



MIT Center for  
Energy and Environmental  
Policy Research

## Cash for Corollas: When Stimulus Reduces Spending

Mark Hoekstra, Steven L. Puller, and  
Jeremy West

April 2015

CEEPR WP 2015-005

# Cash for Corollas: When Stimulus Reduces Spending

Mark Hoekstra

Steven L. Puller

Jeremy West \*

April 2015

## Abstract

The 2009 Cash for Clunkers program aimed to stimulate consumer spending in the new automobile industry, which was experiencing disproportionate reductions in demand and employment during the Great Recession. Exploiting program eligibility criteria in a regression discontinuity design, we show nearly 60 percent of the subsidies went to households who would have purchased during the two-month program anyway; the rest accelerated sales by no more than eight months. Moreover, the program's fuel efficiency restrictions shifted purchases toward vehicles that cost on average \$5,000 less. On net, Cash for Clunkers significantly reduced total new vehicle spending over the ten month period.

JEL: H20, Q40, Q50

Keywords: Fiscal stimulus, multifaceted policy, regression discontinuity

---

\*Hoekstra: Texas A&M University and NBER, [mhoekstra@econmail.tamu.edu](mailto:mhoekstra@econmail.tamu.edu). Puller: Texas A&M University and NBER, [puller@econmail.tamu.edu](mailto:puller@econmail.tamu.edu). West (corresponding author): Massachusetts Institute of Technology, [westj@mit.edu](mailto:westj@mit.edu). We are grateful to Pierre Mouganie for excellent research assistance and to Daron Acemoglu, David Autor, John Cochrane, Amy Finkelstein, Jonathan Gruber, Catie Hausman, Chris Knittel, Jason Lindo, Joshua Linn, Erzo F.P. Luttmer, Jonathan Meer, Gregor Pfeifer, Jim Poterba, Matt Zaragoza-Watkins, Sarah Zubairy, and seminar participants at the Institute for the Study of Labor (IZA), the International Industrial Organization Conference, Louisiana State University, MIT, Texas Stata Empirical Microeconomics Conference, University of Florida, University of Melbourne, University of Michigan, and University of Texas for helpful comments. We gratefully acknowledge financial funding from the National Science Foundation EV-STS. Any errors are our own.

# 1 Introduction

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, better known as “Cash for Clunkers” (CfC), differs from these more general 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. In fact, stimulus programs specifically targeting the auto industry are commonplace around the world: more than 15 countries implemented programs similar to CfC in response to the Great Recession [Haugh et al., 2010].

A major objective of the CfC 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]. The motivation for targeting stimulus towards durable goods manufacturing is straightforward: in the previous 7 recessions dating back to 1960, that sector lost a greater share of jobs (10.6 percent) than any other sector, and on average was responsible for 46 percent of total job losses despite currently making up less than 10 percent of overall employment [Leamer, 2009].

However, another priority for the Obama Administration is to improve the fuel efficiency of the U.S. vehicle fleet, as echoed in President Obama’s 2009 statement that, “Ending our dependence on oil, indeed, ending our dependence on fossil fuels, represents perhaps the most difficult challenge we have ever faced...”<sup>1</sup> As a result, the CfC policy was written to achieve multiple goals: to accelerate the purchase of new vehicles to increase revenues to the auto industry, and to increase the fuel efficiency of the fleet by requiring new vehicles purchased under the program to have sufficiently high fuel economy.

The fuel efficiency restrictions imposed by the program could have either enhanced or undermined the stimulus effect of the policy. On 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

---

<sup>1</sup>Remarks by the President on National Fuel Efficiency Standards made May 19, 2009, retrieved from [www.whitehouse.gov](http://www.whitehouse.gov).

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 CfC 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 CfC, 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.<sup>2</sup>

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 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 Cash for Clunkers significantly increased new car purchases during the two months of the program; barely eligible households were significantly more likely to purchase a new vehicle during that time compared to barely ineligible households. However, this increase was modest relative to the size of the program, as we estimate that nearly 60 percent of the households who purchased under the program would have bought a new vehicle during those two months anyway. Thus, while the program did ac-

---

<sup>2</sup>We focus only on identifying the stimulus impact for the U.S. automobile 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.

celerate some purchases, the majority of buyers under the program would have purchased during July or August even absent the program.

In addition, our estimates indicate that 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. Thus, the roughly 40 percent of subsidized purchases that were pulled forward from the future were accelerated by a maximum of 7 to 9 months. Put differently, 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 Council of Economic Advisors (CEA) or the three years assumed by the National Highway Traffic Safety Administration (NHTSA).<sup>3</sup>

However, perhaps most importantly, 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 automobile 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.<sup>4</sup> 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.<sup>5</sup>

---

<sup>3</sup>See [Council of Economic Advisers, 2009](#) and [NHTSA, 2009](#).

<sup>4</sup>Our 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.

<sup>5</sup>Put 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,

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. The barely eligible households tended to purchase less expensive and smaller vehicles such as the Toyota Corolla, which was the most popular new vehicle purchased under the program. Estimates indicate that each household purchasing under the program spent an average of around \$5,000 less on a new vehicle than they otherwise would have. Thus, while the program increased short-term spending over the period of a few months, that increase came at the cost of reducing cumulative longer-term spending. We calculate that justifying this tradeoff would require an annualized discount rate of 208%. A back-of-the-envelope calculation suggests that Cash for Clunkers – 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 to four 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.

