

ESSAYS ON TRANSPORTATION AND ENERGY EFFICIENCY POLICY

A Dissertation

by

JEREMY DAVID WEST

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Chair of Committee, Steven L. Puller
Committee Members, Li Gan
Yan Liu
Jonathan Meer
Head of Department, Timothy Gronberg

August 2014

Major Subject: Economics

Copyright 2014 Jeremy David West

ABSTRACT

This dissertation examines two questions of consumers' motor vehicle purchase and utilization. Both are related to policy variation induced by the U.S. Car Allowance Rebate System, better known as "Cash for Clunkers."

First, we directly investigate the impact of Cash for Clunkers, which was an economic stimulus program aimed at increasing new vehicle spending by subsidizing the replacement of older vehicles. Using a regression discontinuity design, we show the increase in sales during the two month program was completely offset during the following seven to nine months, consistent with previous research. However, we also find the program's fuel efficiency restrictions induced the purchase of more fuel efficient but less expensive vehicles, thereby reducing industry revenues by three billion dollars over the entire nine to eleven month period. This highlights the conflict between the stimulus and environmental objectives of the policy.

Second, we investigate a related topic. Due to the high political costs of raising the tax rate on gasoline, the United States government combats the negative externalities associated with gasoline consumption by regulating the fuel efficiency of new cars sold. However, the success of these Corporate Average Fuel Economy restrictions depends crucially on whether inducing households to drive more fuel efficient cars causes them to drive more miles, which would offset some or all of the reduction in gasoline consumption. We examine this question by applying a regression discontinuity design to exploit the increase in vehicle fuel efficiency induced among new car buyers in Texas during the Cash for Clunkers program in 2009. While new car buyers whose "clunker" was barely eligible for the subsidy drove a similar number of miles per year prior to the policy and are similar in other ways to barely ineligible new car

buyers, they bought significantly more fuel efficient vehicles. However, they did not respond by driving more miles following the program. As a result, the increased fuel economy reduced gasoline consumption proportionally. This suggests that behavioral responses do not undermine the effectiveness of fuel efficiency standards.

DEDICATION

I dedicate this work to my family and friends, who provided constant encouragement, to my wife, who patiently tolerated many long nights I stayed at the office, and to our dog, who kept her company.

ACKNOWLEDGMENTS

I am deeply grateful to the Lynde and Harry Bradley Foundation, the Private Enterprise Research Center, the S. Charles Maurice Endowment, the Efficient Vehicles and Sustainable Transportation Systems Center, and the Texas A&M University Economics Department for providing generous financial support towards this research.

This work would not have been possible without restricted-use data provided by the Texas A&M Transportation Institute, the Texas Commission on Environmental Quality, the Texas Department of Motor Vehicles, and the Texas Department of Transportation. I additionally thank DataOne Software for data services.

Finally, this work benefited extensively from discussions with the members of my committee, as well as from comments by Hunt Allcott, Antonio Bento, John Cochrane, Paul Ferraro, Ken Gillingham, Mark Hoekstra, Mark Jacobsen, Chris Knittel, Joanna Lahey, Arik Levinson, Shanjun Li, Jason Lindo, Joshua Linn, Erzo F.P. Luttmer, Gregor Pfeifer, Michael Price, Dave Rapson, Arthur van Bentham, Matthew Zaragoza-Watkins, Sarah Zubairy, and seminar participants at Amherst College/University of Massachusetts-Amherst, Berkeley's Camp EI, the University of Colorado, the Federal Trade Commission, Ghent University, University of Guelph, University of Wisconsin-Milwaukee, Resources for the Future, the Stanford Institute for Theoretical Economics, University of Tennessee, Texas A&M University, University of Virginia, the 2013 and 2014 Industrial Organization Conferences, and the 2013 Stata Texas Empirical Microeconomics Conference.

Any errors, inaccuracies, and opinions in this manuscript are of course my own.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
DEDICATION	iv
ACKNOWLEDGMENTS	v
TABLE OF CONTENTS	vi
LIST OF FIGURES	viii
LIST OF TABLES	ix
1. INTRODUCTION	1
2. CASH FOR COROLLAS: WHEN STIMULUS REDUCES SPENDING	3
2.1 Background and Empirical Strategy	7
2.1.1 The Cash for Clunkers Program	7
2.1.2 Empirical Strategy	9
2.2 Data	15
2.3 Results	17
2.3.1 Cash for Clunkers and the Timing of Purchase	17
2.3.2 Effect on New Vehicle Spending	18
2.3.3 Robustness and Threats to Identification	21
2.4 Interpretation and Discussion	25
2.5 Conclusions	26
3. HOW DO HOUSEHOLDS RESPOND TO INCREASES IN FUEL ECONOMY? REGRESSION DISCONTINUITY EVIDENCE	28
3.1 Background and Empirical Strategy	37
3.1.1 The Cash for Clunkers Program	37
3.1.2 Empirical Strategy	39
3.2 Data	43
3.3 Results	46
3.3.1 Tests of the Identification Strategy	46
3.3.2 Driving Outcomes	49

3.3.3 Household-Level Outcomes	51
3.4 Conclusions	52
4. SUMMARY AND CONCLUSIONS	55
REFERENCES	56
APPENDIX A. DATA APPENDIX	63
A.1 Defining a Household's Fleet	63
A.2 Calculating Household VMT	64
APPENDIX B. FIGURES AND TABLES	67
B.1 Figures and Tables for Section 2	67
B.2 Figures and Tables for Section 3	79

LIST OF FIGURES

FIGURE	Page
B.1 Monthly new vehicle sales, annual rate (source: NADA)	67
B.2 Relationship between fuel economy and vehicle price (MSRP)	68
B.3 First-stage: Probability of purchase being subsidized by CfC	69
B.4 Cumulative fraction of households purchasing a new vehicle by period	70
B.5 Purchases during July 2009 - August 2009 (Cash for Clunkers)	71
B.6 Purchases during July 2009 - April 2010 (10 months)	72
B.7 Identification checks: National Household Travel Survey (spring 2009)	73
B.8 Identification checks: Households purchasing vehicles prior to CfC . . .	74
B.9 Identification checks: Fleet characteristics for households purchasing during July 2009 - April 2010 (10 months)	75
B.10 Texas counties included in this study (EPA non-attainment)	79
B.11 Test of identification strategy: Household covariates reported in the National Household Travel Survey 2009	80
B.12 Test of identification strategy: Additional household covariates in the National Household Travel Survey 2009	81
B.13 Falsification test: Household-level results in 2008	82
B.14 Falsification test: Indicators of possible economic shocks in 2009	83
B.15 Selected characteristics of vehicles that were purchased new during CfC in 2009	84
B.16 Main results: 2010 outcomes for the new vehicles purchased	85
B.17 Main results: 2010 outcomes at the household-level	86

LIST OF TABLES

TABLE	Page
B.1 Summary statistics for new vehicle purchases July 2009 - April 2010 .	76
B.2 Discontinuities in vehicle characteristics for purchases during July 2009 - April 2010 (10 months)	77
B.3 Robustness of estimated discontinuities to alternate time windows . .	78
B.4 Summary statistics in 2010: Contrasting vehicle owners to purchasers	87
B.5 Test of identification strategy: Regression discontinuity estimates for households in NHTS-2009	88
B.6 Falsification exercise: Regression discontinuity estimates for household outcomes in 2008	89
B.7 Falsification exercise: Regression discontinuity estimates for indicators of possible economic shocks in 2009	90
B.8 Regression discontinuity estimates for household outcomes in 2010 . .	91

1. INTRODUCTION

The two major sections of this dissertation encompass two articles that I coauthored (Hoekstra, Puller, and West, 2014; West, Hoekstra, Meer, and Puller, 2014). In both articles we study consumer behavior in the transportation sector, particularly in personal motor vehicles. Both are related to policy variation induced by the U.S. Car Allowance Rebate System (CARS), better known as “Cash for Clunkers.”

Section 2 quantifies the economic stimulus from the U.S. Cash for Clunkers program, which was aimed at increasing new vehicle spending by subsidizing the replacement of older vehicles. Using administrative data on the population of household vehicle fleets in Texas, we estimate discontinuities for vehicle purchase timing and show that the increase in sales during the two month program was completely offset during the following seven to nine months, consistent with previous research (e.g. Mian and Sufi, 2012). However, the program successfully induced consumers to purchase alternate – higher fuel economy – vehicles relative to the counterfactual. These more fuel efficient vehicles were also less expensive vehicles, thereby reducing industry revenues by three billion dollars over the entire nine to eleven month period. This highlights the conflict between the stimulus and environmental objectives of the policy.

Section 3 evaluates the extent to which consumer utilization of fuel efficient vehicles (due to a lower price-per-mile of gasoline) undermines policies that promote energy efficiency. Due to the high political costs of raising the tax rate on gasoline, the United States government combats the negative externalities associated with gasoline consumption by regulating the fuel efficiency of new cars sold. However, the success of these Corporate Average Fuel Economy (CAFE) restrictions depends cru-

cially on whether inducing households to drive more fuel efficient cars causes them to drive more miles, which would offset some or all of the reduction in driving and gasoline consumption externalities. We examine this question by applying a regression discontinuity design to exploit the increase in vehicle fuel efficiency induced among new car buyers in Texas during the Cash for Clunkers program in 2009. While new car buyers whose “clunker” was barely eligible for the subsidy drove a similar number of miles per year prior to the policy and are similar to barely ineligible new car buyers, they bought significantly more fuel efficient vehicles. However, despite having a more fuel efficient vehicle fleet, the barely eligible households did not respond by driving more miles following the program. As a result, the barely eligible households reduced fuel consumption. We argue that if fuel economy standards such as CAFE are the relevant policy to evaluate, then our empirical strategy has decided advantages over other approaches used in the literature.

2. CASH FOR COROLLAS: WHEN STIMULUS REDUCES SPENDING

In efforts to boost economic activity via higher consumer and government spending, the U.S. government implemented several fiscal stimulus programs during the last two recessions. These policies typically operate either by reducing tax rates and providing tax rebates, as in the Economic Growth and Tax Relief Reconciliation Act of 2001, or by directly increasing government spending, as in the American Recovery and Reinvestment Act of 2009. The Car Allowance Rebate System (CARS), better known as “Cash for Clunkers,” differs from these other stimulus programs in that it aimed to increase consumer spending on a particular durable good – new vehicles – that had experienced a precipitous drop in sales during the 2009 recession.

A major objective of the program – and arguably the primary one – was to provide economic stimulus to U.S. vehicle and parts manufacturers (and therefore to the U.S. economy) by shifting expenditures “...from future periods when the economy is likely to be stronger, to the present...” [Romer and Carroll, 2010]. However, another priority for President Obama and the administration at that time was to improve the fuel efficiency of the U.S. vehicle fleet. Thus, the CARS policy was written to achieve multiple goals: the program attempted not just to accelerate the purchase of new vehicles to increase revenues to the auto industry, but also to induce those households to purchase more fuel efficient vehicles.

The fuel efficiency restrictions imposed by the program have potentially important implications for the stimulus effect. On the one hand, lowering the relative price of fuel efficient vehicles might induce buyers to increase spending by selecting vehicles with more expensive fuel-saving technologies, such as hybrids. On the other hand, the restrictions could induce households to purchase smaller, less expensive vehicles

in order to meet the fuel efficiency criteria, which would decrease overall new vehicle spending. The net impact of these restrictions on the stimulus effect of the program is an empirical question. The key contribution of this paper is to estimate not only how Cash for Clunkers impacted the timing of consumers' purchases, but also how the program affected total new vehicle spending.

The primary challenge to identifying the impact of any stimulus policy on spending is finding a valid counterfactual: what would have occurred in the absence of the policy. In the case of Cash for Clunkers, this requires determining both the timing and the type of vehicles that would have been purchased absent the policy. A major strength of our study is that we are able to apply a regression discontinuity design that uses the behavior of barely ineligible households as a counterfactual for barely eligible households. Specifically, we exploit the fact that households owning "clunker" vehicles rated at eighteen miles per gallon (MPG) or less were eligible for the program, whereas households with clunker vehicles rated nineteen MPG or higher were not. Although this strategy precludes examining the impact of the program on regional economic outcomes, given how the program was implemented it is difficult for us to think of a more compelling counterfactual.¹

We apply this regression discontinuity design to administrative data on all households in Texas. Intuitively, we compare the purchasing behavior of *all* households barely eligible for the program to that of all barely ineligible households. The identifying assumption of this approach is that all other determinants of purchasing behavior are continuous across the eligibility threshold. There is little reason to doubt this assumption: eligibility for the program was based on the EPA combined fuel economy rating and applied only to consumers who had owned their clunker for

¹We focus only on identifying the stimulus impact for the U.S. auto industry. Though we believe that this policy likely had important consequences for the broader U.S. economy, we do not attempt to quantify the impact of the program on overall economic growth.

at least one year. As a result, there was little scope for the type of manipulation that would invalidate the research design. In addition, we know of no other programs that affected households discontinuously at this cutoff. Thus, it is difficult to construct a plausible mechanism that would undermine the identifying assumption of our research design.

Using this method, we find that although Cash for Clunkers significantly increased new car purchases during the two months of the program, all of this increase represented a shift forward from the subsequent seven to nine months. Specifically, during the two months of the program, the frequency of purchasing a new vehicle was around fifty percent higher for the barely eligible households as compared to the barely ineligible ones, confirming that the program induced households to purchase new vehicles. By seven to nine months after Cash for Clunkers had ended, the barely eligible and barely ineligible households were equally likely to have purchased a new vehicle since the beginning of the program. On net, the program did not result in any more vehicle purchases than otherwise would have occurred over the nine to eleven month period that includes the two program months. This finding represents a slightly longer time to reversal than the six and seven month time horizons found by Li, Linn, and Spiller [2013] and Copeland and Kahn [2013], respectively, and is similar to that found by Mian and Sufi [2012]. As noted in Mian and Sufi [2012], this reversal occurred much more quickly than the five years assumed by the CEA or the three years assumed by the NHTSA (Council of Economic Advisers, 2009; National Highway Traffic Safety Administration, 2009).

However, as discussed earlier, the program's fuel efficiency restrictions could have shifted both the *type* and *price* of vehicles purchased, which would have important implications for the program's effect on auto industry revenues. The primary contribution of this paper is that it is to our knowledge the first to use quasi-experimental

methods to examine the impact of Cash for Clunkers on overall new vehicle spending.² To do so, we apply the same regression discontinuity design. Here, however, we focus only on new car buyers who purchased a vehicle either during the program or in the eight months that followed. This time horizon is constructed such that the probability of purchase is held constant across the cutoff, meaning that the only factor that affects overall spending is the amount spent conditional on purchase. This approach enables us to focus on new car buyers and avoid averaging across all Texas households, more than ninety-five percent of whom did not purchase a new vehicle within this ten month time horizon.³

Strikingly, we find that Cash for Clunkers actually *reduced* overall spending on new vehicles during the period beginning with the first month of the program and ending eight months after the program. Estimates indicate that each household purchasing under the program spent an average of \$4,600 less on a new vehicle than they otherwise would have. Thus, we estimate that this stimulus program – which dispensed three billion dollars in subsidies toward the purchase of 677,000 new vehicles nationally – actually reduced revenues to the auto industry by around three billion dollars over the course of less than one year. This highlights how – even over a relatively short period of time – a conflicting policy objective can cause a stimulus program to instead have a contractionary net effect on the targeted industry.

²Our study joins a broader literature examining the economic stimulus of policies such as tax rebates (e.g. Shapiro and Slemrod, 2003; Johnson, Parker, and Souleles, 2006; Agarwal, Liu, and Souleles, 2007; Parker, Souleles, Johnson, and McClelland, 2013), income tax reductions (House and Shapiro, 2008), and direct government spending on health, education, and infrastructure (Feyrer and Sacerdote, 2011; Wilson, 2012). The most closely related paper is Li and Wei [2013], who use a dynamic discrete choice model to examine the tradeoff between the environmental and stimulus components of the CARS program.

³Put differently, the potential concern with focusing on buyers, rather than all households, is that the policy could induce different types of households to purchase. While this could certainly be true if one were to focus on the two months of the program, it is no longer true over the longer time frame. That is because this longer time period is constructed such that there was no selection of households toward or away from buying a car based on the policy. As a result, the only impact was on the timing and type of purchase made, not the type of household who purchased.

