}_{jk}(\check{\boldsymbol{\theta}}, \mathbf{s}) + \ln s_{jk} - \ln \hat{s}_{jk}\big(\boldsymbol{\xi}^{t-1}(\check{\boldsymbol{\theta}}, \mathbf{s}), \check{\boldsymbol{\theta}}\big), \tag{1.26}$$

where s_{jk} (an element of $\mathbf{s}$) is jk's observed market share, and $\hat{s}_{jk}$ is jk's predicted share from the econometric model. With $\boldsymbol{\xi}(\check{\boldsymbol{\theta}}, \mathbf{s})$ now at hand, we update $\check{\boldsymbol{\theta}}$ based on the numerical maximization algorithm, and then repeat the above process until obtaining MSL estimates $\hat{\boldsymbol{\theta}}$. MSL estimate $\hat{\xi}_{jk}$ is obtained by using the contraction mapping at $\hat{\boldsymbol{\theta}}$.

1.3.4 Estimation Results

Table 1.3 summarizes the results of maximum simulated likelihood estimation described above. Point estimates and standard errors of selected parameters are reported. Each column corresponds to a specific value of K, the multiplier for annual VMT, which is shown in the last row. Household characteristics included are the logarithm of the following variables: household size, population density, the number of vehicles owned and household annual income. As for annual income, I fit a linear spline with two knots ($\text{knot}_1 = \log(46,715)$ and $\text{knot}_2 = \log(73,580)$).[22] Vehicle attributes included are dummy variables for vans, SUVs and pickup trucks (passenger cars as the base type) and the logarithm of the following variables: footprint (the area bounded by the four wheels) and horsepower (hp) divided by weight (lb), which is a measure of acceleration capacity.

First, I explain the implications of the estimates in relation to the discrete vehicle choice. I interpret estimates of $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ here, and results on $\boldsymbol{\gamma}_2$ and α_3 will be discussed in analyzing the estimated VMT demand equation (1.10) below. Note that with the utility function in consideration, α_{1i} is the marginal utility of money, and each element of $\boldsymbol{\beta}_1$ is the elasticity of α_{1i} with respect to the corresponding household characteristic. A negative element of $\boldsymbol{\beta}_1$ means that the marginal utility of money decreases with the corresponding characteristic, so that a household with a larger value of the characteristic is less sensitive to vehicle prices. A positive element of $\boldsymbol{\beta}_1$ means the opposite. Notably, the elasticity of the marginal utility of money with respect to annual household income is estimated -0.31 to -0.35 if annual household income is less than \$46,715; -0.66 to -0.71 if it is between \$46,715 and \$73,580; -2.05 to -2.59 if it is greater than \$73,580.[23] The marginal utility of money decreases with income, and does so faster (in terms of elasticity) for households with higher income.

A positive element of $\boldsymbol{\beta}_2$ means that the effect of the price of driving (r_{ijk}) on

[22] Roughly, a third of the households in my data set have annual income less than \$46,715, another third between \$46,715 and \$73,580, and the other third above \$73,580.

[23] The elasticity for households with income between \$46,715 and \$73,580 is given by summing the coefficients of log(income) and max[log(income/46,715),0]. Similarly, the elasticity for households with income between greater than \$73,580 is given by summing the coefficients of log(income), max[log(income/46,715),0] and max[log(income/73,580),0].

utility from model jk increases with the corresponding household characteristic, so that a household with a larger value of the characteristic is more sensitive to r_{ijk}. A negative element means the opposite. The estimates suggest that, other things equal, households of larger size or in a more populated area are more sensitive to r_{ijk}, and households owning more vehicles or earning higher income are less sensitive to it.

Next, let us focus on the VMT demand equation (1.10). $\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1$ is the coefficients of household characteristics $\boldsymbol{x}_i$. Table 1.4 reports point estimates and standard errors of $\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1$ for the three cases of Table 1.3. Other things equal, households of larger size, in a less populated area, with less vehicles owned, or with higher income plan to drive the vehicle longer distance. Since coefficients of max[log(income/46,715),0] and max[log(income/73,580),0] are not significant, the elasticity of VMT with respect to household income does not differ significantly across different income groups.

In Table 1.4, results on γ_2 suggest that consumers expect to drive SUVs and pickup trucks (in comparison to passenger cars) or relatively large (i.e., large footprint) vehicles (within each vehicle type) for longer distance. In terms of Table 1.3 (that is, the indirect utility function (1.2)-(1.6)), this means that if model jk is an SUV or a pickup truck, and/or it is relatively large within its type, then household i's indirect utility conditional on buying model jk (V_{ijk}) is more sensitive to the price of driving (r_{ijk}), as it is expected to be used more intensively.

In equation (1.10), $\alpha_3 - 1$ is the elasticity of VMT with respect to the price of driving (r_{ijk}) conditional on vehicle choice.[24] In Table 1.4, the point estimate of $\alpha_3 - 1$ ranges from -0.17 to -0.27.[25] This range is overall consistent with previous studies. Bento et al. (2009) find the long-run VMT elasticities for different classes of new vehicles to be around -0.2 to -0.3. Reviewing the literature, Small and Van Dender (2007) also regard the range of -0.2 to -0.3 as compatible with previous studies.

We now look at the overall performance of the econometric model. First, using parameter estimates of cases [1]-[3] of Table 1.3, Table 1.5 shows predicted choice probabilities of two selected vehicle models (the economical Toyota Corolla 1.8 liter engine model and the luxurious Cadillac DeVille 4.6 liter engine model) for three hypothetical households. The three households' annual income is set at the 5th, 50th or 95th percentile of the sample, and other household characteristics are set at the sample average. Predictions in Table 1.5 confirms a reasonable tendency that, relatively speaking, lower income households prefer smaller, cheaper vehicles, while wealthier households prefer larger, more luxurious vehicles.

With the estimates in Table 1.3, we can calculate each vehicle model's elasticity of sales with respect to its own price and fuel efficiency (e.g., gallons per 100 miles) by aggregating the change in each household's choice probability of the vehicle model. Based on parameter estimates of cases [1]-[3] of Table 1.3, Table 1.6 reports the mean and standard deviation of elasticities of sales thus calculated. The average elasticity with respect to vehicle price is estimated around -2 to -2.5, which is similar in magnitude to findings from other studies that use micro cross-sectional data (e.g.,

[24] As in Bento et al. (2009), my estimation is based on cross-sectional data, so elasticity estimates should be interpreted as long-run elasticities.

[25] Equivalently, the point estimate of α_3 ranges from 0.73 to 0.83, which is consistent with assumption (1.6) that requires $0 < \alpha_3 < 1$ to ensure the concavity of the utility function (1.7).

Berry et al., 2004; Bento et al., 2009; Train and Winston, 2007).

The average elasticity with respect to fuel efficiency is around -0.1 to -0.5 and changes with K. The larger K, the more valuable consumers view fuel economy, hence the larger (in magnitude) the impact of a fuel efficiency change on vehicle sales. We also notice that, for all of the three cases, the fuel efficiency elasticity is much smaller (in magnitude) than the price elasticity.

1.4 Supply Side of the New Vehicle Market

This section considers the supply side of the U.S. new vehicle market in model year 2001. I first discuss the CAFE standards at that time, then consider an economic model to analyze each automaker's behavior in the market.

1.4.1 CAFE Standards for Model Year 2001

Under the CAFE standards for model year 2001, an automaker's sales-weighted harmonic average miles per gallon (mpg) in each of the three vehicle categories (domestic passenger cars, import passenger cars, and light-duty trucks) needs to exceed the corresponding standard (27.5 mpg for domestic and import passenger cars, and 20.5 for light-duty trucks).[26] If the automaker does not meet the standard in a category, it must pay a fine of $5 for each 0.1 mile per gallon below the standard, for each vehicle sold in the category.

Table 1.7 shows the sales-weighted harmonic average mpg by automaker and

[26]Equivalently, the sales-weighted arithmetic average gphm (gallons per 100 miles) in each category needs to be below $100/27.5 \approx 3.64$ gphm for domestic and import passenger cars, and $100/20.7$ ≈ 4.83 gphm for light-duty trucks.

category for model years 1999-2003.[27] A red entry implies that the average is under the target for the automaker-category. Since the CAFE standards allowed banking and borrowing of credits within each automaker-category, up to three years, falling short of the standard in one year does not immediately mean that the automaker is in violation and needs to pay fines.

Table 1.7 suggests three types of responses from automakers to the CAFE standards. First, the automaker's category average mpg is consistently above the standard (e.g., Toyota in all categories). Second, the average mpg is consistently below the standard over time (BMW, DaimlerChrysler and Porsche's import passenger cars). In this case, automakers choose to pay fines rather than meeting the standard. In the third situation, the average mpg is above the standard for some years and below for other years, but the automaker uses banking and borrowing to meet the target and avoid fines (e.g., Ford's light-duty trucks).

1.4.2 A Model of Automakers' Profit Maximization

I set up an economic model of automakers in an imperfectly competitive market and under CAFE standards. It is based on the models used in previous studies to analyze the supply side of the vehicle market and related issues such as fuel economy regulations and gasoline taxes (e.g., Berry et al., 1995; Bento et al., 2009; Jacobsen, forthcoming; Coleman and Harrington, 2010; Klier and Linn, 2012; Whitefoot et al., 2012).

[27] For automakers that sell flex-fuel vehicles, the corporate average values for domestic passenger cars and light-duty trucks in Table 1.7 include bonus mpg credits from selling flex-fuel vehicles.

I consider a Nash equilibrium in which automaker n maximizes its profits by optimally setting the price and fuel efficiency of the vehicle models it produces, given other vehicle attributes of n's models, and prices and all attributes of other automakers' models. In the maximization problem, automakers also take account of CAFE standards of model year 2001.

The goal of this section is to derive an estimate of the marginal production cost of improving fuel efficiency for each vehicle model. This is done by exploiting the first order conditions of the maximization problem and using the results from the demand side. The marginal cost estimates thus obtained will be used for simulating counterfactual CAFE standards in Section 1.5.

In considering automakers' maximization under fuel economy regulations, it is important to distinguish their responses to the CAFE standards because different responses lead to different formulations of the maximization problem (Jacobsen, forthcoming; Whitefoot et al., 2012). As discussed above, there are three possibilities regarding how automakers respond to the CAFE standards: unconstrained, violating and constrained. Below I model an automaker's profit maximization in each case. For simplicity, I do not explicitly consider banking and borrowing of credits over years, so the following problems are for a single period.

1.4.2.1 Automakers Unconstrained by CAFE Standards

An automaker is unconstrained in a category if the (sales-weighted harmonic) average mpg of the category consistently exceeds the corresponding target. Suppose automaker n is not constrained by CAFE standards in any category. Profit maximization can be formulated as

$$\max_{\{p_{jk}, e_{jk}\}_{jk \in \mathcal{J}_n}} \sum_{jk \in \mathcal{J}_n} (p_{jk} - c_{jk}) s_{jk}, \tag{1.27}$$

where $\mathcal{J}_n$ is the set of vehicle models produced by automaker n; c_{jk} is the unit production cost of model jk, s_{jk} is the sales volume of jk.

The first order conditions with respect to the price of vehicles produced by automaker n are:

$$\forall jk \in \mathcal{J}_n,$$

$$s_{jk} + \sum_{j'k' \in \mathcal{J}_n} (p_{j'k'} - c_{j'k'}) \frac{\partial s_{j'k'}}{\partial p_{jk}} = 0. \tag{1.28}$$

Note that we have data on market sales (s_{jk}) and vehicle price (p_{jk}) of each model. From the demand side model estimated before,

$$s_{jk} = \sum_i W_i \widetilde{\Pr}_i(jk). \tag{1.29}$$

In (1.29), W_i is the sample weight of household i in the NHTS data set, and $\widetilde{\Pr}_i(jk)$ is the simulated probability of household i choosing model jk:

$$\widetilde{\Pr}_i(jk) = \frac{1}{A} \sum_{a=1}^{A} \Pr_i(jk | \Delta_i^a), \tag{1.30}$$

where $\Pr_i(jk|\Delta_i^a)$ is as given in (1.8). So, $\partial s_{j'k'}/\partial p_{jk}$ is calculated as

$$\frac{\partial s_{j'k'}}{\partial p_{jk}} = \begin{cases} \sum_i W_i \widetilde{\Pr}_i(jk)\{1 - \widetilde{\Pr}_i(jk)\}\frac{\partial V_{ijk}}{\partial p_{jk}}, & \text{if } j'k' = jk; \\ -\sum_i W_i \widetilde{\Pr}_i(jk)\widetilde{\Pr}_i(j'k')\frac{\partial V_{ijk}}{\partial p_{jk}}, & \text{if } jk \neq j'k'. \end{cases} \tag{1.31}$$

Thus, the system of equations (1.28) can be solved for c_{jk} $\forall jk \in \mathcal{J}_n$.

Similarly, the first order conditions with respect to fuel consumption of vehicle models produced by automaker n are:

$$\forall jk \in \mathcal{J}_n,$$
$$-\frac{\partial c_{jk}}{\partial e_{jk}} s_{jk} + \sum_{j'k' \in \mathcal{J}_n} (p_{j'k'} - c_{j'k'}) \frac{\partial s_{j'k'}}{\partial e_{jk}} = 0. \tag{1.32}$$

We have data on market sales s_{jk} and vehicle price p_{jk} of each model, and c_{jk} is obtained as described above for all jk produced by automaker n. $\partial s_{j'k'}/\partial e_{jk}$ is estimated from the demand side model as

$$\frac{\partial s_{j'k'}}{\partial e_{jk}} = \begin{cases} \sum_i W_i \widetilde{\Pr}_i(jk)\{1 - \widetilde{\Pr}_i(jk)\}\frac{\partial V_{ijk}}{\partial e_{jk}}, & \text{if } j'k' = jk; \\ -\sum_i W_i \widetilde{\Pr}_i(jk)\widetilde{\Pr}_i(j'k')\frac{\partial V_{ijk}}{\partial e_{jk}}, & \text{if } jk \neq j'k'. \end{cases} \tag{1.33}$$

Thus, the system of equations (1.32) can be solved for $-\partial c_{jk}/\partial e_{jk}$, the marginal cost of fuel efficiency improvement, for all $jk \in \mathcal{J}_n$.

1.4.2.2 Automakers Violating CAFE Standards

An automaker is violating the CAFE standards in a category if the average mpg of the category is consistently below the corresponding target and the automaker is paying a fine of $5 for each 0.1 mile per gallon below the standard, for each vehicle sold in the category. Violating automakers choose not to meet the standard, but instead choose to pay fines.

Suppose automaker n is violating the standards and paying the fines in category g (but unconstrained in other categories). The automaker's total fine payment in this case is

$$50 \times \Big(\frac{100}{\bar{e}^g} - \frac{100}{(\sum_{jk \in \mathcal{J}_n^g} e_{jk} s_{jk})/(\sum_{jk \in \mathcal{J}_n^g} s_{jk})} \Big) \times \Big(\sum_{jk \in \mathcal{J}_n^g} s_{jk} \Big), \tag{1.34}$$

where $\bar{e}^g$ is the standard for category g (in gphm) (for example, about 4.83 for the light-duty truck category) and $\mathcal{J}_n^g$ is the set of vehicle models produced by automaker n and belonging to category g. In (1.34), $100/\bar{e}^g$ is the category g target in terms of mpg, and $100/\{(\sum_{jk \in \mathcal{J}_n^g} e_{jk} s_{jk})/(\sum_{jk \in \mathcal{J}_n^g} s_{jk})\}$ is the sales-weighted harmonic average fuel economy (in mpg) of models by automaker n and in category g.

So profit maximization can be formulated as

$$\max \sum_{jk \in \mathcal{J}_n} (p_{jk} - c_{jk}) s_{jk} - \Big\{ 50 \times \Big(\frac{100}{\bar{e}^g} - \frac{100}{(\sum_{jk \in \mathcal{J}_n^g} e_{jk} s_{jk})/(\sum_{jk \in \mathcal{J}_n^g} s_{jk})} \Big) \times \Big(\sum_{jk \in \mathcal{J}_n^g} s_{jk} \Big) \Big\}. \tag{1.35}$$

In this case, the first order conditions becomes more complicated as we need to consider the CAFE fine term. But we can still obtain c_{jk} and $-\partial c_{jk}/\partial e_{jk}$ in the same way as above.

1.4.2.3 Automakers Constrained by CAFE Standards

An automaker is constrained by the CAFE standards in a category if the average mpg of the category marginally exceeds the corresponding target.[28] In this case, the automaker chooses vehicle prices and attributes to marginally satisfy the target.

Suppose automaker n is constrained in category g (but unconstrained in other categories). The profit maximization problem can be written as

$$\max_{\{p_{jk},e_{jk}\}_{jk\in\mathcal{J}_n}} \sum_{jk\in\mathcal{J}_n} (p_{jk}-c_{jk})s_{jk} \quad \text{s.t.} \quad \sum_{jk\in\mathcal{J}_n^g} (\bar{e}^{gn}-e_{jk})s_{jk} \geq 0, \tag{1.36}$$

where $\bar{e}^{gn}$ is the sales-weighted average gphm of automaker n's category g fleet. Dividing the constraint by $\sum_{jk\in\mathcal{J}_n^g} s_{jk}$ and rearranging terms, we can see that it is equivalent to requiring the sales-weighted arithmetic average gphm (or sales-weighted harmonic average mpg) of automaker n's category g fleet to be at least as fuel efficient as the value it actually achieved in the market.[29] The constraint is written as in (1.36) in order to be consistent with the notation in Section 1.5 below, which explains more about how to interpret the CAFE constraint.

[28] Since the actual CAFE regulations allow banking and borrowing of credits over time, it may be the case that the category average mpg is well above or below the target for one year, and the automaker is still constrained in the category.

[29] With banking and borrowing allowed in the actual CAFE standards, the constrained fleet's observed average fuel efficiency in each year does not generally equal the target value. Because I am not considering banking and borrowing in modeling profit maximization for the sake of simplicity, I assume that the automaker is required to set the constrained fleet's average fuel efficiency at least at the level actually observed in the market. This makes $\bar{e}^{gn}$ different across n.

The first order conditions with respect to the price of vehicles produced by automaker n are:

$$\forall jk \in \mathcal{J}_n, \quad s_{jk} + \sum_{j'k' \in \mathcal{J}_n} \{p_{j'k'} - c_{j'k'} + I^g_{jk}\lambda^g_n(\bar{e}^{gn} - e_{jk})\} = 0, \tag{1.37}$$

where I^g_{jk} takes 1 if jk is in category g vehicles and 0 otherwise, and λ^g_n is the Lagrange multipliers for the constraints. Similarly, the first order conditions with respect to gphm for the constrained case are:

$$\forall jk \in \mathcal{J}_n, \quad \{-\frac{\partial c_{jk}}{\partial e_{jk}} - I_{jk}\lambda^g_n\}s_{jk} + \sum_{j'k' \in \mathcal{J}_n} \{p_{j'k'} - c_{j'k'} + I^g_{jk}\lambda^g_n(\bar{e}^{gn} - e_{jk})\} = 0. \tag{1.38}$$

Once λ^g_n is given, we can solve the above system of first order conditions (1.37) and (1.38) for c_{jk} and $-\partial c_{jk}/\partial e_{jk}$ in the same way as in the unconstrained case. Estimates of λ^g_n are constructed from the results of Anderson and Sallee (2011) who estimated the shadow cost of CAFE constraints.

1.4.2.4 More Details

In practice, automakers may be in violation and/or constrained in more than one category simultaneously. The profit maximization problem for these cases can be easily formulated as a combination of the above cases.

Investigating the time trend of automakers' corporate average fuel economy and following previous studies (Jacobsen, forthcoming; Anderson and Sallee, 2011), I assume that categories in violation are import passenger cars of DaimlerChrysler, BMW and Porsche, and light-duty trucks of BMW. Likewise, constrained categories are domestic passenger cars and light-duty trucks of General Motors, Ford and DaimlerChrysler. All other categories are assumed unconstrained.

Previous studies that endogenize the choice of fuel efficiency (e.g., Bento et al., 2009; Coleman and Harrington, 2010; Klier and Linn, 2012) use the same approach as above to infer the marginal cost of fuel efficiency improvement for each vehicle model. In these studies, like mine, the inferred marginal costs are essential for counterfactual simulations. That is, when an automaker changes fuel efficiency of its vehicles to maximize profits in a counterfactual simulation, model jk's inferred marginal cost is used to calculate the change in its production cost. More specifically, based on engineering studies (as summarized in National Research Council, 2002) and Coleman and Harrington (2010), the change in model jk's production cost due to a fuel efficiency adjustment from e^0_{jk} to e^1_{jk} is given by:

$$T_{jk}(e_{jk}) = -\frac{\partial c_{jk}}{\partial e_{jk}}\Big|_{e_{jk}=e^0_{jk}} \left\{ \left(\frac{e^0_{jk}}{e^1_{jk}} \right)^2 - 1 \right\}, \tag{1.39}$$

where e^0_{jk} is the actual gphm observed in the market, and $-\partial c_{jk}/\partial e_{jk}\big|_{e_{jk}=e^0_{jk}}$ is the marginal cost obtained above from a system of first order conditions. Thus, $T_{jk}(e_{jk})$ expresses the change in model jk's technology cost due to adjusting fuel efficiency from the actual value observed in the market. This model fits the engineering estimates of incremental technology costs of fuel efficiency improvement very well (Coleman and Harrington, 2010).

1.5 Simulating the Impacts of Counterfactual CAFE Standards

Based on the above analysis of the demand and supply sides, I compare the effectiveness of different features of CAFE standards by simulating a new vehicle market under counterfactual regulations. Specifically, I evaluate the impacts of three policy instruments recently introduced into CAFE standards:

- Instrument A: "footprint-based" functions that assign target values based on vehicle size
- Instrument B: intra-firm transferring of fuel efficiency credits across vehicle categories
- Instrument C: inter-firm trading of fuel efficiency credits

Combining these instruments, I construct four counterfactual policies to be simulated. These policies are summarized as follows.

- Policy 1 mimics the actual CAFE standards for model year 2001
- Policy 2 replaces Policy 1's "flat" functions with Instrument A
- Policy 3 adds Instrument B to Policy 2
- Policy 4 adds Instrument C to Policy 3, mimicking the design of the actual CAFE standards for model years 2012-2016

The next subsection explains the four policies in detail.

The overall framework of the simulation is as follows. Given a counterfactual policy, I find a Nash equilibrium in which each automaker sets vehicle prices and fuel efficiency values of its vehicles to maximize profits under the regulatory constraint. Consumers' demand for different vehicle models changes according to the demand side model estimated above. I obtain various measures (such as fuel efficiency, profits, fuel savings) at the equilibrium, and then compare the four policies in terms of these measures to evaluate the impacts of the three policy instruments.

1.5.1 Counterfactual CAFE Standards

This subsection discusses four counterfactual CAFE standards to be simulated, and how these standards are set.

First, we express the CAFE constraint for category g in a more general form than in (1.36):

$$\sum_{jk \in \mathcal{J}_n^g} (\bar{e}_{jk} - e_{jk}) s_{jk} \geq 0, \tag{1.40}$$

where summation is over all vehicle models that are produced by automaker n and belong to category g (denoted by $\mathcal{J}_n^g$); $\bar{e}_{jk}$ and e_{jk} are target and achieved gphm of model jk, respectively; s_{jk} is U.S. market sales of jk. Note that (1.40) is more general than (1.36) because in (1.40), target value $\bar{e}_{jk}$ may differ across jk, while in (1.36) it is constant for all jk in automaker n's category g fleet.

The constraint can be explained as follows. If model jk of category g is more fuel efficient than its target value (i.e., $\bar{e}_{jk} > e_{jk}$), then for each unit of model jk sold, automaker n earns a positive entry ("credit") of $\bar{e}_{jk} - e_{jk}$ in category g. On the other hand, if model jk is less fuel efficient than its target value (i.e., $\bar{e}_{jk} < e_{jk}$), then for each unit of model jk sold, it earns a negative entry ("debit") of $\bar{e}_{jk} - e_{jk}$ in category g. Automaker n satisfies the standard for category g if the total balance is non-negative after summing all entries in category g.

To focus on primary aspects of the regulations (namely, Instruments A-C above), counterfactual CAFE standards considered below ignore or simplify some features of actual CAFE standards. First, I do not distinguish domestic passenger cars and import passenger cars, so all passenger cars are grouped into a single category. Likewise, I do not consider banking and borrowing of credits over different model years, and bonus credits for selling flex-fuel vehicles.

1.5.1.1 Policy 1: Flat Targets / No Transferring / No Trading

Policy 1 mimics the actual CAFE standards for model year 2001 except for the points explained above. Automakers face two CAFE constraints, one for passenger cars (PCs) and the other for light-duty trucks (LDTs). The constraints are given by letting $\bar{e}_{jk} = 100/27.5 \approx 3.64 \ \forall jk$ for passenger cars, and $\bar{e}_{jk} = 100/20.7 \approx 4.83 \ \forall jk$ for light-duty trucks. Note that within each category, the target value is common to all vehicles, regardless of vehicle size (”flat standards”). Figure 1.1 plots the target values along with actual gphm of all models in model year 2001. Credits earned in one category may not be used to offset in the other category (”no transferring” of credits), and credits may not be sold to or bought from other automakers (”no trading” of credits). If an automaker does not comply with the standard in a category, it must pay a fine of \$5 for each 0.1 mpg short of the standard, for each vehicle sold in the category.

