FOUR PAPERS ON TRANSPORTATION AND THE ENVIRONMENT

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Kevin Daniel Roth

May 2013

FOUR PAPERS ON TRANSPORTATION AND THE ENVIRONMENT

Kevin Daniel Roth, Ph.D.

Cornell University 2013

The main essay of this thesis, found in the first chapter, examines how two policies that are *a priori* equivalent, fuel economy standards and feebates, interact differently with complementary policies that also attempt to improve fuel economy. To examine these interactions I build a general equilibrium model of the automobile market that allows manufacturers to trade off horsepower, weight, and fuel economy of vehicles along a production possibility frontier (PPF). I also estimate household demand for vehicles and miles for a simulation model that includes the used car and scrappage markets. This model allows me to simulate the interaction of a research and development policy that increases the PPF of domestic firms, or a tax credit that increases demand for efficient vehicles, with either a CAFE standard or feebate. I find that vehicle emissions increase under all these interactions but the effects are muted under the feebate because it allows fuel economy to improve by 0.60% to 1.88%, while CAFE, by targeting an average fuel economy, will always offset these uncoordinated complementary policies.

The second essay examines transportation systems with unpriced congestion where single-occupant low-emission vehicles are allowed into high occupancy vehicle (HOV) lanes to encourage their adoption exacerbates congestion costs for carpoolers. The resulting welfare effects of the policy are negative, with environmental benefits overwhelmingly dominated by the increased congestion costs. Exploiting the introduction of the Clean Air Vehicle Stickers policy in Cal-

ifornia with a regression discontinuity design, our results imply a best-case cost of \$124 per ton of reductions in greenhouse gases, \$606,000 dollars per ton of nitrogen oxides reduction, and \$505,000 dollars per ton of hydrocarbon reduction, exceeding those of other options readily available to policymakers.

The third essay examines the 'Energy Paradox.' From previous literature, it can be found that consumers tend to undervalue discounted future energy costs in their purchase decisions for energy-using durables. We show that this finding could, in part, result from ignoring consumer heterogeneity in empirical analyses as opposed to true undervaluation.

The fourth essay examines used vehicle scrappage rates. We find that not only are vehicles lasting longer but that scrap rates are less responsive to changes in vehicle price than previously estimated. These parameters help to refine the parameters used to evaluate public policies like CAFE standards and gasoline taxes that are known to have important effects on used vehicles.

BIOGRAPHICAL SKETCH

Kevin Roth was raised in Ashton, Wisconsin, a small town outside of Madison. He attended St. Peter's Catholic School and Middleton High School. He attended the University of Wisconsin-Madison from 1999 to 2004 where he majored in Mathematics, Economics and German Linguistics and studied for a year in Freiburg, Germany. After undergraduate, he worked for the Department of Justice in Washington D.C. and later for Charles River Associates in Boston. He began his PhD at Cornell in 2007 studying environmental economics.

To my family

ACKNOWLEDGEMENTS

I would like to begin by thanking Antonio Bento, my advisor, who has been an amazing advisor and good friend. I thank him for the many hours he spent educating me on everything in this profession, large and small, from how to ask good research questions to how to properly format a document. I am particularly grateful to Antonio for investing and believing in me at the start of the program. He showed an incredible amount of trust, allowing me to do research when I was an untested 1st year student. His high expectations and belief in my ability have pushed me tremendously as a scholar. I have no doubt that my ability to do research, my current placement, and my future successes are due to his time and energy. It is my hope that some day I can be the advisor he was to me to another student and to pass such dedication on to future students. I could not have chosen a better advisor and have been incredibly lucky to work so closely with such a prominent scholar in the field of environmental economics.

I would also like to thank my other committee members, Steve Coate and Matt Freedman. Steve helped me to sharpen my thinking about models and how to take a real world question and turn it into a simple model with strong predictive power. His encouragement pushed this thesis into sharper focus. Matt was one of the first professors I interacted with closely at Cornell having served as his TA in my first year. He has helped me to navigate many of the issues that young researchers face in the profession and has always been a great person to seek advice from. I thank him for the energy he has put into reading and giving feedback on my work over the years.

I would also like to thank a number of other professors at Cornell and elsewhere. First and foremost Shanjun Li who spent many hours teaching me discrete choice and econometric from which I have profited immensely. I would also like to thank Jordan Matsudaira for all his advice and teaching over the years. I have also learned much from being a TA for George Boyer, Uri Possen, and Bob Hutchens. I would also like to thank many others from whom I have received advice including Emily Owens and Levon Barseghyan as well as Dan Kaffine for being a great mentor and coauthor.

I would also like to thank my family and friends. My parents, Dan and Sharon, my sister, Tracy, my brother, Chris, his wife, Kelsey, as well as my grandparents, Laura and Buddy, Helen and Alois, have always been my biggest supporters. I was extremely fortunate to grow up in a household were my family believed each person was given talents and not to use them was a waste. They encouraged my curiosity and learning in everything from astronomy and biology, to music and history. While always making it clear that every line of work was honorable, they would never allow me to pursue a less ambitious path than I was capable of taking. There are far too many friends to mention whom I have leaned on in graduate school. I would like to thank Sharon Kim for countless hours of fun in grad school and advice that she was uniquely able to give in this process. My friends from home and before grad school who have always kept me grounded: Alison, Bron, Matt and Courtney, Melissa, Lyzz, Sarah, Lauren and Nick, Kelly, Courtney, Rachel, Aquene, and Jim. My incredible support network in Ithaca over the years including, Scott, Erik and Jun, Gerulf and Steve, Sanjay, John S., Teevrat, Laura and Andrew, Jessica B., Mike S., Alex R., Ben and Carrie O., Ellen, Tammy, Andy, Dennis, John Y., Brinda, Mario, Deni, Joe L., John P., Fernanda, Baran, Danielle and Jameel, Clayton, Emily, Lily and Jeremy, Shai and Yardenne, Ken W. and his family. I'd also like to thank the others in my research group, Scarlett, Yiwei, Joel, Rick, Ben, Corey, Andrew, and Dennis for their friendship and feedback over the years. I wanted to do a PhD because

I wanted to do something difficult that I was proud of. I think by doing something difficult I have come to appreciate those on whom I lean and who have made such an endeavor possible. I leave Ithaca with an incredible number of people in my life.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

THE UNINTENDED CONSEQUENCES OF UNCOORDINATED REGULATION: EVIDENCE FROM THE TRANSPORTATION SECTOR

1.1 Introduction

It is often the case that different aspects of a market failure are targeted with different policy instruments when the optimal instrument is not available to regulators. Although these aspects may be closely related, the instruments that address each may not be under the authority of a single agency. But it can be difficult to coordinate the implementation of multiple instruments when each is administered by a different agency. This failure to coordinate may result in the actions of one agency undermining the goals of another. Problems with coordination may appear not only across agencies at the same level of government, but also vertically as state or local agencies pursue their own policy objectives independent of federal agencies.[1] While failure to coordinate may be inevitable, it may be possible to choose instruments that are more robust to coordination failures. Two policies that appear at first equivalent may react very differently to the uncoordinated actions of another regulator pursuing its own policy objectives.

The transportation market is one such sector where the conditions arise for unintended consequences from multiple agencies attempting to correct many aspects related to the under provision of fuel economy. The most far-reaching

[1]Several reasons given for this lack of coordination are heterogeneous preferences resulting in some states desiring stronger regulation than, chosen nationally (Goulder, Jacobsen, & van Benthem, 2012), and federal “gridlock” resulting in decentralized decision making (Fischer & Preonas, 2010; Lyon & Yin, 2010) by various agencies or states.

regulation in this market, fuel economy standards, attempts to reduce the externalities associated with carbon emissions and gasoline use by improving the fuel economy of the vehicles on the road. This is not, however, the only reason that fuel economy may not rise to optimal levels. Firms may underinvest in research and development (R&D) if they are unable to capture spillovers to other firms, a problem which is often addressed with government sponsored R&D grants. On the demand side, policy makers may also subsidize the sales of new technology, such as hybrid vehicles, if they are concerned that information on these new products is slow to diffuse. The purpose of this paper is to examine how the failure to coordinate policies targeting fuel economy may cause interactions that are unanticipated by regulators, and how the inability to improve that coordination may prompt regulators to switch from a fuel economy standard to a closely related policy—a feebate—that appears, in a simple theoretical model, to confer no advantage over a fuel economy standard.

Historically, the Environmental Protection Agency along with the National Highway Traffic Safety Administration have administered fuel economy standards in the United States. Known as the Corporate Average Fuel Economy (CAFE) standard, this standard mandates a minimum average fuel economy for the car and truck fleets produced by each manufacturer. In recent years, both domestically as well as abroad, governments have begun to consider an alternative policy instrument known as a feebate.[2] Feebate, a portmanteau of fee and rebate, is a schedule of government taxes (and subsidies) that change

[2]The most recent feebate was proposed by Senators Jeff Bingaman (D-NM), Olympia Snowe (R-ME), Richard Lugar (R-IN), and John Kerry (D-MA) as part of bill S.1620 The Efficient Vehicle Leadership Act of 2009. Most countries in the EU as well as the United Kingdom have some form of feebate, either on new vehicles as in the French 'bonus/malus' system or on yearly registration fees as in Germany and the Scandinavian countries. See also (Anderson, Parry, Sallee, & Fischer, 2011).

proportionally with the fuel efficiency rating of a vehicle.[3] A feebate policy is characterized by a pivot point and a fee rate. Cars with efficiency above the pivot point receive a rebate and those below it a fee.

Although generally unrecognized by policy makers and economists, a fuel economy standard and a feebate can, in a simple framework, be set to provide identical outcomes.[4] Both policies generate a price wedge shifting preference towards fuel economy, and it does not matter whether that price wedge is set by the government in the case of the feebate, or the firms in the case of fuel economy standards. But the distinction in who sets the price wedge does matter when the baseline policy interacts with other uncoordinated policies administered by different agencies attempting to further improve fuel economy. While these complementary policies are often intended to improve overall fuel economy, the choice of baseline policy can have large implications for the success of these complementary policies. Examples of complementary policies include R&D grants from the Department of Energy on the supply side, or fuel-efficient vehicle tax credits, administered by the Internal Revenue Service, are one of many such policies on the dema. Because CAFE allows the firms to determine the price wedge, complementary policies will prompt firms to relax the price wedge to exactly achieve the standard and, due to more vehicle sales, increase total emissions. The interaction of a feebate with a complementary policy, by contrast, will produced improvements in the average fuel economy but may increase or decrease total emissions depending on how the quantity of sales and vehicle miles traveled (VMT) change. This paper demonstrates this distinc-

[3]More generally feebates can be used for the efficiency rating of products ranging from appliances to homes. For example, the city of Portland, Oregon has considered feebates on residential and commercial construction.

[4](Klier & Linn, 2012c) is, to the best of my knowledge, the only other paper that notes the equivalence between feebates and standards in the first order conditions of producers.

tion in behavior between fuel economy standards and a feebate, and shows that R&D policies may not only be inefficient, but actually counter-productive when pursued in a regulatory framework that guarantees those improvements will be channeled away from the targeted product attribute.

To accurately model how CAFE and feebates interact with these complementary policies I build a detailed model of the automobile market. I model suppliers as endogenously selecting the optimal fuel economy of their products along a production possibility frontier, allowing for more flexibility in the way they meet these regulatory targets. This is the first general equilibrium model to bring a supplier model of this type into a framework that also accounts for VMT, used vehicles, and scrappage markets.

Several prior studies have focused on supplier behavior in the automobile market. Some earlier models of automobile supply, including (Berry, Levinsohn, & Pakes, 1995), (Goldberg, 1998), (Kleit, 2004), (Austin & Dinan, 2005), (Bento, Goulder, Jacobsen, & von Haefen, 2009), and (M. R. Jacobsen, 2012a) assumed either fixed product characteristics or allow for some increase in fuel economy based on a cost-benefit analysis using engineering cost curves. (Knittel, 2011) provides evidence of a production possibility frontier where manufacturers trade off between horsepower, weight, and the fuel economy of vehicles. (Klier & Linn, 2012a), whose approach is most similar to mine, estimate the cost of CAFE when accounting for these product attribute tradeoffs, while (Whitefoot, Fowlie, & Skerlos, 2012) and (Gramlich, 2009) also investigate these costs using a production possibility frontier with fewer tradeoffs. Several papers in this literature have examined feebates specifically. (W. B. Davis, Levine, Train, & Duleep, 1995), (D. L. Greene, Patterson, Singh, & Li, 2005), and (Small,

2012) examine the effects of a US feebate policy using simulation.[5] Only (Small, 2012) compares feebates with a CAFE standard in a simulation model, finding they can produce effects that are similar in magnitude.[6]

This paper also contributes to a very large literature that compares equivalent tax (or price) instruments and quantity controls on environmental disamenities. This literature often, although not exclusively, focuses on how uncertainty affects the optimal choice of policy, as first examined by (Weitzman, 1974). (Adar & Griffin, 1976), (Yohe, 1978), (Stavins, 1996), (Hoel & Karp, 2001) and (Newell & Pizer, 2003) generalized these results to other settings. (Roberts & Spence, 1976) analyze a hybrid system combining price and quantity instruments. Using a political economy model, (Finkelshtain & Kislev, 1997) illustrate how price and quantity instruments are subject to different lobbying effort by firms. (Baron, 1985b) shows that price instruments dominate when the regulator has incomplete information on a firm with market power.

A final group of studies examines the efficiency of overlapping and uncoordinated regulation. (Goulder et al., 2012) find that local policies to improve fuel economy can be partially or entirely offset by a national CAFE standard. (Baron, 1985a) finds that less efficient local regulations will attempt to compensate when a national agency fails to properly regulate pollutants. Several studies (Böhringer, Koschel, & Moslener, 2008; Böhringer & Rosendahl, 2010) have examined overlapping regulation in European energy markets finding that they decrease efficiency. (Levinson, 2010) suggests similar inefficiencies are likely to

[5] Two other papers study the outcomes of feebate policies in Europe. (D'Haultfoeuille, Durrmeyer, & Février, 2012) estimate the fuel economy increase caused by the policy in France. (Klier & Linn, 2012c) use the sales changes in European countries under a feebate to find the shadow cost of compliance with a standard.

[6] One reason for this is that (Small, 2012) chooses a fee rate to meet the same fuel savings as the CAFE standard.

occur in the US market if permit trading overlaps with command and control policies as is often proposed.[7]

Methodologically, my study differs from previous models of supplier behavior by allowing new vehicle manufacturers to adjust product characteristics in a framework that also included used vehicles and scrappage markets. I also model the VMT decisions of households, which allows me to accurately capture changes in total emissions.[8] Because my model incorporates a production possibility frontier, I am able to shock this frontier with improvements from an R&D program and illustrate the effects on all associated markets, which is not possible in earlier models.

To conclude, I perform several simulations that illustrate how a CAFE standard and feebate differ in the presence of two complementary policies.[9] Specifically these simulations compare how a CAFE standard and feebate interact with an efficient vehicle tax credit or a national R&D policy. My model allows me to capture how leakage patterns differ in each of these scenarios in terms of changes to average fuel economy of the fleet and total emissions. I find that while the CAFE standard undermines the intended fuel economy effects of these complementary policies and increases total emssions, the interaction with a feebate results in higher average fuel economy, but these gains are insufficient to offset higher vehicle sales and VMT resulting in higher emissions. This suggest that while feebates may be the optimal baseline policy, better coordination of these policies is required to ensure reduced emissions.

[7] A related literature (Kolstad, Ulen, & Johnson, 1990; Burrows, 1999) considers overlapping legal regulation, finding that it can improve efficiency. For an overview of other overlapping environmental policies see (Fischer & Preonas, 2010).

[8] (Klier & Linn, 2012a) is the most similar but does not model used or scrappage markets and does not model the VMT decision of drivers.

[9] In future work I plan to explore the optimality of each policy focusing instead on uncertainty.

The remainder of the paper is organized as follows. Section 2 derives the theoretical equivalence between a general fuel economy standard and a feebate and discusses why this equivalence fails to hold in the presence of complementary polices. Section 3 details the specification of the model used in estimation and simulation, section 4 details the data used, and section 5 the estimation of demand using that data. Section 6 then presents an overview of the simulation and the results and section 7 concludes.

1.2 Fuel Economy Standards versus a Feebate

1.2.1 Theoretical Equivalence

Although there are many variations on fuel economy standards[10] and feebates, initially I focus on the most simple formulation of each. For the fuel economy standard, the firms must meet a sales-weighted average fuel economy greater than or equal to the level mandated by the government. The simplest formulation for a feebate consists of two elements: a fee rate and a pivot point. These two elements provide for a fee schedule where vehicles are awarded a proportional rebate for efficiency over the pivot point and a proportional fee when their efficiency is less than the pivot point.[11] I limit the scope of my analysis to feebates that are linear in fuel efficiency. Initially I only consider revenue-neutral

[10]I use 'fuel economy standard' to designate a generic standard and CAFE to designate the specific formulation used in the United States, which, for example, has separate standards for cars and trucks.

[11]For example, under a feebate with a pivot point of 0.03 gallons-per-mile and a rate of $1000 per 0.01 gallons per mile, a Lexus ES 350 at 0.045 gallons per mile would be assessed a $1500 fee while a Toyota Prius at 0.020 gallons per mile would receive a $1000 rebate. When policies like this are enacted, they are often not a smooth linear function as modeled here but a schedule of tax 'notches' that can create distortions (Sallee & Slemrod, 2012).

feebates although simulations in later sections allow for non-revenue-neutrality. Because a one unit mile-per-gallon improvement is not an equal improvement in gasoline savings across the entire range of fuel economy ratings, fuel economy standards are calculated using the harmonic mean and feebates are calculated based on the gallon-per-mile rating of the vehicle.[12]

Consider an economy in which K identical firms produce J different types of vehicles for a large population of consumers of size N. Each consumer is endowed with income I_o and owns an equal share of the profits of each firm. Consumers get utility from vehicles and the consumption of a numeraire good x.

Consumers are divided into J groups depending on their preferred vehicle type. Let N_j denote the size of the group of consumers who prefer type-j vehicles and that type-j consumers only consume type-j vehicles, getting zero utility from any other vehicle. To simplify the analysis I assume that utility is linear in the numeraire good. Each consumer is characterized by the pair (w, j), where j is the preferred vehicle type, and w is the willingness to pay for that vehicle, which is distributed uniformly $w \sim U[\underline{w_j}, \overline{w_j}]$.

The economy will produce an allocation of vehicles and numeraire good amongst consumers. Let it be described by the mapping $\{q(w, j), x(w, j)\}$ where:

$$q(w, j) = \begin{cases} 1 & \text{if } (w, j) \text{ owns a vehicle of type } j \\ 0 & \text{otherwise} \end{cases}$$

[12]Using the mpg rating of the vehicle provides non-linear benefits for each unit increase. A one unit mpg improvement to a low-fuel-economy vehicle provides larger gasoline and operating cost reductions than those achieving higher fuel economy (Larrick & Soll, 2008). The harmonic mean is used to calculate the sales weighted fuel economy for CAFE in the United States. Evaluation of efficiency in liters-per-100 kilometers as is done in European and Asian countries requires the use of the arithmetic mean to calculate these averages.

and $x(w, j)$ is consumption of the numeraire. The different policy instruments are evaluated based on the allocations they produce.

Utility maximization implies that a type-j consumer will purchase a type-j vehicle if his willingness to pay exceeds the price of that vehicle, p_j. Given the uniform distribution, the demand for vehicle j is

$$q_j(p_j) = N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} \tag{1.1}$$

Turning to the firms, I assume that each of the K firms, indexed by k, produces with the cost function $c(q_{k1}, \ldots, q_{kn}) = \sum_{j=1}^{n} c_j(q_{kj})$.[13] For each vehicle j, the cost function $c_j(q_j)$ is assumed to be increasing, smooth, and strictly convex. Let $mc_j(q_j) = c'_j(q_j)$.

Assume that the fuel economy is fixed for each vehicle. It can be expressed either as the mile per gallon rating mpg_j or its inverse g_j, the gallon per mile rating, such that $mpg_j = 1/g_j$. Firms operate in a competitive market.

Unregulated Case

Where the firms are not subject to regulation by the government a representative firm optimizes over a vector of quantity $\boldsymbol{q_k} = \{q_{k1}, \ldots, q_{kn}\}$ of goods to sell at market-determined prices $\boldsymbol{p} = \{p_1, \ldots, p_J\}$. The firm's problem is formulated as

$$\max_{\boldsymbol{q_k}} \sum_{j=1}^{J} [p_j q_{kj} - c_j(q_{kj})] \tag{1.2}$$

[13]In a setting with identical firms, tradability of permits is not required but will be if firms are heterogeneous.

This optimization implies that

$$p_j = mc_j(q_{kj}) \text{ or } q_{kj} = mc_j^{-1}(p_j) \tag{1.3}$$

Market supply of vehicle j is

$$q_j^S = Kq_{kj} \tag{1.4}$$

To clear the market for vehicle j, the quantity demanded equals the quantity supplied when

$$q_j^D = N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j) = q_j^S \tag{1.5}$$

Since mc_j is monotonically increasing there exists a unique p_j that sets supply equal to demand.

Fuel Economy Standard

Now let the government subject firms to a fuel economy standard. The sales-weighted harmonic mean of the vehicles produced must be greater than or equal to the mandated fuel economy level, $\overline{mpg}$. Under the standard, firm k chooses a vector of quantities $\boldsymbol{q_k}^c = \{q_{k1}^c, \ldots, q_{kJ}^c\}$ to solve

$$\max_{\boldsymbol{q}} \sum_{j=1}^{J} [p_j q_{kj} - c_j(q_{kj})] \tag{1.6}$$

$$s.t.$$

$$\overline{mpg} \leq \frac{\sum_{j=1}^{J} q_{kj}}{\sum_{j=1}^{J} q_{kj} \frac{1}{mpg_j}}$$

The constraint, which represents the fuel economy standard, can be significantly simplified using the gallon-per-mile rating of the vehicle, g_j and the gallon-per-mile equivalent of the fuel economy standard $\overline{g} = 1/\overline{mpg}$.

$$0 \leq \sum_{j=1}^{J} (\overline{g} - g_j) q_{kj} \tag{1.7}$$

The first order condition for product j is

$$p_j = mc_j(q_{kj}) - \lambda(\overline{g} - g_j) \tag{1.8}$$

where λ is the shadow cost of the fuel economy constraint. This implies that firms choose quantity according to:

$$q^c_{kj} = mc_j^{-1}(p_j + \lambda(\overline{g} - g_j)) \tag{1.9}$$

Equilibrium is determined by:

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j + \lambda(\overline{g} - g_j)) \tag{1.10}$$

The prices that solve this equation are identical for consumers and producers. Denote that vector by $\boldsymbol{p^c} = \{p^c_1, \ldots, p^c_J\}$ and the resulting allocation by $\{q^c(w, j), x^c(w, j)\}$

Feebate

Alternatively let the government subject the firms to a feebate with rate R. Denote the pivot point, in gallons per mile, that results in revenue neutrality as g_0. Consumers face the tax/rebate inclusive price p_j. When the fee $-R(g_0 - g_j)$ is levied on vehicle j, the firm faces the price $p_j + R(g_0 - g_j)$ and chooses $\boldsymbol{q^f_k} = \{q^f_{k1}, \ldots, q^f_{kJ}\}$ to solve

$$\max_{\boldsymbol{q_k}} \sum_{j=1}^{J} [(p_j + R(g_0 - g_j))q_{kj} - c_j(q_{kj})] \tag{1.11}$$

which gives the first order condition

$$q^f_{kj} = mc_j^{-1}(p_j + R(g_0 - g_j)) \tag{1.12}$$

Equilibrium is determined by:

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j + R(g_0 - g_j)) \tag{1.13}$$

Denote the consumer price vector that solves these equations as $\boldsymbol{p^f} = \{p_1^f, \ldots, p_J^f\}$. Note that the producer prices, unlike in the case of the fuel economy standard, are not identical to consumer prices. Denote the resulting allocation as $\{q^f(w, j), x^f(w, j)\}$.

Feebates versus Standards

It can be seen from the first order conditions of the firm under each regulation, equations 1.9 and 1.12, that either instrument can be used to create a price wedge that increases the consumer price on inefficient vehicles and decreases the consumer price on efficient ones. In both cases this results in a higher average efficiency of the vehicles sold in the market. This leads to two propositions:

PROPOSITION 1: *An allocation $\{q^c(w, j), x^c(w, j)\}$ achieved under fuel economy standard $\overline{mpg}$ can be replicated by setting fee rate $R = \lambda$.*

PROPOSITION 2: *An allocation $\{q^f(w, j), x^f(w, j)\}$ achieved under feebate rate R and resulting revenue neutral pivot point g_0 can be replicated using a fuel economy standard by setting a standard of $\overline{mpg} = 1/g_0$.*

Propositions 1 and 2 follow since the price wedge created is equivalent under both policies. This results in identical consumer prices and quantities of vehicles sold. What is less intuitive is the equivalence in the firms' profits and therefore the resulting allocation of the numeraire good. This relies on the revenue neutrality of the feebate and that the constraint under the fuel economy standard is exactly binding on the firms. This equivalence is demonstrated in more detail in Appendix A but can be described graphically.

Diagrammatic Intuition

Figure 4.5 illustrates firm profits for a simplified case of the model above where $j \in \{1, 2\}$. Vehicle 1 is a high performance, low efficiency vehicle, while vehicle 2 is a low performance, high efficiency vehicle. Panel A depicts the market equilibrium where firms are not subject to regulation. The profits from vehicle 1 are depicted by the shaded Region A, while profits from vehicle 2 are depicted by Region B.

Panel B depicts profits where the firms are subject to a fuel economy standard. The price wedge on vehicle 1 results in higher consumer and producer prices and reduces the quantity sold. Profits from this vehicle are now $C + D$. The standard reduces prices for vehicle 2, increasing total sales. This results in profit losses in Region F. Profits from this vehicle are now $G - F$.

Panel C depicts profits where the firms are subject to a feebate. The fees collected on vehicle 1 result in a difference between consumer prices P_1 and producer prices P_1^*. Region I is no longer part of firm profits and instead is collected to finance rebates on vehicle 2. These rebates increase the producer price P_2^* above the consumer price P_2 resulting in total firm profits on this vehicle as $L + J$.

The diagrams show how the price wedge created under either policy results in identical consumer prices and total quantities sold. To see that the total profits are identical under the two regimes, first note that the binding fuel economy standard ensures that $D = E + F$. Similarly the revenue neutrality of the feebate ensures that $I = J + K$. If we remove the Region D profits from vehicle 1 and apply them to vehicle 2, the diagram of profit and losses would be that of vehicle 2 under the feebate: F is no longer a profit loss and becomes Region K, while E becomes increased profits as in Region J.

Despite this equivalence, it is important to note that a key distinction between these policies lies in who controls the creation of the price wedge, and the ability of that entity to react to changes in the market. In the case of the feebate the price wedge is directly created by the government, while in the case of the fuel economy standard, it is generated by the firms to achieve the standard mandated by the government.

1.2.2 Interaction with Complementary Policies

The equivalence shown above is predicated on the complete flexibility by the government to change the target of the baseline policy in response to any event in the market. Federal policies such as CAFE are often set in advance for many years at a time. In the more than 30-year history of CAFE, the standard has only been updated a few times. Once set, these policies are minimally adjusted to account for other developments in the market. Additionally, other agencies or states may choose to enact polices that address the under provision of fuel economy in the market. A multitude of these policies exist at the state and national level, including various subsidies for the production of these vehicles, tax incentives, and special privileges for fuel efficient vehicles. Little thought is given when these policies are implemented as to how they interact with a national fuel economy standard. First I illustrate the effects of a fuel efficient tax credit, inspired by the demand policy shock examined in (Goulder et al., 2012) under CAFE. I then examine how outcomes would differ under a feebate. Finally I briefly discuss the effects of a shock to the production possibility frontier.

To illustrate this difference in outcomes I continue with the special case where firms produce only two vehicles indexed by $j \in \{1, 2\}$. In Panel A of Figure 1.2 either a feebate or fuel economy standard is modeled. Vehicle 1 is the inefficient vehicle and vehicle 2 is the efficient vehicle. For inefficient vehicles, the wedge a-b will be generated by a feebate of $R(g_o - g_1)$, or the equivalent fuel economy standard $\lambda(\overline{g} - g_1)$. For efficient vehicles, wedge c-d will be of size $R(g_o - g_2)$, in the case of a feebate, and $\lambda(\overline{g} - g_2)$ in the case of fuel economy.

Now let an agency enact a tax credit aiming to increase demand for the fuel efficient vehicle. Panel B depicts these effects under a fuel economy standard. To maintain the required average, $\overline{g}$, initially the firm may simply maintain output of the efficient vehicle at Q_2. Should the firm do this, there will be no change in vehicles sold in the market and the firm will entirely capture the subsidy with increased prices. This outcome represents a best case scenario in terms of changes to emissions as there is no increase in sales and the average is maintained. This, however, is not the least cost solution for the firm as the firm has not set λ equal across all products. The increased demand for efficient vehicles allows the firms to decrease the shadow cost of the fuel economy standard constraint on the inefficient vehicle. Now the firm has increased the sales of inefficient products to Q_1' and also of efficient vehicles to Q_2'. Thus total emissions are now larger as more vehicles are on the road with the same average fuel efficiency as before the tax credit.

The effects of the tax credit under a feebate are illustrated in Panel C. The case illustrated here assumes the government cannot change the rate or pivot point in response to the credit. Because neither R nor g_o have changed, $R(g_o - g_2)$ remains the same and the price wedge from e-f is preserved in g-h after the demand increase. No changes occur for the inefficient product as the price wedge remains $R(g_o - g_1)$. Together these results imply that although more cars are on the road, they are of higher average fuel economy. If however the government does shift the pivot point to maintain revenue neutrality, this only improves the outcome. This case (not illustrated here) would shift the pivot point up to maintain revenue neutrality. This shift, towards the more efficient vehicle, would reduce the price wedge on the efficient product and increase it on the inefficient product. This would mute some of the overall quantity expansion, yet still maintain the higher average fuel economy relative to the initial case. Whether the pivot point adjusts or not, it is clear that the overall effect is ambiguous and may result in lower emission while it must be the case that emissions are higher under a fuel economy standard.

The second type of complementary policy I model affects the ability of firms to supply fuel efficiency. To illustrate this case, I examine a government sponsored R&D program that increases the production possibility frontier (PPF) for domestic firms. An example of such a policy occurred as recently as August 2011 when the US Deptartment of Energy awarded $175 million to domestic firms to improve vehicle fuel economy. Assuming these grants are successful, one might expect the fuel efficiency of all vehicles to improve. These gains may not, however, materialize if firms can channel these improvements towards weight and horsepower, attributes that are valued for their safety and performance. This decision is highly dependent upon the baseline policy chosen by the EPA.

Although this type of policy shock shares some similarities to the tax credit, its effects are somewhat subtler. In an unconstrained world these improvements would be channeled in constant ratios towards fuel economy, horsepower, and weight. Both a fuel economy standard and a feebate create a price wedge that increases the valuation towards fuel economy. The crucial distinction is that under the standard, this wedge is set by the firm while under the feebate it is set by the government. When the PPF of the firm expands in the presence of the fuel economy standard constraint, the firm will never exceed the standard unless the expansion is so great that the standard is no longer binding. Firms will channel these improvements towards the characteristics consumers value most. Sales will increase as all products become more attractive, while the fleet average fuel economy is held constant. The feebate does not allow for a reduced price wedge, implying that some portion of these improvements will be channeled towards fuel economy, increasing the fleet average. Any improvements will, of course, make vehicles more desirable and increase sales so the effect of the policy on total emissions is ambiguous: more efficient vehicles may reduce emissions, but more sales of all vehicles may increase them.

The lesson from both of these complementary policies is clear when interacting with a fuel economy standard: they must increase total emissions and will have no effect on the average fuel economy of the fleet. By contrast the outcomes under a feebate are less clear: while both complementary policies result in higher average fuel economy, they may also increase total sales, resulting in an ambiguous effect on total emissions.

It should be noted that these predictions ignore many idiosyncrasies of how fuel economy standards are enacted under CAFE. For example CAFE is only binding on domestic firms, and the results above assume that the complementary policies only target the CAFE constrained, domestic firms. Many policies, like that studied in (Goulder et al., 2012), will target all efficient cars—those produced by constrained firms as well as foreign, unconstrained firms. This simple analysis also ignores the substitution consumers make between the efficient and inefficient vehicles. Several other particular details of how CAFE is constructed affect these conclusions. The different standard for cars versus trucks provides one important example. In cases where a policy changes the ratio of cars to trucks, the overall fuel economy of products sold may change although the individual car and truck fleets attain the same average. Broader market effects may also have implications for both CAFE and feebates. Market power may cause other firms to change their fuel economy or quantity decisions in ways that are hard to predict but will have implications for overall emissions. To simulate all of these interactions, I now turn to building a detailed model of the automobile market.

1.3 Model Specification

The model consists of five agents: new vehicle producers who maximize profits, households who maximize utility, used vehicle dealers, scrappage dealers, and the government. Households make ownership and vehicle miles traveled (VMT) decisions. New vehicle producers choose price, horsepower, and weight of the vehicles they produce to maximize profits. Because they face a production possibility frontier (PPF) when making their decision over product attributes, the choice of horsepower and weight determines the fuel economy of the vehicle. Vehicles are aggregated in the simulation by manufacturer, class, and age.

Within each of the seven manufacturers, there are nine potential classes of vehicles and four age classes, as listed in Table 2. The used car market adjusts prices such that the supply of all used cars (after scrappage) is sold. The scrappage market removes cars based on the age and price of the car using estimates from Bento, Roth, and Zuo (2012). Where a feebate is simulated, the government sets the pivot point to maintain revenue neutrality.

1.3.1 The Supplier Model

Suppliers maximize profits in a Bertrand Oligopoly framework. While this model accounts for some product innovation in the medium run, it does not account for drastically new types of product offerings. It only allows for two types of innovation. The first is changes along the PPF between horsepower, weight, and fuel economy, which are assumed to be within control of the firm. The second type of change that can occur is an expansion of the PPF, which is assumed to occur exogenously. It also has been shown to expand faster when regulatory pressure increases (Knittel, 2011). I also assume that an exogenous R&D policy has the ability to shift the PPF outward.

