

Cash for Corollas: When Stimulus Reduces Spending

Mark Hoekstra

Steven L. Puller

Jeremy West *

July 2016

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 more than 50 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 \$7,600 less. Thus, we estimate that on net, the \$3 billion program reduced total new vehicle spending by more than \$5 billion over that 10 month period.

JEL: H20, Q40, Q50

Keywords: Fiscal stimulus, multifaceted policy, regression discontinuity

*Hoekstra: Texas A&M University and NBER, markhoekstra@tamu.edu. Puller: Texas A&M University and NBER, spuller@tamu.edu. West (corresponding author): Massachusetts Institute of Technology, westj@mit.edu. We are grateful to Pierre Mouganie for excellent research assistance. This work benefited from valuable comments by several anonymous referees as well as Daron Acemoglu, David Autor, John Cochrane, Amy Finkelstein, Jonathan Gruber, Catie Hausman, Stephen Holland, Christopher Knittel, Jason Lindo, Joshua Linn, Erzo F.P. Luttmer, Jonathan Meer, Gregor Pfeifer, Jim Poterba, Matt Zaragoza-Watkins, Sarah Zubairy, and numerous seminar participants. 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 seven 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 was 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...”¹ 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

¹Remarks by the President on National Fuel Efficiency Standards made May 19, 2009, retrieved from 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.²

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.

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.

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. This estimate is similar to the 55 percent estimate reported by [Mian and Sufi \(2012\)](#) and somewhat higher than the 45 and 33 percent estimates reported by [Li, Linn, and Spiller \(2013\)](#) and [Copeland and Kahn \(2013\)](#), respectively. Thus, while the program did accelerate some purchases, we estimate that the majority of buyers under the program would have purchased during the program months

²While 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.

even absent the program.

In addition, our estimates indicate that eight 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 subsidized purchases that were pulled forward from the future (which were roughly 45 percent of all purchases) were accelerated by a maximum of eight months. Put differently, on net the program did not result in any more vehicle purchases than otherwise would have occurred over the ten 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).³

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. To do so, we apply the same regression discontinuity design.⁴ Here, however, we primarily 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 over-

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

⁴While the primary difference between our study and [Li, Linn, and Spiller \(2013\)](#), [Mian and Sufi \(2012\)](#), and [Copeland and Kahn \(2013\)](#) is our focus on total vehicle spending, we also differ in that we use a regression discontinuity design rather than a difference-in-differences approach. [Li, Linn, and Spiller \(2013\)](#) compare the U.S. and Canada over time and estimate that 45 percent of subsidies did not accelerate purchases at all and that the program reduced CO2 emissions at a (high) cost of \$92 to \$288 per ton. [Mian and Sufi \(2012\)](#) exploit regional differences within the U.S. in the eligibility of the vehicle fleet for Cash for Clunkers and estimate that sales were accelerated by no more than 10 months. In addition, one advantage of the approach used by [Mian and Sufi \(2012\)](#) is that it enables them to examine the impact on regional employment, where they find no effect. [Copeland and Kahn \(2013\)](#) compare the production and sales of eligible and ineligible vehicles over time and estimate that the increase in production during the program was less than half of the increase in sales, as the rest were absorbed by inventories. They also conclude that the program had a negligible impact on GDP. Finally, [Li and Wei \(2013\)](#) use a dynamic discrete choice model to examine the tradeoff between the environmental and stimulus components of the CARS program. More broadly, our study also joins a larger 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](#)).

all 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, though we also report similar results from the unconditional analysis.

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 \$8,000 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 213%.

Moreover, we find that the program performed poorly in achieving its secondary objective of improving the fuel economy of the U.S. fleet. We estimate that the program induced households to purchase vehicles that were only 3 MPG more fuel efficient than the vehicles they otherwise would have purchased. This increase in fuel economy, while still notable, does not lead to reductions in environmental damages that justify the program costs.

These findings have significant implications for the national auto industry. We estimate that Cash for Clunkers – which dispensed \$3 billion in subsidies toward the purchase of 677,000 new vehicles nationally – actually reduced revenues to the auto industry by more than \$5 billion over the course of less than one year. Importantly, the estimated impact on national spending is large even if we only extend our regression discontinuity estimate to households that traded in vehicles rated 17 or 18 MPG under the program. Under that conservative assumption, our results still imply that the \$3 billion program reduced new vehicle spending by nearly \$2 billion. 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.⁵ 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.

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 U.S. sales fell from more than sixteen million in 2007 to around ten million in 2009. The shaded region shows the spike that coincided with the two-month stimulus program.

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

⁵Scrapage 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 Bentham, Forthcoming](#)).

barely eligible households.⁶ 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.⁷ 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 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 \$4500.⁸ 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.⁹

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

⁶There 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.

⁷The 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](#)).

⁸As a result, we observe spikes in the distribution of purchases that increased MPG by 5 and 10, as shown in Appendix Figure [B.3](#).

⁹Separate 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

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.2 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. 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 type of vehicle 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.

Eligible_i 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 estimate Equation (1) with least squares, allowing for separate relationships between the running variable and the outcome on each side of the eligibility threshold. 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. We estimate effects using a range of bandwidths that are as large as 5 MPG and as small as 2 MPG.¹⁰

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.

¹⁰Optimal bandwidth techniques, such as that proposed by [Calonico et al. \(2015\)](#), require continuous running variables and are thus not directly applicable to our setting. Instead, we show results from a range of bandwidths, including that which is as small as feasible (2 MPG).

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, there should not be any manipulation around 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

The primary outcome of interest in our study is new vehicle spending. We focus on new vehicle spending, rather than industry profits, because we believe it best measures the mechanism through which any broader stimulus impact on the economy should occur. As discussed earlier, the reason for targeting stimulus on industries such as the new automobile industry is because those sectors lose a disproportionate number of jobs during recessions. Increases in total new vehicle spending result in increases in demand for vehicle components and the workers who make and assemble them. Those workers then continue their normal spending. By comparison, measuring industry profits would not directly capture this increase in demand for labor throughout the supply chain.

Our analyses of the impact of the program on new vehicle spending are somewhat different from our analyses on vehicle purchases. 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.

