



Vehicle miles (not) traveled: Fuel economy requirements, vehicle characteristics, and household driving[☆]



Jeremy West^{a, d, *}, Mark Hoekstra^{b, c}, Jonathan Meer^{b, c}, Steven L. Puller^{b, c, d}

^aUniversity of California at Santa Cruz, United States

^bTexas A&M University, United States

^cNBER, United States

^dThe E2e Project, United States

ARTICLE INFO

Article history:

Received 26 May 2015

Received in revised form 1 September 2016

Accepted 30 September 2016

Available online 18 October 2016

JEL:

L91

Q41

Q48

Keywords:

Fuel economy

Rebound effect

Regression discontinuity

ABSTRACT

A major concern with addressing the negative externalities of gasoline consumption by regulating fuel economy, rather than increasing fuel taxes, is that households respond by driving more. This paper exploits a discrete threshold in the eligibility for Cash for Clunkers to show that fuel economy restrictions lead households to purchase vehicles that have lower cost-per-mile, but are also smaller and lower-performance. Whereas the former effect can increase driving, the latter effect can reduce it. Results indicate that these households do not drive more, suggesting that behavioral responses do not necessarily undermine the effectiveness of fuel economy restrictions at reducing gasoline consumption.

© 2016 Published by Elsevier B.V.

1. Introduction

Negative externalities from gasoline consumption are well-documented, ranging from the local effects of automobile pollution on health (Currie and Walker, 2011; Knittel et al., 2011) to the global impact of vehicle emissions on climate change (Interagency Working Group, 2013). The current level of gasoline taxes in the United States is generally thought to be insufficient to correct for these externalities (McConnell, 2013), but the direct policy solution – increasing these Pigouvian taxes – often faces political resistance. Despite a large body of research recommending using a gasoline tax, U.S.

transportation policy addresses fuel consumption externalities primarily by regulating the fuel efficiency of new vehicles via Corporate Average Fuel Economy (CAFE) requirements.¹

Although CAFE standards remained largely constant for nearly two decades, the federal government has set ambitious new targets for the fuel economy of the future fleet. Regulators project that these new standards will increase the average fleet-wide fuel economy of new light-duty vehicles to 46.2 miles per gallon by 2025, compared to 25.9 miles per gallon in 2010 (NHTSA, 2012). In the absence of behavioral changes, these projections amount to a substantial reduction in gasoline consumption.

However, policy analysts argue that increasing the fuel economy of the vehicle fleet will not necessarily lead to a proportionate reduction in fuel consumption (e.g. National Research Council, 2013). The intuition underlying this concern is straightforward: because vehicles with higher fuel economy travel farther per gallon of fuel, the

[☆] This work benefited from valuable comments by editor Christopher Knittel and several anonymous referees as well as Hunt Allcott, Antonio Bento, Paul Ferraro, Ken Gillingham, Mark Jacobsen, Arik Levinson, Shanjun Li, Joshua Linn, Gregor Pfeifer, Dave Rapson, Arthur van Benthem, Matthew Zaragoza-Watkins, and numerous seminar participants. We gratefully acknowledge financial support from the National Science Foundation EV-ST5 and MIT CEEPR. Any errors are our own.

* Corresponding author.

E-mail addresses: westj@ucsc.edu (J. West), markhoekstra@tamu.edu (M. Hoekstra), jmeer@tamu.edu (J. Meer), spuller@tamu.edu (S. Puller).

¹ See Knittel (2013) for a history of the (lack of) political support for increasing the gasoline tax dating back to the Nixon administration. Extensive research examines the inefficiencies associated with using fuel economy standards rather than a gasoline tax (e.g. Portney et al., 2003; Fischer et al., 2007; Anderson et al., 2011; Jacobsen, 2013).

cost of driving each mile is lower in fuel-efficient vehicles. This lower cost-per-mile may increase the quantity of miles traveled. This has been called the “rebound effect”.

Despite the simplicity of this argument at a conceptual level, researchers have struggled to quantify the size of the rebound effect that arises from an increase in fuel efficiency (Gillingham et al., 2013a). The fundamental challenge has been a lack of exogenous variation in fuel economy. Vehicle owners select the vehicles they purchase in part based on their expected driving behavior, so disentangling the causal impact of fuel economy on driving is empirically problematic. To circumvent these endogeneity issues, most research on the rebound effect exploits variation in fuel prices – rather than fuel economy – to identify the relationship between vehicle miles traveled (VMT) and the price-per-mile of driving. There are several reasons why the impact of fuel prices on consumption may differ from the rebound effect for *fuel economy*, at least in the short run.

The primary difference between rebound effects caused by fuel prices and fuel economy is that in contrast to fuel prices, fuel economy is highly – and typically negatively – correlated with other desirable vehicle attributes, such as vehicle performance (e.g., horsepower) and safety (e.g., vehicle size). Thus, while both gasoline prices and fuel economy alter the *cost* per mile of driving, fuel economy restrictions may also affect the *benefit* per mile traveled. More formally, a change in fuel prices induces movement along the demand curve for VMT because the price-per-mile varies but vehicle characteristics are held constant. However, if a household purchases a more fuel efficient but smaller and lower-performing vehicle, then the change in vehicle characteristics shifts VMT demand *in* and the decrease in the price-per-mile moves the household *down* the demand curve. Therefore, the sign of the effect of fuel economy standards on VMT is theoretically ambiguous. As a result of this logic, we argue that variation in fuel prices is better suited to predicting the efficacy of changing gasoline taxes, but that exogenous variation in fuel economy, coupled with correlated vehicle attributes, is necessary in order to better understand the impact of fuel economy standards.

With this objective, we use household-level administrative data from Texas to study a unique natural experiment in which some households were quasi-randomly induced to buy more fuel-efficient vehicles. We do so by exploiting a discontinuity in the eligibility requirements for the 2009 U.S. “Cash for Clunkers” (CfC) program, which was a two-month program that incentivized eligible households to purchase more fuel-efficient vehicles. Specifically, we use a regression discontinuity design to assess the household driving response to the exogenous increases in new vehicle fuel economy induced by the program. Households that owned a “clunker” with an EPA-rated fuel economy of 18 MPG or less were eligible for the subsidy, while households owning a clunker with an MPG of 19 or more were ineligible. Our empirical strategy is to compare the fuel economy of vehicle purchases and subsequent vehicle miles traveled of barely eligible households to those households who were barely ineligible. The key identifying assumption is that all determinants of fuel economy and miles driven are smooth through the eligibility criteria, with the exception of program eligibility. To our knowledge, this is the first study to use quasi-experimental variation in fuel economy to estimate how household driving behavior and fuel consumption respond to policy-induced improvements in fuel economy. We find this approach to be considerably more compelling than one based on panel data, where one might worry that a change in household fuel economy over time is caused by changes in unobserved income or commute distance, which themselves would affect vehicle miles traveled.