In a simulation under Policy 1, automaker n that chooses to comply with the CAFE standards solve the following profit maximization problem:

$$\max_{\{p_{jk},e_{jk}\}_{jk\in\mathcal{J}_n}} \sum_{jk\in\mathcal{J}_n}(p_{jk}-c_{jk})s_{jk} \quad \text{s.t.} \quad \begin{aligned} &\sum_{jk\in\mathcal{J}_n^c}(100/27.5-e_{jk})s_{jk}\geq 0, \\ &\sum_{jk\in\mathcal{J}_n^t}(100/20.7-e_{jk})s_{jk}\geq 0. \end{aligned} \tag{1.41}$$

There are two CAFE constraints, one for passenger car models (denoted by $\mathcal{J}_n^c$) and the other for light-duty truck models (denoted by $\mathcal{J}_n^t$).

1.5.1.2 Policy 2: Footprint-based Targets / No Transferring / No Trading

Policy 2 sets target values based on vehicle size (footprint). Roughly speaking, compared to Policy 1, Policy 2 gives smaller vehicles more stringent targets, and larger vehicles less stringent targets.

Target value $\bar{e}_{jk}$ varies across jk and is determined as a function of jk's footprint. Specifically, $\bar{e}_{jk}$ for passenger car jk with footprint x_{jk} (square feet) is given by:

$$
\bar{e}_{jk} = \begin{cases} a_c & \text{if } x_{jk} < 41, \\ c_c x_{jk} + d_c & \text{if } 41 \leq x_{jk} \leq 56, \\ b_c & \text{if } 56 < x_{jk}, \end{cases} \tag{1.42}
$$

where $a_c = 3.44$, $b_c = 4.2362$, $c_c = 0.05308$, $d_c = 1.2638$ (Figure 1.2). Except for very small or large cars, whose target (a_c or b_c, respectively) is independent of footprint, $\bar{e}_{jk}$ is given by a linear function of footprint x_{jk}.

Similarly, the target value for light-duty truck jk with footprint x_{jk} is given by:

$$
\bar{e}_{jk} = \begin{cases} a_t & \text{if } x_{jk} < 41, \\ c_t x_{jk} + d_t & \text{if } 41 \leq x_{jk} \leq 66, \\ b_t & \text{if } 66 < x_{jk}, \end{cases} \tag{1.43}
$$

where $a_t = 4.3412$, $b_t = 5.4777$, $c_t = 0.0455$, $d_t = 2.4773$ (Figure 1.2).

Figure 1.2 plots these footprint-based target functions. Comparing Policy 2 to Policy 1, we note that the target is tightened for passenger cars with footprint smaller than 44.70 sq.ft., and relaxed for passenger cars with footprint larger than 44.70 sq.ft. For light-duty trucks, the target is tightened if footprint is smaller than 51.73 sq.ft., and relaxed if footprint is larger than 51.73 sq.ft.

This means that, under Policy 2, selling small (large) cars or trucks does not generate as many positive (negative, respectively) credits as under Policy 1. Therefore, for the purpose of meeting the CAFE standards, selling small (large) cars or trucks is not as advantageous (disadvantageous) under Policy 2 as under Policy 1.

In other words, the required level of corporate average fuel efficiency in each category differs across automakers, and depends on the footprint distribution of their fleet. In order to comply with the standards, automakers focusing on smaller cars or trucks need to achieve higher category average efficiency than those focusing on large cars or trucks.

In (1.42) and (1.43), the slope parameters of the linear part (c_c andc_t) and the threshold footprint values (41, 56 and 66 sq. ft.) are the same as in the actual footprint-based CAFE standards for model years 2012-2016. Based on these, other parameters in each category are determined so that the sales-weighted average gphm of those automakers that are constrained in that category under Policy 2 will be the same both under Policies 1 and 2.[30] As can be seen below, setting parameters this way makes the two policies comparable to each other in the sense that various measures are compared given that the policies achieve (almost) the same market average fuel efficiency.

As in Policy 1, credit transferring between a producer's PC and LDT fleets and credit trading across producers are not allowed. The same fine payment rule applies as under Policy 1 above.

In the simulation under Policy 2, automaker n that chooses to comply with the

[30] In the actual CAFE standards for model years 2012-2016, slope parameters (c_c and c_t) and threshold footprint values remain constant over time, while d_c and d_t become smaller every year, leading to more stringent standards.

CAFE standards solve the following profit maximization problem:

$$\max_{\{p_{jk},e_{jk}\}_{jk\in\mathcal{J}_n}} \sum_{jk\in\mathcal{J}_n}(p_{jk}-c_{jk})s_{jk} \quad \text{s.t.} \quad \begin{aligned} &\sum_{jk\in\mathcal{J}_n^c}(\bar{e}_{jk}-e_{jk})s_{jk}\geq 0, \\ &\sum_{jk\in\mathcal{J}_n^t}(\bar{e}_{jk}-e_{jk})s_{jk}\geq 0, \end{aligned} \tag{1.44}$$

where $\bar{e}_{jk}$ is defined by (1.42) for PCs and by (1.43) for LDTs. This problem is identical to (1.41) except that target value $\bar{e}_{jk}$ varies over jk.

1.5.1.3 Policy 3: Footprint-based Targets / Transferring / No Trading

Policy 3 adds intra-firm credit transferring between PCs and LDTs to Policy 2, so positive credits from one category can be used to offset negative credits in the other category.[31] Except for that, Policies 2 and 3 are the same. Credit transferring provides flexibility to how an automaker achieves required average efficiency.

In the simulation under Policy 3, automaker n that chooses to comply with CAFE standards solve the following profit maximization problem:

$$\max_{\{p_{jk},e_{jk}\}_{jk\in\mathcal{J}_n}} \sum_{jk\in\mathcal{J}_n}(p_{jk}-c_{jk})s_{jk} \quad \text{s.t.} \quad \sum_{jk\in\mathcal{J}_n}(\bar{e}_{jk}-e_{jk})s_{jk}\geq 0, \tag{1.45}$$

where $\bar{e}_{jk}$ is given by (1.42) for PCs, and by (1.43) for LDTs. Credit transferring means that an automaker faces just one unified CAFE constraint that covers both PCs and LDTs.

[31] I have also considered a counterfactual policy that adds intra-firm credit transferring to Policy 1. Thus, this policy is with flat targets and credit transferring, and without credit trading. Simulation results under this policy are very similar to those under Policy 3.

1.5.1.4 Policy 4: Footprint-based Targets / Transferring / Trading

Policy 4 adds inter-firm credit trading to Policy 3. Credit trading between automakers is another flexibility in the standards that might reduce the social cost of achieving a given regulatory goal. This policy includes all three new features of the actual reformed CAFE standards.

In the simulation under Policy 4, automaker n solves the following profit maximization problem:

$$\max_{\{p_{jk},e_{jk}\}_{jk\in\mathcal{J}_n}} \sum_{jk\in\mathcal{J}_n} \{p_{jk} - c_{jk} + z\big(\bar{e}_{jk} - e_{jk}\big)\} s_{jk}, \tag{1.46}$$

where z is the price of one unit (gphm) of fuel efficiency credit and determined in the simulation to clear the credit trading market. If an automaker does not meet its target alone, it is required to purchase credits from other automakers and offset the shortage. Note that there is no explicit CAFE constraint in (1.46) because it is replaced by trading of fuel efficiency credits. Automakers are required to offset a negative balance, if any, by purchasing credits from other automakers.

Target values $\bar{e}_{jk}$ for Policy 4 are lower (that is, more stringent) than those set by (1.42) and (1.43), and used for Policies 2-3 (about 0.086 gphm lower for PCs, and 0.039 gphm for LDTs). These adjustments are necessary because using (1.42) and (1.43) for Policy 4 would result in market average fuel efficiency much worse than under Policies 1-3.[32] Remember that I want to set the market average fuel efficiency

[32]An intuitive explanation is as follows. Under Policies 2 and 3 with functions (1.42) and (1.43), CAFE constraints are far from binding for a number of automakers, but positive credits of these firms cannot be used and have no value. If the same functions (1.42) and (1.43) were used under

to be almost the same under all counterfactual policies in order to make meaningful comparisons of them. The functions are adjusted for this purpose.

1.5.2 Simulation Procedure

Simulations consider 11 automakers (General Motors, Ford, DaimlerChrysler, Toyota, Honda, Nissan, Volkswagen, Hyundai, BMW, Kia, Subaru, Suzuki, Porsche).[33]

Under Policies 1-3, automakers choose if they comply with the standards (either constrained or unconstrained), or if they violate the standards and pay fines. I assume that automakers complying with the actual CAFE standards choose to comply under counterfactual policies as well, and those paying fines choose to do so. Therefore, under Policies 1-3 BMW and Porsche are the only manufacturers that violate the standards, and all other manufacturers comply.[34]

Except for BMW and Porsche under Policies 1-3, all automakers in the following simulations solve a profit maximization problem (1.41), (1.44), (1.45), or (1.46) under a given policy. BMW and Porsche under Policies 1-3 face a profit maximization problem with fine payment, one similar to (1.35) but modified accordingly for each policy.

Under Policies 1-3, I take the following steps to find a Nash equilibrium.

Policy 4, positive credits owned by these automakers could be sold to automakers constrained under Policies 2 or 3, and lower their average fuel efficiency (increase their average gphm). As a result, market average fuel efficiency would be worsened.

[33] Actual CAFE standards consider ownership relations to group manufacturers. My simulations follow this grouping, so General Motors includes Isuzu and Saab; Ford includes Jaguar, Mazda and Volvo; DaimlerChrysler includes Mitsubishi.

[34] Under actual standards, DaimlerChrysler's import passenger car category (mainly, Mercedes-Benz) consistently pays fines. Since the counterfactual standards I consider do not distinguish domestic and import passenger cars, and DaimlerChrysler's import PC sales are much smaller than domestic PC sales, I assume DaimlerChrysler's passenger car category complies with Policies 1-3.

1. Set vectors $\boldsymbol{p}$ and $\boldsymbol{e}$ that contain initial values for prices and fuel efficiency values of all vehicle models in the sample (regardless of the manufacturer).

2. Solve automaker 1's maximization problem with initial values taken from automaker 1's entries of $\boldsymbol{p}$ and $\boldsymbol{e}$.

3. Update $\boldsymbol{p}$ and $\boldsymbol{e}$ by replacing automaker 1's entries in $\boldsymbol{p}$ and $\boldsymbol{e}$ with the solution to step 2.

4. Repeat steps 2 and 3 for automaker 2, then automaker 3, and so on.

5. After solving profit maximization of all automakers, go back to automaker 1 and repeat steps 2, 3 and 4.

6. Continue step 5 until $\boldsymbol{p}$ and $\boldsymbol{e}$ converge.

Because convergence means that each automaker is optimizing given all other automakers' choices, $\boldsymbol{p}$ and $\boldsymbol{e}$ found at the end of this updating process is a Nash equilibrium under the counterfactual CAFE standards.

Under Policy 4, I set an initial value for the price of fuel efficiency credits (z). Then, follow the above steps to find a Nash equilibrium under that price. If there is excess demand for credits in the equilibrium, I increase the credit price, and repeat the above steps to find a Nash equilibrium under the new price. Conversely, if there is excess supply of credits, I lower the credit price and repeat the above steps. This process continues until I find the credit price that clears the market for fuel efficiency credits.

As discussed before, I consider three different possibilities regarding the cost of improving fuel efficiency (each corresponding to K =1, 4 or 7). For each of these cost scenarios, I compare the above counterfactual policies through simulation.

Because a very large number of vehicle models are considered in the simulation and many of them do not sell enough to have a significant impact on the producer's profits (the objective of the optimization problem), different initial vectors $\boldsymbol{p}$ and $\boldsymbol{e}$ lead to slightly different Nash equilibria even under the same cost scenario and policy. Therefore, in order to have a more comprehensive picture, I obtain three Nash equilibria for each cost scenario and policy, starting from different initial vectors.

1.5.3 Simulation Results

This subsection discusses findings from my simulations of the U.S. new vehicle market. Tables 1.8-1.13 and Figures 1.3-1.5 summarize the findings. As I simulate three Nash equilibria for each combination of a cost scenario and a policy, the tables and figures report the average values from the three (slightly different) equilibria. First, I explain what each table and figure shows, then evaluate and compare the four counterfactual CAFE standards in terms of various market-level and automaker-level measures.

To show how footprint-based standards typically work, Table 1.8 reports seven largest automakers' target and achieved (average) fuel efficiency under Policies 1 and 2.[35] For each category (passenger cars, PCs, or light-duty trucks, LDTs) and

[35] Table 1.8 is based on a Nash equilibrium with $K = 4$. I obtain essentially the same values in other equilibria as well.

policy, the columns "Footprint", "Standard" and "gphm" respectively show each manufacturer's sales-weighted average footprint, standard and fuel efficiency. The column "Bind" indicates that the constraint for the corresponding category is binding in the automaker's maximization problem.

Under footprint-based standards, standards vary across automakers depending on the footprint distribution of each automaker's fleet. Roughly speaking, Policy 2 relaxes the standards (compared to Policy 1) for firms that on average sell large cars or trucks, and tightens the standards for those that on average sell small cars or trucks. Table 1.8 shows that General Motors (GM) and Ford, who produce relatively large cars and trucks, face less stringent standards under Policy 2, while other automakers who produce relatively small cars and trucks face more stringent standards.

For selected automakers, Table 1.9 shows Lagrange multipliers (λ^c, λ^t and λ) from profit maximization under Policies 1-3 (see problems (1.41), (1.44) and (1.45)), along with the price of credits (z) and net revenue from credit trading ($\sum_{jk \in \mathcal{J}_n} z\big(\bar{e}_{jk} - e_{jk}\big) s_{jk}$) under Policy 4 (see problem (1.46)). If there are two constraints (under Policies 1 and 2), λ^c is the Lagrange multiplier for the PC constraint, and λ^t is that for the LDT constraint. Under Policy 3, λ is the Lagrange multiplier for the single constraint on the entire fleet of PCs and LDTs. A positive multiplier implies that the corresponding CAFE constraint is binding, and its value is the shadow price of the constraint. We can interpret the Lagrange multipliers as follows: For example, $\lambda = a$ means that if every model's target is increased (i.e., relaxed) by 0.1 gphm, and consequently the automaker's standard, which is obtained as the sales-weighted average of individual targets, is increased by the same amount, then the automaker's profits per vehicle sold goes up by $\$a/10$. The larger the Lagrange multiplier, the more costly the constraint is for the automaker. Note that, if a constraint is relaxed by the introduction of footprint-based targets (GM and Ford), the corresponding Lagrange multiplier under Policy 2 becomes smaller than that under Policy 1. On the other hand, for DaimlerChrysler whose constraints are tightened under Policy 2 as compared to Policy 1, the multipliers are larger under Policy 2.

Tables 1.10-1.12 summarize various market-level, aggregate outcomes from the simulations. They first report the sales-weighted average gphm of PCs, of LDTs, and of the entire fleet (PCs and LDTs); sales-weighted average footprint of PCs and LDTs; and share of PCs in the entire fleet (PCs and LDTs). Tables 1.10-1.12 also show aggregate outcomes relating to the sale/purchase and use of all PCs and LDTs of the model year 2001 fleet, including [1] automakers' profits; [2] consumer surplus;[36] [3] technology costs for fuel efficiency adjustment;[37] [4] estimated fuel use external costs over the vehicles' lifetime;[38] and [5] social welfare (=[1]+[2]-[4]).Values per new vehicle sold are shown in the lower part of the tables. Values for Policies

[36] Consumer surplus is calculated with the method of Small and Rosen (1981).

[37] All counterfactual policies are designed to achieve the market average fuel efficiency of approximately 4.10 gphm, which is about 0.05 gphm more efficient than the actual average in model year 2001. Under a counterfactual policy, automakers adjust each (vehicle) model's fuel efficiency by changing the technology level applied. This adjustment increases or decreases its production cost. I call this a change in the technology cost for fuel efficiency. Tables 1.10-1.12 report the sales-weighted average of these technology cost changes due to fuel efficiency adjustment.

[38] Fuel use external costs over the vehicles' lifetime are calculated with an estimated lifetime VMT schedule and survival probability of the average car and truck (Table 4-3 of Environmental Protection Agency and National Highway Traffic Safety Administration, 2010). A 5% discount rate is applied to future external costs. Fuel use externality is evaluated at \$0.42 per gallon, based on an estimate by Parry et al. (2007). With these assumptions, I estimate fuel use external costs over a vehicle's lifetime by multiplying first year's fuel use external costs obtained from a simulation by 8.457 for passenger cars and by 8.389 for light-duty trucks.

2-4 are given as changes from Policy 1. Based on Tables 1.10-1.12, Figures 1.3-1.5 graphically show changes in [1]-[5] from Policy 1.

Policies 1 and 2 are compared below to see the effects of making each vehicle's target value depend on its footprint. Then, I discuss the impacts of intra-firm credit transferring between PCs and LDTs (Policy 3), and inter-firm credit trading (Policy 4)

1.5.3.1 Effects of Footprint-based Targets

First, I compare outcomes under Policies 1 and 2. Policy 1 is essentially equivalent to the actual CAFE standards for model year 2001. Particularly, each vehicle's target value under Policy 1 depends only on its category (PCs or LDTs), but not on vehicle size. Policy 2 differs from Policy 1 only in how each vehicle's target is set. Target values depend on vehicle size (footprint) and larger vehicles receive less stringent (i.e., larger in terms of gphm) targets.

Comparing Policies 1 and 2 in Tables 1.10-1.12 and Figures 1.3-1.5 suggests that introducing footprint-based targets alone makes little difference at the market level, as long as the two policies realize the (almost) same market average fuel efficiency. Vehicle size (footprint) and the share of passenger cars in the market change only slightly. Changes in measures [1]-[5] are relatively small in magnitude and the same measure sometimes takes opposite signs under different simulations (even under the same K). Coleman and Harrington (2010) also obtain similar results that the impact of footprint-based targets is generally small.

As for changes in profits from under Policy 1 to under Policy 2 at the individual automaker level, the results are not very robust across different simulations (even under the same K). For many automakers, individual profits may increase or decrease under Policy 2 depending on the equilibrium realized. In other words, just making each vehicle's target footprint-dependent does not have an impact that is large enough to bring consistent changes across different equilibria.

1.5.3.2 Effects of Credit Transferring within an Automaker

Policy 3 adds to Policy 2 credit transferring between a firm's PC and LDT fleets, but not credit trading across firms. So automakers face a single, combined constraint, and can use positive credits in one category to offset negative credits in the other. Overall, simulations find that Policy 3 is effective in reducing gasoline consumption.

Tables 1.10-1.12 show that, compared to Policies 1 and 2, Policy 3 makes PCs less fuel efficient (about 2-5%), and LDTs more efficient (about 2-4%), although fuel efficiency of the entire fleet (PCs and LDTs) remains almost the same (as desired). Fuel efficiency improves in the LDT fleet and worsens in the PC fleet because automakers that are constrained in both categories (PCs and LDTs) under Policy 2 (GM, Ford and DaimlerChrysler) choose to earn credits in the LDT category to offset the shortages in the PC category under Policy 3. Since my estimates suggest that it is typically more costly to improve LDTs' fuel efficiency than PCs', the aggregate technology cost to achieve the targeted market average efficiency increases by about \$12 million-\$97 million (compared to Policy 1).

Since improving LDTs' fuel efficiency is more costly (in terms of technology costs), it is counter-intuitive that the Big Three choose to do so to cancel out worsened fuel efficiency of PCs. One reason is that LDT buyers are willing to pay more for gphm improvement than PC buyers. LDTs are generally driven longer distance than PCs, so that LDTs owners receive greater fuel savings from marginal gphm improvement, thus have higher willingness to pay for it. Producers respond to this demand and improve LDTs' fuel efficiency.

More influential is the difference in price sensitivity of demand between PCs and LDTs. PC buyers generally earn lower income, so that demand for PCs are more sensitive to vehicle prices. This demand structure makes it more profitable for automakers to increase gphm of PCs to lower their technology costs and prices, and to decrease gphm of LDTs to raise their technology costs and prices. Note that the share of PCs in the market increases as a result.

In spite of the increase in the aggregate technology cost mentioned above, these changes in fuel efficiency make Policy 3 effective in reducing aggregate fuel consumption (about 0.1%-0.25% compared to Policies 1 and 2). Indeed, among the four counterfactual policies, Policy 3 is most useful for the purpose of curbing gasoline consumption. The lower the production cost (technology cost) for improving fuel efficiency, the more gasoline consumption is reduced because of larger fuel efficiency adjustments in each category. Despite almost no change in market average fuel efficiency, aggregate gasoline consumption decreases drastically because LDTs are on average driven longer distance than PCs both annually and over the vehicle's life, so that gphm changes of LDTs influence aggregate fuel consumption more than those of PCs. Shrinking the fuel efficiency differences between PCs and LDTs is very effective in cutting fuel consumption, and Policy 3 generates this substitution through producers' and consumers' choices in the market.

In terms of social welfare, Policy 3 is clearly superior to Policies 1 and 2 under the low and middle valuation/cost scenarios ($K = 1, 4$), but not under the high cost scenario because Policy 3 is not as effective under the high valuation/cost scenario ($K = 7$) in reducing gasoline consumption. We also note that reduced fuel use externalities account for a significant portion of a change in social welfare (especially for $K = 1, 4$).

1.5.3.3 Effects of Credit Trading across Automakers

Policy 4 adds to Policy 3 inter-firm trading of fuel efficiency credits. Policy 4 mirrors the new CAFE standards because it features all three instruments recently introduced into the standards (footprint-based targets, intra-firm credit transferring and inter-firm credit trading).

Of the four counterfactual policies, Policy 4 achieves the required market average fuel efficiency (approximately 4.10 gphm) at the lowest aggregate technology cost for fuel efficiency adjustment.[39] The aggregate technology cost is reduced by \$109 million-\$136 million, compared to Policy 1. It goes down because automakers that face higher marginal costs of improving their vehicles' fuel efficiency choose not to meet the standards independently, but they instead choose to purchase credits from other automakers that can improve fuel efficiency at lower costs. In one of

[39] Refer to footnote 37 for an explanation of the technology cost for fuel efficiency adjustment.

their simulations, Austin and Dinan (2005) find similar effects of inter-firm trading of fuel efficiency credits on the aggregate technology cost, although their counterfactual policies do not seem to be at the same level of stringency, so their results likely show the total effects of design differences (i.e., with or without credit trading) and stringency differences.

My simulations suggest that a 0.35-0.9% tightening of CAFE standards would increase the aggregate technology cost by about a similar amount. Thus, introducing inter-firm credit trading can offset an increase in the aggregate technology cost due to a 0.35-0.9% tightening of standards annually. This size of tightening lowers aggregate gasoline consumption over the lifetime of a particular year's entire fleet (PCs and LDTs) by about 0.2-0.65%.

To better understand the effects of credit trading, Tables 1.9 and 1.13 report simulation results for individual automakers. Table 1.9 shows the price of fuel efficiency credits (z), and net revenue from sale or purchase of credits ($\sum_{jk\in\mathcal{J}_n} z\big(\bar{e}_{jk} - e_{jk}\big)s_{jk}$). Table 1.13 compares each automaker's standard and average fuel efficiency under Policies 3 and 4.

Under Policy 4, the average efficiency (in gphm) of each of the Big Three is greater than the corresponding standard (Table 1.13), meaning that they are purchasers of credits (Table 1.9). On the contrary, other manufactures in the tables keep the average efficiency (in gphm) smaller than the corresponding standard (Table 1.13), and sell the credits thus earned to other automakers (Table 1.9).

Compared to Policy 3 (and Policies 1-2 as well), the Big Three's average efficiency worsens under Policy 4, while other firms' average efficiency improves. This explains why the aggregate technology costs for fuel efficiency adjustment is lowered under Policy 4. Vehicles from the Big Three on average face higher marginal costs for improving fuel efficiency than vehicles from other automakers. Therefore, offsetting the Big Three's worsened fuel efficiency by improving other automakers' fuel efficiency significantly reduces the aggregate technology cost for fuel efficiency adjustment, compared to the case in which the Big Three achieve their targets by themselves.

This reduction in the aggregate technology cost leads to an increase in social welfare (Tables 1.10-1.12 and Figures 1.3-1.5). The result is consistent across different equilibria I simulated. Indeed, Policy 4 attains the highest social welfare under all equilibria I simulated. Comparing the magnitude of the changes in the aggregate technology cost and social welfare under each scenario implies that a reduction in the technology cost due to inter-firm credit trading can account for a large portion of a social welfare increase.

On the other hand, inter-firm credit trading is not as helpful for reducing gasoline consumption as intra-firm credit transferring. Under Policy 3 that features intra-firm transferring but not inter-firm trading, fuel consumption is reduced through fuel efficiency improvement in the Big Three's LDT fleet. Under Policy 4, buying credits from other firms is a better option for the Big Three, so that their LDTs' fuel efficiency does not improve as much as under Policy 3, and Policy 4 does not attain large fuel savings.

On the other hand, at the individual automaker level, some firms lose and others gain under Policy 4. Although this is partially due to payment or revenue from credit trading, more crucial are changes in market shares. As in Table 1.13, credit trading worsens fuel efficiency of the Big Three's models, and improves that of other firms' models. With these changes, consumers' demand shifts away from the Big Three, decreasing the Big Three's sales and increasing others' sales. The sales changes are large enough that the Big Three's profits fall and other firms' profits increase consistently across different simulations.

1.6 Conclusion

This paper has examined the impacts of three major policy instruments recently introduced into the reformed Corporate Average Fuel Economy (CAFE) standards, namely, footprint-based targets, intra-firm credit transferring and inter-firm credit trading. Previous studies have paid little attention to the design of CAFE standards, as compared to the stringency (i.e., target fuel economy values). This paper has particularly focused on the design and is the first study to compare all of the three instruments simultaneously. I have modeled household vehicle and VMT (vehicle miles traveled) choices, and automakers' decisions in the vehicle market regulated by CAFE standards, then estimated the model with year 2001 U.S. data. Modeling and estimating the VMT choice along with the vehicle choice is important in economic studies of fuel economy, as the value of fuel economy for households (i.e., savings in household fuel spending) crucially depends on VMT. Based on estimation results, I have simulated market equilibria under various counterfactual CAFE policies to see the impacts of the three instruments. Unlike previous studies on the design of CAFE standards, I carefully set these counterfactual standards so that they achieve almost the same market average fuel economy, and thus are at essentially the same level of stringency. This makes it possible to examine the impacts of design differences in the counterfactual standards, without being confounded by the impacts of stringency differences as in previous studies.