The primary reason we examine new vehicle spending for new car buyers, rather than across all households, is because more than 95 percent of the households in our sample did not purchase a new vehicle during the program or in the eight months that followed. Because so few households purchased new vehicles, estimates that do not condition on purchases (discussed later in the robustness section and shown in Appendix Figure B.4 and Table B.1) are much more sensitive to choices regarding polynomial and bandwidth. This is because even large discontinuities become very small when averaging over nearly 20 times as many households that did not purchase any new vehicles, which means small changes in estimated slopes can make large differences in the resulting local average treatment effects. As a result, in our main analysis we focus on the households that purchase new cars during the pull forward window, as these are the households that should be driving any effects on spending.

The identifying assumption for our analysis on new-car-buying households 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. This means that our design requires the selection process into buying a car be similar across the eligibility threshold during the pull forward window. We believe this identifying assumption is likely to hold. For example, while it is likely that barely eligible buyers would be different from ineligible households who bought *during the program*, this should no longer 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.¹¹ Consistent with this identifying assumption, we show that there is no compelling

¹¹An 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, then 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. Similarly, it would be problematic if the subsidy pulled some eligible households into buying new cars when they otherwise would have purchased used cars. In that case, it would be difficult to find a pull forward window at all, since there is no reason the new-vehicle-buying rate of the control group should “catch up”. However, in Appendix Figure B.5 we show that there is no evidence of a discontinuity in the likelihood of purchasing a used car either during the two month program, or over the 10 month pull forward window.

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 4.4.

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 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. Finally, a household we classify as ineligible based on the MPG of their oldest car may trade in a newer “clunker” that has a lower MPG rating.

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. 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 use 2SLS to 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 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.

We measure household new vehicle spending using transaction prices for all new vehicles

sold in Texas, which are reported to the Texas DMV for tax purposes. 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.¹² 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, though our regressions focus on the smaller sample of households who own clunkers with MPG ratings between 14 MPG and 23 MPG. The distribution of households is shown in Appendix Figure C.1. As shown in panel (a), the distribution is lumpy throughout, including a large drop in the number of households at 19 MPG versus 18 MPG. As a result, one might worry that there is manipulation – i.e., households who own vehicles rated at 19 MPG present themselves as having 18 MPG vehicles in order to be eligible for the program. However, panel (b) of Figure C.1 shows that the distribution of households was very similar one year earlier, well before Cash for Clunkers. This suggests that the lumpiness is due to the discrete nature of the running variable and the resulting aggregation of vehicles with different levels of popularity, rather than manipulation around the eligibility threshold. This conclusion is also consistent with our understanding of the implementation and enforcement of the program, as described earlier.

In some specifications, we use demographic data from the Census. These data include

¹²This 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.

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 Section 4.4.

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. Specifically, we use VIN to match 34,817 clunkers from the official CARS transaction data to our Texas households in the DMV data. However, due to the imperfect process of defining and matching households described in Appendix A, we were unable to match all of these CARS transactions to new vehicle purchases made by households in our data. As a result, we adjust our empirical first stage by scaling up the matched set to equate to the full set of CARS transactions made by households in Texas.

Summary statistics for vehicle and household fleet characteristics in 2010 are presented in Table 1. Among the 4.5 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 \$28,160. Table 1 also shows Census Tract characteristics such as demographics and income, which we use as control variables.

Finally, given that our data are from Texas, it is natural to question how representative Texas households are of U.S. households as a whole. While it is difficult for us to speak to this directly, we can compare vehicle sales in Texas to national sales over the relevant time period. Panel (a) of Appendix Figure B.1 shows that as a share of national sales, Texas sales stayed roughly constant from 2007 to 2011. Perhaps more importantly, panel (b) shows that the annual change in vehicle sales in Texas closely tracked the annual change for the United States overall. This suggests that at least with respect to vehicle sales over this volatile period, Texas seems to track the U.S. reasonably well.

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 for 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 visual 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 10 months. The purchase probability is higher for eligible households when focusing only on the program months of July-August 2009, but the ineligible households appear to have “caught up” by spring 2010.

Corresponding regression estimates are shown in Table 2, which takes the same form as the tables that follow it. Column 1 shows estimates from a quadratic specification using a bandwidth of 5 MPG, while columns 2 and 3 show results using a bandwidth of 4 MPG and quadratic and linear specifications, respectively. Column 4 shows a linear specification

with a bandwidth of 3, while column 5 shows a linear specification with a bandwidth of 2. Finally, column 6 shows results from a linear specification and a bandwidth of 2 when controlling for county fixed effects as well as Census Tract population, median age, percent white, percent black, percent Asian, percent Hispanic, household size, housing units, percent owner-occupied housing, median income, and median home value.

Consistent with Figure 2, results from Table 2 indicate that while there was a statistically significant and economically large discontinuity in the likelihood of purchase during the Cash for Clunkers program, it disappeared sometime in the following 7 to 8 months. While the exact month at which the estimated discontinuity reaches zero depends somewhat on bandwidth and functional form assumptions, in nearly all cases it approaches zero 7 to 8 months after the program ended. Thus, consistent with panels (d) through (f) of Figure 2, the estimates from the last two rows of Table 2 suggest that all of the increased sales during the program would have occurred anyway in the seven to eight 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. However, while this ten month window seems to be our best estimate of the true pull forward window based on Figure 2 and Table 2, we emphasize that our findings are robust to alternative pull forward windows. Specifically, in Section 4.4, we show that our results are strongly robust to pull forward time windows ranging from 9 to 14 months. In doing so, 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 spending are robust to different time windows that almost certainly capture the true pull forward window, as shown in Section 4.4.

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 8 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 Fig-

ure 2 to estimate the distribution of counterfactual purchases over time.¹³ Results are shown 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 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. The only difference is that barely eligible households were incentivized to purchase more fuel efficient vehicles (and to do so sooner). 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. We then estimate the difference in the fuel economy and transaction price of the vehicles purchased by barely eligible and barely ineligible households.

Figure 4 shows the discontinuity in the likelihood of receiving the subsidy across the

¹³Specifically, we take the actual number of vehicles sold to CfC-eligible households as of each given date 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 677,000 total subsidies awarded (or, respectively, 43,000 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 276,000 future counterfactual national new vehicle purchases had in actuality already occurred, i.e. that 41% of subsidized purchases were counterfactually yet to be made.

eligibility threshold. As discussed earlier, this discontinuity is less than one for several reasons. First, because we defined the running variable for each household as the MPG of its oldest car, some seemingly ineligible households still use the subsidy during the program by trading in a different vehicle. In addition, many eligible households who buy new vehicles may choose not to trade in their oldest vehicle, and others purchase in the eight months that follow the program and thus could not receive the subsidy. As shown in the first row of Table 3, we estimate the discontinuity as between 21 and 23 percentage points, all of which are statistically significant at the one percent level.