## 2 Background and Empirical Strategy

### 2.1 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.<sup>6</sup> 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, and 43,000 of these purchases were by households in Texas.

---

not the type of household who purchased.

<sup>6</sup>Scrappage policies have been implemented in numerous countries internationally; studies of these programs include Hahn [1995], Alberini, Harrington, and McConnell [1996], Adda and Cooper [2000], Miravete and Moral [2011], Sandler [2012], Busse, Knittel, Silva-Risso, and Zettelmeyer [2012], and Klößner and Pfeifer [2015]. More generally, the literature has investigated the determinants of scrappage decisions, including the effect of gasoline prices and used car resale value (Li, Timmins, and von Haefen, 2009; Jacobsen and van Benthem, Forthcoming).

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 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. Restrictions were placed on both the vehicle being traded in and the vehicle being purchased. The restriction on the trade-in vehicle 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.<sup>7</sup> 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.<sup>8</sup> 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 over the 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

---

<sup>7</sup>There 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.

<sup>8</sup>The 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](#)).

to have a combined EPA fuel economy rating of at least twenty-two miles per gallon in order to receive the subsidy. The amount of the subsidy depended on the difference in fuel economy between the new passenger car and the scrapped clunker. If the difference 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 yielded a \$4500 rebate.<sup>9</sup>

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. In 2009, manufacturers offered hybrid versions of several models with prices about five thousand dollars higher than the standard versions. However, these examples are more the exception than the rule. The general trend is that higher fuel economy vehicles have lower prices. This relationship is shown in Appendix Figure B.1 which shows a negative overall relationship between MPG and vehicle price among the set of vehicles offered to U.S. consumers. Therefore, a second possibility is that consumers could respond to the fuel economy restrictions by purchasing smaller, less expensive vehicles. In fact, the most popular new car purchased under the program was the relatively inexpensive Toyota Corolla.

In this paper, 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.

## 2.2 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 which the likelihood of household purchase of a new vehicle is equal for barely eligible and ineligible households. In addition, we use the estimated differences in the likelihood of purchase for different time periods to calculate the counterfactual distribution of purchases over time for eligible households.

---

<sup>9</sup>Separate 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](http://www.nhtsa.gov/CARS-archive/official-information/day-one.pdf)



Second, having identified the pull forward window, we focus on all purchases during this time window and analyze differences between barely eligible and ineligible households in new vehicle fuel economy and transaction price.

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 timing of purchase and types of vehicles purchased.

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

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

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 new vehicle purchased, and the characteristics of the new vehicles purchased. The probability of purchase is defined over different time periods, including the two months of the program, though for the other outcomes we focus on the ten month time period shown to be the point at which the program no longer had an effect on the likelihood of purchase. In the very rare cases in which a household purchased more than one new vehicle over the time frame in question, we sum the transaction prices to construct a measure of total household new vehicle spending. For other outcomes, such as fuel economy, we average across new vehicle purchases by that household.

*Eligible<sub>i</sub>* is an indicator equal to one if the household is classified as being eligible for the program (i.e., if 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 Section 3. We allow for separate relationships between the running variable and the outcome on each side of the eligibility threshold. We estimate Equation (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  $\beta_3$ , which measures the jump in the outcome when going from just-ineligible to just-eligible for the Cash for Clunkers program.

### 2.2.1 Estimating the Effect on the Timing of Purchase and 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. Intuitively, we begin by estimating the probability that a barely eligible and barely ineligible household purchased during the program in July-August 2009. (Not surprisingly, the barely eligibles were more likely to purchase a new vehicle.) Then, we expand the time window sequentially to include more months (i.e. July-September, July-October, July-November, ...) and estimate when the barely ineligible households “catch up”. More formally, for each time window, we estimate Equation (1) with household-level data for which the dependent variable is an indicator of 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.

### 2.2.2 Estimating the Effect on New Vehicle Spending

Our analyses of the impact of the program on new vehicle spending are somewhat different. Rather than using data on all households, we focus only on households that purchased a new vehicle during the “pull forward” window. 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, more than 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 the outcome were continuous across the eligibility threshold. We find it difficult to construct a mechanism 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.<sup>10</sup> 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 4.3.

We note 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. First, because we examine time windows that extend beyond the two months of the program, many households

---

<sup>10</sup>An example that 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.

purchased vehicles after the program had ended and thus were not eligible for the subsidy at all, regardless of the fuel economy rating of their trade-in. Second, 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.<sup>11</sup>

The fuzzy nature of our regression discontinuity design can be seen in Figure 4. As we describe in our data section below, we use data on all purchases by Texas households that includes information of whether the household received a Cash for Clunkers subsidy. Figure 4 shows the discontinuity in the likelihood that a household that purchased a vehicle during the ten month window that also includes the eight months following the program. The discontinuity is around 37 percentage points.<sup>12</sup> 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 vehicle characteristics including vehicle spending, we estimate Equation (1) using a dependent variable of fuel economy or new vehicle price.<sup>13</sup> 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, we rescale the reduced-form estimates by the discontinuity in the likelihood of treatment (Angrist, Imbens, and Rubin, 1996; Imbens and Lemieux, 2008).

### 3 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

---

<sup>11</sup>One potential third reason for the “fuzzy” design relates to how we classify each household’s so-called clunker, as we describe in Section 3 below. 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. Importantly, this situation does not violate our identifying assumption; it merely implies that the RD is fuzzy.