Counterfactual simulations provide the following key policy implications. First, simply replacing "flat" targets with "footprint-based" targets (Policy 2) have little impact on the vehicle market and aggregate gasoline consumption.

Second, allowing intra-firm credit transferring across vehicle categories, but not inter-firm credit trading (as in Policy 3) is effective in curbing aggregate gasoline consumption. When credit transferring is allowed, the Big Three make their light-duty trucks (LDTs) more fuel efficient and their passenger cars (PCs) less fuel efficient. Since LDTs are on average driven more miles than PCs, this shift in fuel efficiency reduces gasoline consumption from the fleet of (about 0.1%-0.25%). Shrinking the fuel efficiency differences between PCs and LDTs is useful for reducing fuel consumption, and Policy 3 generates this substitution through producers' and consumers' voluntary choices in the market.

The value of reduced gasoline consumption depends on the price of gasoline (about \$1.40 in the simulations) and an estimate of the external costs of gasoline use (\$0.42 in the simulations). The more expensive gasoline is (as in 2012, when the price is at around \$3.60) or the more serious the regulator considers gasoline use externalities to be, the more valuable reducing gasoline consumption becomes for the society. Under these circumstances, allowing intra-firm credit transferring, but not inter-firm credit trading, would become more attractive to policymakers.

Third, further allowing inter-firm credit trading in addition to intra-firm credit transferring (as in Policy 4) attains a given market average fuel efficiency at the lowest aggregate technology costs for fuel efficiency adjustment. Credit trading shifts the role of improving fuel efficiency at additional technology costs to automakers that can do so relatively cheaply, thus significantly lowering the aggregate technology costs for fuel efficiency adjustment. This makes Policy 4 the most effective among the four simulated policies in terms of social welfare. Since Policy 4 mimics the actual new CAFE standards, simulation results are overall in favor of the shifts to the new standards.

Yet, there may be a few drawbacks to Policy 4. By adding inter-firm credit trading to intra-firm credit transferring, Policy 4 nullifies the impact of intra-firm credit transferring on gasoline use reduction. In addition, it seems to affect negatively on the profits of some automakers, notably the Big Three.

This paper has several limitations and possible extensions. First, this paper focused on new vehicles, so the possibility of substitution between new vehicles and used vehicles is ignored. In practice, used cars account for a large part of the vehicle market, and Bento et al. (2009) point out the importance of considering used cars in analyzing the vehicle market. Because the National Household Travel Survey has data on used car purchases, it would be possible in future work to include used cars in the vehicle choice model (not as individual used cars but as bundles of used cars based on age, vehicle type, make, etc). Second, in my framework automakers choose only the price and fuel economy of vehicles. Studies such as Klier and Linn (2012) and

Whitefoot et al. (2012) additionally endogenize the choice of other vehicle attributes such as weight and horsepower, and consider trade-offs between these attributes and fuel economy. Incorporating their approach into my framework would enrich the analysis of this paper. Adding used cars and endogenizing other vehicle attributes are especially important when we focus on significant stringency changes, as these changes have larger impacts on decisions by consumers and producers, and market outcomes. Since I compare policies that are essentially at the same level of stringency, the effects of overlooking these features would be relatively small. Lastly, while this paper examined CAFE standards, other types of policies can be used to improve fuel economy and reduce fuel consumption, such as gasoline taxes and feebates. In future work, I plan to simulate these policies with the framework of this paper and compare the effectiveness of various policies in achieving a given target.

	Mean	Std. dev.	Min	Max
Household Characteristics				
Vehicle miles traveled (100 miles)	153.78	120.83	0.32	2,000
Household size (persons)	2.78	1.29	1	14
Population density persons / square mile	3,000	4,762	50	30,000
# of vehicles owned	2.50	1.17	1	12
Household income ($)	70,715	38,342	6,950	131,150
Regular gasoline price ($/gallon)	1.50	0.07	1.29	1.67
Vehicle Attributes				
Vehicle price ($)	28,234	13,906	9,045	114,645
Fuel efficiency (miles/gallon)	24.38	5.02	14.42	57.57
Fuel efficiency (gallons/100 miles)	4.27	0.85	1.74	6.94
Footprint (sq ft)	46.67	6.79	34.50	66.60
Curb weight (lb)	4,015	762	2,375	6,000
Horsepower (hp)	199	54	70	415
Horsepower/weight (hp/lb)	0.050	0.011	0.022	0.111
Passenger car dummy	0.538	0.499	0	1
Van dummy	0.096	0.295	0	1
SUV dummy	0.234	0.424	0	1
Pickup truck dummy	0.131	0.338	0	1

The number of households is 5884. The number of vehicle models is 457. Fuel efficiency in the table is unadjusted fuel efficiency used for CAFE standards. Fuel eoonomy ratings consumers see are adjusted by the EPA to account for actual driving conditions. Roughly, adjusted values are 15% less fuel efficient. All dollars are in 2001 U.S. dollars.

Table 1.1: Summary Statistics

	2001			2012		2011
	Target (mpg)	CAFE (mpg)	Sales (vehicles)	Target (mpg)	CAFE (mpg)	Sales (vehicles)
Domestic Passenger Cars						
General Motors	27.5	28.3	2,184,214	32.3	33.0	1,238,428
Ford	27.5	27.7	1,309,834	33.2	35.4	829,957
DaimlerChrysler (Chrysler)	27.5	27.9	739,681	32.2	31.7	415,841
Toyota	-	-	-	33.4	36.7	628,512
Honda	27.5	32.7	794,448	30.4	36.2	437,065
Nissan	27.5	27.9	137,253	32.9	35.2	340,309
Import Passenger Cars						
General Motors	27.5	28.4	71,503	-	-	92,404
Ford	27.5	27.9	236,797	31.7	31.0	-
DaimlerChrysler (Chrysler)	27.5	26.5	215,072	-	-	-
Toyota	27.5	30.6	986,390	34.1	50.0	1,103,744
Honda	27.5	29.8	42,271	31.1	43.1	86,021
Nissan	27.5	28.7	257,247	34.1	37.5	327,542
Light Duty Trucks						
General Motors	20.7	20.7	1,902,731	23.8	23.5	1,089,179
Ford	20.7	20.4	1,988,290	24.1	24.4	831,846
DaimlerChrysler (Chrysler)	20.7	20.8	1,812,945	25.7	24.6	846,202
Toyota	20.7	22.1	652,229	25.6	25.3	287,882
Honda	20.7	25.0	281,606	25.3	27.9	486,255
Nissan	20.7	20.7	356,816	26.1	24.7	302,394

The table compares selected automakers' target to be achieved, average fuel economy and sales volume in model years 2001 and 2012 (or 2011). The data are taken from National Highway Traffic Safety Administration (2012). Model year 2012 data are projected values. Honda's MY 2012 standards are strangely low in all vehicle categories and most likely wrong.

Table 1.2: Comparison of Model Years 2001 and 2012 (or 2011)

		[1]	[2]	[3]
$\boldsymbol{\beta}_1$	log(hh size)	0.011	-0.023	-0.040
		(0.045)	(0.047)	(0.046)
	log(pop den)	0.065***	0.090***	0.100***
		(0.021)	(0.017)	(0.018)
	log(# vehicles)	0.008	-0.018	-0.024
		(0.053)	(0.055)	(0.054)
	log(income)	-0.308***	-0.343***	-0.350***
		(0.065)	(0.056)	(0.057)
	max[log(income/46,715),0]	-0.391*	-0.365*	-0.310
		(0.204)	(0.199)	(0.195)
	max[log(income/73,580),0]	-1.353	-1.877**	-1.710**
		(1.038)	(0.860)	(0.805)
δ_1	(mean)	-6.045***	-5.983***	-6.017***
		(0.527)	(0.538)	(0.526)
	(std. dev.)	0.068**	0.064**	0.061**
		(0.027)	(0.028)	(0.028)
$\boldsymbol{\beta}_2$	log(hh size)	0.138***	0.101**	0.080
		(0.049)	(0.051)	(0.049)
	log(pop den)	0.001	0.027	0.038**
		(0.021)	(0.017)	(0.018)
	log(# vehicles)	-0.147**	-0.177***	-0.185***
		(0.057)	(0.058)	(0.057)
	log(income)	-0.150**	-0.190***	-0.200***
		(0.072)	(0.063)	(0.062)
	max[log(income/46,715),0]	-0.313	-0.269	-0.202
		(0.223)	(0.216)	(0.210)
	max[log(income/73,580),0]	-1.401	-1.941**	-1.783**
		(1.042)	(0.863)	(0.807)
$\boldsymbol{\gamma}_2$	log(footprint)	0.420***	0.374***	0.337***
		(0.118)	(0.116)	(0.114)
	log(hp/lb)	-0.089	-0.094	-0.095
		(0.089)	(0.088)	(0.088)
	Van	0.049	0.048	0.048
		(0.041)	(0.041)	(0.040)
	SUV	0.153***	0.135***	0.118***
		(0.033)	(0.032)	(0.031)
	Pickup	0.148***	0.140***	0.132***
		(0.040)	(0.040)	(0.039)
δ_2	(mean)	-3.582***	-2.104***	-1.560**
		(0.826)	(0.814)	(0.790)
	(std. dev.)	0.019	0.016	0.015
		(0.031)	(0.031)	(0.030)
α_3		0.728***	0.784***	0.829***
		(0.079)	(0.074)	(0.072)
K		1	4	7

Estimation results in relation to the indirect utility function, i.e., equations (1.13)-(1.17). Standard errors in parentheses.
*** indicates significance at the 1% level, ** at the 5% level, * at the 10% level. Standard errors in parentheses.

Table 1.3: Estimation Results: Indirect Utility Function

		[1]	[2]	[3]
$\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1$	log(hh size)	0.127***	0.124***	0.121***
		(0.022)	(0.022)	(0.022)
	log(pop. den.)	-0.064***	-0.063***	-0.062***
		(0.006)	(0.006)	(0.006)
	log(# vehicles)	-0.155***	-0.159***	-0.162***
		(0.024)	(0.024)	(0.024)
	log(income)	0.158***	0.152***	0.150***
		(0.033)	(0.033)	(0.033)
	max[log(income/46,715),0]	0.079	0.096	0.108
		(0.092)	(0.092)	(0.092)
	max[log(income/73,580),0]	-0.048	-0.064	-0.073
		(0.109)	(0.109)	(0.109)
$\boldsymbol{\gamma}_2$	log(footprint)	0.420***	0.374***	0.337***
		(0.118)	(0.116)	(0.114)
	log(hp/lb)	-0.089	-0.094	-0.095
		(0.089)	(0.088)	(0.088)
	Van	0.049	0.048	0.048
		(0.041)	(0.041)	(0.040)
	SUV	0.153***	0.135***	0.118***
		(0.033)	(0.032)	(0.031)
	Pickup	0.148***	0.140***	0.132***
		(0.040)	(0.040)	(0.039)
$\alpha_3 - 1$		-0.272***	-0.216***	-0.171**
		(0.079)	(0.074)	(0.072)
μ	(std. dev.)	0.674***	0.675***	0.675***
		(0.007)	(0.007)	(0.007)
K		1	4	7

Estimation results in relation to the continuous VMT choice equation : $\log(m_{ijk}) = -\delta_{1i} + \delta_{2i} + \log(\alpha_3) + (\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)\boldsymbol{x}_i + \boldsymbol{\gamma}_2 \boldsymbol{z}_{jk} + (\alpha_3 - 1)\log(r_{ij}) + \mu_{ijk}$, with $m_{ijk} = K \times m_{ijk1}$, where m_{ijk1} is first year's annual VMT. *** indicates significance at the 1% level, ** at the 5% level, * at the 10% level. Standard errors in parentheses.

Table 1.4: Estimation Results: VMT Demand Equation

		K	Annual income $22,360	$59,020	$131,150
Toyota Corolla, 1.8L, $13,395	[1]	1	2.11%	1.37%	0.61%
	[2]	4	2.06%	1.35%	0.58%
	[3]	7	2.02%	1.33%	0.58%
Cadillac DeVille, 4.6L, $44,101	[1]	1	0.23%	0.27%	1.35%
	[2]	4	0.26%	0.32%	1.32%
	[3]	7	0.29%	0.37%	1.24%

Predicted choice probabilities of two selected vehicle models (the economical Toyota Corolla 1.8 liter engine model and the luxurious Cadillac DeVille 4.6 liter engine model) for three hypothetical households. The three households' annual income is set at the 5th, 50th or 95th percentile of the sample, and other household characteristics are set at the sample average.

Table 1.5: Predicted Probabilities of Choosing Selected Models

		[1]	[2]	[3]
	K	1	4	7
Own-price elast. of sales	(mean)	-2.47	-1.99	-1.94
	(std. dev.)	0.27	0.19	0.20
Own-gphm elast. of sales	(mean)	-0.09	-0.29	-0.49
	(std. dev.)	0.03	0.08	0.13

The mean and standard deviation of elasticities of sales with respect to each vehicle's own price and fuel efficiency (gphm).

Table 1.6: Predicted Elasticities

	1999	2000	2001	2002	2003
Domestic Passenger Cars					
DaimlerChrysler	**27.2**	27.9	27.9	27.7	29.7
Ford	27.6	28.3	27.7	27.9	27.9
General Motors	27.7	27.9	28.3	28.8	28.9
Honda	33.5	31.4	32.7	32.4	34.4
Nissan	29.9	28.1	27.9	28.9	28.9
Toyota	28.3	33.3	-	33.6	28.1
Standard	27.5	27.5	27.5	27.5	27.5
Import Passenger Cars					
BMW	**25.4**	**24.8**	**25.0**	**26.2**	**26.8**
Daewoo	-	28.6	28.6	28.2	29.1
DaimlerChrysler	**26.5**	**25.3**	**26.5**	**26.6**	**26.3**
Ford	28.5	**27.4**	27.9	28.1	28.2
General Motors	**25.5**	**25.4**	28.4	27.8	28.3
Honda	29.4	29.3	29.3	29.8	31.9
Hyundai	30.8	30.7	31.3	31.2	30.4
Kia	30.9	30.0	30.5	29.7	30.4
Nissan	29.9	28.3	28.7	29.5	**27.4**
Porsche	**24.1**	**24.3**	**23.7**	**23.9**	**24.1**
Subaru	27.5	28.0	27.8	27.6	27.6
Suzuki	35.5	35.0	35.1	33.8	33.0
Toyota	29.9	28.9	30.6	29.3	32.4
Volkswagen	28.6	28.8	28.5	29.5	29.8
Standard	27.5	27.5	27.5	27.5	27.5
Light Duty Trucks					
BMW	-	**17.5**	**19.2**	**20.1**	**20.0**
DaimlerChrysler	20.8	21.4	20.8	21.5	22.2
Ford	20.8	21.0	**20.4**	20.7	21.3
General Motors	**20.3**	21.0	20.7	21.2	21.3
Honda	26.1	25.4	25.0	25.4	24.7
Hyundai	-	-	25.0	24.5	24.4
Isuzu	21.1	20.9	21.1	21.0	22.3
Kia	24.4	23.5	23.0	21.4	**19.7**
Nissan	21.2	20.8	20.7	20.7	21.9
Suzuki	23.8	23.0	22.1	21.9	21.8
Toyota	22.9	21.8	22.1	22.1	21.9
Volkswagen	**19.1**	**18.9**	**20.4**	**20.6**	21.3
Standard	20.7	20.7	20.7	20.7	20.7

The table shows the sales-weighted harmonic average mpg by automaker and category, as well as the target for each category. A red entry means that the average is under the target for the automaker-category, though it does not necessarily imply a violation because of banking and borrowing of credits.

Table 1.7: Corporate Average Fuel Economy (in mpg) (1999-2003)

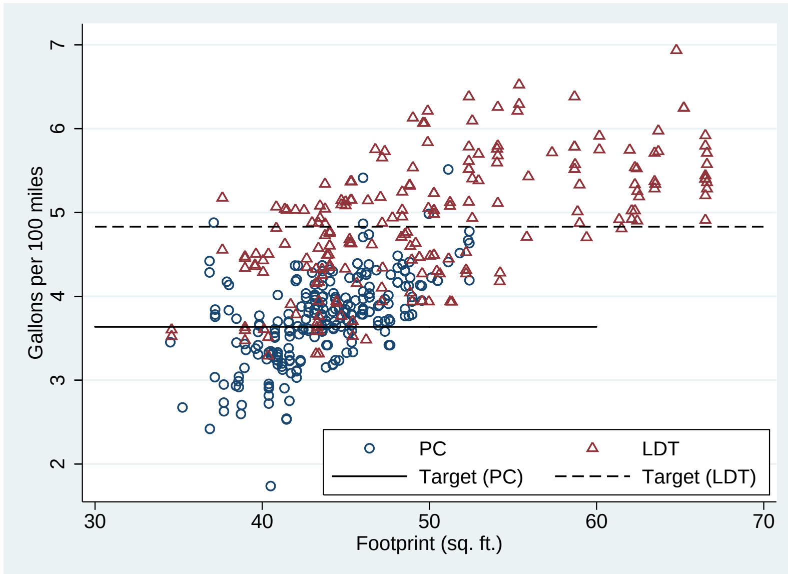

Figure 1.1: Target Values under Flat Standards (Policy 1) and Actual Fuel Efficiency

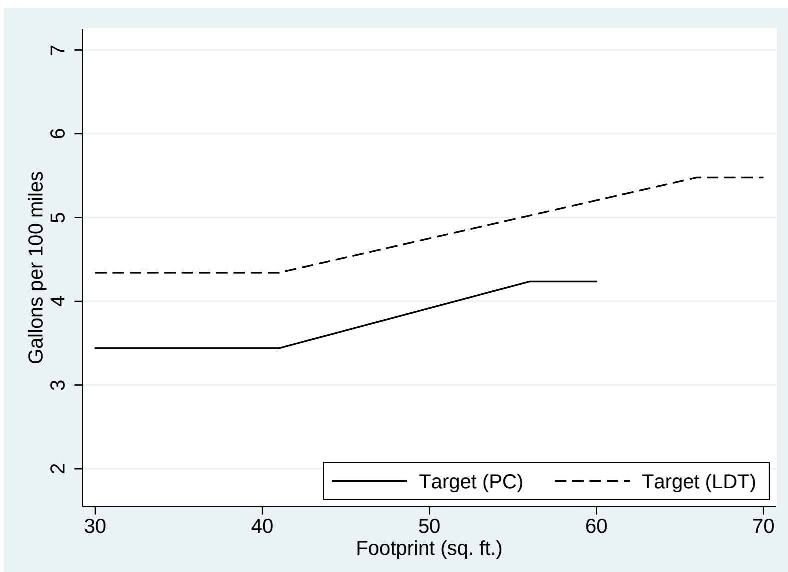

Figure 1.2: Target Values under Footprint-based Standards (Policies 2 and 3)

	Passenger Cars							
	Policy 1				Policy 2			
	Footpr. (sq. ft.)	Target (gphm)	Fuel eff. (gphm)	Bind	Footpr. (sq. ft.)	Target (gphm)	Fuel eff. (gphm)	Bind
General Motors	44.95	3.64	3.64	×	44.95	3.66	3.66	×
Ford	44.85	3.64	3.64	×	44.86	3.66	3.66	×
DaimlerChrysler	44.23	3.64	3.64	×	44.23	3.62	3.62	×
Toyota	42.78	3.64	3.29		42.78	3.57	3.29	
Honda	43.66	3.64	3.16		43.66	3.58	3.16	
Nissan	43.00	3.64	3.56		43.00	3.55	3.55	
Volkswagen	41.77	3.64	3.63		41.77	3.49	3.49	×
All 11 automakers	43.87	-	3.52		43.87	-	3.52	

	Light Duty Trucks							
	Policy 1				Policy 2			
	Footpr. (sq. ft.)	Target (gphm)	Fuel eff. (gphm)	Bind	Footpr. (sq. ft.)	Target (gphm)	Fuel eff. (gphm)	Bind
General Motors	52.61	4.83	4.83	×	52.61	4.87	4.87	×
Ford	52.69	4.83	4.83	×	52.71	4.87	4.87	×
DaimlerChrysler	50.37	4.83	4.83	×	50.40	4.77	4.77	×
Toyota	48.30	4.83	4.57		48.30	4.68	4.57	
Honda	48.85	4.83	4.01		48.85	4.70	4.01	
Nissan	45.19	4.83	4.74		45.20	4.53	4.53	×
Volkswagen	-	-	-		-	-	-	
All 11 automakers	50.91	-	4.76		50.92	-	4.76	

The table reports selected automakers' three sales-weighted averages (footprint, target and fuel efficiency) under Policies 1 and 2. In each of the two vehicle categories (passenger cars and light-duty trucks), CAFE standards require an automaker's average fuel efficiency (gphm) to be below its target.

Table 1.8: Target and Achieved Fuel Efficiency under Policies 1 and 2 (by Automaker)

	$K=1$						
Policy:	1		2		3	4	
Flat (FL)/Ftprnt (FP):	FL		FP		FP	FP	
Intra-firm transferring:	No		No		Yes	Yes	
Inter-firm trading:	No		No		No	Yes	
	λ^c (\$)	λ^t (\$)	λ^c (\$)	λ^t (\$)	λ (\$)	z (\$)	Credit sales (mil.\$)
General Motors	76.0	49.8	75.7	49.3	52.3	-	-20.8
Ford	64.8	54.7	64.4	54.2	57.4	-	-47.9
DaimlerChrysler	82.5	62.5	82.8	63.1	64.9	-	-58.1
Toyota	-	-	-	-	-	-	48.9
Honda	-	-	-	-	-	-	51.0
Nissan	-	-	0.1	1.9	0.9	-	9.5
Volkswagen	-	-	1.4	-	1.4	-	3.0
All 11 automakers	-	-	-	-	-	59.3	0.0

	$K=4$						
	1		2		3	4	
General Motors	74.2	53.2	72.8	50.8	61.0	-	-39.8
Ford	66.8	56.7	65.7	55.1	59.1	-	-30.3
DaimlerChrysler	88.0	66.1	88.9	67.6	71.2	-	-37.5
Toyota	-	-	-	-	-	-	41.5
Honda	-	-	-	-	-	-	43.3
Nissan	-	-	-	6.0	2.4	-	8.0
Volkswagen	-	-	5.9	-	5.9	-	2.5
All 11 automakers	-	-	-	-	-	49.6	0.0

	$K=7$						
	1		2		3	4	
General Motors	72.7	54.8	71.6	52.0	62.8	-	-29.5
Ford	68.5	58.2	66.1	55.6	60.8	-	-21.1
DaimlerChrysler	92.2	68.5	93.5	71.0	75.0	-	-22.9
Toyota	-	-	-	-	-	-	28.1
Honda	-	-	-	-	-	-	33.2
Nissan	-	-	-	9.7	3.6	-	4.3
Volkswagen	-	-	10.0	-	9.1	-	1.2
All 11 automakers	-	-	-	-	-	44.2	0.0

Under Policies 1-3, λ^c, λ^t or λ is the Lagrange multiplier (i.e., shadow price) of the corresponding constraint in an automaker's profit maximization (λ^c for passenger cars (PCs), λ^t for light-duty trucks (LDTs), and λ for the entire fleet of PCs and LDTs). For Policy 4, the price of credits (z) and net revenue from credit trading are shown. All dollars are in 2001 U.S. dollars.

Table 1.9: The (Shadow) Price of CAFE Constraints and Fuel Efficiency Credits

	Policy 1	Change from Policy 1: Policy 2	Policy 3	Policy 4
Flat (FL)/Footprint (FP):	FL	FP	FP	FP
Intra-firm transferring:	No	No	Yes	Yes
Inter-firm trading:	No	No	No	Yes
Fuel efficiency (PC) (gphm)	3.52	0.0001	0.1691	0.0154
Fuel efficiency (LDT) (gphm)	4.76	0.0005	-0.1898	-0.0175
Fuel efficiency (PC + LDT) (gphm)	4.10	0.0002	0.0010	0.0001
Footprint (PC) (sq. ft.)	43.87	0.0015	0.0024	-0.0059
Footprint (LDT) (sq. ft.)	50.91	0.0120	0.0142	0.0073
Market share (PC) (%)	53.22	0.0072	0.0223	-0.0094
Automakers' profits (million $) [1]	-	-26.12	-14.39	28.29
Consumer surplus (mil. $) [2]	-	20.23	11.82	88.43
Technology costs for fuel eff. adjust. (mil. $) [3]	-	-8.66	12.44	-113.16
Lifetime fuel use externalities (mil. $) [4]	42,632	7.30	-96.86	-36.08
Social welfare (= [1] + [2] − [4]) (mil. $) [5]	-	-13.19	94.29	152.80
Automakers' profits / vehicle ($)	-	-1.73	-0.95	1.87
Consumer surplus / vehicle ($)	-	1.34	0.78	5.84
Technology costs for fuel eff. adjust. / vehicle ($)	-	-0.57	0.82	-7.48
Lifetime fuel use externalities / vehicle ($)	2,818	0.48	-6.40	-2.38
Social welfare / vehicle ($)	-	-0.87	6.23	10.10

The table shows simulation results at the market level for the case of $K = 1$. All dollars are in 2001 U.S. dollars.

$K = 1$: low valuation of fuel efficiency by consumers at the time of vehicle purchase, and low marginal production costs for improving fuel efficiency.