The reduced-form impact of Cash for Clunkers on spending is shown in Figure 5, which indicates that barely eligible households spent significantly less on new vehicles than did barely 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 3, which show estimates for bandwidths of 5, 4, 3, and 2 MPG, and for polynomial fits ranging from quadratic to linear. Estimates indicate that eligible households purchased vehicles that cost between \$1,600 and \$2,100 less. All estimates are statistically significant ($p < .001$). These findings indicate that eligible households purchased vehicles that were significantly less expensive 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 3 indicate that the barely eligible households did indeed purchase different, cheaper vehicles over the ten month period. Estimates range from \$1,500 to \$1,800, all of which are statistically significant. Similarly, in the last row we show estimates of the impact of the program on fuel economy, which as discussed earlier is the mechanism through which one would expect the program to affect the type (and price) of vehicle purchased. Estimates indicate that barely eligible households purchased vehicles with MPG ratings that were between 0.64 and 0.81 MPG higher; all estimates are statistically significant at conventional levels.

By combining the effects on transaction price and MSRP, we can also estimate the incidence of the Cash for Clunkers subsidies. To the extent that car dealers do not pass through 100 percent of the subsidy, and assuming an otherwise constant differential between transaction price and MSRP over time, we would expect to see dealers raise the transaction price.

Results in the third row of Table 3 show this is not the case. Rather, the difference between transaction price and MSRP is negative and statistically significant, varying from -\$141 to -\$194 across specifications. This suggests somewhat more than complete passthrough. While this seems somewhat unusual, it could be that the publicity generated by the program and the prospect for sales induced dealerships to discount more heavily than they otherwise would have. In addition, this is similar to a finding by Kaul et al. (2012), who report a similar pass-through rate of more than 100 percent for subsidies in the German version of Cash for Clunkers.

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 21 to 23 percentage points, as shown in Figure 4 and in the first row of Table 3. This implies, for example, that each subsidy induced households to purchase vehicles that were approximately 3 MPG more fuel efficient (e.g., $0.6578/0.2107$) than the counterfactual vehicle. This estimate is somewhat larger than the estimate from Li et al. (2013), though we note that paper uses both a different research design and different data.¹⁴

In Table 4, we show local average treatment estimates of the effect of the Cash for Clunkers subsidy on new vehicle spending. Intuitively, this estimator rescales the reduced-form effect on spending (e.g., \$1,600) by the estimated discontinuity in the likelihood of receiving a subsidy (e.g., 0.21). Estimates range from \$7,600 in column 6 to \$9,200 in column 3; all are statistically significant at conventional levels. Thus, results indicate that on average each buyer under the program spent around \$8,000 less than they otherwise would have.¹⁵

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 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 three-quarters

¹⁴Li et al. (2013) estimate an intent-to-treat effect of 0.21 MPG using the full model and the June - December time period. Given that the approximately 10.8 percent of U.S. transactions over that time period were Cash for Clunkers transactions (677,000/6,270,967), that implies a per-subsidy effect of 1.9 MPG.

¹⁵To 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.

of households downsizing and one-quarter 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 4 MPG higher. This difference represents roughly a downsize from a Chevrolet Equinox SUV to a Toyota Corolla. This shift is also likely to disproportionately lower industry profits, as large SUVs and pickup trucks are widely reported to have higher profit margins than smaller, more fuel efficient vehicles.¹⁶

Finally, we can also show the impact of the program on spending by comparing the distributions of new vehicle purchases by households on either side of the eligibility threshold. This is shown in Figure 6, which compares the prices paid by households with 18 and 19 MPG clunkers. Results indicate that the net effect of the stimulus was to reduce purchases of vehicles priced between \$30,000 and \$50,000 and increase sales of vehicles priced between \$15,000 and \$25,000.

The mechanism underlying this shift toward less expensive cars is shown in Figure 7. Panel (a) shows that barely eligible households shifted toward more fuel efficient vehicles, while Panels (b) through (d) show that these vehicles were also smaller and lighter (as measured by curb weight), lower-performing (as measured by horsepower per pound), and less likely to be an SUV or pickup truck. In short, results indicate that barely eligible households shifted to more fuel efficient vehicles that were also smaller, lower-performing, and less expensive.

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 \$1,600 less expensive, which translates to around \$7,600 less per vehicle purchased under the program. Thus, Cash for Clunkers significantly reduced new vehicle spending over a period of less than a year.¹⁷

¹⁶For example, see <http://blogs.wsj.com/corporate-intelligence/2013/01/29/fords-margins-its-all-about-the-trucks/>.

¹⁷This 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.

4.3 Effect on U.S. Assembly and Value-Added of New Vehicle Purchases

In light of the evidence presented above that the Cash for Clunkers program shifted buyers toward more fuel efficient, less expensive vehicles, one might also wonder about the impact the program had on U.S. value-added. That is, to the extent that the smaller vehicles purchased by consumers were disproportionately imported, whereas larger SUVs would have been produced domestically, the stimulus impact on the U.S. auto industry could be even worse than outlined above. To address this question, we examine the impact on whether the new vehicle purchased was assembled in the U.S., and the estimated amount of production value added in the U.S. or Canada. We do so using data from DataOne Software and the National Highway Traffic Safety Administration’s American Automobile Labeling Act Reports.

Results are shown in Figure 8. There is no evidence that the Cash for Clunkers program shifted purchases either away from vehicles assembled in the U.S. or away from vehicles with relatively high value-added in the U.S. Thus, while Cash for Clunkers did reduce overall new vehicle spending, affecting all parts and vehicle manufacturers including those who produce and assemble domestically, there is no evidence the program had a disproportionate effect on U.S. production.

4.4 Robustness and Threats to Identification

4.4.1 Sensitivity to Time Window

One potential limitation of the analysis presented above is that while we believe there is strong evidence in Figure 2 and Table 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 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 program.

Robustness results are shown in Table 5. This table reports estimates for time windows ranging from 9 months to 14 months. We use a bandwidth of 2 MPG and a linear fit with controls. As one would expect, the first-stage 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 reduced-form estimate on spending also falls as the window is lengthened. However, all estimates remain both statistically significant and economically meaningful.

Perhaps more importantly, the third row of Table 5 shows that the local average treatment effects (i.e., the reduced-form estimate rescaled by the magnitude of the first stage using 2SLS) remains remarkably consistent across the different time windows. For example, the estimated effect of a Cash for Clunkers subsidy on new vehicle spending is -\$7,622 using a window of 9 months and -\$7,535 for a window of 12 months. Estimates are similarly consistent across windows for MSRP. 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 each subsidy induced households to buy significantly more fuel efficient cars that cost around \$7,600 less.