Using data on all Texas households, we find that by May of 2010, the barely eligible and barely ineligible households were equally likely to have purchased a new vehicle, suggesting that the program pulled forward sales from the very near future. Using these estimates

of the pull-forward period, we analyze households that purchased new vehicles between the start of the program in July 2009 and the end of the pull-forward period in May 2010. Thus, by construction there was no effect of the program on the likelihood of purchase during this time window. As we show in Section 4.4, households who purchased during this time window have very similar demographic and previous purchasing and driving characteristics across the eligibility cutoff.

We find a meaningful discontinuity in the fuel economy of new vehicles purchased by CfC-eligible households relative to ineligible households. However, we also find that the more fuel-efficient vehicles purchased by the eligible households were cheaper, smaller, and lower-performing. This suggests that given the current technological limitations and the cost of fuel-saving technologies such as hybrids, households respond to fuel economy restrictions by purchasing vehicles that are more fuel-efficient, but are less desirable along other dimensions.

Analyzing subsequent driving behavior, we find that households induced to purchase more fuel-efficient (but cheaper, smaller, and lower-performing) vehicles do not drive additional miles in the two years after purchase. This is consistent with a shifting in of the VMT demand curve due to changing vehicle characteristics, coupled with a movement down the demand curve for VMT because improved fuel economy reduces the price-per-mile of driving.

Through the application of this quasi-experimental design, this paper makes three primary contributions to the literature. First, we believe this to be the first paper to exploit credibly exogenous variation in household fuel economy to identify the effect on driving behavior. As a result, we are able to obtain estimates that are causal under reasonable assumptions, without the need to impose stronger assumptions required to model vehicle purchase and driving decisions.

Second, our finding of no rebound effect is directly relevant for policies such as CAFE, given that auto manufacturers are likely to “downsize” the new vehicle fleet by selling smaller cars than they otherwise would, in order to comply with the new set of CAFE standards (Knittel (2011)). The NHTSA assumes a 10% rebound effect, based on the existing literature, when calibrating the CAFE standards (NHTSA, 2012). However, much of the existing literature on the rebound effect does not incorporate the effect of downsizing on driving. Our results suggest that if future fuel economy standards require households to downsize vehicles, then estimates of rebound that do not account for changes in vehicle characteristics are likely to be overstated. The extent to which downsizing will mitigate the rebound effect depends on how much future technological innovation softens the tradeoff between fuel economy and desirable vehicle characteristics (Anderson and Sallee (forthcoming)). Moreover, downsizing will be affected by the fact that countries with the largest car markets use attribute/size-based fuel economy standards which can mitigate some of the incentive to downsize.² The implication of our paper is that it is important to understand the extent of downsizing and incorporate it into rebound effects that are built into fuel economy standards.

Finally, these results have implications for evaluating the welfare comparisons that are frequently made between price-based policies such as a gasoline tax and quantity-based regulations such as CAFE. Quantity-based regulations such as fuel economy standards have been criticized as inefficient on the intensive margin for distorting vehicle utilization relative to the first-best policy of imposing a

² See Anderson and Sallee (forthcoming) for a detailed review of the literature and a straightforward theoretical framework that conceptualizes how different types of fuel economy standards and fuel taxes affect fuel economy, vehicle size, and vehicle miles traveled.

Pigouvian tax. This paper makes an important point: extensive margin policies can have countervailing effects on intensive marginal utilization decisions. One effect of increasing fuel economy is captured by a price elasticity of driving – altering the fuel efficiency of the fleet reduces the price-per-mile of driving. A second effect is a vehicle-attribute elasticity of driving – shifting households to fuel-efficient cars with less desirable characteristics can reduce the utility-per-mile of driving and thus the amount of driving. Both of these effects must be captured by a complete welfare analysis to compare a particular policy to first-best.³

This paper is organized as follows. Section 2 reviews the literature on the rebound effect and bolsters our argument regarding the distinction between variation in fuel prices versus fuel economy. Section 3 provides an overview of the U.S. Cash for Clunkers program, describes the data included in our study, and details our empirical strategy. Our findings are presented in Section 4, along with identification checks and falsification exercises. We conclude in Section 5.

2. The energy consumption rebound effect

Personal vehicles are a major target of U.S. energy and environmental policy. Personal light-duty vehicles generate 16% of U.S. greenhouse gas emissions and consume nearly 10% of world petroleum liquids.⁴ It is widely believed that the externalities from gasoline consumption are not fully internalized into gasoline prices (McConnell (2013)). The primary policy in the U.S. since 1978 has been the Corporate Average Fuel Economy (CAFE) standards that set minimum fuel economy requirements on new vehicles. However, many analysts and policymakers have noted that increasing the fuel economy of the fleet can result in increased driving, which can undermine the goal of reducing fuel consumption. This problem, originally called the “Jevons Paradox” and later articulated by Khazoum (1980), is a more general shortcoming of energy efficiency standards. NHTSA assumes a rebound effect of 10% in formulating CAFE standards and academic literature reviews cite rebound effects in the range from 5 to 30%; for example see Gillingham et al. (2013b), Hymel and Small (2013), and Greening et al. (2000).

The rebound effect that we study is more precisely called the “direct rebound” effect.⁵ It measures the effect of improving the energy-efficiency of a durable good on the total energy consumed by that good. To see this more formally, consider a model of a household’s choice of VMT and the resulting consumption of gasoline. Take a household with a vehicle fleet characterized by its fuel economy MPG_i and other characteristics of the vehicle X_i .⁶ VMT is an input to the production of household transportation services, hence it is

a derived demand, given by: $VMT_i = f\left(\frac{\$}{\text{mile}}_i, X_i, W_i\right)$ where $\frac{\$}{\text{mile}}_i$ is the price per mile of driving and W_i are demographic characteristics of household i . The price-per-mile of driving is the price per gallon of gasoline divided by the fuel economy in miles per gallon, so $\frac{\$}{\text{mile}}_i = \frac{P_{gas}}{MPG_i}$. Importantly, we allow for there to be a technological relationship between fuel economy and other vehicle characteristics, $X_i(MPG_i)$.