Table 1.10: Market Level Outcomes ($K = 1$)

	Policy 1	Change from Policy 1: Policy 2	Policy 3	Policy 4
Flat (FL)/Footprint (FP):	FL	FP	FP	FP
Intra-firm transferring:	No	No	Yes	Yes
Inter-firm trading:	No	No	No	Yes
Fuel efficiency (PC) (gphm)	3.52	-0.0005	0.1184	-0.0005
Fuel efficiency (LDT) (gphm)	4.76	0.0004	-0.1325	-0.0009
Fuel efficiency (PC + LDT) (gphm)	4.10	-0.0002	0.0007	-0.0008
Footprint (PC) (sq. ft.)	43.87	0.0012	0.0029	-0.0218
Footprint (LDT) (sq. ft.)	50.91	0.0128	0.0163	-0.0091
Market share (PC) (%)	53.23	0.0138	0.0347	0.0106
Automakers' profits (million $) [1]	-	55.96	63.42	-120.75
Consumer surplus (mil. $) [2]	-	-32.73	-63.18	271.06
Technology costs for fuel eff. adjust. (mil. $) [3]	-	-7.44	93.09	-109.23
Lifetime fuel use externalities (mil. $) [4]	42,618	4.04	-67.93	-18.07
Social welfare (= [1] + [2] − [4]) (mil. $) [5]	-	19.19	68.18	168.39
Automakers' profits / vehicle ($)	-	3.70	4.19	-7.98
Consumer surplus / vehicle ($)	-	-2.16	-4.18	17.92
Technology costs for fuel eff. adjust. / vehicle ($)	-	-0.49	6.15	-7.22
Lifetime fuel use externalities / vehicle ($)	2,817	0.27	-4.49	-1.19
Social welfare / vehicle ($)	-	1.27	4.51	11.13

The table shows simulation results at the market level for the case of $K = 4$. All dollars are in 2001 U.S. dollars.
$K = 4$: medium valuation of fuel efficiency by consumers at the time of vehicle purchase, and medium marginal production costs for improving fuel efficiency.

Table 1.11: Market Level Outcomes ($K = 4$)

		Change from Policy 1		
	Policy 1	Policy 2	Policy 3	Policy 4
Flat (FL)/Footprint (FP):	FL	FP	FP	FP
Intra-firm transferring:	No	No	Yes	Yes
Inter-firm trading:	No	No	No	Yes
Fuel efficiency (PC) (gphm)	3.52	-0.0006	0.0778	-0.0170
Fuel efficiency (LDT) (gphm)	4.76	0.0006	-0.0853	0.0178
Fuel efficiency (PC + LDT) (gphm)	4.10	-0.0001	0.0013	-0.0021
Footprint (PC) (sq. ft.)	43.87	0.0015	0.0033	-0.0242
Footprint (LDT) (sq. ft.)	50.90	0.0140	0.0161	-0.0098
Market share (PC) (%)	53.25	0.0086	0.0230	0.1064
Automakers' profits (million \$) [1]	-	8.45	-6.07	-274.05
Consumer surplus (mil. \$) [2]	-	9.48	-37.86	310.68
Technology costs for fuel eff. adjust. (mil. \$) [3]	-	-5.93	97.15	-136.12
Lifetime fuel use externalities (mil. \$) [4]	42,605	5.29	-36.42	-11.36
Social welfare (= [1] + [2] − [4]) (mil. \$) [5]	-	12.64	-7.51	47.99
Automakers' profits / vehicle (\$)	-	0.56	-0.40	-18.11
Consumer surplus / vehicle (\$)	-	0.63	-2.50	20.53
Technology costs for fuel eff. adjust. / vehicle (\$)	-	-0.39	6.42	-9.00
Lifetime fuel use externalities / vehicle (\$)	2,816	0.35	-2.41	-0.75
Social welfare / vehicle (\$)	-	0.84	-0.50	3.17

The table shows simulation results at the market level for the case of $K = 7$. All dollars are in 2001 U.S. dollars.
$K = 7$: high valuation of fuel efficiency by consumers at the time of vehicle purchase, and high marginal production costs for improving fuel efficiency.

Table 1.12: Market Level Outcomes ($K = 7$)

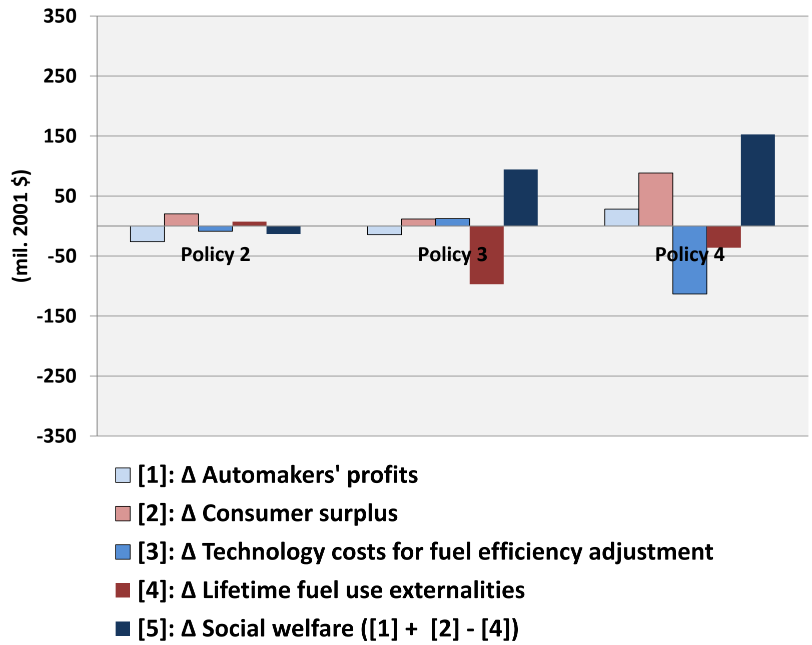

Figure 1.3: Changes from Policy 1 as reported in Table 1.10 ($K = 1$)

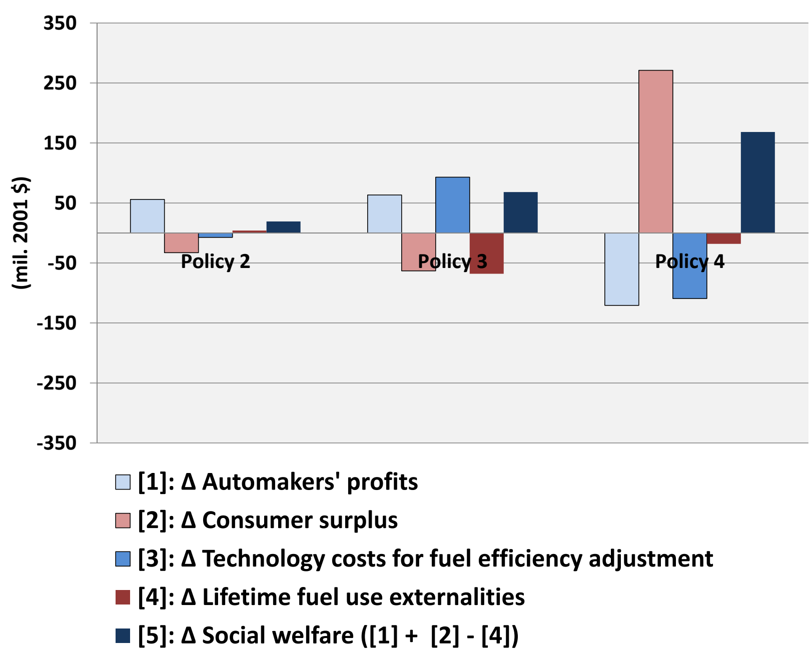

Figure 1.4: Changes from Policy 1 as reported in Table 1.11 ($K = 4$)

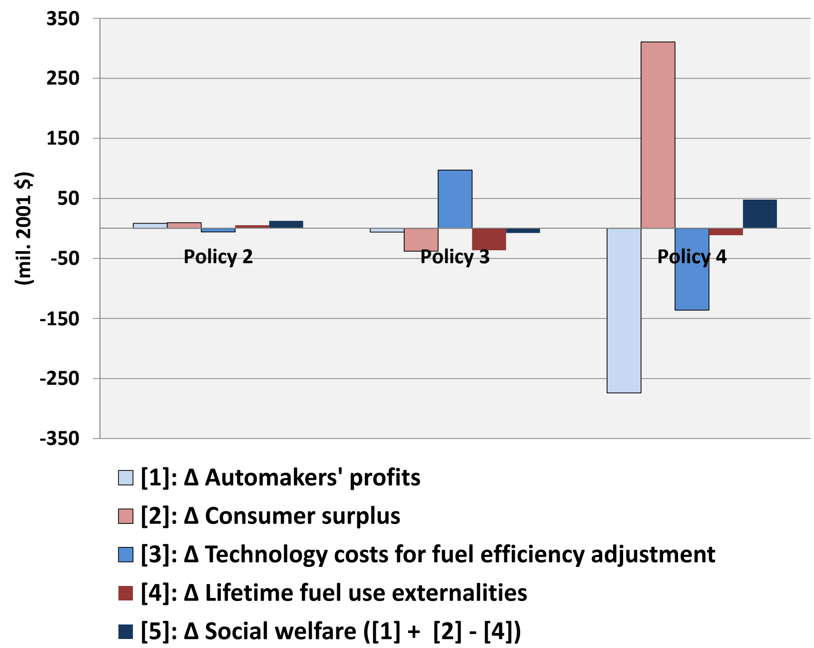

Figure 1.5: Changes from Policy 1 as reported in Table 1.12 ($K = 7$)

	$K = 1$				
	Target (gphm)		Fuel eff. (gphm)		Diff.
Policy:	3	4	3	4	4
	[1]	[2]	[3]	[4]	[2]-[4]
General Motors	4.22	4.16	4.22	4.25	-0.08
Ford	4.33	4.27	4.33	4.51	-0.23
DaimlerChrysler	4.41	4.35	4.41	4.73	-0.38
Toyota	3.99	3.92	3.77	3.40	0.52
Honda	3.83	3.75	3.35	3.02	0.74
Nissan	3.97	3.90	3.97	3.65	0.25
Volkswagen	3.49	3.41	3.49	3.29	0.12
All 11 automakers	4.16	4.10	4.10	4.10	0.00

	$K = 4$				
	Target (gphm)		Fuel eff. (gphm)		Diff.
Policy:	3	4	3	4	4
	[1]	[2]	[3]	[4]	[2]-[4]
General Motors	4.22	4.17	4.22	4.36	-0.19
Ford	4.33	4.28	4.33	4.45	-0.17
DaimlerChrysler	4.41	4.35	4.41	4.65	-0.30
Toyota	3.99	3.92	3.77	3.40	0.52
Honda	3.83	3.75	3.35	3.02	0.74
Nissan	3.97	3.90	3.97	3.65	0.25
Volkswagen	3.49	3.41	3.49	3.29	0.12
All 11 automakers	4.16	4.10	4.10	4.10	0.00

	$K = 7$				
	Target (gphm)		Fuel eff. (gphm)		Diff.
Policy:	3	4	3	4	4
	[1]	[2]	[3]	[4]	[2]-[4]
General Motors	4.22	4.16	4.22	4.32	-0.16
Ford	4.33	4.28	4.33	4.41	-0.14
DaimlerChrysler	4.41	4.35	4.41	4.56	-0.20
Toyota	3.99	3.92	3.77	3.53	0.39
Honda	3.83	3.75	3.35	3.12	0.63
Nissan	3.97	3.90	3.97	3.75	0.15
Volkswagen	3.49	3.41	3.49	3.35	0.06
All 11 automakers	4.16	4.10	4.10	4.10	0.00

The table reports selected automakers' (sales-weighted) average target and fuel efficiency under Policies 3 and 4.

Table 1.13: Automaker Level Outcomes

Chapter 2: Pricing Automobile Fuel Economy: A New Hedonic Approach

2.1 Introduction

A number of papers have used the hedonic method to estimate the relationship between vehicle price and attributes (e.g., Atkinson and Halvorsen, 1984; Ohta and Griliches, 1986; Dreyfus and Viscusi, 1995). Some of these studies include fuel economy as a regressor and analyze the marginal price of fuel economy, i.e., how changes in fuel economy affect vehicle prices, with other attributes held constant. In the standard hedonic approach, the marginal price of fuel economy, or equivalently consumers' marginal willingness to pay (MWTP) for fuel economy is estimated in two steps. First, assuming a reasonable functional form, we regress vehicle price on various vehicle attributes (e.g., weight, horsepower and fuel economy) and obtain the hedonic price function that describes the relationship between vehicle price and attributes. Then, the marginal price of fuel economy is given as the partial derivative of the hedonic price function with respect to fuel economy.

A common problem many previous studies have encountered is that the marginal price of fuel economy (miles per gallon) is often estimated to be only insignificantly positive, or sometimes even negative (e.g., Knittel, 2009; Arguea and Hsiao, 1993; Goodman, 1983; Deaton and Muellbauer, 1980; Hogarty, 1975). This is counter-intuitive because other things equal, a vehicle with better fuel economy, which reduces the owner's fuel cost spending, will be valued more highly and thus should be more expensive. Some studies (e.g. Matas and Raymond, 2009; Espey and Nair, 2005; Murray and Sarantis, 1999) have obtained a statistically significant and correct sign on fuel economy. However, the frequency of obtaining insignificant or wrongly signed estimates casts doubt on the robustness of the standard hedonic approach in analyzing the value of fuel economy.

Previous studies argue that unreasonable estimates result from multicollinearity and/or omitted variables. First, fuel economy is very highly correlated with some vehicle attributes, such as weight and horsepower. Thus, including fuel economy and these variables simultaneously on the of right-hand side of the hedonic price regression may result in multicollinearity and give unstable estimates on fuel economy. Second, there are many vehicle attributes that are difficult to observe and not well represented in the hedonic regression (e.g., interior quality, safety features, reliability). Though these attributes would be technologically independent of fuel economy, we will see in Section 2.2 that they are very likely correlated with fuel economy (and other included variables) through consumer preferences, thus causing omitted variable bias. The results in Section 2.2 imply that omitted variable bias is a more serious problem than multicollinearity.

This paper proposes an alternative approach to estimating how fuel economy is priced in the market and how consumers discount future fuel cost savings. A distinctive aspect of our approach is the use of each vehicle's miles traveled, or how much the vehicle is driven. Based on an optimization problem for consumers' vehicle purchase, we derive an equation relating vehicle miles traveled (VMT) and the marginal price of fuel economy. We use this equation and estimate each vehicle's marginal price of fuel economy and the discount rate for future fuel cost savings. An advantage of our approach is that it can significantly alleviate the omitted variable bias from the attributes unrelated with fuel economy, such as interior quality, safety features, and reliability. This is possible because the equation relating fuel economy and the marginal price of fuel economy is only slightly, if any, affected by the unrelated attributes.

Our approach additionally estimates something previous studies have not estimated: the *total* price of fuel economy. This is the portion of each vehicle's retail price attributable to fuel economy, or how much consumers pay in total for each vehicle's fuel economy (consumers' total willingness to pay for fuel economy). This price is not observed explicitly in the market, although it surely exists. The difference in the total price of fuel economy across vehicles mainly comes from the difference in the cost of achieving different combinations of fuel economy and other attributes that affect fuel economy, such as weight and horsepower.

Theoretically, the key feature of our approach is that we take the steps of the standard hedonic method backward. In the standard approach, using data on vehicle price and attributes, we first estimate the hedonic price function for automobiles by regressing vehicle price on various attributes. Then, the marginal price of an attribute is obtained as the first derivative of the estimated hedonic price function with respect to that attribute. In our approach, we first construct a proxy for each vehicle's marginal price of fuel economy by using data on gasoline prices and the vehicle's estimated annual miles traveled. The logic behind this is as follows. With Rosen's (1974) argument that at a point in the space of product attributes where a transaction occurs, the hedonic price function is tangent to the consumer's bid function, each vehicle's marginal price of fuel economy equals its buyer's MWTP for fuel economy. In turn, his MWTP should equal the present value of expected fuel cost savings over the life of the vehicle due to a marginal fuel economy change. And the present-value expected fuel cost savings depend on gasoline prices and vehicle miles traveled (VMT). Therefore, we can obtain a proxy for each vehicle's marginal price of fuel economy based on gasoline prices and its VMT. We thus observe a proxy for each vehicle's marginal price of fuel economy, in addition to its attributes. Second, we recover the hedonic price function for fuel economy, the function relating the total price of fuel economy with vehicle attributes such as fuel economy and weight, as an envelope of d8ffereht vehicles' (proxied) marginal prices of fuel economy. Therefore, compared to the standard hedonic approach, we are essentially taking the procedure backward. From the estimated hedonic price function, we obtain an estimate of each vehicle's marginal price of fuel economy (which would be more accurate than the proxy variable originally used).

We apply our approach to model year 2001 vehicles sold in the U.S. Information on vehicle attributes such as vehicle weight, horsepower and fuel economy is taken from the U.S. Environmental Protection Agency's fuel economy test data. The National Household Travel Survey provides the make, model, year, and estimated VMT of each vehicle. We match these two data sets and take them to estimation.

Estimation results suggest that our approach works. We estimate that consumers discount future fuel cost savings at the annual rate of 26-43%, much higher than usual rates of return on investment. As for the marginal price of fuel economy or MWTP for fuel economy, a fuel efficiency improvement of 0.1 gallon per 100 miles on average increases vehicle price by $74.7 in 2000 U.S. dollars (for the middle case of the discount rate of 34%). Larger vehicles tend to have higher marginal prices of fuel economy, basically because these vehicles are driven more miles, so buyers of these vehicles are more willing to pay for fuel economy. The average total price of fuel economy is $1,950 (for the case of the discount rate of 34%). Larger vehicles tend to have higher total prices of fuel economy as well, which implies that the cost spent for fuel economy is higher in these vehicles than in smaller vehicles. The estimated total prices of fuel economy suggest that for most vehicles around 5-10% of their retail price is attributable to fuel economy.

The paper is organized as follows. Section 2 applies the standard hedonic approach to our data and analyzes potential problems of the approach. Sections 3 and 4 discuss the theoretical background and framework of the new approach we propose. Section 5 describes the estimation procedure, and Section 6 explains the data sets used. Section 7 reports the results of our estimation. Section 8 checks the robustness of the results. Section 9 concludes.

2.2 Applying the Standard Hedonic Approach to Our Data

To see if the standard hedonic approach works, we run a simple hedonic regression using the data from model year (MY) 2001 vehicles sold in the U.S. (We later apply our new approach to MY 2001 as well.) We estimate a hedonic price function by regressing the retail price (H) of each vehicle on its attributes including fuel economy. The function is estimated with the log-log form:

$$\ln(H_i) = \beta_1 + \beta_2 \ln(e_i) + \beta_3 \ln(w_i) + \beta_4 \ln(a_i) + \boldsymbol{\delta} \cdot (\text{other controls}) + \varepsilon_i, \tag{2.1}$$

where the subscript i indicates vehicle trim i, e_i is fuel economy (measured in gallons per 100 miles), w_i is vehicle weight and a_i is acceleration capacity (horsepower divided by weight). Some other controls are included as discussed below.

Vehicle price (manufacturer's suggested retail price) data is taken from *WARDS Automotive Yearbook*. Vehicle attributes data is from the Environmental Protection Agency's "fuel economy test car list data." We use gasoline-engine vehicles only, so diesel-engine or hybrid-engine vehicles are excluded from the sample. The unit of analysis is at the vehicle trim level (1177 observations).[1] Table 2.1 gives summary statistics. Section 2.6 explains more about the data sources.

Table 2.2 reports estimation results. Columns (1) and (2) are estimated with ordinary least squares (OLS), and columns (3) and (4) are with weighted least squares (WLS), where weights are given by the sales volume of each trim. Dummy vari-

[1] A trim is a subcategory of a model. For example, Toyota Camry CE is one of the trims under the model Toyota Camry.

Variable	Mean	Std. Dev.	Min	Max
Retail price ($) ($H$)	26,555	12,797	9,045	129,595
Fuel economy (gallons/100 miles) (e)	5.24	1.04	2.56	8.16
Weight (lb) (w)	4,156	783	2,250	6,000
Horsepower/Weight (hp/lb) (a)	0.049	0.011	0.031	0.120
Light duty trucks (LDT)	0.625	0.484	0	1
Manual transmission (MT)	0.449	0.498	0	1
All wheel drive (AWD)	0.286	0.452	0	1
Rear wheel drive (RWD)	0.409	0.492	0	1

The number of observations is 1177.

Table 2.1: Summary Statistics

ables are included to control for the vehicle's drive system (FWD/RWD/AWD),[2] transmission type (AT/MT),[3] and light duty truck status (LDT).[4] Columns (2) and (4) additionally include a dummy variable for luxury vehicles (Luxury) and another variable (ABS/TC) that controls for the level of safety features.[5]

Most regressors have significant coefficients with a reasonable sign in all columns. The results suggest that other things equal, increasing weight (w) or acceleration capacity (a) raises vehicle price. Additionally, other things equal, light duty trucks and manual transmission trims are less expensive, while all or rear wheel-drive, the luxury status, and safety features represented by the anti-lock brake system and traction control increase vehicle price. These results are consistent across all columns (and

[2]Front-Wheel-Drive/Rear-Wheel-Drive/All-Wheel-Drive. FWD is treated as the base category.

[3]Automatic Transmission/Manual Transmission. AT is treated as the base category.

[4]Vehicles included in the regression (light duty vehicles) fall into two large categories: passenger cars and light duty trucks. Pickup trucks, sport utility vehicles and vans are in the light duty truck (LDT) category. Other vehicles are in the passenger car category.

[5]The variable "Luxury" takes 1 if the trim is classified as a luxury model in *WARDS Automotive Yearbook*, and 0 otherwise. The variable "ABS/TC" is the sum of two dummy variables: it takes 2 if the trim has the anti-lock brake system and traction control as standard features, 1 if either of them is a standard feature, 0 otherwise.

	(1)	(2)	(3)	(4)
Estimator:	OLS	OLS	WLS	WLS
ln(e)	-0.20*	0.17**	-0.44***	-0.15**
	(0.10)	(0.082)	(0.083)	(0.069)
ln(w)	1.00***	0.52***	1.38***	0.96***
	(0.095)	(0.077)	(0.080)	(0.068)
ln(a)	0.89***	0.55***	0.78***	0.46***
	(0.048)	(0.040)	(0.045)	(0.039)
LDT	-0.27***	-0.11***	-0.12***	-0.056***
	(0.025)	(0.021)	(0.019)	(0.016)
MT	-0.12***	-0.11***	-0.14***	-0.12***
	(0.016)	(0.013)	(0.015)	(0.012)
AWD	0.23***	0.13***	0.16***	0.13***
	(0.025)	(0.020)	(0.021)	(0.017)
RWD	0.14***	0.053***	0.100***	0.058***
	(0.022)	(0.018)	(0.019)	(0.016)
Luxury		0.39***		0.28***
		(0.019)		(0.016)
ABS/TC		0.079***		0.076***
		(0.0099)		(0.0075)
Constant	4.96***	7.11***	1.78***	3.71***
	(0.66)	(0.53)	(0.58)	(0.48)
Observations	1177	1177	1177	1177
R^2	0.635	0.777	0.687	0.792

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2.2: Results of Estimating Equation (2.1)

with most previous studies).

The coefficient on fuel economy (fuel consumption) e is of our primary interest. We expect its coefficient, or the elasticity of retail price with respect to fuel consumption, to be negative, as consumers are less willing to pay for a less fuel efficient vehicle (with a larger e), everything else equal. Comparing columns (1)-(4), we discuss two findings regarding the estimated coefficient of $\ln(e)$.

First, while the estimate has a reasonable negative sign in columns (1), (3) and (4), it is positive and significantly so at the 5% level in column (2). That is, only OLS with Luxury and ABS/TC gives a counter-intuitive sign for the coefficient of fuel consumption. The unreasonable sign in column (2) may be a result of using OLS. With OLS, trims with relatively small sales, which are more likely to be outliers in the sample, have a larger impact on the estimates than with sales-weighted least squares. They may be the main force causing the unreasonable coefficient on $\ln(e)$. Indeed, WLS with the same regressors (column (4)) provides a reasonable coefficient (-0.15). This implies the advantage and importance of sales-weighting in obtaining a reasonable parameter estimate on fuel efficiency, as sales-weighted least squares can prevent the estimation to be driven too much by vehicles with small sales and outlying attributes. Previous hedonic studies of the automobile market mostly use OLS. Table 2.2 suggests that those studies obtaining an unreasonable estimate on fuel efficiency by OLS could have avoided it by applying WLS.

Second and more importantly, including variables for the luxury status and safety features changes the estimates drastically. There are various vehicle attributes that are desirable for consumers and affect the retail price, but are not well represented in columns (1) and (3), such as interior equipment, safety features, comfort and reliability. The variables Luxury and ABS/TC control for these attributes (at least to some extent). Regardless of the estimator used (OLS or WLS), the inclusion of Luxury and ABS/TC significantly reduces the magnitude of most of the other coefficients, and changes the sign of the OLS estimate on $\ln(e)$. This clearly implies strong correlations between attributes in columns (1) or (3) such as fuel economy, weight and horsepower (divided by weight) and attributes represented by Luxury and ABS/TC (e.g., interior equipment, safety features, comfort and reliability). Thus, regressions without variables to control for the latter attributes are most likely biased, and inferring MWTP for fuel economy (or other attributes) based on these regressions is misleading, as it cannot estimate the effect of changing fuel economy (or another attribute) only. Some previous studies (e.g., Arguea, Hsiao, and Taylor, 1994; McManus, 2007) discuss the marginal price of fuel economy based on regressions without these variables, but the results in Table 2.2 imply the possibility of omitted variable bias in the estimates from these studies.