There are K firms indexed by $k = \{1, \ldots, K\}$. Products are indexed by $j = \{1, \ldots, J\}$, with the subset of products produced by firm k denoted by J_k. The vector of prices is denoted by $\boldsymbol{p} = \{p_1, \ldots, p_J\}$, horsepower by $\boldsymbol{h} = \{h_1, \ldots, h_J\}$, weight by $\boldsymbol{w} = \{w_1, \ldots, w_J\}$, and fuel economy, in gallons-per-mile, by $\boldsymbol{g} = \{g_1, \ldots, g_J\}$

Firm f maximizes profits by choosing price and product attributes for the products it produces:

$$\max_{p_j,h_j,w_j,g_j} \sum_{j\in J_k}^{J}(p_j - c_j(g_j, h_j, w_j))q_j(\boldsymbol{p}, \boldsymbol{h}, \boldsymbol{w}, \boldsymbol{g}) \tag{1.14a}$$

$$s.t.$$

$$ln(g_j) = \omega_1 ln(h_j) + \omega_2 ln(w_j) + G_j \tag{1.14b}$$

$$ln(c_j) = \mu_1 ln(h_j) + \mu_2 ln(w_j) + \mu_3 ln(g_j) + C_j \tag{1.14c}$$

Marginal cost, c_j, is a function of the vehicle product attributes, but is constant in quantity. Quantity of product j demanded, q_j, is a function of the prices and attributes of all products produced in the market. Constraint 1.14b defines the PPF faced by the firm and determines the tradeoff between horsepower, weight and fuel economy. The ω values are estimated using historical data and the G_js are calibration parameters. Constraint 1.14c defines the marginal cost of producing this vehicle. The μ parameters are based on estimates from (Berry, Kortum, & Pakes, 1996) and the C_js are calibration parameters.[14]

I distinguish between seven producers. Following (M. R. Jacobsen, 2012a), firms are heterogeneous and produce products of differing fuel economy. Under CAFE, these differences expose them to different regulatory environments. Some firms, like Toyota, produce sales weighted fuel economy above the standard required by CAFE and therefore face the unregulated producer problem as outlined in maximization problem 1.14. Other firms, in an unconstrained world, would produce fuel economy averages below the standard. Historically, these have been the domestic manufacturers Ford, GM, and Chrysler. These firms are

[14] I choose to allow marginal costs to adjust to changes in product characteristics, although the values for these parameters vary substantially across the literature. In particular it is unclear which direction decreased weight should change marginal cost. Weight can be decreased by reducing the amenities (and cost) of the car, or it can increase cost by replacing metal parts with more costly, lightweight materials (Cheah, 2010). (Berry et al., 1996) found that cost decreased as weight was reduced when CAFE was altered in the past.

subject to the profit maximization problem:

$$\max_{p_j,h_j,w_j,g_j} \sum_{j\in J_k}^{J}(p_j - c_j(g_j, h_j, w_j))q_j(\boldsymbol{p}, \boldsymbol{h}, \boldsymbol{w}, \boldsymbol{g}) \tag{1.15a}$$

$$s.t.$$

$$ln(g_j) = \omega_1 ln(h_j) + \omega_2 ln(w_j) + G_j \tag{1.15b}$$

$$ln(c_j) = \mu_1 ln(h_j) + \mu_2 ln(w_j) + \mu_3 ln(g_j) + C_j \tag{1.15c}$$

$$\sum_{j\in J_{Car}}^{J} q_j(\boldsymbol{p}, \boldsymbol{h}, \boldsymbol{w}, \boldsymbol{g})[\overline{g_{Car}} - g_j] = 0 \tag{1.15d}$$

$$\sum_{j\in J_{Truck}}^{J} q_j(\boldsymbol{p}, \boldsymbol{h}, \boldsymbol{w}, \boldsymbol{g})[\overline{g_{Truck}} - g_j] = 0 \tag{1.15e}$$

The terms $\overline{g_{Car}}$ and $\overline{g_{Truck}}$ are the fuel economy levels targeted by CAFE. Historically, CAFE has had separate targets for cars and trucks. For most of the last 30 years, this has been set at 27.5 mpg (0.036 gpm) for cars and 20.7 mpg (0.048 gpm) for trucks. These firms face three types of constraints. Constraint 1.15b is the PPF, constraint 1.15c is the marginal cost equation, and the constraints 1.15d and 1.15e are the CAFE constraint for cars and trucks. The Lagrange multiplier for constrains 1.15d and 1.15e are the shadow costs of the regulation, which are equal for all cars or trucks. The problem of the constrained firms can also be used for the problem of the unconstrained firms; however, because their fuel economy exceeds the standard, the multiplier is zero effectively removing the CAFE constraints.

Recently, the CAFE standard has been reformulated to adjust based on the footprint of the vehicle fleet of the manufacturer. I do not adjust the standard for these changes for two reasons. First, I do not model changes to the footprint of each vehicle, preventing me from precisely predicting the new target. Second, there are additional allowances for environmentally beneficial changes such as reductions in hydrofluorocarbons, which change the standard.

Under a feebate, each vehicle is assessed a fee or rebate, and that is added to the price of each vehicle. The manufacturers' problem in this regulatory environment is given as

$$\max_{p_j,h_j,w_j,g_j} \sum_{j\in J_k}^{J}(p_j - R(g_o - g_j) - c_j(g_j,h_j,w_j))q_j(\boldsymbol{p},\boldsymbol{h},\boldsymbol{w},\boldsymbol{g}) \tag{1.16a}$$

$$s.t.$$

$$ln(g_j) = \omega_1 ln(h_j) + \omega_2 ln(w_j) + G_j \tag{1.16b}$$

$$ln(c_j) = \mu_1 ln(h_j) + \mu_2 ln(w_j) + \mu_3 ln(g_j) + C_j \tag{1.16c}$$

where R is the rate of the feebate and g_o is the pivot point, both established by the government.

1.3.2 Household Demand

The households in the model derive utility from their choice of vehicles and VMT each year. To account for multicar households, each household is modeled as facing multiple choice occasions.[15] On each choice occasion, the household chooses a vehicle or the outside good. The outside good includes any transportation option not explicitly modeled, including public transportation. The budget of each household used to purchase vehicles includes exogenous income as well as the value of the initial car endowment. Utility is modeled according to (Bento et al., 2009):

$$V_{ij} = \frac{-1}{\alpha}exp\left(-\alpha\left(\frac{y_i}{T_i} - r_{ij}\right)\right) + \frac{1}{\beta}exp\left(-\beta p_{ij}^M + \gamma_{ij}D_iX_j\right) + \tau_{ij}D_iX_j + \psi_jX_j \tag{1.17}$$

[15]The use of the term 'choice occasion' simply implies that households face multiple opportunities to buy vehicles. Because these choice occasions are independent of one another they do not allow for vehicle choices on one choice occasion to influence the utility on another choice occasion. (Spiller, 2010) relaxes this assumption and finds that it produces a slight downward bias in the gasoline elasticity of VMT although the effect is smaller than the bias from aggregation of the product space and from omitting vehicle fixed effects, which I account for.

Parameters α,β, γ_{ij}, τ_{ij}, and ψ_j are estimated parameters.[16] Income, y_i, is divided among T_i choice occasions and r_{ij} is the rental rate, which includes depreciation and the foregone interest of borrowing the value of the car for a year. Income also includes the household's share of producer profits, government transfers (particularly from a non-revenue neutral feebate), and any capital gains or losses based on used vehicle holdings. Interactions between household characteristics, D_i, and vehicle characteristics, X_j, enter both the linear and non-linear portion of the utility function. The operating cost of the vehicle, p_{ij}^M, is the price per mile to drive the vehicle and includes gasoline costs as well as any government taxes.

Applying Roy's identity to the indirect utility function gives the demand for VMT.

$$VMT_{ij} = exp\left(\alpha\left(\frac{y_i}{T_i} - r_{ij}\right) + \beta p_{ij}^M + \gamma_{ij} D_i X_j\right) \tag{1.18}$$

where the terms α,β, and γ_{ij} are parameters identical to those in the utility function.

1.3.3 Used Vehicles, Scrappage Markets, and the Government

The used car market consists of all cars that are not new or scrapped. Vehicles are scrapped on a probabilistic basis. The fraction of scrapped cars of type j is given by θ_j and adjusts according to $\theta_j = b_j(p_j)^\eta$. Scrappage consists of two components, technical scrappage and scrappage due to price. Technical scrappage, b_j, is scrappage due to aging and is calibrated to match the technical scrappage rates estimated in Bento, Roth, and Zuo (2012). The scrappage due to price

[16]The parameters α and β are restricted to be positive. These restrictions are made to ensure that income is positively related to utility and that price is negatively related to utility.

adjusts based on the price elasticity parameter, η. This value, also taken from Bento, Roth, and Zuo (2012) is estimated at -0.71.

For a feebate, the government must balance its budget. It can do this two ways. Either it can choose a pivot point g_o such that.

$$\sum_{j=1}^{J}(R(g_o - g_j)q_j(\boldsymbol{p}, \boldsymbol{h}, \boldsymbol{w}, \boldsymbol{g})) = 0 \tag{1.19}$$

or it can raise revenue to make up a deficit, or rebate the excess. Because the behavior I am interested in assumes the government cannot adjust the CAFE standard or feebate, I do not allow the government to change the pivot point after balancing its budget in the initial condition; rather, it must come from a uniform lump sum tax.

The solution to the model is a vector of used car prices, new car prices, new car horsepower, new car weight, and a pivot point. In equilibrium these values ensure that manufacturers have maximized profits, each non-scrapped car has a buyer, and, in the case of feebates, that the policy remains revenue neutral. To implement the model, the parameters, particularly of the demand model, must be estimated using observed data.

1.4 Data

The data used to estimate the demand model come primarily from two sources, the National Household Transportation Survey (NHTS) and Ward's Automotive Yearbooks. Additional data sources include National Automobile Dealers Association (NADA) used car price data, gasoline price data from the American Chamber of Commerce Researchers Association (ACCRA) and national level sales data from R.L. Polk. Table 3 presents descriptive statistics for these data sets.

Information on household demographics and their vehicle choices come from the 2001 and 2009 waves of the NHTS. These data contain 150,147 households and detail information on family size, education, the number of workers, children, and drivers in the household, as well as the location of the household, allowing for classification into rural or urban status and region of the country. The data identifies the vehicle choices made by the household based on make, model, and vehicle age.[17] I exclude households with missing income or gender as well as residents of dormitories. I also drop vehicles whenever they are unidentified or mileage is missing. After cleaning, 150,134 households remain, together owning 316,164 vehicles.

Vehicle attribute data come from Wards Automotive Yearbooks 1972-2010. For the primary years used in the demand estimation (2000-2006), it provides price, horsepower, weight, and fuel economy. For these years there is provide information on 1,523 vehicles. The full sample of vehicles used to estimate the PPF contains information on 43,839 vehicles classified by make, model, and trim.

Vehicle data are supplemented with several other data sets. National sales data from R.L. Polk provide annual new car sales from 2000-2006. These sales data are used to create annual national shares of each product. NADA used car data are used to recover the yearly depreciation rates. This data set contains 6,776 vehicles at the model level for model years 1990 through 2010. Percent de-

[17]As vehicles age, the identification becomes decreasingly credible and this motivates me to generate used car quantities and characteristics based on new car samples in a method described below.

preciation is constructed as the percent of value lost over the course of the year. Although I allow this value to vary based on the make and class of the vehicle, I assume it is constant over the lifetime of the vehicle. The mean depreciation rate across all vehicles is calculated as 16.2%. The rental rate of the vehicle is calculated as $r_j = D_j + \rho P_j$ where D_j is the depreciation rate and ρP_j is the forgone interest of owning the car for a year. I set ρ at 3.21%, the average return on t-bills from 2000-2006. Finally, ACCRA data on monthly average gasoline price by state are used to calculate per-mile gasoline costs. These data, taken from 2001 and 2009, are averaged across months, and the average price of gasoline is taken as the price for gasoline in the state where the household lives.

1.5 Estimation

1.5.1 Empirical Model

To simulate counterfactual polices, I require a model of demand. The goal of this estimation is to specify as parsimonious a demand specification as possible, while still addressing several welfare calculation and estimation concerns. These concerns are specifically 1) to jointly estimate the vehicle choice and VMT demand in a welfare consistent framework, 2) to include sufficient heterogeneity among the agents to capture the relevant substitution patterns, and 3) to control for the endogeneity between price and other unobserved product characteristics. To address these concerns, the method of demand estimation follows two literatures. The first jointly estimates vehicle choice and VMT (Bento et al., 2009; Feng, Fullerton, & Gan, 2005; Spiller, 2010; Gillingham, 2010). The second literature estimates vehicle choice from market level data (Berry et al., 1995; Nevo, 2001; Petrin, 2002). Consumer demand starts with the conditional indirect utility model

$$V_{ij} = \frac{-1}{\alpha} exp\left(-\alpha\left(\frac{y_i}{T_i} - r_{ij}\right)\right) + \frac{1}{\beta} exp\left(-\beta p_{ij}^{M} + \gamma_{ij} D_i X_j\right) + \tau_{ij} D_i X_j + \psi_j X_j + \varepsilon_{ij} \quad (1.20)$$

$$where$$

$$\alpha = exp(\tilde{\alpha})$$

$$\beta = exp(\tilde{\beta})$$

Each household chooses between the available products and the outside good for each choice occasion. The outside good j=0 includes all other transportation options. The inclusion of this outside good allows for the total demand for all vehicles to be downward sloping, an important margin of adjustment for policies aimed at reducing fuel use.[18] J denotes the choice subset in a given year. Each household makes multiple discrete-continuous choices of vehicle and VMT each year. The number of choice occasions, denoted T_i, are taken to be the number of drivers in the household plus one. The error term, ε_{ij}, is a random taste shock that has a Type I, extreme-value distribution. To capture preference heterogeneity among the agents, I include interactions of household and vehicle characteristics in the spirit of (Goldberg, 1995). In particular, I try to create sufficient heterogeneity in the product space I am most concerned with by interacting horsepower and weight with numerous household characteristics.[19]

[18]At this stage I do not estimate used car demand due to lack of data on market shares, but I extend these estimates to the used car market in the simulation by a method detailed below.

[19]Fuel economy appears as part of the cost per mile of operation. Fuel economy cannot be interacted with other household characteristics to maintain the welfare consistent formulation that allows for the recovery of demand for VMT. Alternatively, I could use a random coefficient specification to model this heterogeneity, but even using efficient methods such as Halton sequences or quadrature would drastically increase the computational burden of the simulation below.

Take $\theta = \{\alpha, \beta, \gamma, \tau, \delta\}$, the estimated preference parameters. With the utility specification given above, the probability that household i chooses product j is given as

$$P_{ij} = P_i(j|p, X, D, \theta) = \frac{exp(V_{ij})}{\sum_h exp(V_{ih})} \tag{1.21}$$

The choice of VMT is assumed to be observed with a random error ϵ_{ij}, which is distributed standard normal:

$$VMT_{ij} = exp\left(\alpha\left(\frac{y_i}{T_i} - r_{ij}\right) - \beta p_{ij}^M + \gamma_{ij} D_i X_j + \epsilon_{ij}\right) \tag{1.22}$$

which gives rise to the likelihood function

$$l(\widehat{VMT}_{ij}|j \in J, j \neq 0) = \frac{1}{(2\pi)^{1/2\sigma_i}} exp\left(-\frac{1}{2}\left(\frac{\widehat{VMT}_{ij} - VMT_{ij}}{\sigma_i}\right)^2\right) \tag{1.23}$$

The joint likelihood that a household choice occasion i chooses vehicle j and a given VMT is given by

$$L_i = \prod_{j=1}^{J} P_{ij}(j)^{1_{ij}} \prod_{j=1}^{J} l(\widehat{VMT}_{ij}|j \in J, j \neq 0)^{1_{ij}} \tag{1.24}$$

where 1_{ij} is an indicator variable that equals 1 if j is chosen on that choice occasion and 0 otherwise.

1.5.2 Estimation

The above vehicle choice and VMT demand model is estimated in two steps. The first step uses maximum likelihood estimation with the nested contraction mapping algorithm from (Berry et al., 1995). Demand estimation follows from (Berry, Levinsohn, & Pakes, 2004), but is estimated using maximum likelihood similar to (Train & Winston, 2007) and (Langer, 2010). Let θ_2 be the set of non-linear parameters $\{\alpha, \beta, \gamma, \tau\}$.[20] Predicted market shares $\hat{S}_J(\theta_2, \delta)$ are matched to observed market shares S_j rather than observed sample shares to reduce variance.[21] This is particularly important with the 2009 wave of the NHTS in which some regions of the country are overrepresented and representative population weights are not available.[22] (Berry, 1994) shows that for a given θ_2, a unique δ exists that matches observed market shares with predicted market shares.

$$S_j = \hat{S}_J(\theta_2, \delta(\theta_2)) = \frac{\sum_i P_{ij}(\theta_2.\delta(\theta_2))}{N_i} \tag{1.25}$$

The term δ, which represents utility derived from the automobile that is invariant across households, is specified as $\delta_j(\theta_2, S_j) = \psi_j X_j + e_j$, where e_j represents utility not captured by observable characteristics.

Estimation of θ_2 proceeds in an iterative fashion. For a trial value of non-linear parameters θ_2, the contraction mapping recovers the mean utility δ_j for each product. Having recovered the household invariant utility, the non-linear parameters are solved by maximizing the likelihood function. The VMT equation is also estimated in this maximum likelihood stage. VMT demand is provided in the data only for the year of the survey. Therefore, the gas price used is taken from the survey year (rather than the model year, which is used for the vehicle choice).[23]

When θ_2 has been maximized and δ_j recovered, the linear parameters ψ_j are

[20]Some regressors in γ are not interactions with household characteristics. Because these regressors are exponentiated with other terms that do interact with household characteristics and enter independently into the VMT estimation, they are non-linear rather than linear parameters.

[21]To calculate observed shares, the total sales are divided by the number of choice occasions nationally. To approximate the national number of drivers per household (plus one) I add the adult population plus the number of total households from census data.

[22]Although this oversampling is not included in the documentation, it is clear from the summary statistics of the data.

[23]Although cross sectional variation does exist between states in the 2009 survey, a second cross section helps identify these VMT parameters with more precision. For this reason the 2001 survey is included in the estimation.

estimated from the vector $\boldsymbol{\delta}$. This set of moment conditions is based on the exogeneity assumption that the vector of product characteristics X_j is independent of the error term e_j. Of particular concern is that the error term e_j can be decomposed as $\xi_j + \varepsilon_j$, where ξ_j is correlated with price. In this case the exogeneity assumption fails and the bias leads to an underestimate of the price coefficient. This is a frequently raised concern because unobserved quality increases both price and market share. To identify θ_2, I instrument for price using product characteristics of vehicles within and outside the firm. As in (Whitefoot et al., 2012), I choose instruments that are less flexible in the short run, including vehicle footprint and length as well as the number of products in the same market segment, both within and outside of the firm.[24] I estimate demand for each model year in 2000-2006 as a separate market, having been limited to these years by the availability of data for new vehicle market shares. This allows households to choose from over 200 new cars in any given year. Having estimated the demand side, I can then solve for the marginal costs of each product as described below in the simulation section.

1.5.3 Estimation Results

Tables 1.3 and 1.5 report the results of the demand estimation. The household-vehicle interaction coefficients are presented in Table 1.3. The results from this estimation are generally intuitive and in line with expectations. For example, interpreting the γ-parameters that enter the VMT decision, households with more

[24]These instruments are unlikely to be perfectly exogenous, but as my model is predicated on the fact that price is chosen simultaneously with horsepower, weight, and fuel economy, I do not use these three characteristics as instruments, as is traditionally done. As can be noted from the price setting condition of the firms in equation 1.9, CAFE requires the firms to correlate price and fuel economy. However imperfect, the instruments I use represent an improvement over the OLS coefficient estimates, which result in upward sloping demand curves.

workers, and those in suburban and rural locations demand more VMT. The τ-parameters, which enter only the vehicle decision, are also generally intuitive. Households with at least one college degree prefer higher horsepower, lower weight vehicles. Households in the South and Midwest prefer domestic vehicles, while households on the coasts prefer European vehicles. One surprising, although statistically insignificant, effect is the interaction between Asian vehicles and Southern households being positive. This may be a result of the NHTS 2009 survey over sampling the urban areas of Texas, which may bias southern preference towards urban preferences. Alternatively, this may be due to the increased manufacturing presence of these companies in Southern states. Many of the terms are statistically insignificant, likely due to the large number of household interactions with horsepower and weight; however, α and β, which control the response to vehicle price and the price of driving, are statistically significant.

The household invariant parameters are presented in Table 1.5. Horsepower and weight have positive coefficients, which matches expectations that consumers will value these attributes (for their performance and safety, respectively). This relationship is also consistent with the tension necessary for suppliers who face an opportunity cost when increasing fuel economy at the expense of these attributes.

Table 1.6 presents several elasticities of interest. The sales-weighted price elasticity of demand is estimated at -7.5, which lies within the generally accepted range for these estimations. Gasoline-price elasticity of VMT is estimated at -0.45, which is higher than (Bento et al., 2009) estimate of -0.32, but lower than (Spiller, 2010) estimate of -0.62. It also lies within the range given by (Graham & Glaister, 2002), who found long-run estimates of the gasoline-price elasticity of

VMT ranging from -0.23 to -0.80. The results in Table 1.6 also suggest that urban drivers, who have more public transportation options, have more elastic demand than suburban or rural drivers, and that working households, who have fewer discretionary trips, have more inelastic demand than retired households.

1.6 Simulation

1.6.1 Assumptions and Calibration

To bring the estimated results to the simulation, several changes must be made. The demand model estimated above was estimated for new cars only and for more than 200 new car products per year. To use this level of disaggregation in the new and used car market is computationally challenging. Thus, products are aggregated into the class categories listed in Table 1.1. The attributes of these vehicles are the weighted sum of observed models in those classes. This provides the relevant product characteristics with the exception of the product specific fixed effect, which controls the market share of each product. Matching the demanded level of vehicles with the observed market shares generates these values.

The modeling of the used car market poses several challenges. The first involves the extension of the utility model to used vehicles. I make the assumption that the utility parameters for used cars are the same as new cars. This may result in the model predicting used cars as being more substitutable with new cars than they actually are. Over-substitutability would provide an upper bound on potential leakage to the used car market, as new car makers would have more market power than the model would predict. To determine the product fixed effect for these used vehicles, initial quantities are required.[25] To obtain these quantities for a given class for a previous model year, I assume the class had the same total sales in previous years but decrease those quantities to account for scrappage due to aging.[26]

Some calibration of the suppliers is also necessary. Two sets of cost are needed: the shadow cost of CAFE and the marginal cost of products. I take the shadow costs of CAFE from (Anderson & Sallee, 2011b), who estimate these values using the flex-fuel loophole. As this loophole decreases the standard for these firms, I calibrate at this lower level, but increase it to the normal value for my initial equilibrium. To recover the marginal costs of each product, I use the first order conditions of each firm with respect to price. To find these values, I solve the system of first order conditions such that they are all equal to zero.[27]

[25]Two other options are possible here, using the observed used cars in the NHTS dataset or using national shares observed. Identifying older vehicles in the NHTS data is unreliable. Additionally, current class distinctions become difficult to impose on older vehicles. Estimates of quantities of used vehicles are also unreliable and prohibitively expensive.

[26]Although this yields unrealistic quantities of some vehicles (for example, large numbers of older SUVs) it also has advantages. The assumption stabilizes the baseline simulation particularly if I simulate outcomes in future years and allow the fleet to age. In the cases where a product had not existed in the past and a new vehicle ages into that product category, price has to drop drastically to match the increasing supply with the nonexistent demand. By smoothing initial quantities, price fluctuations in later years that are unrelated to the regulatory environment would be minimized.

[27]These values will solve all the firms FOCs including those with respect to horsepower and weight. Because the FOCs with respect to horsepower and weight will not be exactly zero I find the residual and subtract it from each FOC for my starting conditions. This ensures that the initial starting equilibrium is stable. I could also perform this exercise using the FOC with respect to horsepower or weight instead of price but I opt for price as pricing is likely the most flexible for the firms and therefore the most likely to be exact.

1.6.2 Simulation Results

Efficient Vehicle Tax Credit for Domestic Cars

The first set of simulations models an efficient vehicle tax credit. This credit is modeled after the hybrid tax credits created as part of the Energy Policy Act of 2005. They range from $1,700 to $3,400. These credits were designed to support new technologies that risk-averse consumers with limited information may be hesitant to adopt. These credits were limited to the first 60,000 hybrid cars or trucks produced by each manufacturer. As the Asian manufacturers have already exhausted these credits, the vast majority currently go to domestic producers and I model this policy as a uniform $2,000 credit on the most efficient car produced by each domestic manufacturer.[28]

The results of this simulation are given in Tables 1.7 and 1.8. The interactions of the tax credit under the CAFE baseline are examined in Panel A of Table 1.7. As predicted above, the overall quantity of domestic vehicles sold under a CAFE standard increases. The individual fuel economy ratings of cars and trucks also remain constant before and after the policy. Despite this, the overall fuel economy of the domestic manufacturers increases by 0.47% to 0.58% as the policy, targeted at cars, increases the car fleet relative to the truck fleet. This result would reverse if the tax credit had instead targeted the most fuel efficient trucks. Market competition decreases overall sales for the foreign manufacturers, but it slightly increases the share of cars relative to trucks resulting in small improvements in the fleet fuel economy of foreign firms. Panel C of Table 1.8 displays the changes to VMT and gasoline and calculates welfare changes from

[28] I assume the revenues to finance this tax are raised from a uniform lump-sum tax.

these values. The interaction of the tax credit with CAFE results in an increase in total VMT and gasoline consumption.The first set of welfare calculations, which follows (Small & Rosen, 1981), incorporates changes in utility due to VMT and vehicle choice and is broken into changes in welfare for consumers of domestic and foreign vehicles. These large improvements for domestic firms occur because the subsidy on the most efficient vehicles relaxes the CAFE constraint for cars and allows manufacturers to channel the fuel economy of all cars, not just those targeted with the tax credit, towards other more valued characteristics. The associated externalities are calculated using the percentage changes in VMT and gallons of gasoline applied to the U.S. VMT and gasoline consumption in 2011 using costs from (Parry & Small, 2005). As both VMT and gasoline consumption increase, all externalities increase. These large costs overwhelm all other welfare gains associated with improved vehicle choice resulting in total welfare losses of $22.25 million.

The interactions of the complementary tax credit with a feebate baseline are examined in Panel B of Table 1.7 and Panel D of Table 1.8. I choose a rate of $500 per 0.01 gpm as this is a common benchmark rate chosen for the evaluation of feebate policies in the literature. The results in Panel B show that the increases in total quantity are on a similar order as those under CAFE in Panel A. What is very different, and one of the key results of my paper, is that the average fuel economy increases for the targeted car class. Combined with the increased share of cars, at 0.45% the overall fuel economy of the domestic manufacturers is nearly double of that under CAFE. The feebate either preserves or increases the fuel economy rating of every manufacturer and every class of vehicle in the presence of the efficient vehicle tax credit. This result does not, however imply that overall gasoline use decreases or that welfare increases under this interaction. The changes to VMT, gasoline consumption and welfare are depicted in Panel D of Table 1.8. While the complementary policy does increase the average fuel economy of the fleet, and a model holding VMT constant might suggest there would be gasoline savings, I find that the increased vehicle sales and rebound effect, result in higher total VMT and gasoline consumption. These increases are, however, smaller than those under CAFE. The total welfare improvement from vehicle choice is less than the interaction with CAFE, but the aggregation of associated externalities is considerably less as well. The total welfare loss is $149.01 million. Although the feebate is successful in increasing fuel economy and mitigates the emissions increases under this interaction, it is not necessarily better for welfare. It is noteworthy that total impact on welfare is due largely to changes in congestion, accidents, and local pollutants, all of which are VMT related externalities. Any policy that allows for an increase in VMT is therefore likely to be welfare reducing.

Research and Development Policy

In the final set of simulations, a complementary national R&D program expands the PPF of domestic manufacturers. R&D policies such as this may be enacted when private R&D spending is less than optimal due to positive information spill overs. While these policies aim to improve fuel economy, they are uncoordinated with the EPA and are not contingent upon meeting higher standards. For example, in August 2011 the Department of Energy awarded $175 million dollars to domestic firms (many of them automakers) attempting to "improve the fuel efficiency of next generation vehicles."[29] The amount by which these

[29]http://energy.gov/articles/department-energy-awards-more-175-million-advanced-vehicle-research-and-development. Although this particular set of grants was timed near

grants could expand the PPF is difficult to predict, but I choose an increase that allows manufacturers to produce 1% higher fuel economy vehicles while maintaining current levels of horsepower and weight.

The results of these simulations are compared in Tables 1.9 and 1.10. Table 1.9 shows how the R&D policy affects individual firm fleets. Panel A shows that under the CAFE standard, the three domestic manufacturers maintain the sales weighted fuel economy of both cars and trucks at pre-policy levels. One reason for this is that my model allows the manufacturers to channel this fuel economy improvement towards horsepower and weight. This, in turn, makes these vehicles more attractive to consumers, and increases overall sales for domestic firms by nearly 1%. In the absence of a policy prompting people to drive less, this will promote emissions leakage. One idiosyncrasy of CAFE is that it has different standards for cars and trucks. Although one might expect the net effect of the policy on fuel economy of these firms to be zero, it is not. The expansion of the PPF changes the relative composition of each firm's fleet between cars and trucks. In this simulation, the policy has a slight negative effect on fuel economy for Ford and GM. This is primarily because the truck sales are more important for these manufacturers and CAFE binds tighter on their truck than car fleets. As that constraint is relaxed due to the increase of the PPF, the truck segment expands faster than the car segment reducing overall fuel economy. This also affects the fuel economy of the foreign firms. These firms have a small overall increase in fuel economy as their sales shift toward cars.[30] The net result for new

the announced increase in the new CAFE standard, the firms awarded the money were not required to reach a higher target and no changes to CAFE have been proposed in response to the outcomes of this research.

[30]This confirms a result found in other papers (e.g. (Whitefoot et al., 2012)) that foreign, unconstrained firms increase sales of trucks when CAFE is tightened. In this simulation, the expansion of the PPF moves the market closer to the unconstrained equilibrium where domestic firms produce more trucks and foreign firms more cars.

vehicles is a decrease in fuel economy of 0.02%.

Panel C of Table 1.10 shows the VMT, fuel use, and welfare consequences of the policy. As might be expected, welfare increases for consumers of domestic vehicles. More surprising is the large decrease in welfare for consumers of foreign vehicles. This effect seems to arise because the domestic firms move in the product space allowing for foreign firms to capturing consumers left in the vacated product space. These shifts are particularly large under CAFE as the domestic firms make large changes to the profitable attributes of horsepower and weight. While consumers of domestic vehicles are better off, consumers of foreign vehicles are now purchasing a product less tailored to their individual tastes. Welfare is also negatively impacted by the externalities associated with VMT and gasoline consumption, which increase by 0.047% and 0.043% respectively. The total welfare losses due to the interaction of CAFE with the R&D policy are $276.79 million.

The results under a feebate are strikingly different and are reported in Panel B of Table 1.9. I find that fuel economy increases for the individual car and truck fleets of the domestic manufacturers engaged in the R&D program. The increase in fuel economy is nearly 0.60% for each of these manufacturers. I also find that quantity increases substantially less than under CAFE.

Panel D of Table 1.10 documents the VMT and gasoline consumption changes as well as the welfare consequences of this interaction. As noted above, domestic firms channel the improvements largely towards fuel economy resulting in profit increases between 0.05-0.07%, far smaller than under CAFE. Consumers may value fuel economy, but they value it less than horsepower and weight, which implies that the welfare improvements for consumers of these vehicles are more modest than under CAFE. Forcing these improvements towards fuel economy also mutes the effect of the interaction in other ways. The feebate provides less opportunity for shifting in the horsepower and weight dimensions of the product space. This means that very few welfare changes are experienced by consumers of foreign vehicles. It is perhaps surprising that, despite an increase in overall fuel economy of 0.37% and an almost negligible increase in sales quantity of 0.03%, overall VMT and gasoline use increase. When individuals drive more fuel efficient vehicles this lowers the price of driving, implying that VMT (and all related externalities) increase. I find that total VMT increases by 0.035% but the higher fuel economy of the fleet works to reduce the total gasoline consumed to drive those miles. Despite this, the higher quantity of sales and the extra miles travel produces an increase in gasoline consumption of 0.017%. While this fuel economy increase helps to limit the externalities associated with carbon, they are dominated by welfare losses arising from increased VMT. In total, welfare decreases $79.52 million when the R&D policy interacts with the feebate. In this simulation I find that the feebate is not only better at increasing fuel economy and mitigating emissions expansions, but also reduces the welfare losses implied by the interaction.

1.7 Conclusion

This paper develops a model to examine the relative merits of a CAFE standard versus a feebate in the presence of complementary but uncoordinated policies. Although *a priori* identical, they react differently to complementary policies. CAFE targets the average fuel economy of the fleet and allows the manufacturers to establish the price wedge, while a feebate targets the price wedge and allows the market to determine the average price. To compare their interaction with complementary policies, I build a model of the automobile market that pays particular attention to supplier decisions and incorporates used vehicle and scrappage markets. My model is the first to incorporate a production possibility frontier into the supplier decision in a general equilibrium model that also accounts for the used vehicle and scrappage markets. This allows me to examine how an increase in the production possibility frontier caused by a research and development program interacts with a CAFE standard or feebate, which has not been possible in previous general equilibrium models. I am also able to examine how a tax credit on efficient vehicles interacts with a feebate or CAFE standard. I find that while both complementary policies increase rather than decrease emissions, the increase is smaller under the baseline feebate policy because it improves the fuel economy of the fleet while a baseline CAFE standard does not.

To build the model used in these simulations I first estimate a demand model and calibrate the supply model using a variety of data sources. Parameters for vehicle and mileage demand are estimated simultaneously in a maximum-likelihood framework that accounts for correlation between unobserved product attributes and vehicle price. These parameters are estimated using detailed 2009 NHTS data that links household characteristics with vehicle choices. Vehicle data from Wards Automotive Yearbooks provide information on product characteristics and are used in both demand and supply estimation. The production possibility frontier of the suppliers, estimated using vehicle data from Wards Automotive Yearbooks between 1971 and 2010, allows manufacturers to endogenously choose horsepower, weight, and fuel economy of the vehicles they produce.

I simulate how several complementary policies attempting to further increase the fuel economy of the fleet interact with either a national CAFE standard or a feebate. I find that the tax credit increases VMT and fuel consumption under CAFE by 0.13% and 0.11% respectively. Comparatively, the feebate increases in VMT by 0.12%, but because it allows for average fuel economy improvements, it increases fuel consumption by only 0.08%. I find similar results for a complementary R&D program. While the expansion of the PPF for domestic manufacturers could allow them to produce vehicles of a higher fuel efficiency, I find that it is not necessarily in their best interest to do so. Under a single standard there would be no improvement in average fuel economy but under the current CAFE standard, which is differentiated by class, binds tighter on trucks than cars; the R&D policy relaxes this constraint implying that average fuel economy is, surprisingly, reduced by 0.02%, increasing gasoline consumption by 0.043%. By contrast, feebates channel nearly 60% of the possible R&D improvements to fuel economy, but these improvements are not enough to result in lower VMT or fuel consumption, which increase by 0.035% and 0.017% respectively.

These findings underscore several lessons. Because different agencies may target closely related aspects of a market failure using different policies, these policies may interact in ways that give rise to outcomes that were unintended by the regulators. Without coordination, these interactions may not only result in welfare losses but also increase total emissions. This is particularly concerning in the case of R&D spending to improve technology, which is often believed to be a promising channel to reducing emissions without substantial sacrifice. Moreover, the lack of enthusiasm to approve national climate change legislation suggests coordination is unlikely to improve in the near future. Thus it may be useful to implement a baseline policy that is more robust to uncoordinated, complementary policies.

Table 1.1: Potential Vehicle Types

Manufacturer	Class	Age
Ford	Small Car	New
General Motors	Medium Car	2-5 years
Chrysler	Large Car	6-23 years
Toyota	Small SUV	24 years or older
Honda	Medium SUV	
Other Asian	Large SUV	
European	Van	
	Small Truck	
	Large Truck	

Table 1.2: Descriptive Statistics

Household Data			Vehicle Data		
Variable	Mean	(SD)	Variable	Mean	(SD)
Family Size	2.32	*1.15*	Price	$ 28,445.00	*14426*
Number Working	1.09	*0.95*	Weight (lbs)	3652	*835.8*
Number Children	0.38	*0.79*	Horsepower	198	*59.70*
Cars per Household	2.03	*0.83*	Fuel Economy (gpm)	0.047	*0.010*
VMT (per car)	11,179	*8050*	Depreciation	$ 5,034.00	*4091*
Household Composition	Percent		Variable	Model Count	
Urban	61%		Domestic	638	
Suburban	10%		Asian	597	
Rural	30%		European	288	
Children in Household	20%		Car	842	
Retired Person in Household	41%		Van	128	
Northeast	15%		SUV	460	
Midwest	11%		Truck	93	
South	53%				
West	21%		Total	1523	
California	14%				

Notes: Household data from NHTSA 2009. Vehicle data from Wards Automotive Yearbook.

Table 1.3: First Stage Parameter Values: Vehicle and Household Interaction Terms

Coefficient		Estimate	Standard Error
alpha		-2.032	*0.187*
beta		1.770	*0.314*
Gamma Parameters			
Vehicle	Household		
constant	constant	1.657	*0.243*
	# working	0.047	*0.026*
	suburban	0.057	*0.089*
	rural	0.181	*0.092*
hp		-4.508	*2.400*
hp	adult age	7.261	*4.719*
wt		4.310	*1.553*
wt	adult age	-3.926	*2.623*
van		0.025	*0.205*
van	child	0.163	*0.182*
van	retired	0.139	*0.173*
suv		-0.041	*0.135*
suv	child	0.109	*0.142*
suv	retired	-0.117	*0.119*
truck		0.264	*0.148*
truck	child	-0.131	*0.167*
truck	retired	-0.365	*0.171*

Table 1.4: First Stage Parameter Values: Vehicle and Household Interaction Terms

Tau Paramters			
Vehicle	Household		
hp	subruban	0.541	*0.266*
hp	rural	0.150	*0.209*
hp	4year degree	0.287	*0.111*
hp	child	-0.017	*0.166*
wt	subruban	-0.271	*0.149*
wt	rural	-0.098	*0.124*
wt	4year degree	-0.146	*0.059*
wt	child	-0.053	*0.095*
van	adult age	-0.154	*0.097*
van	child	0.342	*0.189*
suv	adult age	-0.118	*0.059*
suv	child	0.025	*0.165*
truck	adult age	-0.059	*0.082*
truck	child	0.133	*0.215*
US	midwest	0.314	*0.128*
US	west	-0.059	*0.139*
US	south	0.171	*0.112*
Asian	midwest	-0.589	*0.194*
Asian	west	0.262	*0.150*
Asian	south	0.105	*0.131*
Euro	midwest	-0.800	*0.499*
Euro	west	0.243	*0.387*
Euro	south	-0.103	*0.363*

Table 1.5: Second Stage Parameter Values

Coefficient	Estimate		Standard Error
Constant	-10.330	***	*0.284*
Price	-0.886	***	*0.126*
Horsepower	0.762	***	*0.161*
Weight	0.593	***	*0.107*
European Firm	0.911	***	*0.179*
Asian Firm	-0.300	***	*0.081*
SUV	0.255	**	*0.109*
Truck	-0.100		*0.201*
Van	-0.443	**	*0.171*

Notes: Second stage parameters are estimated using an IV framework in the spirit of Berry et al. (1995) that instruments for price with vehicle length, footprint and number of products in class both within and outside of the firm. Domestic Firms and Cars are the ommitted categories of the estimation.