<sup>12</sup>By comparison, the discontinuity in the likelihood of receiving a subsidy is around 75 percentage points during the two months of the program.

<sup>13</sup>As discussed earlier, in the rare cases in which households purchased more than one new vehicle, we sum the purchase prices, such that our measure captures total new vehicle spending by each household.

and Sandler [2011], we restrict our analysis to households that owned no more than seven vehicles in June 2009 (a very minor restriction). 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.<sup>14</sup> 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 with an EPA combined rating of between ten and twenty-seven miles per gallon, inclusive, which spans the largest bandwidth (9 MPG) 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 (NHTS). Although the NHTS does not include identifying information that allows us to match to the administrative data at the household-level, we can use the rich demographic information in NHTS to test our identifying assumption, as we show in

---

<sup>14</sup>This 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.

Section 4.3.

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 1. Among the 5.5 million households in our data, 4.1% purchased a new vehicle in the time window from July 2009-April 2010. The average fuel economy rating of vehicles purchased by households in our sample was 21.7 MPG, while the average transaction price was \$27,600. Table 1 also shows Census Tract characteristics such as demographics and income, which we use as control variables.

## 4 Results

### 4.1 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 2, 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 2 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 purchases by at least some 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 2 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” at least some 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 spring 2010. Put differently, panels (d) through (f) of Figure 2 suggest that all of the increased sales during the program would have occurred 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].

In our estimation of revenue effects below, we define the “pull forward” window to end in April 2010, which is a ten month period including the two-month program. As we show in Section 4.3, our results on the effect of spending are strongly robust to pull forward time windows ranging from 9 to 14 months. In 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 our study finds a pull forward window very similar to that found in other studies, and that our results on the effect of the program on the spending are robust to different time windows, as shown in Section 4.3.

One downside of the results shown in Figure 2 is that while it is clear that at least some purchases were pulled forward from the 7 to 9 months after the program, it is less clear how many of the subsidized purchases would have been made during the two month program anyway. Yet, this is critical to an informed evaluation of the program. For example, if only a few buyers accelerated purchases by, say, several years, but the rest would have purchased during the program anyway, then the pull-forward window would be a poor measure of how much the program accelerated sales.

To provide a more informative measure of the extent to which Cash for Clunkers was successful in pulling forward sales, we use the discontinuity estimates corresponding to Figure 2 to estimate the distribution of counterfactual purchases over time.<sup>15</sup> Results are shown

---

<sup>15</sup>Specifically, we take the actual number of vehicles sold to CfC-eligible households as of each given date

in Figure 3, which indicates that fully 59 percent of subsidies went toward purchases that would have occurred in July or August of 2009 even absent the program. Thus, the majority of the households who purchased under the program would have bought new vehicles during that time anyway. Figure 3 also indicates that roughly 10 percent of program purchases were pulled forward from each of the subsequent two-month periods, ending with March/April.

Overall, these findings suggest that the program had a relatively small impact on the timing of sales. Nearly 60 percent of the subsidies went to households that would have purchased during July and August even absent the program, and the other households would have purchased vehicles no later than April of 2010.

## 4.2 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 change in the composition of vehicles purchased, which has potentially important implications for the stimulus effect on the industry.

To estimate the effect on the type and price of the vehicle purchased, we focus on the households who bought during the 10 month pull forward window. All of these households would have purchased by the end of April 2010 even absent the program, since by construction that was the time frame over which the program had no 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.<sup>16</sup> We then estimate the difference in the fuel economy and transaction price of the vehicles purchased by barely eligible and barely ineligible households.

As shown in Figures 5 and 6, vehicles purchased by eligible households during the ten

---

and use the estimated discontinuity in cumulative purchase frequency on that date to determine the counterfactual number of vehicles sold to these households. The difference between the actual and counterfactual vehicle sales reflects the number of outstanding “pulled-forward” purchases as of that date. Knowing that there were 677 thousand total subsidies awarded (or, respectively, 43 thousand in Texas), we can use these estimates to determine the date at which each subsidized purchase would counterfactually have occurred. For example, as of August 31, 2009, our estimates indicate 276 thousand future counterfactual national new vehicle purchases had in actuality already occurred, i.e. that 41% of subsidized purchases were counterfactually yet to be made.

<sup>16</sup>That is because Section 4.1 suggests that 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.



month pull-forward period were more fuel efficient and less expensive than vehicles purchased by ineligible households. This is consistent with [West, Hoekstra, Meer, and Puller \[2015\]](#), who show that vehicles purchased by the barely eligible tend to be smaller and have less horsepower per pound of vehicle weight. Corresponding regression estimates are shown in the first two rows of [Table 2](#), which show estimates for bandwidths of 9, 7, and 5 MPG, and for polynomial fits ranging from cubic to linear. Estimates indicate that vehicles purchased by eligible households were rated 0.6 to 1.0 MPG higher than vehicles purchased by ineligible households. All estimates are statistically significant ( $p < .001$ ). Estimates in the second row also indicate that the transacted price of vehicles purchased by eligible households was an average of \$1,400 to \$2,200 less than for vehicles purchased by ineligible households during the program. All estimates are statistically significant at conventional levels. These findings indicate that eligible households purchased vehicles that were both less expensive and more fuel efficient than those they would have purchased absent the program.

To confirm that the types of vehicles purchased did change – that this finding of reduced spending is one of substitution across vehicles, not merely an artifact of subsidy incidence – we show estimates of the effect on vehicle MSRP. The estimates using MSRP shown in the third row of [Table 2](#) indicate that the barely eligible households did indeed purchase different, cheaper vehicles over the ten month period. Estimates range from \$1,200 to \$2,000, all of which are statistically significant.