Given this setup, we can find how total gasoline consumption changes when there is an (exogenous) increase in fuel economy. A household’s total gasoline consumption is $gallons_i \equiv \frac{VMT_i \left(\frac{P_{gas}}{MPG_i}, X_i(MPG_i), W_i\right)}{MPG_i}$. Taking logs and differentiating with respect to MPG yields the elasticity of gasoline consumption with respect to fuel economy ($E_{gallon-MPG}$):

$$E_{gallon-MPG} = -1 + \underbrace{-E_{VMT-\frac{\$}{\text{mile}}}}_{\text{standard rebound}} + \underbrace{\frac{1}{gallons_i} \cdot \frac{\partial VMT_i}{\partial X_i} \cdot \frac{\partial X_i}{\partial MPG_i}}_{\text{attribute-based adjustment}} \quad (1)$$

If the two terms in braces are zero, then an increase in fuel economy leads to a one-for-one proportionate decrease in fuel consumption – there is no rebound effect. The two terms capture different behavioral adjustments that create a response that is not one-for-one. The first term captures the amount that driving increases when the price-per-mile falls but vehicle characteristics X_i remain constant. This term – which is positive – has been the focus of much of the literature that estimates rebound.

However, a second behavioral adjustment can occur that captures complementarities between vehicle attributes and energy consumption. There are a variety of channels through which vehicle characteristics can be complementary to driving. First, larger vehicles are more spacious and can make driving a more comfortable experience. Second, passengers in heavier vehicles experience lower fatality rates in the event of an accident (Anderson and Auffhammer (2014)). Finally, consumers value the improved acceleration that comes from vehicles with higher horsepower-per-pound, and generally horsepower-per-pound is lower in more fuel-efficient vehicles. As we show below, fuel economy is negatively correlated with a number of vehicle characteristics that are complementary to driving.

Visually, this decomposition of the fuel consumption response to energy efficiency improvements corresponds to both a movement along and shifting of the derived demand for gasoline. Fig. 1 provides an example. Consider a vehicle that is more energy efficient but also provides lower “performance”. The effect of the efficiency improvement $MPG'_i > MPG_i$ is to reduce the price-per-mile of driving, which shifts households down the derived demand function

³ Note that even if the net effect of the price-per-mile and vehicle-attribute elasticity is zero, this does not imply that a gasoline tax is equally efficient as a fuel economy standard, as shown in detail in Anderson and Sallee (forthcoming). Also, our empirical approach places strong emphasis on identifying causal impacts of fuel economy by exploiting quasi-random variation in fuel economy, which to our knowledge is new to the literature. A limitation of this approach is that we are not in a position to estimate the relative magnitudes of these two elasticities or to calculate welfare measures. However, our analysis does suggest that one effect can mitigate the other.

⁴ See Environmental Protection Agency (2015) and U.S. Energy Information Administration (2014).

⁵ The literature also has studied the “indirect rebound” effect which incorporates the effect of changing the efficiency of one durable good on the energy consumed by other durable goods that the household owns. See Borenstein (2015) for a detailed discussion of the different components of the total rebound effect. In this paper, we do not explore whether households receiving the subsidy and purchasing less expensive vehicles, increased energy consumption via consumption outside personal vehicle transportation.

⁶ For simplicity of exposition, assume that households own only one vehicle, but our empirical analysis will allow for multi-vehicle fleets. In addition, assume for exposition that vehicle characteristics X_i are a scalar, though more generally X_i could represent a vector of characteristics.




Fig. 1. Illustration of two components of policy-induced improvement in fuel economy.

(holding characteristics constant). But the lower performance characteristics $X_i' \prec X_i$ cause a shift 'in' of the derived demand. Depending on the size of the two effects, the net effect on fuel consumption is ambiguous.

In many formulations of the rebound effect that are used for empirical analysis, it is assumed that the energy efficiency improvement does not change any of the other attributes of the service delivered by the durable good. This implicitly assumes that this second term – the attribute-based adjustment – is zero.⁷

However, this "attribute-based adjustment" is likely to be negative in the case of vehicles, which would mitigate the size of the standard rebound effect. The more fuel efficient cars offered by manufacturers tend to have different, arguably less desirable, characteristics. As we show in Section 4.2, more fuel efficient vehicles are smaller, have less horsepower, and generally are less valuable as proxied by sales price. These tradeoffs are driven by technology – Knittel (2011) documents with historical data that improvements in fuel economy require sacrificing vehicle characteristics such as horsepower, size, and weight.⁸

If gasoline taxes were the relevant policy instrument, then the standard rebound effect is most relevant. This effect captures the impact of raising the price of gasoline via a tax, while keeping drivers in cars with the same characteristics. However, fuel economy standards are likely to change the characteristics of households' vehicle fleets. Manufacturers are likely to comply with fuel economy standards by selling vehicles that have less powerful engines, are less spacious, and are lighter.⁹ Consequently, an understanding of the effect of fuel economy standards on gasoline consumption needs to account for both the standard rebound effect and attribute-based adjustments.

The existing empirical literature has focused on estimating the standard rebound effect.¹⁰ Much of this literature exploits variation in the price of gasoline, which generates variation in the price-per-mile of driving holding vehicle characteristics constant. (In part, the rationale for exploiting changes in gasoline prices is that it provides quasi-random variation in the price-per-mile of driving, while sources of credibly exogenous variation in fuel economy are difficult to find.) Thus, the existing empirical literature on rebound, while speaking to the effects of gasoline taxes, is not well-positioned to assess the impact of fuel economy policies on driving behavior and fuel consumption.

In this paper, we estimate the net effect of both the standard rebound effect and attribute-based adjustments in the years immediately after an exogenous increase in fuel economy. This is a different form of "rebound" that addresses a different policy question than the rebound effect estimated in much of the existing literature. Gillingham et al. (2013b) refer to this form of rebound as a "policy-induced improvement" and argue that the size of this effect is more relevant for understanding the effects of energy efficiency policy such as CAFE.¹¹

It is important to note that the size of the driving response that we estimate should not be interpreted as estimating the welfare implications of energy efficiency improvements. Even if households were to respond by driving more miles, a full welfare calculation would need to account for the utility of the additional driving. Ultimately, the welfare implications depend upon whether the household response to increased energy efficiency mitigates distortions from first-best levels of driving, which is beyond the scope of this paper. This paper documents how a factor not receiving attention in the literature – the vehicle-attribute elasticity of driving – can counteract any price-per-mile elasticity of driving. This should be incorporated into both welfare analyses and to policy design that targets gasoline consumption with fuel economy standards.

3. Background and empirical strategy

3.1. The Cash for Clunkers program

We exploit the Cash for Clunkers program as a quasi-random source of variation in the fuel economy of a household's vehicle fleet. The program, formally known as the Consumer Assistance to Recycle and Save (CARS) Program, created incentives for households to replace used, fuel-inefficient vehicles with new, fuel-efficient vehicles. The program lasted for eight weeks during the summer of 2009 and offered households a rebate of \$3500 or \$4500 towards the purchase of the new fuel-efficient car when they scrapped their "clunker." A requirement of the program was that the clunker had to be taken off the road and scrapped; thus the rebate could be viewed as the trade-in value of the old car from the perspective of the household. Due to the scrappage requirement, the program attracted relatively older and low value vehicles. The average age of scrapped clunkers was 13.8 years.