Table 1.6: Elasticities

Panel A: Price Elasticity of Demand	
Vehicle Average	-3.02
Sales Weighted	-7.46
Panel B: Elasticity of Gasoline Use with Respect to Price	
Full Sample	-0.45
Urban	-0.54
Suburban	-0.36
Rural	-0.46
Working	-0.44
Retired	-0.46
Poor	-0.44
Rich	-0.48

Table 1.7: Tax Credit: Quantity and Fuel Economy Outcomes

Panel A: Fuel Economy Standard—CAFE

	Quantity	GPM	MPG	MPG Cars		MPG Trucks	
			Total	PrePolicy	PostPolicy	PrePolicy	PostPolicy
Ford	2.46%	-0.47%	0.47%	27.5	27.5	20.7	20.7
GM	2.85%	-0.48%	0.48%	27.5	27.5	20.7	20.7
Chrysler	2.92%	-0.58%	0.58%	27.5	27.5	20.7	20.7
Toyota	-0.13%	-0.03%	0.03%	33.2	33.2	22.7	22.7
Honda	-0.14%	-0.04%	0.04%	33.1	33.1	24.6	24.6
Other Asian	-0.13%	-0.04%	0.04%	28.7	28.7	22.8	22.8
European	-0.09%	-0.03%	0.03%	27.7	27.7	20.7	20.7
Total New	1.50%	-0.20%	0.20%				

Panel B: Feebate (£500 per .01 gpm)

	Quantity	GPM	MPG	MPG Cars		MPG Trucks	
			Total	PrePolicy	PostPolicy	PrePolicy	PostPolicy
Ford	2.16%	-0.87%	0.86%	26.7	27.1	20.1	20.1
GM	2.67%	-0.95%	0.94%	26.7	27.1	20.1	20.1
Chrysler	2.64%	-0.92%	0.91%	26.6	27.1	20.2	20.2
Toyota	-0.12%	-0.02%	0.02%	33.5	33.5	22.8	22.8
Honda	-0.12%	-0.03%	0.03%	33.2	33.2	24.7	24.7
Other Asian	-0.10%	-0.03%	0.03%	28.7	28.7	22.9	22.9
European	-0.07%	-0.02%	0.02%	27.1	27.1	20.3	20.3
Total New	1.39%	-0.45%	0.45%				

Table 1.8: Tax Credit: Welfare

Panel C: Fuel Economy Standard—CAFE	
Percent Change in VMT	0.13%
Percent Change in Gallons of Gasoline	0.11%
Welfare (Millions of Dollars)	
Vehicle Choice and VMT	
Domestic Vehicles	346.14
Foreign Vehicles	-30.78
Total from Vehicles	315.09
Associated Externalities	
Carbon	-7.53
Congestion and Accidents	-255.21
Local Pollutants	-74.60
Total Externalities	-337.34
Total Welfare Change	-22.25
Panel D: Feebate (£500 per .01 gpm)	
Percent Change in VMT	0.12%
Percent Change in Gallons of Gasoline	0.08%
Welfare (Millions of Dollars)	
Vehicle Choice and VMT	
Domestic Vehicles	152.70
Foreign Vehicles	-3.60
Total from Vehicles	149.07
Associated Externalities	
Carbon	-5.49
Congestion and Accidents	-226.41
Local Pollutants	-66.18
Total Externalities	-298.08
Total Welfare Change	-149.01

Table 1.9: Research and Development Policy: Quantity and Fuel Economy Outcomes

Panel A: Fuel Economy Standard—CAFE

	Quantity	GPM	MPG	MPG Cars		MPG Trucks	
			Total	PrePolicy	PostPolicy	PrePolicy	PostPolicy
Ford	0.99%	0.01%	-0.01%	27.5	27.5	20.7	20.7
GM	0.82%	0.01%	-0.01%	27.5	27.5	20.7	20.7
Chrysler	0.80%	0.00%	0.00%	27.5	27.5	20.7	20.7
Toyota	0.10%	-0.06%	0.06%	33.2	33.2	22.7	22.7
Honda	0.11%	-0.06%	0.06%	33.1	33.1	24.6	24.6
Other Asian	0.10%	-0.05%	0.05%	28.7	28.7	22.8	22.8
European	0.07%	-0.07%	0.07%	27.7	27.7	20.7	20.7
Total New	0.53%	0.01%	-0.01%				

Panel B: Feebate (£500 per .01 gpm)

	Quantity	GPM	MPG	MPG Cars		MPG Trucks	
			Total	PrePolicy	PostPolicy	PrePolicy	PostPolicy
Ford	0.07%	-0.59%	0.59%	26.7	26.8	20.1	20.2
GM	0.05%	-0.60%	0.59%	26.7	26.8	20.1	20.2
Chrysler	0.05%	-0.59%	0.59%	26.6	26.8	20.2	20.3
Toyota	0.00%	0.00%	0.00%	33.5	33.5	22.8	22.8
Honda	0.00%	0.00%	0.00%	33.2	33.2	24.7	24.7
Other Asian	0.00%	0.00%	0.00%	28.7	28.7	22.9	22.9
European	0.00%	0.00%	0.00%	27.1	27.1	20.3	20.3
Total New	0.03%	-0.37%	0.37%				

Table 1.10: Research and Development Policy: Welfare

Panel C: Fuel Economy Standard—CAFE	
Percent Change in VMT	0.041%
Percent Change in Gallons of Gasoline	0.043%
Welfare (Millions of Dollars)	
Vehicle Choice and VMT	
Domestic Vehicles	61.16
Foreign Vehicles	-237.98
Total from Vehicles	-173.70
Associated Externalities	
Carbon	-2.83
Congestion and Accidents	-77.58
Local Pollutants	-22.68
Total Externalities	-103.09
Total Welfare Change	-276.79
Panel D: Feebate (£500 per .01 gpm)	
Percent Change in VMT	0.035%
Percent Change in Gallons of Gasoline	0.017%
Welfare (Millions of Dollars)	
Vehicle Choice and VMT	
Domestic Vehicles	8.23
Foreign Vehicles	-0.09
Total from Vehicles	8.13
Associated Externalities	
Carbon	-0.84
Congestion and Accidents	-67.18
Local Pollutants	-19.64
Total Externalities	-87.65
Total Welfare Change	-79.52

Figure 1.1: Comparison of Firm Profits

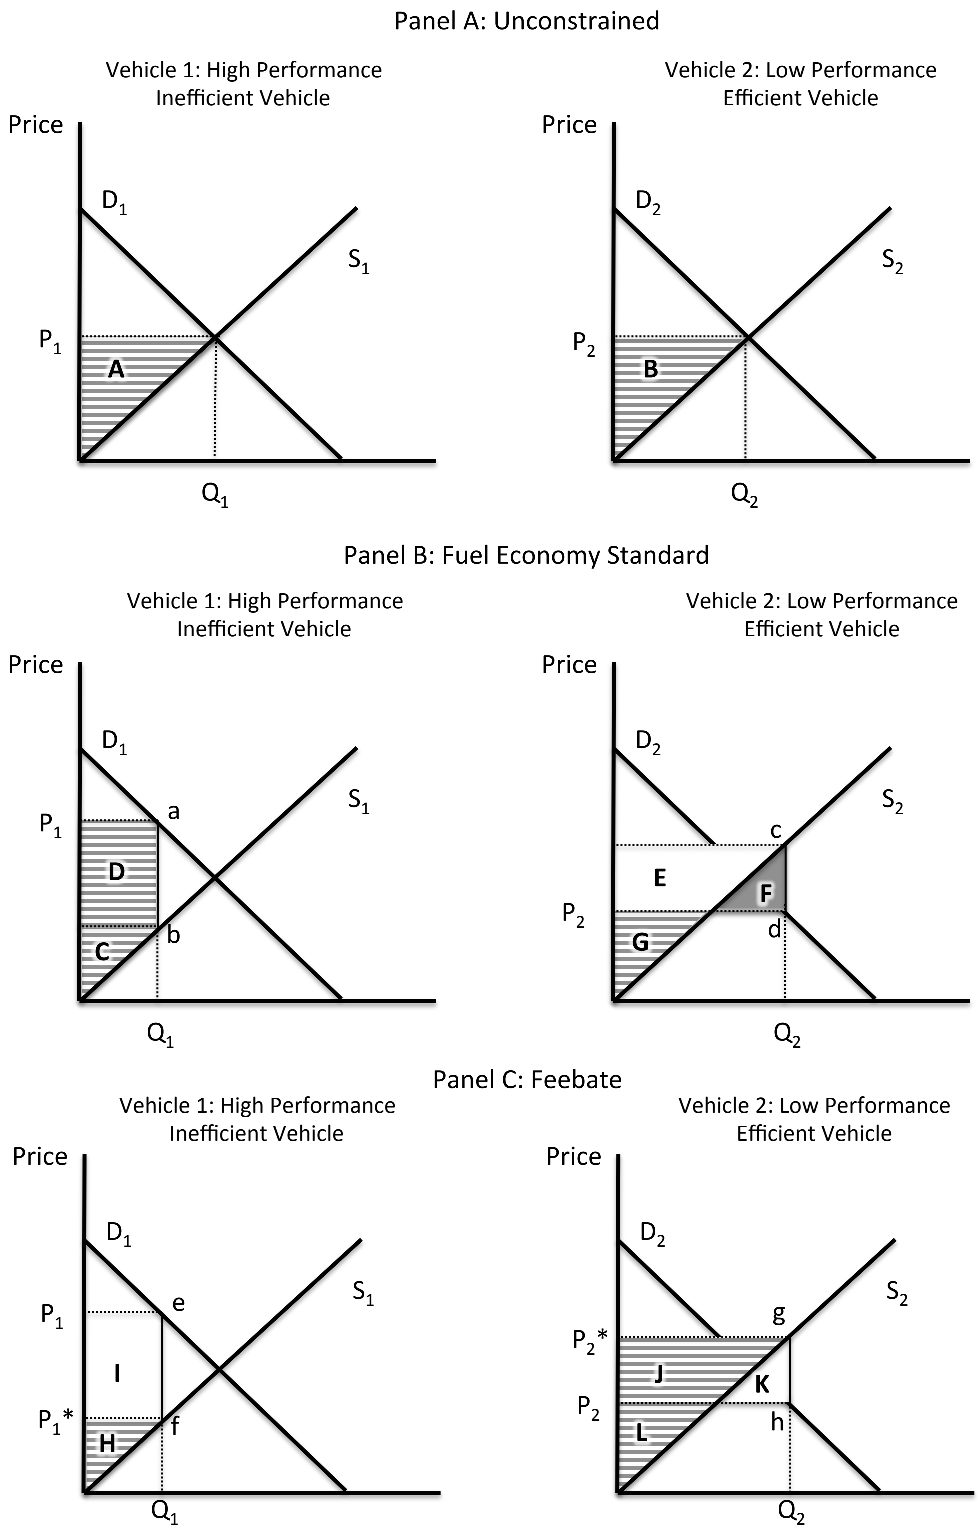

Figure 1.2: Efficient Vehicle Tax Credit

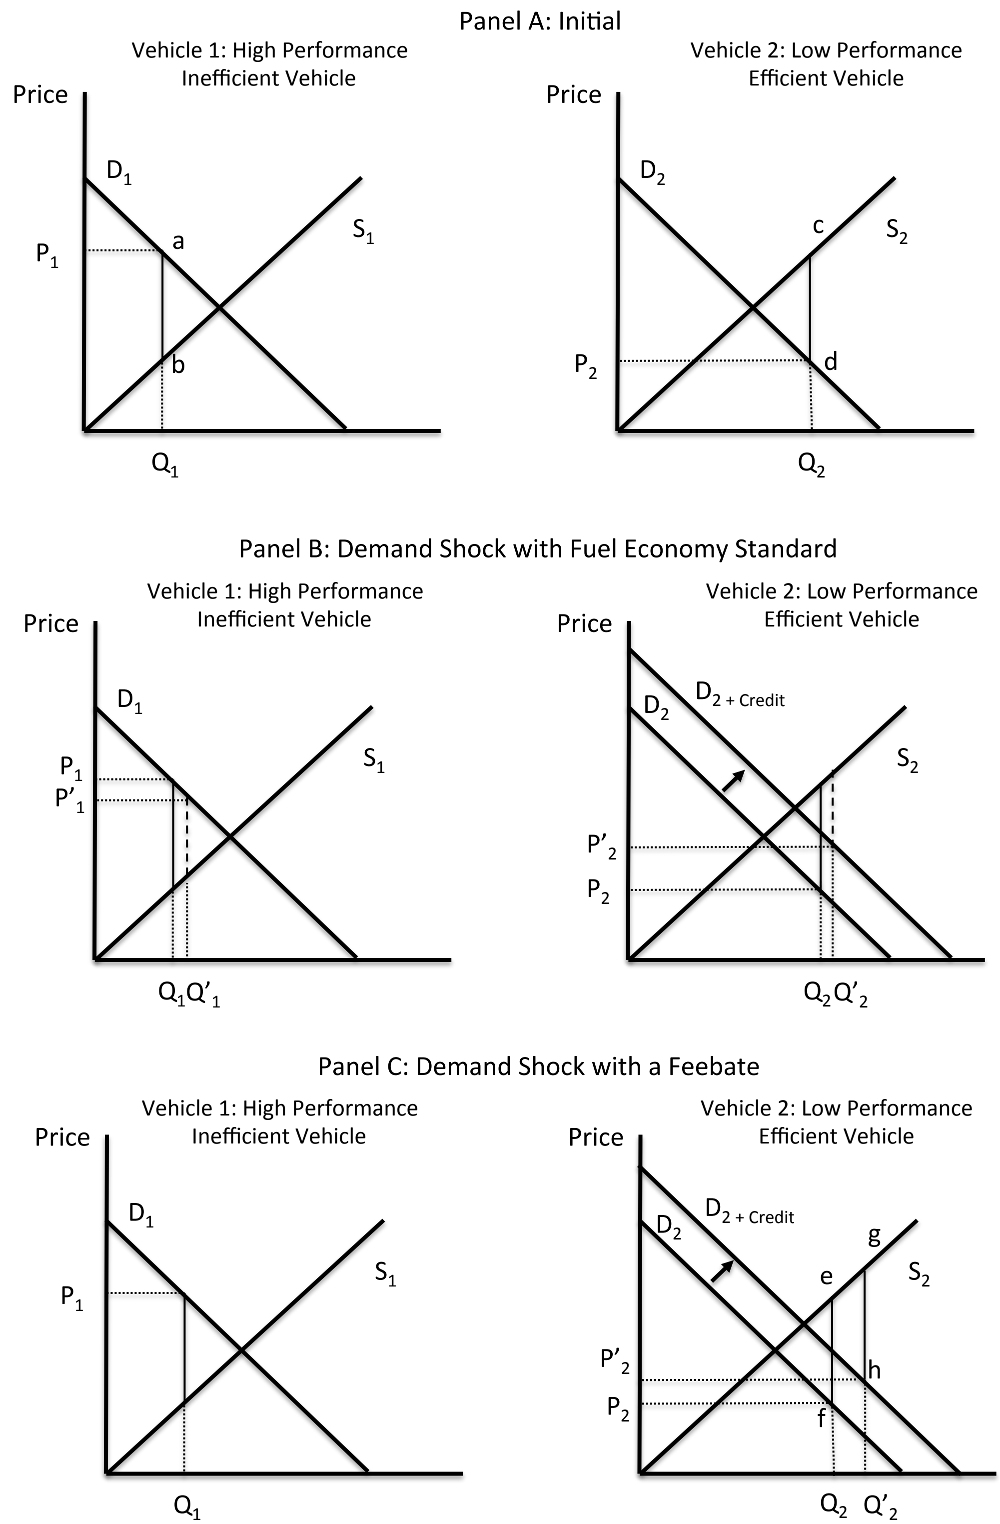

CHAPTER 2

THE EFFECTS OF REGULATION IN THE PRESENCE OF MULTIPLE UNPRICED EXTERNALITIES: EVIDENCE FROM THE TRANSPORTATION SECTOR

2.1 Introduction

For reasons discussed in (Harberger, 1974), the estimation of the overall welfare effects of government interventions to correct externalities is more challenging than first outlined by (Pigou, 1932). In the presence of unpriced externalities or other pre-existing distortions, policies levied to correct an externality can exacerbate or alleviate these other distortions in related markets. A priori, theory cannot shed light on the relative importance of the primary welfare effect of the policy defined by the welfare gain from correcting the externality addressed by the policy and the interaction effects defined as the welfare effect that results from the interaction of the new policy with other unpriced externalities. Multiple unpriced externalities are particularly prevalent in the transportation sector, where unpriced congestion and air pollution interact in nontrivial ways across space and time.[1] Recently, in an attempt to reduce automobile-related emissions, policymakers have introduced policies to stimulate the demand for ultra-low-emission vehicles (ULEVs) such as gas-electric hybrids.[2] A popular

[1] Automobiles are major contributors to local and global air pollution, including carbon monoxide (CO), nitrogen oxides (NOx), and hydrocarbons (HC), and the transportation sector accounts for 20% of greenhouse gas (GHG) emissions in the United States (Ep, 2007). Automobile use also leads to significant congestion costs. In 2003 these were estimated at $63 billion dollars, with drivers in our area of focus losing 93 hours annually to congestion delays (Schrank & Lomax, n.d.).

[2] Policies to promote the purchase of these new hybrid vehicles are also motivated by the perceived lag in adoption that often results from a lack of acceptability of new technologies by consumers (D. Greene, of Transportation, Assessment, & Division, n.d.; Helfand & Wolverton, 2011).

policy, in place in nine states and under consideration in six others, consists of allowing solo-hybrid drivers access to high occupancy vehicle (HOV) lanes on major freeways.[3] In this paper, we take advantage of the introduction of this policy in Los Angeles, California to study interactions between multiple unpriced externalities. We demonstrate the first-order importance of the interaction effect between the policy and unpriced congestion and show that it generates substantial welfare losses, dominating the expected primary welfare gain of the policy. While it may not be too surprising that allowing solo-hybrid drivers into HOV lanes is likely an inefficient policy to promote the adoption of green technologies, we stress the remarkable variation of the interaction effect between unpriced congestion and the policy across space and time. Congested locations and peak travel periods make the HOV lane exemption attractive to hybrid drivers, but create the greatest congestion costs for existing carpoolers. While adding a single hybrid to any HOV lane at 2AM creates zero social costs of congestion, adding one daily hybrid driver at 7AM to a very congested road in our study area (the I-10W) generates $4500 in annual social costs. On these exceptionally congested roads, HOV lane traffic may be up to 30% above socially optimal levels, implying significant congestion costs from allowing hybrid access. As such, estimates of the effects of the policy that rely on average values would underestimate the impact of a marginal hybrid on HOV travel times, leading to erroneous estimates of the welfare effects of the policy. Our findings imply a best-case cost of $124 per ton of reductions in greenhouse gas emissions, $606,000 dollars per ton of nitrogen oxides (NOx) reduction, and $505,000 dollars per ton of hydrocarbon reduction in the most optimistic calculations. These costs exceed those of other options readily available to policymakers. Further,

[3]Of nearly twenty states with HOV lanes, nine allow hybrid vehicles to drive in HOV lanes: Arizona, California, Colorado, Florida, New Jersey, New York, Tennessee, Utah and Virginia.

a policy that was perceived as free was far from free. We find that it costs carpoolers \$3-\$9 for every \$1 of benefit transferred to hybrid drivers. If instead a tax-credit were implemented, funded through a broad-based tax, the cost to taxpayers would have been only \$1.4 per dollar transferred. To measure the magnitude of these interaction effects, we have assembled a rich dataset that includes real time data from the Freeway Performance Measurement System (PeMS) in California. PeMS reports hourly travel time for major routes and traffic flow for detectors located in HOV and mainlines. We use both travel time on a single route (I-10W) and detector level traffic flow for District 7, which corresponds to the greater Los Angeles metropolitan area and more than 36% of all HOV lanes in California. The analysis controls for possible confounding factors through the use of a regression discontinuity (RD) design where travel time and traffic flow in the HOV and mainline lanes are compared before and after the start of the policy. A common tension within the RD literature is the use of short run RD estimates to establish long run welfare effects. As a robustness check, we also take advantage of the recent ending of the policy to test the persistence of these changes in travel time. Similar changes at the end of the policy suggest that there is a sustained, upward shift in congestion levels, validating the use of our estimates for the welfare analysis. This paper contributes broadly to the literature on environmental policy in a second-best setting. When examining the welfare effects of environmental policies in a second best setting, (Bovenberg & de Mooij, 1994), (Bovenberg & Goulder, n.d.), and (Parry & Small, n.d.), emphasize the importance of considering the interactions between preexisting distortions caused by distortionary taxes in factor markets and the new environmental policy.[4] As in (Parry & Small, n.d.), here the interactions occur between the

[4]In the context of automobile-related policies, (Parry & Small, n.d.) derive rules for the optimal second-best gasoline tax and examine the interactions between a gasoline tax and preexisting distortions caused by labor taxes and unpriced externalities within the transportation

policy and unpriced congestion externalities. In contrast with nearly all prior work in this area, which has typically relied on analytical and calibrated simulation models, our analysis provides one of the first econometric estimates of the interaction effect between an environmental motivated policy and a competing unpriced externality.[5] These econometric estimates are incorporated into a general equilibrium welfare framework, which is used to calculate the overall welfare and distributional impacts of the policy. A key insight of the second-best literature is that revenue-raising instruments can alleviate some of the welfare losses associated with interaction effects (Goulder, Parry, & Burtraw, 1997). In this context, we compare the welfare effects of this non-revenue raising policy with alternative revenue raising policies, including auctioning of the hybrid stickers and high-occupancy toll lanes (HOT) that vary with fuel economy. We also compare this policy against tax incentives, as recently examined in (Sallee, n.d.).

2.2 Background on the Policy

California has long had a reputation for being at the forefront of environmental policy. Because of perceived costs, many environmental policies typically face resistance from taxpayers and industry. In contrast, the Clean Air Vehicle Sticker (CAVS) policy was popular among nearly all interested parties. In the words of Assemblywoman Fran Pavley (D-Agoura Hills), author of the measure, "This is a win, win, win – cleaning up our air, reducing dependence on foreign oil and

system, such as congestion, air pollution and accidents with a simulation model.

[5]The exception is (West & Williams III, 2007) who estimate the parameters necessary to calculate the optimal second-best gasoline tax using household level data.

saving money at the pump."[6] Beginning August 10, 2005 and ending June 30th, 2011, owners of hybrid vehicles achieving 45 miles-per-gallon (mpg) or better were able to apply for a special sticker that allowed them access to HOV lanes regardless of the number of occupants in the vehicle. The goal of the CAVS policy was to stimulate the demand for highly fuel-efficient vehicles, particularly of ultra-low-emission vehicles (ULEV), such as the Honda Insight, Honda Civic Hybrid, and Toyota Prius. The sticker had limited transferability and was allocated to the vehicle and not the driver. The original bill allowed for the issuance of 75,000 stickers, but later legislation eventually increased the limit to 85,000. Stickers were available for $8 dollars to all owners of eligible vehicles, including the owners of the over 65,000 estimated hybrids already registered in California at the start of the policy. The rationale for this decision rested on the idea of not penalizing earlier adopters. In Los Angeles County, a total of 27,228 stickers were distributed over the course of the program. By August 20th, more than 12,000 applications for stickers were submitted to the DMV, implying that a substantial number of hybrids entered the HOV lanes at the beginning of the program.[7] By February 2007, all 85,000 stickers had been issued. While the original CAVS hybrid policy expired June 30th, 2011, it is nonetheless crucial to understand the effects of these programs as California introduced a new HOV exception program with 40,000 stickers for electric, hydrogen fuel cell, and plug-in hybrid vehicles on January 1st 2012.[8]

[6]Salladay, Robert. 2004. "Hybrids Move Closer to Using Carpool Lanes," Los Angeles Times, May 7, 2004.

[7]Gledhill, Lynda. 2005. "Drivers race for carpool permits for hybrids: At 1,000 applicants a day, some predict gas-saver gridlock," SF Chronicle, August 20.

[8]The CAVS program was initially designed to expire on December 31st, 2008, however, as a result of the popularity of the program, an organized group of hybrid drivers has successfully lobbied for subsequent extensions of the program. Barringer, Felicity. 2011. "Hybrid Owners Seek to Extend Carpool Privilege" New York Times, May 18.

2.3 Data

Data on major highways in California is collected by the Freeway Performance Measurement System (PeMS), a joint effort by the California Department of Transportation (Caltrans), the University of California, Berkeley, and the Partnership for Advanced Technology on the Highways (PATH). PeMS obtains real-time 30-second loop detector data across 12 Caltrans districts. Each detector compiles data on traffic flow and lane occupancy, which are then used to calculate traffic speed.[9] We use both travel time on a single route and detector level traffic flow for District 7, corresponding to the greater Los Angeles metropolitan area.

2.3.1 Travel Time on the I-10W

Because commuters are primarily concerned with the time it takes to commute along a particular route, our initial analysis focuses on hourly travel time over a single freeway route. Routes are defined as a segment of the freeway system from a fixed starting point to a fixed destination and are predetermined by PeMS. A route level measure of travel time combines information from multiple detectors that a commuter would typically drive.[10] The data set obtained from PeMS reports the hourly travel times along a 17.5-mile section of the I-10W for

[9] Lane occupancy is the fraction of time the detector is on due to the presence of a vehicle. Based on average vehicle length and this lane occupancy measure, the speed of traffic is computed. See PeMS FAQ for more information: http://pems.eecs.berkeley.edu/?dnode=Help&content=help_faq.

[10] For example, drivers commuting to downtown LA from West Covina typically use the I-10 route, while those commuting from Thousand Oaks to downtown LA use US Hwy 101. PeMS computes travel time over a freeway segment by dividing the length of the segment by the calculated traffic speed at that detector and summing travel time across the segments that form the route.

the HOV lane and each of the four mainline lanes (Map 1). With the exception of a three-plus occupant requirement during peak travel times in the HOV lane, this route is fairly representative in terms of size and design for the Los Angeles metropolitan area.[11] In addition to the I-10W, travel time data is also collected for the I-210W to broadly capture demand on competing freeways. While a large window of data around the policy is desirable, our window of analysis is limited by the availability of I-210W data to be January 2004 through December 2007. Finally, as weekend and holiday travel demand is substantially different from weekday demand, these observations are removed from the initial analysis resulting in a total of 34,980 hourly travel time observations by lane. Figure 1 plots the 2004-2007 average travel times for the HOV and mainline lanes of the I-10W across the hours of the day. Morning peak is defined by Caltrans as 5 A.M. to 9 A.M. and afternoon peak as 4 P.M. to 7 P.M. The mid-day off-peak corresponds to 10 A.M. to 4 P.M. and the night off-peak from 8.P.M. to 4 A.M.[12] The figure reveals substantial variation in travel times over the course of the day, with maximum travel time levels of over 35 minutes in the mainline reached during the morning peak. This figure also underscores the large differential between mainline and HOV lane travel times during the morning peak, with a maximum difference of nearly 10 minutes at 7 A.M. While the route level data allows us to explore the heterogeneity of the CAVS policy across various times of the day, it does not allow for exploration of the effect across various locations of the LA metropolitan area. To explore the spatial effects of the program and generality of the results, we also consider a detector level dataset of traffic flow

[11]The I-10W, westbound from West Covina to Los Angeles, was selected on the criteria of data availability, data on a competing route, and high detector density. See Appendix A for more detail on the selection of the I-10W and further discussion of the data used.

[12]See Tables E.2 and E.3 in Appendix E for weekday travel time averages for each lane and route, including the I-210W, during the four peak and off-peak periods. The westbound direction implies peak demand in the morning period.

spanning the LA area.

2.3.2 Detector Level Data

We estimate the citywide effect of the CAVS policy during peak hours using a comprehensive dataset of 677 detectors collected by PeMS. These detectors record hourly observations of traffic flow for HOV and mainline lanes on 18 freeways in Los Angeles. The wide spatial distribution of the final set of detectors used in the analysis can be noted in Map 1. Three months of data, July-September 2005, provide hourly traffic flow observations for 1,750 detectors across the 18 freeways, which includes all freeways with HOV lanes.[13] In addition, we also collect three months of data for 331 detectors around the end of the policy, June 30th, 2011.[14] Detectors located at on- and off-ramps are deleted from the analysis. Each detector is also required to have at least 50 observations after all deletions.[15] As is the case for the I-10W, many routes have a dominant commuting pattern such that only one peak time of day experiences congestion. We estimate the effect of the CAVS policy for each detector during the particular peak period corresponding to maximum traffic flow for that detector, during the three-month period.[16] This ultimately yields 677 individual detector level esti-

[13] A larger window than three months would strain the assumptions of the local linear regressions used below.

[14] To ensure comparability with the start of the policy, we limit the detectors used in this analysis to those active at the start of the policy. Removal of decommissioned detectors and those with insufficient data further reduces the sample size.

[15] These deletions include hours outside of the specified peak period, weekend observations and observations labeled as less than 100 percent observed. Where detectors are not properly functioning, PeMS imputes missing values. By dropping all observations where Percent Observed is less than 100, all data with PeMS imputation are removed from the analysis.

[16] This method implies that for nearly all of the I-10W detectors, the morning peak is selected as the most congested time of day, as expected. Detectors for which the maximum traffic flow occurs outside the peak periods are also excluded, as a detector with maximum flow occurring at 2 A.M. is of questionable quality.

mations of the effect of the CAVS policy on traffic flow, 200 of which are located in HOV lanes for our main analysis of the start of the policy.[17]

2.3.3 Other Covariates

The PeMS data is supplemented with hourly measures of weather from the National Weather Service at nine airports in the Los Angeles area. These measures include rainfall in inches, visibility in miles, cloud cover as a percentage of the sky, temperature in degrees Fahrenheit, and wind speed in miles per hour. The Fullerton airport station is closest to the I-10W, and data from this station is matched to the travel time data.[18] Nominal weekly retail gasoline prices (regular reformulated) for Los Angeles from 2004 to 2011 were obtained from the Energy Information Administration.

2.4 Empirical Strategy

We begin by describing the empirical strategy utilized to estimate the effect of the CAVS policy on travel time on the I-10W. We employ a regression discontinuity (RD) design where logged hourly travel time in lane i at hour h on date t, TT^i_{ht}, is regressed separately by lane on $1(Hybrid_t)$, an indicator variable for observations after the implementation of the CAVS policy, a vector of covariates

[17] For the 331 detectors that comprise the end of policy analysis, 82 are HOV lane detectors.

[18] For the Fullerton station, of the 35,064 total observations, 840 had at least one weather measure missing. These missing weather measures are imputed from the other stations in the Los Angeles area (Chino, Hawthorne, Hollywood, Long Beach, Los Angeles, Ontario, Santa Monica and Van Nuys), following the algorithm used in (Auffhammer & Kellogg, n.d.). See Appendix A.

X_{ht}, and a flexible nth-order polynomial in date $f(Date_t)$:

$$TT^i_{ht} = \alpha^i + \beta^i 1(Hybrid_t) + \gamma^i X_{ht} + f(Date_t) + \epsilon^i_{ht}. \tag{2.1}$$

The coefficient of interest, β^i, is the treatment effect of the CAVS policy on travel time in lane i.[19] Policy date is taken to be August 20th, 2005, which is when stickers first became available. The vector of covariates X_{ht} includes indicator variables for month of the year interacted with day of the week, and indicators for hour of the day. Additional controls include weather variables (linear and quadratic rainfall, linear and quadratic visibility, and indicators for cloud cover in the central specification), logged gas prices, and travel time on competing routes (I-210W). [20] Finally, as people often choose freeway routes based on travel updates in the hour before they leave home, we include travel time on the I-210W lagged by one hour.[21] While the introduction of the stickers identifies the short run-effect of the policy, our interest is on the overall welfare effects and distributional impacts of the policy, typically a long run calculation. Therefore, the welfare estimates presented below are based on the assumption that the unobserved, latent travel-time function is a vertical translation of the observed travel-time function. However, one may be concerned with the non-instantaneous adjustment for the rest of the stickers, which were distributed over a longer period of time. To address this concern, we also estimate the effect

[19]As the dependent variable is logged, a one-unit increase in $1(Hybrid_t)$ would imply a percentage increase in travel times of $e^(\beta^i - 1)$. Because this transformation does not significantly change any of our results, we ignore it and simply discuss parameter estimates in the results for the sake of exposition.

[20]Several robustness checks also include measures of temperature, (sustained) wind speed and wind gusts. Temperature is included as three indicator variables, below 80 degrees, 80 to 100 degrees, and above 100 degrees. Wind is included as an indicator variable for sustained wind speeds above 20 miles per hour and similarly an indicator for wind gusts indicates gusts above 20 miles per hour.

[21]While the inclusion of the I-210W is justified by economic theory, as a robustness check, estimations without the inclusion of the I-210W are also performed, yielding results similar to the key findings presented below. See Appendix E.

of the policys recent termination date of July 1st, 2011.[22] Estimating the policy effects at both the beginning and end of the policy provides insight into the transition and evolution of the policy effect.

2.4.1 RD: Global Polynomial

The potential for omitted time-varying factors to confound our estimation make observations substantially before or after the introduction of the policy less informative about the effect of the policy on travel time. Without controlling for these time-varying factors, the error term may be correlated with time, and thus with $1(Hybrid_t)$, producing biased estimates of β^i. Under reasonable assumptions, regression discontinuity methods yield consistent estimates of β^i in the presence of time-varying omitted variables. (J. Hahn, Todd, & Klaauw, n.d.) show that nonparametric identification of a constant treatment effect with a sharp RD design requires that the conditional mean function is continuous at the threshold. In other words, provided that all other factors affecting travel time besides the CAVS policy are continuous at the policy date, the RD design will yield a consistent estimate of the effect of the policy.[23] Equation (2.1) includes a single, flexible nth order polynomial in date, $f(Date_t)$, which controls for unobserved, time-varying factors that evolve smoothly and may influence travel

[22] As discussed below, due to data window limitations arising from the recent termination of the policy, we are restricted in terms of the types of analysis we can perform at the end of the policy. As such, most of our analysis focuses on the effects estimated at the introduction of the policy. We interpret the estimates as a conservative estimate of the congestion effects of the policy, and where possible, also incorporate the end of policy estimates to support our analysis.

[23] If other discontinuous, unobserved changes occurred at the policy date, the effect of those unobservables would be indistinguishable from the effect of the policy. For example lane closures due to construction would cause discontinuities in traffic flow. To reduce the likelihood that the effects observed below are caused by other such scenarios we perform several robustness checks including an examination of weekend travel time that would likely be affected by construction but not the CAVS policy.

times but are unrelated to the policy.[24] Following the approach of (DiNardo & Lee, 2004) and (L. W. Davis, n.d.), an eighth-order polynomial was selected as the most parsimonious specification that adequately describes the underlying time trend with a reasonable degree of smoothness, with estimates for sixth through tenth-order polynomials also reported below.[25] As standard tests suggest serial correlation is of some concern, robust standard errors clustered at the week level are calculated for all regressions.[26] Finally, we note that we do not have a sufficient time window to estimate a global polynomial model for the end of the policy estimates.