In connection with later sections, it is very important to note that correlation between fuel economy and attributes like interior equipment and safety features does not necessarily mean that production of fuel economy is *technologically* related with production of these attributes. Indeed, it would make more sense to consider that they are technologically independent (if the effect of weight increase from these attributes is accounted for by including weight as a regressor). Even though two attributes are technologically independent, consumer preferences may make them correlated in marketed vehicles. For example, fuel economy and safety features such as the anti-lock brake system and traction control seem technologically independent (once weight is included in the regression). On the other hand, everything else equal, consumers with higher demand for driving would prefer more fuel efficient vehicles, and they may also want more safety features because more driving increases the risk of involving in an accident. If this is the case, other attributes equal, more fuel efficient vehicles will be equipped with more safety features, and we will observe partial correlation between fuel economy and safety attributes. Section discusses technological independence of attributes in more detail.

In addition, Luxury and ABS/TC (and similar variables) may not be able to remove the omitted variable bias well enough. Our Luxury and ABS/TC are simple discrete variables that take 0 or 1 (Luxury) or 0, 1 or 2 (ABS/TC). Indeed, data for attributes like interior equipment, safety features, comfort and reliability is mostly given, if any, in the form of dummy variables, in contrast to those attributes that can be expressed as continuous variables, such as fuel economy, weight and horsepower. These simple discrete variables may not be effective enough to take away the bias due to the strong correlations between attributes like interior equipment, safety features, comfort and reliability, and fuel economy (and others). Thus, the result in columns (2) or (4) may still contain substantial omitted variable bias, and the estimated marginal price of fuel economy may still be misleading as well.

These arguments show the difficulty of separating the effect on retail price of fuel economy from that of such attributes as interior equipment, safety features, comfort and reliability.

2.3 Background: Hedonic Cost Function

For understanding our approach, it will be useful to start from the vehicle production process. Thus, this section briefly considers the hedonic cost function of automobile production and discusses how vehicle attributes are related in the cost function.

A hedonic cost function is the supply-side counterpart of a hedonic price function. It describes the cost of producing a heterogeneous good as a function of its attributes. Consider the following hedonic cost function of producing a vehicle.

$$c = h(e, \mathbf{q}), \tag{2.2}$$

where c is the total cost of producing the vehicle, e is fuel economy (measured in gallons per 100 miles) and $\mathbf{q}$ is a vector of other vehicle attributes.[6]

We decompose the hedonic cost function $h(e, \mathbf{q})$ into the sum of a function that involves e and another function that does not:

$$h(e, \mathbf{q}) = f(e, \mathbf{q_1}) + g(\mathbf{q_1}, \mathbf{q_2}), \tag{2.3}$$

where $\mathbf{q} = [\mathbf{q_1}, \mathbf{q_2}]$.[7]

$f(e, \mathbf{q_1})$ represents the part of c that is technologically related with fuel economy, and we call f the (total) cost of fuel economy. A vector of its arguments $\mathbf{q_1}$

[6] As in most studies using the hedonic cost function approach, we assume constant returns to scale (i.e., production volume) in vehicle production, so that c does not depend on how many vehicles are produced by the firm, plant or production line.

[7] There is no loss of generality in this decomposition because it is possible to set $\mathbf{q} = \mathbf{q_1}$ and $g(\cdot) = 0$ if none of the elements in $\mathbf{q}$ is additively separable from e.

are not (completely) additively separable from e in the cost function, so that the marginal cost of fuel economy improvement ($-\frac{\partial f}{\partial e}$) depends on $\mathbf{q_1}$.

$g(\mathbf{q_1}, \mathbf{q_2})$ represents the other part of c that is technologically independent of e and thus does not vary with e. It has two types of arguments: $\mathbf{q_1}$ and $\mathbf{q_2}$. $\mathbf{q_2}$ does not show up in f, so it is completely additively separable from e. On the other hand, $\mathbf{q_1}$ is also observed in f, so g captures the part of the effect of $\mathbf{q_1}$ on c that is additively separable from e.

An example of attributes in $\mathbf{q_2}$ would be safety devices such as air bags. The cost of marginally improving fuel economy is unlikely affected by whether or not the vehicle is equipped with air bags (once vehicle weight is included in $\mathbf{q_1}$ and thus the possible effect of the weight increase from the air bags is accounted for). Thus, air bags do not enter f. On the other hand, the cost of installing a safety device is unlikely related with the level of fuel economy, once the effect of other attributes in $\mathbf{q}$ is controlled for. Thus, e does not enter g.

Vehicle weight would be a good example of attributes in $\mathbf{q_1}$. Marginally improving fuel economy would be more costly with a vehicle of 4,000 lb and 20 miles per gallon than with a vehicle with 3,500 lb and 20 miles per gallon, as weight negatively affects fuel efficiency. This illustrates the dependence of f on vehicle weight. On the other hand, vehicle weight would also affect g because it simultaneously represents other attributes than weight itself. In virtually all studies (hedonic or discrete choice) that use vehicle weight as a regressor (including the regression in Section 2.2), this variable implicitly represents other attributes as well, such as vehicle size (and sometimes even the level of optional equipment). This multiple representation is necessary because in practice it is impossible to include all vehicle attributes as regressors. If we take vehicle size as an example, heavier vehicles are generally larger, so that they need more materials for producing the body, and have higher material costs. These material costs should be included in g, not f, because fuel economy is affected by weight itself, and what it comes from does not matter. For instance, a weight increase from loading rocks (found somewhere on the road) into the vehicle's trunk has essentially the same effect on fuel economy as an equivalent weight increase due to an increase in vehicle size and material use (if the potential effect of a change in aerodynamic drag is ignored). Thus, it is possible that an attribute in $\mathbf{q_1}$ appears both in f and g. f accounts only for how the attribute changes the cost of providing fuel economy, while g accounts for how it changes the production cost in ways unrelated with fuel economy.

Other than vehicle weight, $\mathbf{q_1}$ would include such attributes as horsepower, torque, body styles, drive systems (front-wheel drive, rear-wheel drive or four-wheel drive) and transmission types (manual or automatic). For each of the elements in $\mathbf{q_1}$, it is possible to give an engineering explanation as to how it is related with fuel efficiency. What is important to note in connection with later sections is that these attributes affects the marginal cost of fuel economy improvement.

2.4 Theoretical Framework

This section discusses the theoretical framework of our approach to estimating the hedonic price function for fuel economy.

Let us consider a consumer who is buying a new automobile. The consumer's choice problem is formulated as:

$$\max_{z,\mathbf{q},m,e} u(z,\mathbf{q},m;\boldsymbol{\omega}) \quad s.t. \quad y = z + H(e,\mathbf{q}) + p \cdot m \cdot e. \tag{2.4}$$

In the utility function $u(\cdot)$, z is consumption of the numéraire good. $\mathbf{q}$ is a vector of vehicle attributes. m is the distance (measured in 100 miles) the consumer expects to travel with the vehicle (expected vehicle miles traveled (VMT)), and $\boldsymbol{\omega}$ is a vector of her characteristics. In the budget constraint, y is her income; $H(\cdot)$ is the hedonic price function for automobiles, one of whose arguments is e, fuel economy measured in gallons per 100 miles; and p is expected gasoline price per gallon. Thus, the term $p \cdot m \cdot e$ is the total fuel cost the consumer expects to pay (\$/gallon × miles × gallons/mile). For the moment, we consider only a single period in the problem. We will later introduce multiple periods to reflect the fact that consumers own vehicles over years and likely consider total fuel cost over the lifetime of their vehicles. Exogenous to the choice problem are $\boldsymbol{\omega}$, y, p, and the shape of $u(\cdot)$ and $H(\cdot)$. The consumer chooses z, $\mathbf{q}$, m and e.

Note that m enters $u(\cdot)$, but not $H(\cdot)$. That is, how much the consumer will drive the vehicle after purchase affects her utility, but not its price. On the other hand, I assume that e enters $H(\cdot)$, but not $u(\cdot)$: The level of fuel economy influences vehicle prices, but good fuel economy does not directly increase her utility on its own. Consumers prefer better fuel economy only because it lowers fuel costs $p \cdot m \cdot e$. Although some consumers who are environmentally conscious may obtain utility directly from owning fuel efficient vehicles, this assumption will be valid for most consumers. It is this property of e that enables the estimation approach we propose.

First order conditions of the consumer's choice problem are:

$$
\begin{aligned}
u_z(z, \mathbf{q}, m; \boldsymbol{\omega}) &= \lambda, \\
u_{q_k}(z, \mathbf{q}, m; \boldsymbol{\omega}) &= \lambda H_{q_k}(e, \mathbf{q}) \quad \forall k, \\
u_m(z, \mathbf{q}, m; \boldsymbol{\omega}) &= \lambda p e, \\
pm &= -H_e(e, \mathbf{q}), \qquad (2.5) \\
y &= z + H(e, \mathbf{q}) + p \cdot m \cdot e,
\end{aligned}
$$

where λ is the Lagrange multiplier for the budget constraint. The rational consumer's vehicle choice should satisfy these equations in equilibrium.

Among these conditions, this study focuses on equation (2.5), which is the first order condition with respect to e. pm is the fuel cost savings from a marginal improvement in fuel economy (i.e., MWTP for fuel economy). Similarly, $-H_e(\mathbf{q}, e)$ is the price increase from a marginal improvement in fuel economy (i.e., marginal price of fuel economy). The equality implies that the marginal willingness to pay and the marginal price are equalized at the optimum.

As in the case of the hedonic cost function $h(e, \mathbf{q})$ above, the hedonic price function $H(e, \mathbf{q})$ can be decomposed into the sum of a function that involves e and another function that does not:

$$H(e, \mathbf{q}) = F(e, \mathbf{q}) + G(\mathbf{q}).^{8} \tag{2.6}$$

How are F and f (or G and g) related with each other? If the automobile market is perfectly competitive, the hedonic price function matches the hedonic cost function, so that $F(e, \mathbf{q}) = F(e, [\mathbf{q_1}, \mathbf{q_2}]) = f(e, \mathbf{q_1})$ and $G(\mathbf{q}) = G([\mathbf{q_1}, \mathbf{q_2}]) = g(\mathbf{q_1}, \mathbf{q_2})$. In reality, the auto market is oligopolistic, so F and f (or G and g) will differ in accordance with automakers' markup-setting strategies. Therefore, it may be possible that F is affected not only by e and $\mathbf{q_1}$, but also by $\mathbf{q_2}$. Still, it would be the case that e and $\mathbf{q_1}$ affect F much more significantly than $\mathbf{q_2}$, because e and $\mathbf{q_1}$ affect F directly through the production technology $f(e, \mathbf{q_1})$, while $\mathbf{q_2}$ does not affect F technologically, but only indirectly through the automaker's pricing strategy.

In the empirics below, we will analyze the effect of $\mathbf{q_2}$ on f by first estimating f without variables representing $\mathbf{q_2}$, and then including them. We will see that including these variables changes the result only slightly, implying small effects of $\mathbf{q_2}$ on f.

With the decomposition of H, equation (2.5) now becomes,

$$pm = -F_e(e, \mathbf{q}). \tag{2.7}$$

Up to this point we have assumed that the consumer uses the vehicle in a single

[8] There is no loss of generality in this decomposition because we can set $G(\mathbf{q}) = 0$ if none of the elements in $\mathbf{q}$ is additively separable from e.

period. In reality, each vehicle is used over years, so the consumer's willingness to pay for a marginal fuel economy improvement will depend on the present value of fuel cost savings over the life of the vehicle. With this consideration, the left-hand side of equation (2.7) is replaced with the the present value of fuel cost savings from a marginal fuel economy improvement, which we model as

$$\sum_{t=0}^{L} d^t \cdot p_t \cdot s^t m, \tag{2.8}$$

where $L + 1$ is the length of the vehicle's life (in years), d is the annual discount factor, p_t is the expected gasoline price at time t, m is the expected VMT for the first year and s is one minus the annual rate of VMT reduction. We assume that consumers expect future gasoline prices to stay at the current level.[9] Then, equation (2.7) is replaced with

$$\frac{1-(ds)^{L+1}}{1-ds} pm = -F_e(e, \mathbf{q}). \tag{2.9}$$

pm is multiplied by a factor ($A \equiv \frac{1-(ds)^{L+1}}{1-ds}$) that accounts for the vehicle's lifetime and the consumer's discounting of future fuel cost savings. We assume that d and s, and thus A, are common to all consumers.

2.5 Estimation Procedure

We will use equation (2.9) to estimate the marginal and total price of fuel economy ($-\frac{\partial F}{\partial e}$ and F) as a function of fuel economy and other vehicle attributes.

[9]The current price is the best predictor if gasoline prices follow a random walk.

Suppose that we have data on $\mathbf{q}$, e and Apm for each vehicle in the data set.Then, we can estimate $F_e(e, \mathbf{q})$ by making reasonable assumptions on its functional form and finding parameter values that best fit the observed market outcomes ($\mathbf{q}$, e and Apm) to equation (2.9).[10] Using the estimated parameters, we calculate the fitted value of the marginal price of fuel economy of each model, or the average of marginal willingness to pay for fuel economy among consumers choosing the model. After estimating $F_e(e, \mathbf{q})$, we use the parameter estimates to recover $F(e, \mathbf{q})$, which gives us the total price of fuel economy, or the portion of the vehicle price attributable to fuel economy, as a function of fuel economy e and other vehicle attributes $\mathbf{q}$.

This procedure contrasts with the standard hedonic approach. In the standard approach, using data on vehicle price and attributes, we first estimate the hedonic price function for automobiles, $H(e, \mathbf{q})$, by regressing vehicle price on various attributes. Then, the marginal price of fuel economy is obtained as the first derivative of the estimated hedonic price function with respect to fuel economy. We have already argued that this approach does not work very well in evaluating the value of fuel economy.

Our approach discussed above proceeds in the opposite order. Fuel economy price, $F(e, \mathbf{q})$, is not explicitly observed. Instead, marginal willingness to pay for fuel economy Apm is observed. Thus, we start from using the first order condition with respect to e (equation (2.9)), and then recover the information on the hedonic price function for fuel economy, $F(e, \mathbf{q})$. Essentially, F is recovered as an envelope of

[10]You might be concerned about the fact that $\mathbf{q}$, e and Apm are endogenous in the sense that consumers choose these values simultaneously. But the hedonic price function (and hence the marginal hedonic price function) is a locus of equilibrium points and not a behavioral function, so simultaneity is not a problem here. See Bockstael and McConnell (2006, p.175) for more details.

numerous consumers' observed marginal willingness to pay for fuel economy. Note that this is possible because the left-hand side of equation (2.9) does not depend on (the derivative of) the utility function $u(\cdot)$ because e does not enter u but appears only in the budget constraint. Likewise, the left-hand side is independent of $z(= y - H - pme)$ and $\boldsymbol{\omega}$, so we need not include individual (or household) income and characteristics in the estimation. In usual hedonic models, the attribute we are interested in (e.g., environmental quality) directly enters the utility function, so that first order condition with respect to that attribute involves (the derivative of) $u(\cdot)$, z and $\boldsymbol{\omega}$. As is widely known, this dependence on $u(\cdot)$, z and $\boldsymbol{\omega}$ significantly complicates the estimation procedure. In this study, we exploit an unusual situation that equation (2.9) is independent of $u(\cdot)$, z and $\boldsymbol{\omega}$ and estimate the unobserved total price of fuel economy $F(e, \mathbf{q})$.

With this method of starting from the first order condition with respect to e, we can extract information only on $F(\cdot)$, the part of vehicle price associated with fuel economy, because $G(\cdot)$ drops off through differentiation and does not show up in equation (2.9). This is impossible with the standard approach, since attributes other than e, especially $\mathbf{q_1}$, affect both F and G, and there is no way separating their effect on F from that on G. If our focus is on fuel economy related issues, $F(e, \mathbf{q})$ includes sufficient information. And information on G is not only redundant but possibly even confusing. Our approach makes it possible to separate F from unwanted G.

In the following estimation, I assume that the hedonic (total) price function for fuel economy $F(\cdot)$ takes the translog form as follows:

$$\begin{aligned}\ln(F_i) =& \beta_1 + \beta_2 \ln(e_i) + \beta_3 \ln(w_i) + \beta_4 \ln(a_i) + \frac{1}{2}\beta_5[\ln(e_i)]^2 + \frac{1}{2}\beta_6[\ln(w_i)]^2 + \frac{1}{2}\beta_7[\ln(a_i)]^2 \\ &+ \beta_8 \ln(e_i)\ln(w_i) + \beta_9 \ln(e_i)\ln(a_i) + \beta_{10} \ln(w_i)\ln(a_i) + \boldsymbol{\delta} \cdot (\text{other controls}),\end{aligned} \tag{2.10}$$

where the subscript i indexes vehicle model i, e_i is fuel economy (gallons per mile), w_i is vehicle weight and a_i is acceleration capacity (horsepower divided by weight). Other controls include dummy variables for drive systems (FWD/RWD/AWD), transmission types (AT/SAT/MT), and vehicle categories (passenger car/light duty truck).

With this specification, equation (2.9) becomes,

$$Ap_i m_i = -\{\beta_2 + \beta_5 \ln(e_i) + \beta_8 \ln(w_i) + \beta_9 \ln(a_i)\}\frac{F_i}{e_i}.^{11} \tag{2.11}$$

Eliminating F_i in equation (2.11) by using equation (2.10), rearranging terms and adding the error ε_i, we have the following equation to be estimated by nonlinear least

[11] Gasoline prices depend on i because different vehicles require different types of gasoline.

squares:

$$\begin{aligned}
Ap_i m_i = & -\{\beta_2 + \beta_5 \ln(e_i) + \beta_8 \ln(w_i) + \beta_9 \ln(a_i)\} \\
& \times \exp\{\beta_1 + (\beta_2 - 1)\ln(e_i) + \beta_3 \ln(w_i) + \beta_4 \ln(a_i) + \frac{1}{2}\beta_5[\ln(e_i)]^2 + \frac{1}{2}\beta_6[\ln(w_i)]^2 \\
& + \frac{1}{2}\beta_7[\ln(a_i)]^2 + \beta_8 \ln(e_i)\ln(w_i) + \beta_9 \ln(e_i)\ln(a_i) + \beta_{10}\ln(w_i)\ln(a_i) \\
& + \boldsymbol{\delta} \cdot (\text{other controls})\} + \varepsilon_i.
\end{aligned} \tag{2.12}$$

As explained in detail below, we construct (a proxy for) m_i from a survey data set. That is, m_i is the sample average VMT of all model i vehicles in the data set. The frequency observed in the survey differs across i. For example, a model with a relatively large market share will be observed more frequently in the survey. m_i is likely to be more accurate (i.e., close to the population average VMT of model i) if it is based on more observations. Therefore, the difference in frequency across i should be reflected in the heteroskedasticity of the error term ε_i. To take account of this, we estimate equation (2.12) with weighted nonlinear least squares, where weight comes from each model's frequency of observations in the sample.

The multiplicative factor A, which accounts for the length of vehicle life and the consumer's discounting of future fuel cost savings, affects only β_1 and does not change the relative magnitude of the marginal price of fuel economy across different vehicle models. At first, we ignore A and estimate the relative marginal price $(-\tilde{F}_e(e, \mathbf{q}) \equiv$

$-F_e(e, \mathbf{q})/A)$ by regressing just pm on e and $\mathbf{q}$. That is, we estimate

$$
\begin{aligned}
p_i m_i = &- \{\beta_2 + \beta_5 \ln(e_i) + \beta_8 \ln(w_i) + \beta_9 \ln(a_i)\} \\
&\times \exp\{\tilde{\beta}_1 + (\beta_2 - 1)\ln(e_i) + \beta_3 \ln(w_i) + \beta_4 \ln(a_i) + \frac{1}{2}\beta_5[\ln(e_i)]^2 + \frac{1}{2}\beta_6[\ln(w_i)]^2 \\
&+ \frac{1}{2}\beta_7[\ln(a_i)]^2 + \beta_8 \ln(e_i)\ln(w_i) + \beta_9 \ln(e_i)\ln(a_i) + \beta_{10}\ln(w_i)\ln(a_i) \\
&+ \boldsymbol{\delta} \cdot (\text{other controls})\} + \tilde{\varepsilon}_i,
\end{aligned}
\tag{2.13}
$$

where $\tilde{\beta}_1 = \beta_1 - \ln(A)$ and $\tilde{\varepsilon}_i = \varepsilon_i/A$. In other words, estimating (2.13) gives the marginal value of fuel economy in case consumers take only the fuel costs of the first year into account.

After estimating equation (2.13), we estimate the multiplicative factor A and the discount factor d by combining the predicted relative marginal price of fuel economy and findings from National Research Council [NRC] (2002) and Environmental Protection Agency [EPA] (2009). Based on engineering estimates of benefits and costs of various fuel efficient technologies summarized in NRC (2002), and market penetration rates of different technologies given in EPA (2009), we estimate the marginal price of fuel economy for the average vehicle. Then, we estimate A, by dividing this engineering-based estimate of the average vehicle's marginal price of fuel economy by the average vehicle's predicted relative marginal price of fuel economy ($-\hat{\tilde{F}}_e(\bar{e}, \bar{\mathbf{q}})$, where $\bar{e}$ and $\bar{\mathbf{q}}$ are sales-weighted average values). Plugging this $\hat{A}$, along with parameter values of s and L, into the relation $A = \frac{1-(ds)^{L+1}}{1-ds}$, we obtain an estimate of the discount factor d.

Finally, we estimate the absolute magnitude of model i's marginal price of fuel economy by

$$-\hat{F}_e(e_i, \mathbf{q}_i) = -\hat{\tilde{F}}_e(e_i, \mathbf{q}_i)\hat{A}. \tag{2.14}$$

Similarly, using $\hat{A}$ and estimates from equation (2.13), we can estimate the absolute magnitude of model i's total price of fuel economy, $\hat{F}(e_i, \mathbf{q}_i)$.

2.6 Data

We apply the above model to model year (MY) 2001 gasoline-engine passenger cars and light duty trucks marketed in the United States.[12] We need data on vehicle attributes (e and $\mathbf{q}$), vehicle owners' expectation on vehicle use (m, s and L) and gasoline prices (p). Lastly, in order to estimate A, we will use engineering-based estimates of the marginal price of fuel economy.

Though in the theoretical model m is the distance each vehicle is expected to be driven for the first year, no data is available on the owner's expected VMT. Therefore, we use a surveyed vehicle's estimated VMT as a proxy for m, which is obtained from the 2001 National Household Travel Survey (NHTS), a national survey conducted by the Federal Highway Administration. This survey contains information on the vehicle(s) each surveyed household owns (such as make, model and model year) and the estimated annual VMT of the vehicle(s). We will look at only MY 2001 vehicles. Two types of VMT estimates will be used in the estimation. The first is self-reported

[12] Non-gasoline engine vehicles such as hybrid vehicles, which were introduced to the U.S. market at around that time, and diesel engine vehicles are excluded from the sample. Flex fuel vehicles are included.

Variable	Mean	Std. Dev.	Min	Max
Self-reported VMT	14,942	9,493	500	76,000
Odometer-based VMT	14,660	8,805	45	73,452
Fuel economy (gallons/100 miles)	4.97	1.02	2.85	7.55
Weight (lb)	4,056	781	2,375	6,000
Horsepower/Weight (hp/lb)	0.048	0.006	0.033	0.099
LDT	0.518	0.499	0	1
MT	0.086	0.138	0	1
SAT	0.003	0.031	0	0.483
AWD	0.237	0.294	0	1
RWD	0.238	0.319	0	1

The number of observations is 1616.

Table 2.3: Summary Statistics

VMT (SVMT) (Table 2.3, row 1), which is based on the owner's recollection on his vehicle's annual VMT. The second is odometer-based VMT (OVMT) (Table 2.3, row 2), which is derived from two odometer readings of the same vehicle on two different dates (usually a few months apart from each other). Like the approach taken in the NHTS, we exclude from the sample 78 observations (passenger cars or light trucks) whose two VMT measures are extremely different. This leaves us 1,616 observations.[13] Since vehicles are identified only up to the vehicle model level in the 2001 NHTS and equation (2.13) is estimated at the this level, we calculate the average VMT by vehicle model.

For vehicle attributes (e and $\mathbf{q}$), we will use the Environmental Protection Agency (EPA)'s "Fuel Economy Test Car List Data" for MY 2001. This is the original fuel economy test data administered by the EPA and is used to determine the fuel

[13] Observations are excluded if $|\text{SVMT-OVMT}| > 10,000$ miles and (SVMT$>$ 4$\times$OVMT or SVMT$<$ 0.25$\times$OVMT).

economy label values available to consumers. The data set covers more than 1,000 vehicle configurations and provides information on vehicle attributes associated with fuel economy or vehicle emissions (e.g., vehicle weight, engine characteristics, transmissions, drive systems, emission control systems), and test results (fuel economy values and pollutant emissions). Table 1 summarizes variables from this data set (rows 3-8). In the following regressions, the data is aggregated to the vehicle model level, as vehicles are identified only up to this level in the 2001 NHTS. Model level average values are calculated by using sales at the configuration level as weights.

Gasoline price data are are obtained from the Energy Information Administration. The annual average price for year 2000 is used as a proxy for the expected future gasoline price p that consumers held when purchasing a MY 2001 vehicle.[14] We assume homogeneous consumers, so that all consumers are assumed to have the same expectation about future gasoline prices. For vehicles requiring premium gasoline, its annual average price ($1.639) is used. For others, the average regular gasoline price ($1.462) is used.

An engineering estimate of the marginal price of fuel economy for the average vehicle (this estimate is denoted by K) is constructed with findings from National Research Council [NRC] (2002) and Environmental Protection Agency [EPA] (2009). Based on meetings and interviews with representatives of automotive manufacturers and component and subsystem suppliers, and through published engineering studies, NRC (2002) estimates the rate of fuel economy improvement and the incremental retail price from (separately) applying various fuel efficient technologies. EPA (2009)

[14]Generally, MY 2001 vehicles began to be in the market in late (calendar year) 2000.

gives statistics on market penetration rates of these technologies over time. As explained in detail in Appendix A, based on these studies we estimate that K is \$35-\$45.