2.1 Background and Empirical Strategy

2.1.1 *The Cash for Clunkers Program*

The Cash for Clunkers program, formally known as the Consumer Assistance to Recycle and Save (CARS) Act, was a nationwide vehicle scrappage program.⁴ Signed into law on June 24, 2009, the program incentivized households to replace used, fuel inefficient vehicles with new, fuel efficient vehicles. Specifically, the program offered consumers a rebate of \$3,500 or \$4,500 towards the purchase of a new fuel efficient car provided they scrapped a used vehicle. Transactions became eligible for rebates on July 1, 2009 and ended on August 24, 2009. Over the eight weeks of the program, Congress allocated a total of \$3 billion toward the subsidies. More than 677,000 vehicles were purchased under the program, 43,000 of which were in Texas.

As with most vehicle scrappage programs, the subsidy could only be used toward the purchase of a new vehicle; used vehicles did not qualify for the rebate. This requirement was driven by the major goal of the program: to accelerate the sale of new vehicles and provide fiscal stimulus to the auto industry and the broader economy. The program was largely motivated by the precipitous drop in vehicle sales during the 2008-2009 recession. This drop is depicted in Figure B.1, which shows that the seasonally-adjusted annualized number of sales fell from more than sixteen million in 2007 to around ten million in 2009.

However, the program also aimed to reduce the environmental costs imposed by the national vehicle fleet. It did this by placing restrictions on both the vehicle being traded in and the vehicle being purchased. The restriction on the trade-in vehicle

⁴Scrapage policies have been implemented around the world, and studies of these programs include Hahn [1995], Alberini, Harrington, and Virginia [1996], Adda and Cooper [2000], Miravete and Moral [2011], Sandler [2012], and Busse, Knittel, Silva-Risso, and Zettelmeyer [2012]. More generally, the literature has investigated the determinants of scrapage decisions, including the effect of gasoline prices and used car resale value (Li, Timmins, and von Haefen, 2009; Jacobsen and van Bentham, 2013).

is critical to our research design: the subsidy was only available to consumers who could trade in a vehicle rated by the EPA at a combined eighteen miles per gallon or less. This feature of the program enables us to use the purchasing behavior of barely ineligible households as a counterfactual for the barely eligible households.⁵ The program required that this traded-in vehicle be taken off the road and scrapped, meaning that the program attracted primarily older, low value vehicles.

If this restriction on the fuel efficiency rating of the trade-in vehicle were the only environmental component of the program, the theoretical impact of the program on new vehicle spending would be straightforward. The subsidy would lower the price of vehicles purchased during the program relative to those purchased in the future, which would accelerate the timing of sales. In addition, assuming that new vehicle characteristics such as vehicle size, performance, and interior amenities are normal goods, the income effect of the subsidy would result in purchases of somewhat more expensive vehicles.⁶ As a result, we would expect to see higher new vehicle spending during the program, and an increase in total revenues to the auto industry during the medium- to long-run.

However, the program also had a second environmental feature aimed at inducing households to purchase more fuel efficient vehicles than they otherwise would have. It did this by offering subsidies that lowered the relative price of fuel efficient vehicles compared to other vehicles. Specifically, if the new vehicle purchased were a passenger vehicle, it was required to have a combined EPA fuel economy rating of at

⁵There were additional requirements that the clunker be in drivable condition, no more than 25 years old, and continuously insured and registered in the same owner's name for one year prior to the transaction. These criteria appear to have been strictly enforced. The National Highway Traffic Safety Administration (the agency that administered the program) required legal documentation of registration histories and operated the computer system which determined vehicle-specific eligibility.

⁶The extent of the income effect would be moderated by imperfect pass-through, but the literature generally finds that dealerships passed on nearly 100% of the rebates to customers (e.g. Busse, Knittel, Silva-Risso, and Zettelmeyer, 2012).

least twenty-two miles per gallon. If the difference in fuel economy between the new passenger car and clunker was between four and nine MPG, the rebate was \$3500, and if the difference was ten MPG or more, the rebate was \$4500. If the new vehicle was a Category 1 Truck (e.g. SUV or small to medium pickup truck), a two to four MPG difference between the new truck and clunker generated a \$3500 rebate while an improvement of five or more MPG provided a \$4500 rebate.⁷

Although it is clear that the net effect of these restrictions on the vehicle purchased was to lower the relative price of fuel efficient vehicles, the effect of these restrictions on the composition of vehicles purchased is ambiguous *a priori*. One possibility is that these restrictions would induce consumers to spend more money on relatively expensive fuel-saving technologies, such as hybrid electric vehicles. On the other hand, there is a negative overall relationship between MPG and vehicle price among the set of vehicles offered to U.S. consumers, as shown in Figure B.2. Therefore, it is possible that consumers could respond by purchasing smaller, less expensive vehicles. We focus both on how the program shifted the timing of consumer purchases and on how it affected overall spending by changing the composition of vehicles purchased.

2.1.2 Empirical Strategy

Our empirical strategy consists of two steps, both of which make use of household-level data to estimate the effect of the Cash for Clunkers program on purchase behavior. First, we estimate the “pull forward” period induced by the stimulus program. Beginning with the first month of the program, we estimate the time window for

⁷Separate criteria applied to Category 2 (large pickups or large vans) and Category 3 trucks (work trucks), but we do not discuss those here because there were comparatively few of these vehicles purchased. For a complete set of eligibility criteria, see the NHTSA rules in the Federal Register available at:

www.nhtsa.gov/CARS-archive/official-information/day-one.pdf

which the frequency of household purchases of new vehicles is equal for the barely eligible and barely ineligible households. Second, having identified this pull forward window, we focus on all purchases during this time window and analyze differences in the fuel economy and prices paid for households that were barely eligible versus barely ineligible.

To estimate the effect of the Cash for Clunkers program, we use a regression discontinuity design that compares households that were barely eligible for the program to those that were barely ineligible. That is, we compare households whose clunkers were barely above the CARS eligibility cutoff of eighteen miles per gallon to those who barely qualified. We use this regression discontinuity strategy both to identify the pull forward window and to analyze the effect of the program on the types of vehicles purchased.

To formally estimate the reduced-form discontinuities at the eligibility threshold, we use the following equation:

$$\begin{aligned} \text{Outcome}_i = & \beta_0 + \beta_1 * f(\text{distance-to-cutoff}_i) * \text{eligible}_i + \\ & \beta_2 * f(\text{distance-to-cutoff}_i) * (1 - \text{eligible}_i) + \beta_3 * \text{eligible}_i + \epsilon_i \end{aligned} \tag{2.1}$$

where the outcomes include indicators for whether the household received the subsidy and whether the household purchased a new vehicle, the log of the price of the vehicle purchased, and the characteristics of the new vehicles purchased. These outcomes are defined both during the two months of the program, as well as during broader time spans including months afterward.

Eligible_i is an indicator equal to one if the household is classified as being eligible for the program (i.e., the most trade-in-likely vehicle had an MPG rating of eighteen or less). We describe how our data identify a household's eligibility status in Sec-

tion 2.2. We allow for separate relationships between the running variable and the outcome on each side of the eligibility threshold. We estimate equation (2.1) with least squares and standard errors are clustered at the level of the running variable [Lee and Card, 2008]. The coefficient of interest is β_3 , which measures the jump in the outcome when going from just-ineligible to just-eligible for the Cash for Clunkers program.

2.1.2.1 Identifying the “Pull Forward” Window

First, we estimate the number of months after the beginning of the two month program for which the probability of purchasing a new vehicle is equalized across the eligibility threshold. We estimate Equation (2.1) for a dependent variable indicating whether the household purchased a new vehicle during the time window. Importantly, when examining the impact of the program on vehicle purchasing behavior, we use data on all Texas households. Thus, the identifying assumption is that all other determinants of car purchasing behavior among the population of Texas households is smooth across the cutoff. Under that assumption, any discontinuity in the fraction of households purchasing a new vehicle can properly be interpreted as the causal effect of the program.

We view this assumption as likely to hold for several reasons. First, vehicle owners were required to show proof of ownership of their eligible vehicles for one year prior to the start of the program, which is before the policy was being discussed. In addition, the eligibility cutoff was based on the EPA combined fuel economy rating, as opposed to some other, more subjective, rating. As a result, it is difficult to imagine how households could have manipulated where they were relative to the eligibility cutoff. In addition, we know of no other policies that had discontinuous impacts across this eighteen MPG threshold. Collectively, these factors imply that

because we focus on all households, there is little ex ante reason to believe that those with vehicles rated at or just below eighteen MPG are different from those just above the cutoff. We also show empirical evidence consistent with this assumption. For example, we use survey data from the National Household Travel Survey to show that household characteristics such as income and demographics were similar across this threshold. Thus, there is little evidence that policymakers deliberately chose this cutoff because of a discontinuous change in some household characteristic.

2.1.2.2 Estimating Effect of Program on New Vehicle Spending

Our analysis of the impact of the program on new vehicle spending is somewhat different. Rather than using data on all households, we focus only on households that purchased a new vehicle. Crucially, we do this for a time period constructed such that the program did not have an impact on the probability of purchase. Because the net effect on spending depends on both the probability of purchase and the amount spent conditional on purchase, once we hold the probability of purchase constant, the only factor driving the impact on overall spending is the amount spent on the new vehicle. In addition, by focusing on the new car buyers we can avoid averaging spending across all Texas households, over 95 percent of whom did not purchase a new vehicle during the program or in the eight months that followed.

Our identifying assumption for this analysis requires that for households purchasing a vehicle over a period of time during which there was no discontinuity in the probability of purchase, all household-level determinants of new vehicle spending were continuous across the eligibility threshold. We find it difficult to construct a story that would violate this assumption. For example, while it is possible to imagine why barely eligible buyers would be different from ineligible households who bought *during the program*, it is hard to think why this would be true over this longer

time horizon. By construction this longer time horizon contains a similar number of new vehicle buyers across the cutoff – the only difference is that some of those with clunkers rated at eighteen MPG or below were incentivized to purchase earlier during that time window than the other households.⁸ Consistent with this identifying assumption, we show that there is no compelling evidence of discontinuities with respect to the purchasing choices made by households above and below the cutoff a year before Cash for Clunkers. Similarly, households that purchased vehicles during the program or in the months that followed look similar across the cutoff with respect to the characteristics of their non-clunker vehicles, as we show in section 2.3.3.

We emphasize that this research design is a “fuzzy” regression discontinuity design. That is, while the likelihood of receiving a subsidy changes sharply and discontinuously at the eligibility cutoff, it is less than one. This is due to several factors. The first relates to the way in which we classify each household’s so-called clunker. Consider a household whose oldest vehicle is rated at nineteen MPG. As described below, this will lead us to designate that vehicle as the household’s clunker, even though that household may also have owned and traded in a newer car rated at eighteen MPG or less. In addition, households with vehicles rated eighteen MPG or less could choose to purchase vehicles without trading in that vehicle under the program. This could be because they wished to keep that vehicle, or because they wanted to buy a less fuel efficient vehicle that did not qualify under the program. Finally and perhaps more importantly, because we examine time windows that extend beyond the two months of the program, many households purchased vehicles after the program had ended and thus were not eligible for the subsidy at all, regardless

⁸An example which would violate the identifying assumption is if the program were to accelerate some purchases by (say) two years, while simultaneously causing a similar number of eligible households to delay their purchases by more than a year. If that were the case – and it does seem far-fetched – the rate at which households bought vehicles over the ten month window might be similar across the cutoff, even though household characteristics would be different.

of the fuel economy rating of their trade-in.

The fuzzy nature of our regression discontinuity design can be seen in Figure B.3, which shows the discontinuity in the likelihood of receiving the subsidy during the two months of the program, as well as during the ten month window that also includes the eight months following the program. As expected, while the discontinuity in the likelihood of receiving the subsidy was around 75 percentage points during the two months of the program, it is considerably smaller (though still visually clear) over the ten month period. This is consistent with what one would expect; the majority of households purchasing over the ten month period did so after the program had ended, and thus could not use the subsidy regardless of the fuel economy of their clunker.

In order to estimate the impact of the program on new vehicle spending, we estimate equation (2.1) where the dependent variable is either fuel economy or new vehicle price. Formally, this is the reduced-form estimate of the impact of the program, or the intent-to-treat effect. Given the fuzziness of the regression discontinuity design discussed above, in order to recover the local average treatment effect (LATE) measuring the impact of receiving the subsidy, this estimate must be rescaled by the discontinuity in the likelihood of treatment (Angrist, Imbens, and Rubin, 1996; Imbens and Lemieux, 2008). Thus, as shown in Figure B.3, reduced-form estimates for longer time windows will be scaled up more than estimates for shorter time windows.⁹

⁹To adjust for the fact that we were unable to match all 42,354 CARS trades in Texas to households in our data, we rescaled the fraction of sales accordingly. This implicitly assumes a take-up rate for households we could not match similar to those whose clunker was rated the same and whom we could match. This aspect of the data precludes rescaling via two-stage least squares, so we instead manually rescale intent-to-treat estimates to recover the local average treatment effect.

2.2 Data

Our data include all households in Texas. We use confidential administrative records maintained by the Texas Department of Motor Vehicles to determine household level vehicle fleets and changes in the composition of the fleet. For each household, we have information on cars in the household fleet and when the household purchased each vehicle. Following Knittel and Sandler [2011], we restrict our analysis to households that owned no more than seven vehicles in June 2009. For further details on the construction of the database for household vehicle fleets, see Appendix A.1.

Transaction prices for all vehicles sold (new or used) in Texas are reported to the Texas DMV for tax purposes, and we use transaction prices to measure revenue to the auto industry. These prices include any amount of subsidy if the transaction fell under the Cash for Clunkers program, so we are accurately quantifying the revenue received by the industry. Importantly, the DMV administrative records also include the unique vehicle identification number (VIN) for each registered vehicle. We decode each VIN using a database obtained from DataOne Software to determine vehicles' fuel economy and other vehicle characteristics.

We use a simple approach to classify each household's distance from the CARS eligibility cutoff – the running variable in our regression discontinuity design. Our goal in doing so is to determine which vehicle in a household's fleet is most likely to be removed from the fleet when a new car is purchased, and use the fuel economy of that “clunker” to classify the household relative to the eligibility cutoff. We expect these vehicles to be older, lower-value vehicles given the requirement that they be scrapped under the program. We define the clunker for each household as the oldest vehicle that the household owns, measured by the vehicle model year, as of June 30, 2009. In the rare case that a household owns two vehicles with the same model year,

we use the vehicle that the household has owned for the most days.¹⁰ Because the household was required to scrap the clunker and the maximum subsidy was \$4500, we require that the household’s clunker be at least five (model) years old to exclude higher value vehicles that were unlikely to be scrapped. We restrict our sample to households that owned, as of June 2009, a potential clunker obtaining an EPA combined rating of between ten and twenty-seven miles per gallon (inclusive), which spans the largest bandwidth used in our regression discontinuity specifications.

In some specifications, we use demographic data from the Census. These data include Census tract-level economic and demographic characteristics from the 2000 decennial Census, which we link using address information in the administrative database. Finally, in tests of the identification strategy, we use a separate dataset from the spring 2009 National Household Travel Survey [U.S. Department of Transportation, 2009]. Although the NHTS does not include information that allows us to match to our data at the household-level, we can use the rich demographic information in NHTS to support our identifying assumption, as we show in section 2.3.3.

To facilitate our first-stage, we are able to identify transactions that occurred under the Cash for Clunkers program by matching our administrative data to the NHTSA database archive of all program transactions. There were 42,354 official CARS transactions in Texas, of which we match most by VIN to Texas households in the DMV data. The match rate is imperfect, however, due to typos and related database errors. For our empirical first-stage, we scale up the matched set to equate to the full set of CARS transactions.

Summary statistics for vehicle and household fleet characteristics in 2010 are presented in Table B.1. The average fuel economy rating of vehicles purchased by

¹⁰This simple method of defining clunkers yields remarkably similar predictions as that using a more complex propensity score method, while requiring less completeness of data on vehicle characteristics.

households in our sample was 21.7 MPG, while the average transaction price was \$27,600. Table B.1 also shows Census Tract characteristics such as demographics and income, which we use as control variables.