2.4.2 RD: Local Linear

In a local linear regression discontinuity design, time varying factors are controlled for with a linear trend within some local bandwidth of the policy discontinuity (G. W. Imbens & Lemieux, 2008). In Appendix E, robustness checks on the route level analysis using a local linear method - where $f(Date_t)$ is linear and interacted with the policy variable - are performed, yielding similar results to the global polynomial estimates reported below. The local linear specification

[24] An alternative approach to the regression discontinuity design would be to use a difference-in-differences approach. However, it is difficult to construct an appropriate control group. As all freeways in California were subject to the policy, no freeway in California could be considered untreated. Comparing against a freeway outside of California would strain the necessary assumptions of the difference-in-difference approach.

[25] The most common method of polynomial selection in the literature chooses the order that smoothly describes the underlying trend in the data, while presenting estimates for alternative polynomial orders (D. S. Lee & Lemieux, n.d.). As this involves a substantial element of modeler discretion, misspecification is a concern. As (DiNardo & Lee, 2004) note, misspecification of the order can lead to biased estimates of the discontinuity and erroneous interpretations of statistical significance. We report results for other polynomial orders as well as the order chosen by the Bayesian Information Criterion (BIC) (Matsudaira, n.d.). Graphs of other polynomial orders are included in Figures D.1-11 in Appendix D.

[26] Clustered standard errors are reported for our results as they yielded more conservative estimates of the standard errors than Newey-West standard errors; Newey-West standard errors are presented in robustness checks in Appendix E.

is also used for the detector level analysis of the effect of the CAVS policy on traffic flows both at the introduction of the policy in 2005 and its expiration in 2011.[27] The estimating equation is similar to equation (2.1), with logged-hourly traffic flow as the dependent variable and weather covariates as well as hourly fixed effects included. Traffic flow effects are estimated for each detector using an Epanechnikov kernel with a 30-day bandwidth.[28] To determine the citywide effect of the policy on traffic flows in HOV and mainline lanes, the RD estimates across detectors are averaged. Standard errors are calculated using 5,000 bootstrap samples.[29]

2.5 Route Level Results

2.5.1 The Effect of the CAVS Policy on Travel Time

Figure 2 illustrates the regression discontinuity strategy for estimating the effect of the CAVS policy on travel time. Panel (a) depicts travel time residuals in the HOV lane during the morning peak and panel (b) for the mainline. Similarly, panel (c) depicts the HOV lane residuals during the afternoon peak and panel (d) displays the residuals for the mainline during the afternoon peak. Panels

[27]Local linear regression is utilized in the detector analysis because of the challenges associated with correctly specifying global polynomial controls for each of the 677 detectors.

[28]A 30-day bandwidth may not be optimal for all detectors; a robustness check Table E.19 in Appendix E presents results using the Silverman Rule for bandwidth selection and Table E. 21 presents results using the method outlined in (G. Imbens & Kalyanaraman, 2012).

[29]We also correct these bootstrapped errors for spatial correlation and the imprecision of the detector level estimates. To account for the imprecision of the estimates from the detector level regressions, we first generate 5,000 sets of detector-level effects using the estimated means and standard deviations on each detector. These effects are then spatially partitioned by route-direction into disjoint blocks. For each of the 5,000 bootstrap renditions of the data we sample blocks with replacement to account for the spatial correlation in the estimates. These samples are then used to generate the mean and standard deviation of the estimated effect of the policy.

(e) (h) present travel time residuals during off peak periods. The plotted points represent the averaged, biweekly residuals of log hourly travel time regressed against the covariate vector X_t. These residuals should reveal any underlying, time-varying trends as well as any discontinuous changes in travel time at the policy date.[30] The fitted lines are the predicted values of a regression of these residuals on the eighth-order polynomial time trend and the CAVS policy variable. Panels (a) and (c) reveal that travel time in the HOV lane increased during peak hours due to the CAVS policy, and that these effects are larger during the morning peak than in the afternoon.[31] Figure 2 also shows no policy effect on travel time in the mainline.[32] Figure 2 reveals that congested lanes and times of day display noisier plots than less congested conditions. This is consistent with (Schrank & Lomax, n.d.) who document that over half of congestion delays, especially in peak periods, are the result of nonrecurring events.

Regression Discontinuity Estimates

Table 1 presents the regression discontinuity estimates of the effect of the CAVS policy on travel time.[33] For each time window and lane, the table reports the

[30]These underlying trends can produce biased estimates when using OLS. For example OLS finds a statistically significant increase of 4.6 percent for travel time in the mainline during the morning peak, despite the absence of a discontinuity in panel (b) of Figure 2. See Table E.4 in Appendix E.

[31]When the time differential between the HOV and mainline lanes is the largest, stickered hybrids will have the strongest incentive to move into the HOV lane - from the I-10W mainline and also alternate routes without an HOV lane option. During off-peak hours, traffic flow is often below the threshold level where congestion occurs (Vickrey, 1969). In these free flowing periods, travel time will be invariant to the addition of hybrids and the policy should have small effect.

[32]Slight differences between Figure 2 and the estimates in Table 1 are due to the fact that the polynomial and discontinuity are fit to the residuals in Figure 2, while they are included jointly in the regression in Table 1.

[33]These regressions include gas price, lagged I-210W travel time, weather covariates, indicator variables for hour of day, and for month-day of week or day of week as noted. Robust standard errors clustered by week are included in parenthesis for all regressions. Robustness

point estimate and standard error for the coefficient of interest, the percentage effect of the CAVS policy on travel time, broken down by morning peak, afternoon peak and off-peak periods. The preferred specification of the eighth-order polynomial is presented in column III. Consistent with Figure 2, the results in column III confirm that the increase in travel time during the morning peak on the HOV lane is 9.0 percent and is statistically significant at the 1 percent level; this effect corresponds to an increase of travel time of 2.2 minutes. Columns I-VI of Table 1 also present results for sixth-order through tenth-order polynomials, as well as the results for the polynomial order chosen under the BIC.[34] Across all polynomial orders, the estimate of the effect of the CAVS policy on travel time in the HOV lane is relatively stable, ranging between 8.4 and 11.6 percent, and is statistically significant at the 1 percent level.[35] The point estimates in the afternoon peak, 5.7 to 6.7 percent, are smaller than those estimated during the morning peak, consistent with the fact that congestion is most severe in the morning peak. Table 1 also reports the estimated effect of the CAVS policy on mainline travel times. Broadly speaking, these point estimates are statistically insignificant and unstable (-6.4 to 7.0%) across polynomial orders and

checks of the global polynomial regression discontinuity estimates are presented in Tables E.8-E.12 in Appendix E, with generally consistent with those in Table 1. These robustness checks test the sensitivity of our main specification by altering covariate specifications (fixed-effects, controls for competing routes, additional weather covariates, inclusion of holidays and gas prices), removal of days near the beginning of the policy, estimating separate polynomials on either side of the discontinuity, Newey-West standard errors, various data windows and altering weather aggregation methods. Local linear regression estimates are generally consistent with the global polynomial estimates and are reported in Tables E.13-E.17. All estimated coefficients are included in Tables E.5-E.7.

[34]The BIC selects the 8th order polynomial as the preferred specification in the mainline during the morning peak. For the morning peak HOV lane, a 10th order polynomial specification is chosen.

[35]As a check on the size of the policy effect, we note that Caltrans occasionally performs vehicle counts detailing the number, type, and occupancy of cars passing a point on major freeways in California, including the I-10W. A two-hour car count conducted between 6:30 A.M. and 8:30 A.M. on a weekday in 2006 found 167 single occupancy hybrids traveling in the HOV lane. This represents 5.8 percent of total vehicles counted in the HOV lane by Caltrans during that time. Several other plausibility checks are detailed in Table E.22 in Appendix E.

times of day. These provide little evidence that the CAVS policy affected mainline travel times. As we discuss in section 2.7the mainline estimates do not enter in our welfare calculations, but are nonetheless useful to establish that travel time changes are unique to the HOV lane. While suggestive that the policy had no effect in the mainline, we recognize that the mainline estimates are not as precise as would be ideal. In particular, the fact that the mainline and HOV lane estimates cannot be statistically distinguished raises concerns that the increase in HOV travel times is driven by something other than the CAVS policy.[36] We examine this concern and several others related to the mainline below and in the detector level results in Section 2.6.

Effects of the Policy by Time of Day

Thus far, we have examined the effects of the CAVS policy at peak and off-peak periods and calculated the average treatment effect. However, it is likely that these effects will vary even within peak periods, especially during the morning peak. Small (1982) finds that individuals will adjust work-trip departure times in response to changes in congestion and that such behavior can result in heterogeneous responses during peak hours. Hybrid drivers previously commuting in the mainline may depart closer to their preferred time, given access to the less congested HOV lane. Table 2 presents the effect of the policy by hour on HOV lane travel time, under the preferred specification of an eighth-order global polynomial. The distribution of magnitudes mimics the congestion levels noted in Figure 1, such that congested times of day are most affected by the

[36]We provide p-values for tests of the null hypothesis that the HOV and mainline effect are identical as well as a test for induced demand. This no induced demand test examines the null hypothesis that the true mainline effect is the decrease in travel time expected (one-fourth the flow increase of the HOV lane) if all hybrid drivers had originated in the mainline lanes prior to the policy. These tests are generally inconclusive given the mainline standard errors.

CAVS policy. Point estimates are insignificant at 5 A.M. and at their largest from 8 to 9 A.M. These magnitudes grow as rush hour progresses from 8.8 percent at 6 A.M. to 12.2 percent at 9 A.M.[37] The effect of the policy on travel time is again most pronounced during the congested evening peak hours of 5 and 6 P.M. with effects of 6.0 and 8.2 percent respectively.

2.5.2 Further Exploration

The estimates presented above provide strong evidence that the CAVS policy increased travel times during peak hours on the HOV lane while mainline travel times remained unchanged. A skeptic could still argue, however, that the pattern of effects found here are the result of standard seasonal changes such as school year effects, or that overall demand for driving increased around the CAVS policy date, or that there are changes in commuting demand unrelated to the policy, perhaps driven by macroeconomic factors such as unemployment. Here we briefly discuss several robustness checks presented in Table 3. First, one may be concerned that the effects found in the HOV lane are the result of seasonal variation. Table 3 columns II and III present the regression estimates of placebo tests using global regression discontinuity, where in column II, the policy variable becomes active on August 20, 2004, and for column III, it becomes active on August 20, 2006.[38] None of the point estimates are significant at the 10

[37]The asymmetric nature of travel times is also noted by (Arnott, De Palma, & Lindsey, 1993) and (Small, Winston, & Yan, 2005) who document that travel times grow until late in the peak due to persistence of events earlier in the day.

[38]One concern is that the global polynomial results in an unbalanced number of observations on either side of the placebo policy date. In Table E.13in Appendix E, local linear results confirm the global polynomial analysis. In Table E.10 in Appendix E we also present regression estimates of the 2005 policy effect using a traditional difference-in-differences analysis where 2004 or 2006 serve as the control year. Under these specifications, HOV lane point estimates are significant at the 1 percent level and are qualitatively similar to the point estimates from our

percent level in the HOV lane. Second, it is possible that total demand for driving increased around the time of the policy implementation. Table 3 column IV presents global RD estimates using only weekend observations. These estimates show no evidence of an increase in travel time on weekends, suggesting that our results are not driven by a general increase in demand. Table 3 column V shows that adding unemployment to the model has almost no effect on the estimates, suggesting that the polynomial trend is capturing any work-week-specific demand changes related to employment. Finally, Table 3 column VI pools HOV and mainline observations with a single polynomial trend, yielding point estimates similar to our central specification with tighter standard errors. Paired t-tests find that the peak estimates in the HOV and mainline lanes can be rejected as identical, suggesting there was not an across-the-board increase in commuting demand.[39]

2.6 Detector Level Results

Thus far, we have established convincing evidence of the effect of the CAVS policy on travel time on the I-10W. It is unclear, however, if these changes are unique to the I-10W. The detector-level analysis, which is performed across a larger set of roads at both the start and end of the policy, can help generalize these results. Furthermore, the larger data set allows for more precise estimates of the policy effect on mainline lanes. Figure 3 panel (a) plots the kernel den-

central specifications.

[39]The paired t-test, accounting for the covariance between the HOV and mainline estimates, yields t-statistics of 5.07 in the morning peak and 5.26 in the afternoon peak. The detector level regressions below also allow us to statistically distinguish the HOV and mainline estimates in the most congested parts of the city where the effect is most pronounced, but without the restrictive assumptions of a single polynomial trend.

sity smoothed citywide distribution of the effects of the CAVS policy on hourly traffic flow for mainline and HOV lanes at the start of the policy. This figure is generated by separately estimating policy effect coefficients for the hundreds of detectors across the city. While the mainline detectors indicate little evidence of an effect of the policy on traffic flow, as shown in Table 4 column I, there is a citywide positive increase of 5.7 percent in the HOV lane flow, statistically significant at the 1 percent level. Exploring the heterogeneity of the policy effect across the city, Table 4 columns II-IV report the average estimated effect of the policy on traffic flow by distance from downtown LA.[40] For detectors within 10 miles of downtown LA, column II reports a mean effect in the HOV lanes of 9.1 percent, while the mean estimated effect in the mainline lanes is a statistically insignificant 1.5 percent. Figure 3 panel (b) plots the smoothed distribution of hourly flow effects for different spatial subsets, including the subset of detectors on the I-10W. Despite the different HOV lane passenger restrictions, the distribution of effects found on the I-10W is similar to those for detectors located on other freeways near downtown Los Angeles.[41] Furthermore, for detectors located 0-10 miles from downtown LA, this estimation has sufficient power to statistically distinguish the effect of the policy on HOV and mainline lanes, allowing us to reject the possibility that there was a common increase in traffic across both HOV and mainline lanes as a result of increased aggregate demand. Columns III and IV increase the distance from downtown LA to 20 and 30 mile rings. Moving further from downtown LA, the HOV lane effect drops to zero

[40] As noted by (Anas, Arnott, & Small, 1998), L.A. has multiple Central Business Districts (CBD). Here, downtown LA was chosen to be the intersection of the I-10 and I-5 freeways. This corresponds to the area near Union Station in LA.

[41] The I-10W detectors also confirm the route level analysis (detectors between 3 and 20 miles from downtown LA). The traffic flow increase of 9.6 percent observed for the I-10W HOV detectors implies that travel time in the I-10 HOV lane would have increased 6.7, which is statistically indistinguishable from the estimate of 7.2 percent found in the route level analysis using local linear regression (Table E.13 column II in Appendix E).

and becomes statistically insignificant, suggesting that the increases observed are closely tied to congestion and not due to a general increase in demand. Next, we turn to the estimates at the end of the policy. We note that though the city-wide effect is somewhat larger, it is not statistically distinguishable from the start of policy estimates. Nonetheless, the increase in the point estimate is consistent with expectations that the effect of the policy would increase as additional stickers were distributed over time.[42] Finally, we note that the estimates in Table 4 reveal no evidence of a change in flow in mainline lanes. While policy makers may have expected congestion decreases in the mainline to be a potential benefit of the policy, these results are suggestive of the presence of induced demand per the fundamental law of highway congestion (Downs, 1962; Vickrey, 1969; Duranton & Turner, 2011). Although our results are consistent with the presence of induced demand, it is not possible to determine whether it is the result of new Vehicle Miles Traveled or of diverted demand from other routes or times of day. The source of this induced demand is important, as new VMT generated by individuals commuting more frequently or switching from other modes of travel can increase emissions, undermining a stated goal of the policy.[43] In the best-case scenario for emissions reductions, diverted demand from other routes or times of the day is the source of induced demand. In the welfare analysis below, we argue that our final conclusions are robust to any of these scenarios.

[42]The larger point estimate at distances far from the CBD at the end of the policy is also consistent with the idea that commuters in more congested areas of the city would have been more inclined to acquire stickers at the start of the policy, while more remote commuters may have delayed their acquisition of stickers.

[43]New VMT from public transportation users is likely to be small as only 4.1 percent of commuters in Los Angeles use bus and 0.7 percent use rail (State of the Commute Report, 1999).

2.7 Welfare

2.7.1 Conceptual Framework

Overall Welfare EffectsIn the spirit of the literature on taxation in a second best setting (Harberger, 1974; Bovenberg & Goulder, n.d.), here we outline the key welfare effects of the CAVS policy with the aid of a welfare formula and diagrams. We provide more details and mathematical formulas used to calculate the effects in Appendix B. To begin, consider a classical transportation network (Vickrey, 1969) where a fixed number of agents N select between the HOV and mainline lane. In both HOV and mainline lanes markets, distortions stem from the fact that agents ignore external congestion and pollution costs when making driving decisions, generating a wedge between the marginal private cost and the marginal social cost of using a vehicle. By allowing hybrid vehicles into the HOV lane, the CAVS policy generates a tension between congestion relief benefits in the mainline lane and congestion costs in the HOV lane.[44] Let W be the total social cost, including congestion costs imposed on carpoolers as well as carpool formation costs, congestion costs of the HOV lane on single occupant hybrids, congestion costs of travel in the mainline lane, emissions costs of travel in hybrids and regular vehicles, and private net cost of driving hybrid and regular vehicles. Let the social cost function be defined as:

$$
\begin{aligned}
W =& 3C_C v_C T_H(C_H) + 3C_C \tau_C + \overline{H} v_{\overline{H}} T_{\overline{H}}(C_H) + N_{ML} v_{ML} T_{ML}(N_{ML}) + \\
& E_R(C_C + N_{ML}) + E_{\overline{H}} \overline{H} + B_R(C_C + N_{\overline{ML}}) + B_{\overline{H}} \overline{H}
\end{aligned}
$$

[44]The higher level of congestions in the mainline lane suggests larger congestion relief benefits, while the triple occupancy in the HOV lane suggests larger costs. Which effect dominates is an empirical matter to be resolved below.

where C_C is the number of 3-occupant carpool vehicles, $\overline{H}$ is the number of stickered hybrids, and N_{ML} is the number of mainline drivers. Value of time for each type of agent is v_i. $T_H(C_H)$ is HOV travel time as a function of the number of HOV lane vehicles where $C_H = C_C + \overline{H}$, τ_C is the transaction cost of carpool formation paid by carpoolers, and $T_{ML}(N_{ML})$ is mainline travel time as a function of the number of mainline drivers. E_R is the external emissions cost of regular vehicles, E_H is the smaller external emissions cost of hybrid vehicles, B_R is the private net cost of driving a regular vehicle and $B_{\overline{H}}$ is the private net cost of driving a hybrid vehicle. Figure 4 depicts the equilibrium in the HOV lane and the mainline lanes, which we assume to be the only distorted markets in the economy. It also displays the market for vehicles with stickers. Suppose a sticker is issued that allows an additional driver of a qualified hybrid vehicle into the HOV lane without having to carpool, such that $\frac{(dC_H)}{(d\overline{H})} = 1$ and $\frac{(dN_{ML})}{(d\overline{H})} = -1$. The marginal welfare effect of an additional hybrid vehicle using the HOV lane can be determined via the change in total social cost:

$$\frac{dW}{(d\overline{H})} = (E_H - E_R) + (3C_C v_C + \overline{H} v_{\overline{H}})\frac{(dT_H)}{(dC_H)} + v_{\overline{H}}(T_H - T_{ML}) - v_{ML} N_{ML} \frac{(dT_{ML})}{(dC_{ML})} \quad (2.2)$$

The first term $(E_H - E_R)$ is the marginal primary welfare gain, arising from the fuel economy improvements associated with the adoption of hybrid vehicles induced by the CAVS policy. Integrating over the number of adopted hybrids gives the total effect, equal to the reduction in the external cost of emissions of mainline drivers (given by the area abcd in Figure 4) net of the increased external costs of emissions of hybrids in the HOV lane, (given by the area efgh). The second term $(3C_C v_C + \overline{H} v_{\overline{H}})\frac{(dT_H)}{(dC_H)}$ is the marginal cost-side congestion interaction effect, defined as the welfare loss for HOV lane drivers that results from the policys goal of inducing hybrid vehicles into the HOV lane. The total effect equals the area fgji and reflects the value of the travel time delays in the HOV lane caused by the additional hybrids. The third term $v_{\overline{H}}(T_H - T_{ML})$ is the marginal rent effect equal to the travel time cost-savings for the hybrid vehicle from HOV lane access. It represents the maximum willingness to pay for a hybrid vehicle with a sticker, where given the fixed supply of stickers, rents will be generated per (Bento & Jacobsen, n.d.).[45] The total effect is given by the rectangle *klmn*. This rectangle is effectively equal to the difference between the area *padq* (in the mainline market) and the area *rehs* (in the HOV lane market). As we shall see below, a priori, it is not obvious whether hybrid drivers appropriate this rent. The final term $-v_{ML}N_{ML}\frac{(dT_{ML})}{(dC_{ML})}$ is the congestion relief in the mainline lanes. In reality, the number of agents is not fixed, and congestion relief in the mainline may induce demand from other transportation options.[46] The term above is the marginal partial equilibrium congestion interaction effect, which we argue represents the upper bound of the congestion relief benefits for all other drivers in the freeway system. The total effect is given by area *btuc*.[47] We explicitly note that the estimates of the effect of the policy on mainline travel times do not enter the calculation of this source of welfare below. This welfare effect is exclusively calculated from the estimated number of hybrid drivers that leave the mainline

[45]We assume that there are a sufficient number of drivers who are approximately indifferent between a hybrid vehicle and their next best alternative, such that the private net costs of driving a regular or a hybrid vehicle are equal ($B_R = B_{\overline{H}}$). Therefore, the only cost incurred to gain access to the HOV lane is the minimal administrative cost of the sticker. To the extent this not true, there would be second order welfare loss that results from individuals choosing a less preferred vehicle. This would manifest as a downward sloping willingness to pay for the travel time savings of a stickered hybrid, reducing the magnitude of the rent effect. Because the program was relatively small, we also abstract from any potential general equilibrium changes in the value of all other vehicles.

[46]For example, drivers on less congested alternative routes will now replace the exiting hybrids on the congested travel route, dissipating congestion relief for the original drivers. In turn, they may be replaced by drivers from backroads, or even new VMT. As noted in Table 5, we assume new VMT accounts for 15% of these trips.

[47]In Appendix C, we demonstrate the conditions under which the partial equilibrium congestion interaction effect derived above will upper-bound the system-wide congestion interaction effect. The intuition is simple. The hypothetical congestion relief benefit is larger the greater the external costs of congestion. Therefore in other freeways or travel options not as congested as the mainline, the potential benefit cannot be as large as the benefit in the mainline.

for the HOV lane, and the initial level of congestion in the mainline. Creating congestion relief benefits within the freeway system leads to re-allocation of agents across freeways and potentially new trips and vehicle miles travelled (Hymel, Small, & Dender, n.d.). To the extent that policy induces new vehicle miles travelled, it may generate an additional negative source of welfare corresponding to the value of the external costs of emissions associated with these trips. This effect is represented by the area *abvw*. While our welfare calculations are general equilibrium in nature, the partial equilibrium calculations derived from the estimates in the preceding sections, in particular the effects of the policy on travel time in the HOV lane and the implied number of drivers leaving the mainline to the HOV lane, serve as the key parameters for the welfare analysis. Distributional Impacts: Who appropriates the rents generated by the program?While it is obvious that carpoolers will be made worse off by the policy, a priori it is not obvious who benefits from the policy. We note that for a mainline driver to move into the HOV lane, he must experience some gain. If the agent already owns a qualified hybrid, he will appropriate the overall benefits of travel time reductions and will only pay for the cost of the sticker; alternatively, if the agent does not own a qualified hybrid, his maximum willingness to pay for a hybrid vehicle with a sticker will reflect the benefits of travel time reductions he would experience (given by the difference between *dq* and *hs*). This is to say that other agents in the system, including dealers of new hybrid vehicles or used hybrid sellers, have the potential to extract part of this willingness to pay for the sticker and appropriate some of the rents generated by the program.

2.7.2 Welfare Effects

Table 6 displays estimates and confidence intervals of the annual and present value welfare effects of the CAVS policy on the I-10W, broken down by welfare source.[48] It underscores the following key results. First, the net welfare effect of the policy was negative and equal to -$1.6 million dollars with congestion relief benefits for other drivers included and -$3.3 million dollars without, and the confidence interval does not include zero. Over the nearly six years of the policy, this represents a discounted net welfare loss of -$8 million to -$18 million dollars. The primary welfare gain from the policy is roughly $28,000 per year, representing the emissions savings benefits if all hybrid vehicles on the I-10W were purchased because of the CAVS policy. However, as (Shewmake & Jarvis, 2011) note, roughly two-thirds of all stickers were distributed to hybrid vehicles registered before the start of the CAVS policy, suggesting the primary welfare gain may be smaller.[49] By contrast, the interaction effects of the policy are substantially larger. The cost-side congestion interaction effect from increased HOV lane congestion is substantial and is estimated to be approximately -$4.0 million per year.[50] The rent effect associated with HOV lane access

[48]The welfare calculations rely on the estimates from Sections 2.5 and 2.6, additional PeMS and Caltrans data, as well as parameters from the literature. Table 5 displays these parameters, with additional discussion of parameter choices in Appendix B. Travel time and flow are linked following (Burger & Kaffine, n.d.) and discussed in Appendix B. Confidence intervals are generated from the standard errors of the hourly travel time effects in Table 2. All calculations in Table 6 (including emissions effects) are for a one-way commute on the I-10W. The results in Section 2.6 suggest that the policy would have a similar effect on travel time and emissions on other core freeways throughout Los Angeles. The end-of-policy results in Section 2.6 are slightly larger than those from the beginning of the policy. We use the beginning-of-policy estimates as a lower bound and report welfare effects using the end-of-policy estimates in Appendix E Table E.26. At the conclusion of this section, we present back-of-the-envelope calculations of the statewide effect of the CAVS policy.

[49]Furthermore, (Diamond, n.d.) and (Gallagher & Muehlegger, 2011) find no evidence that the CAVS policy stimulated hybrid purchases, implying emissions savings and the primary welfare gain may be near zero.

[50]These values are an upper bound on congestion costs if some marginal carpoolers broke their carpools. However, it is unlikely that this represents a significant effect, as our estimates

for hybrids is $672,000 per year.[51] The system-wide congestion interaction effect for all other drivers in the transportation system is upper-bounded at $1.7 million dollars per year, representing the partial equilibrium congestion relief benefit for I-10W mainline drivers.[52] Finally, the emissions interaction effects arising from induced VMT are very small, at -$6,275 annually for greenhouse gas emissions, -$714 for NOx emissions, and -$251 for hydrocarbons.[53] Thus, the congestion costs arising from the policy interaction dominate the overall welfare effect. Given the large negative effects of the policy on congestion, sensitivity analysis in Appendix B finds that varying relevant parameters within reasonable distributions does not affect the key finding that the overall welfare impacts are negative. Because congestion is not priced, policies that exacerbate the congestion externality increase the cost of clean technology promotion. This is particularly true in congested areas like LA, as adding a car to the HOV lane requires either an occupancy of 10 people or a solo hybrid driver with a value of time of $200 in order to compensate for the congestion externality generated. Second, Table 6 illustrates the distributional costs of the policy. The CAVS policy reduced welfare for the average carpooler by $176 dollars per year. In contrast, dividing the rent effect by the 904 hybrids per day estimated to use the I-10

---

of the number of hybrids entering the HOV lane based on increased travel time are consistent with physical hybrid counts conducted by Caltrans.

[51]Small, Winston and Yan (2005) find that drivers may be willing to pay for increased reliability of travel times. Following their procedure, an additional reliability benefit of $100,000-$150,000 per year is created. On the other hand, hybrid owners entering the HOV lane will decrease reliability for carpoolers. Due to the relatively small magnitude of benefits and offsetting effects, we exclude reliability benefits and focus on travel time.

[52]This upper-bound is likely to be a substantial overestimate. Estimates in Section 2.6 find evidence of induced demand, implying that either hybrids originated from less congested transportation options or drivers re-optimized their travel decisions and replaced exiting hybrids.

[53]Greenhouse gas emissions from induced new VMT (agents induced to drive by the absence of hybrid vehicles in the mainline) are calculated to be 300 tons, a small increase equivalent to the yearly emissions of 50 average fuel-economy vehicles. Similarly, NOx and hydrocarbon emissions increased by a slight 0.05 and 0.06 tons.

HOV gives a maximum rent per sticker of $743 dollars per year.[54] These rents were most likely appropriated by hybrid owners. Two-thirds of stickers were distributed to hybrids registered before the policy, and despite the supply constraint, the remaining rents were likely captured by hybrid purchasers rather than dealers or manufacturers, similar to the concurrent tax credits studied by (Sallee, n.d.). We highlight several key findings related to the distributional impacts: First, while total carpooler congestion cost substantially outweighs hybrid benefits, the cost per individual carpooler is relatively small and less than the rents generated per sticker. Second, as hybrid owners are wealthier than the average carpooler, this policy is likely to be regressive. Finally, the diffuse costs spread across nearly 22,000 daily carpoolers and the concentrated benefits of the policy may have enabled policy approval and subsequent extension. Third, we also calculate the costs of transferring $1 dollar to hybrid owners. The transfer ratio is at minimum 3.31 with a maximum of 8.87, implying a striking cost of roughly $3-9 dollars to transfer $1 dollar of benefit per hybrid.[55] Alternatively, hybrid tax credits could be used to stimulate hybrid purchases. Such a credit could be financed through taxes, at a cost of approximately $1.40 per dollar transferred, at a standard labor tax marginal excess burden of 0.4 (Browning, n.d.). While access to HOV lanes is very valuable, the large number of HOV lane users and heavy preexisting congestion implies substantial costs are cre-

[54] To validate our assumption that there are a large number of households who are indifferent between hybrid and regular vehicle ownership prior to the policy, we investigated what this value would imply for the premium a household is willing to pay for a hybrid with a sticker. Doubling the $743 dollars a year benefits (for a two-way commute) and privately discounting it (5 percent) over the initially proposed life of the policy gives a net present value of the sticker of $4,800 on the I-10W. Our derived value is similar to that presented in (Shewmake & Jarvis, 2011), who estimate an average premium of $3,200 for a stickered hybrid, as well as suggestions of a $3,000-$5,000 premium for a stickered hybrid from some in the auto industry (http://hffo.cuna.org/12433/ article/2599/html).

[55] The lower bound of 3.31 assumes that hybrid drivers appropriated all the rents and includes congestion relief, while the upper bound of 8.87 assumes that manufacturers appropriated one-third of the rent and excludes congestion relief. Using the end of policy estimates, the bounds increase to 4.20 and 11.26.

ated when transferring benefits to hybrids via the CAVS policy. Finally, under the most optimistic scenario for the CAVS policy every sticker stimulated a hybrid purchase and induced demand does not occur we generate back-of-the-envelope-calculations of the statewide cost per ton of emission reductions. Because the I-10W is a particularly congested freeway, it may be inappropriate to simply apply the estimated welfare effects calculated above. Appendix B presents our methodology for the back-of-the-envelope calculation and sensitivity analysis on key parameters. We find that the best-case cost per ton of GHG emissions reductions is $124 per ton, with costs per ton of NOx and hydrocarbon reductions of $606,000 dollars and $505,000 dollars.[56] Even under this extremely generous scenario, this is roughly an order of magnitude larger than estimates of the marginal social cost of GHG emissions and substantially larger than other emission control options.[57]

2.8 Conclusion

This paper employs a regression discontinuity design to estimate the interaction effect between the Clean Air Vehicle Sticker policy and unpriced congestion in

[56] To the extent that the end-of-policy estimates in Section 2.6 are suggestive of larger policy impacts on flow and thus congestion costs, scaling up the cost of GHG reductions accordingly yields a GHG cost per ton of $160. Accounting for the fact that it is unlikely that all stickers stimulated purchases of hybrid vehicles will raise the cost per ton; removing the two-thirds of stickers received by preregistered hybrids increases the GHG cost to $482 per ton. On the other hand, using mean estimates of the effect of early adoption on mainstreaming hybrid adoption (0.6 additional purchases of hybrids per early adopter) from (Heutel & Muehlegger, 2010), the cost per ton would fall to $310.

[57] California Assembly Bill 32 Scoping Plan is a comprehensive study of the cost of reducing greenhouse gases prepared by the California Air Resources Board. The Scoping Report considered a wide-range of policies, with estimated costs per ton of emissions ranging from -$300 for greenhouse gas standards for vehicles, to $300 for additional solar water heaters. (Chandra, Gulati, & Kandlikar, n.d.) estimate that the cost of GHG emissions reductions from tax rebates for hybrid drivers was $195 per ton, while (Li, Linn, & Spiller, n.d.) estimate that the Cash-for-Clunkers program reduced greenhouse gas emissions at a cost of $91-$301 per ton.

Los Angeles. Although policies that allow single-occupant hybrid vehicles in HOV lanes are viewed as a free method to stimulate hybrid demand, we provide evidence that the CAVS policy in California resulted in substantial welfare losses. We show that the losses from the interaction between the policy and unpriced congestion overwhelmingly dominated the primary welfare gain from increased environmental benefits associated with the adoption of new hybrid vehicles. Our results also underscore the remarkable variation of the interaction effect across space and time, whereby adding one daily hybrid driver at 7AM on the I-10W generates $4500 in annual social costs. When incorporated into a welfare analysis, our econometric estimates imply a best-case cost of $124 per ton for reductions in greenhouse gas emissions, $606,000 dollars per ton of nitrogen oxides (NOx) reduction, and $505,000 dollars per ton of hydrocarbon reduction in the most optimistic case. These are substantially above other readily available options to policymakers. The results presented here have important implications for policy. Given the substantial cost of the CAVS policy, at a minimum it is worth considering alternative policies that may have achieved similar goals at lower cost. While a hybrid tax credit would cost taxpayers $1.40 per dollar transferred to hybrid owners, this cost increases to $3.31-$8.87 under the CAVS policy. Therefore tax credit incentives funded through distortionary taxes would be preferred to the CAVS policy despite its perception of being free. Earlier literature on environmental policy in a second best setting highlighted the superiority of revenue-raising instruments over nonrevenue raising instruments (Goulder et al., 1997). In our context, a natural option would be to auction the stickers to hybrid drivers. In this case, recycling auctioned revenues broadly by cutting preexisting distortions would only reduce the costs per dollar transferred to hybrid owners to $2.91-$8.47. If revenues were instead used to compensate the carpoolers in full, hybrid drivers value of time would need to be in excess of $200/hour to raise enough revenues to offset carpoolers for the value of lost travel time. Whether auctioned or not, the major source of the inefficiency of the CAVS stickers comes from the fact that this policy is blind to the heterogeneity of the external costs of congestion across time and space. Alternatively, policymakers could ration HOV access via congestion pricing. Ideally, a High Occupancy Toll (HOT) could consider both the external costs of congestion and air pollution. While the emissions discount for hybrid vehicles would be invariant across space and time (roughly 0.7 cents/mile), the congestion fee itself would still adjust to reflect prevailing congestion conditions, as in (Keeler & Small, 1977). For example, on the I-10W during peak periods, the congestion fee would be roughly 45 cents/mile.[58] The fact that the congestion fee would be at least 60 times higher than the discount to hybrid vehicles underscores the significance of the interaction effect studied here. Moving forward, policymakers are already replacing CAVS with new policies to promote the adoption of plug-in hybrids. Starting January 2012 40,000 stickers were issued to plug-in hybrid vehicles allowing them to drive in HOV lanes. They have also allowed for an unlimited number of stickers for electric, zero emissions vehicles. More broadly, our findings imply that, even if these vehicles were truly zero-emission, promoting their adoption at the expense of exacerbating congestion will still generate substantial welfare losses. In contrast, promoting the use of buses in HOV lanes, although a far less celebrated technology, may represent the win-win in terms of pollution and congestion that policymakers were hoping with the CAVS policy.

[58]The discount per mile for hybrid drivers reflects the fact that the emissions reduction benefits they generate are time and space invariant. The congestion fee per mile varies with the level of congestion, and is calculated as the cost per mile delay imposed on other drivers.