As discussed earlier, these estimates represent intent-to-treat estimates. In order to estimate the local average treatment effect, we rescale these intent-to-treat estimates by the discontinuity in the probability of receiving the subsidy over that ten month time period. That discontinuity was 37.3 percentage points, as shown in Panel (b) of [Figure 4](#). Rescaling the reduced-form estimate of -\$1,875 as shown in row 1 and column 2 of [Table 2](#) results in a local average treatment effect of -\$5,026 ( $-\$1,875/0.373$ ). Thus, results across specifications indicate that on average each buyer under the program spent between \$3,800 and \$5,900 less than they otherwise would have.<sup>17</sup>

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

---

<sup>17</sup>To 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.

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 \$10,000 cheaper with a fuel economy that is 5 MPG higher.<sup>18</sup> 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 nearly 60 percent of the subsidies went to households who would have purchased during the two months of the program anyway, and the others accelerated sales by no more than eight months. Second, and more importantly, the fuel efficiency restrictions of the program led to a substantial change in the type of vehicles purchased. During the two month program and in the eight months that followed, eligible households purchased vehicles that were an average of \$1,800 less expensive, which translates to around \$5,000 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.<sup>19</sup>

## 4.3 Robustness and Threats to Identification

### 4.3.1 Sensitivity to Time Window

One potential limitation of the analysis presented above is that while we believe there is strong evidence in Figure 2 that Cash for Clunkers had no net effect on the likelihood of purchase over some time period around ten months, a time period consistent with prior estimates, 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., transaction 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.

Robustness results are shown in Table 3. This table reports estimates for time windows

---

<sup>18</sup>These figures represent the estimated treatment effects in model (2) rescaled by 0.37 and divided by 0.5.

<sup>19</sup>This 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.

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. For example, estimates in column 3 of the transaction price estimates indicate that eligibility induced households to purchase vehicles that cost \$1,500 to \$2,100 less than the vehicles they would have purchased.

Perhaps more importantly, column 3 shows that the local average treatment effects (i.e., the reduced-form estimate rescaled by the magnitude of the first stage) remains remarkably consistent, varying only from 2.5 to 2.7 MPG for fuel economy, and -\$4,896 to -\$5,171 for transaction price. Thus, even if one were to believe that it took 14 months rather than 10 months for the sales effect of the program to be completely offset, it does not change the conclusion that each subsidy induced households to buy significantly more fuel efficient cars that cost around \$5,000 less.

### 4.3.2 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 mechanism 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 three 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.<sup>20</sup> 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 Appendix Figure B.2, there is no evidence of visually compelling discontinuities in any of these variables, consistent with the identifying assumption. In Appendix Table B.1, we test formally for discontinuities in nine demographic characteristics, and find that none are statistically different from zero.

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. Specifi-

---

<sup>20</sup>The National Household Travel Survey includes information on household demographic characteristics, vehicle ownership, and travel information for a representative sample of U.S. households. We focus only on Texas households.

cally, we analyze the households that purchased during July and August 2008, including the same calendar months in the year prior to the program. Appendix Figure B.3 shows results for the probability of purchasing a new vehicle, the fuel economy rating of those purchases, and the transaction price of those purchases. 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.<sup>21</sup> 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). Thus, there is little evidence that underlying preferences regarding whether to purchase a new vehicle and what type of vehicle to purchase (i.e., price and fuel economy) vary discontinuously across the eligibility threshold.

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 Appendix Figure B.4. 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 have a tendency to purchase vehicles that are less expensive than barely ineligible households, as shown by the lack of a discontinuity in the MSRP of non-clunker vehicles in panel (c).

---

<sup>21</sup>We note, however, that due to the large sample size, even discontinuities of small economic magnitude such as those shown in Figure B.3 are statistically significant. For example, using a cubic fit and a bandwidth of 9 MPG, the estimated discontinuity in panel (a) of Appendix Figure B.3 is two-tenths of a percentage point, and is statistically significant at conventional levels.

## 5 Interpretation and Discussion

### 5.1 The Mechanism Behind the Effect on New Vehicle Spending

As we discuss above, our findings indicate that Cash for Clunkers reduced revenues due to the high fuel economy required for a new vehicle to receive the subsidy. In short, the program reduced the relative prices of fuel efficient cars and thus shifted purchases towards higher fuel economy (and less expensive) cars.

One might conjecture that another mechanism is possible: the program shifted purchases forward to a time period in which consumer confidence was lower.<sup>22</sup> While this seems somewhat unlikely given the majority of subsidized buyers did not accelerate their purchases at all due to the program, a substantial fraction of subsidized buyers did accelerate their purchases. If consumer confidence was lower in months closer to summer 2009, households may have been more cautious about making expensive durable goods purchases and thus more likely to purchase cheaper vehicles. Under this mechanism, the program's effects on revenues would have acted not only through the fuel economy restriction, but also through the *timing* of purchases. While this would not affect the validity of our reduced-form estimate of the impact of the Cash for Clunkers program on spending, it would suggest an alternative mechanism.