4.4.2 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 manipulation around the cutoff seems unlikely given how the program was implemented, 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.¹⁸ 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 C.2, there is no evidence of visually compelling discontinuities in any of these variables, consistent with the identifying assumption.

In addition, we examine whether 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 between July of 2008 and April of 2009, which is exactly one year earlier. Appendix Figure C.3 shows results for the probability of purchasing a new vehicle, the fuel economy rating of those purchases, and the transaction price of those

¹⁸The 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.

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.

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 C.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. In addition, there is perhaps some evidence that barely eligible households favor vehicles with somewhat lower MSRPs (panel (c)), and come from neighborhoods with slightly lower income (panel (d)), though we note the result on income appears to be driven entirely by the households with 19 MPG clunkers.¹⁹ There are no differences in the demographics (median age or percent white) of households across the cutoff.

4.4.3 Estimates Unconditional on New Vehicle Purchase

Another potential concern with our empirical approach is that we have estimated the effect of Cash for Clunkers on spending using only the sample of households that purchased new vehicles over the 10 month period. Here, we instead estimate the effects on spending using the full population of households, regardless of whether or not they purchased a new vehicle over the ten month period.

¹⁹We also note that in results not shown there exists a similar pattern with respect to income in 2008, but in contrast to 2009, there is no evidence of a significant drop in spending in 2008, as shown in Figure C.3.

Results are shown in Appendix Figure B.4, with corresponding regression estimates shown in Table B.1. As shown in Figure B.4, there is a visually compelling discontinuity in the likelihood of receiving a subsidy at the cutoff. Estimates in the first row of Table B.1 range from 0.9 to 1 percentage point, all of which are statistically significant at the one percent level. Panel (b) of Figure B.4 shows household spending across the eligibility threshold, where there also appears to be a discontinuity. Estimates of this reduced-form effect are shown in the second row of Table B.1. While all of the estimates are statistically significant at the one percent level, they are sensitive to bandwidth and functional form assumptions. For example, the estimate in column (3) corresponding to a bandwidth of 4 MPG and a linear specification is -\$39, while the estimate in column (2) is more than three times as large at -\$129. This variation in estimates leads to huge variation in estimated local average treatment effects, which as shown in the third row vary from -\$3,914 to -\$14,091. These estimates are statistically significant at the one percent level and are centered around our estimates of -\$8,000 to -\$9,000 from our main analysis above. As our result, we conclude that while the unconditional analysis results in a much wider range of estimates due to sensitivity to functional form and bandwidth, the overall findings are quite similar to our main analysis in which we focus on households that purchased new vehicles over the ten month period.

5 Discussion

5.1 The Role of Fuel Economy Restrictions Versus MSRP Restrictions and Timing of Purchase

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.

However, the interpretation that the effect on spending is driven entirely by the fuel economy restrictions ignores the fact that households were specifically prohibited from purchasing vehicles whose MSRPs exceeded \$45,000. This seems unlikely to drive all of the substitution, particularly given the distribution of purchase prices shown in Figure 6, where it is clear that many buyers shifted away from vehicles priced at less than \$45,000. Still, the question remains as to how much of the overall spending effect can be attributed to the restriction on high-priced vehicles. In Appendix Figure C.5 we show the fraction of vehicles

that were purchased with $\text{MSRP} > \$45,000$. As shown in that figure, there is a small discontinuity of around 0.02; barely eligible households were two percentage points less likely to purchase a vehicle that cost more than \$45,000 than barely ineligible households. A back-of-the-envelope calculation indicates that this restriction can explain just less than 16 percent of the overall reduction in spending across the cutoff.²⁰

In addition to the MSRP restriction, one might also conjecture that another mechanism is possible: the program shifted purchases forward to a time period in which consumer confidence was lower.²¹ While this seems somewhat unlikely given the majority of subsidized buyers did not accelerate their purchases at all due to the program, some buyers did accelerate their purchases by a few months. 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 β_3 will reflect the difference in Consumer Sentiment at time of purchase that was caused by the program pulling forward sales.

Results are shown in panel (a) of Appendix Figure C.5, and indicate that barely eligible households purchased vehicles when the Consumer Sentiment Index was 0.2 points lower, on average. 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 difference is also small compared to overall changes in the index over this time period. For example, from January to December of 2007, the index fell from 96.9 to 75.5. Similarly, the index went from 78.4 to 60.1 from

²⁰Barely ineligible households who purchased a vehicle priced higher than \$45,000 spent an average of \$57,700, or roughly \$12,700 more than they would have absent the restriction. Because only two percent of buyers were affected at the cutoff, this implies a spending difference of \$254 ($0.02 * \$12,700$), which is 16 percent of our estimated discontinuity of \$1,601.

²¹It 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.

January to December of 2008. Because it seems implausible that such a small change in consumer sentiment could cause a \$1,500 reduction in new vehicle spending, we conclude that the reduction in spending was driven primarily by the fuel economy restrictions, rather than differences in timing of purchase.

5.2 Alternative Justifications for the Reduction in Spending: Acceleration of Purchases and Environmental Impact

One might argue that to the extent increasing new vehicle spending was very important in July and August of 2009, perhaps the short-term increase in spending could justify the longer-term reduction in cumulative new vehicle spending. The effect of the program on cumulative U.S. new vehicle spending by CfC-eligible households is shown in Figure 9.²² 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 the middle of March the counterfactual expenditures becomes larger and by April has grown to be \$5 billion more than actual expenditures under the program. It is difficult to make the case that the brief acceleration in spending justifies the loss of \$5 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 213 percent.²³ Given the expected (and realized) duration of the recession, it seems difficult to argue in favor of such a discount rate.²⁴ 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 also could argue that this decline in industry revenue over less than a year could be justified if the program offered cost-effective environmental benefits. However, previous research has shown this program to be a costly means to reduce environmental damage (Knittel (2009), Li et al. (2013)). We use our estimated discontinuities to conduct a back-

²²Discontinuities for cumulative purchase frequency are estimated daily using a bandwidth of 2 MPG with controls and applied to purchase counts by eligible households within the full bandwidth. The per-vehicle reduced form spending discontinuity is taken from Column (6) of Table 3. Scaling from Texas to the aggregate national level is done using a scale factor of 11.83, as Texas accounted for 8.45 percent of U.S. sales during 2009-2010.