In order to estimate the discount factor d, we need to assume parameter values of s, one minus the annual rate of VMT reduction, and $L+1$, the length of vehicle life. Following NRC (2002), we take $r = 0.955$. We use three values of $L+1$: 10, 15 and 20. We will see that for L large enough (e.g., $L \geq 9$), changing L affects d only slightly.

2.7 Results

2.7.1 Estimating the Relative Magnitude of Marginal Prices of Fuel Economy

Table 2.4 reports the result of weighted nonlinear least squares estimation of equation (2.13) for different dependent variables and specifications. Regression (1) is our base model, and regressions (2)-(4) will be discussed later for robustness checks.

Regression (1) uses self-reported VMT (SVMT) to construct the dependent variable. It includes dummy variables that seem technologically related with fuel economy as discussed in Section 2.3: the light duty truck status, transmission types and drive systems. Several coefficients are statistically significant (although with the translog specification, coefficients are often jointly significant even when they are not individually). In particular, most coefficients of variables related with fuel economy (e) and vehicle weight (w) are statistically significant at the 5 or 10% level (β_2, β_3,

β_5, β_6, β_8 and β_9).

Rather than analyzing the relative marginal price implied by Table 2.4 in detail, we will first estimate the multiplicative factor A and the discount factor d. Using $\hat{A}$, we will estimate the absolute marginal price of fuel economy and analyze how it differs across vehicles.

2.7.2 Estimating the Multiplicative Factor and Discount Rate

For the (hypothetical) vehicle with the sales-weighted average attributes, we have two estimates of the incremental price of a 1% fuel economy improvement. Our model, using VMT and vehicle attributes, gives $A[-\hat{\tilde{F}}_e(\bar{e}, \bar{\mathbf{q}})]\frac{\bar{e}}{100}$, where $-\hat{\tilde{F}}_e(\cdot)$ is the predicted relative marginal price of fuel economy and $\bar{e}$ and $\bar{\mathbf{q}}$ are sales-weighted average values. The engineering estimate (K) based on NRC (2002) and EPA (2009) is \$35-\$45, as explained in detail in Appendix A. We estimate the multiplicative factor A by equating these two estimates and solving for A. With the engineering-based estimate of \$35-\$45, we obtain A for three different values of K (35, 40 and 45). The discount factor d is then estimated using the relationship $A = \frac{1-(ds)^{L+1}}{1-ds}$, and the discount rate r is given by $r = \frac{1}{d} - 1$.

Table 2.5 shows the estimates of A, d and r for different values of the average vehicle's marginal price of fuel economy (K) and the length of vehicle life ($L + 1$), with $s = 0.955$. Table 2.5 also shows WP, the ratio of willingness to pay (WTP) for a marginal fuel economy improvement (or the marginal price of fuel economy) to the expected present value (PV) of fuel cost savings from the same improvement. WTP is given by $\frac{1-(ds)^{L+1}}{1-ds}pm$, and PV by $\frac{1-(vs)^{L+1}}{1-vs}pm$, where v is the real interest rate, which we assume to be 7%. Thus,

$$WP \equiv WTP/PV = \frac{1-(ds)^{L+1}}{1-ds} \Big/ \frac{1-(vs)^{L+1}}{1-vs}. \tag{2.15}$$

WTP/PV$<$ 1 implies consumers' undervaluation of fuel economy, while WTP/PV$>$ 1 implies overvaluation. The lower WTP/PV($<$ 1), the larger the degree of consumers' undervaluation is.

The multiplicative factor A is estimated around 3-4, so for a marginal fuel economy improvement, consumers are on average willing to pay three to four times more than the fuel cost savings from that improvement rewarded in the first year.

The estimates of the discount rate r and the WTP/PV ratio imply that consumers discount future fuel cost savings very fast. The discount rate r is estimated to be 26-43%, depending on K and L. This range is much higher than the real interest rate we use (7%). We also find that the choice of L (the length of vehicle life) has almost no effect on r (or d). Future fuel cost savings are discounted so fast that savings realized after 10 years or later have little value. The WTP/PV ratio is around 0.35-0.60. Consumers are willing to pay for only 35-60% of the present value of total fuel cost savings over the life of the vehicle. These numbers are consistent with the so-called Energy Paradox that consumers significantly undervalue long-term energy cost savings from energy efficiency improvements.

2.7.3 Estimating the Absolute Magnitude of Marginal Prices of Fuel Economy

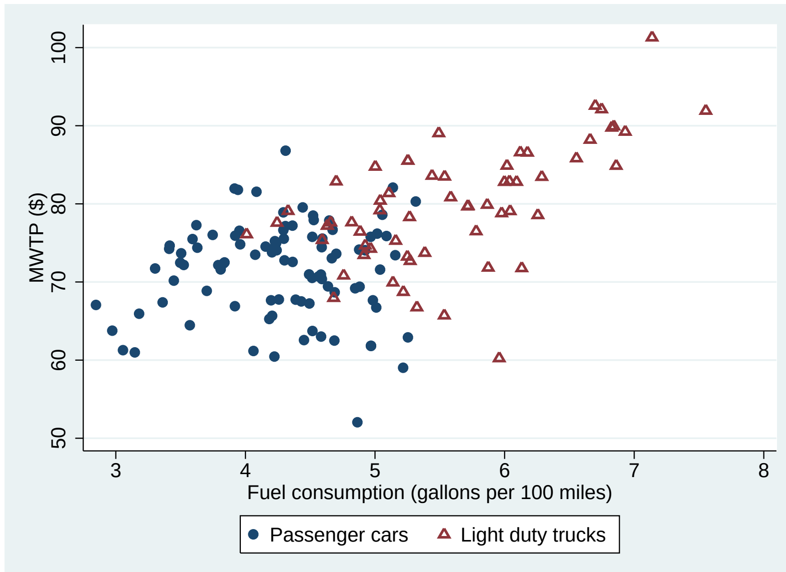

Figure 2.1: Marginal Willingness to Pay for Fuel Economy and Fuel Economy

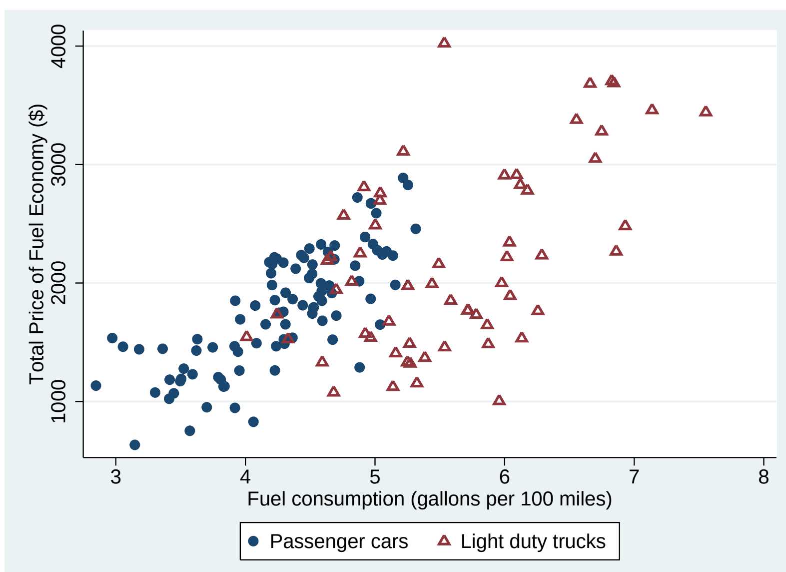

Figure 2.2: Total Price of Fuel Economy and Fuel Economy

With the predicted relative marginal price of fuel economy $(-\hat{\tilde{F}}_e(e, \mathbf{q}))$ derived from column (1) of Table 2.4 and an estimate of the multiplicative factor $\hat{A}$ at

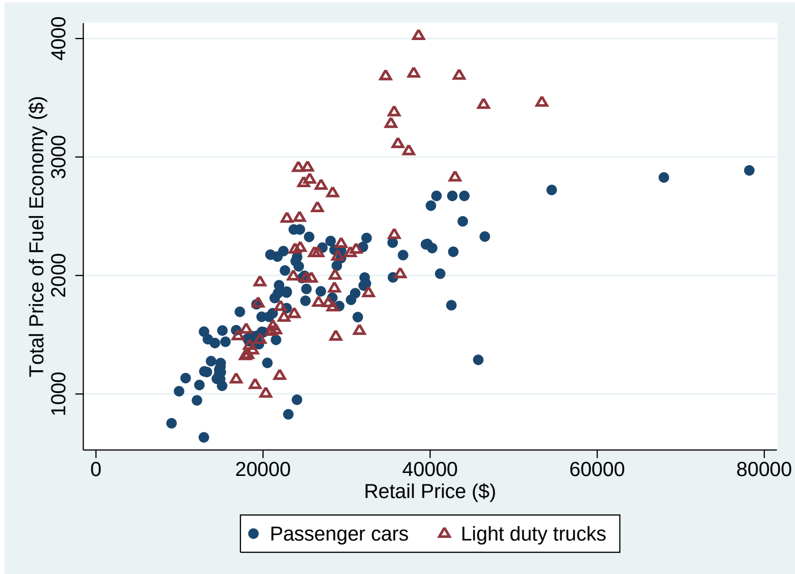

Figure 2.3: Total Price of Fuel Economy and Retail Price

hand, we now estimate the absolute magnitude of each model's marginal price of fuel economy by multiplying them together. In the following we show the results based on $\hat{A}$ from the middle case of $K = 40$ ($\hat{A} = 3.46$). The (model-level, unweighted) average price of fuel economy improvement of 0.1 gallon per 100 miles is estimated to be \$74.7 (in 2000 U.S. dollars), and the standard deviation of \$7.8. In other words, on average, consumers are willing to pay \$74.7 for an improvement of 0.1 gallon per 100 miles. Figure 2.1 plots the estimated marginal price ($-\hat{F}_e$) against fuel economy (gallons per 100 miles) to roughly see what models tend to face larger marginal willingness to pay for fuel economy and thus have larger marginal prices of fuel economy due to higher VMT or the use of premium gasoline. Figure 2.1 shows that less fuel efficient ($\approx$ larger and heavier) vehicles tend to have higher marginal prices. The average marginal price for passenger cars only is \$71.5, while that for light duty trucks is \$79.8.

The total price of fuel economy (F) can be recovered using equation (2.10), the estimated coefficients in column (1) of Table 2.4 and $\hat{A} = 3.46$. Column (1) of Table 2.6 reports selected percentiles, the mean and standard deviation of the total price of fuel economy (F) predicted by the base regression (column (1) of Table 2.4), treating a vehicle model (e.g., Toyota Camry) as one observation. The values in parentheses are the standard errors of $\hat{F}$ for the corresponding vehicle models that result from randomness in the estimated coefficients of Table 2.4. The (model-level, unweighted) average is estimated to be \$1,942 (in 2000 U.S. dollars), and the standard deviation is \$642. As in the case of $-\hat{F}_e$ above, Figure 2.2 plots $\hat{F}$ against fuel economy (gallons per 100 miles) to roughly see what vehicles likely have larger F. Generally, less fuel efficient (≈ larger and heavier) vehicles tend to have larger F. That is, they tend be priced higher for fuel economy.

Figure 2.3 plots the relationship between the estimated total price of fuel economy $\hat{F}$ and the retail price. It shows a strong positive correlation between the two prices. This is interesting because we do not use any information on retail prices in estimating $\hat{F}$. Generally, 5-10% of the retail price is estimated to be attributable to fuel economy. Moreover, while light duty trucks show a relatively proportional relationship between the two prices, passenger cars present a nonlinear relationship in the sense that expensive luxury cars do not have proportionally high total prices of fuel economy. This suggests that luxury and non-luxury cars do not differ so much in terms of F, so that retail price differences between them mostly come from differences in the portion unrelated with fuel economy, G.

2.8 Robustness

We check the robustness of the above results in two respects. First, we estimate the same regression by using a different VMT measure. Second, we check whether the inclusion of variables that are technologically unrelated with fuel economy drastically changes the results.

2.8.1 Using Odometer-based VMT

We construct the dependent variable from odometer-based, rather than self-reported, VMT and do the same procedure as in the last section. Figure 2.4 compares self-reported VMT (SVMT) and odometer-based VMT (OVMT) for the 1,616 vehicles in our sample. Clearly, the two measures are positively correlated (the correlation coefficient is 0.66), but for many observations there is a wide disparity between them. Therefore, estimating with OVMT will provide a robustness check of the above results with SVMT.

Column (3) of Table 2.4 shows the result of estimating equation (2.13) with using OVMT to construct the dependent variable. This regression is comparable to column (1). Obviously, OVMT gives less precise estimates than SVMT. None of the coefficients of regressors associated with e or w, which are mostly significant with SVMT, are significant at the 10% level. (Their p-values are mostly between 0.1 and 0.2.) This difference implies that SVMT is a better proxy for (the first year's) expected VMT at the time of purchase (i.e., m in the theoretical model of Section 2.4).

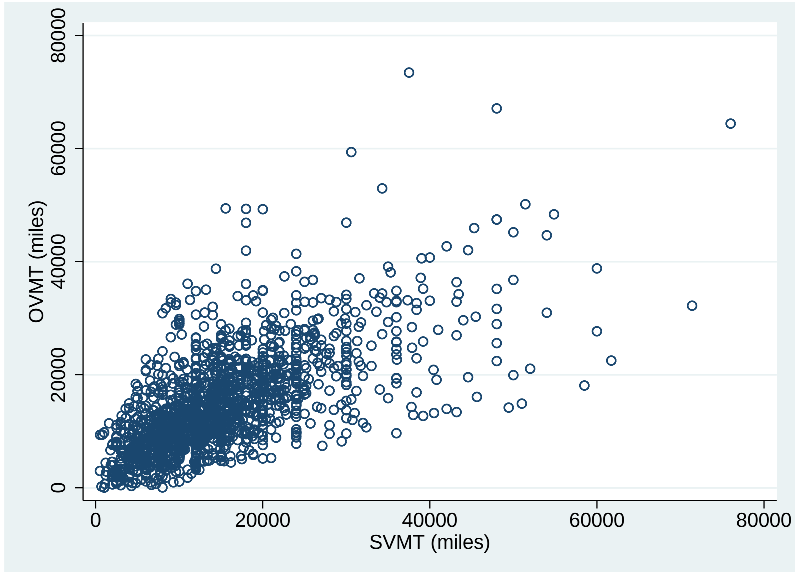

Figure 2.4: Self-reported VMT (SVMT) and Odometer-based VMT (OVMT)

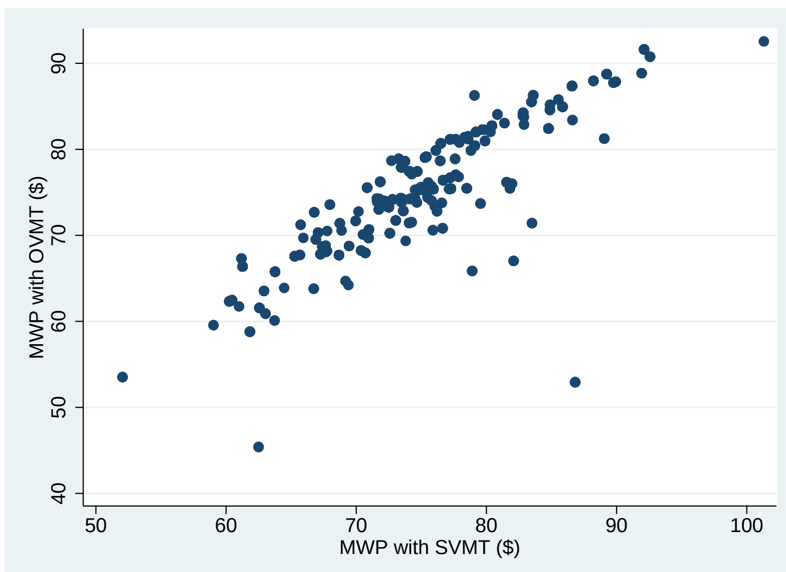

Figure 2.5: Marginal Willingness to Pay for Fuel Economy with SVMT and OVMT (2000 U.S. $)

Though less precise, OVMT-based predicted values and estimates are in many cases very close to SVMT-based predicted values and estimates. First, OVMT estimates $\hat{A} = 3.56$ and $\hat{d} = 0.76$ for $K = 40$ and $L+1 = 14$, while SVMT gives $\hat{A} = 3.46$ and $\hat{d} = 0.75$. Using these estimates of A, Figure 2.5 compares the OVMT-based

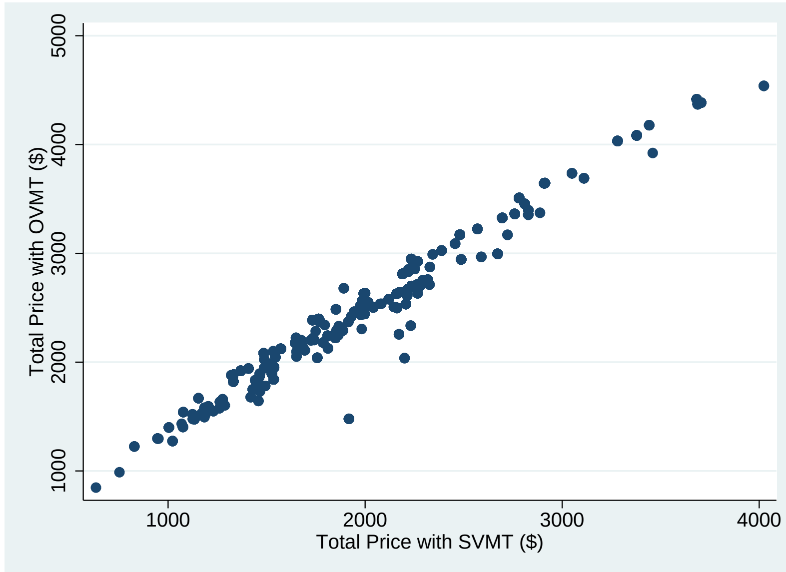

Figure 2.6: Total Price of Fuel Economy with SVMT and OVMT (2000 U.S. $)

predicted marginal price of fuel economy and the SVMT-based marginal price. It shows that for most vehicles these two predictions are similar. The (model-level, unweighted) average from OVMT is \$74.6, while that from SVMT is \$74.7. We also compare the total price of fuel economy predicted by OVMT and SVMT in Table 2.6 and Figure 2.6. Each column of Table 2.6 reports selected percentiles, the mean and standard deviation of the total price of fuel economy (F) calculated from the estimates in the corresponding column of Table 2.4. OVMT gives larger $\hat{F}$ than SVMT, as well as larger standard errors (given in parentheses). Yet, Figure 2.6, which plots OVMT-based $\hat{F}$ against SVMT-based $\hat{F}$, shows that they are very highly correlated.

These comparisons of OVMT-based and SVMT-based estimates confirm the robustness of the results discussed in the last section, and also suggest that SVMT is a better measure to use in estimating how fuel economy is valued in the market.

2.8.2 Including Variables Technologically Unrelated with Fuel Economy

For the standard approach discussed in Section 2.2, we see that the result is very sensitive to the inclusion of those variables that seem to be technologically unrelated with fuel economy, such as the luxury status and safety features. Table 2.2 shows that including these regressors drastically changes the coefficients of $\ln(e)$, $\ln(w)$ and $\ln(a)$.

For our approach, the theoretical model suggests that the variables technologically independent of fuel economy ($\mathbf{q_2}$) unlikely have a large effect on the results based on equation (2.9). This is because $\mathbf{q_2}$ does not affect F through production technology, but only indirectly, if any, through the automaker's pricing strategy. We test whether our approach is also sensitive to attributes in $\mathbf{q_2}$ by additionally including variables for the luxury status (Luxury) and safety features (ABS/TC).

Columns (2) and (4) of Table 2.4 show the result of estimating equation (2.13) with Luxury and ABS/TC. Column (2) uses SVMT and column (4) uses OVMT for constructing the dependent variable. While Luxury is estimated to have a significantly negative effect,[15] other coefficients do not change significantly from columns (1) or (3). This makes a clear contrast with the standard approach in Table 2.2.

Figure 2.7 plots the predicted total price of fuel economy $\hat{F}$ calculated from column (1) of Table 2.4 against $\hat{F}$ from column (2). Also, columns (2) and (4) of

[15] The negative coefficient of Luxury implies that other things equal, a buyer of a luxury vehicle is on average less willing to pay for a marginal improvement of the vehicle's fuel economy basically because it is (on average) driven less than a non-luxury vehicle.

Table 2.6 show some statistics of $\hat{F}$ obtained from the corresponding columns of Table 2.4. From Figure 2.7 and Table 2.6, we find that $\hat{F}$ is robust to including Luxury and ABS/TC.

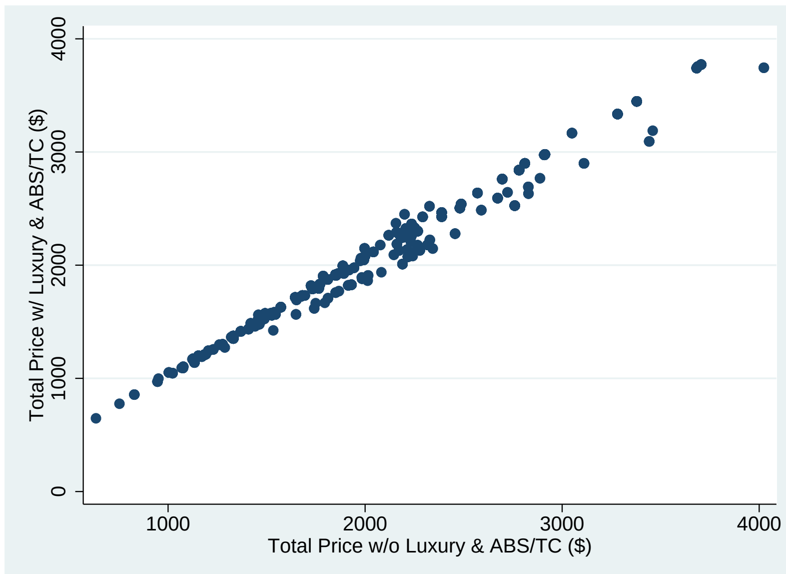

Figure 2.7: Total Price of Fuel Economy with and without Luxury and ABS/TC (2000 U.S. $)

These observations provide support for the validity of our approach. Unlike the estimates from the standard approach, the results in Section 2.7 are insensitive to including Luxury and ABS/TC, representatives of attributes that are technologically independent of fuel economy (i.e., $\mathbf{q_2}$). This insensitivity is consistent with our theoretical framework and suggests that our approach has succeeded in separately analyzing the value of fuel economy (the "F" part) without the complication from the price of other attributes (the "G" part).

2.9 Conclusion

This paper has proposed an alternative hedonic approach to estimating how fuel economy is valued in the market. The basic idea is that we can observe a proxy for a consumer's marginal willingness to pay for her vehicle's fuel economy by using its VMT (and gasoline prices). With the theoretical prediction that marginal willingness to pay equals marginal price at an optimum, this is equivalent to observing each vehicle's marginal price of fuel economy with error. By taking the steps of the standard hedonic method backward, we estimate the marginal and total price of fuel economy as a function of vehicle attributes such as fuel economy and weight.

An important advantage of our approach over the standard hedonic approach is that ours is much less likely affected by omitted variables bias from attributes technologically unrelated with fuel economy such as interior quality and safety features (i.e., attributes denoted by $\mathbf{q_2}$). As discussed in Section 2.2, fuel economy is strongly correlated with $\mathbf{q_2}$, largely due to consumer preferences (though not due to production technology). Thus, omitting these attributes in the standard hedonic regression of vehicle price on vehicle attributes, which usually happens because it is difficult to represent these attributes well enough in regressions, results in a biased estimate of marginal price of fuel economy. Our approach makes it possible to separate the portion of vehicle price that varies with fuel economy (F) from the portion that does not (G). This is because in our approach we first estimate $\frac{\partial F}{\partial e}$, and then recover F using information on $\frac{\partial F}{\partial e}$. The effect of a vehicle attribute on the fuel-economy-unrelated portion of vehicle price (G) can be separated because it does not affect $\frac{\partial F}{\partial e}$. We have shown in Section 2.8 that the estimates and predictions from our approach are robust to including variables that represent $\mathbf{q_2}$.

We have applied this procedure to MY 2001 new vehicles sold in the U.S. With additionally using engineering-based estimates of the average MY 2001 vehicle's marginal price of fuel economy, we estimate that consumers discount future fuel cost savings at the annual rate of 26-43%, much higher than usual rates of return on investment. A fuel efficiency improvement of 0.1 gallon per 100 miles is estimated to increase vehicle price by, on average, \$74.7 in 2000 U.S. dollars (for the middle case of the discount rate of 34%). Larger vehicles tend to have higher marginal prices of fuel economy, basically because these vehicles are driven more miles, so buyers of these vehicles are more willing to pay for fuel economy. The average total price of fuel economy is estimated at \$1,950 (for the case of the discount rate of 34%). Larger vehicles tend to have higher total prices of fuel economy as well, which implies that the cost spent for fuel economy is higher in these vehicles than in smaller vehicles. The estimated total prices of fuel economy suggest that for most vehicles around 5-10% of their retail price is attributable to fuel economy.

.1 Appendix on Engineering Estimates of the Marginal Price of Fuel Economy

Here we analyze two reports on automobile fuel economy (National Research Council, 2002; Environmental Protection Agency, 2009) and estimate the marginal price of fuel economy for the average vehicle from an engineering, rather than economic, point of view. We will consider a number of fuel efficient technologies that were available or expected to be available in 2001. Among these technologies we want to identify "marginal" technologies at that time which automakers chose to or not to use in their models for a marginal fuel economy adjustment. Then, engineering estimates of the benefit and cost of the "marginal" technologies allow us to estimate the marginal price of fuel economy for the average vehicle.