Figure 2.1: PeMS Detectors in District 7

Figure 2.2: Average Travel Time by Hour

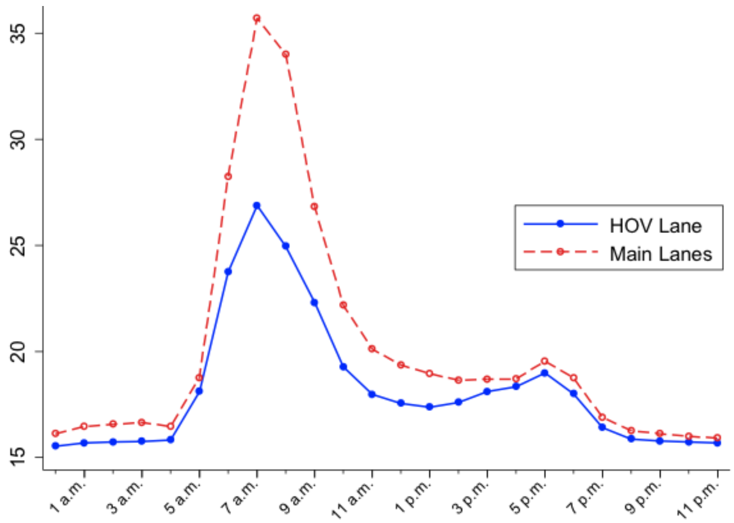

Figure 2.3: Interstate 10 West Travel Time

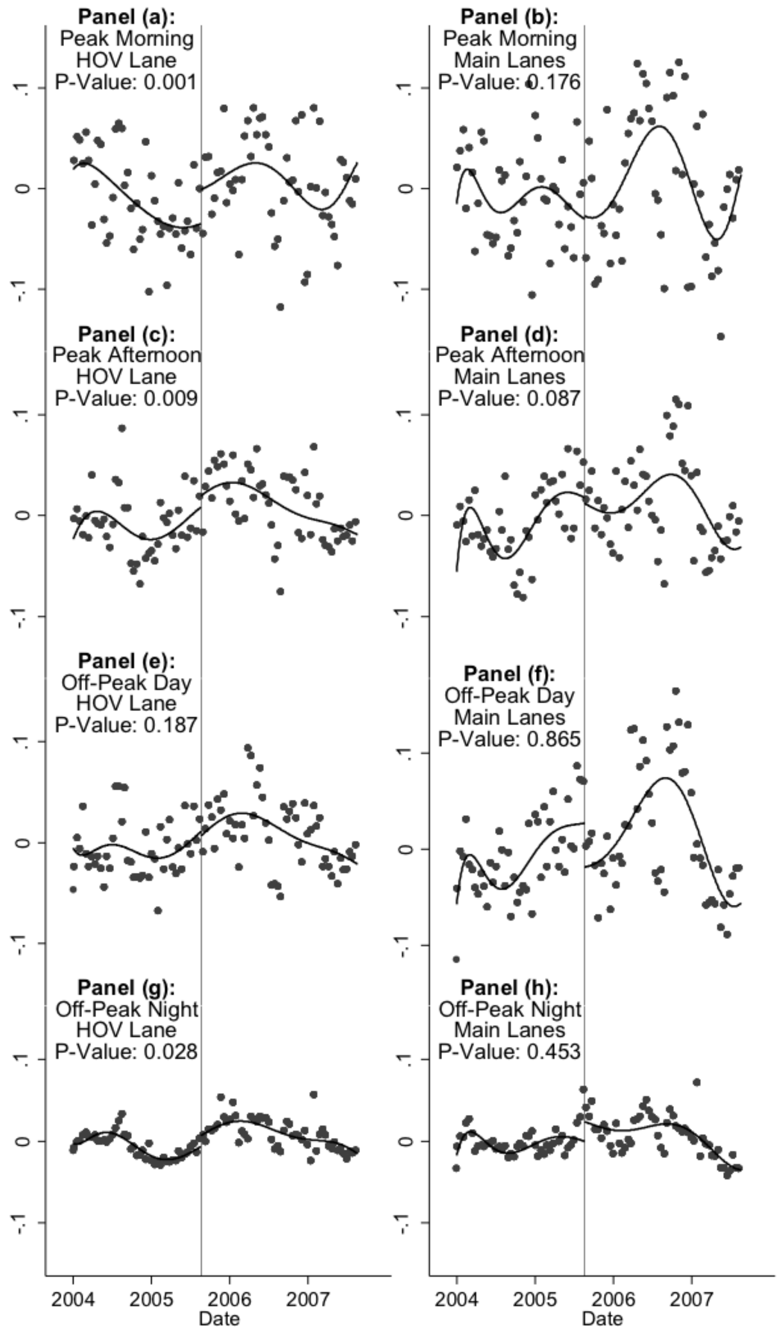

Figure 2.4: Distribution of Detector Level RD Estimates for Flow

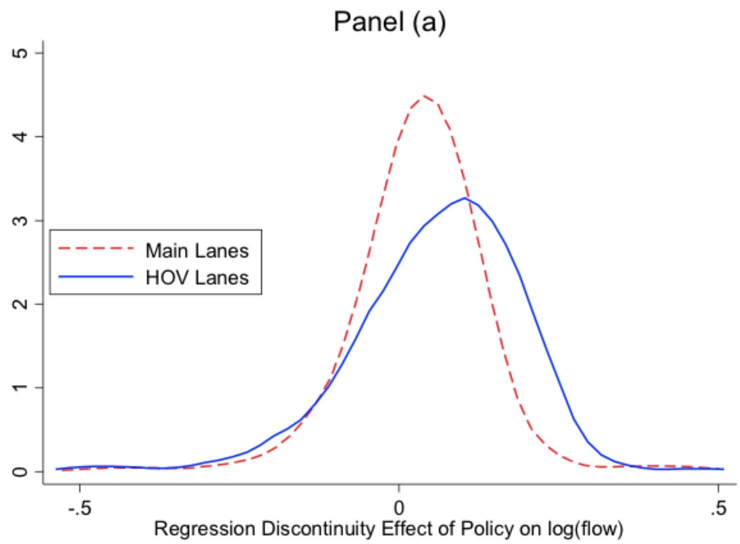

Figure 2.5: Distribution of Detector Level RD Estimates for Flow

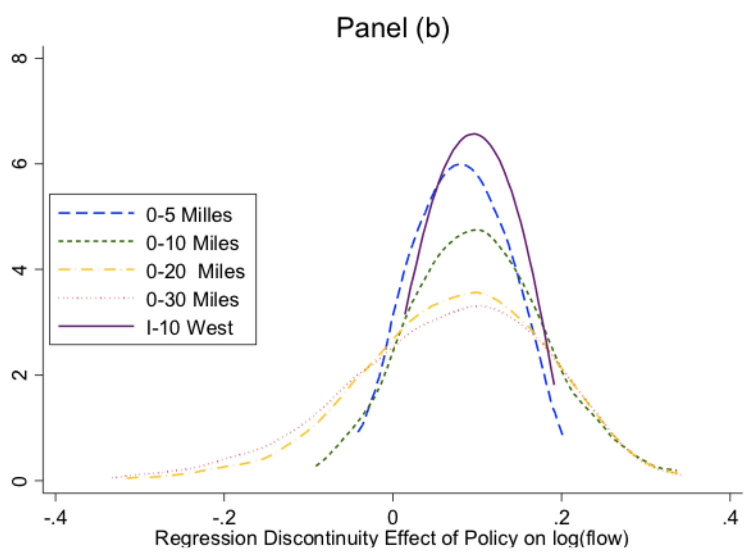

Figure 2.6: Welfare Effects of the CAVS Policy

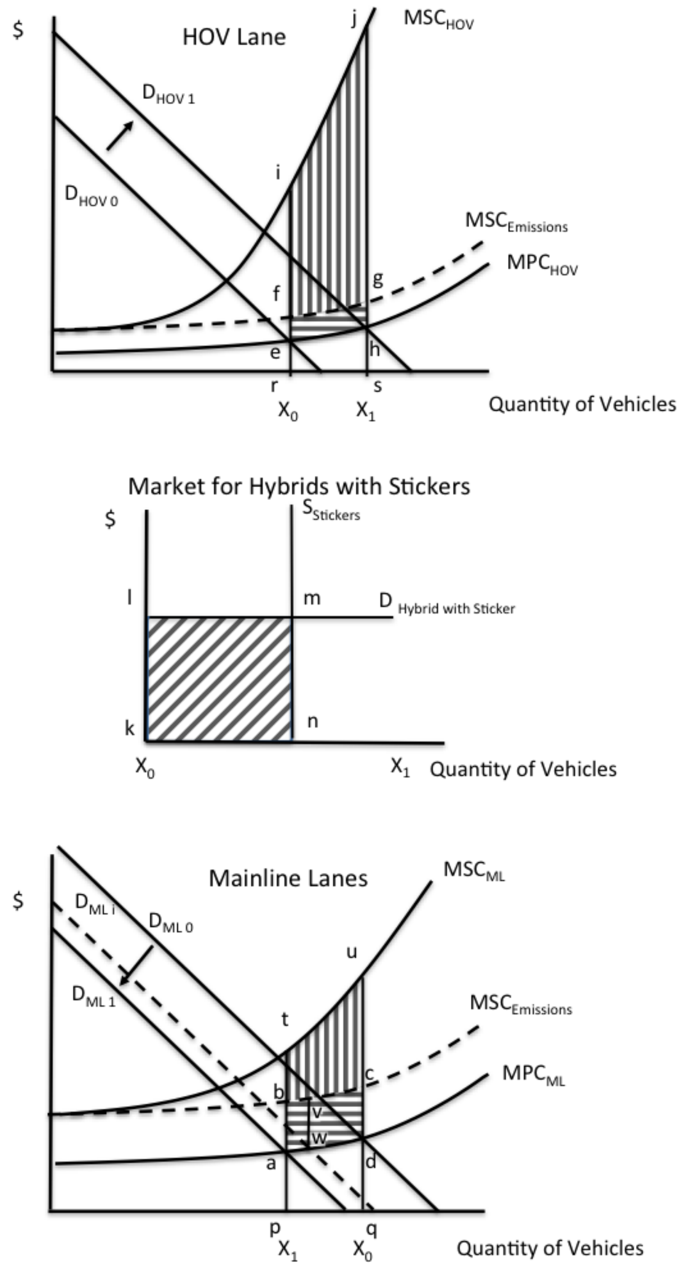

Table 2.1: Regression Discontinuity Estimates: Global Polynomial Results

	I	II	III	IV	V	VI
Polynomial order	6	7	8	9	10	BIC
			Morning peak			
CAVS policy/ HOV	0.088***	0.084***	0.090***	0.090***	0.116***	0.116***
	-0.026	-0.028	-0.028	-0.027	-0.026	-0.026
CAVS policy/ Mainline	-0.011	-0.022	0.06	0.06	0.070*	0.06
	-0.049	-0.053	-0.051	-0.051	-0.049	-0.051
Observations	4944	4944	4944	4944	4944	4944
P-Value for Induced Demanda	0.823	0.989	0.112	0.146	0.052	0.093
P-Value for Test of Difference between HOV and Mainline	0.084	0.088	0.607	0.77	0.41	0.332
			Afternoon peak			
CAVS policy/ HOV	0.063***	0.066***	0.057***	0.057***	0.067***	0.062***
	-0.021	-0.021	-0.021	-0.022	-0.024	-0.015
CAVS policy/ Mainline	-0.001	0.004	0.059*	0.059*	0.055	0.059*
	-0.031	-0.028	-0.036	-0.036	-0.037	-0.036
Observations	3952	3952	3952	3952	3952	3952
P-Value for Induced Demanda	0.64	0.472	0.051	0.051	0.062	0.046
P-Value for Test of Difference between HOV and Mainline	0.1	0.087	0.963	0.973	0.782	0.934

Table 2.2: Regression Discontinuity Estimates: Global Polynomial Results

	Mid-day off-peak					
CAVS policy/ HOV	0.048**	0.050**	0.027	0.027	0.041	0.041**
	-0.021	-0.021	-0.021	-0.021	-0.021	-0.021
CAVS policy/ Mainline	-0.064	-0.063	-0.008	-0.009	-0.007	-0.009
	-0.042	-0.042	-0.049	-0.051	-0.051	-0.05
Observations	5927	5927	5927	5927	5927	5927
P-Value for Induced Demand	0.227	0.24	0.975	0.969	0.955	0.975
P-Value for Test of Difference between HOV and Mainline	0.021	0.021	0.511	0.518	0.386	0.363
	Night off-peak					
CAVS policy/ HOV	0.033***	0.033***	0.016**	0.016**	0.022***	0.022***
	-0.008	-0.008	-0.007	-0.007	-0.008	-0.008
CAVS policy/ Mainline	-0.008	-0.005	0.016	0.017	0.023	0.023
	-0.019	-0.017	-0.022	-0.021	-0.021	-0.021
Observations	8894	8894	8894	8894	8894	8894
P-Value for Induced Demand	0.976	0.836	0.361	0.343	0.197	0.197
P-Value for Test of Difference between HOV and Mainline	0.062	0.055	0.971	0.964	0.992	0.992

Table 2.3: Hourly Global Polynomial Regression Discontinuity Estimates

	I	II	III	IV	V	VI
			Morning peak			
	5 A.M.	6 A.M.	7 A.M.	8 A.M.	9 A.M.	10 A.M.a
Policy (HOV)	0.027	0.088**	0.096**	0.125***	0.122	0.072*
	-0.025	-0.04	-0.042	-0.037	-0.038	-0.038
Observations	989	989	989	988	989	988
			Afternoon peak			
	4 P.M.	5 P.M.	6 P.M.	7 P.M.	8 P.M.a	
Policy (HOV)	0.045	0.060**	0.082***	0.038	0.007	
	-0.027	-0.024	-0.028	-0.025	-0.017	
Observations	988	988	988	988	988	

Table 2.4: Regression Discontinuity Estimates: Further Robustness Checks

	I	II	III	IV	V	VI
			Morning Peak			
	True Date	Placebo		Weekend	Macroeconomic	
Policy (HOV)	0.090***	-0.006	-0.073*	0.002	0.091***	0.102***
	-0.028	-0.034	-0.039	-0.014	-0.027	-0.035
Observations	4944	4944	4944	1980	4944	
Policy (Mainline)	0.06	-0.008	-0.024	0.011	0.06	0.048
	-0.045	-0.032	-0.06	-0.02	-0.045	-0.036
Observations	4944	4944	4944	1980	4944	9888
			Afternoon Peak			
	True Date	Placebo		Weekend	Macroeconomic	
Policy (HOV)	0.057***	-0.032	-0.006	0.023	0.056***	0.081***
	-0.022	-0.025	-0.029	-0.043	-0.021	-0.026
Observations	3952	3952	3952	1584	3952	
Policy (Mainline)	0.059*	-0.042**	0.04	0.023	0.059*	0.034
	-0.035	-0.021	-0.048	-0.046	-0.034	-0.027
Observations	3952	3952	3952	1584	3952	7904
Policy date	8/20/05	8/20/04	8/20/06	8/20/05	8/20/05	8/20/05
L.A. Unemployment	N	N	N	N	Y	Y
Pooled, single trend	N	N	N	N	N	Y

Table 2.5: Regression Discontinuity Estimates: Further Robustness Checks

	I	II	III	IV	V	VI
			Mid-Day Off Peak			
	True Date	Placebo		Weekend	Macroeconomic	
Policy (HOV)	0.027	-0.017	-0.002	-0.02	0.027	0.028
	-0.021	-0.027	-0.03	-0.031	-0.021	-0.031
Observations	5928	5928	5928	2376	5928	
Policy (Mainline)	-0.008	-0.028	0.048	-0.02	-0.008	-0.009
	-0.049	-0.026	-0.058	-0.038	-0.048	-0.036
Observations	5927	5927	5927	2376	5927	11855
			Night Off Peak			
	True Date	Placebo		Weekend	Macroeconomic	
Policy (HOV)	0.016**	-0.003	-0.008	0.008	0.015**	0.039***
	-0.007	-0.01	-0.012	-0.019	-0.007	-0.02
Observations	8894	8894	8894	3554	8894	
Policy (Mainline)	0.016	-0.005	-0.014	0.002	0.017	-0.007
	-0.022	-0.006	-0.022	-0.021	-0.021	-0.014
Observations	8894	8894	8894	3554	8894	17788
Policy date	8/20/05	8/20/04	8/20/06	8/20/05	8/20/05	8/20/05
L.A. Unemployment	N	N	N	N	Y	Y
Pooled, single trend	N	N	N	N	N	Y

Table 2.6: Average of Local Linear Treatment Effects by Distance from CBD

	I	II	III	IV
Distance from CBD	0-10	10-20	20-30	0-30
		Start of Policy		
CAVS policy/ HOV	0.091**	0.058**	-0.009	0.055**
	-0.028	-0.029	-0.102	-0.025
Detectorsa	50	124	26	200
Observationsb	8,297	20,662	4,110	33,069
Flowc	915	845	686	842
CAVS policy/ Mainline	0.015	0.008	0.021	0.012
	-0.017	-0.03	-0.022	-0.017
Detectorsa	152	254	71	477
Observationsb	24,461	40,863	10,682	76,006
Flowc	5,108	5,248	4,616	5,109
P-value test of difference in HOV and Mainline	0.019	0.127	0.35	0.186
Implied Number of Vehicles Removed from Mainlined	-83	-49	6	-46
Mainline Null Hypothesis without Induced Demande	-0.016	-0.009	0.001	-0.009
P-value for test of Induced Demandf	0.07	0.556	0.393	0.211

Table 2.7: Average of Local Linear Treatment Effects by Distance from CBD

	I	II	III	IV
Distance from CBD	0-10	10-20	20-30	0-30
		End of Policy		
CAVS policy/ HOV	0.088***	0.048*	0.116***	0.071***
	-0.025	-0.024	-0.043	-0.018
Detectorsa	21	46	15	82
Observationsb	3,579	7,677	2,409	13,665
Flowc	917	885	557	833
CAVS policy/ Mainline	0.003	0.008	0.012	0.007
	-0.008	-0.005	-0.013	-0.005
Detectorsa	74	131	44	249
Observationsb	12,674	22,441	7,040	42,155
Flowc	5,078	5,195	4,568	5,049
P-value test of difference in HOV and Mainline	0.001	0.102	0.021	0.001
Implied Number of Vehicles Removed from Mainlined	-81	-42	-65	-59
Mainline Null Hypothesis without Induced Demande	-0.016	-0.008	-0.014	-0.012
P-value for test of Induced Demandf	0.02	0.002	0.048	0

Table 2.8: Additional Welfare Analysis Parameters

	Value	Source
I-10W route length	17.5 miles	PeMS
Value of time (Hybrid drivers)	\$32.86	Small, Winston, and Yan (2005)
Value of time (Carpoolers)	\$20.87	Small, Winston, and Yan (2005)
I-10W HOV lane occupancy per vehicle	3.1	Caltrans
I-10W mainline occupancy per vehicle	1.1	Caltrans
Elasticity of new VMT (short-run)	0.15	Hymel, Small, and Van Dender (2010)
Hybrid fuel efficiency	45 mpg	EPA
Fleet fuel efficiency	20 mpg	EPA
Hybrid NOx emissions per mile	0.02 grams	CA SULEV standards
Hybrid hydrocarbon emissions per mile	0.01 grams	CA SULEV standards
Fleet NOx emissions per mile	0.07 grams	EPA Tier II standards
Fleet hydrocarbon emissions per mile	0.09 grams	EPA Tier II standards
Marginal social damage of GHG emissions	\$21/ton	US Interagency Working Group
Marginal social damage of NOx emissions	\$15,000/ton	Small and Kazimi (1995)
Marginal social damage of hydrocarbon emissions	\$4,100	Small and Kazimi (1995)

Table 2.9: Welfare Effects of the CAVS Policy for the I-10 West

Panel A Welfare Effects	Annual	Present Value
Primary welfare gain	$28,127	$283,449
Cost-side congestion interaction effect	-3,990,620	-21,495,008
	[-$7,024,200, -$936,100]	
Rent effect	671,882	2,363,240
	[$186,650, $1,153,750]	
System-wide congestion interaction effect	1,744,620	9,397,209
	[$423,500, $3,057,000]	
Emissions interaction effect	-7,240	-38,998
Net welfare effect of the CAVS policy for the I-10W	-1,553,225	-8,366,274
	[-$2,776,700, -$321,200]	
Excluding system-wide congestion interaction effect	-3,297,846	-17,763,483
	[-$5,829,700, -$750,000]	
Panel B Distributional Effects		
Carpoolers using I-10W HOV lane (daily)	21,943	-
Hybrids using I-10W HOV lane (daily)	904	-
Congestion cost per carpooler	-176	-948
Rent per stickera	$743	$4,002
Transfer ratio- lower bound	3.31	-
	[2.86, 3.46]	
Transfer ratio- upper bound	8.87	-
	[8.09, 9.07]	

CHAPTER 3

IS THERE AN ENERGY PARADOX IN FUEL ECONOMY? A NOTE ON THE ROLE OF CONSUMER HETEROGENEITY AND SORTING BIAS

3.1 Introduction

Although economic theory suggests that rational consumers should be willing to pay \$1.00 more for a vehicle that saves them \$1.00 in discounted future fuel costs, a growing body of literature finds a marginal willingness-to-pay (MWTP) for reduced discounted future fuel costs ranging from \$0.35 to \$0.79 (Helfand & Wolverton, 2011; D. Greene et al., n.d.). This perceived undervaluation of future fuel costs is an example of an energy paradox in the automobile market. Energy paradox is a general concept used to explain the unexpectedly slow diffusion of apparently cost-effective, energy-efficient technologies that involve similar trade-offs between up-front capital costs and future operating costs (Jaffe & Stavins, 1994). Such a paradox may exist in automobiles and other energy-using durables (Hausman, 1979).[1]

With increasing concerns related to climate change, energy security, and local pollution, many have used this potential market failure to justify policies that promote efficiency-improving technology. Policies that encourage even a small correction in this paradox have the potential to result in sizable decreases in energy use and its related externalities. The magnitude and sources of this paradox have broad implications for any technology that uses energy. In the automobile sector, some of this interest has focused on the debate between the

[1]The typical magnitude of the energy paradox in appliances requires discount rates of 20% although estimates vary widely (see (Hausman, 1979), (Dubin & McFadden, 1984), (Ruderman, Levine, & McMahon, 1987)).

gasoline tax and corporate average fuel economy (CAFE) standards, as well as the design of future CAFE standards. Precise estimates of the MWTP for reduced discounted future fuel costs are central to this debate (Parry et al., 2010). If consumers correctly value future fuel costs, gasoline taxes are found to be less costly than CAFE standards in achieving targeted fuel reductions (Fischer, Harrington, & Parry, 2007; M. R. Jacobsen, 2012b). However, the opposite is true if consumer undervaluation is sufficiently large. Although we focus on the automobile market, this note has implications for the valuation of energy efficiency in a very broad category of purchases.

Our concern with prior literature is that it has often examined the energy paradox ignoring the underlying consumer heterogeneity in MWTP for future reductions in fuel costs. If consumers are heterogeneous in their MWTP, they will sort into vehicles based on vehicle fuel efficiency: those with high MWTP for reduced fuel costs will sort into fuel-efficient vehicles and those with low MWTP will sort into fuel-inefficient ones. We show in this paper that ignoring consumer heterogeneity in the MWTP for future fuel cost in a (multinomial) logit specification could result in heteroskedasticity and bias the estimate of the MWTP toward zero, suggesting spurious undervaluation. The purpose of this note is not to argue whether there is undervaluation of fuel economy or not. Rather our point is that an empirical analysis that ignores consumer heterogeneity may overstate the magnitude of undervaluation. Similar concerns of bias due to sorting were raised in a recent study of the value of a statistical life using labor market data (DeLeire, Khan, & Timmins, 2013).

In Section 3.2, we analytically illustrate the potential for bias from ignoring consumer preference heterogeneity in future fuel cost in a simplified multinomial logit framework. In Section 3.3, we provide further evidence with simulations in a richer model of vehicle demand. In doing the simulations, we first generate data from an equilibrium model of the automobile market and then try to recover the average MTWP for fuel cost using a logit model and a random coefficient logit model.[2] Our analysis shows that, when undervaluation of fuel costs is not present in the data-generating mechanism, the logit model could erroneously suggest significant undervaluation, whereas the random coefficients logit model recovers the true average MWTP.

3.2 Bias Analysis from Ignoring Preference Heterogeneity

, In the context of vehicle demand, we assume that each consumer chooses to buy a new vehicle, from among J models or products, or not to make any purchase (labeled choosing the outside good) in a given period. For ease of exposition, we assume that the utility of consumer i from vehicle choice j only depends on a single dimension of vehicle characteristics, fuel cost (fc). We relax this assumption in the simulations below. The utility of consumer i from vehicle j is

$$u_{ij} = \beta_i fc_j + \epsilon_{ij} \tag{3.1}$$

where the heterogeneous preference β_i has a mean $\overline{\beta}$ and variance σ_β . ϵ_{ij} has an i.i.d. type I extreme value distribution (conditional on fc_j) with a variance of $\sigma_\epsilon = \pi^2/6$. The utility function can be rewritten as:

$$u_{ij} = \overline{\beta} fc_j + \beta_i fc_j + \epsilon_{ij} = \overline{\beta} fc_j + e_{ij} \tag{3.2}$$

[2]Given that consumers have multiple vehicle models from which to choose, the empirical methods are multinomial logit models, but we suppress the word multinomial to save space throughout our paper.

where the variance of the composite error e_{ij}, $var(e_{ij}|fc_j) = \sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2$. Because e_{ij} exhibits heteroskedasticity, the underlying i.i.d. assumption of e_{ij} would be violated if one is to estimate β using the multinomial logit. To analyze bias of the parameters estimated from the multinomial logit model due to heteroskedasticity, we scale the utility and rewrite (2) as

$$\frac{u_{ij}\sigma_\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} = \frac{\overline{\beta} fc_j \sigma_\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} + \frac{e_{ij}\sigma_\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} \tag{3.3}$$

$$u_{ij} = \frac{\overline{\beta} fc_j \sigma\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} + e_{ij} \tag{3.4}$$

$$= \overline{\beta} fc_j + \left[\frac{\overline{\beta} fc_j \sigma\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} - \overline{\beta} fc_j \right] + e_{ij} \tag{3.5}$$

$$= \overline{\beta} fc_j + \overline{\beta} fc_j \left[\frac{\sigma\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} - 1 \right] + e_{ij} \tag{3.6}$$

The normalization makes e_{ij} homoscedastic with a variance of $\pi^2/6$. Equation (3.4) has re-casted the issue of heteroskedasticity to that of omitted variable in a discrete choice model: estimating $\overline{\beta}$ using the multinomial logit model ignoring heteroskedasticity based on equation (3.2) leads to the same problem as estimating $\overline{\beta}$ based on the last line in equation (3.4) while ignoring $z_j = \overline{\beta} fc_j \left[\frac{\sigma\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}} - 1 \right]$. The omitted variable z_j is positive as long as σ_β^2 is not equal to zero (assuming $\overline{\beta}$ to be negative). It is positively correlated with Following (L.-F. Lee, 1982) and (Yatchew & Griliches, 1985) which analyze the omitted variable bias in discrete choice models, the estimate of $\overline{\beta}$ from the multinomial logit model would be biased upward (toward zero). Moreover, a larger σ_β^2 implies a smaller $\frac{\sigma\epsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\epsilon^2}}$ (closer to zero), and a stronger correlation between z_j and fc_j, therefore, the bias would be larger as well. The following two issues would add complications to the above analysis. First, the omitted variable introduces possible misspecification into the multinomial logit model. Second, a utility function with more vehicle characteristics as defined in equation (3.7) in the next section adds confounding factors to the bias in $\overline{\beta}$ when other vehicle characteristics are correlated with the fuel cost variable. Nevertheless, the bias analyzed above is likely to be the dominant issue in addressing our research question. Because the bias does not have a closed-form solution due to the non-linear nature of the model even in the simple specification analyzed above, next we use Monte Carlo simulations that incorporate those two issues to provide further support.

3.3 Simulations

3.3.1 The Equilibrium Model of Automobile Market

The equilibrium model for data generation is composed of a demand side and a supply side. In the demand side, the utility of consumer i from vehicle j is defined as

$$u_{ij} = \alpha_i p_j + \beta_i fc_{ij} + \beta_i fc_j + \gamma_i x_j + \epsilon_{ij} \tag{3.7}$$

where α, *beta*, and γ are individual-specific taste parameters. We define $\theta = \{\alpha_i, \beta_i, \gamma_i\}$. p_j is price of model j. fc_{ij} is the present value of the total expected discounted fuel cost of the vehicle; it is defined by

$$fc_{ij} = \sum_{t=0}^{T_j} \delta_i^t * AVMT_{itj} * gp_{it}^e / MPG_j \tag{3.8}$$

where T_j is the expected lifetime of vehicle model j, δ_i is an individual-specific discount factor, $AVMT_{itj}$ is annual vehicle miles of travel in year t (which usually decreases with the vehicles age), and gp^e_{it} is the expected gasoline price at year t of consumer i. Heterogeneity can arise from any of the elements used to calculate the lifetime fuel cost of the vehicle. x_j is a vector of other vehicle attributes, and ϵ_{ij} is assumed to have a type I extreme value distribution. We normalize the utility from the outside good, u_{i0}, to zero. The probability of household i choosing vehicle j is given by

$$P_{ij} = \frac{exp(\overline{u_{ij}})}{1 + \sum_{jt} exp(\overline{u_{ij}})} \tag{3.9}$$

where $\overline{u_{ij}} = u_{ij} - \epsilon_{ij}$. Given individual choice probabilities, the aggregate demand can be obtained through summation. The supply side is composed of several firms, each producing multiple vehicle models. They engage in Bertrand competition in that each firm chooses prices to maximize its total profit in a given year, taking the products available as fixed. Following the literature, we assume that the marginal cost of each product is constant. The total profit of firm f is

$$\pi^f = \sum_{f \in F} [(p_j - mc_j) q_j(p, \theta)] \tag{3.10}$$

where F is the set of all products produced by firm f, mc_j is the marginal cost, and q_j is the aggregate demand. p is the price vector and it is obtained through the first-order conditions in equilibrium

$$p = mc + \Delta^{-1} q(p, \theta) \tag{3.11}$$

where the element of Δ, Δ_{jr} is zero if j and r are produced by different firms. Otherwise, it is equal to $\partial q_r / \partial p_j$. Given the underlying preference parameters and marginal cost, this equation can be used to compute equilibrium prices and sales.

3.3.2 Data Generation

Through our data generation approach, we aim to mimic the U.S. auto market. Vehicle information comes from the 2001 Wards Automotive Yearbook; vehicle characteristics include miles per gallon (MPG), horsepower, weight, and manufacturer. We construct marginal cost, a function of MPG, horsepower, and weight, for each model based on estimates from (Berry et al., 1996).[3] We randomly choose a set of vehicle models (25 in the baseline simulation) and assume that these models are available in each year from 2001 to 2006, the time span for our analysis. For ease of exposition, we make several demand-side assumptions. For preference parameters, we assume that all consumers have the same preference on all characteristics except fuel costs. In calculating fuel costs, we assume that the discount factor δ, annual vehicle miles of traveled $AVMT$, and expected gasoline price gp^e are all constant across consumers for any given vehicle. We assume a 10% yearly discount rate. Vehicle lifetime and age-specific annual miles of travel for passenger cars and light trucks are from (Lu, 2006). We further assume that expected gasoline prices during a vehicles lifetime are equal to current annual gasoline price (i.e., gasoline price follows a random walk (Anderson, Kellogg, Sallee, & Curtin, 2011)). Annual gasoline prices during 20012006 are from the Energy Information Administration. These simplifying assumptions, innocuous for our conclusion, imply that consumer heterogeneity is manifested only through the consumer-specific taste parameter on fuel cost, β_i. In the baseline simulation, we assume that β_i has a uniform distribution; the range of the distribution affects the degree of consumer heterogeneity. We choose two levels of dispersion for the taste parameter [-4, 0], and [-

[3] We also add a random error term to the marginal cost of each attribute and to the marginal cost of each product based on the standard errors estimated by (Berry et al., 1996). All costs are converted to 2001 dollars.

3,-1]. (Anderson, Kellogg, et al., 2011) using survey data examine the dispersion of predicted gasoline price, defined as the standard deviation of the predictions divided by the mean. They find this dispersion ranges from 30 to 60% in recent years, which roughly correspond to the dispersion of the two uniform distributions, noting that our distribution assumptions are different. Heterogeneity on discount rates, VMT and vehicle lifetime will further increase the dispersion on the parameter, β_i.

We generate data in two steps. First, we generate equilibrium prices for each model, assuming the whole market with 50,000 consumers in each year. Second, based on equilibrium prices, we generate vehicle choices for 20,000 consumers in each year. The choices of these consumers as well as equilibrium prices are taken as data for the estimation.

3.3.3 Estimation

The goal of the estimation is to recover the underlying preference parameters and to obtain consumers MWTP for reduced fuel costs. For ease of exposition, we assume that the econometrician observes all vehicle characteristics relevant to consumers.[4] We employ two methods: a logit model and a random coefficients logit model. The logit model is estimated using the standard maximum likelihood method. As discussed in Train (2003), the appeal of the random coefficients model comes from its ability to incorporate unobserved consumer heterogeneity, which in our context avoids sorting bias. This model is estimated using the simulated maximum likelihood method. To conduct numerical inte-

[4] In real applications, it is important to control for unobserved product attributes. Most recent literature on the energy paradox has explicitly dealt with this issue.

gration in the simulated method, we employ Halton sequences, which are more efficient than direct Monte Carlo sampling.

3.3.4 Results

We find three main results from the Monte Carlo analysis. Result 1: In the presence of heterogeneity, the logit model suggests undervaluation of the MWTP for reduced future fuel costs, even when undervaluation is not present in the data. Support: Panel A in Table 1 shows that consumers undervalue fuel costs by 29%. The parameter estimates on vehicle price and fuel cost implies that consumers are only willing to pay \$0.71 for a \$1.00 reduction in discounted future fuel costs. The bias comes from individuals sorting into vehicles based on their MWTP: those very averse to fuel costs (e.g., with very negative MWTP) purchase vehicles with low fuel costs. The correlation between fuel cost and the average MWTP among consumers who purchase corresponding vehicles is depicted on the left panel of Figure 1 (the correlation coefficient is 0.83). We believe that at least part of the undervaluation found in prior literature could be attributable to this type of sorting bias.

Result 2: The random coefficients logit model correctly identifies the MWTP. Support: Table 1, Panel A shows that, by explicitly modeling consumer heterogeneity, the random coefficients logit model is able to recover the underlying parameters on vehicle price and fuel cost. The implied MWTP is 1, indicating that consumers are willing to pay \$1.00 for a \$1.00 reduction in discounted future fuel costs, consistent with our model assumption.

Result 3: The greater the heterogeneity, the larger the bias from the logit model. Support: The underlying data-generating process in Panel A of Table 1 implies twice the heterogeneity of Panel B. As a consequence, the undervaluation for the logit model in Panel A, 29%, is larger than the 10% undervaluation in Panel B. Table 2 presents Monte Carlo results for alternative specifications. Panel A suggests that increased market power magnifies the bias from the logit model, with the undervaluation going to 37% from 29% in the baseline model in Table 1. Increasing the number of vehicle draws (Table 2, Panel B) slightly decreases the undervaluation from 29% to 27%. The three findings discussed above still hold when the distribution of MWTP takes a log-normal distribution (Table 2, Panel C).

3.4 Conclusion

Our analysis shows that, if not accounted for, unobserved consumer heterogeneity can significantly affect the estimated MWTP for discounted future fuel costs. We believe that this may partly explain consumer undervaluation of future fuel costs and the wide range of estimates found in the literature. Here we have modeled consumer heterogeneity through the valuation of fuel economy. However, this is only one of many potential ways of representing consumer heterogeneity. For example, the heterogeneity could also arise from expected future fuel costs. While ignoring this source of heterogeneity would create a similar bias as the one identified in this paper, the implications for policy (whether or not there are consumer mistakes that constitute market failure) may be different. To properly evaluate the existence, source and magnitude of the energy paradox, further econometric analysis that explicitly models consumer heterogeneity from multiple sources by using random coefficient models in either a discrete choice or hedonic framework (e.g., (Berry et al., 1995); (Bajari & Benkard, 2005)), are needed.