To assess this possibility, we use data from the University of Michigan's Consumer Sentiment Index, which is a widely used measure of consumer confidence in the economy. We estimate our regression discontinuity specification (Equation (1)) where  $Outcome_i$  is the Consumer Sentiment Index in the month that household  $i$  made its purchase. Because barely eligible households tended to purchase earlier than the barely ineligible, the RD estimate of  $\beta_3$  will reflect the difference in Consumer Sentiment at time of purchase that was caused by the program pulling forward sales.

Given both the predicted and realized duration of the recession, we find that barely eligible households purchased vehicles when the Consumer Sentiment Index was 0.2 points lower, on average. These results are shown in Appendix Figure B.5. To put this effect in perspective, the average level of the index was 66.3 in 2009, implying a relative change of 0.3 percent. This small change seems unlikely to explain the large difference in spending between barely eligible and ineligible households. As an illustration, assume that i) the 0.2

---

<sup>22</sup>It is worth noting that this effect would have to offset the effect of the additional income that barely eligible households receive due to the subsidy, which would lead those households to purchase larger, more expensive vehicles.

point difference did explain the entire \$2,000 difference in intent-to-treat intensive-margin spending, and ii) the effect of consumer sentiment on new vehicle spending is linear. Those assumptions suggest that the change in consumer sentiment from 2009 to 2010 (from 66.3 to 71.8) would imply an increase in vehicle purchase price of \$55,000 per vehicle. Similarly, the assumptions would imply that conditional on purchase, consumers would buy vehicles that cost \$178,000 more in 2014 (CSI=84.1) than in 2009 (CSI=66.3). Such effects would not only be implausible, but impossible.

Thus, it seems that even during the Great Recession, consumer sentiment did not change enough over short periods of time to explain a meaningful amount of the program's effect on vehicle spending. This underscores that if the program had been implemented *without the fuel economy restriction*, then it is very unlikely that Cash for Clunkers would have caused households to purchase less expensive vehicles.

## 5.2 The Impact of Cash for Clunkers on National New Vehicle Spending

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 automobile 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 a struggling industry – significantly reduced industry revenues over a period of less than a year.

The back-of-the-envelope calculation above extrapolates a local average treatment effect beyond the 18 MPG households, and it is possible that the treatment effect on households with lower MPG clunkers differs from the 18 MPG households.<sup>23</sup> Nevertheless, we can

---

<sup>23</sup>For example, it could be that households with less efficient clunkers would have purchased larger and more expensive new vehicles absent the program compared to their counterparts at the 18 MPG threshold, implying

provide lower bound estimates of the overall revenue effect by assuming that our estimate of the LATE only generalizes to households very close to the discontinuity. For example, assume that the revenue effect for households with 17 or 18 MPG clunkers is given by our estimates in Section 4.2, and that the revenue effect for households with clunkers below 17 MPG *is \$0*. (Here we are being *very* conservative by assuming that households with very inefficient clunkers did not purchase cheaper vehicles than they would have purchased absent the program, which is almost certainly not the case.) Because 38% of Cash for Clunkers transactions involved a clunker with either 17 or 18 MPG, this lower bound estimate is that the program reduced industry revenues by \$1.0 billion to \$1.5 billion, depending on specification. If we extrapolate the LATE to eligible households with clunkers 15-18 MPG (which account for 75% of all Cash for Clunkers transactions), then the lower bound estimates of reduced revenues are \$1.9-\$3.0 billion.

The upshot of these calculations is that even if one is only willing to extrapolate the LATE to households very close to the eligibility cutoff, the program significantly reduced the very consumer spending that it intended to increase.

### 5.3 Potential Justifications for the Program’s Longer-Term Reduction in Spending

One might argue that to the extent increasing new vehicle spending was really important in July and August of 2009, perhaps that short-term increase in spending could justify the longer-term reduction in cumulative new vehicle spending. The effect of the program on cumulative new vehicle spending by CfC-eligible households is shown in Figure 7.<sup>24</sup> The figure shows actual spending and estimates of counterfactual spending if there had been no CfC program. Cumulative spending under the CfC program was larger than counterfactual spending for the months immediately after the program. However, by February 1 the counterfactual expenditures becomes larger and by April has grown to be \$4.0 billion more than actual expenditures under the program. It is difficult to make the case that the brief

---

that we understate the reduction in national spending. Alternatively, if low-MPG households were less constrained or more inframarginal than their counterparts with more efficient clunkers, we would overstate the reduction in national new vehicle spending.

<sup>24</sup>The “Actual spending” curve in Figure 7 is computed by scaling the actual number of purchases by CfC-eligible households in Texas up to the national level, and multiplying this sales count times \$26,930, the average vehicle expenditure for the eligible side just at the discontinuity. The “Counterfactual” line is determined by multiplying the sales count (for each date) times the proportional reduction in sales indicated by the discontinuity estimate, then multiplying this times \$28,805 (which is \$26,930 plus the estimated \$1875 difference in spending from Column (2) of Table 2).

acceleration in spending justifies the loss of \$4.0 billion in revenues to the auto industry, for two reasons. First, we calculate that in order to justify the estimated longer-term reduction in cumulative spending to boost spending for a few months, one would need a discount rate of 208 percent.<sup>25</sup> Given the expected (and realized) duration of the recession, it seems difficult to argue in favor of such a discount rate.<sup>26</sup> Second, we note that Cash for Clunkers seems especially unattractive compared to a counterfactual stimulus policy that left out the environmental component, which also would have accelerated purchases for some households without reducing longer-term spending.

One could also argue that this decline in industry revenue over less than a year 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].