2.3 Results

2.3.1 *Cash for Clunkers and the Timing of Purchase*

First, we examine the impact of the program on the likelihood that a household purchased a new vehicle. Graphical results are shown in Figure B.4, plotting the probability that the household purchased a new vehicle during the time window against the fuel economy of the household's clunker. Markers show the local average for each level of clunker MPG, and marker sizes are proportional to the number of households in the MPG bin. Households just to the left of the vertical line are the eligible households who owned a clunker with a fuel economy below eighteen MPG, while the households to the right of the vertical line are the ineligible households. Because a household has a low probability of purchasing a new car in any given month, the baseline fraction of households purchasing is small over any short time horizon.

Panel (a) of Figure B.4 shows the probability that the household purchased a new vehicle *during the two months of the Cash for Clunkers program*. There is a compelling discontinuity at the cutoff, suggesting that the program increased the likelihood of purchasing a new vehicle by more than one half of a percentage point. This increase is economically significant, and translates to more than a 50 percent increase in the likelihood of purchase during the program. Thus, it is clear that Cash for Clunkers accelerated the timing of new car purchasing by the eligible households.

Importantly, this increase in sales during the program appears to have been offset entirely in the following seven to nine months. Panels (b) through (f) of Figure B.4

show the *cumulative* likelihood of new vehicle purchase over seven to eleven month time frames, including the two months of the program. These panels show compelling evidence that Cash for Clunkers “pulled forward” purchases from the months immediately following the program. While it is clear from panels (b) and (c) that there was still a visually compelling discontinuity in likelihood of purchase after 7 to 8 months, this no longer is the case after 9 to 11 months. The purchase probability is higher for eligible households when focusing only on the program months of July-August 2009, but the ineligible households have “caught up” by March, April, and May of 2010. Put differently, panels (d) through (f) of Figure B.4 suggest that the increase in sales during the program represented an acceleration of sales that would have happened anyway in the seven to nine months after the program ended. This is similar to findings reported by Mian and Sufi [2012], Li et al. [2013], and Copeland and Kahn [2013].¹¹

2.3.2 *Effect on New Vehicle Spending*

As discussed earlier, however, the Cash for Clunkers program also changed the relative prices that consumers faced by offering subsidies that could be used only for the purchase of relatively fuel efficient vehicles. Thus, we now ask whether this environmental component of the program resulted in a lasting change in the composition of vehicles purchased, which has potentially important implications for the stimulus effect on the industry.

As shown in Figure B.5, vehicles purchased by eligible households during the two months of the program were both more fuel efficient and less expensive than

¹¹In principle, we could create time windows that are even more refined than months (e.g. weeks or days) in order to explore equalization of purchase probability. However, our goal is not to define the exact time window for Texas, as other states could have slightly longer or shorter time windows. Rather, our goal is to show that we find results very similar to other studies, and show that our results on stimulus spending are robust to slightly longer and shorter monthly time windows, as we do in section 2.3.3.

vehicles purchased by ineligible households. Corresponding regression estimates are shown in Panel A of Table B.2, which shows estimates for bandwidths of 9, 7, and 5 MPG, and for polynomial fits ranging from cubic to linear. Estimates in Panel A indicate that vehicles purchased by eligible households were rated 1.3 to 1.7 MPG higher than vehicles purchased by ineligible households. All estimates are statistically significant ($p < .001$). Estimates also indicate that the transacted price of vehicles purchased by eligible households was an average of \$2,500 to \$3,000 less than for vehicles purchased by ineligible households during the program. Thus, it is clear that vehicles purchased under the program were both less expensive and more fuel efficient than those purchased by ineligible households during the program.

There are two potential explanations for this finding. On the one hand, the program could have induced households to purchase vehicles that were less expensive and more fuel efficient than they otherwise would have. On the other hand, the differences shown in Figure B.5 could be driven entirely by the acceleration of purchases by households who otherwise would have purchased the same fuel efficient and less expensive vehicles in the months following the program.

To assess these explanations, we examine the characteristics of vehicles purchased *over the entire ten month window*, during which there was no net impact on the likelihood of purchase. Thus, while there is reason to believe that the composition of new car buyers may be different across the cutoff during the two months of the program, there is little reason to believe so over the ten month window. That is because at the start of the program, all of these buyers were going to buy a new vehicle in the next ten months. The only difference is that the barely eligible households were incentivized to purchase more fuel efficient vehicles (and to do so sooner) than barely ineligible households.

Results are shown in Figure B.6, which shows that even over the ten month period

during which there was no net effect on the likelihood of purchase, there are visually compelling discontinuities indicating that barely eligible households purchased vehicles that were more fuel efficient, but less expensive. Corresponding regression discontinuity estimates are shown in Panel B of Table B.2. Estimates are shown for bandwidths of 9, 7, and 5 MPG, and for fits ranging from cubic to linear. Estimates indicate that over this ten month period, eligible households were induced to purchase vehicles rated between 0.6 and 1.0 MPG higher than vehicles they would have purchased absent the program.

Importantly, these estimates show that in order to purchase more fuel efficient vehicles, the barely eligible households spent between \$1,400 and \$2,200 *less* per vehicle than barely ineligible households. All estimates are statistically significant at conventional levels. This is consistent with West, Hoekstra, Meer, and Puller [2014], who show that vehicles purchased by the barely eligible tend to be smaller and have less horsepower per pound of vehicle weight. To confirm that the types of vehicles purchased did change, we show estimates of the effect on vehicle MSRP. The RD estimates using MSRP show that the barely eligible households purchased different, cheaper vehicles over the ten month period.

As discussed earlier, these estimates represent intent-to-treat estimates. Given the discontinuity in the likelihood of receiving the subsidy of around 37 percentage points shown in Panel (b) of Figure B.3, this suggests that each buyer under the program spent on average between \$3,800 and \$5,900 less on a new vehicle than they otherwise would have.

We can provide some perspective for this estimate of reduced spending in terms of specific vehicles. Our estimated average treatment effect represents averaging across some households for whom the fuel efficiency restriction was binding and some for whom it was not. Thus, it is likely that some households treated with the subsidy

“downsized” while others purchased the same vehicle that they would have purchased absent the subsidy. For illustration purposes, suppose that our estimated average spending reduction represents half of households downsizing and half not changing the vehicle purchased. Under this scenario, the downsizing households purchase a vehicle that is around \$9200 cheaper with a fuel economy that is 4.3 MPG higher (the estimated treatment effects divided by 0.5). This difference represents roughly a downsize from a Chevrolet Equinox SUV to a high trim line Toyota Corolla.

In summary, our analysis yields two primary results. The first is that consistent with previous research, we show that the increase in purchases during the two month Cash for Clunkers program was entirely offset by a reduction in purchases over the following seven to nine months. Second, and more importantly, the fuel efficiency restrictions of the program led to a substantial change in the type of vehicles purchased. We show that during the two month program and in the 8 months that followed, eligible households purchased vehicles that were an average of \$1,700 less expensive, which translates to around \$4,600 less per vehicle purchased under the program. Thus, the fuel efficiency restrictions of the program appeared to significantly reduce new vehicle spending over a period of less than a year.¹²

2.3.3 Robustness and Threats to Identification

2.3.3.1 Sensitivity to Time Window

One potential limitation of the analysis presented above is that while we believe there is strong evidence in Figure B.4 that Cash for Clunkers had no net effect

¹²This finding is even more stark than the results of Li and Wei [2013], who estimate structural parameters from a dynamic discrete choice model of vehicle ownership to quantify the tradeoffs between objectives of “green stimulus” programs. Their model parameters imply that more vehicles would have been sold under an alternative policy that subsidized scrappage but did not attach fuel economy restrictions on the new vehicles. In contrast to Li and Wei, our paper finds that the stimulus effect was not only smaller but actually *negative* under the Cash for Clunkers policy with fuel economy restrictions.

on the likelihood of purchase over some time period around ten months, the exact time window is somewhat ambiguous. Consequently, we test the robustness of our main results on both fuel economy and new vehicle spending (i.e. price) to various windows. In doing so, we focus primarily on the local average treatment effect, which represents how much less each household that was subsidized by the program spent as a result of the fuel efficiency restrictions. We find that the LATE estimate of reduced vehicle spending is very robust to the time window.

Robustness results are shown in Table B.3. This table reports estimates for time windows ranging from 9 months to 14 months. We use a bandwidth of 9 MPG and a cubic fit with controls. As one would expect, the reduced-form estimate decreases as the time window expands because we are adding months in which the program was not in effect. For the same reason, the first-stage estimate falls as the window is lengthened. However, estimates remain both statistically significant and economically meaningful, and indicate that eligibility induced households to purchase vehicles that were rated 0.6 to 1 MPG higher than the vehicles they otherwise would have purchased.

Perhaps more importantly, column 3 shows that the local average treatment effect (i.e., the reduced-form estimate reweighted by the magnitude of the first stage) remains remarkably consistent, varying only from 2.6 to 2.7 MPG. Thus, even if one were to believe that it took 14 months rather than 9 months for the sales effect of the program to be completely offset, it does not change the conclusion that the policy induced households to buy significantly more fuel efficient vehicles.

Results for both transaction price and MSRP are similarly robust. While we observe that as expected both reduced-form and first-stage estimates decline as the window is lengthened from 9 to 14 months, the local average treatment estimates remain stable at around -\$5,000. This demonstrates that regardless of the exact time

period during which one thinks the sales effect of the program was offset, there is robust evidence that the program significantly shifted purchases toward less expensive vehicles.

2.3.3.2 *Identifying Assumption of the Regression Discontinuity Design*

Another potential concern with the above analysis is whether the identifying assumption of the regression discontinuity design is valid. For example, while it is difficult for us to construct a plausible story as to how households could have manipulated where they were relative to the eligibility threshold, one might be concerned that policymakers endogenously selected the cutoff based on household characteristics such as income.

We address this issue in two ways. First, we use the spring 2009 National Household Travel Survey to ask whether there are discontinuities in potentially important household characteristics that determine vehicle purchasing behavior.¹³ Assessed characteristics include the number of adults living in the home, the number of weekly travel days, the log of household income, the proportion living in an urban area, the proportion living in a house, and proportion white. As shown in Figure B.7, there is no evidence of visually compelling discontinuities in any of these variables, consistent with the identifying assumption.

In addition, we examine whether the vehicle purchasing behavior was different for barely eligible households than barely ineligible ones in *the year before* Cash for Clunkers. Specifically, we analyze the households that purchased during July and August 2008, including the same calendar months in the year prior to the program. Figure B.8 shows results for the probability of purchasing a new vehicle, the fuel economy rating of those purchases, and the transaction price of those purchases.

¹³The National Household Travel Survey includes information on household demographic characteristics, vehicle ownership, and travel information for a representative sample of U.S. households.

Panel (a) of the Figure shows no evidence of a visually compelling discontinuity, suggesting that households on either side of the cutoff do not differ meaningfully in their underlying propensity to purchase a new vehicle.¹⁴ Likewise, households on either side of the cutoff purchased vehicles in 2008 that were similarly priced and had similar fuel economy, as shown in panels (b) and (c). Importantly, these three graphs for 2008 – the year prior to the program – are distinctly different from the corresponding figures for the two months of the Cash for Clunkers program in 2009 (Figure B.4(a) and Figure B.5(a) and (b)).

Lastly, we examine the characteristics of households who bought during Cash for Clunkers or in the eight months that followed. Our identifying assumption requires that these households be similar across the cutoff. Our administrative data allow us to compare characteristics of household vehicle fleets for households whose clunker is just above and below the eligibility cutoff. Results are shown in Figure B.9. As shown in panel (a), there is little evidence of a difference in the number of vehicles owned by these households. Also, we compare the fuel economy of the *other vehicles* – the non-clunker vehicles – in the households’ fleets (we exclude the clunker because it is by definition smooth through the eligibility threshold). As shown in panel (b), if anything, barely eligible households owned vehicles that were slightly *less* fuel efficient. Assuming this represents a persistent difference in household preferences, it suggests that our main estimates may slightly understate the increase in fuel efficiency induced by the program. Importantly, there is little evidence that barely eligible households purchase less expensive vehicles generally than barely ineligible

¹⁴We note, however, that due to the large sample size, even discontinuities of small economic magnitude such as those shown in Figure B.8 are statistically significant. For example, using a cubic fit and a bandwidth of 9 MPG, the estimated discontinuity in panel (a) of Figure B.8 is two-tenths of a percentage point. This estimate is significant at conventional levels and of approximately the same magnitude as the discontinuity in the probability of purchase for the eleven month post period from July 2009 - May 2010.

households, as shown by the lack of a discontinuity in the MSRP of non-clunker vehicles in panel (c).

2.4 Interpretation and Discussion

As described above, the main finding of our paper is that while Cash for Clunkers did accelerate the timing of purchases, it also reduced new vehicle spending. Specifically, we find that over a period of less than one year, eligibility for the program is associated with a reduction in spending of between \$1,400 and \$2,200 per new-car-buying household. This scales to a \$3,800 to \$5,900 reduction in spending per household that purchased a vehicle with the subsidy.

To translate these estimates into the effect of the program on the national auto industry, we perform a straightforward back-of-the-envelope calculation. There were a total of 677,238 clunker trades in the U.S. Under the assumption that our local average treatment effect also represents the average treatment effect for all Cash for Clunker purchases nationally, this suggests that the CARS program reduced new vehicle spending by \$2.6 billion to \$4.0 billion. Thus, our estimates indicate that the Cash for Clunkers program – while designed to provide stimulus to the auto industry – actually reduced industry revenues by around \$3 billion over a period of less than a year.

One could argue that this decline in industry revenue over the medium run could be justified to the extent the program offered a cost-effective environmental benefit. Unfortunately, the existing evidence overwhelmingly indicates that this program was a costly way of reducing environmental damage. For example, Knittel [2009] estimates that the most optimistic implied cost of carbon reduced by the program is \$237 per ton, while Li et al. [2013] estimate the cost per ton as between \$92 and \$288. These implied cost of carbon figures are much larger than the social costs of

carbon of \$33 per ton (in 2007 dollars) estimated by the IWG on the Social Cost of Carbon [Interagency Working Group, 2013].

2.5 Conclusions

In this paper, we examine the stimulus impact of the Cash for Clunkers program on new vehicle purchases and overall new vehicle spending. We do so by using a regression discontinuity design that compares households barely eligible for the program to barely ineligible households.

Consistent with the existing literature, we show that while the program significantly increased the number of vehicles sold during the two months of the program, this entire increase represented a shift from sales that would have occurred in the following seven to nine months. Thus, over a 9 to 11 month period, the program had no impact on the number of vehicles sold.

Strikingly, however, we show that over a 9 to 11 month period, including the 2 months of the program, Cash for Clunkers actually reduced the amount of money spent on new cars by 2 to 4 billion dollars. We attribute this to the fuel efficiency restrictions imposed on new vehicles that could be purchased with the subsidy, which induced households to buy smaller and less expensive vehicles. In short, by lowering the relative price of smaller, more fuel efficient vehicles, the program induced households to purchase vehicles that cost between \$4,000 and \$6,000 less than the vehicles they otherwise would have purchased.

Thus, while the stimulus program did increase revenues to the auto industry in the short run, the environmental component of the bill actually lowered new vehicle spending over the medium run by inducing people to buy more fuel efficient but less expensive cars. More generally, our findings highlight the difficulty of designing policies to achieve multiple goals, and suggest that in this particular case, environmental

objectives undermined and even reversed the stimulus impact of the program.

3. HOW DO HOUSEHOLDS RESPOND TO INCREASES IN FUEL ECONOMY? REGRESSION DISCONTINUITY EVIDENCE

The negative externalities associated with gasoline consumption have been well-documented, and range from the effects of vehicle emissions on health (e.g. Currie and Walker [2011] and Knittel et al. [2011]) to national security concerns (e.g. National Research Council [2013]) to the impact of carbon emissions on climate change (e.g. Interagency Working Group [2013]). The current level of gasoline taxes in the U.S. is generally thought to be insufficient to correct for these externalities (McConnell [2013]).