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

The criteria for eligibility provide us with a cutoff for our regression discontinuity research design. The clunker must have had a combined EPA fuel economy rating of 18 MPG or less.¹² The vehicle purchased must have been a new vehicle; used vehicles did not qualify for the rebate. If the new vehicle was a passenger vehicle, it must have a combined fuel economy of at least 22 MPG. In the case of passenger vehicles, if the difference in fuel economy between the new passenger car and clunker was between 4 and 9 MPG, the rebate was \$3500, and if the difference was 10 MPG or more, the rebate was \$4500. If the new vehicle was a Category 1 Truck (e.g. SUV or small to medium pickup truck), a 2–5 MPG difference between the new truck

⁷ In some settings that have been studied this assumption may be valid, as in the case of water heaters where a more energy efficient model has more upfront cost to improve efficiency but still delivers the same volume and temperature of hot water (Allcott and Sweeney (2015)).

⁸ We should note that it is not the case that all higher fuel economy cars are smaller vehicles with less desirable characteristics. For example, the Tesla Model S (with the 2015 sticker price \$69,900) is a high performance vehicle, so purchasing a Tesla could both move down and shift out the derived demand for VMT. While there may exist some vehicles like the Tesla where increased fuel economy does not come with less desirable characteristics (other than price), there is a strong tradeoff among the vehicles that sell in large numbers. Moreover, we show that when provided with subsidies to purchase more fuel efficient vehicles during Cash for Clunkers, most households chose to downsize.

⁹ See Knittel (2011) and Klier and Linn (2012) for an analysis of the technological tradeoffs of fuel economy standards.

¹⁰ Estimates of rebound that receive considerable policy attention are from recent papers by Small and van Dender (2007) and Hymel and Small (2013). These papers use a representative consumer model that is aggregated to match state-level panel data and simultaneously model the choice of vehicles, vehicle miles traveled, and fuel economy. Surveys of research on the rebound effect include Gillingham et al. (2013b), Austin (2008) and Greening et al. (2000). In addition, a rich literature has modeled the choice and utilization of vehicles in the process of addressing a host of other policy questions; for example see Mannering and Winston (1985), Goldberg (1998), West (2004), Fullerton and Gan (2005), Bento et al. (2009), Gillingham (2012), and Allcott and Wozny (2014).

¹¹ See Gillingham et al. (2013b) for a thorough discussion of the definitions, estimation, and caveats of interpreting rebound effects.

¹² There were additional requirements that the clunker be in drivable condition, no more than 25 years old, and continuously insured and registered in the same owner's name for one year prior to the transaction.

and clunker generated a \$3500 rebate while an improvement of 5 or more MPG generated a \$4500 rebate.¹³ Busse et al. (2012) find that dealerships passed on nearly 100% of the rebates to customers.

We should note that we do not evaluate the CfC program directly; rather we use the program design as a source of quasi-random variation in fuel economy. A separate literature has evaluated how well CfC achieved program objectives with the general consensus that CfC failed both to reduce emissions in a cost-effective manner and to stimulate revenue for the auto industry. Specifically, Knittel (2009) estimates that the cost per ton of CO₂ emission reductions exceeded \$365/ton. Li et al. (2013) use a difference-in-differences with Canada to estimate a cost per ton of CO₂ ranging from \$92–\$288. Using a similar methodology as in this paper, our companion paper (Hoekstra et al., forthcoming) estimates that each subsidy averaging \$4210 reduced environmental damages by \$253 with nearly all of the benefits from CO₂. On stimulus grounds, Mian and Sufi (2012) use an instrumental variables strategy and find that the program induced more sales during the two months of the program, but those increased purchases were offset by fewer purchases in the 10 months after the program expired. Also, Hoekstra et al. (forthcoming) show that because the program primarily subsidized households who would have purchased in the near future anyway and induced those households to purchase less expensive vehicles, the program reduced total new vehicle spending by more than \$5 billion in the near term. Finally, Copeland and Kahn (2013) find minimal effects of the program on manufacturer production.

3.2. Empirical strategy

We use a regression discontinuity design to estimate the impact on vehicles miles traveled of an exogenous shift of households to more fuel efficient vehicles. We compare differences in the behavior of households whose “clunker” was barely eligible for the CfC subsidy to households whose clunker was barely ineligible. Intuitively, households that are barely eligible and barely ineligible are very similar in their preferences and driving characteristics except that the program induced barely eligible households to purchase more fuel-efficient, and as it turns out “downsized”, vehicles. As we document below, the barely eligible and barely ineligible households are very similar in a number of characteristics, which supports our identifying assumption. Importantly, we focus on new car buyers, rather than all car owners. We do this because we otherwise cannot disentangle the effect of driving a more fuel efficient car from the effect of driving a new car.¹⁴

Our empirical strategy has two steps, both of which use household-level data on vehicle ownership and utilization that we describe in Section 3.3. First, we identify the set of households who, over some time period, would have purchased a new vehicle independent of the Cash for Clunkers program. Our rationale is the following: the program may have induced some barely eligible households to accelerate their purchases to the two-month program period in order to take advantage of the subsidy. In contrast, the program did not have such an effect on the barely ineligible households. As a result, if we were to study only the households who purchased during the program, one might be concerned that the

set of barely eligible purchasers is not similar to the set of barely ineligible purchasers.

We overcome this problem by first estimating the “pull-forward window”. Beginning at the start of the program, we find the period of time over which barely eligible and barely ineligible households were equally likely to purchase a new vehicle. Because the probability of purchase over that time period was similar for barely eligible and ineligible households – by construction all were going to buy new vehicles during the window – there is little reason to expect any pre-existing differences in the composition or preferences of those barely eligible and ineligible buyers. We then focus our subsequent empirical analysis on the households that purchased during this “pull-forward window”.

The second step is to take the set of households purchasing during this pull-forward time window and compare the purchasing and subsequent driving behavior of barely eligible and barely ineligible households. The barely eligible serve as our intent-to-treat group and the barely ineligible serve as our control group. Specifically, we measure the extent to which the program induced households to purchase vehicles that are more fuel-efficient, but also smaller and lower-performance. And then we test whether the households induced to purchase these different types of vehicles subsequently drove more miles after purchase.