As a result, while it may be possible to make a case for a generic subsidy program that pulls forward new vehicle purchases during a recession for a subset of subsidy recipients, it seems difficult to justify the inclusion of an environmental component in that program.<sup>27</sup> At least in the case of Cash for Clunkers, that environmental component both failed to meet its environmental objectives in a cost-effective way and had a contractionary effect on the industry targeted for stimulus.

## 6 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.

Results indicate that the majority of buyers under the two-month program would have

---

<sup>25</sup>To calculate this annualized discount rate, we find the monthly discount rate that equalizes the present discounted value of the two flows of expenditures in Figure 7, and convert to an annual rate.

<sup>26</sup>For example, the unemployment rate even in April of 2010 was 9.9%, compared to 9.5% in July of 2009.

<sup>27</sup>Under the assumption that new vehicles are normal goods, an unconditional subsidy program would accelerate purchases while at the same time perhaps slightly increasing (due to an income effect) total spending over the longer term.



bought a new vehicle during those two months even without the program. Furthermore, results indicate that the other households who bought under the program - and were responsible for the increase in sales during that time - would have otherwise purchased new vehicles in the following seven to nine months. Thus, over a nine to eleven month period, the program had no impact on the number of vehicles sold.

Strikingly, however, we show that over that nine to eleven month period, including the two months of the program, Cash for Clunkers actually reduced new vehicle spending by around \$5,000 per subsidy. 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 significantly less than the vehicles they otherwise would have purchased.

Thus, while the stimulus program did increase revenues to the auto industry during the two-month program, the environmental component of the bill actually lowered total new vehicle spending over less than a year by inducing people to buy more fuel efficient but less expensive cars. These 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.

## 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 McConnell. Estimating an emissions supply function for accelerated vehicle retirement programs. *Review of Economics and Statistics*, 78(2):251–265, May 1996.
- 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.
- 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.
- 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.
- 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.
- Robert W. Hahn. An economic analysis of scrappage. *RAND Journal of Economics*, 26(2):222–242, Summer 1995.
- David Haugh, Annabelle Mourougane, and Olivier Chatal. The automobile industry in and beyond the crisis. OECD Economics Department Working Paper No. 745, January 2010.
- 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.
- 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 and Arthur A. van Benthem. Vehicle scrappage and gasoline policy. *American Economic Review*, Forthcoming.
- David S. Johnson, Jonathan A. Parker, and Nicholas S. Souleles. Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610, December 2006.
- Stefan Klößner and Gregor Pfeifer. Synthesizing Cash for Clunkers: Stabilizing the car market, hurting the environment. Working paper, February 2015.
- 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 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.
- Edward E. Leamer. *Macroeconomic Patterns and Stories*. Springer, 2009.
- 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 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.
- 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.
- NHTSA. CARS report to Congress. Technical report, U.S. Department of Transportation: National Highway Traffic Safety Administration, December 2009.


- 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.
- Christina Romer and Christopher Carroll. Did ‘Cash-for-Clunkers’ work as intended? White House commentary, Council for Economic Advisers, April 2010.
- James M. Sallee. The surprising incidence of tax credits for the Toyota Prius. *American Economic Journal: Economic Policy*, 3:189–219, May 2011.
- 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.
- Jeremy West, Mark Hoekstra, Jonathan Meer, and Steven L. Puller. Vehicle miles (not) traveled: Why fuel economy requirements don’t increase household driving. MIT Working Paper, 2015.
- 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.

Figure 1: U.S. monthly new vehicle sales, seasonally-adjusted annual rate (Source: NADA)




Figure 2: Cumulative fraction of households purchasing any new vehicle by time period


(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 3: Counterfactual distribution of purchase timing for subsidized purchases




Figure 4: First-stage: Probability of purchase being subsidized by Cash for Clunkers  
Purchases during July 2009 - April 2010 (ten months)




Figure 5: Reduced-form: Fuel economy  
Purchases during July 2009 - April 2010 (ten months)




Figure 6: Reduced-form: Transaction price  
Purchases during July 2009 - April 2010 (ten months)






Figure 7: Cumulative Expenditures Beginning in July 2009

Estimated actual and counterfactual expenditures by U.S. CfC-eligible households for households that purchased a vehicle during July 2009 - April 2010 (ten months)




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

	Median	Mean	Std. Dev.
Total number of households		5,523,973	
<b>Sample: purchased new vehicle</b>			
Number of households		239,559	
Fraction of households		0.041	
<b>Characteristics of new vehicles</b>			
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
<b>Census Tract characteristics</b>			
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 2: Reduced-form 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)
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	Cubic	Cubic	Quadratic	Cubic	Quadratic	Quadratic	Linear
Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	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  $\beta_3$  in Equation (1). Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses. Figures 5 and 6 present graphs corresponding to MPG and transacted price.

Table 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: The first column presents coefficients from regressing each dependent variable (in rows) on an indicator for CARS eligibility, including purchases occurring within the duration of the indicated time window. The LATE estimates are the reduced-form intent-to-treat estimates re-scaled by the first-stage discontinuity in the probability of receiving the subsidy. Standard errors are omitted from the table to enable better comparison; every one of these coefficients is statistically significant with  $p < 0.001$ . The “10 months” rows correspond to the main results as presented in Column (2) of Table 2.

# A Data Appendix - For Online Publication

## 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 (e.g. tractor trailers).

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.<sup>28</sup> 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.