Figure 3.1: Fuel Cost and Average Marginal Willingness to Pay among Buyers

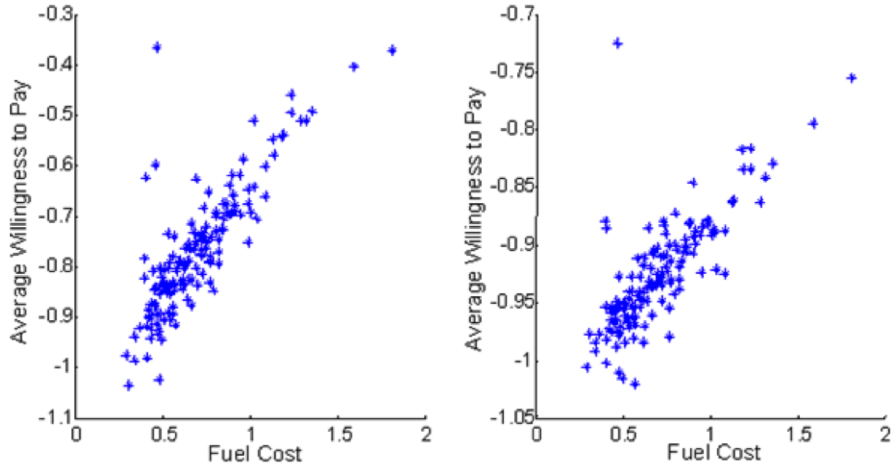

Table 3.1: Monte Carlo Results

	TRUE	Estimates			
Panel A: baseline model		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Constant	1	0.6	0.05	1.05	0.07
Price	2	2.02	0.01	2.00	0.01
Fuel cost	2	1.43	0.03	2.01	0.07
Weight	4	4.49	0.15	3.83	0.17
Horsepower	8	7.68	0.14	8.18	0.15
Sigmaa	4			4.18	0.26
Log-likelihood		228,335		228,268	
Implied valuation for $1 drop in fuel cost		$0.71		$1.00	
Implied undervaluation		29%			
Panel B: smaller heterogeneity		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Onstant	1	0.93	0.05	1.08	0.06
Price	2	2.01	0.01	2.01	0.01
Fuel cost	2	1.82	0.03	2.03	0.06
Weight	4	3.99	0.15	3.8	0.16
Horsepower	8	8.05	0.14	8.21	0.14
Sigmaa	2			2.31	0.31
Log-likelihood		225,942		225,933	
Implied valuation for $1 drop in fuel cost		$0.90		$1.01	
Implied undervaluation		10%			

Table 3.2: Robustness Checks

	TRUE	Estimates			
Panel A: monopoly instead of oligopoly		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Constant	1	0.62	0.06	1.1	0.09
Price	2	2.02	0.02	2.01	0.02
Fuel cost	2	1.27	0.03	2.14	0.14
Weight	4	4.38	0.19	3.8	0.21
Horsepower	8	7.86	0.18	8.32	0.19
Sigma	4			4.55	0.42
Log-likelihood		158,480		158,437	
Implied valuation for $1 drop in fuel cost		$0.63		$1.07	
Implied undervaluation		37%			
Panel B: 50 vehicle models instead of 25		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Constant	1	0.62	0.03	1	0.05
Price	2	2.02	0.01	2.00	0.01
Fuel cost	2	1.48	0.03	1.99	0.06
Weight	4	4.18	0.12	3.99	0.13
Horsepower	8	7.75	0.11	8.01	0.11
Sigma	4			4.01	0.21
Log-likelihood		326,436		326,351	
Implied valuation for $1 drop in fuel cost		$0.73		$1.01	$0.99
Implied undervaluation		27%			

CHAPTER 4

VEHICLE LIFETIME TRENDS AND SCRAPPAGE BEHAVIOR IN THE U.S. USED CAR MARKET

This study examines scrappage patterns in passenger cars and light trucks in the United States between 1969 and 1999. The transportation sector accounts for 29% of all greenhouse gas emissions in the U.S., of which 79% are from passenger vehicles (EPA, 2012). The magnitude of these emissions has prompted regulators to implement gasoline taxes and increase the average fuel economy of new vehicles using the Corporate Average Fuel Economy (CAFE) standard. Because 95% of all vehicles on the road are used vehicles, these policies have the potential to have large consequences, intended or unintended, on emissions from used vehicles. To properly evaluate the potential changes to the used vehicle market, two parameters are needed: the average vehicle lifetime and an elasticity of scrappage with respect to vehicle price. With relatively few studies of the used vehicle market, the parameters often chosen ignore how they may change over time or are not well suited to model these types of policies. Our study fills this gap by estimating these parameters and examines which are stable over time and which are not. This helps to guide future analysis as our estimates suggest that ignoring some dynamics in the used car market may be more problematic than others. Our study estimates that vehicle lifetime has increased, nearly 13% from 1969 to 1987, and that the scrappage elasticity with respect to vehicle price is -0.83, which has remained relatively stable over time. While the lifetime of vehicles has been estimated elsewhere, its increase has generally been ignored in policy analysis. Our estimate of the scrappage elasticity with respect to vehicle price is also important, as the values of this parameter commonly used are an order of magnitude more elastic. These omissions can have implications for many studies of the used vehicle market.

For example, while economists have advocated for federal carbon or gasoline taxes to reduce emissions from automobiles, policy makers continue to rely upon the CAFE Standard, which imposes limits on the minimum average fuel economy of the cars and trucks sold by each firm.[1] While there is considerable variation in the estimated cost of the CAFE standard, (M. R. Jacobsen, 2012b; Klier & Linn, 2012b; Anderson & Sallee, 2011a) it is relatively expensive compared with other options because it only regulates the new vehicle market and relies on time to increase the fuel economy of the used vehicle market. The increase in vehicle lifetime we estimate suggests that changes to the new car market caused by CAFE will take longer to affect the used car market than expected. General equilibrium price adjustments also play a role in the comparison of these policies. The CAFE standard decreases the supply of inefficient vehicles increasing their value and unintentionally lengthening their lifetime (Gruenspecht, 1982). Gasoline taxes, by contrast, disproportionally increase the operating cost of inefficient vehicles, reducing their value and shortening their lifespan (Bento et al., 2009). Comparing these policies requires a scrappage elasticity with respect to vehicle price and with no direct estimates available in the literature, these studies generally adopt values between -3 and -6.[2] Our substantially more inelastic estimate of this parameter suggests a diminished role

[1]For much of the past few decades the standard for cars has been 27.5 mpg while it has been 20.7 for trucks. The standard must be met separately for the foreign and domestic fleets produced by a manufacturer. Recent changes to the CAFE standard allow for more flexible targets based on the footprint of the vehicle and allow for some trading between firms. Problematically CAFE standards can reduce the scrappage of inefficient used vehicles as substitution from the more tightly regulated new car market increases their value (Gruenspecht, 1982). This may be a particular concern because the recently tightened CAFE standard will likely increase the price of inefficient, high performance new vehicles (Klier & Linn, 2012b).

[2]This value is derived from (Alberini, Harrington, & McConnell, 1998), which finds that a $1000 bounty (equivalent to 67% of the average vehicle value) causes a 193% increase in the number of vehicles scrapped. The scrappage elasticity with respect to vehicle price implied by these values is approximately -3 and -6 is adopted as a robustness test.

for these price effects, bringing these policies somewhat closer together in terms of efficiency than previous studies might suggest.

Vehicle lifetimes are also important for studies that examine how consumers value technology that reduces future gasoline costs (Helfand & Wolverton, 2011; D. Greene et al., n.d.). These studies often rely on estimates of vehicle lifetime that are many decades old (e.g. Lu, 2006). Our study allows us to place bounds on the error introduced to these calculations by using these outdated vehicle lifetimes. We find that omitting these lifetime increases may result in over or undervaluation of fuel-saving technology by nearly 7%. While sensitive to assumptions, we are also able to directly estimate this undervaluation through a novel channel. By estimating the response of scrappage to a change in gasoline price or an equivalent vehicle price change, our results are suggestive of undervaluation. We find that consumers recognize between \$0.53 and \$0.73 of a \$1 increase in operating cost.[3] These estimates lie within in the range of other studies of that use alternative methodologies (Allcott & Wozny, 2012).

Our study uses the most comprehensive publicly available data on passenger car and light truck populations by calendar year in the United States, covering vehicle model-years 1969 to 1999. The benefit of this aggregate data is that it covers a long period of time and allows for more time series variation. The data allows us to calculate mean scrappage rates at each age, which we fit to a logistic curve using nonlinear least squares. The resulting scrappage rate curve allows us to calculate the expected lifetime for a vehicle that faces those scrappage rates. We note that vehicle lifetime, particularly in the case of passenger cars, has continued to rise, likely due to technology changes or demographic

[3]The valuation suggested by a scrappage elasticity with respect to vehicle price in the range of -3, which is traditionally used for simulation, suggest a valuation of \$0.15.

shifts. Following Walker (1968) we then explain deviations from these mean scrappage rates due to a variety of cyclical factors including changes in vehicle prices, maintenance and repair costs, GDP, and steel prices. These regressions allow us to estimate a scrappage elasticity with respect to vehicle price. We also explore the relationship between gasoline prices and scrappage rates allowing us to compare our study with several others that measure the relationship between gasoline prices and scrappage in depth using more recent data (M. Jacobsen & van Benthem, 2013; Knittel & Sandler, 2010).

Our paper complements a large body of literature on scrappage behavior and the used car market. Several papers have examined how policies such as CAFE and gasoline taxes can influence vehicle prices and therefore, intentionally or unintentionally, affect the lifetime of used vehicles (Gruenspecht, 1982; Bento et al., 2009; M. R. Jacobsen, 2012b; Goulder et al., 2012).[4] Other studies have focused on the response to policies directly targeting used vehicles, including inspection and maintenance programs (Ando, McConnell, & Harrington, 2000), national vehicle retirement programs such as Cash-for-Clunkers (Miravete & Moral Rincón, 2009; Li, Timmins, & von Haefen, 2011; Schiraldi, 2011), and local scrappage subsidies targeted to a specific state or city (Alberini, Harrington, & McConnell, 1995; R. W. Hahn, 1995; Alberini et al., 1998; Adda & Cooper, 2000; Sandler, 2012).[5]

[4]Others have examined CAFE but omit scrappage or the used car market (Klier and Linn, 2012; Goldberg, 1998). One implication of our findings is that although important general equilibrium price effects remain in the used vehicle market, they are perhaps less important than would be suggested by alternative parameters.

[5](Ando et al., 2000) find I/M programs are limited in their ability to reduce emissions due to costs. (Miravete & Moral Rincón, 2009) find Cash-for-Clunkers programs can accelerate the adoption of new technology while Li, Linn and Spiller (2010) find it is an expensive method to improve overall fuel economy. (R. W. Hahn, 1995) estimates a scrappage curve from a local policy in California, while (Alberini et al., 1998) examine scrappage resulting from a program in Delware. (Sandler, 2012) finds adverse selection played a major role in the high initial scrappage rates for a scrappage policy in San Francisco.

Another series of recent papers have attempted to measure the marginal willingness-to-pay for technology that improves the fuel efficiency of vehicles. These studies suggest that consumer may undervalue this technology because they are more sensitive to vehicle prices than the future gasoline costs of operating that vehicle (Kilian & Sims, 2006; Allcott & Wozny, 2012; Busse, Knittel, & Zettelmeyer, 2013; Sallee, West, & Fan, 2011). Vehicle lifetime is an important component in calculating vehicle operating costs and existing studies often rely on estimates of vehicle lifetime that are for vehicles several decades old (e.g. (Lu, 2006)) which ignore the possibility that technology has extended vehicle lifetime.

Finally our paper contributes to a literature that examines the determinants of vehicle scrappage. These studies have focused on the role of technology (Walker, 1968; D. L. Greene & Chen, 1981), climate (B. W. Hamilton & Macauley, 1999) and gasoline prices (Li et al., 2011; M. Jacobsen & van Benthem, 2013; Knittel & Sandler, 2010). Several of these studies have noted the increase in vehicle lifetime but have not noted the broader implications of these changes. Those that have focused on gasoline prices have found divergent results for the effect of gasoline prices. While higher gasoline prices always compositionally shift the used vehicle fleet towards higher average fuel economy, the aggregate scrappage rate may increase or decrease.

Our paper contributes to these literatures by estimating vehicle lifetime and scrappage elasticities with respect to vehicle price. We also examine the scrappage elasticity with respect to gasoline price. Rather than focusing on the compositional effects of gasoline price changes, we examine the aggregate scrappage, or scale effect, that gasoline prices may have on scrappage. We then show the implications of these estimates for issues ranging from the value consumer place on technology that improves fuel efficiency; to the effect our elasticities have for important public policies like CAFE standards and gasoline taxes.

The rest of the paper is arranged as the following. Section 4.1 describes the data and the empirical strategy, section 4.2 presents results, and section 4.3 discusses the policy implications and section 4.4 concludes.

4.1 Basic Model and Data

4.1.1 Basic Model

The econometric model is based on earlier models of the automobile scrappage ultimately deriving from (Walker, 1968). The first stage of the model fits a logistic curve to mean scrappage rates at each age, which largely captures engineering scrappage arising from mechanical failure.[6] The second stage explains deviations from the mean scrappage function estimated in the first stage, generally resulting from cyclical factors such as changes in vehicle price or maintenance and repair costs.

The first stage uses nonlinear least squares to fit a logistic curve to the mean of scrappage rates by vehicle age. The logistic function, which has been shown to fit scrappage data well (Walker, 1968; Parks, 1977; D. L. Greene & Chen, 1981;

[6](Parks, 1977) describes engineering scrappage as the failure of vehicle components, which gradually become increasingly expensive as the vehicle ages. The rate at which these failures occur depend on the durability of the vehicle, which may be influenced by decisions made by the manufacturer or the environment where the vehicle drives.

Feeney & Cardebring, 1988), is given by:

$$M_a = \frac{1}{L + B * e^{(-k*a)}} + \epsilon_a \tag{4.1}$$

where is age of a vehicle in years since model year, M_a is the mean scrappage rate of vehicles at age a.[7] We estimate the parameters L, B, and k to capture the shape of the logistic function. L controls the level of the "asymptotic scrappage rate." If we take the limit of equation 4.1 as age approaches infinity, the function asymptotes towards the scrappage rate $1/L$.[8] The error term, ϵ_a, is assumed to be normally distributed.

This estimation allows us to calculate vehicle lifetime. To calculate the expected lifetime for vehicles, LT, we follow (D. L. Greene & Chen, 1981):

$$LT = \sum_a a * M_a \prod_{i=1}^{a-1} (1 - M_i) \tag{4.2}$$

where a is vehicle age and M_a is mean scrappage rate at age a conditional on surviving until then as in equation 4.1. $(1 - M_i)$ is the survival rate of vehicles aged , hence $M_a \prod_{i=1}^{a-1}(1 - M_i)$ gives the probability of a vehicle being scrapped at age a.

The second stage captures the deviations from the predicted average scrappage rate in a given calendar year:

$$S_t = \alpha_0 R_t^{\alpha} P_t^{\beta} K_t M_t^* \tag{4.3}$$

where,

$$M_t^* = \sum_{a=1}^{14} \frac{K_{at}}{K_t} M_a \tag{4.4}$$

[7]Precisely M_a is the mean proportion of vehicles surviving years that are scrapped, on the average, prior to their $(a - 1)$th birthday which are average across calendar years.

[8]B and k determine when the scrappage rate starts to increase rapidly and enter the exponential and mature phases. Ceteris Paribus, increasing B (or decreasing k) postpones when the exponential and mature phases occur.

Equation 4.3 defines the structural relationship of total scrappage to both engineering and cyclical factors. S_t is the number of vehicles scrapped in calendar year t across all ages. R_t is the turnover rate of vehicle ownership. This term is included because the decision to scrap a vehicle is made by used vehicle dealers and will be subject to the volume of trade-in vehicles each year.[9] P_t is the vehicle price ratio index (the vehicle price divided by the maintenance and repair costs), capturing the value of holding a vehicle, and K_t is the total number of vehicles in operation in calendar year . M_t^* is the predicted scrappage rate arising from engineering factors, for the total population, which is weighted using the age distribution of vehicles in year calendar t. It is calculated based on equation 4.4 and is a weighted average of the expected age-specific scrappage rates, M_a, estimated using equation 4.1. The weights used are the number of vehicles of age a in calendar year t, K_{at}, over the total number of vehicles in calendar year t.

To empirically estimate equation 4.3 we transform the equation using logarithms. This allows us to estimate this second-stage relationship using ordinary least squares (OLS). In our central specification, this equation becomes:

$$ln\frac{S_t}{K_t} - ln(M_t^*) = const + \alpha * ln(R_t) + \beta * ln(P_t) + \epsilon_t \tag{4.5}$$

The coefficient of interest is β, which measures the elasticity of scrappage with respect to vehicle price. This elasticity is important for simulation models of the used vehicle market where scrappage adjusts to changes in vehicle price. The first term on the left hand side, $ln\frac{S_t}{K_t}$, is the log of the observed scrappage rate in calendar year t, the second term, $ln(M_t^*)$, is the predicted scrappage rate from engineering factors related to aging.[10] The residual scrappage rate is the difference

[9] If a new vehicle entering the system pushes an old vehicle to be scrapped α will be 1.

[10] Because the estimates from the first stage only enter as dependent rather than independent variables, the uncertainty of those estimates will not affect the standard errors of the second stage parameters.

of these terms. It is explained by a constant, a proxy for turnover rate, and the vehicle price ratio index. Finally we assume the error term, ϵ_t, is normally distributed and newey-west standard errors are estimated in all regressions. See Appendix A for further discussion of the advantages of the 2-stage approach adopted here relative to other potential models.

We also use several econometric models to examine the relationship between scrappage and gasoline prices. The first is a linear specification similar to that used by (M. Jacobsen & van Benthem, 2013). This model does not average observations by vehicle age or calendar year as in the 2-stage specification above but rather controls for these factors with fixed effects as follows:

$$y_{amt} = \alpha_a D_a + \alpha_m D_m + \beta * gp_t + \epsilon_{amt} \tag{4.6}$$

In this equation y_{amy} is the scrappage rate of vehicles that are of age a and model year m scrapped in calendar year t, $D_i(i \in \{a, m\})$ is the vector of dummy variables for age and model year, gp_t is the gasoline price in year calendar t. The coefficient β provides the marginal effect of gasoline price on scrappage rate.

We also estimate the scrappage elasticity with respect to gasoline price by modifying Walker's 2nd stage equations as follows:

$$ln\frac{S_t}{K_t} - ln(M_t^*) = const + \alpha * ln(R_t) + \beta * ln(gp_t) + \epsilon_t \tag{4.7}$$

This equation is identical to equation 4.5 except $ln(gp_t)$, gasoline price in calendar year t, is substituted for the vehicle price ratio index.[11] It also allows for

[11] We use the gasoline price index to be consistent across the estimation of the two nonlinear specifications, although using real gasoline price does not significantly alter the estimates.

more direct comparison with the hazard model run by (Knittel & Sandler, 2010), which assumes nonlinearity in scrappage rates.

4.1.2 Data

The data used for our regressions primarily comes from publicly available counts of automobiles collected by R.L. Polk for Ward's Automotive. These data, covering model years 1969-1999, provide annual counts of U.S. passenger cars and light trucks by model year and are reported in Ward's Automotive Yearbooks (1981-2002). While covering a long time period, these data only distinguish between passenger cars and light trucks and do not provide model or class counts. For more detailed vehicle count data we supplement these data with household level data on vehicle holdings from National Household Transportation Surveys in 1995, 2001, and 2009.

Each year Ward's reports the number of vehicles in operation as of July 1st for 14 model years, allowing us to calculate scrappage rates for a model year at each age. For example in 1990 Ward's provides the count of passenger cars and light trucks in operation for model years 1976 through 1990. Population counts from the 1991 Ward's Yearbook allow for calculation of the number of vehicles that were scrapped in the interim. The scrappage rate is calculated as the number of vehicles removed from operation at age a divided by the number of vehicles of that model year in operation at the previous age, $a-1$. The long time span enables us to compare with previous studies (Walker, 1968; D. L. Greene & Chen, 1981) and establish how vehicle lifetimes and scrappage elasticity with respect to vehicle price have changed over time.

The censorship of population counts beyond 14 years of age limits our ability to observe the tail of the scrappage curve, but because it asymptotes at older ages, scrappage rates beyond 14 years can be inferred from the pattern established before this cut-off. Scrappage rates at younger ages can also be problematic. Occasionally vehicle counts increase in the first and second year of vehicle lifetime implying a negative scrappage rate. New vehicle models tend to enter the market ahead of the calendar year, and are often sold through the next calendar year; therefore, vehicles populations for the first two years are removed from our analysis allowing us to calculate scrappage rates for ages 2 through 14.[12] This results in 650 scrappage rate observations.[13] Some population data is given for model years 1969 through 2002 but only model years 1969-1987 give complete population counts for all 14 years of age. More recent model years do not have population counts at older ages. Vehicles from 1999 are the most recent with population counts at 2 years of age and are the last model year for which we use data in any of our regressions.

Table 1 shows average scrappage rates of cars and trucks at various ages. These are calculated for three sets of model years: 1969-1979, 1980-1987, and the full sample, 1969-1987 to examine decadal changes in vehicle lifetime. Previous studies (Walker, 1968; Parks, 1977; D. L. Greene & Chen, 1981; Feeney & Cardebring, 1988) have used the logistic curve to fit these scrappage rates because they grow slowly for the first several years, increase rapidly around 6 years of age and finally after 10 years begin to asymptote towards a high, but stable level.

[12]For example, a 2000 model year vehicle may appear in calendar year 1999 and may still be sold as new in 2001. Scrappage for extremely new vehicles are usually very low, due only to serious accidences. Therefore, the data we observe for 0 and 1-year-old vehicles often change mostly due to the sales of these inventories. Including these first two years does not significantly change the point estimates of our regressions but does increase the standard errors in some specifications.

[13]Model Years 1969-1987 each have 13 calendar years worth of model counts, 1988 has 12, 1989 has 11, until 1999 which is the last model year with an age of more than 2.

Because the scrappage rates asymptote towards levels near 7-20%, rather than 100%, explains the presence of some extremely old vehicles in the current fleet.[14] Table 1 also demonstrates that scrappage rates for trucks are consistently lower at any given age than they are for cars.

In recent decades dramatic compositional changes have occurred in the light truck. The category of light truck contains not only pickup trucks but also minivans, SUVs and CUVs and these vehicles have grown as a share of the light truck market. As we will document later, light truck lifetime appears to stagnate over time. We suspect this stagnation is related to the expansion of SUVs market share within the light truck category. To explore this hypothesis, we use data from the National Household Travel Survey (NHTS) to obtain counts of vehicles by the subcategories of car, SUV, and pickup truck. These surveys provide vehicle counts by class, model year, and age up to 25 years old. NHTS exists for five discrete calendar years, 1983, 1990, 1995, 2001 and 2009. Because the survey is not annual and sample sizes vary over time, we cannot calculate yearly scrappage rates from this data.[15] In section 3.1 we show how these data allow us to compare the order of scrappage rates for various vehicle classes with one another. To compare car, SUV and pickup truck scrappage rates, we use vehicle counts for these three classes for the NHTS survey years in which the category of SUV is recorded: 1995, 2001 and 2009.

We use a variety of data from other sources to construct key variables that af-

[14]The small but increasing number of extremely old vehicles on the road can also be noted from subsequent NHTS surveys. In the 1995 NHTS, for example, 1.6% vehicles are over 25 years old. This number grows to 3.1% in 2001 NHTS and 3.7% in the 2009 NHTS.

[15]Although we would ideally be able to infer information about national level scrappage rates from these samples, we found the survey population weights did not provide vehicle counts similar to the Wards data. For some years, the 1990 NHTS in particular, the implied vehicle populations can be quite different from Ward's report. Hence, we do not attempt to use these data to determine lifetime scrappage curves but only to compare vehicle classes to one another.

fect scrappage rates. To examine the scrappage elasticity with respect to vehicle price, we require data not only on used vehicle prices, but also on maintenance and repair costs. These variables will affect the reservation value for scrapping a vehicle. Studies that look to (Walker, 1968) create a vehicle price ratio index by dividing a used vehicle price index by a maintenance and repair cost index. This assumes that these variables will have equal but opposite effects on scrappage: as used vehicle prices increase or maintenance and repair costs decrease, consumers will scrap vehicles at a lower rate. Our main specification also imposes this restriction but we examine this assumption in our robustness checks by including each of the indexes individually. The used vehicle price index and the motor vehicle maintenance and repair cost index, gathered by the Bureau of Labor Statistics, are subcategories used in the calculation of the Consumer Price Index. Both indexes are seasonally adjusted with the base period of 1982 to 1984. In the construction of the used vehicle price index the BLS averages vehicle auction prices from National Automobile Dealers Association (NADA) and prices published by Kelly Blue Book.[16]

Figure 1 plots the logged vehicle price ratio index and the aggregate observed scrappage rate from 1970 to 2000. Consistent with (Gruenspecht, 1982), vehicle prices seem to decrease after the 1980's when CAFE standards would increase demand for relatively scarce inefficient used vehicles, while vehicle prices adjusted to higher levels.

As noted by (Walker, 1968) vehicle turnover rate may affect scrappage rates and will depend on many factors including credit availability, income, and assets. Following this literature we proxy for the rate of turnover with the ratio of new vehicle registrations to total vehicle ownership. The total number of new

[16]For more detail see (Pashigian, 2001).

vehicles is taken as the number of age 0 vehicles from Ward's Automotive Yearbooks. The total number of vehicles in operation for each calendar year is provided by Ward's Motor Vehicle Facts and Figures. Although new car purchases are made by consumers, while scrappage is decided on by used vehicle dealers, it may be possible that this proxy for turnover rate is endogenous. We therefore examine another proxy for turnover rate: annual GDP, taken from International Financial Statistics.

Two measures of gasoline prices are also used. First we use annual gasoline price data from the Department of Energy. We also use the gasoline price index gathered by the Bureau of Labor Statistics. The BLS seasonally adjusts this index and uses a base period of 1982 to 1984.

Finally, for robustness tests, we use the annual average U.S. steel scrap price per metric ton from the U.S. Geological Survey,[17] and U.S. imports vehicle sales data from Ward's yearbooks. The percentage of vehicles imported is constructed by dividing these values by the number of new vehicles sales. Further details, including descriptive statistics, can be found in Appendix Table A.1.

4.1.3 Time Series Properties of the Data

In the second stage of the regression, standard tests fail to reject the presence of a unit-root for scrappage rates as well as vehicle prices. With only 30 years of data, unit-root tests of the residuals are often marginal and sensitive to specification. A Dickey-Fuller test strongly suggests the residuals are stationary above

[17] 1969-1998 data are from: http://minerals.usgs.gov/minerals/pubs/metal_prices/, http://minerals.usgs.gov/minerals/pubs/commodity/iron_&_steel_scrap/360798.pdf. And 1999-2001 data are from http://minerals.usgs.gov/minerals/pubs/commodity/iron_&_steel_scrap/index.html#m

the 1% level, while the Elliot-Rothenberg-Stock test cannot reject a unit-root.[18] Autocorrelation plots are displayed in Figures 2 and 3. While the evidence is not decisive for stationary or nonstationary residuals, the cyclicality of the residuals in these figures can be the result of an AR(2) process (Harvey, 1981).

In our basic specifications, we view the model as a cointegrated model, as is implicitly assumed by prior work in this area. Although not traditionally modeled as an autoregressive process, we address the possibility that the residuals are non-stationary with AR(1) and AR(2) models, wherever possible, for both the vehicle price and gasoline price.

4.2 Results

4.2.1 Scrappage by Vehicle Age

Table 2 reports the results of estimation of equation 4.1, which fits a logistic curve to mean scrappage rates by age. The first panel of Table 2 shows the results for passenger cars across three periods as well as the comparison with model years 1966 through 1977 estimated by (D. L. Greene & Chen, 1981) and post-World War II models estimated by (Walker, 1968). Figure 4 displays these estimated scrappage rate curves for passenger cars for two sets of model years along with the post-war curve estimated by (Walker, 1968). Figure 4 and Table 2 show that the asymptotic scrappage rate has been declining over the last century and that scrappage rates at any given age have progressively decreased for more

[18]The estimated value of rho in AR-1 and AR-2 regressions is less than 0.4, which cannot be judged as statistically distinct from 1 with a sample size of 30 years worth of data.

recent vehicle model years. The asymptotic scrappage rate is 26.32% for post-war vehicles, it decreases to 21.51% for the 70s, and is 18.8% for the 80s. Table 2 also calculates the expected lifetime following equation 4.6. The lifetime for passenger cars has increased from 10 years for post-war cohort, to 12.5 for the 70s cohort, and 14 years for the 80s cohort.

The second panel of Table 2 shows the results for light trucks. We first confirm that light trucks display longer lifetimes than passenger cars. The scrappage profile is, however, complex. Rather than decreasing, asymptotic scrappage rates increase from 9.25% in the 70s to 20% in the 80s, yet overall vehicle lifetimes remain around 15- and 16-year-old.

Previous literature has noted the increased lifetime for passenger cars or the fleet as a whole but it has generally been attributed to events that are unlikely to explain why it continues to increase, or why it should be observed in passenger cars and not light trucks. Reasons that have been suggested in the past include a change in post-war technology (Walker, 1968), and an increase in the share of light trucks (D. L. Greene & Chen, 1981). (B. W. Hamilton & Macauley, 1999), who only examine passenger cars, suggest shifts in population to the Sunbelt are the source of this increased lifetime. This explanation seems plausible but should also increase the lifetime of light trucks. Another explanation is that on-board diagnostic systems and other technology improvements may also extend vehicle lifetime by catching small mechanical failures before they become more expensive (EPA, 2002), which would also explain why similar trends are noted in the vehicle fleets of other countries like Sweden (Feeney & Cardebring, 1988).

We speculate that light truck lifetime may also be increasing, either due to population shifts or technology improvement, but the observed stagnation in light truck lifetime is due to opposing trends in the light truck market. The first trend is for vehicles classified as light trucks to become more similar to cars over time. Vans, SUVs, and, recently, CUVs are all categorized as light trucks but may be more like cars, in terms of technology and patterns of use, than pickup tucks. The second trend is the increasing share of SUVs at the expense of pickup trucks. Under these circumstances, 'Simpson's paradox' may appear. This is a paradox in which a trend present in all subgroups, for example increasing vehicle lifetime, is reversed when the groups are combined[19]. Such reversals are possible when subgroups have differing base rates and the share of one subgroup is increasing. Thus if SUVs and pickup trucks are both increasing in vehicle lifetime but SUVs have shorter lifetimes overall, increasing the share of SUVs may undermine the overall increase in vehicle lifetime when the two are aggregated.

Without decades of disaggregated data, we cannot definitively prove that Simpson's paradox is responsible for this stagnation in vehicle lifetime, but we can show that the conditions that give rise to it are present. These two conditions are that the share of SUVs is increasing over time, and that SUVs have scrappage rates between that of passenger cars and pickup trucks, depressing the lifetime for the broader category of light trucks. The growth of SUVs as a share of the light trucks is well known and can be calculated from NHTS data. The market share of SUVs in the NHTS data has increased from 7% in 1995 to 18% in 2009.[20] To provide evidence for the second condition we attempt to show that the scrappage rate of SUVs for any given model year, lies between that of passenger cars and pickup trucks. For example between 1995 and 2001

[19]This paradox, named after Edward Simpson who first described it in 1951, occurs in many manifestations including gender discrimination (Bickel, Hammel, OConnell, et al., 1975) and smoking death rates (Appleton, French, & Vanderpump, 1996).

[20]See Appendix Table A.2 for descriptive statistics on each wave of NHTS survey.

the scrappage rate of 1992 model year passenger cars, SUVs and pickup trucks was 21%, 20% and 13%, respectively, while the same values for 1983 model year vehicles was 69%, 42% and 39%. Because the surveys are not from consecutive years, we cannot estimate the scrappage curves at each age but we can test the relative order of the scrappage rates. To perform this test, we normalize the scrappage rate of passenger cars to 1 and pickup trucks to 0, and rescale the SUV scrappage rate accordingly. The density of the rescaled SUV scrappage rates, smoothed with an Epanechnikov kernel using a bandwidth of 0.15, is plotted in Figure 5. The average is 0.25, which, as predicted, lies in between that of passenger cars and pickup trucks. This means that while SUV scrappage rates are most similar to pickup trucks, they are not identical and are lower than passenger cars. A one tailed t-test on these values suggests that SUV scrappage rates are statistically different from those of pickup trucks at the 10% level.[21] Although we do not have yearly scrappage rates of SUVs or pickup trucks over time to confirm the trend is due to Simpson's paradox, the patterns noted here are highly suggestive that the growth in the SUV share may be undermining the within class increase of trucks.

4.2.2 Elasticity of Scrappage with Respect to Used Vehicle Price

Table 3 column I reports the OLS estimates of equation 4.5 for passenger cars. For each regression the table reports the coefficients and standard errors. Our estimate of -0.83, using our basic specification in Column I, is statistically indistinguishable from that of Walker at -0.66 given in column IX.[22] These estimates

[21]Based on 28 observations, the one-tailed t-statistics is 1.64. Two outliers, in which cases the rescaled SUV scrappage rate has an absolute value greater than 3, are excluded.

[22]Turnover rate is higher although statistically indistinguishable from earlier estimates.

of the scrappage elasticity with respect to vehicle price are, however, far lower than in temporary and local programs (R. W. Hahn, 1995; Alberini et al., 1995; Alberini, Harrington, & McConnell, 1996; Alberini et al., 1998). These studies find elasticities between -1.7 and -3.[23] Because these programs are geographically limited and often short in duration, vehicle owners may change their scrappage decisions to take advantage of the program. The geographic and temporal limits of these policy likely result in a substantially larger scrappage elasticity with respect to vehicle price than would be expected by a permanent, national policy like CAFE standards or a gasoline tax. Further, as documented in (Sandler, 2012), these programs suffer from adverse selection, which may further overstate the scrappage elasticity with respect to vehicle price.

Table 3 columns II through IX examine the robustness of the scrappage elasticity with respect to vehicle price estimated in Column I. Following (Walker, 1968), the vehicle price ratio index regressor is the log of the ratio of the used vehicle price index and the maintenance and repair cost index, which assumes that coefficients of these two variables are equal in magnitude but opposite sign. We separately estimating these coefficients in Column II and find that the negative of the estimate on 'Maintenance and Repair Cost Index' at 0.74 is indistinguishable from that on the 'Used Vehicle Price Index' at -0.78 supporting this assumption. We also examine several other proxies for turnover rate including $ln(GDP)$ in column III finding this measure only further reduce the scrappage elasticity with respect to vehicle price to -0.61.

[23]Table 6 in (Alberini et al., 1995) suggests the average scrappage elasticity with respect to vehicle price is -1.7 for waivered vehicles and -2.56 for non-waivered vehicles. Table 4 in (Alberini et al., 1996) implies that the scrappage elasticity with respect to vehicle price is -1.8 at the mean vehicle value. (Alberini et al., 1998) Figure 3b. suggests scrappage rises from approximately 70 to 210 vehicles for a \$1000 bounty, which is 65% of the average vehicle value of \$1535.58 given in Appendix A 2.1, implying a scrappage elasticity of -3. Table 2 in (R. W. Hahn, 1995) implies an average scrappage elasticity of -1.75.

Column IV examines several other potentially important covariates, including steel price, GDP, and percentage imported vehicles. Steel price is included to capture changes in value of the scrapped vehicles. Because it is also possible that the entrance of foreign competitors may affect vehicle lifetime (B. W. Hamilton & Macauley, 1999), we also control for the percentage of the fleet that is imported. These additional covariates result in a point estimate of -0.76 for the scrappage elasticity with respect to vehicle price, a small and statistically insignificant decrease from the initial specification.

Columns V through VIII examine the robustness of our estimates to AR(1) and AR(2) models. The point estimates range from -0.71 to -0.87 and all are statistically indistinguishable from the point estimate in column I.

As demonstrated in Section 4.2.1 there have been substantial changes to trucks over the vehicles that comprise the light truck category. While consistent with our results for passenger cars, we leave these regressions to Appendix Table A.5.

4.2.3 Results from Other Specifications

Table 4 reports estimates of the marginal effect of gasoline prices on scrappage rates using equation 4.7. Controlling for only model year and age columns I through IV estimate, which are occasionally statistically different from zero, but in all cases suggest an inelastic response to changing gasoline prices. These point estimates imply scrappage elasticities that range from 0.07 to 0.31 and the 95% confidence interval rejects an elasticity grater that 0.51. Columns V and VI perform an unbalance panel regression using an AR(1) process following

(Baltagi & Wu, 1999). The marginal effect estimated in column VI is 1.506 and is statistically significant at the 5% level and suggests an elasticity of 0.28.