More formally, we compare households whose clunker was barely above the CfC eligibility cutoff of eighteen miles per gallon to those who barely qualified. We estimate the reduced-form discontinuities at the eligibility threshold using the following equation:

$$\begin{aligned} \text{Outcome}_i = & \beta_0 + \beta_1 * f(\text{distance-to-cutoff}_i) * \text{eligible}_i \\ & + \beta_2 * f(\text{distance-to-cutoff}_i) * (1 - \text{eligible}_i) \\ & + \beta_3 * \text{eligible}_i + \epsilon_i \end{aligned} \quad (2)$$

where *eligible_i* is an indicator equal to one if the household is classified as being eligible for the program (i.e., the most trade-in-likely vehicle had an MPG rating of eighteen or less). We describe how our data identify a household’s eligibility status in Section 3.3. We allow 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 barely-ineligible to barely-eligible for the Cash for Clunkers program. We use this specification to estimate both the pull-forward window and the effect of the program on the cars purchased and miles driven.

3.2.1. “Pull-forward” window

In order to estimate the pull-forward window, we follow the approach in Hoekstra et al. (forthcoming). We use a sample of all households in Texas. We estimate the number of months after the beginning of the two-month program for which the probability of purchasing a new vehicle is equalized across the eligibility threshold. We 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 during the two program months.) 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 Eq. (2) with household-level data where the dependent variable is an indicator of whether the household purchased a new vehicle during the time window. Our pull-forward window is defined as the shortest period beginning in July 2009 for which the probability of purchasing a new vehicle is equalized between the barely eligible and barely ineligible.

¹³ Separate criteria applied to Category 2 (large pickups or large vans) and Category 3 trucks (work trucks), but we do not discuss those here because there were so few of these vehicles. For a complete set of eligibility criteria, see the NHTSA rules in the Federal Register available at: <http://www.nhtsa.gov/CARS-archive/official-information/day-one.pdf>

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

Once we define this pull-forward window, the households that purchase during this window serve as the households that we include in our primary analysis. For these households, because the purchase probability is equalized, it is reasonable to assume that the Cash for Clunkers program did not affect *whether* the household purchased a new vehicle but only the *timing* and *type* of purchase within this window. Thus, there is little reason to expect differences in the underlying vehicle preferences and driving behavior of the new-car-buying households on either side of the cutoff. We provide empirical support for this assumption in Section 4.4. The pull-forward window that we estimate in Section 4.1 is 11 months. This pull-forward period is very similar to findings in other studies including Mian and Sufi (2012), Li et al. (2013), Copeland and Kahn (2013), and Hoekstra et al. (forthcoming).¹⁵

3.2.2. VMT effects of owning smaller and more fuel-efficient cars

After isolating households that purchase a new vehicle during the pull-forward window, we measure discontinuities in the types of vehicles purchased and the subsequent driving. We do so by estimating Eq. (2) with different outcome variables. First, we estimate the effect on types of cars purchased by defining the outcome variable as fuel economy and various vehicle characteristics such as horse power, curb weight, size, number of cylinders, engine displacement, and four-wheel drive. This will estimate the extent to which the program quasi-randomly shifted households into more fuel-efficient and smaller, lower performance vehicles. Second, we estimate the effect on the number of miles driven by defining the outcome variable to be annual vehicle miles traveled by the household (across all vehicles).

The identifying assumption of our analysis is that for households purchasing a vehicle over a period of time when there is no discontinuity in the probability of purchase, all household-level determinants of vehicle miles traveled after 2009 are continuous across the eligibility threshold. Under that assumption, any discontinuity in vehicle miles traveled at the cutoff is properly interpreted as the causal effect of shifting households into more fuel-efficient and downsized vehicles.

We find this identifying assumption to be reasonable for several reasons. First, the nature of the program makes manipulation very unlikely. Because households were required to own the “clunker” for one year prior to trade-in, there was little scope for households to manipulate where they were relative to the cutoff. Moreover, the fuel economy that determines eligibility is determined by the vehicle’s EPA fuel economy rating and is independent of any driving behavior by the household.

Second, we find it difficult to construct a mechanism that would violate this assumption. For example, while it is possible to imagine why barely eligible households would be different from ineligible households who bought *during the program*, it is hard to think why this would be true over this longer time horizon. By construction this longer time horizon contains a similar number of new vehicle buyers across the cutoff – the only difference is that some of those with clunkers rated at eighteen MPG or below were incentivized to purchase earlier during that time window than the other households.¹⁶

¹⁵ We should note that in this paper we use a slightly longer pull-forward window than our other paper (Hoekstra et al., forthcoming). We do so to be conservative in our estimates and ensure smoothness of unobservables across the discontinuity. By extending our window, at worst we add never-takers to our sample, which should not affect inference. We note that in our other paper, we illustrate robustness to slightly longer and shorter pull-forward windows and show that results are unchanged.

¹⁶ An example that would violate the identifying assumption is if the program were to accelerate some purchases by (say) two years, while simultaneously causing a similar number of eligible households to delay their purchases by more than a year. If that were the case – and it does seem far-fetched – the rate at which households bought vehicles over the pull-forward window might be similar across the cutoff, even though household characteristics would be different.

The identifying assumption is also consistent with empirical evidence. We show in Section 4.4 that there is no compelling evidence of discontinuities with respect to household characteristics or pre-treatment purchase and driving behavior.

3.3. Data

Our empirical setting is Texas, the second largest state in the U.S. as measured either by population or consumption of gasoline for transportation.¹⁷ We use several large administrative databases in Texas for our study.

3.3.1. Household fleets

To determine household-level vehicle fleets over time, we use confidential vehicle registration records maintained by the Texas Department of Motor Vehicles (DMV). This database allows us to identify the vehicles in a household’s fleet and when the household purchased each vehicle so that we can trace the evolution of each household’s fleet. We provide details on this process in Appendix A. In addition to providing a measure of fleets, these records include the unique vehicle identification number (VIN) for each registered vehicle. The VIN information in the DMV data allow us to measure a variety of characteristics for each vehicle, such as EPA-rated fuel economy and horsepower. We use a database obtained from DataOne Software to “decode” the VIN of each car in our sample. Importantly, our data on fuel economy is the same information that was used to determine eligibility for the CFC program.