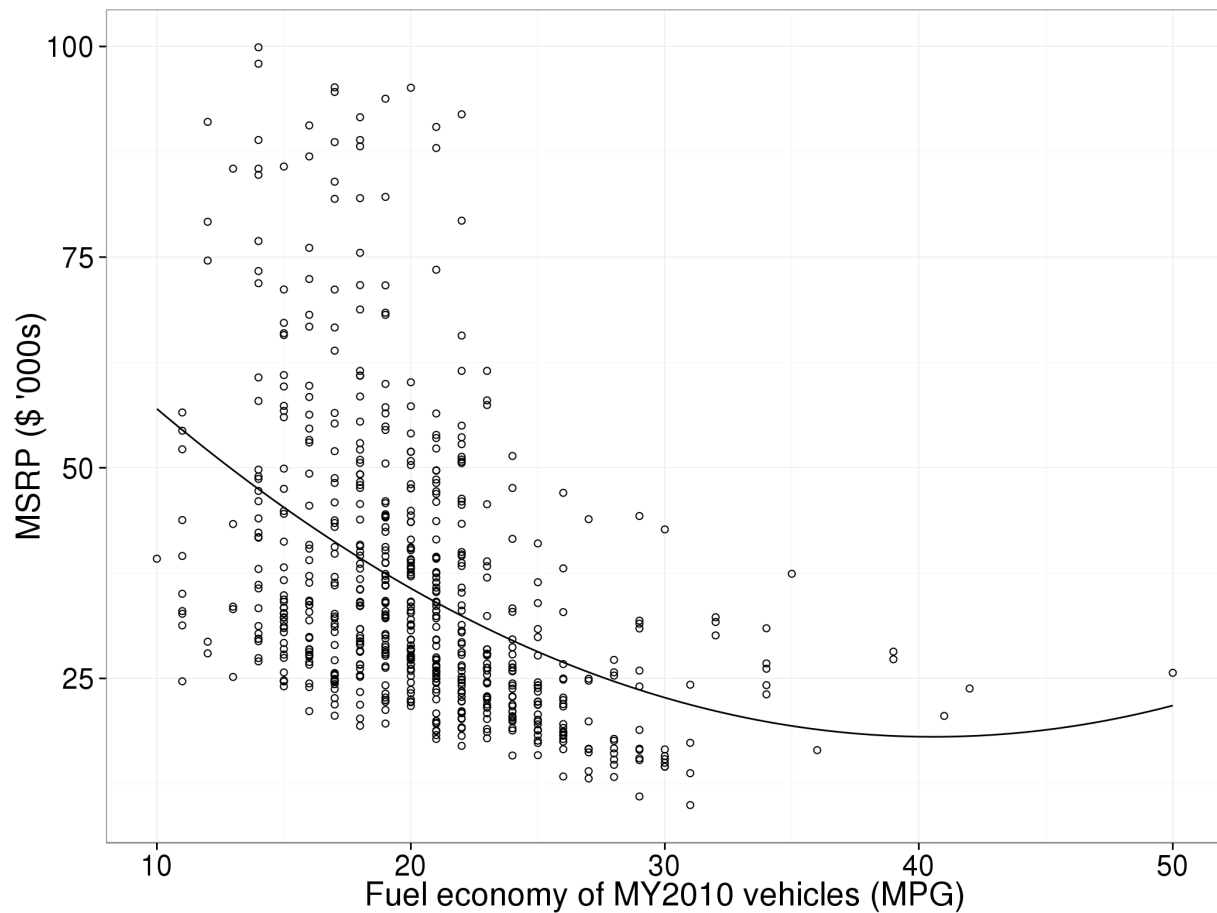
Because DMV registrations are better suited for tracking vehicle purchases than exits from a household’s fleet, we make two additional adjustments to households’ duration of vehicle ownership. We remove a vehicle from a household’s fleet if the latest observed registration (in Texas) has lapsed by six 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.

---

<sup>28</sup>We 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.

## B Additional Figures and Tables - For Online Publication

Figure B.1: Relationship between fuel economy and vehicle price (MSRP)



Notes: This figure plots the fuel economy and manufacturer suggested retail price (MSRP) for all sub-\$100,000 passenger vehicles offered in model year 2010. Data source: DataOne Software.

Figure B.2: Identification checks: National Household Travel Survey (spring 2009)