²³To 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 9, and convert to an annual rate.

²⁴For example, the unemployment rate even in April of 2010 was 9.9%, compared to 9.5% in July of 2009.

of-the-envelope calculation of the environmental benefits of the program, and we also find that the environmental benefits do not justify the program costs.

To see this, consider the effect of the program – it induced households who otherwise would have purchased a new vehicle during the next ten months to purchase a new vehicle with higher fuel economy and to scrap the clunker. Thus the environmental benefits come from two sources: (1) the program induced households to buy new cars that have different lifetime emission costs, and (2) the program scrapped used vehicles that otherwise would have been driven for their remaining lifetimes.

Our paper estimates several key parameters of this calculation. First, we estimate that the program induced households to purchase a new car that was 3.1 MPG more fuel efficient than the household otherwise would have purchased. This parameter allows us to calculate the counterfactual gasoline consumption and resulting emissions of the new cars. Second, we estimate that households purchasing under the program would have purchased a new vehicle anyway in the very near future. This suggests that the first-order impact is the difference in fuel economy of the new vehicles rather than the timing of replacing a clunker with a new vehicle. Finally, a companion paper ([West et al. \(2015\)](#)) estimates the impact on driving behavior and finds no evidence that households drive more miles in response to purchasing more fuel efficient (but smaller) vehicles, so we assume zero rebound effect.

In order to estimate the first effect, we calculate the difference in lifetime emissions between the new cars that were purchased under Cash for Clunkers and the new cars that counterfactually would have been purchased absent the program. We consider the same five pollutants as [Holland et al. \(forthcoming\)](#) - CO₂ and the local criteria pollutants NO_x, PM_{2.5}, VOCs, and SO₂. Over the average vehicle lifetime of 152,137 miles ([Lu \(2006\)](#)), the change in fuel economy from 22.0 to 25.1 reduces gasoline consumption by 859 gallons. The marginal environmental damages of CO₂ and SO₂ are proportional to gasoline consumption. Emissions of NO_x, PM_{2.5} and VOCs are determined by the emission standards in place for new cars in 2009, and we follow the approach of [Holland et al. \(forthcoming\)](#) using these standards as the estimated emission rates. The emission per mile standards are identical for the typical 2009 vehicles with 22 and 25 MPG, yielding no changes in emissions of those three local pollutants.²⁵ Therefore, the environmental benefits of the program on the new vehicle fleet are reductions in CO₂ and SO₂ emissions over the lifetime of the new car.

Second, consider the impact of the program feature that the clunkers were scrapped.

²⁵In fact, following our alternative scenario that one-quarter of households made no change in purchase and three-quarters switched from the Equinox to the Corolla, there would still be no changes in these local pollutants because the Equinox and Corolla have identical (bin 5) emission standards.

Absent scrappage, the vehicles may have been sold in the used car market and driven for their remaining lifetimes. A full analysis would require a complete modeling of the dynamics of how households replace used cars which is beyond the scope of this paper. It is likely that some households who would have purchased the used clunkers instead purchase even older and higher emission vehicles while other households instead purchase newer used vehicles with lower emissions. We make the simplifying assumption that the clunkers that were scrapped are replaced by equally polluting vehicles and driven the same number of remaining lifetime miles as the scrapped clunkers. Thus the effect of early scrappage had no net effect on emissions. Our estimates are unlikely to be sensitive to this assumption because most scrapped clunkers were near the end of their lifetimes anyway – the average mileage of the scrapped vehicles was 160,155 miles.

We estimate environmental damages by combining estimated emission rates from the Energy Information Administration and [Holland et al. \(forthcoming\)](#) along with estimated marginal damages from [Interagency Working Group \(2013\)](#) and [Muller and Mendelsohn \(2009\)](#).²⁶ We find that each subsidized vehicle reduced environmental damages by \$253 (with nearly all of the damages from CO₂). Comparing the benefits to the average subsidy of \$4,210, it is clear from this back-of-the-envelope calculation that program benefits were relatively small compared to fiscal costs.

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 minority of subsidy recipients, it seems difficult to justify the inclusion of an environmental component in that program. 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.

5.3 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 \$1,600 and \$2,100 per new-car-buying household. This scales to a \$7,600 to \$9,200 reduction in spending per household that purchased a vehicle

²⁶In particular, we use a social cost of carbon of \$33 per ton (in 2007 dollars) ([Interagency Working Group, 2013](#)) and marginal damages of SO₂ of \$970/ton.

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 \$5.1 billion to \$6.2 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.²⁷ Nevertheless, we can 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 \$1.9 billion to \$2.4 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 \$3.8-\$4.7 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 \$3 billion program significantly reduced the very consumer spending that it intended to increase.

²⁷For 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 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.

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 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 eight months. Thus, over a nine to ten month period, the program had no impact on the number of vehicles sold.

Strikingly, however, we show that over that ten month period, including the two months of the program, Cash for Clunkers actually reduced new vehicle spending by around \$8,000 per subsidy. We attribute this primarily 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.
- Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, November 2015.
- 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.
- Stephen P. Holland, Erin Mansur, Nicholas Z. Muller, and Andrew Yates. Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review*, forthcoming.

- 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 Bentham. 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.
- Ashok Kaul, Gregor Pfeifer, and Stefan Witte. The incidence of cash for clunkers: An analysis of the 2009 car scrappage scheme in Germany. University of Zurich Working Paper No. 68, April 2012.
- 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.
- 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.
- S. Lu. Vehicle scrappage and travel mileage schedules. Technical report, National Highway Traffic Safety Administration, January 2006.

- 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.
- Nicholas Z. Muller and Robert Mendelsohn. Efficient pollution regulation: Getting the prices right. *American Economic Review*, 99(5):1714–1739, December 2009.
- 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.
- 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



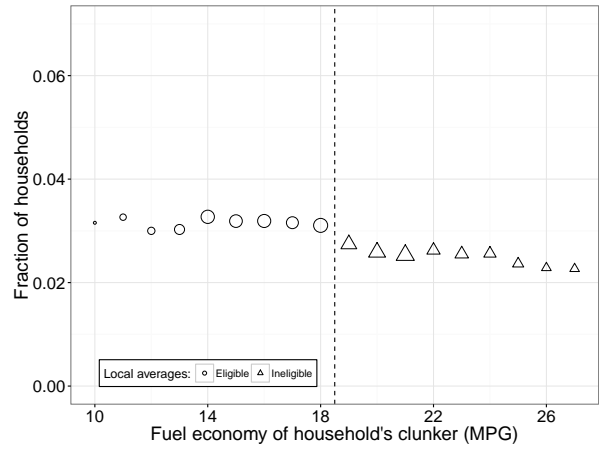
Note: Shaded region indicates Cash for Clunkers. Data source: National Automobile Dealers Association.