3.3.2. Vehicle miles traveled

We compute vehicle miles traveled (VMT) for each VIN from odometer readings recorded in annual vehicle emissions tests. By comparing odometer readings between sequential emission tests of each vehicle, we calculate the miles driven over the period between tests. We use this mileage and the number of days between tests to compute the “average daily VMT” and apply this metric to each calendar date between tests. After performing this calculation for each vehicle in a household’s fleet, we add across vehicles for each calendar date to measure household-level VMT. For our analysis, we study the year beginning just after the end of the Pull-Forward Window, so we analyze VMT from June 1, 2010 to May 31, 2011. As a simple example, consider a vehicle that has two odometer readings: a value of zero at the date of sale on July 1, 2009 and a value of 27,500 at its first emissions test 730 days later. The average daily VMT is 37.67 miles per day. Thus, we estimate that the VMT is 13,750 ($= 365 \times 37.67$) for the period from June 1, 2010 through May 31, 2011. For the rest of the paper, we refer to this metric as “annual VMT”. We provide a complete discussion of VMT calculations, including several examples, in Appendix A.¹⁸

Our data on household VMT are quite complete – we observe annual VMT for over 98% of households that purchased new vehicles during the 11-month pull-forward window. To compute each vehicle’s gasoline usage, we divide its VMT by the EPA combined fuel economy of the vehicle. We sum VMT and gasoline usage over vehicles in a household’s fleet to obtain annual measures of household driving and fuel consumption.

An important institutional feature is that we study urban areas in Texas. Emissions tests are required annually in seventeen EPA non-attainment counties in Texas for each vehicle older than two years,

¹⁷ Measures of state-level gasoline consumption by end use are available from the U.S. Energy Information Administration at http://www.eia.gov/state/seds/sep_fuel/html/pdf/fuel_mg.pdf.

¹⁸ In the event that a household sold a vehicle during this time period, we only attribute miles to the household that were driven during the household’s ownership of the vehicle, which is straightforward to do as we observe the odometer reading at each sale date.

a more stringent requirement than that mandated by many states. These counties include the areas surrounding Houston, Dallas-Fort Worth, Austin, and El Paso.¹⁹ Although Texas is sometimes stereotyped as having more trucks and heavy vehicles than other states, we show in Section 4.5 that households and vehicle fleets of these urban areas of Texas do not substantially differ from urban areas across the U.S.

3.3.3. Running variable classification (clunker MPG)

We use a simple approach to classify each household's distance from the CfC 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 to qualify for a CfC subsidy. 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. This simple method of defining clunkers yields remarkably similar predictions as that using a more complex propensity score method, while requiring less completeness of data on vehicle characteristics.

3.3.4. Sample restrictions and sociodemographic covariates

In addition, we impose several sample restrictions. Because the focus of our study is on household drivers, rather than institutional fleets, we follow Knittel and Sandler (2011) in excluding a small number of households that owned more than seven vehicles as of June 2009 (just before CfC). Because CfC offered a maximum subsidy of \$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 include only households that had owned their clunker since at least July 2008 because one condition for CfC transactions was that the vehicle had been owned by the household for at least a full year.

In some specifications, we use demographic data from the Census. These data include Census tract-level economic and demographic characteristics from the 2000 decennial Census, which we link using address information in the administrative database. Finally, in tests of the identification strategy, we use a separate dataset from the spring 2009 National Household Travel Survey (NHTS). Although the NHTS does not include information allowing for direct matching to our data at the household-level, it includes a random sample of the households in Texas, so we can use the rich survey information in NHTS to test our identifying assumption.

We estimate discontinuities for households that purchased a new vehicle during the 11-month pull-forward window – the period spanning from the start of CfC in July 2009 through May 2010. As we show in Section 4.1, the barely-eligible and barely-ineligible households were equally likely to purchase a new vehicle during this time window. Summary statistics for this sample are presented in Table 1. There are 126,147 households purchasing new vehicles in our sample. The mean rated fuel economy of the new vehicles is 22.0 MPG. As far as driving behavior, the mean annual VMT for a household summed across all vehicles in the household is 31,540 miles and the mean annual gasoline consumption is 1646 gal. This table also summarizes Census Tract characteristics such as demographics and income, which we use as control variables.

Table 1
Summary statistics for new vehicle purchases July 2009–May 2010.

	Median	Mean	St. dev.
Total number of households		2,525,771	
<i>Sample: purchased new vehicle</i>			
Number of households		126,147	
Fraction of households		0.05	
<i>Characteristics of new vehicles</i>			
Fuel economy (MPG)	21	22.02	5.88
Book value (MSRP \$ '000s)	26.25	28.19	11.00
Curb weight (lbs.)	3,689	3,906	952.4
Footprint (ft ²)	56.12	61.56	14.35
Size (ft ³)	499.01	562.55	176.65
Horsepower	230	229.24	77.24
Horsepower per 1000 pounds	55.62	58.18	12.95
Engine displacement (L)	3.30	3.34	1.28
Six or more cylinders	1	0.56	0.50
4WD or AWD	0	0.31	0.46
<i>Household driving outcomes</i>			
New vehicle miles traveled ('000s)	12.76	13.75	8.14
Total annual household VMT ('000s)	27.49	31.54	21.68
Household annual fuel consumption (gal)	1,403	1,646	1,207
<i>Census Tract characteristics</i>			
Population	6,079	6,764	3,348
Median age	33.50	34.07	4.69
White (%)	81.90	76.08	18.31
Black (%)	4.20	8.96	13.94
Asian (%)	2.60	4.53	5.53
Hispanic (%)	11.60	19.75	20.50
Household size	2.92	2.86	0.42
Housing units	2,246	2,485	1,187
Owner-occupied (%)	81.00	74.64	20.08
Median income (\$)	60,192	63,380	25,970
Median home value (\$ '000s)	114.70	133.55	84.59

Notes: Statistics reported for Texas households living in a EPA nonattainment county that purchased a new vehicle either during Cash for Clunkers or during the subsequent nine months (from July 2009 through May 2010 in total). Only households with a clunker between 14 and 23 MPG (bandwidth of five) are included. The Census Tract-level characteristics are from the 2000 Decennial Census.

4. Results

4.1. Pull-forward window

The first step of our empirical analysis is to estimate the time period for which the Cash for Clunkers program did not affect the probability that a household purchased a new vehicle. The program likely induced some households that would soon be in the market for a new car to pull the sales forward so as to qualify for the subsidy. We estimate this pull-forward window and use the sample of households purchasing during this time window in our primary analysis. We have a priori reasons to believe that this set of households is very likely to satisfy our identification assumption, and we show evidence of the identification assumption in Section 4.4.

Intuitively, we find the time window, beginning with the first month of the two-month program, where households with barely eligible clunkers are equally likely to purchase a new vehicle as households with barely ineligible clunkers. Thus we start with a dataset that includes *all households* in Texas (in EPA non-attainment counties) and investigate the probability that a household purchases a new vehicle.