However, one possible policy to address these externalities - increasing Pigouvian taxes - is widely considered to be politically infeasible.¹

As a result, United States transportation policy primarily targets the fuel efficiency of vehicles sold by imposing Corporate Average Fuel Economy (CAFE) requirements.²

While these fuel economy standards have largely remained constant for last two decades, the federal government recently set ambitious new standards for the fuel economy of new cars. It is projected that the new standards will increase the average fleet-wide fuel economy of new vehicles to over 50 miles per gallon by 2025, compared to 29 miles per gallon in 2011.

However, increasing the fuel economy of the vehicle fleet will not necessarily lead to a proportionate reduction in fuel consumption. An increase in fuel economy

¹See Knittel [2013] for a history of the (lack of) political support for increasing the gasoline tax dating back to the Nixon administration.

²Extensive research has studied the inefficiencies associated with using fuel economy standards rather than a gasoline tax; see for example Jacobsen [2013], Fischer et al. [2007] and Portney et al. [2003].

reduces the price per mile of driving, which is the price per gallon of fuel divided by the miles per gallon fuel economy. If households respond to the lower marginal cost of driving by increasing vehicle miles traveled, then the effectiveness of this policy in reducing fuel consumption is undermined.³

This problem, originally called the Jevons paradox, is a more general shortcoming of energy efficiency standards. In the case of transportation, this “rebound” or “take-back” effect implies that an increase in fuel efficiency will cause a less-than-proportional decrease in fuel consumption.⁴

Unfortunately, the existing empirical literature on the rebound effect, while speaking to the effects of gasoline taxes, is not well-positioned to assess the impact of increases in fuel economy on driving behavior. Much of this literature uses variation in the price of gasoline to estimate changes in vehicle miles traveled, with the goal of understanding the effect of changing the price per mile of driving while holding characteristics of a household’s vehicles constant. The policy analog of this effect is raising the price of gasoline, perhaps via a tax, while keeping drivers in cars with the same vehicle characteristics.

Understanding the effects of policies that increase fuel economy standards requires one to estimate something different. Requiring that vehicle manufacturers

³Estimates of rebound that receive considerable policy attention are from recent papers by Small and van Dender [2007] and Hymel and Small [2013]. These papers use a representative consumer model that is aggregated to match state-level panel data and simultaneously model the choice of vehicles, vehicle miles traveled (VMT), and fuel economy. Surveys of research on the rebound effect include Gillingham, Rapson, and Wagner [2013b], Austin [2008] and Greening, Greene, and Difiglio [2000]. In addition, a rich literature has modeled the choice and utilization of vehicles in the process of addressing a host of other policy questions; for example see Mannering and Winston [1985], Goldberg [1998], West [2004], Fullerton and Gan [2005], Bento, Goulder, Jacobsen, and von Haefen [2009], Gillingham [2012], and Allcott and Wozny [Forthcoming].

⁴The rebound effect that we address is often referred to as the “direct rebound effect”, addressing the use of the energy-utilizing good that has become more energy efficient. Other income and substitution responses could lead to indirect effects, and macroeconomic effects could exist as well. See Borenstein [2013] for a detailed theoretical discussion and Gillingham et al. [2013b] for a review of the empirical literature.

sell more fuel efficient vehicles is also likely to change the *other characteristics* of the vehicles that are offered. Improvements in the energy efficiency likely leads to reductions in power, weight, and size.⁵ Thus, an understanding of the likely effects of CAFE standards necessitates estimating how households respond to vehicles that are both more fuel efficient *and* smaller and less powerful, as dictated by the technological tradeoffs of vehicle manufacturing. This is a different form of “rebound” that addresses a different policy question than the rebound effect estimated in much of the existing literature. Gillingham et al. [2013b] refer to this form of rebound as a “policy-induced improvement” and argue that the size of this effect is more relevant for understanding the effects of energy efficiency policy such as CAFE.⁶

Estimating the driving response to fuel economy standards is a challenge without credibly exogenous variation in fuel economy of the vehicle fleet. Absent such variation, researchers must turn to data on purchase decisions and driving behavior to estimate the driving response to fuel economy. But credibly estimating this effect is difficult because of the selection problems inherent in individuals’ choice of vehicle. If, for example, drivers with (unobservably) longer commutes choose to drive vehicles with higher fuel economy compared to owners of less-efficient vehicles, rebound estimates from this comparison will be overstated. On the other hand, if low-income buyers choose to drive smaller, less expensive vehicles with higher fuel economy compared to high-income buyers who both drive more and own less-efficient vehicles, then

⁵See Knittel [2011] and Klier and Linn [2012] for an analysis of the technological tradeoffs of fuel economy standards.

⁶Put slightly differently, an MPG-induced change in the price per mile of driving is likely to create both ‘movement along’ and ‘shifting in’ of the derived demand curve for VMT. See Gillingham et al. [2013b] for a thorough discussion of the definitions, estimation, and caveats of interpreting rebound effects. They note that many caveats are necessary when interpreting results from the existing literature, including that estimates do not capture changes in characteristics that accompany efficiency improvements. With those caveats in mind, they summarize estimates of the (direct) rebound effect ranging from 5 percent to 40 percent with most estimates in the range of 5 to 25 percent.

rebound may be understated. In addition, within-household substitution of driving across vehicles can also lead to biased estimates when using vehicle-level data not linked at the household level. For example, if a household replaces a medium-MPG minivan with a high-MPG small sedan, it may well substitute miles toward its other vehicle – say, a low-MPG SUV – which would cause the researcher with vehicle-level data, unable to observe this shift, to overstate the fuel savings. On the other hand, the household may instead substitute miles from the low-MPG SUV to the high-MPG sedan, which would yield larger fuel savings than expected. As a result, it is clear that any attempt to estimate the rebound effect must overcome both of these problems.⁷

Ideally, identification would be based upon plausibly exogenous variation in fuel economy. But given the absence of such variation, researchers have relied instead on another empirical strategy. In a theoretical model of driving, households choose the number of miles driven in response to the price per mile of transport – the price of gasoline divided by fuel economy. Given the absence of exogenous variation in fuel economy, much of the existing literature has exploited variation in the price of gasoline.

This approach using gasoline price variation is limited in its ability to evaluate the effects of fuel economy standards. First, as stated above, it is more suited to addressing gasoline taxes than fuel economy standards because it does not account for changes in vehicle characteristics that accompany increases in fuel economy. Second, setting aside the issue of other vehicle characteristics, households may not respond symmetrically to increases in fuel economy as to decreases in gasoline price. Indeed, there are several reasons to believe that households will respond differently to changes

⁷Knittel and Sandler [2013] show evidence of within-household substitution of miles between vehicles.

in fuel economy and gasoline prices. For example, changes in fuel economy are more certain than changes in gasoline prices, and may therefore induce larger behavioral adjustments.⁸ On the other hand, households may respond more to changes in gasoline prices than to equivalent changes in fuel economy because gasoline prices are more salient [Gillingham et al., 2013a].⁹ For these two reasons, rebound effects derived in part from using variation in gasoline prices may not be informative as to the effect of CAFE standards, the primary way the U.S. attempts to reduce gasoline consumption.

In this study, we estimate the effects of fuel economy on driving behavior using quasi-random variation in the fuel economy of a household’s vehicle fleet. To motivate our source of quasi-random variation, it is useful to begin by considering the ideal experiment. We would identify a set of households who were definitely going to purchase a new car in the next so many months. Then we would randomly assign a subsidy to purchase a high-MPG car, and compare the subsequent driving in households who were treated and untreated.

We identify the rebound effect by exploiting a discrete cutoff in eligibility for the 2009 Cash for Clunkers program that subsidized the purchase of new fuel efficient vehicles when households scrapped fuel inefficient “clunkers”. Households that owned clunkers with a fuel economy of 18 MPG or less were eligible for the subsidy, while households owning clunkers with an MPG of 19 or more were ineligible. This program could have two effects, both of which are discontinuous at the cutoff. First,

⁸For example, Linn [2013] illustrates this possibility with a model of adjustment costs to changing VMT over time.

⁹In fact, Li, Linn, and Muelegger [2012] show that even when restricting attention to gasoline price variation, households respond differently to the tax versus the non-tax component of gasoline price. They find that this differential response affects gasoline consumption but not VMT, which suggests even more nuanced heterogeneity in household behavior. More generally, a broader literature suggests that consumers may not respond symmetrically to all components of a price [Chetty, Looney, and Kroft, 2009].

the program could change the fuel economy of new cars that are purchased. Second, the program could change the set of households who purchase new vehicles by inducing some households to purchase who otherwise would not be in the market for a new car. The ideal experiment measures the first effect.

If the second effect is present and the set of households induced to purchase by the program differs in driving behavior from those who were already in the market for a car, then this approach could create bias. Therefore, as we describe below, we perform a number of placebo tests to show that any bias is likely to be minimal.

We do so by using data on vehicle registrations in Texas matched to data on annual odometer readings taken during state emissions tests. We focus on households that purchased a new vehicle in 2009, including those who purchased under the program and those who did not. We show that new car buyers whose “clunker” was barely eligible for the subsidy purchased substantially more fuel efficient vehicles than new car buyers whose “clunker” barely exceeded the 18 miles per gallon maximum fuel economy allowed to qualify for the program. Importantly, given the nature of the program, we have every reason to believe that this difference in household fuel efficiency arose as a direct result of the restrictions imposed by the Cash for Clunkers program, rather than from differences in unobserved driving habits or preferences. As a result, we find this approach to be considerably more compelling than one based on panel data, where one might worry that a change in household fuel economy over time is caused by changes in unobserved income or commute distance, which themselves would affect vehicle miles traveled.

The identifying assumption of our study is that all other determinants of vehicle miles traveled in 2010 are continuous across the eligibility cutoff. Under that assumption, any discontinuity in vehicle miles traveled at the cutoff is properly interpreted as the causal effect of the change in household fuel economy. Several factors give us

confidence that this identifying assumption holds. First, Mian and Sufi [2012] and Hoekstra, Puller, and West [2014], among others, show that Cash for Clunkers did not fundamentally alter the total number of car purchases but rather led to an intertemporal shift ahead of purchases by about ten months. As a result, there is little reason to believe that those who were replacing vehicles barely eligible for the Cash for Clunkers program were meaningfully different types of consumers from those replacing vehicles that were barely ineligible for the program. Second, we show that in the year prior to the program, new car buyers barely eligible for the program drove a similar number of miles as those barely ineligible for the program. This implies that the two groups of households were similar with respect to their underlying driving behavior, consistent with the identifying assumption. Third, we find no evidence that new car buyers barely eligible for the program experienced worse negative income shocks - as proxied by the level or change in the likelihood of financing the new car - than the other new car buyers. Similarly, we show that there was no discontinuity in the likelihood of moving to a new residence after 2008 across the eligibility threshold, which indicates there is no evidence of time-varying changes in commutes. In short, we find that new car buyers on either side of the threshold were similar with respect to both time-invariant and time-varying determinants of vehicle miles traveled.

Our results show that despite the significant increase in fuel economy, households do not respond by driving more miles. Point estimates are small, negative, and not statistically different from zero. As we show below, the households increasing fuel economy purchased vehicles that were smaller, lighter, and had less horsepower. This suggests that the net effect of households owning more fuel efficient (but smaller and less powerful) cars is no change in the amount of total driving. Importantly, our estimates are sufficiently precise as to rule out large responses in vehicle miles traveled.

This paper makes two contributions to the literature. First, to our knowledge, this is the first paper to exploit credibly exogenous variation in household fuel economy to identify the effect on driving behavior. As a result, we are able to obtain estimates that are causal under reasonable assumptions, without the need to impose stronger assumptions required to model vehicle purchase and driving decisions.

Second, our estimate better captures the policy-relevant parameter for policy-makers who target fuel economy standards, rather than gasoline prices, in their attempt to reduce consumption. We identify the rebound effect using variation from a program that mirrors the policies imposed by CAFE. That is, just as the CAFE requirements are intended to reduce fuel consumption by changing the composition of vehicles purchased, Cash for Clunkers similarly restricted the choice set of vehicles for eligible buyers. As a result, we view our estimates as more informative as to the impact of CAFE on driving relative to approaches that rely on other types of variation in driving cost.

We emphasize that the driving response we identify is different from the rebound effect measured in much of the existing literature. Our measured effect includes *both* changes in the price per mile and other characteristics, as both are affected if households choose more fuel efficient cars. In contrast, the rebound effect in much of the existing literature imagines a change in fuel economy without sacrificing size, power, comfort and safety. This suggests that at the very least, using rebound estimates that hold vehicle characteristics constant can overstate the VMT response to fuel economy standards. This has important policy implications, as the National Highway Traffic Safety Administration (NHTSA) explicitly accounted for a rebound effect when it was designing the 2012 Corporate Average Fuel Economy (CAFE) standards.¹⁰

¹⁰NHTSA assumed a rebound effect of 10% – that is a 10% increase in fuel economy causes

Our findings suggest that, while there are plausible reasons why fuel consumption may not fall as much as predicted in response to changes in fuel economy, the empirical evidence suggests this is not the case. Rather, our results indicate that when households are induced to buy more fuel efficient vehicles, they do not respond by driving more miles in the year after improving fuel economy.

These findings have two important implications for energy efficiency policy. First, we quantify a parameter of considerable policy interest – the effect of downsizing the vehicle fleet on vehicle utilization. We argue that our estimate is a more policy-relevant metric than many estimates in the existing literature for understanding how the downsizing of an individual household’s fleet affects vehicle miles traveled. Inducing households to choose more fuel efficient cars among the current set of vehicle offerings is not likely to increase a household’s total driving and exacerbate driving-related externalities. This should give policymakers some cause for optimism, as it suggests that second-best strategies such as CAFE used to combat the negative externalities associated with gasoline consumption are more effective than previously thought.

Second, our results have implications for evaluating the welfare comparisons that are frequently made between first-best policies such as a gasoline tax and actual energy efficiency policies such as CAFE. A standard view is that gasoline taxes induce households to purchase vehicles with the optimal level of fuel economy and then to drive the fleet the optimal number of miles. Fuel economy standards have been criticized as inefficient on the intensive margin by distorting vehicle utilization relative to first-best. This paper makes an important point – extensive margin policies can have countervailing effects on intensive marginal utilization decisions. One effect of

households to drive 1% more, thus “taking back” some of the potential fuel savings. See U.S. Department of Transportation [2010].

increasing fuel economy is captured by a price elasticity of driving – altering the fuel efficiency of the fleet reduces the price per mile of driving. A second effect is a vehicle elasticity of driving – shifting households to fuel efficient cars with different characteristics can reduce the utility per mile of driving and thus the amount of driving. Both of these effects must be captured by a complete welfare analysis to compare a particular policy to first-best.

Our empirical approach places strong emphasis on identifying causal impacts of fuel economy by exploiting quasi-random variation in fuel economy, which to our knowledge is new to the literature. A limitation of this approach is that we are not in a position to estimate the relative magnitudes of these two elasticities or to calculate welfare measures. However, our analysis does suggest that the joint effect of these two elasticities is zero in the short-run.

3.1 Background and Empirical Strategy

3.1.1 The Cash for Clunkers Program

We exploit the Cash for Clunkers program as a quasi-random source of variation in the fuel economy of a household’s vehicle fleet. The program, formally known as the Consumer Assistance to Recycle and Save (CARS) Program, created incentives for households to replace used, fuel inefficient vehicles with new, fuel efficient vehicles. The program lasted for eight weeks during the summer of 2009 and offered households a rebate of \$3,500 or \$4,500 towards the purchase of the new fuel efficient car when they scrapped their “clunker.” A requirement of the program was that the clunker had to be taken off the road and scrapped; thus the rebate could be viewed as the trade-in value of the old car from the perspective of the household. Due to the scrappage requirement, the program attracted relatively older and low value vehicles. The average age of scrapped clunkers was 13.8 years.

The CARS Act was signed into law on June 24, 2009 and transactions first became eligible for rebates on July 1, 2009. Initial take-up of the program was substantial, and the \$1 billion that was allocated under the law quickly ran out. Congress allocated an additional \$2 billion on August 7, and those funds quickly were exhausted as well. The program ended on August 24 with over 677,000 vehicles purchased, 44,000 of which were in Texas.