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


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




(b) July 2009 - December 2009 (6 months)




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



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



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



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

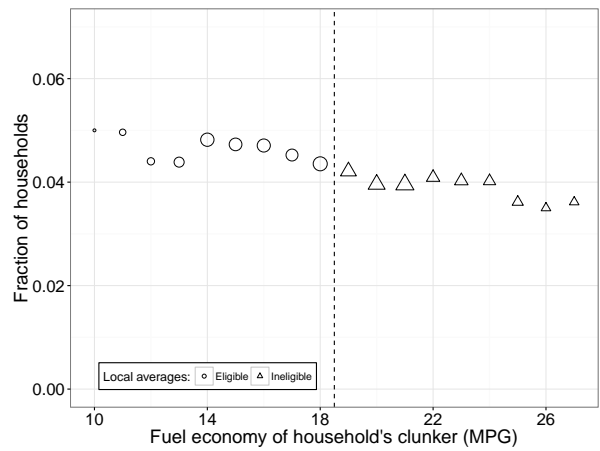




Figure 3: Counterfactual timing of CfC-subsidized vehicle purchases




Note: 10 month time window July 2009 - April 2010.

Figure 4: First-stage: Households subsidized by CfC for new vehicle purchases




Note: 10 month time window July 2009 - April 2010.

Figure 5: Reduced-form: Household spending on new vehicle purchases (transaction price)



Note: 10 month time window July 2009 - April 2010.


Figure 6: Distribution of purchase prices for households with 18 or 19 MPG clunkers




Note: 10 month time window July 2009 - April 2010.

Figure 7: Reduced-form: Selected vehicle attributes of new vehicle purchases


(a) Fuel economy (MPG)




(b) Safety/comfort (curb weight)



(c) Performance (horsepower per pound)




(d) Body type of SUV or pickup truck




Note: 10 month time window July 2009 - April 2010.

Figure 8: Reduced-form: U.S. manufacturing value added and assembly of new vehicle purchases

(a) Production value added in U.S. or Canada

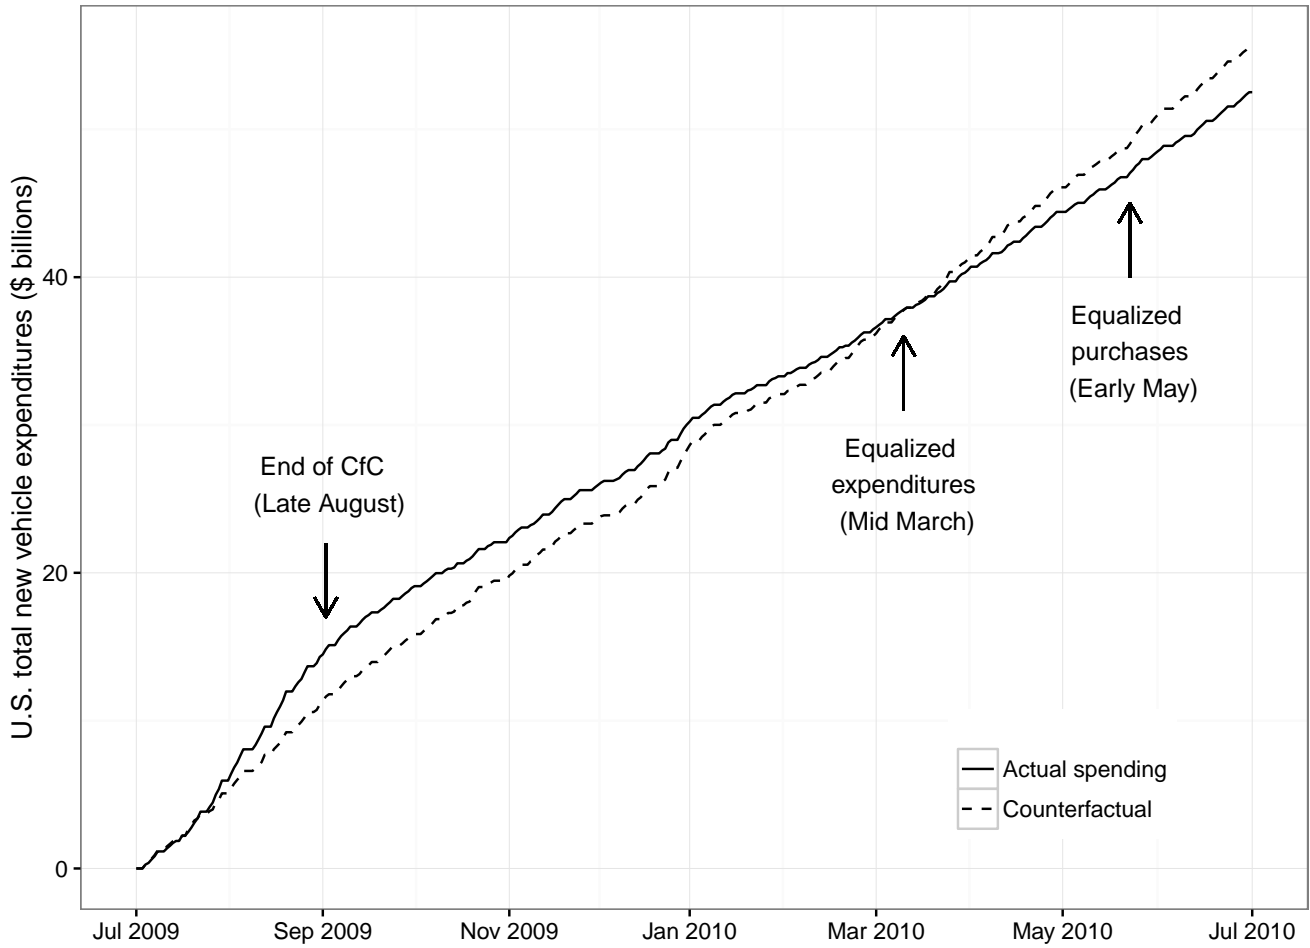


(b) U.S. plant final assembly of vehicle



Note: 10 month time window July 2009 - April 2010. Data sources: DataOne Software, National Highway Traffic Safety Administration's American Automobile Labeling Act Reports.