Results are shown in Fig. 2, which take the same form as subsequent figures. The x-axis shows the running variable of the MPG of the household's clunker, and the y-axis shows the outcome variable. Households just to the left of the vertical line own clunkers with fuel economy of 18 MPG and are barely eligible, while households just to the right of the vertical line are barely ineligible. The circles and triangles represent local averages, where the marker size corresponds to the number of households in the MPG bin.

Panel (a) of Fig. 2 shows the probability that a household in Texas purchased a new vehicle during the two months of the Cash

¹⁹ The Texas Commission on Environmental Quality (TCEQ) provided us with emissions test records for vehicles in EPA non-attainment counties in Texas. These counties include four of the largest metropolitan areas and nearly 60% of the state's population.




Fig. 2. Cumulative fraction of households purchasing any new vehicle by time period.

for *Clunkers* program. There is a clear discontinuity at the eligibility cutoff, suggesting that the program increased the likelihood of purchasing a new vehicle by more than one half of a percentage point. Thus, it is clear that Cash for Clunkers accelerated the timing of new car purchasing by the eligible households.

However, as one can see from the other panels in Fig. 2 that show progressively longer time windows, the ineligible households have an equal purchase probability by May 2010. The purchase probability is nearly equalized for the time window July 2009–March 2010 (panel (d)), and appears to be fully equalized by the late spring of 2010 (panels (e)–(f)).

Formal estimates of the discontinuities in purchase probability are shown in Table 2. In time windows up to 8 and 9 months after the

start of the program, the barely eligible are still more likely to have purchased a new vehicle. However, by 11 months after the start of the program, the barely eligible and barely ineligible households do not have statistically different purchase probabilities. Thus, for our primary analysis, we use a pull-forward window of 11 months spanning July 2009–May 2010. However, in Table 5 we show robustness of all of our major outcomes to windows ranging from 9 months to 14 months.

This method of calculating the pull-forward window assumes that the program did not have important interactions with the used car market. It is possible that households who otherwise would be in the market for a used car during the pull-forward window might take advantage of the subsidy and purchase a (perhaps inexpensive) new


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	0.0069***	0.0059***	0.0069***	0.0067***	0.0058***	0.0063***
(2 months)	(0.0005)	(0.0005)	(0.0003)	(0.0004)	(0.0005)	(0.0005)
7 months	0.0036***	0.0029***	0.0042***	0.0038***	0.0030***	0.0044***
	(0.0007)	(0.0008)	(0.0005)	(0.0006)	(0.0008)	(0.0008)
8 months	0.0028***	0.0022**	0.0035***	0.0031***	0.0023***	0.0038***
	(0.0007)	(0.0009)	(0.0005)	(0.0006)	(0.0008)	(0.0008)
9 months	0.0017**	0.0010	0.0028***	0.0023***	0.0010	0.0027***
	(0.0008)	(0.0009)	(0.0006)	(0.0006)	(0.0009)	(0.0008)
10 months	0.0009	0.00002	0.0022***	0.0016**	0.0001	0.0020**
	(0.0008)	(0.0010)	(0.0006)	(0.0007)	(0.0009)	(0.0009)
11 months	-0.0002	-0.0012	0.0014**	0.0007	-0.0010	0.0011
	(0.0009)	(0.0010)	(0.0006)	(0.0007)	(0.0009)	(0.0009)
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	2,525,771	2,085,825	2,085,825	1,686,940	1,074,239	1,074,239

*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 Eq. (2). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

car. Policy-induced substitution between the new and used vehicle market could have implications for overall gasoline consumption and emissions. If the program induced households to switch from buying a used to a new car, we would expect the barely eligible households to exhibit a lower probability of purchasing a used car than the barely ineligible during the pull-forward window. However, as we discuss in Online Appendix C, we find no evidence of this type of substitution. In Online Appendix Fig. C.1 and the correspond-

ing Table C.1, we show estimates of the discontinuity for different windows. Even during the two months of the Cash for Clunkers program, the barely eligible and barely ineligible purchased used cars at similar rates, suggesting that the program did not induce substitution between the used and new car market. And, important for our identification strategy, over the 11-month pull-forward window, the barely eligible and barely ineligible purchased used cars at very similar rates.



Note: 11 month time window July 2009 - May 2010.

Fig. 3. Reduced-form: Selected vehicle attributes of new vehicle purchases.

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

	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Fuel economy (MPG)	0.7937*** (0.1038)	0.7198*** (0.1221)	0.7569*** (0.0751)	0.7734*** (0.0853)	0.6579*** (0.1136)	0.6716*** (0.1131)
MSRP (dollars)	-1,917*** (193)	-1,796*** (227)	-2,162*** (139)	-2,000*** (159)	-1,956*** (221)	-1,660*** (213)
Curb weight (lbs.)	-175.82*** (16.45)	-160.38*** (19.02)	-159.57*** (11.70)	-162.30*** (13.10)	-153.22*** (17.22)	-150.46*** (17.11)
Footprint (ft ²)	-1.6031*** (0.2533)	-1.0035*** (0.2929)	-1.4073*** (0.1801)	-1.2881*** (0.1995)	-1.0037*** (0.2591)	-1.0961*** (0.2579)
Size (ft ³)	-20.630*** (3.104)	-15.411*** (3.589)	-17.481*** (2.207)	-17.393*** (2.451)	-13.819*** (3.199)	-15.111*** (3.180)
Horsepower	-13.659*** (1.346)	-11.941*** (1.576)	-15.380*** (0.969)	-14.586*** (1.101)	-11.565*** (1.470)	-10.740*** (1.451)
Horsepower/1000 lbs.	-1.0616*** (0.2279)	-0.7720*** (0.2699)	-1.8197*** (0.1660)	-1.5291*** (0.1910)	-0.7944*** (0.2590)	-0.6219*** (0.2564)
Engine displacement (L)	-0.1802*** (0.0222)	-0.1525*** (0.0257)	-0.1971*** (0.0158)	-0.1871*** (0.0178)	-0.1473*** (0.0234)	-0.1459*** (0.0233)
6+ cylinders	-0.0930*** (0.0086)	-0.0832*** (0.0102)	-0.1054*** (0.0063)	-0.0995*** (0.0072)	-0.0835*** (0.0095)	-0.0795*** (0.0094)
4WD or AWD	-0.0475*** (0.0081)	-0.0504*** (0.0094)	-0.0279*** (0.0058)	-0.0382*** (0.0065)	-0.0400*** (0.0085)	-0.0424*** (0.0085)
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	126,147	103,671	103,671	83,628	53,417	53,417

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$ 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 Eq. (2). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

4.2. New vehicle characteristics

Our regression discontinuity analysis shows strong evidence that the Cash for Clunkers program induced households to purchase vehicles that were both more fuel-efficient and “downsized”. Results are shown in Fig. 3. Panel (a) shows visually compelling evidence that the barely eligible purchased vehicles that were more fuel-efficient than the barely ineligible. Corresponding regression estimates are shown in the first row of Table 3. We estimate effects using bandwidths ranging from 5 MPG down to 2 MPG with both quadratic and linear functional forms.²⁰

The first row of Table 3 shows estimates of the impact of program eligibility on the fuel economy of the new vehicle. We find that eligibility induced households to purchase more fuel efficient vehicles, as one would expect due to the monetary incentive to purchase higher fuel economy cars. We find that eligibility increased fuel economy by 0.66–0.79 MPG which corresponds to about a 3.0–3.6% increase in fuel economy. This increase in fuel economy is robust to changes in the bandwidth and the inclusions of demographic covariates.