In Table 5 column I we estimate this elasticity using equation 4.7. This specification accounts for the nonlinearity of scrappage rates, and suggests that aggregate scrappage does not react very much to gasoline prices with a point estimate at -0.13. This specification is most similar to (Knittel & Sandler, 2010) who estimate that for a 5 cent per mile increase in operating cost, scrappage decreases by 12 percent suggesting a scrappage elasticity with respect to gasoline price of -0.21, which is statistically indistinguishable from our estimate.[24] This point estimate is the opposite sign of what we estimate using the linear specification but is not statistically different from zero. Columns II through IV examine the robustness of this estimate and generally suggest negative and inelastic response to gasoline price. The decrease in the point estimate when including GDP is particularly large suggesting that omitting income effects may bias this coefficient towards elastic values. Without these controls, gasoline prices will capture the substitution towards used vehicles that occurs during recessions. Columns V through VIII show the results from AR(1) and AR(2) models, which, although still imprecise, are statistically indistinguishable from the basic model in Column I and range from -0.04 to 0.14. Generally, the results from Tables 4 and 5 show that the estimated elasticities are particularly sensitive to the model used but are always statistically different from 1 suggesting that scrappage is inelastic with respect to gasoline prices.

[24]For this calculation we use values given in Table 1 of (Knittel & Sandler, 2010) for 2004. Without disaggregate data, we cannot replicate the estimation strategy of (Li et al., 2011), however these authors employ a logistic model similar to ours.

4.2.4 Comparison of the Elasticities of Scrappage with Respect to Vehicle and Gasoline Prices

According to economic theory, there should be a direct relationship between the scrappage elasticities with respect to vehicle and gasoline price. A change in gasoline price will change the operating cost of the vehicle. If consumers are rational gasoline price changes will be capitalized into the used vehicle price. This allows us to compare our estimated elasticities. Our scrappage elasticity with respect to gasoline price provides the scrappage response to a $0.10 increase in gasoline cost. This increase in gasoline cost will imply a decrease in used vehicle prices. The used vehicle price decrease or the gasoline price increase should imply the same scrappage increase if consumers fully capitalize operating cost changes into the value of the vehicle. Testing this valuation requires the use of a scrappage elasticity with respect to gasoline price that is positive. We adopt the value 0.24 from our simplest linear specification in Table 4 Column I. To calculate the expected change in operating cost due to this price change, we must make fairly strong assumptions about the fuel economy, vehicle price and age of the average used vehicles during the time of our study. Values estimated by Polk and the Bureau of Transportation Statistics put the mean fuel economy of used vehicles at 23.8 mpg, price at $8,786 and age at 11 years for the current fleet.[25] Using these numbers we find that consumers underreact to changes in gasoline price although we cannot reject full valuation. The estimated scrappage elasticity with respect to vehicle price of our basic model, -0.83, suggests that consumers recognize only $0.53 of a $1 increase in operating cost. Using

[25]The discount rate used was 5% and year VMT comes from (Lu, 2006). Mean fuel economy and price comes from BTS, mean age comes from Polk. Vehicle lifetime uses our estimates from 1980-1987.

the lower estimate of -0.61 from our specification in Table 3 column III suggest they recognize only $0.73 of a $1 increase.[26]

If we assume full valuation, the error must lie either in the average statistics gathered by Polk and BTS, or in the estimation of one of the two elasticities used in our calculation. If the error is in the estimation of the scrappage elasticity with respect to vehicle price the correct value must be -0.437, lower than our estimate and an order of magnitude lower than that from local scrappage programs.[27] If the error is in the scrappage elasticity with respect to gasoline price, the value must be 0.41, larger than our estimates using either specification, as well as the estimates implied by (Li et al., 2011), (M. Jacobsen & van Benthem, 2013) or (Knittel & Sandler, 2010). These calculations are, of course, sensitive to assumptions and the error may not be isolated to one element. They are, however, similar to other estimates of this undervaluation studied through alternative methods such as that of Allcott and Wozny (2012) who find consumers only recognize $0.76 of a $1.00 increase in operating cost.

4.2.5 Identification

The capitalization of gasoline price into vehicle price also raises potential concerns that other factors, particularly macroeconomic events, such as recessions, may affect the scrappage decision and gasoline prices or vehicle prices simul-

[26]Using our largest scrappage elasticity with respect to gasoline price and smallest scrappage elasticity with respect to vehicle price suggests a valuation of $1.13. Using the smallest, positive scrappage elasticity with respect to gasoline price and largest with respect to vehicle price gives a valuation of $0.15.

[27]This value is not, however, statistically different than our lowest possible estimate of this parameter in Table 4 Column IV. This specification, which substitutes the new vehicle price index as a proxy for turnover rate, is our most imprecise measure of the scrappage elasticity with respect to used vehicle price.

taneously. If this is the case, one may be concern that our estimates are biased. Ideally we would be able to instrument for vehicle or gasoline prices to address these concerns. Finding such an instrument is difficult for at least two reasons. First, very few valid instruments for gasoline price have been found in the literature and those that do use them use recent supply shocks that are temporary (e.g. (Hughes, Knittel, & Sperling, 2008)).[28] Our data spans an earlier time period and is annual, which prevents identification using short-term disruptions in gasoline price. Second there is evidence that consumers anticipate the return of gasoline prices to earlier levels during these particularly salient shocks (Anderson, Kellogg, et al., 2011). While consumers may make short-run adjustments to VMT in response to these fluctuations, they seem less likely to make major capital investments based on transitory shocks.

To some extent the relationship established between our estimated elasticities with respect to vehicle and gasoline prices helps to gain insight into how problematic these biases may be. If the previous implied scrappage elasticity with respect to vehicle price of -3 is adopted, consumers must only recognize $0.15 of a $1.00 increase in operating costs, significantly less than most estimates of this undervaluation. Some omitted factors such as macroeconomic shocks like recessions can be examined with the addition of other controls. Our robustness tests that include GDP[29] shown in Table 3 Columns III and IV, Table 4 Columns III and IV, and Table 5 Columns III and IV suggest omitting these factors does not dramatically affect these estimated elasticities. While disag-

[28] These authors use supply disruptions from Hurricane Katrina as an instrument for gasoline price. While such temporary price shocks may encourage drivers to temporarily decrease the miles they drive, they seem less likely to have the scrappage effect that a permanent increase of the same magnitude would.

[29] (J. Hamilton, 2009) documents that GDP and gasoline demand are positively correlated. As noted during the latest downturn, recessions decrease demand for gasoline lowering its price. Simultaneously, drivers tend to substitute away from new vehicles towards used vehicles decreasing scrappage rates.

gregate data may allow us to control for some of these effects, such data is not available for nearly as many years as our aggregate data. The longer duration of our data provides more variation in vehicle and gasoline prices over time but precludes us from using differential effects across fuel economy levels in identification.[30] Nevertheless our estimates are nearly identical to those that use more recent disaggregate data, for example (M. Jacobsen & van Benthem, 2013) who also estimate this elasticity at -0.7 to -0.8.

While there is little we can definitively do to address identification in these regressions, we argue that our results, which are robust to various models and additional covariates, represent a substantial improvement on the values that are traditionally used in policy simulation.

4.3 Further Discussion

Estimates presented above have implications for nearly all studies of the used vehicle market. Here we predominantly focus on the implications for two particularly active areas of research. First our parameters are central to evaluating policies, such as gasoline taxes or CAFE standards, aimed at reducing gasoline consumption and vehicle emissions. Simulation of these policies requires a value for vehicle lifetime and a scrappage elasticity with respect to vehicle price. These values help to predict scrappage changes due to general equilibrium price effects in the used car market. A second area of active research ex-

[30]This strategy uses the differential effect of gasoline price shocks on high versus low fuel economy vehicles to identify this elasticity. It requires that shocks to GDP, which are correlated with gasoline price changes, affect high and low fuel economy vehicles identically. This exclusion restriction may fail if, for example, low-income individuals sort into fuel-efficient vehicles and hold onto vehicles longer during recessions due to credit constraints.

amines the possibility that consumers may undervalue technology that reduces fuel use. Such undervaluation may guide the choice of optimal policy (Fischer et al., 2007). The incentive to improve fuel economy when gasoline taxes increase may be undermined if consumers do not value or understand those future fuel costs. But to calculate how much that technology reduces discounted future fuel costs requires a measure vehicle lifetime. Here we provide several back-of-the-envelope calculations attempting to show the importance of our estimates for these studies.

4.3.1 Implications for CAFE Standards

To predict the potential gasoline savings of CAFE, it is important to understand the speed at which aging removes old vehicles from the road. Problematically, these vehicle lifetimes may be based off of scrappage curves that are several decades old. In order to estimate vehicle lifetimes, one must observe the full scrappage curve of a specific model year until most of the vehicles have been scrapped, which may take several decades. The estimates of vehicle lifetime presented in section 3.1 suggest that vehicles several decades old may have considerably shorter lifetimes than those vehicles produced today. Longer vehicle lifetime could substantially impede the diffusion of new vehicles influenced by CAFE into the used vehicle fleet.

To illustrate how longer vehicle lifetimes may impede this diffusion we predict the fuel economy profile of the used vehicle market using two scrappage curves estimated 30 years apart.[31] This gives some indication of how much

[31] Specifically we use Walker's Postwar (1952-1957) estimates for the old scrappage curve and our estimates from 1980-1987 as the new scrappage curve.

discrepancy may occur between the predicted levels of fuel economy based on outdated scrappage profiles and the true scrappage profiles that recognize the technology change during the intervening years. To simplify this calculation we focus on the passenger car segment, which has a separate, higher standard than the light truck segment under CAFE. Historically CAFE has mandated that manufacturers achieve 27.5 mpg on average for passenger cars or pay fines based on the shortfall.[32] The CAFE standard is, however, scheduled to increase over the next decade. For our simulation, we generate an initial fleet that uniformly meets the 27.5-mpg standard and predict how quickly a new fleet produced at a uniform 40 mpg affects the used vehicle market.[33] Using a shorter vehicle lifetime will imply these changes in the new vehicle fleet will change the used vehicle fleet faster than when using a longer vehicle lifetime. The fuel economy of the average used vehicle over time is presented in Table 7 and plotted in Figure 6. The dashed blue line in Figure 6 projects the average fuel economy using the older scrappage curves, which imply shorter vehicle lifetimes, while the red solid line shows the outcomes under the newer scrappage curve, which imply longer vehicle lifetimes. The old curve suggests that the higher CAFE standard will affect the used vehicle market much faster than the new curve and is over optimistic about the speed at which CAFE can affect the fuel economy of the used vehicle fleet. Table 6 shows that some intermediate targets, like 35 mpg, can take a full four years longer to achieve using our new scrappage profiles.

A second back-of-the-envelope calculation shows this delay from another

[32]For passenger cars the old CAFE standard required a minimum average fuel economy of 27.5 mpg. This standard has been changed and started to increase in 2011. Additional changes to the standard allowed for a flexible target based on the footprint of the vehicle. These changes may affect magnitudes but not the qualitative conclusions of this calculation.

[33]We generate this population by projecting a fleet back in time assuming each year 10,000 vehicles are produced and are reduced by the estimated scrappage curve at each age.

perspective. Policy makers may want to know what CAFE standard is needed to increase the fuel economy of the whole fleet above a particular target within a set time frame. In Table 7 we calculate the CAFE standard required to achieve a variety of fleet targets within 10 or 15 years assuming that the fleet starts at an average of 27.5 mpg. For example using the old scrappage curves it will take a CAFE standard of 42.7 to achieve a fleet average of 40 within ten years. But the new scrappage rates suggest that the correct CAFE standard to achieve this goal is 50.7. Increasing the standard by 8 miles per gallon is likely to be expensive. Extrapolating the engineering cost curves for the subcompact category of car, a car that can achieve that target at lower cost than other vehicles (NRC 2002), suggest this 8-mpg improvement would cost at least $1,000 more per vehicle.[34] As the target becomes more aggressive, the discrepancy becomes larger. To achieve a fleet average of 50 mpg within 10 years, the older scrappage curves would only require a CAFE standard of 56.3 while the newer curve would require a standard in excess of 80 mpg. The final pair of columns shows that as the time frame is extended to 15 years, an age where most vehicles have been scrapped, this discrepancy decreases but is still surprisingly large. For example to achieve 40 mpg in 15 years the old curve would require the CAFE standard to be 40.5 while the new curve suggests the standard would need to be 42.9.

There are additional effects that may slow this response to higher CAFE standards further. Because CAFE reduces the supply of inefficient new vehicles, consumers may substitute towards inefficient used vehicles increasing their price (Gruenspecht, 1982). This will reduce their scrappage rate. Because our scrappage elasticity with respect to vehicle price is lower than previous esti-

[34] Alternatively the cost of this higher target can be calculated using the coefficients from the hedonic cost study by Berry, Kortum, and Pakes (1996) and converting to 2012 dollars. The higher target would cost $1,400 more per vehicle using these estimates rather than $1000 estimated from the engineering cost curves in the (Council, 2012) study.

mate, there will be less leakage due to this general equilibrium price effect than previous studies would suggest.

4.3.2 Implications for Gasoline Taxes

Our elasticities of scrappage with respect to vehicle price also have implications for the ability gasoline taxes to improve fuel economy. Equilibrium models of the automobile market have generally used scrappage elasticities with respect to vehicle price that range from -3 to -6, while our estimates suggest this parameter is inelastic at -0.83. The elastic values that are traditionally used are based off of studies that, although well executed, were never intended to estimate this parameter but have been used in the absence of another source. This elasticity is particularly important for gasoline taxes because a major effect of gasoline taxes in the used vehicle market is to increase the operating cost of inefficient vehicles and decrease their price. By preferentially scrapping inefficient used vehicles, the average fuel economy of this market rises faster than would occur with aging alone.

To illustrate this we show how this parameter choice can affect predicted policy outcomes in Table 8. (Bento et al., 2009) calculates the increased scrappage from a 25-cent gasoline tax using a price elasticity of -3 under three revenue-recycling methods. This value is based on a local scrappage policy studied in Alberini et al. (1998). Table 9 presents results for each revenue-recycling method using our lower price elasticity of -0.83. We find scrappage would have increased by only 0.10%, rather than the 0.35% simulated in the original study. This lower scrappage would imply 120,000 more vehicles on the road than predicted and, using average vehicle mileage and fuel economy, 48 million fewer gallons of gasoline saved in a year than predicted.[35]

This calculation illustrates that a low scrappage elasticity with respect to vehicle price reduces the ability of gasoline taxes to influence the used vehicle market and increases the role of scrappage due to aging. This implies that the CAFE standards and gasoline taxes are somewhat closer in terms of efficiency than previous literature suggests.[36]

4.3.3 Implications for Studies of the Energy Paradox

As discussed in section 3.4 above, economic theory suggests a rational consumer will pay \$1.00 more for a vehicle that reduces discounted future fuel costs by \$1.00. In many studies, including our own, consumers seem to undervalue these reductions in future fuel cost, a phenomenon often referred to as an 'energy paradox.' Empirically there has been considerable disagreement of the magnitude of this undervaluation (Helfand & Wolverton, 2011; D. Greene et al., n.d.) ranging from \$0.25 (Kilian & Sims, 2006) to \$0.76 (Allcott & Wozny, 2012) to full valuation of \$1.00 (Busse et al., 2013; Sallee et al., 2011). As shown above our estimates of the scrappage elasticity with respect to vehicle and gasoline prices, while unable to reject full valuation, suggest consumers only \$0.53 to \$0.73 of a \$1 reduction in future fuel costs. But our estimates of the increase

[35]This calculation only captures savings due to scale effects that reduce the total number of cars on the road and does not capture any compositional effects of scrappage. It also omits any equilibrium price effects that may occur.

[36]It is important to note that this does not imply that the two policies are identical in terms of gasoline savings. The CAFE standard increases the use of vehicles because improving the fuel economy of the fleet without increasing the price for driving results in more vehicle miles traveled, a phenomenon commonly referred to as the rebound effect (Small & Van Dender, 2007). The two policies also produce different results in the new vehicle market as CAFE implicitly taxes inefficient vehicles and subsidizes efficient ones (Kwoka, 1983).

in vehicle lifetime also have implications for this debate.

To properly evaluate the benefit of efficiency improving technology, both consumers and the researcher must specify how long a vehicle is likely to last. Longer vehicle lifetimes will increase the likelihood that consumers will realize the returns of a technology that improves fuel efficiency. Generally, researchers have applied scrappage rates to these calculations that are several decades old (Lu, 2006). To give some insight into the magnitude of error this may generate in these calculations we examine the value to the consumer of a technology that increases fuel economy from 20 to 30 mpg under two scrappage curves estimated 30 years apart.[37] When using the older scrappage curve with shorter vehicle lifetime, we find that the value of the technology is $5663.50 while using the newer curve with a longer vehicle lifetime provides a benefit of $6075.10. If the researcher uses a shorter vehicle lifetime to assess the value of this technology while consumers use longer vehicle lifetimes, the researcher will bias the results 7% towards overvaluation. Conversely consumers may form their expectations of vehicle lifetime off their last (shorter lifetime) vehicle, while the researcher uses newer longer vehicle lifetimes. Given the example above, this would result in 7% under valuation.

4.4 Conclusion

Despite the large size of the used vehicle market and its importance to policies such CAFE and gasoline taxes, relatively little attention has been paid to it. Our paper shows that the lifetime of passenger cars has continued to increase for

[37] For this calculation we assume a gasoline price of $3 and annual VMT according to (Lu, 2006).

much of the past century. We show that this lifetime increase is unlikely to be due to one-time events and suggests that technology improvement is the likely cause. Moreover we find that scrappage elasticities with respect to vehicle and gasoline prices are quite inelastic, particularly compared with the high elasticities estimated from temporary or local scrappage incentives.

Using a nonlinear specification we estimate that the average lifetime of passenger cars has increased from about 10 years in the 1950s to about 14 in the 1980s. We also estimate the scrappage elasticity with respect to vehicle price at -0.83, far less elastic than found in previous studies that estimate this parameter using local and temporary scrappage programs. Our estimates compared with earlier studies suggest that this elasticity has remained stable for most of the past century. Our findings have important policy implications. The increased vehicle lifetimes we estimate imply that the updated CAFE standards may take several years longer to affect the used vehicle fleet than otherwise predicted. We also show that our estimate of the scrappage elasticity with respect to vehicle price reduces the ability of gasoline taxes to remove used vehicle by 71%. Our estimates are also useful for examining the value consumers place on technology that reduces the future fuel cost of vehicles. Our point estimates of the scrappage elasticity with respect to vehicle and gasoline prices suggest that consumers may only recognize $0.53 to $0.73 for every $1 change in future gasoline costs. We also show that failing to account for the increase in vehicle lifetime when calculating the discounted future fuel costs for vehicles may result in over or underestimates of the energy paradox of 7%.

Our results also have implications for a variety of other programs. Longer vehicle lifetime suggests standards for local pollutants, which are placed on new vehicles only and therefore have relatively high costs (Small & Kazimi, 1995), will take longer to affect the entire fleet of vehicles. Inspections and Maintenance programs (I/M), which focus on pollution reductions for the oldest vehicles (Ando et al., 2000), may become more important for achieving emission reductions. Alternatively policy makers may seek to incentivize manufacturers to build vehicles with emissions reducing technology that lasts the life of the vehicle, which other authors have found to be particularly effective (Harrington, McConnell, & Ando, 2000). Our evidence of increasing vehicle lifetime and low scrappage elasticity with respect to vehicle price also suggest that it may be difficult to use policy to remove the large quantity of fuel inefficient vehicles that built up over the low gasoline prices of the past decades. If the past is any indicator many policies will be implement to reduce emissions from used vehicles and the parameters estimated here may help policy makers to accurately evaluate these programs in the coming years.

Figure 4.1: ln(Vehicle Price Index) and ln(Residual Scrappage Rates)

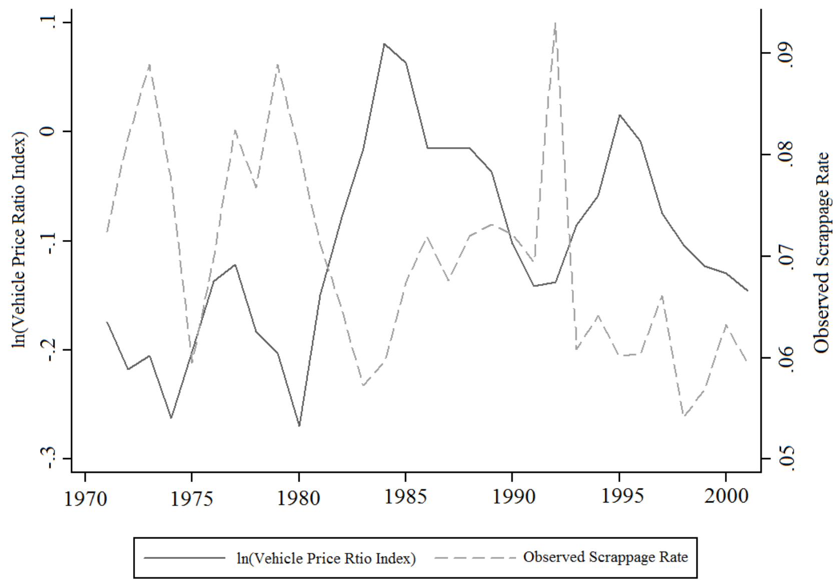

Figure 4.2: Autocorrelation Plot of Vehicle Price Regression Residuals

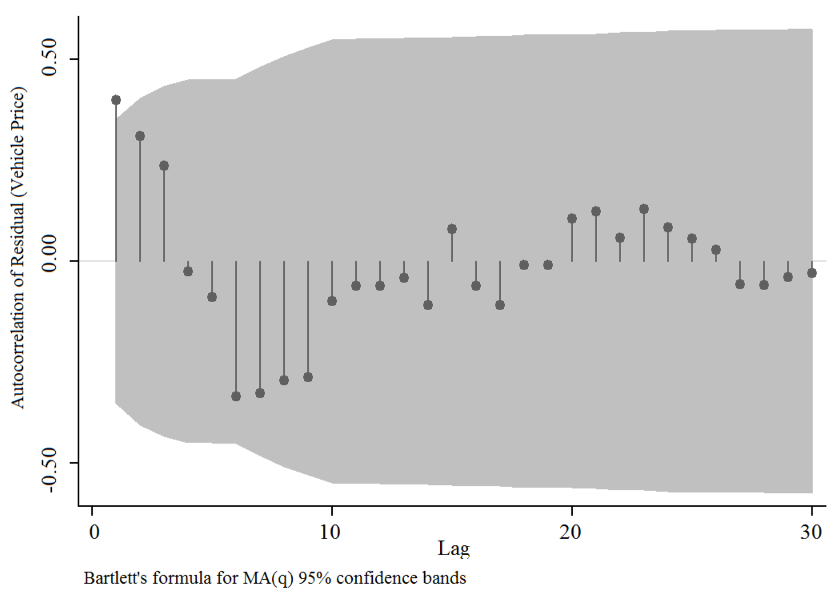

Figure 4.3: Autocorrelation Plot for Gasoline Price Regression Residuals

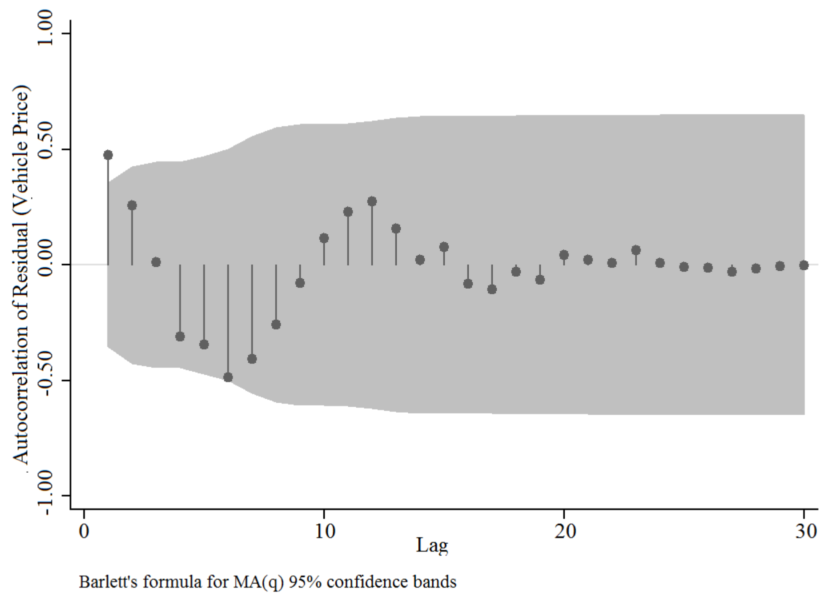

Figure 4.4: Scrappage Rate Curves

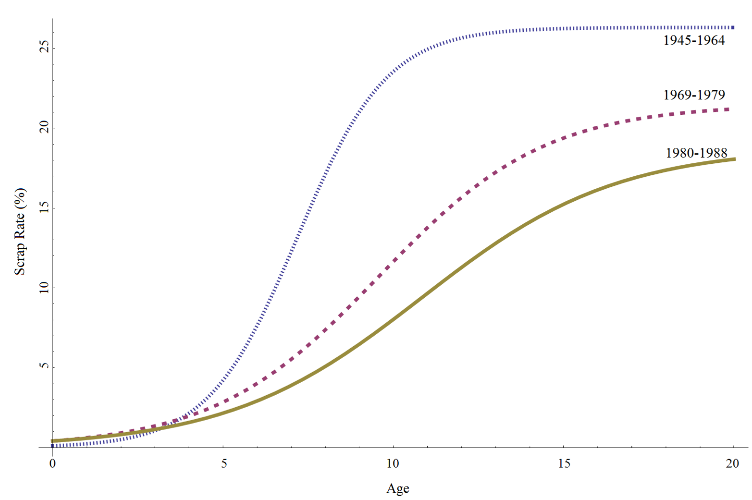

Figure 4.5: Rescaled SUV Scrappage Rates

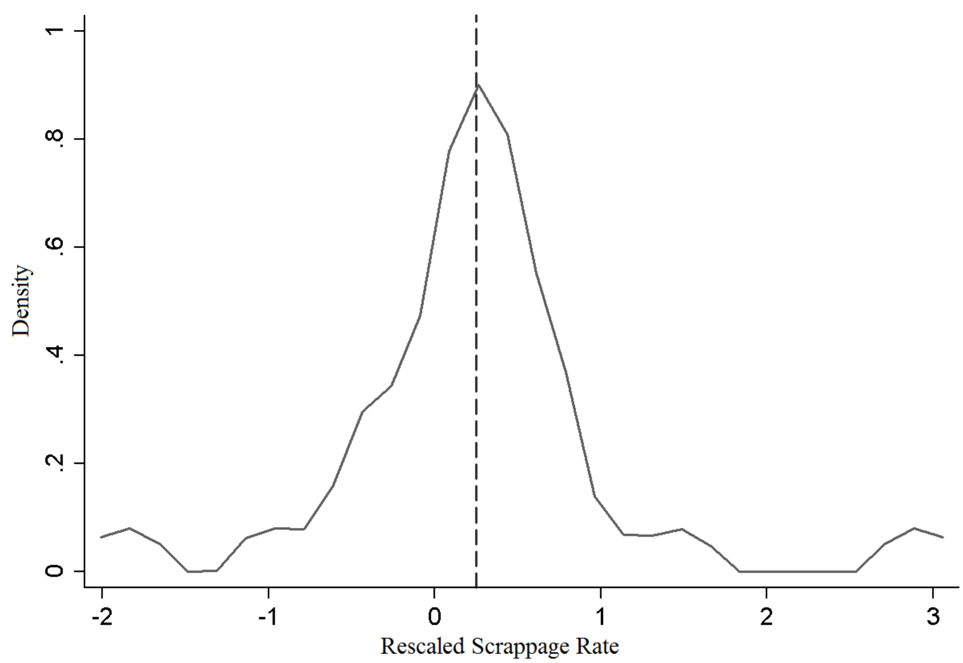

Table 4.1: Percent Scrappage Rate by Age for Passenger Cars and Light Trucks

Model Year	Passenger Cars			Light Trucks		
	1969-1979	1980-1987	1969-1987	1969-1979	1980-1987	1969-1987
Vehicle Age						
2	1.51%	1.09%	1.25%	0.78%	0.86%	1.33%
3	1.84%	1.47%	1.55%	1.32%	0.58%	1.46%
4	2.03%	1.12%	1.52%	1.27%	1.02%	1.58%
5	2.56%	2.03%	1.97%	1.61%	1.34%	1.81%
6	3.79%	3.09%	3.01%	2.20%	1.85%	2.47%
7	5.30%	3.80%	3.98%	2.99%	2.35%	2.85%
8	7.17%	5.16%	5.47%	3.75%	2.95%	3.58%
9	9.40%	6.58%	7.20%	3.77%	3.64%	3.88%
10	11.75%	8.22%	9.29%	5.85%	5.05%	5.57%
11	13.84%	9.66%	11.25%	5.33%	4.93%	5.35%
12	15.65%	11.41%	13.31%	6.57%	7.00%	6.93%
13	17.18%	12.77%	15.06%	7.08%	6.86%	7.09%
14	18.29%	14.26%	16.59%	7.53%	9.40%	8.31%

Table 4.2: Estimates of Logistic Parameters of Engineering Scrappage

	(69-79)	(80-87)	(69-87)	(66-77)1	Post-War2
Passenger Cars	I	II	III	IV	V
L	4.65	5.32	4.57	3.48	3.8
	(0.144)***	(0.26)***	(0.203)***	(0.0651)***	(0.0105)***
B	237.08	236.21	264.09	1661.04	890.48
	(30.41)***	(28.35)***	(32.12)***		(33.83)***
k	0.41	0.35	0.37	0.73	0.76
	(0.017)***	(0.016)***	(0.016)***	(0.0476)***	(0.0059)***
Obs3	13	13	13		
R-Squared	0.99	0.99	0.99		
Asymptotic Scrappage Rate	21.51%	18.80%	21.88%	28.74%	26.32%
Expected Lifetime4	12.47	14.08	13.45	10.9	10
Light Trucks					
L	10.81	5	6.78	7.59	
	(1.261)***	(2.508)*	(1.877)***	(0.4944)***	
B	229.55	236.38	135.82	251.72	
	(66.647)***	(74.774)**	(25.814)***		
k	0.32	0.26	0.23	0.34	
	(0.0455)***	(0.047)***	(0.034)***	(0.0261)***	
Obs3	13	13	13		
R2	0.99	0.99	0.99		
Asymptotic Scrappage Rate	9.25%	20.00%	14.75%	13.18%	
Expected Lifetime4	16.25	15.06	16.33	16.4	

Table 4.3: Elasticity of Scrappage with Respect to Vehicle Price Walker 2nd Stage

	Basic Model			
	I	II	III	IV
Ln(Price Ratio Index)	-0.83		-0.61	-0.76
	(0.25)***		(0.26)**	(0.31)**
Ln(Used Vehicle Price Index)		-0.78		
		(0.31)**		
Ln(Maintenance and Repair Cost Index)		0.74		
		(0.38)*		
Ln(GDP)			-0.27	-0.05
			(0.05)***	-0.11
Ln(Percent Imported)				0
				-0.04
Ln(Steel Price)				0.01
				-0.11
Ln(Turnover Rate)	0.96	0.88		0.84
	(0.19)***	(0.17)***		(0.22)***
Constant	2.83	2.82	2.7	2.88
	(0.46)***	(0.45)***	(0.43)***	(0.44)***
rho 1	-	-	-	-
rho 2	-	-	-	-
Obs	31	31	31	31

Table 4.4: Elasticity of Scrappage with Respect to Vehicle Price Walker 2nd Stage

	AR(1)		AR(2)		Walker
	V	VI	VII	VIII	IX
Ln(Price Ratio Index)	-0.77	-0.78	-0.85	-0.87	-0.66
	(0.31)**	(0.35)**	(0.22)***	(0.32)**	(0.170)***
Ln(GDP)		0		-0.01	
		-0.11		-0.1	
Ln(Percent Imported)		0.02		-0.04	
		-0.14		-0.14	
Ln(Steel Price)		-0.01		-0.07	
		-0.12		-0.1	
Ln(Turnover Rate)	0.98	0.99	1.01	0.98	0.663
	(0.16)***	(0.27)***	(0.15)***	(0.17)***	(0.156)***
Constant	2.88	2.94	2.96	3.04	0.624
	(0.40)***	(0.65)***	(0.38)***	(0.71)***	(0.142)***
rho 1	0.41	0.42	0.34	0.35	-
rho 2	-	-	0.2	0.26	-
Obs	31		31	31	

Table 4.5: Robustness Test for Gasoline Price Elasticity Linear Specification

	Basic Model				AR(1)	
	I	II	III	IV	V	VI
Gasoline Price	1.29		0.69	0.38	1.27	1.506
	-0.912		-0.522	-1.131	-0.67	(0.608)*
Gas Price Index		0.024				
		(0.011)*				
GDP			-0.003			-0.004
			(0.0000)***			(0.001)***
% of Import			0.029			0.028
			(0.008)***			(0.006)***
Steel Price			7.067			7.62
			(2.336)**			(2.594)**
Turnover Rate				-53.710*		
				-22.583		
Constant	13.653	12.235	54.835	17.889	16.642	36.884
	(1.721)***	(1.800)***	(7.703)***	(2.851)***	(0.425)***	(2.381)***
Rho 1					0.64	0.44
Obs	312	312	312	312	282	282
Elasticity	0.24	0.31	0.13	0.07		

Table 4.6: Elasticity of Scrappage with Respect to Gasoline Price Walker 2nd Stage

	Basic Model			
	I	II	III	IV
Ln(Gas Price)	-0.13		-0.12	-0.08
	-0.11		-0.109	-0.144
Ln(Gas Price Index)		-0.12		
		-0.107		
Ln(GDP)			-0.25	-0.11
			(0.079)***	-0.087
Ln(% of Import)				0
				-0.102
Ln(Steel Price)				0.04
				-0.09
Ln(Turnover Rate)	0.89	0.88		0.69
	(0.304)***	(0.319)***		-0.421
Constant	2.76	3.25	2.58	2.94
	(0.752)***	(0.373)***	(0.652)***	(0.786)***
rho 1	-	-	-	-
rho 2	-	-	-	-
Obs	31	31	31	31

Table 4.7: Elasticity of Scrappage with Respect to Gasoline Price Walker 2nd Stage

	AR(1)		AR(2)	
	V	VI	VII	VIII
Ln(Gas Price)	0.07	0.16	-0.04	0.14
	-0.155	-0.208	-0.11	-0.165
Ln(Gas Price Index)				
Ln(GDP)		-0.07		-0.1
		-0.131		-0.095
Ln(% of Import)		-0.03		-0.05
		-0.202		-0.122
Ln(Steel Price)		-0.06		-0.08
		-0.142		-0.947
Ln(Turnover Rate)	1.13	1.07	1.05	1.01
	(0.254)***	(0.309)***	(0.243)***	(0.213)***
Constant	3.33	4.04	3.16	4.18
	(0.637)***	(1.111)***	(0.599)***	(1.054)***
rho 1	0.57	0.59	0.48	0.51
rho 2	-	-	0.05	0.103
Obs	31	31	31	31

Table 4.8: Average MPG of Fleet after CAFE Increase

Year	Short Lifetime	Long Lifetime
1	28.39	28.15
2	29.33	28.8
3	30.33	29.48
4	31.38	30.17
5	32.47	30.88
6	33.58	31.61
7	34.67	32.34
8	35.7	33.07
9	36.61	33.8
10	37.38	34.51
11	38	35.19
12	38.49	35.83
13	38.87	36.42
14	39.16	36.96
15	39.37	37.44
16	39.54	37.86
17	39.66	38.22
18	39.75	38.53
19	39.81	38.79
20	39.86	39

Table 4.9: The Required CAFE Standard to Achieve Various Fleet MPG Targets

MPG Target	10 years		15 years	
	Short Lifetime	Long Lifetime	Short Lifetime	Long Lifetime
30	30.4	31.4	30.1	30.5
31	31.6	33	31.2	31.6
32	32.8	34.7	32.2	32.8
33	34	36.4	33.2	34
34	35.2	38.2	34.3	35.3
35	36.4	40.1	35.3	36.5
36	37.6	42	36.4	37.7
37	38.9	44.1	37.4	39
38	40.1	46.2	38.4	40.3
39	41.4	48.4	39.5	41.6
40	42.7	50.7	40.5	42.9
41	44	53.1	41.6	44.2
42	45.3	55.5	42.7	45.5
43	46.6	58.2	43.7	46.9
44	47.9	60.9	44.8	48.3
45	49.3	63.7	45.8	49.7
46	50.7	66.7	46.9	51.1
47	52	69.9	48	52.5
48	53.4	73.2	49	53.9
49	54.8	76.7	50.1	55.4
50	56.3	80.4	51.2	56.8

Table 4.10: Impact on Used Car Retention Simulating a Gasoline Tax

	Bento et al. 2009 Change in Used Car Ownership	Implied Percent Change in Price with Elasticity of -3	Change in Used Car Ownership with Elasticity of -0.83
Flat Recycling of $0.25 gas tax	-0.35%	0.12%	-0.10%
			[-0.13, -0.07]
Income-based recycling of $0.25 gas tax	-0.37%	0.12%	-0.10%
			[-0.13, -0.07]
VMT-based recycling of $0.25 gas tax	-0.39%	0.13%	-0.11%
			[-0.14, -0.08]

APPENDIX A

EQUIVALENCE OF STANDARDS AND FEEBATES

A.0.1 Initial Equilibrium

Demand

Consider an economy in which K identical firms produce J different types of vehicles for a large population of consumers of size N. Each consumer is endowed with income I_o and owns an equal share of the profits of each firm. Consumers get utility from vehicles and the consumption of a numeraire good x.