The criteria for eligibility provide us with cutoffs for our regression discontinuity research design. The clunker must have had a combined EPA fuel economy of 18 MPG or less.¹¹ The vehicle purchased must have been a new vehicle; used vehicles did not qualify for the rebate. If the new vehicle was a passenger vehicle, it must have a combined fuel economy of at least 22 MPG. In the case of passenger vehicles, if the difference in fuel economy between the new passenger car and clunker was between 4 and 9 MPG, the rebate was \$3500, and if the difference was 10 MPG or more, the rebate was \$4500. If the new vehicle was a Category 1 Truck (e.g. SUV or small to medium pickup truck), a 2-5 MPG difference between the new truck and clunker generated a \$3500 rebate while an improvement of 5 or more MPG generated a \$4500 rebate.¹² Busse, Knittel, Silva-Risso, and Zettelmeyer [2012] find that dealerships passed on nearly 100% of the rebates to customers.

These criteria create a discontinuous eligibility threshold – households who owned clunkers that had fuel economy of 18 MPG or less were eligible for CARS rebates whereas households with 19 or more MPG clunkers were not eligible. Below, we describe how we use our data to classify each household’s eligibility status.

¹¹There were additional requirements that the clunker be in drivable condition, no more than 25 years old, and continuously insured and registered in the same owner’s name for one year prior to the transaction.

¹²Separate criteria applied to Category 2 (large pickups or large vans) and Category 3 trucks (work trucks), but we do not discuss those here because there were so few of these vehicles. For a complete set of eligibility criteria, see the NHTSA rules in the Federal Register available at: <http://www.nhtsa.gov/CARS-archive/official-information/day-one.pdf>

CARS transactions resulted in an increase in the fuel economy of the vehicle fleet for those households that purchased under the program. The average fuel economy of the scrapped clunker was 15.8 MPG while the average fuel economy of new cars purchased under the program was 24.9 MPG.¹³ We should note that we do not evaluate the CARS program directly; rather we use the program design as a source of quasi-random variation in fuel economy. A separate literature has evaluated how well CARS achieved program objectives (for example, see Knittel [2009], Copeland and Kahn [2013], Busse et al. [2012], Mian and Sufi [2012], Li, Linn, and Spiller [2013], and Hoekstra, Puller, and West [2014]).

3.1.2 Empirical Strategy

We use a regression discontinuity design to estimate the impact of an exogenous increase in fuel economy on vehicle miles traveled. Intuitively, we compare households whose “clunkers” were barely above the CARS eligibility cutoff of 18 miles per gallon to those whose barely qualified. Importantly, we focus our analysis on new car buyers, rather than all car owners. We do this because we otherwise cannot disentangle the effect of driving a more fuel efficient car from the effect of driving a newer car.¹⁴ In section 3.3.1, we present evidence that this does not confound our estimates.

Formally, we estimate the following equation:

$$\text{Outcome}_i = \beta_0 + \beta_1 * f(\text{distance-to-cutoff}_i) * \text{eligible}_i + \beta_2 * f(\text{distance-to-cutoff}_i) * (1 - \text{eligible}_i) + \beta_3 * \text{eligible}_i + \epsilon_i$$

where the outcomes are the logarithms of MPG, VMT, and gallons of fuel consumption by household i in the year after the program (2010).

¹³C.A.R.S. Program Statistics, 2009, are available from the NHTSA.

¹⁴In addition, the number of new cars purchased under the Cash for Clunkers program is small relative to the total stock of vehicles in Texas, making the increase in fuel efficiency across all households at the eligibility cutoff statistically and economically undetectable.

$Eligible_i$ is an indicator equal to 1 if the household is classified as being eligible for the program (i.e., the most trade-in-likely vehicle had an MPG rating of 18 or less). We control for a different quadratic relationship between the running variable and the outcome on each side of the eligibility threshold, though we show results of linear specifications as well.

The coefficient of interest is β_3 , which measures the jump in the outcome when going from just-ineligible to just-eligible for the Cash for Clunkers program.

To be sure, this analysis is not sufficient to estimate the complete set of effects of future CAFE standards on driving for several reasons. First, our approach estimates the treatment effect for households with “clunker” fuel economy around 18 MPG and does not speak to how other households respond to increases in fuel economy. We should note however that 18 is not far below the average fuel economy for older vehicles in the fleet and that the potential fuel savings from efficiency improvements is larger for lower fuel economy cars. Second, we capture changes in driving behavior in the year after fuel economy improvements occur. Other behavioral changes such as work or residential location that may take longer to occur are not captured. And third, our analysis measures household response given the current technological tradeoffs between fuel economy and other vehicle characteristics. Future technological innovation by automakers could change the vehicle features that must be sacrificed to improve fuel efficiency.

The identifying assumption of this approach is that all other determinants of VMT in 2010 varied smoothly across the Cash for Clunkers eligibility cutoff. There are several potential threats to identification, however as we show below in section 3.3.1, we do not find evidence of these identification problems.

First, one might be concerned that the 18 MPG cutoff was chosen by the government in a strategic manner, such that households who were barely eligible were

different from those barely ineligible. For example, the government could have chosen the cutoff to target certain types of vehicles or certain vehicle manufacturers, which might lead to vehicle owners have different characteristics on either side of the eligibility threshold. However, we find no evidence of this form of manipulation – we show that households vary smoothly through the discontinuity across a large number of demographic characteristics.¹⁵

A second, and potentially more serious, concern is that the program changed the set of households who purchase a car, and that the households induced to buy a car are different in a dimension that affects driving behavior. Specifically, our sample includes all households who purchased a car during the months when the Cash for Clunkers Program was underway. The “barely ineligible” set consists of households who purchased a car without any program inducement. However, the “barely eligible” set consists of some households who would have purchased a car anyway (but chose a more efficient car due to the program incentive), *and* possibly a set of households who only chose to buy a new vehicle because of the program. Are the households induced to purchase cars different from those households who would have purchased in absent the program? Because we find that the barely eligible households did not drive more in the year after the program, one might be concerned that the set of households induced to purchase were households who either drove less on average or were subject to more negative economic shocks that affect driving.

However, as we show below, the empirical evidence suggests that these concerns are not warranted. In the year prior to the program, the barely eligible households were very similar in the fuel economy of their fleets, and if anything drove slightly

¹⁵Manipulation by the households was even less plausible. Because households were required to own the “clunker” for one year prior to trade-in, there was little scope for households to manipulate where they were relative to the cutoff.

more prior to the program, which would bias against our findings. Moreover, we find that the barely eligible households are not more likely to have faced time-varying economic shocks that impact driving such as moving to a different residence or financing the new vehicle.

These findings suggest that selection into our sample does not bias us towards finding that households do not drive more miles in response to owning more fuel efficient vehicles. This is not surprising given recent research that evaluated the Cash for Clunkers Program. Mian and Sufi [2012] show that new vehicles purchased under Cash for Clunkers represented an intertemporal shift ahead of total purchases by less than a year. As a robustness test, we expand our sample to include all households that purchased a vehicle from the beginning of the Cash for Clunkers program for the next 12 months. We show that this definition of a purchase window implies that among all households in Texas, the barely eligible are just as likely to buy a new car as the barely ineligible, consistent with Cash for Clunkers causing merely an intertemporal shift ahead of purchases. When we use this 12 month sample window, we obtain estimates that are quantitatively similar to our primary findings. These identification tests, taken together, suggests that many of the households induced to purchase under Cash for Clunkers may have been in the market for a new car in the near year anyway, and thus were not meaningfully different from the households who purchased without program inducement.

In summary, while we cannot be absolutely certain that the identifying assumption of our research design is valid, we note that we find little evidence to refute it, and considerable evidence to support it.

3.2 Data

Our empirical setting is Texas, the second largest U.S. state as measured by either population or consumption of gasoline for transportation.¹⁶ We use several large administrative databases in Texas for our study. First, to determine household-level vehicle fleets and ownership spans, we use confidential vehicle registration records maintained by the Texas Department of Motor Vehicles (DMV). In addition to providing a measure of a household’s vehicle fleet, these records include the unique vehicle identification number (VIN) for each registered vehicle.

Our measure of vehicle-level miles traveled (VMT) is computed from odometer readings recorded during annual vehicle emissions tests, which we link by VIN. An important institutional feature for our study is that emissions tests are required annually for each vehicle older than two years in seventeen EPA non-attainment counties, a more stringent requirement than that mandated in many states. These counties include the areas surrounding Houston, Dallas-Fort Worth, Austin, and El Paso.¹⁷ Although Texas is sometimes stereotyped as having more trucks and heavy vehicles than other states, the mix of vehicles in these four urban areas is very similar in terms of fuel economy to that in many urban areas across the U.S. (see, e.g., Busse, Knittel, Silva-Risso, and Zettelmeyer [2012]’s Figure 9 on fuel economy for each Census tract in the country). From these two databases, we calculate household vehicle ownership, VMT, and annual fuel consumption. We provide details on this process in Appendix A.

¹⁶Measures of state-level gasoline consumption by end use are available for 2011 from the U.S. Energy Information Administration at http://www.eia.gov/state/seds/sep_fuel/html/pdf/fuel_mg.pdf. Combined, California and Texas account for approximately twenty percent of U.S. population and gasoline consumption.

¹⁷The Texas Commission on Environmental Quality (TCEQ) provided us with emissions test records for vehicles in the seventeen EPA non-attainment counties in Texas, which are shown in Figure B.10. While these counties comprise only 7% of the 254 counties in Texas, they include four of the largest metropolitan areas and nearly 60% of the state’s population.

To determine vehicles’ fuel economy and other vehicle characteristics, we decode each VIN using a database obtained from DataOne Software. We compute a household’s annual fuel economy by averaging the EPA combined miles per gallon of its vehicle fleet, weighted by the days within the year that the household owned each vehicle.¹⁸ Importantly, we do not weight by VMT of the different vehicles because the treatment that we want to evaluate is the effect of making a household’s fleet comparatively more efficient. In addition, since the focus of our study is on household drivers, rather than institutional fleets, we follow Knittel and Sandler [2011] in restricting our analysis to households that owned no more than seven vehicles in 2010.

We classify each household’s distance from the CARS eligibility cutoff – the running variable in our regression discontinuity design – via a propensity score estimation. Intuitively, we want to know which vehicle in a household’s fleet is most likely to be removed from the fleet when a new car is purchased, and use the MPG of that “clunker” to classify the household relative to the eligibility cutoff. Because the CARS program required that the clunker be scrapped in return for the subsidy of \$3500-\$4500, we would expect the scrapped vehicles to be older, lower value vehicles that the household has owned for a longer period of time.

To classify each household, we begin by merging VINs that were actually scrapped under the CARS program in 2009 to the CARS database available from the National

¹⁸It is not a priori obvious which metric of average household fuel economy is most suitable. If a multi-vehicle household drives each vehicle the same number of miles, then the harmonic mean may be appropriate. However, experimental research has shown that consumers do not properly understand the non-linear relationship between fuel expenditures and MPG (Larrick and Soll [2008]), so households may mentally calculate the average MPG across vehicles in a manner that differs from the harmonic mean. Other possible metrics that could capture actual decision making include the log of the arithmetic average and average of the log MPG. Given the literature on ‘MPG illusion’, it is not obvious which metric is most appropriate. Therefore, we estimate our models separately using each metric, and find that the estimates do not meaningfully differ across definitions. In this paper, we report estimates using the log of the arithmetic average, but results with the two other metrics are available upon request.

Highway Transportation Safety Administration.¹⁹ Then, similar to Busse, Knittel, Silva-Risso, and Zettelmeyer [2012], we use the stock of vehicles as of 2008 for households we match to actual CARS trades to estimate which vehicle characteristics increase the likelihood of a vehicle being a CARS trade. We regress an indicator equal to one for a CARS-traded vehicle on dummies for vehicle type, body style, number of doors, and all- or four-wheel drive, along with the linear, square, and cubic of manufacturer suggested retail price (MSRP), curb weight, horsepower, vehicle age as of 2009, latest odometer reading prior to 2009, and the number of days that the household owned the vehicle as of December 31, 2008. Importantly, we do not include MPG in this propensity model.

Based on the results from this estimation, we predict the likelihood of each vehicle in our sample being scrapped. Then we use the vehicle in each household's fleet that has the largest p-score – i.e. the car most likely to be scrapped – to classify the household's distance from the CARS eligibility cutoff. Consider a specific example – a household with a relatively new Ford F150 (MPG of 16) and a relatively old Toyota Corolla (MPG of 28). Under a scrappage program providing subsidies of \$3500-\$4500, the household would only consider parting with the low value Corolla; it would not make sense to scrap a high value F150 for a mere \$4500 subsidy. Thus, we want to classify this household as ineligible based upon the high fuel economy of the Corolla. Any vehicle with a value substantially higher than the maximum subsidy does not realistically make the household eligible. We have confirmed that our p-score approach successfully excludes newer and high value cars from the eligibility criteria of households and instead classifies households by the old/low value cars in each household's fleet. These vehicles appear to match quite well the actual vehicles

¹⁹Transaction data from the Car Allowance Rebate System (CARS) program are available from the NHTSA and include the full seventeen-digit VIN for traded vehicles and the eleven-digit VIN for purchased vehicles.

scrapped under CARS.

Because we leverage the fuel economy shock provided by the 2009 CARS program, the time frame of our study is 2010. Our empirical sample is restricted to households that: (1) owned, as of December 2008, a potential “clunker” obtaining an EPA combined rating of between ten and twenty-seven miles per gallon (inclusive); and (2) purchased a new vehicle during the CARS time period in 2009.²⁰ Finally, to provide additional control variables in several of our specifications, we link households to tract-level economic and demographic characteristics in the 2000 decennial Census. Summary statistics for vehicle and household fleet characteristics in 2010, along with Census 2000 tract-level covariates, are presented in Table B.4.

3.3 Results

3.3.1 Tests of the Identification Strategy

The identifying assumption of our study is that all other determinants of VMT in 2010 varied smoothly across the Cash for Clunkers eligibility cutoff. This assumption may be violated if, for example, politicians endogenously selected the eligibility threshold based on the types of vehicles or owners who would qualify. Therefore, we test this assumption in several ways, first focusing on the entire population of vehicle owners. Specifically, we use data from the 2009 National Household Travel Survey to test whether there are discontinuities in household characteristics or demographics across the eligibility cutoff. Results are shown in Figures B.11 and B.12, which take the same form as figures following them. The x-axis shows the running variable of the MPG of the most-likely “clunker” owned by the household, and the y-axis shows

²⁰We classify as new any vehicle that was not previously registered in Texas, was purchased for more than \$4000, and had fewer than 500 miles on its odometer at the time of sale. The Texas DMV registration database records the date on which the title was officially transferred rather than the date of a vehicle’s sale, and there is often a lag time of 30-60 days between these dates. To better correspond to the timing of CARS, our sample includes new vehicles for which the title transferred between August 1st and October 31st, 2009.

the outcome variable. Households at 18 MPG and below are “barely eligible”, and households just to the right of the vertical line are “barely ineligible”. The circles and triangles represent local averages, where size corresponds to the number of observations in that cell. We also show fitted lines corresponding to a regression in which we allow for separate fits of the running variable on either side of the eligibility cutoff.

As shown in Figures B.11 and B.12, there is little compelling visual evidence of discontinuities in vehicle owner characteristics such as household income, weekly travel days, or the number of adults, consistent with the identifying assumption. Corresponding regression results are shown in Table B.5; consistent evidence of a statistically significant discontinuity is found for only 1 of 12 measures.^{21,22}

As described earlier, our analysis focuses on new car buyers, rather than on all car owners. We do this because we otherwise cannot disentangle the effect of driving a more fuel efficient car from the effect of driving a newer car. As a result, our identifying assumption requires that, for those who bought new cars in 2009, all other determinants of VMT in 2010 varied smoothly across the eligibility cutoff. Our first test of this is to examine whether barely eligible buyers had different driving behavior prior to the program in 2008.

These estimates are shown in Table B.6, with the graphs in Figure B.13 providing visual evidence. The first row of Table B.6 shows the discontinuity in the log of MPG, while the second and third rows show estimates for total VMT and total gallons of fuel consumed, respectively. Estimates are shown for bandwidths of 9 MPG, 8 MPG, 7 MPG, 6 MPG, and 5 MPG.²³ Column 1 shows estimates using a bandwidth of 9

²¹Households with vehicles that barely qualify for the program are 2 percentage points more likely to be Hispanic.