Various technologies are available to improve fuel economy. For example, NRC (2002) gives a comprehensive list of new technologies, and also engineering estimates of fuel economy gains and retail price increases from applying each technology.[16] These estimates are based on engineering, rather than economics, in the sense that they are constructed through meetings and interviews with representatives of automotive manufacturers and component and subsystem suppliers, and through published engineering studies. Also, EPA (2009) provides data on market penetration

[16] NRC's (2002) list of fuel efficient technologies includes: engine friction reduction; low friction lubricants; multi-valve, overhead camshafts; variable valve timing; variable valve timing and lift; cylinder deactivation;engine accessory improvement; supercharging and downsizing; five-speed (or six-speed) automatic transmissions; continuously variable transmissions; aerodynamic drag reduction; improved rolling resistance; intake valve throttling; camless valve actuation; variable compression ratio; automated shift manual transmissions; integrated starter generators; 42 volt electrical systems; electric power steering.

trends of a number of old and new fuel efficient technologies.

Among these technologies, we want to identify "marginal" technologies for model year (MY) 2001 vehicles. Automakers can adjust the level of fuel economy of a vehicle by choosing a combination of fuel efficient technologies used in the vehicle. Marginal technologies are those that automakers likely add to or remove from the combination when they marginally adjust the vehicle's fuel economy. Efficiency gains and price increases from these marginal technologies determine the marginal price of fuel economy.

What kind of technologies are marginal? Relatively old technologies (e.g. front-wheel-drive, port fuel injection and lockup transmissions) are so mature and common that their penetration rates have been very high and stable since the mid-1990s at the latest (EPA, 2009). They have become the default settings and it is unlikely that these technologies are added to or removed from a vehicle for the purpose of marginal fuel economy adjustment. On the other hand, new technologies that were rarely observed in the market in 2001 could not be marginal technologies, either. New technologies will still be more costly or unstable. It is unlikely that these new technologies were applied to a wide variety of vehicles for marginally improving fuel economy. Thus, marginal technologies are those that were used in not a few vehicles in 2001, but not so commonly that they had become default settings. For vehicles with these technologies, they were one of the last fuel efficient technologies applied. For vehicles without these technologies, they would be used if a marginal fuel economy improvement was needed.

Based on EPA (2009), we consider three technologies (multi-valve, overhead camshaft valve trains; variable valve timing; and five-speed automatic transmissions) as "marginal" in MY 2001. Statistics from EPA (2009) shows that 49% of new cars and trucks in MY 2001 were equipped with a multi-valve, overhead camshaft valve train; 20% were with a variable valve timing system, which is generally added to models with a multi-valve, overhead camshaft valve train; 11% (of MY 2001 new cars and trucks with an automatic transmission) were with a five-speed automatic transmission. Other technologies discussed in NRC (2002) were rarely observed in the market at that time, or we do not have data to show how widely they were used. Therefore, at the time of MY 2001, multi-valve, overhead camshaft valve trains, variable valve timing, and five-speed automatic transmissions were likely to be among those technologies that were applied to or removed from each model for marginally adjusting fuel economy.

Table 7 shows engineering estimates of efficiency gains and price increases from the marginal technologies, provided by NRC (2002). The first two columns give estimated ranges of the rate of fuel economy improvement and of the incremental retail price from applying each of the three marginal technologies. Remember that these estimates are based on engineering, rather than economics, in the sense that they are constructed through meetings and interviews with representatives of automotive manufacturers and component and subsystem suppliers, and through published engineering studies. NRC (2002) gives estimated ranges because the effect of each technology differs across vehicle models, depending on various factors, especially vehicle attributes. For example, it may be the case that variable valve timing is more effective and less costly for smaller cars. Unfortunately, further information is not available on how the rate of fuel economy improvement and the incremental retail price are related with vehicle attributes, so we cannot observe engineering estimates for each model. Alternatively, we assume that the effect of each technology on the (hypothetical) vehicle with the average attributes is well approximated by the midpoints of the ranges.[17] Dividing the average incremental price by the average rate of fuel economy improvement, column (3) of Table 7 gives our engineering estimate of the price of a 1% fuel economy improvement for the average vehicle. The marginal price estimates from the three technologies should not be so different from one another. Indeed, column (3) shows that they stay in a $10 range ($35-$45), providing support for our approach. Based on these results, we use $35-$45 as our (interval) estimate of the marginal price of fuel economy for the average MY 2001 vehicle.

[17]NRC (2002) also uses the midpoints of the ranges in deriving its estimates.

	(1)	(2)	(3)	(4)
VMT type:	SVMT	SVMT	OVMT	OVMT
$\tilde{\beta}_1$	-128.7	-122.2	-81.0	-69.0
	(89.0)	(90.9)	(88.0)	(90.0)
β_2	-39.5**	-37.5*	-26.3	-23.3
	(19.4)	(19.7)	(18.6)	(18.7)
β_3	41.5*	40.3*	26.2	24.0
	(23.3)	(23.7)	(22.5)	(22.9)
β_4	11.4	13.4	6.27	9.50
	(11.6)	(11.5)	(12.5)	(12.2)
β_5	-6.29**	-6.02**	-4.26	-3.85
	(2.75)	(2.78)	(2.64)	(2.63)
β_6	-6.41*	-6.21*	-4.16	-3.83
	(3.29)	(3.34)	(3.16)	(3.18)
β_7	-0.042	0.53	-1.34	-0.42
	(1.63)	(1.57)	(1.89)	(1.77)
β_8	6.32**	6.01**	4.24	3.76
	(2.86)	(2.88)	(2.73)	(2.73)
β_9	1.61*	1.52*	1.20	1.06
	(0.82)	(0.84)	(0.81)	(0.82)
β_{10}	-1.64	-1.65	-1.42	-1.45
	(1.33)	(1.34)	(1.34)	(1.34)
δ_{LDT}	0.0019	-0.014	0.065	0.045
	(0.063)	(0.063)	(0.061)	(0.061)
δ_{RWD}	-0.023	-0.034	0.042	0.027
	(0.076)	(0.077)	(0.075)	(0.074)
δ_{AWD}	0.0021	0.012	-0.022	-0.010
	(0.089)	(0.089)	(0.088)	(0.088)
δ_{MT}	0.16	0.17	0.0024	0.013
	(0.17)	(0.17)	(0.17)	(0.17)
δ_{SAT}	0.17	0.15	-0.74	-0.78
	(0.50)	(0.50)	(0.74)	(0.74)
δ_{Luxury}		-0.12*		-0.16**
		(0.068)		(0.068)
$\delta_{ABS/TC}$		0.015		0.015
		(0.040)		(0.039)
Observations	158	158	158	158

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Weighted nonlinear least squares estimation of equation (2.12). The dependent variable is based on self-reported VMT in regressions (1) and (2), and odometer-based VMT in regressions (3) and (4).

Table 2.4: Estimation of Equation (2.13)

K:	35				40				45			
$L+1$	A	d	r	WP	A	d	r	WP	A	d	r	WP
10	3.03	0.71	0.41	0.48	3.46	0.76	0.32	0.55	3.89	0.80	0.26	0.62
15	3.03	0.70	0.42	0.40	3.46	0.75	0.34	0.45	3.89	0.78	0.28	0.51
20	3.03	0.70	0.43	0.36	3.46	0.74	0.34	0.41	3.89	0.78	0.28	0.47

K: engineering estimate of marginal price of fuel economy, L+1: length of vehicle life, A: multiplicative factor, d: discount factor, r: discount rate, WP: WTP/PV

Table 2.5: Estimates of A, d, r and WTP/PV

	(1)	(2)	(3)	(4)
5%	1069	1091	1405	1521
	(362)	(366)	(616)	(772)
25%	1485	1537	1893	2000
	(587)	(635)	(726)	(847)
50%	1887	1904	2342	2408
	(745)	(646)	(939)	(1167)
75%	2251	2289	2812	2886
	(640)	(755)	(1063)	(1285)
95%	3281	3167	3922	3816
	(967)	(1031)	(1583)	(1817)
Mean	1942	1951	2407	2474
Std. dev.	641	620	726	701
Observations	158	158	158	158

The total price of fuel economy (F) is estimated using equation (2.10), the estimated coefficients in the corresponding column of Table 2.4 and the corresponding $\hat{A}$. The table reports the mean, standard deviation and selected percentiles of $\hat{F}$, treating a vehicle model as one observation. In parentheses are the standard errors of $\hat{F}$ for the corresponding vehicle models that result from randomness in the estimated coefficients of Table 2.4.

Table 2.6: Statistics of Estimated Total Price of Fuel Economy ($\hat{F}$)

	(1) Efficiency gain (%)	(2) Price increase ($)	(3) Mean $ /Mean %
Multi-valve, overhead camshaft	2-5	105-140	35.0
Variable valve timing	2-3	35-140	35.0
Five-speed automatic transmission	2-3	70-154	44.8

Table 7: “Marginal” Fuel Efficient Technologies

Chapter 3: Environmental Policy and Induced Innovation: Evidence from Automobile Fuel Economy Regulations

3.1 Introduction

With oil price hikes in recent years and the need to reduce carbon dioxide emissions, developing fuel efficient vehicles has become an important issue worldwide. A common tool used in many countries to take this challenge is fuel economy standards that require automakers to achieve a certain level of average fuel efficiency in their fleet. Many countries have been tightening fuel economy regulations lately. For example, in 2009 the Obama administration established new automobile fuel economy standards that require each automaker to achieve the average fuel economy of 39 miles per gallon (mpg) for passenger cars, and 30 mpg for light trucks by 2016. The new rule will bring a significant improvement in fuel economy, considering the 2009 Corporate Average Fuel Economy (CAFE) standards of 27.5 mpg for passenger cars and 23.1 mpg for light trucks.

These regulations will bring huge changes in many aspects of automobiles and the auto industry. Average fuel economy will be improved, of course. Smaller cars may become more common because they are more fuel efficient. Vehicle prices are likely to increase as well, as evidenced by a U.S. government estimate that the new U.S. standards will raise the production cost of the average vehicle by $1,300. Lastly, new technologies may be developed for automakers to comply with the regulations. This paper focuses on this last point.

In environmental economics, government regulations, along with resource prices, are considered to be a possible source of technological progress (the "induced innovation hypothesis").[1] According to the hypothesis, firms that are required to comply with environmental standards will allocate more R&D activities to develop technologies useful for meeting the standards. This leads to technological progress in environmental technologies. In other words, environmental regulations "induce" technological progress. The hypothesis indicates that the tightening of fuel economy regulations in recent and coming years may induce innovations in automakers' technology for providing fuel efficiency, which will in turn lower the cost of achieving the regulatory targets.

There are studies finding empirical evidence of induced innovation from environmental regulations.[2] For instance, Newell et al. (1999) analyze technological progress in energy efficiency of air conditioners and water heaters, finding that the direction of innovation is responsive to energy price changes for some products, and that government energy efficiency standards also have a significant impact on the average energy efficiency of the product menu. Popp (2002) uses patent data to study the impact of energy prices on innovations in energy-saving technology, find-

[1] Jaffe et al. (2003) survey the theoretical and empirical literature on induced innovation.

[2] Vollebergh (2007) gives a thorough survey of the recent empirical literature on induced innovation.

ing that a strong, positive impact of energy prices on new innovations. Popp (2006) uses patent data from the United States, Japan and Germany and examines both innovation and diffusion of air pollution control equipment. He finds evidence that innovations respond to environmental regulatory pressure.

There are only few studies on induced innovation in the auto industry. Berry et al. (1996) analyze automakers' production cost function during 1972–1982, when gasoline prices increased rapidly and emissions and fuel economy standards were tightened. They find that after controlling for changes in vehicle characteristics, vehicle production costs increased over time during the period. They also find that patent applications related to combustion engines increased significantly, implying fast technological progress. An important implication of Berry et al. (1996) is that with rapidly changing gasoline prices and regulations, production cost increase and technological progress happen simultaneously, both contributing to fuel economy improvement. Knittel (2009) examines how fuel efficiency technology has developed over time using an estimation model similar to the one of this paper. He finds evidence of technological progress in fuel efficiency taking place in the U.S. in the early 1980s, his estimate for this period might be overestimated because an increase in production costs as suggested by Berry et al. (1996) is not considered.

This paper investigates the possibility of innovation induced by fuel economy regulations. Specifically, I estimate how Japanese automakers' technology has improved since the 1990s. Unlike the U.S. CAFE standards which had not received a significant revision since the late 1980s until after 2005, Japanese fuel economy regulations started to be tightened in the 1990s. I estimate technological progress in fuel efficiency achieved by Japanese automakers before and after regulations were introduced, and see whether technological progress is accelerated under regulations.

An issue in examining innovation in fuel efficiency is how to distinguish technological progress from the effect of an increase in production costs on various fuel-economy-improving systems or devices. The findings of Berry et al. (1996) imply that it is likely that both technological progress and cost increase occurred simultaneously after fuel economy regulations were introduced in Japan. However, it is very difficult to observe the production cost spent on fuel-economy-improving systems or devices used in each vehicle model. This paper intends to avoid this problem by analyzing cars and trucks produced by Japanese automakers and sold in *the United States*. For these vehicles, the production costs pertaining to fuel efficiency likely remained stable during the 1990s and early 2000s. This is, firstly, because U.S. fuel economy regulations were not binding for Japanese automakers during the period. Secondly, gasoline prices in the U.S. were very stable in this period, and so is consumer demand on fuel economy. This stability of fuel efficiency related production costs allows us to observe the effect of technological progress.

The estimation result provides strong evidence on induced innovation in fuel efficiency technology. I regress fuel economy of automobiles on vehicle characteristics such as weight and acceleration (or horsepower) and time. The coefficient on the time variable is interpreted as the rate of technological progress for the sample period. I run this regression for different periods, and observe how the rate has changed over time. I find that the rate of technological progress accelerated significantly after 2000, implying inducement from fuel economy regulations established in Japan in the 1990s.

This paper is organized as follows. Section 2 provides a simple model to understand how automobile fuel economy is determined. Section 3 discusses fuel economy regulations and the trend of average fuel economy in Japan. Section 4 explains the empirical framework of the paper, and section 5 defines estimation equation and explains the data. Section 6 shows estimation results. Section 7 concludes and discusses future work.

3.2 Defining Fuel Efficiency Technology

Let us consider the following simple model of automaker i's "fuel efficiency cost" function at time t:

$$c_j = F_t^i(\mathbf{X}_j, e_j). \tag{3.1}$$

Here, j and t represent vehicle j and time t, respectively. e_j is fuel efficiency of the vehicle, expressed in miles per gallon. $\mathbf{X}_j$ is (a vector of) vehicle characteristics that affect fuel economy, such as vehicle weight, engine size/type, horsepower, acceleration, body style, drivetrain and transmission. c_j is the (minimum) cost of producing a vehicle with a characteristics bundle of $(\mathbf{X}_j, e_j)$. Given $\mathbf{X}_j$, improving fuel economy e_j requires installing more or better efficiency-improving systems/devices, which increase the fuel efficiency cost c_j (i.e., $\partial F_t^i(\mathbf{X}_j, e_j)/\partial e_j > 0$).[3] For example,

[3]Specifically, National Research Council (2002) lists three categories of fuel economy improving devices/systems. The first category improves the energy efficiency of engines by reducing friction and other mechanical losses or by improving the processing and combustion of fuel and air (e.g., variable valve timing, cylinder deactivation and direct injection engines). The second category improves the efficiency of the transmission system where power is transmitted from the engine to the drive shaft or axle (e.g., six-speed automatic transmission and continuously variable transmission). The third category relates to other ways of improving fuel economy such as aerodynamic drag

a vehicle with a six-speed transmission would achieve higher fuel efficiency than a vehicle of the same characteristics except for carrying a five-speed transmission. But installing more or better efficiency-improving systems will incur additional costs.

Solving equation (3.1) for e_j (assuming it is solvable), we obtain:

$$e_j = G_t^i(\mathbf{X}_j, c_j). \tag{3.2}$$

$F_t^i(\cdot, \cdot)$ and $G_t^i(\cdot, \cdot)$ of course contain equivalent information. Below I will mainly work with $G_t^i(\cdot, \cdot)$.

$G_t^i(\cdot, \cdot)$ (or $F_t^i(\cdot, \cdot)$) is a function that represents the level of automaker i's fuel efficiency technology at time t, and is of our primary interest. It describes how fuel efficiency e_j relates to vehicle characteristics such as weight and horsepower, and efficiency-improving systems as summarized in c_j.[4] In other words, $G_t^i(\cdot)$ represents automaker i's production possibility frontier (PPF) related to fuel economy. In this framework, technological change is equivalent to a shift of $G_t^i(\cdot)$ over time. The induced innovation hypothesis claims that there are two sources of technological change: One is "autonomous" technological change that occurs even without any changes in the conditions producers are facing, and the other is "induced" technological change that results from changes in market conditions such as resource prices and regulations. Therefore, induced technological change is represented as a faster shift of $G_t^i(\cdot)$ after fuel economy regulations are tightened or gasoline prices go up.

reduction, rolling resistance reduction and integrated starter/generator systems. The Council also includes novel vehicle concepts such as hybrid electric vehicles in the third category.

[4]The theoretical background on including cost c_j in the technology function can be found in, for example, Alexander and Mitchell (1985), Triplett (1985) and Newell (1997).

This model indicates that automobile fuel economy is determined by three factors: vehicle characteristics affecting fuel economy ($\mathbf{X}_j$), the cost spent on improving fuel economy (c_j), and the level of the producer's fuel efficiency technology (the shape and location of $G_t^i(\cdot)$). Fuel economy that we observe is a combined outcome of these three different factors. It is important to distinguish the first two factors from the last: A change in e_j (Δe_j) due to a change in $\mathbf{X}_j$ or c_j ($\Delta\mathbf{X}_j$ or Δc_j) is a shift *along* the PPF, while Δe_j due to a change in $G_t^i(\cdot)$ ($\Delta G_t^i(\cdot)$) is a shift *of* the PPF.

Fuel economy regulations and changes in gasoline prices are two important factors that (indirectly) change fuel economy through affecting $\mathbf{X}_j$, c_j and $G_t^i(\cdot)$. Facing new regulations or consumers' strong demand for fuel economy due to high gasoline prices, automakers improve fuel economy of their models by adjusting $\mathbf{X}_j$ or c_j or advancing $G_t^i(\cdot)$. Reducing the size and power of the vehicle is the easiest way to improve fuel economy, although it may not be welcomed by consumers. Indeed, the introduction of the CAFE Standards in the U.S. in 1978 brought a significant reduction in vehicle weight and power (**?**). In addition, new regulations or high gasoline prices likely increase efficiency-improvement cost c_j. For instance, the new fuel economy regulations established by the Obama administration in 2009 are expected to increase the cost of the average car by $1,300. Also, Berry et al. (1996) estimate automakers' cost function for 1972-1982 and find that tighter emission standards and increasing gasoline prices likely moved production costs upward. Third, as induced innovation hypothesis predicts, the regulatory pressure or higher demand for fuel efficiency may induce research and development in fuel efficiency and shift the PPF, $G^i(\cdot)$, upward. In order to empirically analyze the effect of fuel economy regulations on technological progress, we need to separate these three effects and extract information on the shift of $G_t^i(\cdot)$ alone.

Estimating induced technological progress would be easy if we have data on e_j, $\mathbf{X}_j$ and c_j for vehicles of different model years. We can estimate $G_t^i(\cdot)$ by regressing e_j on $\mathbf{X}_j$ and c_j for year t models, and analyzing technological progress by observing how $G_t^i(\cdot)$ shifts over time.

The trouble is that c_j is most likely unobservable. It is practically impossible to find out how much is spent for fuel economy improvement in each vehicle model. If we ignore c_j and just regress e_j on $\mathbf{X}_j$ for different model years when c_j is likely to be changing over time as well, we can still control for the effect of changes in $\mathbf{X}_j$ ($\Delta\mathbf{X}_j$), but cannot distinguish the effect of Δc_j and the effect of $\Delta G_t^i(\cdot)$. We cannot know whether fuel economy improvement after controlling for changes in vehicle characteristics comes from technological progress or from increasing the efficiency improvement cost. Under tightened regulations, ignoring c_j is likely to overestimate technological progress because c_j will also increase in order to meet the regulations.[5]

3.3 Fuel Economy Regulations and Trends in Japan

Figure 1 summarizes Japanese fuel economy standards set in the 1990s. Target values are based on weight classes, so each automaker's average fuel economy in each class is required to exceed the corresponding target value. The range of a "flat" part of a line in Figure 1 corresponds to a weight class. In 1993, the government established new

[5]Therefore, Knittel (2009)'s estimate of "technological progress" in the early 1980s is probably overestimated because c_j likely increased over time in the early 1980s due to tightened regulations and high gasoline prices but he did not include any measure of c_j in the estimation.

standards, which came into effect in 2000. These standards aimed at improving fuel economy of the total fleet by 8.5% in 2000, relative to 1990. In 1999, new standards to be satisfied in 2010 were enacted, targeting to improve average fuel economy by about 23% by 2010, relative to 1995.[6]

Have these standards actually helped improve fuel economy? Table 1 summarizes the rate of fuel economy improvement in each weight class for different periods. For example, the value of 0.2 in weight class (1) for 1990-1993 means that sales-weighted average fuel economy of vehicles in that class improved by 0.2% between 1990 and 1993. The last column shows the change in sales-weighted average fuel economy of the total fleet.[7] Table 1 shows that fuel economy improved much faster in the last three periods than the first two. Especially, improvement rates jumped up in the 1996-1999 period in most weight classes. As gasoline prices in Japan was slowly going down in the 1990s (Figure 2), these jumps should be attributed to the introduction of fuel economy standards set in 1993. Although the standards came into effect in 2000 and automakers were not under any obligations until that time, improvement accelerated earlier because new fuel efficient technologies are mostly introduced when each vehicle model experiences a model change, which usually occurs every five to seven years. That is, rather than having a huge jump in year 2000, fuel economy improved gradually since the mid-1990s.

Although Table 1 shows the effect of regulations on improving fuel economy, it does not necessarily imply induced technological progress. As discussed before, when

[6]Recently, targets for 2015 were established in 2007, which are expected to improve average fuel economy by 23.5% in 2015, relative to 2004.

[7]Weight classes are defined differently for 1990-1993, so values for weight classes (4)-(9) are not available.

regulations are tightened, fuel economy improvement may come from three factors: $\Delta\mathbf{X}_j$, Δc_j and $\Delta G_t^i(\cdot)$. As Table 1 is based on weight classes and weight is the most important factor affecting fuel economy, the effect of $\Delta\mathbf{X}_j$ over time is (partially) controlled. The improvement rates in Table 1 still contain the effects of both Δc_j and $\Delta G_t^i(\cdot)$. We cannot be sure from the table alone how much of the improvement after 1996 comes from technological progress ($\Delta G_t^i(\cdot)$).

3.4 Empirical Framework

This paper attempts to obtain a reliable evidence of technological progress induced by Japanese fuel economy regulations set in the 1990s by looking at Japanese cars sold in *the United States*. As explained below, efficiency improvement cost c_j for Japanese cars sold in the U.S. between the late 1980s and 2004 are likely to be stable over time, making the effect of Δc_j on fuel economy negligible. This allows us to find evidence of induced innovation regardless of unobservability of c_j.

There are two key observations for this strategy. First, Japanese automakers practically have faced no fuel economy regulations in the U.S. market. Second, gasoline prices remained stable since the mid-1980s until the early 2000s, making consumers' demand for fuel economy relatively stable compared to the mid and late 2000s, when gasoline prices changed significantly.

Figure 3 shows the historical trend of the CAFE Standards since 1978, when they first came into effect. Under the CAFE Standards, for each of specified categories (currently, imported passenger cars, domestically produced passenger cars, and light trucks with a gross vehicle weight rating (GVWR) of 8,500 pounds or less) of new vehicles sold in the United States in a model year, a manufacturer is required to keep the sales-weighted average fuel economy of its fleet above the level set by the National Highway Traffic Safety Administration (NHTSA). It needs to pay a penalty if it fails to meet the CAFE target for a category.[8] As can be seen in the figure, the standards were being tightened consistently until the mid-1980s. Since then, however, they remained at almost the same levels for two decades. Since the mid-2000s, the light truck target has been tightened again at a fast pace.[9]

Moreover, the CAFE standards have not been binding for most Japanese automakers. Except for some cases, their average fuel economy values have been well above the target values for all years since the introduction of the CAFE standards. Therefore, the CAFE regulations have had little or no influences on Japanese automaker's strategy in the U.S. market. Practically, Japanese automakers have faced no fuel economy regulations in the U.S. Hence, they do not need to increase c_j of their models sold in the U.S. in order to comply with the U.S. regulations.

Next, we look at another factor that affects c_j significantly: gasoline prices and consumers' demand for fuel economy. As gasoline becomes more expensive, consumers will become more willing to pay for fuel economy because it will reduce the operating cost of owning a vehicle. Higher demand for fuel efficient vehicles

[8]Currently, the penalty is \$5.50 for every 0.1 mile per gallon under the target value times the total volume of the vehicles in the fleet.