Consumers are divided into J groups depending on their preferred vehicle type. Let N_j denote the size of the group of consumers who prefer type j vehicles and that type j consumers only consume type j vehicles, getting zero utility from any other vehicle. To simplify the analysis we assume that utility is linear in the numeraire good. Each consumer is characterized by the pair (w, j), where j is the preferred vehicle type and w is the willingness to pay for that vehicle and is distributed uniformly $w \sim U[\underline{w_j}, \overline{w_j}]$. Utility maximization implies that a type j consumer will purchase a type j vehicle if his willingness to pay exceeds the price of that vehicle, p_j. Given the uniform distribution, the demand for vehicle j is

$$q_j(p_j) = N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} \tag{A.0.1}$$

Consumer surplus will be

$$CS = \sum_j N_j \frac{(\overline{w_j} - p_j)^2}{2\overline{w_j} - \underline{w_j}} + NI_o + \pi \tag{A.0.2}$$

Supply

I assume that each of the K firms, indexed by k, produces with the cost function $c(q_{k1}, \ldots, q_{kn}) = \sum_{j=1}^{n} c_j(q_{kj})$.[1] For each vehicle j, the cost function $c_j(q_j)$ is assumed to be increasing, smooth, and strictly convex and let $mc_j(q_j) = c'_j(q_j)$. The fuel economy is fixed for each vehicle. It can be expressed either as the mile per gallon rating mpg_j or its inverse g_j, the gallon per mile rating, such that $mpg_j = 1/g_j$. Firms operate in a competitive market. A representative firm optimizes over a vector of quantity $\boldsymbol{q_k} = \{q_{k1}, \ldots, q_{kn}\}$ of goods to sell at market-determined prices $\boldsymbol{p} = \{p_1, \ldots, p_J\}$. The firms' problem is formulated as

$$\max_{\boldsymbol{q_k}} \sum_{j=1}^{J} [p_j q_{kj} - c_j(q_{kj})] \tag{A.0.3}$$

This optimization implies that

$$p_j = mc_j(q_{kj}) \text{ or } q_{kj} = mc_j^{-1}(p_j) \tag{A.0.4}$$

Market supply of vehicle j is

$$q_j^S = Kq_{kj} \tag{A.0.5}$$

[1] In a setting with identical firms, tradability of permits is not required but will be if firms are heterogeneous.

Market Clearing

To clear the market for vehicle j, the quantity demanded equals the quantity supplied when

$$q_j^D = N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j) = q_j^S \tag{A.0.6}$$

Since mc_j is monotonically increasing there exists a unique p_j that sets supply equal to demand.

Consumer surplus will be

$$CS = \sum_j N_j \frac{(\overline{w_j} - p_j)^2}{2\overline{w_j} - \underline{w_j}} + NI_o + \pi \tag{A.0.7}$$

An allocation in this economy is a mapping $\{q(w, j), x(w, j)\}$ where:

$$q(w, j) = \begin{cases} 1 & \text{if } (w, j) \text{ owns a vehicle of type } j \\ 0 & \text{otherwise} \end{cases}$$

$$x(w, j) = \text{consumption of the numeraire}$$

The allocation $\{q(w, j), x(w, j)\}$ is feasible if:

$$NI_o = \sum_{j=1}^{J} \int_{\underline{w_j}}^{\overline{w_j}} x(w, j) \frac{dw}{\overline{w_j} - \underline{w_j}} + \sum_{j=1}^{J} c_j \left(\int_{\underline{w_j}}^{\overline{w_j}} q(w, j) \frac{dw}{\overline{w_j} - \underline{w_j}} \right) \tag{A.0.8}$$

A.0.2 Fuel Economy Standards

Take a given fuel economy standard $\overline{mpg}$, set by the government, such that $\overline{mpg}$ lies between the minimum and maximum fuel economy of vehicles produced.[2]

[2]This ensures that the standard can be achieved. To ensure the standard is binding, it must also be the case that $\overline{mpg}$ is more efficient than the unconstrained level of average fuel economy.

Let $\overline{g} = 1/\overline{mpg}$. The firm subject to this standard chooses a vector of quantities $\boldsymbol{q_k} = \{q_{k1}, \ldots, q_{kJ}\}$ to solve

$$\max_{\boldsymbol{q_k}} \sum_{j=1}^{J} [p_j q_{kj} - c_j(q_{kj})] \tag{A.0.9}$$

$$s.t.$$

$$\overline{mpg} \leq \frac{\sum_{j=1}^{J} q_{kj}}{\sum_{j=1}^{J} q_{kj} \frac{1}{mpg_j}}$$

The constraint, which represents the sales-weighted harmonic mean of the vehicles produced by the firm, can be rewritten as

$$0 \leq \sum_{j=1}^{J} (\overline{g} - g_j) q_{kj} \tag{A.0.10}$$

The first order conditions for product j are then

$$p_j = mc_j(q_{kj}) - \lambda(\overline{g} - g_j) \tag{A.0.11}$$

where λ is the shadow cost of the fuel economy constraint. This implies that firms choose quantity according to:

$$q^c_{kj} = mc_j^{-1}(p_j + \lambda(\overline{g} - g_j)) \tag{A.0.12}$$

The market clearing conditions, equation A.0.6, are now:

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = K mc_j^{-1}(p_j + \lambda(\overline{g} - g_j)) \tag{A.0.13}$$

The prices that solve this equation are identical for consumers and producers and denote that vector by $\boldsymbol{p^c} = \{p^c_1, \ldots, p^c_j, \ldots, p^c_J\}$. Denote the resulting allocation as $\{q^c(w, j), x^c(w, j)\}$

Profits in the economy, π^c, are:

$$\pi^c = K \sum (p^c_j q^c_{kj} - c_j(q^c_{kj})) \tag{A.0.14}$$

A.0.3 Feebate

Under a feebate with rate R and revenue neutral pivot point g_0 households face the tax/rebate inclusive price p_j. The fee $-R(g_0 - g_j)$ is levied on vehicle j, the firm faces the price $p_j + R(g_0 - g_j)$ and chooses $\boldsymbol{q_k} = \{q_{k1}, \ldots, q_{kJ}\}$ to solve

$$\max_{\boldsymbol{q_k}} \sum_{j=1}^{J} \{[p_j + R(g_0 - g_j)]q_{kj} - c_j(q_{kj})\} \tag{A.0.15}$$

which gives the first order condition

$$p_j = mc_j(q_{kj}) - R(g_0 - g_j) \text{ or } q_{kj} = mc_j^{-1}(p_j + R(g_0 - g_j)) \tag{A.0.16}$$

The market clearing conditions, equation A.0.6, are now:

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j + R(g_0 - g_j)) \tag{A.0.17}$$

Denote the consumer price vector that solves these equations as $\boldsymbol{p^f} = \{p_1^f, \ldots, p_j^f, \ldots, p_J^f\}$. Note that the producer prices, unlike in the case of the fuel economy standard, are not identical to consumer prices. Denote the resulting allocation as $\{q^f(w, j), x^f(w, j)\}$.

Total profits in the economy, π^f, are:

$$\pi^f = K \sum \{[p_j^f + R(g_0 - g_j)]q_{kj}^f - c_j(q_{kj}^f)\} \tag{A.0.18}$$

Also note that revenue neutrality implies that

$$0 = K \sum_{j=1}^{J} R(g_0 - g_j)q_{kj}^f \tag{A.0.19}$$

A.0.4 Equivalence

Equivalence of Consumer Prices

As noted in equation A.0.13, the price setting condition for the fuel economy standard is

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j + \lambda(\overline{g} - g_j))$$

while the price setting condition for the feebate, equation A.0.17, is

$$N_j \frac{\overline{w_j} - p_j}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j + R(g_0 - g_j))$$

To make the two policies equivalent we set $R = \lambda$ and $\overline{g} = g_0$. Because mc_j is single valued, it must be the case that prices faced by the consumer under each policy are identical. While $p_j^f = p_j^c$ this does not mean the price faced by the manufacturers is the same. Manufacturers face price p_j^c under the fuel economy standard and face price $p_j^f + R(g_0 - g_j)$ under the feebate.

Generating the Fuel Economy Standard Vehicle Allocation with a Feebate

Take a given vehicle allocation $\{q^c(w, j)\}$ achieved under a fuel economy standard $\overline{g}$ and resulting in shadow cost λ. This allocation can be achieved using a feebate with fee rate $R = \lambda$ and revenue neutral pivot point $g_0 = \overline{g}$. As shown above, this will imply that $p_j^f = p_j^c$. For product j the quantity produced with fee rate λ will be

$$q_j^S = Kmc_j^{-1}(p_j + \lambda(\overline{g} - g_j))$$

Where mc_j^{-1} is single valued, the quantity produced will be $q_j^S = Kq_{kj}^c$. Because there are identical consumer prices and quantity supplied under this feebate as in the fuel economy standard, the same households will purchase vehicles under each policy. The final requirement is that the feebate be revenue neutral. First note that the revenue generated by the government for product j produced by firm m is

$$R(g_0 - g_j)q^c_{kj}$$

Summed across all products and firms, this is

$$K\sum_{j=1}^{n} R(g_0 - g_j)q^c_{kj} = K\lambda\sum_{j=1}^{n}(g_0 - g_j)q^c_{kj} = 0$$

The first equality substituted in the values of R and g_0. The final uses the fact that q^c_{kj} solves the fuel economy standard constraint for the fuel economy level $\overline{g}$ for each firm, that is, equation A.0.10.

Generating the Feebate Vehicle Allocation with a Fuel Economy Standard

Take a given allocation $\{q^f(w, j)\}$ achieved under fee rate R with pivot point g_0. This allocation can be achieved using a fuel economy standard of $\overline{g} = g_0$. With this required average, the first order conditions of the firm will be

$$q^c_{kj} = mc_j^{-1}(p_j + \lambda(g_0 - g_j)).$$

It remains to be shown that under this constraint the firm must choose to produce q^f_{kj} with a shadow cost of $\lambda = R$. Intuitively, this proof shows that a multiplier that is too low will produce too many fuel inefficient cars and not meet the fuel economy standard, while one that is too high will over produce fuel efficient cars. It is important that mc_j^{-1} be monotonically increasing such that multiple values of the multiplier are not able to produce the same quantity.

Case 1

First assume that $\lambda < R$. For vehicles where $g_0 - g_j > 0$, their fuel economy is more efficient than the pivot point. First note that:

$$\lambda(g_0 - g_j) < R(g_0 - g_j) \tag{A.0.20}$$

For the market clearing condition, equation A.0.17, this implies that:

$$N_j \frac{\overline{w_j} - p_j^f}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j^f + R(g_0 - g_j)) > Kmc_j^{-1}(p_j^f + \lambda(g_0 - g_j))$$

To clear the market, the price must increase such that $p_j^c > p_j^f$. Because demand is downward sloping, this implies that quantity demanded under the fuel economy standard equilibrium will be less than that demanded under the feebate. Therefore $q_{kj}^c < q_{kj}^f$ and

$$q_{kj}^c(g_0 - g_j) < q_{kj}^f(g_0 - g_j) \tag{A.0.21}$$

For the inefficient vehicles where $g_0 - g_j < 0$:

$$\lambda(g_0 - g_j) > R(g_0 - g_j) \tag{A.0.22}$$

For the market clearing condition, equation A.0.17, this implies that:

$$N_j \frac{\overline{w_j} - p_j^f}{\overline{w_j} - \underline{w_j}} = Kmc_j^{-1}(p_j^f + R(g_0 - g_j)) < Kmc_j^{-1}(p_j^f + \lambda(g_0 - g_j))$$

Therefore the fuel economy standard prices must decrease and $p_j^c < p_j^f$. Downward sloping demand implies that $q_{kj}^c > q_{kj}^f$, that is, too many inefficient cars are produced. Multiplying each by the negative term $(g_0 - g_j)$ we have

$$q_{kj}^c(g_0 - g_j) < q_{kj}^f(g_0 - g_j) \tag{A.0.23}$$

Summing across all product using equations A.0.21 and A.0.23 we have

$$\sum_{j=1}^{J} Kq_{kj}^c(g_0 - g_j) < \sum_{j=1}^{J} Kq_{kj}^f(g_0 - g_j) \tag{A.0.24}$$

Because of the revenue neutrality of g_0 under fee rate R

$$0 = \sum_{j=1}^{J} R(g_0 - g_j)q_{kj}^{f} = \sum_{j=1}^{J} (g_0 - g_j)q_{kj}^{f} \tag{A.0.25}$$

Therefore allocation $\boldsymbol{q}^c$ with shadow cost $\lambda < R$ cannot satisfy the fuel economy standard g_0 as

$$\sum_{j=1}^{J} (g_0 - g_j)q_{kj}^{c} < 0 \tag{A.0.26}$$

Case 2

The second case, where $\lambda > R$ mirrors the first. Equation A.0.21 for efficient products, $g_0 - g_j > 0$, becomes

$$q_{kj}^{c}(g_0 - g_j) > q_{kj}^{f}(g_0 - g_j) \tag{A.0.27}$$

while equation A.0.23, for inefficient products, where $g_0 - g_j < 0$, becomes

$$q_{kj}^{c}(g_0 - g_j) > q_{kj}^{f}(g_0 - g_j) \tag{A.0.28}$$

Summing across all product using equations A.0.27 and A.0.28, and using the revenue neutrality of the allocation $\boldsymbol{q}^f$ under fee rate R and pivot point g_0,

$$\sum_{j=1}^{J} (g_0 - g_j)q_{kj}^{c} > \sum_{j=1}^{n} (g_0 - g_j)q_{kj}^{f} = 0 \tag{A.0.29}$$

This implies that the allocation $\boldsymbol{q}^c$ cannot satisfy the fuel economy standard g_0 with shadow cost $\lambda > R$.

Therefor, it must be the case that in order to meet fuel economy standard g_0 firms must do so at shadow cost $\lambda = R$ and produce the same supply of vehicles at the same prices. This means that the same households will buy vehicles resulting in the same allocation $\boldsymbol{q}^c = \boldsymbol{q}^f$.

Equivalence of Consumer Surplus and Numeraire Allocation

Consumer surplus in the economy, defined by equation A.0.7, is

$$\sum_j N_j \frac{(\overline{w_j} - p_j)^2}{2\overline{w_j} - \underline{w_j}} + NI_o + \pi \tag{A.0.30}$$

Because the prices faced by consumer are the same under both the fuel economy standard and the feebate, the only potential difference lies in the potential for a difference in firm profits, which are then distributed to consumers. Firm profits under a feebate, equation A.0.18, can be simplified as:

$$\begin{aligned}
\pi^f &= K \sum \{[p_j^f + R(g_0 - g_j)]q_{kj}^f - c_j(q_{kj}^f)\} && \text{(A.0.31)}\\
&= K \sum [p_j^f q_{kj}^f - c_j(q_{kj}^f)] + K \sum [R(g_0 - g_j)q_{kj}^f] && \text{(A.0.32)}\\
&= K \sum [p_j^f q_{kj}^f - c_j(q_{kj}^f)] && \text{(A.0.33)}\\
&= K \sum [p_j^c q_{kj}^c - c_j(q_{kj}^c)] && \text{(A.0.34)}\\
&= \pi^c && \text{(A.0.35)}
\end{aligned}$$

The simplification from A.0.32 to A.0.33 follows from the revenue neutrality, equation A.0.19. The final simplification makes use of the points proven above: that under either a fuel economy standard or a feebate, the quantity produced is the same, $Kq_{kj}^f = Kq_{kj}^c$, as well as the prices faced by the consumer, $p_j^f = p_j^c$. Therefore the consumer surplus is identical since $\pi^f = \pi^c$. Note that this does not imply that it is equal to the unconstrained level of consumer surplus.

Given the equivalence of p_j it follows that the individual vehicle choices of the agents will be identical. Because profits and therefore dividends are the same under each policy, the remaining budget dedicated to buying the numeraire good will be the same and will produce identical allocation mappings, $\{q^c(w, j), x^c(w, j)\} = \{q^f(w, j), x^f(w, j)\}$.

REFERENCES

Adar, Z., & Griffin, J. M. (1976, October). Uncertainty and the choice of pollution control instruments. *Journal of Environmental Economics and Management*, *3*(3), 178-188.

Adda, J., & Cooper, R. (2000). Balladurette and juppette: A discrete analysis of scrapping subsidies. *Journal of Political Economy*, *108*(4), 778–806.

Alberini, A., Harrington, W., & McConnell, V. (1995). Determinants of participation in accelerated vehicle-retirement programs. *The Rand Journal of Economics*, 93–112.

Alberini, A., Harrington, W., & McConnell, V. (1996). Estimating an emissions supply function from accelerated vehicle retirement programs. *The Review of Economics and Statistics*, 251–265.

Alberini, A., Harrington, W., & McConnell, V. (1998). *Fleet turnover and old car scrap policies* (Tech. Rep.). Resources for the Future.

Allcott, H., & Wozny, N. (2012). *Gasoline prices, fuel economy, and the energy paradox* (Tech. Rep.). National Bureau of Economic Research.

Anas, A., Arnott, R., & Small, K. A. (1998). Urban spatial structure. *Journal of Economic Literature*, *36*(3), pp. 1426-1464.

Anderson, S. T., Kellogg, R., Sallee, J. M., & Curtin, R. T. (2011). Forecasting gasoline prices using consumer surveys. *American Economic Review*, *101*(3), 110.

Anderson, S. T., Parry, I. W. H., Sallee, J. M., & Fischer, C. (2011, Winter). Automobile fuel economy standards: Impacts, efficiency, and alternatives. *Review of Environmental Economics and Policy*, *5*(1), 89-108.

Anderson, S. T., & Sallee, J. M. (2011a). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *The American Eco-*

nomic Review, *101*(4), 1375–1409.

Anderson, S. T., & Sallee, J. M. (2011b, June). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review*, *101*(4), 1375-1409.

Ando, A., McConnell, V., & Harrington, W. (2000). Costs, emissions reductions, and vehicle repair: Evidence from arizona. *Journal of the Air & Waste Management Association*, *50*(4), 509–521.

Appleton, D. R., French, J. M., & Vanderpump, M. P. (1996). Ignoring a covariate: An example of simpson's paradox. *The American Statistician*, *50*(4), 340–341.

Arnott, R., De Palma, A., & Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, *83*(1), 161–179.

Auffhammer, M., & Kellogg, R. (n.d., September). Clearing the air? the effects of gasoline content regulation on air quality. *American Economic Review*(6), 2687-2722.

Austin, D., & Dinan, T. (2005). Clearing the air: The costs and consequences of higher cafe standards and increased gasoline taxes. *Journal of Environmental Economics and Management*, *50*(3), 562 - 582.

Bajari, P., & Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, *113*(6).

Baltagi, B. H., & Wu, P. X. (1999). Unequally spaced panel data regressions with ar (1) disturbances. *Econometric Theory*, *15*(6), 814–823.

Baron, D. P. (1985a, Winter). Noncooperative regulation of a nonlocalized externality. *RAND Journal of Economics*, *16*(4), 553-568.

Baron, D. P. (1985b, November). Regulation of prices and pollution under incomplete information. *Journal of Public Economics*, *28*(2), 211-231.

Bento, A. M., Goulder, L. H., Jacobsen, M. R., & von Haefen, R. H. (2009, June). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review*, *99*(3), 667-99.

Bento, A. M., & Jacobsen, M. (n.d.). Ricardian rents, environmental policy and the 'double-dividend' hypothesis. *Journal of Environmental Economics and Management*(1), 17 - 31.

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

Berry, S., Kortum, S., & Pakes, A. (1996). Environmental change and hedonic cost functions for automobiles. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(23), pp. 12731-12738.

Berry, S., Levinsohn, J., & Pakes, A. (1995, July). Automobile prices in market equilibrium. *Econometrica*, *63*(4), 841-90.

Berry, S., Levinsohn, J., & Pakes, A. (2004, February). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, *112*(1), 68-105.

Bickel, P. J., Hammel, E. A., OConnell, J. W., et al. (1975). Sex bias in graduate admissions: Data from berkeley. *Science*, *187*(4175), 398–404.

Böhringer, C., Koschel, H., & Moslener, U. (2008, June). Efficiency losses from overlapping regulation of eu carbon emissions. *Journal of Regulatory Economics*, *33*(3), 299-317.

Böhringer, C., & Rosendahl, K. (2010). Green promotes the dirtiest: on the interaction between black and green quotas in energy markets. *Journal of Regulatory Economics*, *37*, 316-325.

Bovenberg, A. L., & de Mooij, R. A. (1994). Environmental levies and distortionary taxation. *The American Economic Review*, *84*(4), 1085–1089.

Bovenberg, A. L., & Goulder, L. H. (n.d.). Optimal environmental taxation in the presence of other taxes: General- equilibrium analyses. *The American Economic Review*(4), pp. 985-1000.

Browning, E. K. (n.d.). On the marginal welfare cost of taxation. *The American Economic Review*(1), pp. 11-23.

Burger, N. E., & Kaffine, D. T. (n.d., 2013/04/30). Gas prices, traffic, and freeway speeds in los angeles. *Review of Economics and Statistics*(3), 652–657.

Burrows, P. (1999, June). Combining regulation and legal liability for the control of external costs. *International Review of Law and Economics*, *19*(2), 227-244.

Busse, M. R., Knittel, C. R., & Zettelmeyer, F. (2013). Are consumers myopic? evidence from new and used car purchases. *The American Economic Review*, *103*(1), 220–256.

Chandra, A., Gulati, S., & Kandlikar, M. (n.d.). Green drivers or free riders? an analysis of tax rebates for hybrid vehicles. *Journal of Environmental Economics and Management*(2), 78 - 93.

Cheah, L. W. (2010). *Cars on a Diet: The Material and Energy Impacts of Passenger Vehicle Weight Reduction in the U.S.* Unpublished doctoral dissertation, Massachusetts Institute of Technology.

Council, N. R. (2012). *Effectiveness and impact of corporate average fuel economy (cafe) standards* (Tech. Rep.). National Academy Press.

Davis, L. W. (n.d.). The effect of driving restrictions on air quality in mexico city. *Journal of Political Economy*(1), pp. 38-81.

Davis, W. B., Levine, M. D., Train, K., & Duleep, K. G. (1995). *Effects of feebates on vehicle fuel economy, carbon dioxide emissions, and consumer surplus,* (Tech.

Rep. No. DOE/PO-0031). Department of Energy.

DeLeire, T., Khan, S., & Timmins, C. (2013). Roy model sorting and nonrandom selection in the valuation of a statistical life. *International Economic Review*, *54*(1), 279–306.

D'Haultfoeuille, X., Durrmeyer, I., & Février, P. (2012). *The willingness to pay for global warming reduction: Lessons from the french automobile market.*

Diamond, D. (n.d.). The impact of government incentives for hybrid-electric vehicles: Evidence from us states. *Energy Policy*(3), 972 - 983.

DiNardo, J., & Lee, D. S. (2004). Economic impacts of new unionization on private sector employers: 1984–2001. *The Quarterly Journal of Economics*, *119*(4), 1383–1441.

Downs, A. (1962). The law of peak-hour expressway congestion. *Traffic Quarterly*, *16*(3).

Dubin, J. A., & McFadden, D. L. (1984). An econometric analysis of residential electric appliance holdings and consumption. *Econometrica: Journal of the Econometric Society*, 345–362.

Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from us cities. *The American Economic Review*, *101*(6), 2616–2652.

Ep, A. (2007). *Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990 – 2005*. U.S. Environmental Protection Agency.

Feeney, B. P., & Cardebring, P. (1988). Car longevity in sweden: a revised estimate. *Transportation Research Part A: General*, *22*(6), 455–465.

Feng, Y., Fullerton, D., & Gan, L. (2005, August). *Vehicle choices, miles driven, and pollution policies* (Working Paper No. 11553). National Bureau of Economic Research.

Finkelshtain, I., & Kislev, Y. (1997, February). Prices versus quantities: The political perspective. *Journal of Political Economy*, *105*(1), 83-100.

Fischer, C., Harrington, W., & Parry, I. (2007). Do market failures justify tightening corporate average fuel economy (cafe) standards. *The Energy Journal*, *28*(4), 1–30.

Fischer, C., & Preonas, L. (2010). Combining Policies for Renewable Energy: Is the Whole Less Than the Sum of Its Parts? *International Review of Environmental and Resource Economics*, *4*, 51–92.

Gallagher, K. S., & Muehlegger, E. (2011). Giving green to get green? incentives and consumer adoption of hybrid vehicle technology. *Journal of Environmental Economics and Management*, *61*(1), 1–15.

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

Goldberg, P. K. (1995, July). Product differentiation and oligopoly in international markets: The case of the u.s. automobile industry. *Econometrica*, *63*(4), 891-951.

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

Goulder, L. H., Jacobsen, M. R., & van Benthem, A. A. (2012). Unintended consequences from nested state and federal regulations: The case of the pavley greenhouse-gas-per-mile limits. *Journal of Environmental Economics and Management*, *63*(2), 187 - 207.

Goulder, L. H., Parry, I. W., & Burtraw, D. (1997). Revenue-raising versus other approaches to environmental protection: The critical significance of preexisting tax distortions. *The RAND Journal of Economics*, 708–731.

Graham, D. J., & Glaister, S. (2002, January). The demand for automobile fuel: A survey of elasticities. *Journal of Transport Economics and Policy*, *36*(1), 1-25.

Gramlich, J. (2009). *Gas prices, fuel efficiency, and endogenous product choice in the u.s. automobile industry* (Tech. Rep.).

Greene, D., of Transportation, U. S. E. P. A. O., Assessment, A. Q., & Division, S. (n.d.). How consumers value fuel economy: A literature review.

Greene, D. L., & Chen, C. (1981). Scrappage and survival rates of passenger cars and light trucks in the us, 1966–1977. *Transportation research part A: general*, *15*(5), 383–389.

Greene, D. L., Patterson, P. D., Singh, M., & Li, J. (2005). Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy. *Energy Policy*, *33*(6), 757 - 775.

Gruenspecht, H. K. (1982). Differentiated regulation: The case of auto emissions standards. *The American Economic Review*, *72*(2), 328–331.

Hahn, J., Todd, P., & Klaauw, W. V. d. (n.d.). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*(1), pp. 201-209.

Hahn, R. W. (1995). An economic analysis of scrappage. *The RAND Journal of Economics*, 222–242.

Hamilton, B. W., & Macauley, M. K. (1999). Heredity or environment: Why is automobile longevity increasing? *The Journal of Industrial Economics*, *47*(3), 251–261.

Hamilton, J. (2009). Understanding crude oil prices. *Energy Journal*, *30*(2), 179–206.

Harberger, A. (1974). *Taxation and welfare*. Chicago, Ill.: University of Chicago Press.

Harrington, W., McConnell, V., & Ando, A. (2000). Are vehicle emission inspection programs living up to expectations? *Transportation Research Part D:*

Transport and Environment, *5*(3), 153–172.

Harvey, A. (1981). *The econometric analysis of time series, second ed.* Pearson Education.

Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics*, 33–54.

Helfand, G., & Wolverton, A. (2011). Evaluating the consumer response to fuel economy: A review of the literature. *International Review of Environmental and Resource Economics*, *5*, 103–146.

Heutel, G., & Muehlegger, E. J. (2010). *Consumer learning and hybrid vehicle adoption.*

Hoel, M., & Karp, L. (2001, October). Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics*, *82*(1), 91-114.

Hughes, J. E., Knittel, C. R., & Sperling, D. (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal*, *29*(1), 113–134.

Hymel, K. M., Small, K. A., & Dender, K. V. (n.d.). Induced demand and rebound effects in road transport. *Transportation Research Part B: Methodological*(10), 1220 - 1241.

Imbens, G., & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies*, *79*(3), 933–959.

Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, *142*(2), 615–635.

Jacobsen, M., & van Benthem, A. (2013). *Elasticity in the supply of used cars: estimating the scrap decision* (Tech. Rep.).

Jacobsen, M. R. (2012a). *Evaluating u.s. fuel economy standards in a model with* *producer and household heterogeneity.*

Jacobsen, M. R. (2012b). *Evaluating u.s. fuel economy standards in a model with producer and household heterogeneity.*

Jaffe, A. B., & Stavins, R. N. (1994). The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics*, *16*(2), 91–122.

Keeler, T. E., & Small, K. A. (1977). Optimal peak-load pricing, investment, and service levels on urban expressways. *The Journal of Political Economy*, 1–25.

Kilian, L., & Sims, E. R. (2006). *The effects of real gasoline prices on automobile demand: a structural analysis using micro data* (Tech. Rep.). Working Paper, University of Michigan.

Kleit, A. N. (2004). Impacts of long-range increases in the fuel economy (cafe) standard. *Economic Inquiry*, *42*(2), pp. 279-294.

Klier, T., & Linn, J. (2012a). New-vehicle characteristics and the cost of the corporate average fuel economy standard. *The RAND Journal of Economics*, *43*(1), 186–213.

Klier, T., & Linn, J. (2012b). New-vehicle characteristics and the cost of the corporate average fuel economy standard. *The RAND Journal of Economics*, *43*(1), 186–213.

Klier, T., & Linn, J. (2012c). *Using vehicle taxes to reduce carbon dioxide emissions rates of new cars: Evidence from france, germany, and sweden.*

Knittel, C. R. (2011, December). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review*, *101*(7), 3368-99.

Knittel, C. R., & Sandler, R. (2010). *Carbon prices and automobile greenhouse gas emissions: The extensive and intensive margins* (Tech. Rep.). National Bureau of Economic Research.

Kolstad, C. D., Ulen, T. S., & Johnson, G. V. (1990, September). Ex post liability for harm vs. ex ante safety regulation: Substitutes or complements? *American Economic Review*, *80*(4), 888-901.

Kwoka, J. E. (1983). The limits of market-oriented regulatory techniques: The case of automotive fuel economy. *The Quarterly Journal of Economics*, *98*(4), 695–704.

Langer, A. (2010). *Demographic preferences and price discrimination in new vehicle purchases.* (Tech. Rep.).

Larrick, R. P., & Soll, J. B. (2008). The mpg illusion. *Science*, *320*(5883), 1593-1594.

Lee, D. S., & Lemieux, T. (n.d., September). Regression discontinuity designs in economics. *Journal of Economic Literature*(2), 281-355.

Lee, L.-F. (1982). Specification error in multinomial logit models: Analysis of the omitted variable bias. *Journal of Econometrics*, *20*(2), 197–209.

Levinson, A. (2010, June). *Belts and suspenders: Interactions among climate policy regulations* (NBER Working Papers No. 16109). National Bureau of Economic Research, Inc.

Li, S., Linn, J., & Spiller, E. (n.d.). Evaluating "cash-for-clunkers": Program effects on auto sales and the environment. *Journal of Environmental Economics and Management*(2), 175 - 193.

Li, S., Timmins, C., & von Haefen, R. H. (2011). How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy*, *1*(2), 113–37.

Lu, S. (2006). *Vehicle survivability and travel mileage schedules* (Tech. Rep.).

Lyon, T. P., & Yin, H. (2010). Why do states adopt renewable portfolio standards?: An empirical investigation. *The Energy Journal*, *0*(Number 3), 133-158.

Matsudaira, J. D. (n.d.). Mandatory summer school and student achievement. *Journal of Econometrics*(2), 829 - 850. (¡ce:title¿The regression discontinuity design: Theory and applications¡/ce:title¿)

Miravete, E., & Moral Rincón, M. J. (2009). Qualitative effects of cash-for-clunkers programs.

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

Newell, R. G., & Pizer, W. A. (2003). Regulating stock externalities under uncertainty. *Journal of Environmental Economics and Management*, *45*(2, Supplement), 416 - 432.

Parks, R. W. (1977). Determinants of scrapping rates for postwar vintage automobiles. *Econometrica: Journal of the Econometric Society*, 1099–1115.

Parry, I. W. H., & Small, K. A. (n.d., September). Does britain or the united states have the right gasoline tax? *American Economic Review*(4), 1276-1289.

Parry, I. W. H., & Small, K. A. (2005, September). Does britain or the united states have the right gasoline tax? *American Economic Review*, *95*(4), 1276-1289.

Pashigian, P. (2001). *The used car price index: a checkup and suggested repairs* (Tech. Rep.). BLS Working Paper 338.

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

Pigou, A. (1932). *The economics of welfare*. Macmillan and Company, Limited.

Roberts, M. J., & Spence, M. (1976). Effluent charges and licenses under uncertainty. *Journal of Public Economics*, *5*(3-4), 193-208.

Ruderman, H., Levine, M. D., & McMahon, J. E. (1987). The behavior of the market for energy efficiency in residential appliances including heating and cooling equipment. *The Energy Journal*, *8*(1), 101–124.

Sallee, J. M. (n.d., September). The surprising incidence of tax credits for the toyota prius. *American Economic Journal: Economic Policy*(2), 189-219.

Sallee, J. M., & Slemrod, J. (2012). Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics*, *96*(11-12), 981 - 999.

Sallee, J. M., West, S. E., & Fan, W. (2011). The effect of gasoline prices on the demand for fuel economy in used vehicles: Empirical evidence and policy implications.

Sandler, R. (2012). Clunkers or junkers? adverse selection in a vehicle retirement program. *American Economic Journal: Economic Policy*, *4*(4), 253–281.

Schiraldi, P. (2011). Automobile replacement: a dynamic structural approach. *The RAND journal of economics*, *42*(2), 266–291.

Schrank, D., & Lomax, T. (n.d.). *The 2005 Urban Mobility Report* (Tech. Rep.). Texas Transportation Institute.

Shewmake, S., & Jarvis, L. (2011). Hybrid cars and hov lanes. *Available at SSRN 1826382*.

Small, K. A. (2012). Energy policies for passenger motor vehicles. *Transportation Research Part A: Policy and Practice*, *46*(6), 874 - 889.

Small, K. A., & Kazimi, C. (1995). On the costs of air pollution from motor vehicles. *Journal of Transport Economics and policy*, 7–32.

Small, K. A., & Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica*, *49*(1), pp. 105-130.

Small, K. A., & Van Dender, K. (2007). Fuel efficiency and motor vehicle travel: the declining rebound effect. *The Energy Journal*, 25–51.

Small, K. A., Winston, C., & Yan, J. (2005). Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica*, *73*(4), 1367–1382.

Spiller, E. (2010). *Household vehicle bundle choice and gasoline demand: A discrete-continuous approach.* (Tech. Rep.).

Stavins, R. N. (1996, March). Correlated uncertainty and policy instrument choice. *Journal of Environmental Economics and Management*, *30*(2), 218-232.

Train, K. E., & Winston, C. (2007, November). Vehicle choice behavior and the declining market share of u.s. automakers. *International Economic Review*, *48*(4), 1469-1496.

Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review*, *59*(2), 251–260.

Walker, F. V. (1968). Determinants of auto scrappage. *The Review of Economics and Statistics*, *50*(4), 503–506.

Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, *41*(4), pp. 477-491.

West, S. E., & Williams III, R. C. (2007). Optimal taxation and cross-price effects on labor supply: Estimates of the optimal gas tax. *Journal of public Economics*, *91*(3), 593–617.

Whitefoot, K. S., Fowlie, M., & Skerlos, S. J. (2012). *Product design response to policy intervention: Evaluating fuel economy standards using an engineering model of product design.*

Yatchew, A., & Griliches, Z. (1985). Specification error in probit models. *The Review of Economics and Statistics*, 134–139.

Yohe, G. W. (1978, June). Towards a general comparison of price controls and quantity controls under uncertainty. *Review of Economic Studies*, *45*(2), 229-38.