Figure 9: Estimated aggregate U.S. actual and counterfactual cumulative new vehicle expenditures by CfC-eligible households



Notes: Discontinuities for cumulative purchase frequency are estimated daily using a bandwidth of 2 MPG with controls and applied to purchase counts by eligible households within the full bandwidth. The per-vehicle reduced form spending discontinuity is taken from Column (6) of Table 3. Scaling from Texas to the aggregate national level is done using a scale factor of 11.83, as Texas accounted for 8.45 percent of U.S. sales during 2009-2010.

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

	Median	Mean	St. Dev.
Total number of households		4,525,057	
Sample: purchased new vehicle			
Number of households		197,745	
Fraction of households		0.04	
Characteristics of new vehicles			
Transaction price (\$ '000s)	26.04	28.16	11.17
MSRP (\$ '000s)	26.25	27.97	10.38
Fuel economy (MPG)	21	21.68	5.87
Curb weight (lbs.)	3,760	3,976	1,000
Horsepower per 1000 lbs.	55.62	57.73	12.84
Pickup truck or SUV	1	0.50	0.50
Value added in U.S. or Canada	0.60	0.50	0.28
Final assembly plant in U.S.	1	0.56	0.50
Census Tract characteristics			
Population	5,783	6,361	3,082
Median age	34.20	34.53	5.09
White (%)	82.30	77.75	16.56
Black (%)	3.80	7.97	12.52
Asian (%)	1.40	3.10	4.65
Hispanic (%)	13.90	24.18	24.92
Household size	2.83	2.83	0.42
Housing units	2,215	2,403	1,102
Owner-occupied (%)	78.90	74.11	18.01
Median income (\$)	48,308	53,955	24,710
Median home value (\$ '000s)	93.20	111.86	74.05

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). Only households with a clunker of between 14 and 23 MPG (bandwidth of five) are included. The Census Tract-level characteristics are from the 2000 Decennial Census.

Table 2: Estimated discontinuities for cumulative fraction of households purchasing new vehicle by time period

	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Cash for Clunkers (2 months)	0.0052*** (0.0003)	0.0044*** (0.0004)	0.0054*** (0.0002)	0.0051*** (0.0003)	0.0044*** (0.0004)	0.0049*** (0.0004)
6 months	0.0029*** (0.0005)	0.0019*** (0.0006)	0.0039*** (0.0004)	0.0031*** (0.0004)	0.0026*** (0.0005)	0.0036*** (0.0005)
7 months	0.0024*** (0.0005)	0.0014** (0.0006)	0.0035*** (0.0004)	0.0026*** (0.0004)	0.0022*** (0.0006)	0.0034*** (0.0006)
8 months	0.0016*** (0.0006)	0.0005 (0.0006)	0.0029*** (0.0004)	0.0020*** (0.0004)	0.0013** (0.0006)	0.0026*** (0.0006)
9 months	0.0006 (0.0006)	-0.0007 (0.0007)	0.0023*** (0.0004)	0.0012** (0.0005)	0.0003 (0.0006)	0.0017*** (0.0006)
10 months	-0.0004 (0.0006)	-0.0018** (0.0007)	0.0018*** (0.0004)	0.0004 (0.0005)	-0.0007 (0.0007)	0.0008 (0.0007)
Bandwidth	5 MPG	4 MPG	4 MPG	3 MPG	2 MPG	2 MPG
Polynomial	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Controls	No	No	No	No	No	Yes
Observations	4,525,057	3,717,845	3,717,845	2,985,445	1,897,837	1,897,837

*p<0.1; **p<0.05; ***p<0.01 Each coefficient represents a separate regression of the dependent variable (indicator for new vehicle purchase) on an indicator for CARS eligibility, which yields an estimate of β_3 in Equation (1). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

Table 3: Reduced-form estimated discontinuities for new vehicle purchase characteristics

	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Subsidized (percent)	0.2342*** (0.0062)	0.2198*** (0.0073)	0.2284*** (0.0045)	0.2299*** (0.0052)	0.2106*** (0.0073)	0.2107*** (0.0073)
Spending (dollars)	-1,938*** (161)	-1,906*** (189)	-2,091*** (116)	-2,030*** (133)	-1,905*** (184)	-1,601*** (177)
MSRP (dollars)	-1,803*** (149)	-1,712*** (175)	-1,950*** (108)	-1,870*** (124)	-1,718*** (171)	-1,450*** (165)
Spending-MSRP (dollars)	-135** (57)	-194*** (67)	-141*** (41)	-160*** (46)	-187*** (61)	-151** (61)
Fuel economy (MPG)	0.8104*** (0.0852)	0.6563*** (0.1002)	0.7472*** (0.0615)	0.7347*** (0.0701)	0.6375*** (0.0935)	0.6578*** (0.0929)
Bandwidth	5 MPG	4 MPG	4 MPG	3 MPG	2 MPG	2 MPG
Polynomial	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Controls	No	No	No	No	No	Yes
Observations	197,745	160,918	160,918	127,869	81,118	81,118

*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, which yields an estimate of β_3 in Equation (1). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

Table 4: Two-stage least squares estimates for new vehicle spending during July 2009 - April 2010 (10 months)

	Estimated discontinuity (2SLS)					
	(1)	(2)	(3)	(4)	(5)	(6)
Spending (dollars)	-8,274*** (671)	-8,673*** (843)	-9,157*** (499)	-8,832*** (569)	-9,047*** (853)	-7,600*** (818)
Bandwidth	5 MPG	4 MPG	4 MPG	3 MPG	2 MPG	2 MPG
Polynomial	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Controls	No	No	No	No	No	Yes
Observations	197,745	160,918	160,918	127,869	81,118	81,118

*p<0.1; **p<0.05; ***p<0.01 Each coefficient represents a separate two-stage least squares regression of the dependent variable (new vehicle spending) on an indicator for CARS subsidy, instrumented for by CARS eligibility, which yields a local average treatment effect estimate. Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