²²We note that in Figure B.11(f) owners of vehicles rated at 19 MPG appear to have idiosyncratically high income. As a result, we have estimated our main results excluding this group of individuals, and find similar results.

²³We do not examine bandwidths larger than 9 MPG because there were only 3 so-called “clunker” cars in our data with ratings of 9 mpg; all 3 were Ferrari Enzos.

MPG, allowing for a quadratic function that is allowed to differ on either side of the cutoff. Column 2 is the same, except that it includes controls including county of residence fixed effects and the 2000 tract-level covariates from Table B.4 measuring population density, income, and other demographics. Columns 3 and 4 are similar to Columns 1 and 2, except they allow for linear functions of the running variable.

Estimates of the differences in household fuel economy shown in Table B.6 are statistically significant but economically very small. Estimates are close to 1 percent, which translates to a difference of only 0.19 miles per gallon. Estimates of the discontinuity in total VMT in 2008 are around 5 percent and are marginally significant. As a result, estimates are consistent with the visual evidence in Figure B.13 that households that later bought new cars in 2009 and were barely eligible for the program were, if anything, possibly predisposed toward driving more miles. This implies that our estimates may overstate the rebound effect somewhat.

The overall similarity in driving patterns across these two groups of new car buyers in 2008 indicates that for a confounding factor to lead us to understate a rebound effect in 2010, there must have been a time-varying shock that led barely-eligible new car buyers to drive less than they would have otherwise. We test for such a shock in two ways. First, we examine whether there was a discontinuity in the likelihood of moving to a new residence since 2008. For example, if those barely eligible for the program were more likely to move closer to work than those barely ineligible, our rebound estimates may be biased downward. Figure B.14(a) shows the results, with corresponding regression estimates in Row 1 of Table B.7. In short, there is little visual or statistical evidence of a discontinuity in the likelihood of moving, consistent with our identifying assumption.

In addition, to proxy for income shocks over time, we ask whether there is a discontinuity in the likelihood of financing the new car, relative to the household's

historical likelihood of financing new car purchases. Estimates are shown in Row 3 of Table B.7. Consistent with the raw data shown in Figure B.14(c), results indicate that those barely eligible for the program were actually somewhat less likely to finance a car in 2009 than they were earlier. While this is not a perfect test, given the potential for the program to directly affect financing through both the subsidy as well as the MSRP of the eligible car,²⁴ it does provide evidence that barely-eligible households were no more likely to suffer a negative income shock between 2008 and 2010 than barely-ineligible households.

3.3.2 *Driving Outcomes*

We next turn to whether those barely eligible for the Cash for Clunkers program – that is, those whose “clunkers” had ratings of 18 MPG or less – were more likely to buy a more fuel efficient new car than those buyers who were barely ineligible for the program. Results are shown in Figure B.16(a), and indicate that buyers who were barely eligible for the program bought significantly more fuel efficient vehicles relative to those barely ineligible for the program. This suggests that the subsidy and restrictions of the program did have their intended effect, in that they induced individuals to buy more efficient vehicles than they otherwise would have. However, as shown in Figure B.16(b), there is no evidence that those barely eligible for the program responded by driving more miles; if anything, it appears that they drove fewer miles. As a result, Figure B.16(c) shows that fuel consumption fell.

Corresponding regression estimates are shown in Table B.8. Consistent with the raw data shown in Figure B.16(a), estimates in Column 1 of Table B.8 indicate that Cash for Clunkers eligibility induced buyers to purchase vehicles that were between

²⁴As shown in Figure B.15 in the appendix, vehicles purchases by households eligible for the subsidy were on average smaller, less powerful, and lower-priced than those purchased by households barely ineligible for the subsidy.

4 and 6 percent more fuel efficient, with all estimates statistically significant at the 1 percent level. These estimates provide evidence that as intended, program eligibility did induce drivers to buy more efficient vehicles than they otherwise would have purchased. Importantly, the magnitude of the estimates is unaffected by the inclusion of controls, which is consistent with our identifying assumption. For example, the estimate in Column 1 is 0.0618, while adding controls in Column 2 leaves the coefficient nearly unchanged at 0.0587.

The second row of Table B.8 shows regression discontinuity estimates of the impact of Cash for Clunkers eligibility on vehicle miles traveled. Estimates are small and statistically indistinguishable from zero across all specifications, ranging from -0.0062 to 0.0006. As with MPG, estimates are similar across specifications with and without controls, as well as across bandwidths and polynomials. In short, consistent with Figure B.16(b), there is no evidence that inducing drivers to buy more fuel efficient vehicles causes them to drive more miles one year later.

While this finding may seem counterintuitive, it is plausible given that the more efficient vehicles selected by consumers had less desirable vehicle attributes. The new vehicles purchased by the barely eligible households were both smaller and less powerful than the new cars purchased by the barely ineligible. Figure B.15 in Appendix B illustrates that the more fuel efficient cars purchased by the barely eligible had discretely less horsepower and a smaller cubic volume. Put differently, given available fuel-saving technologies and their costs, policy-induced consumers chose to purchase cars that likely gave them lower utility than the vehicles purchased by the barely ineligible. As a result, the rational response to drive less may have offset any rebound effect arising from the lower marginal cost of driving.

We also examine the effect of program eligibility on the log of gallons of gasoline consumed. Our measure of gallons consumed is calculated based upon each vehi-

cle's VMT and EPA-rated fuel economy; it does not reflect any heterogeneities in a household's realized fuel economy. Given the mechanical relationship between MPG, VMT, and fuel consumption, there is little surprise that estimates in row 3 are negative, ranging from -0.043 to -0.0675. Estimates corresponding to specifications that control for quadratic functions of the running variable are statistically significant at the 1 percent level, while the others are not quite significant at the 10 percent level.

3.3.3 Household-Level Outcomes

As noted earlier, however, the purchase of a more fuel efficient vehicle could lead to substitution across vehicles within households. On one hand, a household may have an incentive to shift miles to the more efficient vehicles in its fleet. However, the vehicles with higher fuel economy also have product characteristics that may lower the utility of driving (horsepower and size). To directly test for the presence of substitution in either direction, we estimate the impact of program eligibility on the fraction of household VMT driven by the new car. Results are shown in the fourth row of Table B.8. Estimates are economically small and statistically insignificant, and range from -0.0016 to 0.0144. In short, there is little evidence that households shifted driving toward or away from the more fuel efficient new vehicle.

We also examine the impact of program eligibility on driving outcomes at the household level. Figure B.17 shows the raw data, and corresponding estimates are in the last three rows of Table B.8. Unsurprisingly given the lack of evidence of within-household substitution, results are overwhelmingly similar to those at the new vehicle level. Regression discontinuity estimates for household MPG range from 3 to 5 percent, all of which are statistically significant at the 1 percent level. As before, however, there is no evidence of an increase in total household VMT: estimates are negative and statistically indistinguishable from zero, and range from 0 to 3

percent. Thus, our findings suggest that potential increases in driving due to the lower marginal cost are offset by the effect of driving a car that is smaller and less powerful than less fuel efficient cars.

Ideally, we would be able to decompose household response into the response to the lowered marginal cost of driving and the response to driving a vehicle with different characteristics. However, there is no significant variation in household fuel economy that is orthogonal to vehicle characteristics. For example, while there is a clear discontinuity in unconditional household miles per gallon at the eligibility threshold as shown in Figure B.17(a), the discontinuity disappears if one conditions on new vehicle characteristics such as size and performance. This reflects the set of vehicles that manufacturers offer and consumers purchase – there are not (enough) vehicles purchased that are the same size with the same horsepower that also have substantial differences in fuel economy. Thus, we are unable to separately identify these two effects, though it appears that they largely offset each other.

In summary, we find that new car buyers barely eligible for the Cash for Clunkers program were significantly more likely to purchase more fuel efficient vehicles. However, despite the fact that these drivers face lower marginal costs of driving, households do not respond by driving more miles, and thus consume less gasoline.

3.4 Conclusions

The question of whether increases in fuel economy will increase miles driven and thus the externalities that are associated with driving and gasoline consumption is crucial for policy. To our knowledge, this is the first paper to address this question using quasi-random variation in a household's fuel economy. Specifically, we show that while households who were barely eligible for the subsidy under Cash for Clunkers purchased significantly more fuel efficient vehicles, they did not respond by

driving more miles. As a result, we find that there is no evidence of a “rebound” effect that would offset some or all of the reduction in fuel consumption that would arise when households are induced to drive more efficient vehicles.

This contrasts significantly with the existing literature, which has generally found larger rebound effects. While we cannot rule out the possibility of heterogeneous treatment effects across the different populations studied as a potential explanation, we believe there are two more likely explanations. First, the existing literature has primarily estimated the elasticity of driving with respect to the cost per mile, where much or all of the variation in cost per mile is due to variation in gasoline prices. It may well be that drivers respond less to changes in fuel economy than to changes in gasoline prices, as the latter are much more salient. In addition, estimates exploiting gasoline price variation hold vehicle characteristics constant. As we discuss above, the existing literature is better-suited to address the effects of gasoline taxes. In this paper, we identify effects solely off of variation in fuel economy induced by eligibility for the Cash for Clunkers program. Our measured effect includes *both* changes in the price per mile and other characteristics, as both are affected if households choose more fuel efficient cars. Our estimated driving response captures the tradeoffs that consumers face when choosing among product choices offered by vehicle manufacturers. One interpretation of our findings is that any increase in driving caused by the reduction in the price of driving is offset on average by a reduction in driving due to reduced consumer utility associated with driving a smaller vehicle with worse performance. If one is concerned with the effects of fuel economy standards such as CAFE, then estimating a rebound effect that holds vehicle characteristics constant can overstate the VMT response.

Our findings suggest that some optimism is in order regarding the effectiveness of fuel economy standards at reducing the externalities associated with gasoline con-

sumption. That is, even though the policy of choice by politicians is the second-best solution of increasing fuel economy standards through CAFE rather than increasing Pigouvian taxes on gasoline, there is little evidence to suggest that the reduction in gasoline consumption is offset by increased driving. One of the major shortcomings of fuel economy standards – the possibility of rebound – is negligible when accounting for changes in both the price of driving and vehicle characteristics that accompany fuel economy improvements.

4. SUMMARY AND CONCLUSIONS

This manuscript reflects two primary findings of importance for energy policy and for the automobile and transportation fuels sectors.

First, we find that the Cash for Clunkers stimulus program did increase revenues to the auto industry in the short run, but that the environmental component of the bill actually lowered new vehicle spending over the medium run by inducing people to buy more fuel efficient but less expensive cars. More generally, our findings suggest that in this particular case, environmental objectives undermined and even reversed the stimulus impact of the program.

Second, we show some optimism is in order regarding the effectiveness of fuel economy standards at reducing the externalities associated with gasoline consumption. That is, even though the policy of choice by politicians is the second-best solution of increasing fuel economy standards through CAFE rather than increasing Pigouvian taxes on gasoline, there is little evidence to suggest that the reduction in gasoline consumption is offset by increased driving. One of the major shortcomings of fuel economy standards – the possibility of rebound – is negligible when accounting for changes in both the price of driving and vehicle characteristics that accompany fuel economy improvements.

Taken together, the findings hold mixed implications for transportation energy policy. On the one hand, we find that fuel economy policy may conflict with other policy objectives, such as economic stimulus. However, we find that a common concern related to fuel economy standards – the so-called rebound effect – appears to be of minimal importance and seems unlikely to significantly undermine this type of energy policy.

REFERENCES

- Jerome Adda and Russell Cooper. Balladurette and Juppette: A discrete analysis of scrapping subsidies. *Journal of Political Economy*, 108(4):778–806, August 2000.
- Sumit Agarwal, Chunlin Liu, and Nicholas S. Souleles. The reaction of consumer spending and debt to tax rebates – Evidence from consumer credit data. *Journal of Political Economy*, 115(6):986–1019, December 2007.
- Anna Alberini, Winston Harrington, and Virginia. Estimating an emissions supply function for accelerated vehicle retirement programs. *Review of Economics and Statistics*, 78(2):251–265, May 1996.
- Hunt Allcott and Nathan Wozny. Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics*, Forthcoming.
- Joshua Angrist, Guido Imbens, and Donald Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- David Austin. Effects of gasoline prices on driving behavior and vehicle markets. Congressional Office Report 2883, January 2008.
- Antonio M. Bento, Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. Distributional and efficiency impacts of U.S. gasoline taxes. *American Economic Review*, 99(3):667–699, June 2009.
- Severin Borenstein. A microeconomic framework for evaluating energy efficiency rebound and some implications. Energy Institute at Haas Working Paper 242R, May 2013.
- Meghan R. Busse, Christopher R. Knittel, Jorge Silva-Risso, and Florian Zettelmeyer. Did “Cash for Clunkers” deliver? The consumer effects of the Car

- Allowance Rebate System. Working paper, November 2012.
- Raj Chetty, Adam Looney, and Kory Kroft. Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4):1145–1177, September 2009.
- Adam Copeland and James Kahn. The production impact of “Cash for Clunkers”: Implications for stabilization policy. *Economic Inquiry*, 51(1):288–303, January 2013.
- Council of Economic Advisers. Economic analysis of the Car Allowance Rebate System (“Cash for Clunkers”). Technical report, Executive Office of The President, September 2009.
- Janet Currie and Reed Walker. Traffic congestion and infant health: Evidence from E-Zpass. *American Economic Journal: Applied Economics*, 3(1):65–90, January 2011.
- James Feyrer and Bruce Sacerdote. Did the Stimulus stimulate? Real time estimates of the effects of the American Recovery and Reinvestment Act. NBER Working Paper No. 16759, February 2011.
- Carolyn Fischer, Winston Harrington, and Ian W. H. Parry. Should automobile fuel economy standards be tightened? *Energy Journal*, 28(4):1–29, December 2007.
- Don Fullerton and Li Gan. Cost-effective policies to reduce vehicle emissions. *American Economic Review Papers and Proceedings*, 95(2):300–304, May 2005.
- Kenneth Gillingham. Selection on anticipated driving and the consumer response to changing gasoline prices. Yale University Working paper, 2012.
- Kenneth Gillingham, Matthew J. Kotchen, David S. Rapson, and Gernot Wagner. Energy policy: The rebound effect is overplayed. *Nature*, 493:475–476, January 2013a.
- Kenneth Gillingham, David Rapson, and Gernot Wagner. The rebound effect and energy efficiency policy. Working paper, 2013b.