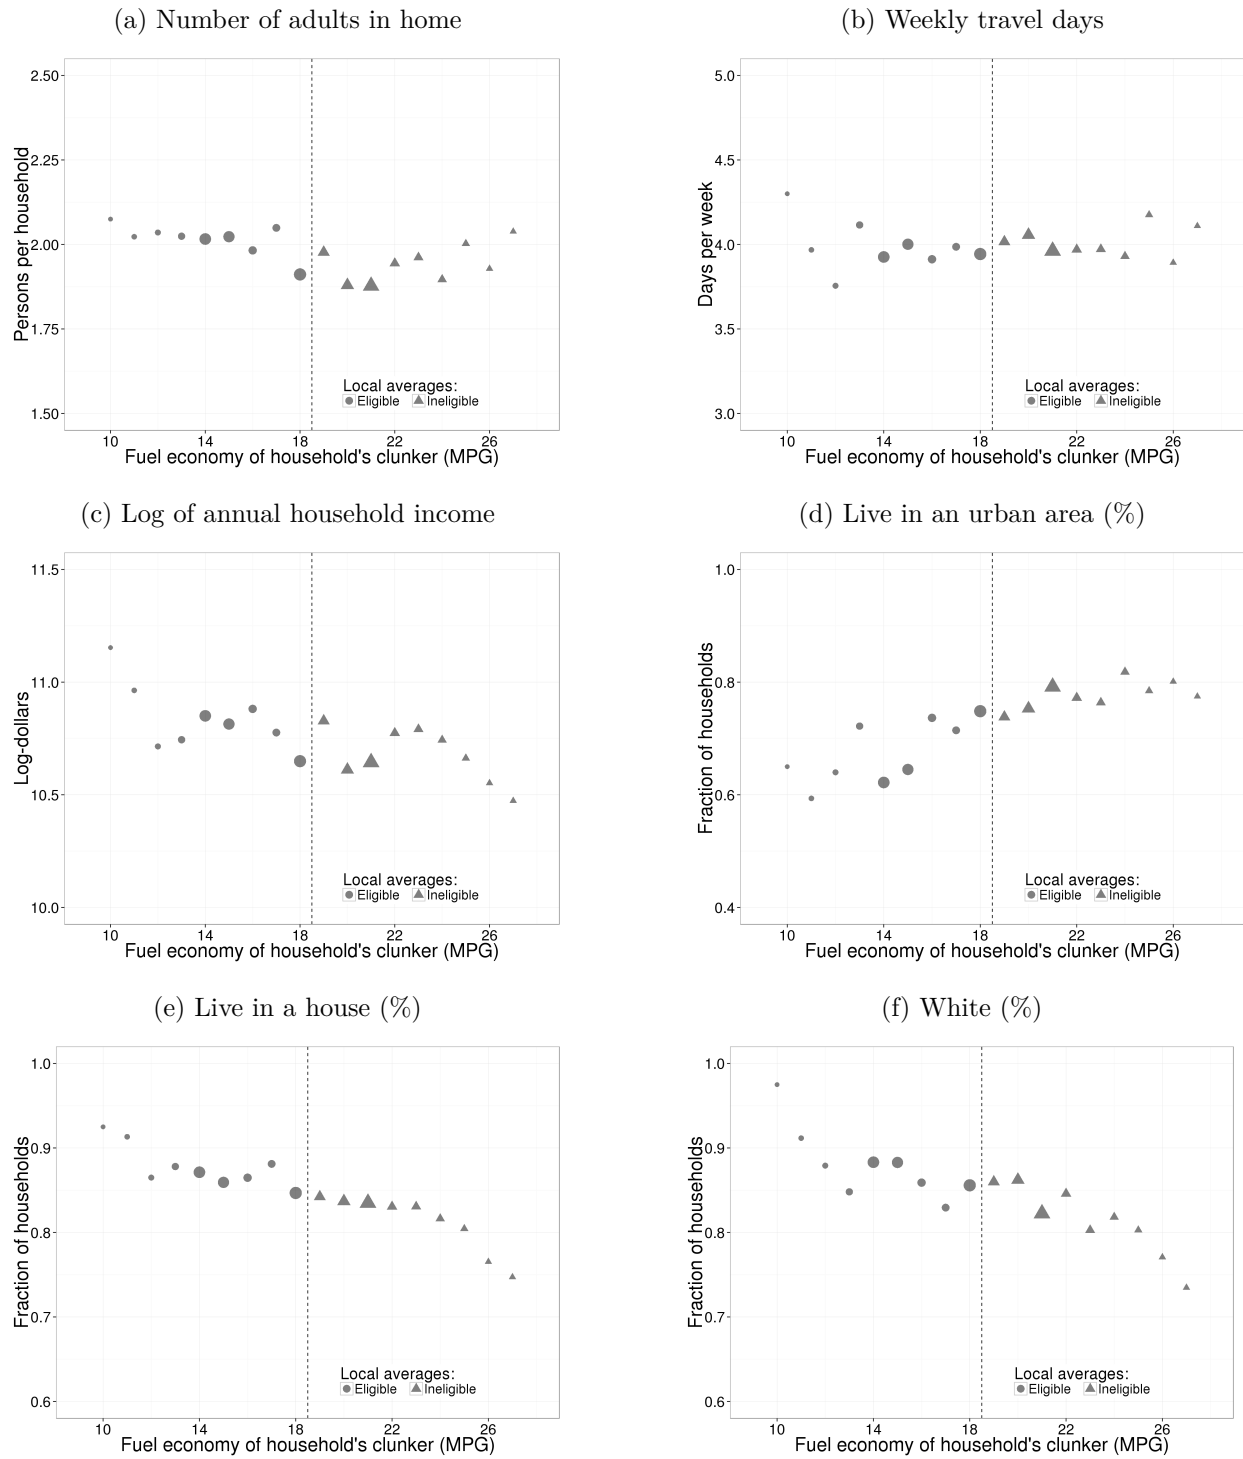




Figure B.3: Identification checks: Households purchasing vehicles prior to CfC

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




Figure B.5: Discontinuity in consumer sentiment index




Table B.1: Identification check: Discontinuities in Vehicle Owner Demographics

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)
Polynomial	Linear	Linear	Linear	Linear	Linear
Observations	6335	6252	6060	5763	5303

Notes: \*  $p < 0.1$  \*\*  $p < 0.05$  \*\*\*  $p < 0.01$  Data are from 2009 NHTS for households in Texas. Each coefficient represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility, which is  $\beta_3$  in Equation (1). Heteroskedasticity-robust standard errors, clustered on the running variable, are reported in parentheses.