[9]In March 2009, the NHTSA announced that the passenger car target would be tightened for model year 2011 and become more stringent than the 1985 target of 27.5 miles per gallon for the first time. Under the new rule set by the Obama administration in 2009, each automaker will be required to achieve the total fleet (passenger cars and light trucks) average of 35.5 miles per gallon by 2016. This will bring a significant improvement in fuel economy, considering the 2009 targets of 27.5 mpg for passenger cars, and 23.1 mpg for light trucks.

induces firms to improve e_j by changing c_j and $\mathbf{X}_j$. Figure 4 shows regular unleaded gasoline real prices in the U.S. since the 1970s. As we see in the figure, gasoline prices were very stable since the late 1980s until 2004 or so, implying that consumers' demand for fuel economy was also relatively stable during the period. Figure 5 plots the trend of sales-weighted average weight of all vehicles sold in the U.S. There is a clear change of trend in 2005: The pace of weight increase slowed down after 2005. This is consistent with the fact that gasoline prices went up significantly in 2005. Consumers started to demand fuel economy in 2005. In other words, gasoline prices had only a negligible impact on their demand for fuel economy until 2004, so that c_j of 2004 models and earlier are likely to be free from consumers' stronger demand for fuel economy due to increasing gasoline prices. I assume that consumer's demand for fuel efficiency remained almost constant between 1988 and 2004.

Based on the above observations, I assume in this paper that c_j of Japanese cars sold in the U.S. market between 1988 and 2004 remained so stable over time that improvement in fuel economy attributable to Δc_j is negligible.

On the other hand, if tightening of fuel economy regulations in Japan induces technological progress in fuel efficiency technology of Japanese automakers, this progress should be reflected even in Japanese cars sold in the U.S. As rational producers, automakers would try to produce vehicles as close to the technological frontier $G_t^i(\cdot)$ as possible. Therefore, wherever a vehicle is produced or sold, its producer has an incentive to use better fuel efficiency technology to produce it. Of course, even within the same company, it would probably be impossible to transplant all new technologies to all factories instantaneously. Thus, if we find evidence of induced technological progress from Japanese fuel economy regulations by observing Japanese cars sold in the U.S., it may be more appropriate to interpret it as the lower bound for the magnitude of induced technological progress.

3.5 Estimation Equation and Data

Using data on Japanese cars sold in the U.S., I estimate the following equation separately for every two year period between 1988 and 2006 (i.e., 1988-1990, 1989-1991, . . ., 2004-2006):

$$\ln mpg_j = \delta t + \mathbf{X}_j \boldsymbol{\beta}. \tag{3.3}$$

Subscript j represents vehicle configuration j.[10] mpg_j is fuel economy (miles per gallon) and $\mathbf{X_j}$ is a vector of vehicle attributes that affect fuel economy: log of weight ($\ln w_j$), log of horsepower-to-weight ratio ($\ln a_j$), which measures acceleration capacity, transmission dummy (automatic [AT], semi-automatic [SAT] or manual [MT]), drivetrain dummy (front wheel drive [FWD], rear wheel drive [RWD] or all wheel drive [AWD]), light-duty truck dummy and a constant. In each estimation, t takes zero if configuration j is observed in the first year of the period, and takes one if it is observed in the second year, and so on.[11] Thus, δ (or more precisely, $\exp(\delta) - 1$) is the annual rate of fuel efficiency improvement after controlling for vehicle characteristics $\mathbf{X}_j$, which may contain the effects of both Δc_j and $\Delta G_t^i(\cdot)$.

[10] A vehicle model (e.g. Toyota Camry) in a given model year usually has a number of configurations with differences in weight, engine displacement, transmission type, and so on. The data used in the estimation generally include several configurations per model.

[11] For instance, suppose configuration j is observed in the data set for year 2000. If the estimation period is 2000-2002 (1999-2001), then $t = 0$ ($t = 1$) for this observation.

Since c_j can reasonably be considered stable from the 1988 to at least 2004, δ is a good estimate of the rate of technological progress during this period.

I use the EPA's annual "Fuel Economy Test Car List Data." This data set contains information on all fuel economy tests performed each year, providing data on e_j and $\mathbf{X}_j$ for almost all vehicle configurations sold in the U.S. each year. In the following estimation, I use only gasoline-engine models, so hybrid and diesel models are excluded. Each year I have around 300-500 configurations used in the regressions.

I estimate Equation (3.3) by Weighted Least Squares using the above data. Weights are calculated as follows. All configurations produced by automaker i in model year t receive the same weight which equals the total sales of i at t divided by the number of configurations produced by i at t. This weighting prevents small automakers like Suzuki and Mazda, which tend to have disproportionately many observations compared to their sales, from being overrepresented.

3.6 Estimation Results

Table 2 summarizes the size and standard error of the rate of fuel efficiency improvement after controlling for changes in vehicle characteristics, calculated as $(\exp(\hat{\delta}) - 1) \times 100$, where $\hat{\delta}$ is the estimate of δ, for all two year periods between 1988 and 2006. Figure 6 plots the values in Table 2.The horizontal axis represents the initial year of a two year period. As discussed above, from period 1988-1990 to period 2002-2004, fuel economy improvement after controlling for vehicle characteristics mostly results from technological progress. Thus, the rate of fuel economy improvement can be reasonably interpreted as the rate of technological progress in fuel efficiency. For the last two periods (2003-2005 and 2004-2006), a large portion of fuel economy improvement might result from an increase in fuel efficiency cost c_j.

Table 2 and Figure 6 provide evidence of both autonomous and induced technological progress. From the initial period of 1988-1990 until the mid 1990s, the average rate is stable around 1.5% a year. During the mid to late 1990s, the rate appears to be on the decrease, reaching the lowest of 0.4% for the 1998-2000 period. After that, it increases sharply to 2.0% for 2001-2003, and remains above 2%. The rate of technological progress has been positive for all periods including the late 1980s and early 1990s, when Japanese automakers faced no binding fuel economy regulations in Japan and the U.S., and gasoline prices were so stable that consumers' demand for fuel economy was likely to be stable as well. Thus, technological progress in fuel efficiency occurs autonomously even without any pressure from fuel economy regulations and consumers' demand for fuel economy. More importantly, the jump in the rate in the 2000s is evidence of induced innovation: Japanese fuel economy regulations accelerated the speed of technological progress in fuel efficiency.

3.7 Conclusion

Environmental regulations are considered a potential factor to accelerate technological progress (the "induced innovation hypothesis"). According to the hypothesis, firms that are required to comply with environmental standards will allocate more R&D activities to develop technologies useful for meeting the standards. This leads to technological progress in environmental technologies. In other words, environmental regulations “induce” technological progress.

This paper analyzes whether automobile fuel economy regulations induce technological progress in fuel efficiency by looking at past experiences of Japanese automakers. Japanese fuel economy regulations started to be tightened in the 1990s. I estimate technological progress in fuel efficiency achieved by Japanese automakers before and after new regulations were introduced, and see whether technological progress is accelerated under the regulations, as indicated by the induced innovation hypothesis.

An issue in estimating technological progress in fuel efficiency is how to distinguish technological progress from the effect of an increase in production costs spent on various fuel-economy-improving systems or devices. Both better technology and increased production costs can increase fuel economy. Findings by Berry et al. (1996) with U.S. data imply that both of these factors were likely to be in effect simultaneously after fuel economy regulations were introduced in Japan. However, it is difficult to observe the cost spent on fuel-economy-improving systems or devices used in each vehicle model. To deal with this issue, this paper focuses on Japanese cars and trucks sold in the U.S. Vehicles produced by Japanese automakers and sold in the U.S. are likely to have little changes in the production cost relating to fuel efficiency during the 1990s and early 2000s for two reasons. First, U.S. fuel economy regulations were not binding for Japanese automakers during the period. Additionally, gasoline prices in the U.S. were fairly stable in this period, and so is consumer demand on fuel economy. These observations imply that production costs relating to fuel efficiency would remain stable for Japanese cars and trucks sold in the U.S. market, allowing us to identify the effect of technological progress.

I regress fuel economy of different vehicle models on vehicle characteristics (e.g., weight and horsepower) and time. The coefficient on the time variable gives the rate of technological progress for the sample period. I estimate the rate of technological progress for different periods between 1988 and 2006, and analyze how the rate has changed over time. Estimation results suggest that technological progress has accelerated significantly after 2000, implying that fuel economy regulations introduced in Japan in the 1990s induced innovations in fuel efficiency technology used by Japanese automakers.

.1 Figures and Tables

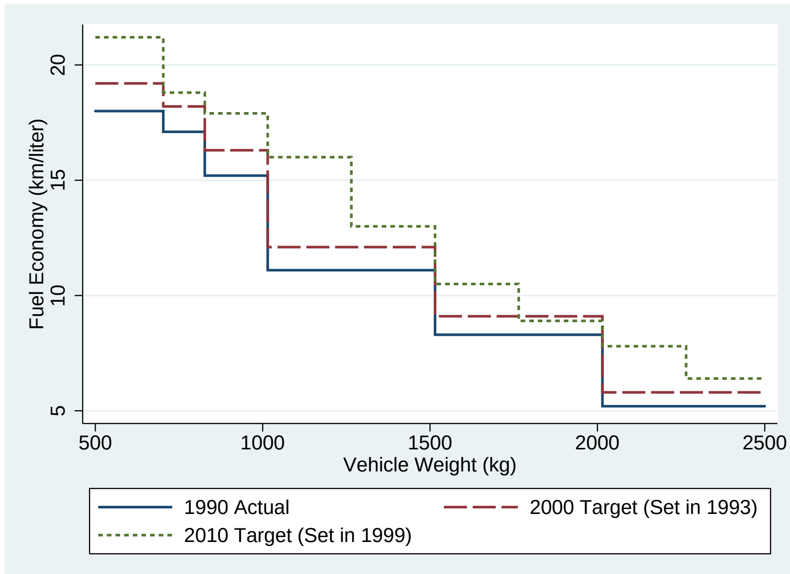

Figure 1: Japanese Fuel Economy Standards Established in the 1990s

Period\Class	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	Total
1990-1993	0.2	-0.3	0.1							-1.6
1993-1996	1.0	-0.1	0.4	0.8	0.4	0.0	0.1	-0.3	0.1	-0.2
1996-1999	0.4	1.5	1.0	1.0	1.1	1.4	0.7	1.3	0.3	1.1
1999-2002	0.7	1.0	1.6	1.9	1.1	0.8	0.8	0.3	0.1	1.4
2002-2005	4.3	1.0	0.4	0.9	0.6	0.8	0.7	0.5	0.3	0.5

Table 1: Fuel Economy Improvement in Japan by Weight Class

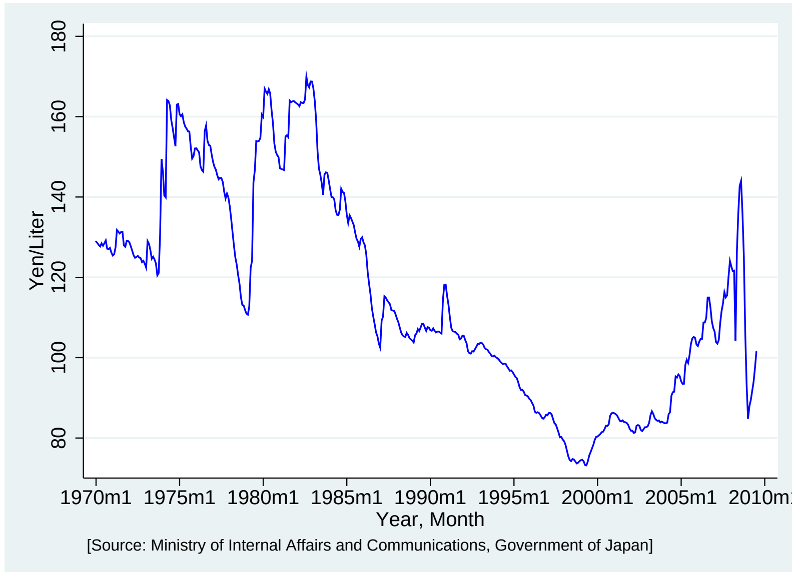

Figure 2: Regular Gasoline Prices per Liter in Japan (in 2005 Yen)

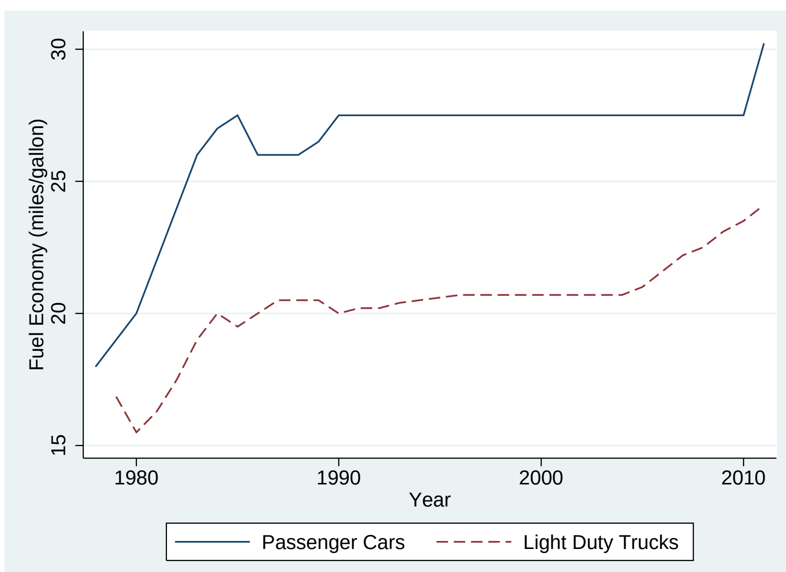

Figure 3: U.S. CAFE Standards

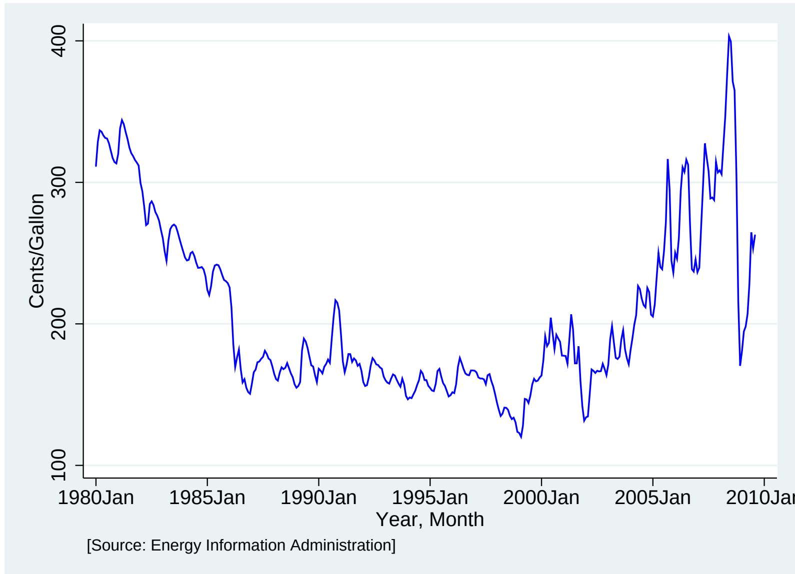

Figure 4: Regular Gasoline Prices per Gallon in the U.S. (in August 2008 Dollars)

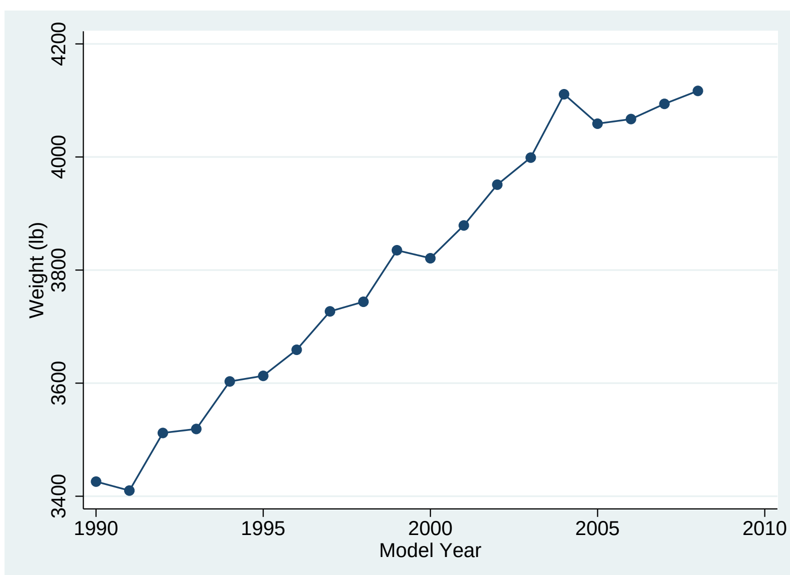

Figure 5: Average Vehicle Weight in the U.S.

Period	Rate (%)	Std. Err.
1988-1990	1.57	0.33
1989-1991	1.11	0.31
1990-1992	1.29	0.30
1991-1993	1.52	0.32
1992-1994	1.41	0.34
1993-1995	1.47	0.32
1994-1996	1.36	0.31
1995-1997	0.88	0.30
1996-1998	1.12	0.29
1997-1999	1.12	0.27
1998-2000	0.40	0.25
1999-2001	0.71	0.26
2000-2002	1.26	0.28
2001-2003	2.03	0.30
2002-2004	2.15	0.28
2003-2005	2.25	0.27
2004-2006	2.25	0.26

Table 2: The Rate of Technological Progress in Fuel Efficiency

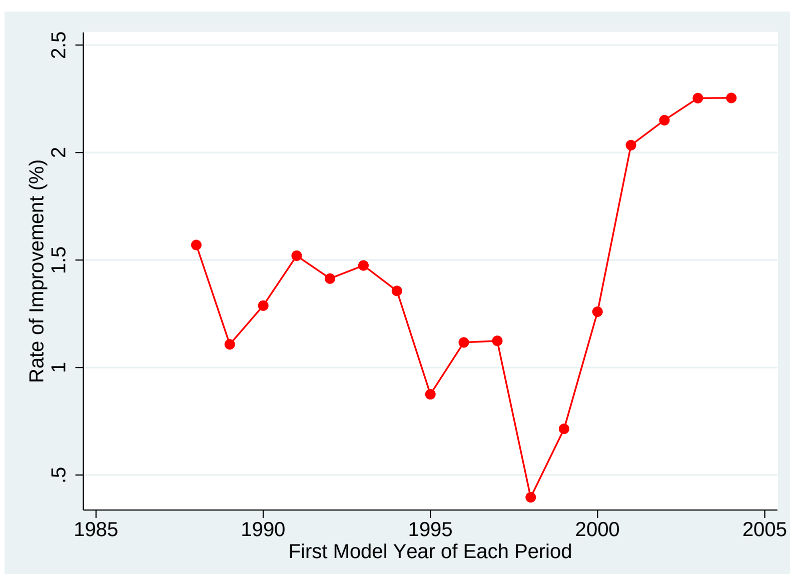

Figure 6: Rate of Technological Progress in Fuel Efficiency

Bibliography

Alexander, Arthur J. and Bridger M. Mitchell (1985), "Measuring technological change of heterogeneous products." *Technological Change and Social Change*, 27, 161–195.

Anderson, Soren T. and James M. Sallee (2011), "Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards." *American Economic Review*, 101, 1375–1409.

Arguea, N. M., C. Hsiao, and G.A. Taylor (1994), "Estimating consumer preferences using market data-an application to US automobile demand." *Journal of Applied Econometrics*, 9, 1–18.

Arguea, Nestor M. and Cheng Hsiao (1993), "Econometric issues of estimating hedonic price functions : With an application to the U.S. market for automobiles." *Journal of Econometrics*, 56, 243–267.

Atkinson, Scott E. and Robert Halvorsen (1984), "A new hedonic technique for estimating attribute demand: An application to the demand for automobile fuel efficiency." *Review of Economics and Statistics*, 66, 417–426.

Austin, David and Terry Dinan (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, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen (2009), "Distributional and efficiency impacts of increased U.S. gasoline taxes." *American Economic Review*, 99, 667–699.

Berry, Steven, Samuel Kortum, and Ariel Pakes (1996), "Environmental change and hedonic cost functions for automobiles." *Proceedings of the National Academy of Sciences*, 93, 12731–12738.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile prices in market equilibrium." *Econometrica*, 63, 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes (2004), "Differentiated products demand systems from a combination of micro and macro data: The new car market." *Journal of Political Economy*, 112, 68–105.

Bockstael, Nancy and Kenneth E. McConnell (2006), *Environmental and Resource Valuation with Revealed Preferences*. Springer.

Coleman, Conrad and Winston Harrington (2010), "A simulation model of the new vehicle market." *mimeo*.

Deaton, Angus and John Muellbauer (1980), *Economics and Consumer Behavior*. Cambridge University Press.

Dreyfus, Mark K. and W. Kip Viscusi (1995), "Rates of time preference and consumer valuations of automobile safety and fuel efficiency." *Journal of Law and Economics*, 38, 79–105.

Dubin, Jeffrey A. and Daniel L. McFadden (1984), "An econometric analysis of residential electric appliance holdings and consumption." *Econometrica*, 52, 345–362.

Environmental Protection Agency (2009), *Light-Duty Automotive Technology and Fuel Economy Trends: 1975 Through 2009*. URL `http://www.epa.gov/oms/fetrends.htm`.

Environmental Protection Agency (2012), "Inventory of U.S. greenhouse gas emissions and sinks: 1990-2010." `http://www.epa.gov/climatechange/Downloads/ghgemissions/US-GHG-Inventory-2012-Main-Text.pdf`.

Environmental Protection Agency and Department of Transportation (2010), "Light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards; final rule." *Federal Register*, 75, 25324–25728.

Environmental Protection Agency and National Highway Traffic Safety Administration (2010), "Joint technical support document: Rulemaking to establish light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards." `http://www.nhtsa.gov/staticfiles/rulemaking/pdf/cafe/Final_Joint_TSD.pdf`.

Espey, Molly and Santosh Nair (2005), "Automobile fuel economy: What is it worth?" *Contemporary Economic Policy*, 23, 317–323.

Gillingham, Kenneth (2011), "How do consumers respond to gasoline price shocks? heterogeneity in vehicle choice and driving behavior." *mimeo*.

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

Goodman, Allen C. (1983), "Willingness to pay for car efficiency." *Journal of Transport Economics and Policy*, 17, 247–266.

Hogarty, Thomas F. (1975), "Price-quality relations for automobiles: A new approach." *Applied Economics*, 7, 41–51.

Jacobsen, Mark R. (forthcoming), "Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity." *American Economic Journal: Economic Policy*.

Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins (2003), "Technological change and the environment." In *Handbook of Environmental Economics* (Karl-Göran Mäler and Jeffrey R. Vincent, eds.), volume 1, Elsevier.

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 (2012), "New vehicle characteristics and the cost of the corporate average fuel economy standards." *RAND Journal of Economics*, 43, 186–213.

Knittel, Cristopher R. (2009), "Automobiles on steroids: product attribute trade-offs and technological progress in the automobile sector." *NBER Working Paper 15162*.

Mannering, Fred and Clifford Winston (1985), "A dynamic empirical analysis of household vehicle ownership and utilization." *RAND Journal of Economics*, 16, 215236.

Matas, Anna and Josep-Lluis Raymond (2009), "Hedonic prices for cars: An application to the spanish car market, 1981-2005." *Applied Economics*, 41, 2887–2904.

McManus, Walter (2007), "The link between gasoline prices and vehicle sales." *Business Economics*, January 2007, 53–60.

Murray, Jonathan and Nicholas Sarantis (1999), "Price-quality relations and hedonic price indexes for cars in the united kingdom." *International Journal of the Economics of Business*, 6, 5–27.

National Highway Traffic Safety Administration (2012), "Summary of fuel economy performance march 2012." `http://www.nhtsa.gov/staticfiles/rulemaking/pdf/cafe/March_2012_Performance_Summary.pdf`.

National Research Council (2002), *Effectiveness and Impact of Corporate Average Fuel Economy Standards*. National Academy Press.

Newell, Richard G. (1997), *Environmental Policy and Technological Change: The Effects of Economic Incentives and Direct Regulation on Energy-saving Innovation.* Ph.D. thesis, Harvard University.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1999), "The induced innovation hypothesis and energy-saving technological change." *Quarterly Journal of Economics*, 114, 941–975.

Ohta, Makoto and Zvi Griliches (1986), "Automobile prices and quality: Did the gasoline price increases change consumer tastes in the U.S.?" *Journal of Business & Economic Statistics*, 4, 187–198.

Parry, Ian, Margaret Walls, and Winston Harrington (2007), "Automobile externalities and policies." *Journal of Economic Literature*, 45, 374–400.

Popp, David (2002), "Induced innovation and energy prices." *American Economic Review*, 92, 160–180.

Popp, David (2006), "International innovation and diffusion of air pollution control technologies: The effects of NO_X and SO_2 regulation in the US, Japan, and Germany." *Journal of Environmental Economics and Management*, 51, 46–71.

Small, Kenneth A. and Harvey S. Rosen (1981), "Applied welfare economics with discrete choice models." *Econometrica*, 49, 105–130.

Small, Kenneth A. and Kurt Van Dender (2007), "Fuel efficiency and motor vehicle travel: The declining rebound effect." *Energy Journal*, 28, 25–51.

Spiller, Elisheba and Heather M. Stephens (2012), “The heterogeneous effects of gasoline taxes: Why where we live matters.” *RFF Discussion Paper*, 12-30.

Train, Kenneth (2009), *Discrete Choice Methods with Simulation*, 2nd edition. Cambridge University Press.

Train, Kenneth E. and Clifford Winston (2007), “Vehicle choice behavior and the declining market share of u.s. automakers.” *International Economic Review*, 48, 1469–1496.

Triplett, Jack E. (1985), “Measuring technological change with characteristics-space techniques.” *Technological Change and Social Change*, 27, 283–307.

Vollebergh, Herman R. J. (2007), “Differential impact of environmental policy instruments on technological change: A review of the empirical literature.” Tinbergen Institute Discussion Paper.

West, Sarah E. (2004), “Distributional effects of alternative vehicle pollution control policies.” *Journal of Public Economics*, 88, 735757.

Whitefoot, Kate, Meredith Fowlie, and Steven Skerlos (2012), “Product design response to industrial policy: Evaluating fuel economy standards using an engineering model of endogenous product design.” *Energy Institute at Haas Working Paper*.