- Pinelopi Koujianou Goldberg. The effects of the Corporate Average Fuel Efficiency standards in the U.S. *Journal of Industrial Economics*, 46(1):1–33, March 1998.
- Lorna A. Greening, David L. Greene, and Carmen Difiglio. Energy efficiency and consumption, the rebound effect: A survey. *Energy Policy*, 28(6-7):389–401, June 2000.
- Robert W. Hahn. An economic analysis of scrappage. *RAND Journal of Economics*, 26(2):222–242, Summer 1995.
- Mark Hoekstra, Steven L. Puller, and Jeremy West. Cash for Corollas: When stimulus reduces spending. Texas A&M University Working Paper, June 2014.
- Christopher L. House and Matthew D. Shapiro. Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, 98(3):737–768, June 2008.
- Kent M. Hymel and Kenneth A. Small. The rebound effect for automobile travel: Asymmetric response to price changes and other quirks of the 2000s. University of California at Irvine Working Paper, 2013.
- Guido W. Imbens and Thomas Lemieux. Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635, February 2008.
- Interagency Working Group. Technical update on the social cost of carbon for regulatory impact analysis - under executive order 12866. Technical report, Interagency Working Group on the Social Cost of Carbon, May 2013.
- Mark R. Jacobsen. Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy*, 5(2):148–187, May 2013.
- Mark R. Jacobsen and Arthur A. van Benthem. Vehicle scrappage and gasoline policy. University of California at San Diego Working Paper, March 2013.
- David S. Johnson, Jonathan A. Parker, and Nicholas S. Souleles. Household expen-

- diture and the income tax rebates of 2001. *American Economic Review*, 96(5): 1589–1610, December 2006.
- Thomas Klier and Joshua Linn. New-vehicle characteristics and the cost of the Corporate Average Fuel Economy standard. *RAND Journal of Economics*, 43(1): 186–213, Spring 2012.
- Christopher R. Knittel. The implied cost of carbon dioxide under the Cash for Clunkers program. Center for the Study of Energy Markets Working Paper, August 2009.
- Christopher R. Knittel. Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review*, 101(7):3368–3399, December 2011.
- Christopher R. Knittel. Transportation fuels policy since the OPEC embargo: Paved with good intentions. *American Economic Review Papers and Proceedings*, 103(3): 344–349, May 2013.
- Christopher R. Knittel and Ryan Sandler. Cleaning the bathwater with the baby: The health co-benefits of carbon pricing in transportation. NBER Working Paper No. 17390, November 2011.
- Christopher R. Knittel and Ryan Sandler. The welfare impact of indirect pigouvian taxation: Evidence from transportation. NBER Working Paper No. 18849, February 2013.
- Christopher R. Knittel, Douglas Miller, and Nicholas J. Sanders. Caution, drivers! children present. traffic, pollution, and infant health. NBER Working Paper No. 17222, 2011.
- Richard P. Larrick and Jack B. Soll. The MPG illusion. *Science*, 320(5883):1593–1594, June 2008.
- David S. Lee and David Card. Regression discontinuity inference with specification

- error. *Journal of Econometrics*, 142(2):655–674, February 2008.
- Shanjun Li and Chao Wei. Green stimulus: A dynamic discrete analysis of vehicle scrappage programs. Working paper, April 2013.
- Shanjun Li, Christopher Timmins, and Roger H. von Haefen. How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy*, 1(2): 113–137, August 2009.
- Shanjun Li, Joshua Linn, and Erich Muelegger. Gasoline taxes and consumer behavior. NBER Working Paper No. 17891, March 2012.
- Shanjun Li, Joshua Linn, and Elisheba Spiller. Evaluating Cash-for-Clunkers: Program effects on auto sales and the environment. *Journal of Environmental Economics and Management*, 65(2):165–193, March 2013.
- Joshua Linn. The rebound effect for passenger vehicles. RFF Discussion Paper 13-19, November 2013.
- Fred Mannering and Clifford Winston. A dynamic empirical analysis of household vehicle ownership and utilization. *RAND Journal of Economics*, 16(2):215–236, Summer 1985.
- Virginia McConnell. The new cafe standards: Are they enough on their own? RFF Discussion Paper 13-14, May 2013.
- Atif Mian and Amir Sufi. The effects of fiscal stimulus: Evidence from the 2009 “Cash for Clunkers” program. *Quarterly Journal of Economics*, 127(3):1107–1142, August 2012.
- Eugenio J. Miravete and Maria Moral. Qualitative effects of ‘cash for clunkers’ programs. University of Texas Working Paper, October 2011.
- National Highway Traffic Safety Administration. CARS report to Congress. Technical report, U.S. Department of Transportation, December 2009.
- National Research Council. *Transitions to Alternative Vehicles and Fuels*. The

- National Academies Press, 2013.
- Jonathan A. Parker, Nicholas S. Souleles, David S. Johnson, and Robert McClelland. Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–2553, October 2013.
- Paul R. Portney, Ian W.H. Parry, Howard K. Gruenspect, and Winston Harrington. The economics of fuel economy standards. *Journal of Economic Perspectives*, 17(4):203–217, Fall 2003.
- Christina Romer and Christopher Carroll. Did ‘Cash-for-Clunkers’ work as intended? White House commentary, Council for Economic Advisers, April 2010.
- Ryan Sandler. Clunkers or junkers? adverse selection in a vehicle retirement program. *American Economic Journal: Economic Policy*, 4(4):253–281, November 2012.
- Matthew D. Shapiro and Joel Slemrod. Consumer response to tax rebates. *American Economic Review*, 93(1):381–396, March 2003.
- Kenneth A. Small and Kurt van Dender. Fuel efficiency and motor vehicle travel: The declining rebound effect. *Energy Journal*, 28(1):25–51, 2007.
- U.S. Department of Transportation. National Household Travel Survey. Federal Highway Administration, 2009.
- U.S. Department of Transportation. Final regulatory impact analysis: Corporate Average Fuel Economy for MY 2012 - MY 2016 passenger cars and light trucks. Technical report, U.S. Department of Transportation National Highway Traffic Safety Administration, March 2010.
- Jeremy West, Mark Hoekstra, Jonathan Meer, and Steven L. Puller. How do households respond to increases in fuel economy? Regression discontinuity evidence. Texas A&M University Working Paper, June 2014.
- Sarah E. West. Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics*, 88(3-4):735–757, March 2004.

Daniel J. Wilson. Fiscal spending jobs multipliers: Evidence from the 2009 American Recovery and Reinvestment Act. *American Economic Journal: Economic Policy*, 4(3):251–282, August 2012.

APPENDIX A

DATA APPENDIX

A.1 Defining a Household's Fleet

The Texas Department of Motor Vehicles (DMV) provided us with confidential access to all Texas vehicle registrations for the years spanning our study. From these records, we attribute individual vehicles to households as follows. First, we used ESRI's ArcMAP software to geocode the population of entered registration addresses to the North American Address Locator database. Of importance, this process additionally returns the standardized postal address for each specific matched location, thereby correcting for database entry errors. For these standardized addresses, we drop records at any address to which more than 700 unique vehicles (VIN17) were registered within a single calendar year, as these are almost exclusively commercial or institutional registrants. For similar reasons, we drop records for which the last name consists of some variation of a commercial, industrial, or other non-household registrant (e.g. corporation, association, dealer, school, etc.). We drop another roughly one percent of DMV records for the following reasons: (1) we could not match the record to a standardized postal address; (2) the record is missing a sale date; or (3) the record is missing a last name in both last name fields. Finally, we drop records for non-consumer vehicle identification numbers that are not included in EPA fuel economy data.

We attribute a pair of vehicles to the same household if either of the following sets of conditions are met: (1) the pair of vehicles is sequentially and jointly registered at multiple locations (i.e. a household moves to a new address); or (2) the pair of vehicles

is registered at the same address to the same “fuzzy” last name.¹ After determining pairs of vehicles belonging to the same household, we chain these connections to allocate the population of vehicles to households for each date included in our data.

Because DMV registrations are better suited for tracking vehicle purchases than exits from a household’s fleet, we make two additional adjustments to households’ duration of vehicle ownership. We remove a vehicle from a household’s fleet if the latest observed registration (in Texas) has lapsed by six months. And, because car dealerships often do not appear in the same DMV registration database as households, we backdate a vehicle’s end date for a household if: (1) the vehicle is later sold by a used car dealership, and (2) the former registered household purchased a new vehicle within six months preceding this sale date. This treats the former registrant’s new vehicle purchase transaction date as a trade-in date for the used vehicle.

A.2 Calculating Household VMT

We calculate vehicle miles traveled for each unique vehicle (VIN17) using three sources of odometer readings. Primarily, we use data from vehicle emissions tests conducted in the seventeen EPA non-attainment counties in Texas, which were provided to us by the Texas Commission on Environmental Quality (TCEQ) for January 1, 2004 through August 20, 2012. In these counties, Texas law requires personal vehicles to undergo emissions testing annually beginning at the vehicle’s second year.² New residents are allowed thirty days to obtain a vehicle emissions test. We augment

¹We use a dynamic Levenshtein distance metric to match last names. First, we trim each of the two last name fields to fifteen letters. Then, we match them pairwise using a Levenshtein critical value of 0.34. The most common entry errors for names in the database are omitted letters (an L-distance of one) and transposed letters (an L-distance of two). For a six letter last name, an L-distance of two requires a critical value of 0.34 to correct. A nine letter last name is allowed three transformations under this critical value.

²The annual emissions inspection requirement is waived for vehicles older than twenty-four years. More information on Texas emissions testing requirements is provided by the Texas Department of Public Safety at <http://www.txdps.state.tx.us/InternetForms/Forms/VI-51.pdf>

these odometer readings with data from the Texas DMV database, which reports the odometer value for each vehicle transaction involving a Texas buyer. Finally, for a fairly small set of vehicles we append odometer readings reported to the U.S. DOT for vehicles scrapped in the Car Allowance Rebate System.³

We determine the temporal duration and total VMT between each sequential pair of odometer readings for each VIN. As many of the odometer readings were at some point manually entered into a database, we attempt to correct for entry errors using several types of adjustments: (1) multiply the reported odometer value by ten; (2) divide the reported odometer value by ten; (3) drop the leading digit of the reported odometer value; (4) subtract one from the leading digit of the reported odometer value; or (5) leave the reported odometer unadjusted. We allow for the adjustment to be made to either the first or the second reading in every sequential pair of odometer values. As a selection metric, for each possible transformation we iteratively compute the equally-weighted average of the absolute value differences between the previous and current, and the current and following readings. In essence, this metric seeks the smoothest path within each set of consecutive three readings. Following this, we drop approximately three percent of remaining readings that imply negative VMT or a daily VMT of less than one or greater than 700. Additionally, at this point we drop readings of fewer than fifty miles apart (which are likely retests of failed inspections) and vehicles for which we observe only a single odometer reading.

Matching on VIN to the DMV registration database, we aggregate vehicle-level VMT to an annual household-level based on vehicle ownership. Within each calendar year, we sum total observed miles driven by the household as well as total days of observed VMT. From these, we compute the average daily VMT per-vehicle for each household for each calendar year. Then, we multiply this value by the total number

³The CARS data are available from the National Highway Safety Traffic Safety Administration.

of “vehicle ownership days” for the household over the calendar year, thus measuring total household-level annual VMT.⁴

Similarly, we calculate the quantity of gasoline consumed by each household within each calendar year. For each vehicle owned by a household, we divide the total observed miles driven in that vehicle within a year by the vehicle’s EPA rating for combined fuel economy. Then, we divide this value by the number of observed days of VMT within the year to obtain the gallons consumed per day for each vehicle. We multiply this by the number of days in the year for which the household owned the vehicle, and sum across the household’s set of vehicles to determine total gallons consumed per year.

⁴This approach does extrapolate VMT for some vehicles, but the nature of this calculation restricts extrapolation to within a calendar year. In light of non-compliance, households moving out of emissions testing counties, and other factors precluding odometer observations, we view this as a reasonable trade-off. The overall fraction of VMT determined using such extrapolation is relatively small and is mostly concentrated in the book-ending years of our data, which are not included in our empirical study.

APPENDIX B

FIGURES AND TABLES

B.1 Figures and Tables for Section 2



Figure B.1: Monthly new vehicle sales, annual rate (source: NADA)






Figure B.2: Relationship between fuel economy and vehicle price (MSRP)




(a) During July 2009 - August 2009 (Cash for Clunkers)




(b) During July 2009 - April 2010 (10 months)


Figure B.3: First-stage: Probability of purchase being subsidized by CfC




(a) July 2009 - August 2009 (Cash for Clunkers)




(b) July 2009 - January 2010 (7 months)




(c) July 2009 - February 2010 (8 months)



(d) July 2009 - March 2010 (9 months)




(e) July 2009 - April 2010 (10 months)




(f) July 2009 - May 2010 (11 months)

Figure B.4: Cumulative fraction of households purchasing a new vehicle by period




(a) New vehicle fuel economy (MPG)




(b) New vehicle price (transaction)

Figure B.5: Purchases during July 2009 - August 2009 (Cash for Clunkers)




(a) New vehicle fuel economy (MPG)




(b) New vehicle price (transaction)


Figure B.6: Purchases during July 2009 - April 2010 (10 months)




(a) Number of adults in home



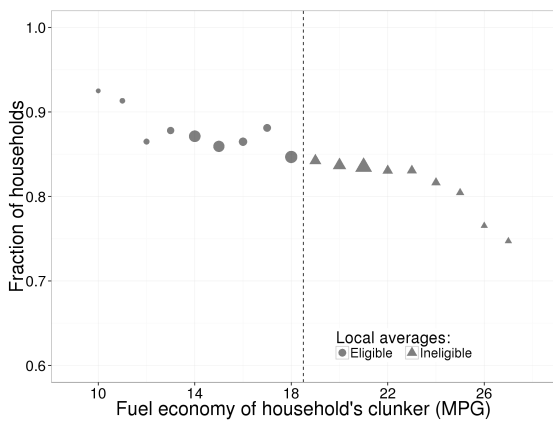
(b) Weekly travel days




(c) Log of annual household income



(d) Live in an urban area (%)

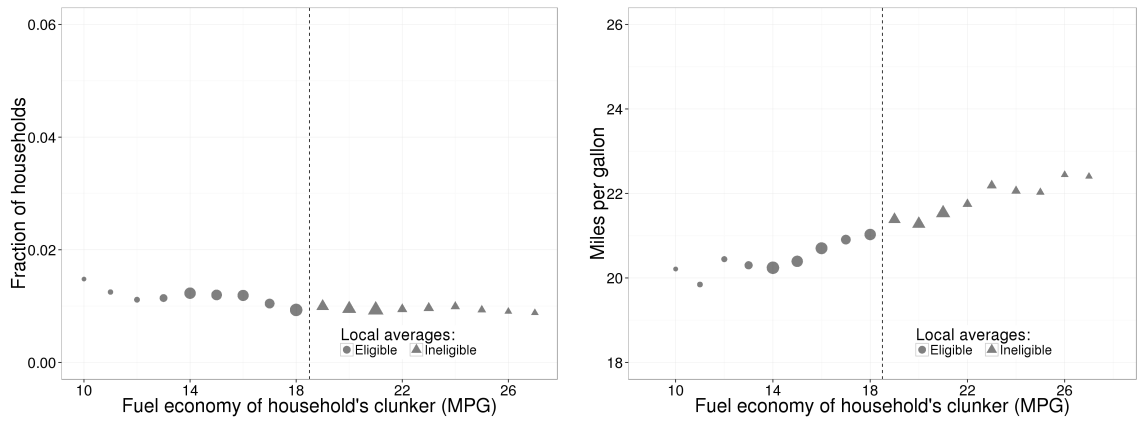


(e) Live in a house (%)




(f) White (%)

Figure B.7: Identification checks: National Household Travel Survey (spring 2009)




(a) Purchased any new vehicle (July-Aug. 2008) (b) Fuel economy of purchases (July-Aug. 2008)




(c) Price of purchases (July-Aug. 2008)


Figure B.8: Identification checks: Households purchasing vehicles prior to CfC



(a) Non-clunker fleet fuel economy (June 2009)



(b) Non-clunker fleet MSRP (June 2009)



(c) Number of vehicles owned (June 2009)

Figure B.9: Identification checks: Fleet characteristics for households purchasing during July 2009 - April 2010 (10 months)

Table B.1: Summary statistics for new vehicle purchases July 2009 - April 2010

	Median	Mean	Std. Dev.
Total number of households		5,523,973	
Sample: purchased new vehicle			
Number of households		239,559	
Fraction of households		0.041	
Characteristics of new vehicles			
Fuel economy (MPG)	21	21.7	5.887
Transaction price (\$ '000s)	25.5	27.6	11.01
MSRP (\$ '000s)	25.6	27.5	10.33
Census Tract characteristics			
Population	5820	6390.7	3089.9
Median age	34.1	34.5	5.071
White (%)	82.3	77.6	16.59
Black (%)	3.80	7.99	12.57
Asian (%)	1.40	3.12	4.670
Hispanic (%)	14.0	24.6	25.24
Household size	2.85	2.83	0.423
Housing units	2214	2404.9	1099.7
Owner-occupied (%)	79.2	74.3	17.82
Median Income (\$ '000s)	48.7	54.2	24.77
Median Home value (\$ '000s)	93.2	112.2	74.29

Notes: Statistics reported for Texas households that purchased a new vehicle either during Cash for Clunkers or during the subsequent eight months (from July 2009 through April 2010 in total). The Census Tract-level characteristics are from the 2000 Decennial Census.