Table 5: Robustness of estimated discontinuities to alternate time windows

	Time window in months					
	9 months	10 (main)	11 months	12 months	13 months	14 months
Subsidized (percent) [First-stage]	0.2254*** (0.0079)	0.2107*** (0.0073)	0.1968*** (0.0067)	0.1843*** (0.0062)	0.1734*** (0.0058)	0.1627*** (0.0054)
Spending (dollars) [Reduced-form]	-1,718*** (186)	-1,601*** (177)	-1,464*** (169)	-1,305*** (162)	-1,319*** (158)	-1,226*** (153)
MSRP (dollars) [Reduced-form]	-1,568*** (173)	-1,450*** (165)	-1,338*** (158)	-1,190*** (152)	-1,193*** (148)	-1,113*** (143)
Spending (dollars) [2SLS]	-7,622*** (800)	-7,600*** (818)	-7,435*** (838)	-7,081*** (859)	-7,605*** (891)	-7,535*** (919)
MSRP (dollars) [2SLS]	-6,955*** (747)	-6,882*** (763)	-6,798*** (782)	-6,459*** (802)	-6,881*** (833)	-6,837*** (860)
Bandwidth	2 MPG	2 MPG	2 MPG	2 MPG	2 MPG	2 MPG
Polynomial	Linear	Linear	Linear	Linear	Linear	Linear
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	74,228	81,118	88,712	96,075	102,816	110,556

*p<0.1; **p<0.05; ***p<0.01 Each coefficient in rows 1-3 represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility. Each coefficient in rows 4-5 represents a separate two-stage least squares regression of the dependent variable (in rows) on an indicator for CARS subsidy, instrumented for by CARS eligibility. Columns vary the time window of included sales. Standard errors are reported in parentheses.

A Data Appendix - For Online Publication

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.²⁸ 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.

²⁸We 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 Figures and Tables for Online Publication

Figure B.1: Representativeness of Texas new vehicle sales

(a) Texas proportion of national new vehicle sales volume




(b) Annual change in new vehicle sales volume




Data source: National Automobile Dealers Association.

Figure B.2: Relationship between fuel economy and sale price



Note: includes all MY2010 vehicles with a sub-\$100,000 MSRP. Data source: DataOne Software.


Figure B.3: Distribution of improvements in fuel economy for actual CARS trades



Note: includes all passenger cars and trucks traded in nationally under CARS program.
Data source: National Highway Traffic Safety Administration.

Figure B.4: First-stage subsidy and reduced-form spending unconditional on purchasing vehicle

(a) Cash for Clunkers subsidy rate for all Texas households



(b) New vehicle spending per Texas household (July 2009 - April 2010)






Figure B.5: Cumulative fraction of households purchasing any **used** vehicle by time period


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




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



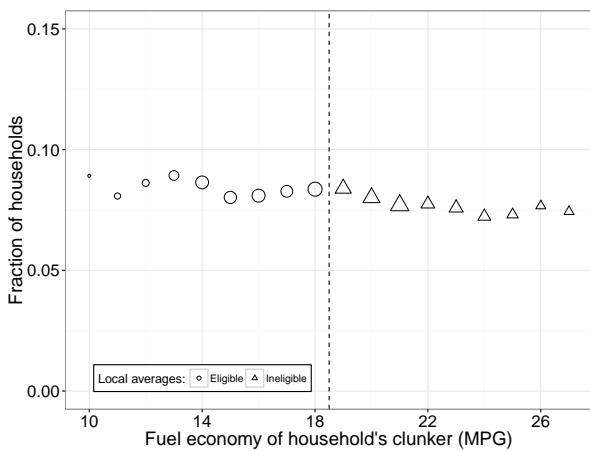
(c) July 2009 - February 2010 (8 months)



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



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



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

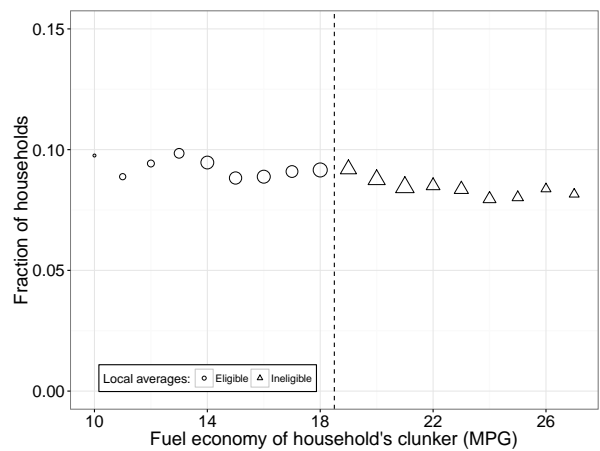


Table B.1: Estimated discontinuities for Texas households unconditional on purchasing vehicle


	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Subsidized (percent) [First-stage]	0.0100*** (0.0002)	0.0091*** (0.0003)	0.0101*** (0.0002)	0.0099*** (0.0002)	0.0089*** (0.0003)	0.0093*** (0.0003)
Spending (dollars) [Reduced-form]	-91*** (19)	-129*** (22)	-39*** (13)	-74*** (15)	-100*** (20)	-47** (20)
Spending (dollars) [2SLS]	-9,106*** (1,907)	-14,091*** (2,478)	-3,914*** (1,325)	-7,453*** (1,547)	-11,211*** (2,330)	-5,092** (2,180)
Bandwidth	5 MPG	4 MPG	4 MPG	3 MPG	2 MPG	2 MPG
Polynomial	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Controls	No	No	No	No	No	Yes
Observations	4,525,057	3,717,845	3,717,845	2,985,445	1,897,837	1,897,837

*p<0.1; **p<0.05; ***p<0.01 Each coefficient in rows 1-2 represents a separate regression of the dependent variable (in rows) on an indicator for CARS eligibility. Each coefficient in row 3 represents a separate two-stage least squares regression of the dependent variable (spending) on an indicator for CARS subsidy, instrumented for by CARS eligibility. Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

C Figures for Online Publication

Figure C.1: Fuel economy of clunkers in Texas fleet

(a) As of June 2009



(b) As of June 2008






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


(a) Number of adults in home




(b) Weekly travel days




(c) Log of annual household income



(d) Live in urban area (%)



(e) Live in single family home (%)



(f) White (%)






Figure C.3: Identification checks: Characteristics of buyers prior to CfC (July 2008-April 2009)

(a) Purchased any new vehicle



(b) Fuel economy of purchases



(c) Price of purchases






Figure C.4: Identification checks: Characteristics of buyers during July 2009 - April 2010


(a) Number of vehicles owned (June 2009)




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




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



(d) Census Tract median income



(e) Census Tract median age



(f) Census Tract white (%)






Figure C.5: Exploring alternate mechanisms for treatment effect

(a) Purchase price of vehicle greater than \$45,000



(b) Consumer sentiment at time of purchase



Data source for consumer sentiment index: University of Michigan.