Table B.2: Discontinuities in vehicle characteristics for purchases during July 2009 - April 2010 (10 months)

Bandwidth	9 MPG		7 MPG		5 MPG		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: July 2009-August 2009 (Cash for Clunkers Program Months)							
Fuel economy (MPG)	1.647*** (0.138)	1.623*** (0.124)	1.495*** (0.161)	1.393*** (0.116)	1.591*** (0.119)	1.478*** (0.114)	1.290*** (0.180)
Price (transacted)	-2910.2*** (218.1)	-2815.1*** (194.9)	-2758.9*** (160.1)	-2575.9*** (169.6)	-2798.2*** (131.2)	-2688.6*** (119.5)	-2668.0*** (91.70)
Price (MSRP)	-2630.9*** (196.6)	-2556.6*** (171.3)	-2627.7*** (147.2)	-2291.2*** (127.1)	-2642.7*** (150.9)	-2475.1*** (105.1)	-2529.2*** (81.67)
Panel B: July 2009-April 2010 (10 months - Cash for Clunkers + 8 Subsequent Months)							
Fuel economy (MPG)	0.989*** (0.130)	0.973*** (0.118)	0.821*** (0.123)	0.679*** (0.0574)	0.964*** (0.113)	0.808*** (0.0734)	0.624** (0.131)
Price (transacted)	-2191.5*** (278.5)	-1874.5*** (214.9)	-1855.8*** (146.5)	-1395.4*** (127.3)	-1987.6*** (224.3)	-1656.8*** (39.96)	-1776.3*** (129.2)
Price (MSRP)	-2000.6*** (261.4)	-1710.7*** (203.7)	-1739.8*** (151.9)	-1250.9*** (97.32)	-1858.6*** (235.3)	-1532.5*** (52.13)	-1647.1*** (136.1)
Polynomial Controls	Cubic No 72,461	Cubic Yes 72,461	Quadratic Yes 72,461	Cubic Yes 69,507	Quadratic Yes 69,507	Quadratic Yes 60,492	Linear Yes 60,492
Observations (Panel A)	72,461	72,461	72,461	69,507	69,507	60,492	60,492
Observations (Panel B)	239,559	239,559	239,559	228,098	228,098	198,144	198,144

Notes: * $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$ Each coefficient represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility, which is β_3 in equation (2.1). Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses. Figures B.5 and B.6 present graphs corresponding to these results.

Table B.3: Robustness of estimated discontinuities to alternate time windows

	Reduced-form	First-stage	LATE	Observations
Fuel economy (MPG)				
9 months	1.061	0.398	2.666	217,548
10 months (main)	0.973	0.373	2.609	239,559
11 months	0.912	0.347	2.628	264,152
12 months	0.831	0.325	2.557	287,887
13 months	0.795	0.307	2.590	311,172
14 months	0.748	0.288	2.597	337,918
Price (transacted)				
9 months	-2058.4	0.398	-5171.9	217,548
10 months (main)	-1874.5	0.373	-5025.5	239,559
11 months	-1728.6	0.347	-4981.6	264,152
12 months	-1591.3	0.325	-4896.3	287,887
13 months	-1580.8	0.307	-5149.2	311,172
14 months	-1453.4	0.288	-5046.5	337,918
Price (MSRP)				
9 months	-1882.7	0.398	-4730.4	217,548
10 months (main)	-1710.7	0.373	-4586.3	239,559
11 months	-1588.1	0.347	-4576.7	264,152
12 months	-1472.9	0.325	-4532.0	287,887
13 months	-1461.5	0.307	-4760.6	311,172
14 months	-1359.4	0.288	-4720.1	337,918
Bandwidth (MPG)				
	9	9	9	9
Polynomial				
	Cubic	Cubic	Cubic	Cubic
Controls				
	Yes	Yes	Yes	Yes

Notes: See Hoekstra et al. [2014] for table notes.

B.2 Figures and Tables for Section 3



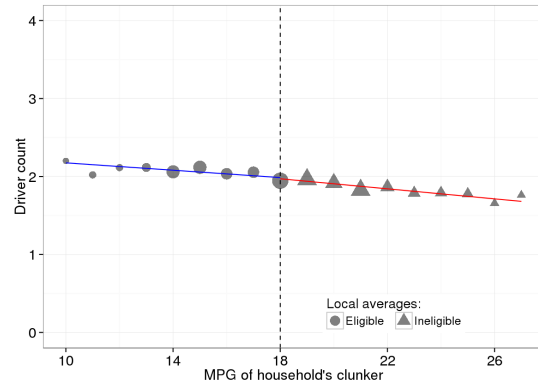



Figure B.10: Texas counties included in this study (EPA non-attainment)




(a) Number of vehicles




(b) Number of drivers



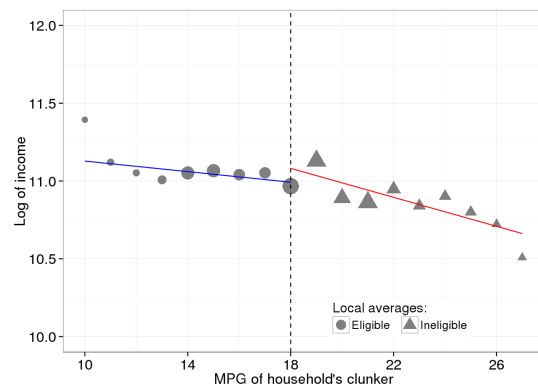
(c) Number of working members



(d) Weekly travel days

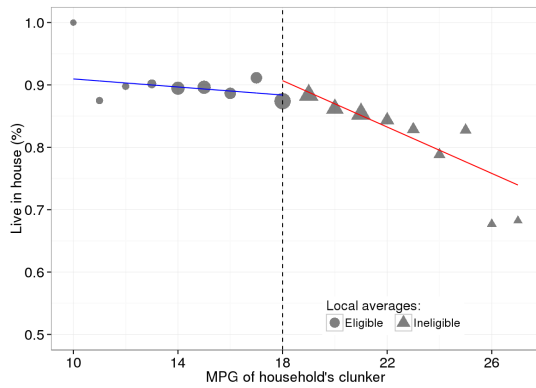


(e) Number of adults




(f) Log of household income


Figure B.11: Test of identification strategy: Household covariates reported in the National Household Travel Survey 2009




(a) Live in house




(b) Live in urban area




(c) Percent White



(d) Percent Black




(e) Percent Asian




(f) Percent Hispanic


Figure B.12: Test of identification strategy:
Additional household covariates in the National Household Travel Survey 2009



(a) Log of MPG




(b) Log of VMT




(c) Log of gallons used


Figure B.13: Falsification test: Household-level results in 2008



(a) Recently moved




(b) Financed CfC-time purchase




(c) Change in new vehicle financing from 2008 stock


Figure B.14: Falsification test: Indicators of possible economic shocks in 2009




(a) Horsepower



(b) Horsepower/Lb.-1000s




(c) Cubic volume




(d) MSRP (\$000s)


Figure B.15: Selected characteristics of vehicles that were purchased new during CfC in 2009




(a) Log of MPG



(b) Log of VMT




(c) Log of gallons used




(d) Fraction of household's VMT


Figure B.16: Main results: 2010 outcomes for the new vehicles purchased



(a) Log of MPG



(b) Log of VMT



(c) Log of gallons used

Figure B.17: Main results: 2010 outcomes at the household-level

Table B.4: Summary statistics in 2010: Contrasting vehicle owners to purchasers

	All vehicle owners		Empirical sample	
	Mean	Std. Dev.	Mean	Std. Dev.
<i>Vehicles purchased during CfC</i>				
Miles per gallon			22.4	6.077
VMT ('000s)			13.8	8.365
Gasoline consumed (gal)			667.9	480.4
<i>All of household's vehicles</i>				
Vehicle count	1.74	1.054	2.66	1.235
Average MPG	19.6	4.273	20.9	4.114
VMT ('000s)	18.5	18.66	30.5	20.71
Gasoline consumed (gal)	1008.0	1096.1	1573.4	1149.9
<i>Census 2000 Tract-level</i>				
Total population	6415.2	3133.8	6752.2	3326.4
Population per km ²	1115.8	1134.5	997.5	941.0
Number of housing units	2393.7	1134.1	2480.3	1173.2
Average household size	2.85	0.463	2.87	0.411
Median income (\$000s)	54.4	23.72	62.3	25.35
Median home value (\$000s)	114.1	73.48	130.4	81.61
Owner-occupied (%)	69.8	22.19	75.1	19.51
Married (%)	59.4	11.78	62.7	10.85
Older than 62 (%)	9.87	5.840	9.49	5.788
Bachelor's degree (%)	29.0	19.14	33.9	19.47
White (%)	72.3	20.57	76.4	18.08
Black (%)	11.3	16.87	8.82	13.66
Asian (%)	3.80	5.023	4.34	5.448
Hispanic (%)	24.1	23.13	20.1	20.80
Below poverty line (%)	8.35	8.457	6.24	6.846
Receive public assistance (%)	2.19	2.524	1.70	2.058
Household observations	4,620,943		67,824	

Notes: See West et al. [2014] for table notes.

Table B.5: Test of identification strategy:
Regression discontinuity estimates for households in NHTS-2009

Bandwidth	9 MPG	8 MPG	7 MPG	6 MPG	5 MPG
Vehicle count	0.0510 (0.0697)	0.0382 (0.0684)	0.0148 (0.0622)	0.00579 (0.0646)	-0.0267 (0.0598)
Driver count	0.0155 (0.0380)	0.00965 (0.0377)	0.00508 (0.0347)	-0.00753 (0.0321)	-0.0217 (0.0295)
Worker count	0.00983 (0.0773)	0.0128 (0.0806)	0.0122 (0.0847)	0.0158 (0.0895)	-0.0319 (0.0782)
Weekly travel days	0.0198 (0.0881)	0.00182 (0.0855)	-0.0599 (0.0666)	-0.0983 (0.0619)	-0.136*** (0.0411)
Number of adults	0.0259 (0.0374)	0.0191 (0.0369)	0.0153 (0.0375)	-0.00173 (0.0342)	-0.0228 (0.0287)
Log of income	-0.0888 (0.0787)	-0.0727 (0.0852)	-0.0699 (0.0908)	-0.0787 (0.0990)	-0.123 (0.0931)
Live in house (%)	-0.0231 (0.0139)	-0.0172 (0.0148)	-0.00893 (0.0116)	-0.0167 (0.0104)	-0.0108 (0.0104)
Live in urban area (%)	-0.00264 (0.0180)	-0.000640 (0.0166)	-0.000193 (0.0175)	-0.000623 (0.0196)	0.00103 (0.0250)
White (%)	-0.0191 (0.0153)	-0.0214 (0.0157)	-0.0229 (0.0170)	-0.0356* (0.0188)	-0.0419 (0.0235)
Black (%)	0.00404 (0.0134)	0.00565 (0.0128)	0.00562 (0.0132)	0.0128 (0.0124)	0.0228** (0.00710)
Asian (%)	0.00646 (0.00875)	0.00662 (0.00901)	0.00910 (0.0102)	0.0160 (0.0128)	0.0174 (0.0165)
Hispanic (%)	0.0172** (0.00639)	0.0191*** (0.00646)	0.0200*** (0.00633)	0.0219*** (0.00697)	0.0189** (0.00740)
Polynomial	Linear	Linear	Linear	Linear	Linear
Observations	6335	6252	6060	5763	5303

Notes: * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ See West et al. [2014] for table notes.

Table B.6: Falsification exercise: Regression discontinuity estimates for household outcomes in 2008

Bandwidth	9 MPG			8 MPG		7 MPG	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Household-level</i>							
Log of MPG	0.0093** (0.004)	0.0076** (0.003)	0.0113*** (0.003)	0.0106*** (0.003)	0.0072* (0.004)	0.0104*** (0.003)	0.0106*** (0.003)
Log of total VMT	0.0575** (0.027)	0.0483* (0.025)	0.0543* (0.029)	0.0492* (0.028)	0.0535* (0.028)	0.0489 (0.028)	0.0486 (0.029)
Log of total gals.	0.0489* (0.026)	0.0413 (0.024)	0.0408 (0.026)	0.0363 (0.025)	0.0469 (0.027)	0.0365 (0.026)	0.0371 (0.028)
Polynomial	Quadratic	Quadratic	Linear	Linear	Quadratic	Linear	Linear
Includes controls	No	Yes	No	Yes	Yes	Yes	Yes
Observations	66,301	66,280	66,301	66,280	65,526	65,526	63,974

Notes: * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Each coefficient represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility. Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses. The outcomes are defined over 2008, the year *prior to* that of the CARS program. “Controls” include county of residence fixed effects as well as the full set of Census 2000 Tract-level covariates shown in Table B.4. Figure B.13 presents graphs corresponding to these results.

Table B.7: Falsification exercise: Regression discontinuity estimates for indicators of possible economic shocks in 2009

Bandwidth	9 MPG			8 MPG		7 MPG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Household-level</i>						
Recently moved	0.0028 (0.003)	0.0036 (0.003)	0.0012 (0.003)	0.0016 (0.002)	0.0061 (0.004)	0.0012 (0.003)
New vehicles financed						
CfC-time purchase	0.0155 (0.009)	0.0020 (0.008)	0.0211 (0.013)	0.0084 (0.010)	0.0002 (0.010)	0.0086 (0.010)
Δ from 2008 fleet	-0.0219** (0.009)	-0.0225** (0.009)	-0.0160** (0.006)	-0.0171*** (0.006)	-0.0279** (0.011)	-0.0161** (0.006)
Polynomial	Quadratic	Quadratic	Linear	Linear	Quadratic	Linear
Includes controls	No	Yes	No	Yes	Yes	Yes
Observations	67,491	67,465	67,491	67,465	66,683	65,091
Observations	67,491	67,465	67,491	67,465	66,683	65,091
Observations	47,615	47,596	47,615	47,596	47,057	45,866

Notes: * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Each coefficient represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility. Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses. “Recently moved” is a binary indicator for whether the household moved to a new Census tract between January 1st and June 30th of 2009. Financing for the new vehicle in 2009 is an indicator for whether the household used any financing for the purchase made during the CARS time period in 2009. From this, we subtract the fraction of the (originally purchased by the household as new) vehicles in the household’s fleet as of 2008 that were purchased using any financing. This forms the “difference from the 2008 fleet” financing proportion. “Controls” include county of residence fixed effects as well as the full set of Census 2000 Tract-level covariates shown in Table B.4. Figure B.13 presents graphs corresponding to these results.

Table B.8: Regression discontinuity estimates for household outcomes in 2010

Bandwidth	9 MPG			8 MPG		7 MPG	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>New vehicles</i>							
Log of MPG	0.0618*** (0.010)	0.0587*** (0.009)	0.0413*** (0.011)	0.0409*** (0.010)	0.0570*** (0.008)	0.0427*** (0.011)	0.0438*** (0.010)
Log of VMT	-0.0057 (0.010)	-0.0056 (0.008)	-0.0058 (0.020)	-0.0062 (0.018)	-0.0016 (0.012)	-0.0056 (0.019)	0.0006 (0.016)
Log of gallons	-0.0675*** (0.010)	-0.0643*** (0.011)	-0.0471 (0.029)	-0.0472 (0.027)	-0.0587*** (0.013)	-0.0483 (0.028)	-0.0432 (0.026)
Fraction of VMT	0.0140 (0.011)	0.0139 (0.011)	-0.0005 (0.010)	-0.0016 (0.010)	0.0144 (0.011)	0.0001 (0.010)	0.0010 (0.011)
<i>Household-level</i>							
Log of MPG	0.0457*** (0.006)	0.0433*** (0.006)	0.0290*** (0.008)	0.0286*** (0.007)	0.0433*** (0.005)	0.0297*** (0.008)	0.0314*** (0.007)
Log of total VMT	-0.0334 (0.031)	-0.0341 (0.026)	-0.0017 (0.032)	-0.0014 (0.029)	-0.0309 (0.030)	-0.0048 (0.029)	-0.0008 (0.031)
Log of total gals.	-0.0814*** (0.027)	-0.0795*** (0.023)	-0.0323 (0.037)	-0.0315 (0.034)	-0.0757** (0.027)	-0.0362 (0.034)	-0.0331 (0.035)
Polynomial	Quadratic	Quadratic	Linear	Linear	Quadratic	Linear	Linear
Includes controls	No	Yes	No	Yes	Yes	Yes	Yes
Observations	67,824	67,797	67,824	67,797	67,011	67,011	65,410

Notes: * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Each coefficient represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility. Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses. The outcomes are defined over 2010, the year following that of the CARS program. "Controls" include county of residence fixed effects as well as the full set of Census 2000 Tract-level covariates shown in Table B.4. Figures B.16 and B.17 present graphs corresponding to these results.