This suggests that a standard rebound effect may be present if there were no corresponding changes in vehicle attributes. However, the vehicles purchased by the barely eligible did not only differ in fuel economy – the vehicles were downsized relative to the purchases of the barely ineligible. We use various metrics of vehicle characteristics to illustrate the downsizing. First, we use Book Value as a composite measure of the value of the vehicle. Panel (b) of Fig. 3 shows that program eligibility induced households to purchase vehicles that are distinctly cheaper. Corresponding regression estimates in the second row of Table 3 show that the Manufacturer Suggested Retail Price (MSRP) for vehicles purchased by the barely eligible was between \$1600 and \$2200 lower than the vehicles purchased by the barely ineligible.

We can assess the specific vehicle characteristics that comprised the downsizing. Panels (c) and (d) of Fig. 3 illustrate two of these characteristics. The cars purchased by the barely eligible have a lower curb weight, which as shown in Anderson and Auffhammer (2014), increases the fatality risk in the event of an accident. In addition, the cars purchased by the barely eligible have less horsepower-per-pound, a proxy for driving performance.

Table 3 shows regression estimates for the effect of program eligibility on a full set of vehicle characteristics. One dimension of downsizing involves characteristics associated with comfort and safety. The barely eligible purchase vehicles with a curb weight that is 150–175 pounds lighter (see row 3 of Table 3). The footprint of the wheelbase is slightly over 1 square foot smaller (row 4). And the vehicle size, as measured by height \times width \times length, is smaller by 13–20 cubic feet.

A more complete set of performance-related characteristics related to downsizing are shown in the remaining rows of Table 3. Consistent with panel (d) of Fig. 3, horsepower-per-pound and engine displacement are significantly lower among the barely eligible car buyers. Also, the barely eligible are around 8–10% less likely to purchase a vehicle with at least 6 cylinders, and 3–5% less likely to purchase a vehicle that is 4-wheel or all-wheel drive.

These results offer compelling evidence that program eligibility caused households to purchase vehicles that are both more fuel efficient and also have characteristics that are associated with less comfort, safety, and vehicle performance. It would be interesting to understand which vehicle characteristics lead to the greatest change in derived demand for VMT. Ideally we would like to see variation in desirable car characteristics that are not strongly correlated with MPG and test whether households drive fewer miles in the vehicles with less desirable characteristics. Unfortunately, fuel economy and various car characteristics are highly correlated, especially among the most popular vehicles sold in 2009. Therefore, it is not possible to statistically disentangle which vehicle characteristics are associated with lower VMT, after controlling for fuel economy.

However, the upshot is that both individual characteristics and a composite measure of value suggest that the barely eligible downsized. If this set of vehicle attributes is complementary to the utility

²⁰ Optimal 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).




Fig. 4. Reduced-form: driving outcomes for households with new vehicle purchases.

of driving ($\frac{\partial VMT_i}{\partial X_i} > 0$ in Eq. (1)), then the increase in fuel economy could cause an attribute-based adjustment that counteracts the standard rebound effect. In the next section, we estimate the joint effect of increased fuel economy and downsizing on household vehicle miles traveled.

4.3. Household driving outcomes

We next turn to whether the households barely eligible for Cash for Clunkers were likely to drive more miles after purchasing a new vehicle, as compared to households that were barely ineligible. We analyze the effect of purchasing a new vehicle on all driving by the household.²¹ Results are shown in Fig. 4. Panel (a) shows discontinuities in the total number of annual household miles driven by all vehicles in a household’s fleet. The barely eligible households do not appear to drive more total household miles than the barely ineligible. In fact, if there is any impact, the barely eligible households drive fewer miles after purchasing a relatively more efficient and downsized vehicle.

Regression estimates of the discontinuity confirm the visual evidence that barely eligible households do not increase VMT rela-

tive to the barely ineligible, as shown in the first row of Table 4. The estimated discontinuity varies somewhat depending upon the bandwidth. But, importantly, in no specification is the effect on household VMT statistically positive, as one would expect based upon standard rebound logic. Rather, households if anything reduce VMT in response to purchasing vehicles that are more fuel efficient and downsized.

These results suggest that any standard rebound effect in Eq. (1) is counteracted by an attributed-based adjustment that reduces VMT. This can be seen even more directly by focusing solely on miles driven in the newly purchased vehicles. Panel (b) of Fig. 4 compares the VMT of the new cars purchased by the barely eligible versus the barely ineligible. Standard rebound would suggest that the barely eligible drive more miles, but the figure shows clearly this is not the case. Corresponding regression estimates of the discontinuity are shown in the second row of Table 4. We find robust evidence that the mileage of the new cars of the barely eligible (which are more fuel efficient but smaller) is less than the new cars of the barely ineligible (which are less fuel efficient but larger). This finding is consistent with an “attributed-based adjustment” that counteracts any standard rebound effect.

We also estimate the effect of program eligibility on household gasoline consumption. Given that program eligibility increases fuel economy without increasing vehicle miles traveled, the gasoline response to increases in fuel economy will not be less than proportional. The estimated effect of program eligibility on total household annual gasoline consumption is shown in the second row of Table 4. Our measure of gasoline consumption is calculated as each vehicle’s annual VMT divided by the EPA rated fuel economy and then summed across all vehicles in the household. Given the mechanical relationship between VMT, MPG, and fuel consumption, there is little surprise

²¹ A small number of households purchased more than one vehicle during the pull-forward window; we analyze the driving outcomes summed across all vehicles. Also, one might worry that our ability to measure a vehicle’s VMT could differ across the discontinuity because the barely eligible are likely to purchase earlier in the pull-forward window, which impacts when the vehicles are ‘due’ for their emission tests. We think this is unlikely, but test whether the availability of annual VMT data changes across the discontinuity. The fourth row of Table 4 shows no evidence of differences in data completeness across the eligibility threshold.

Table 4
Reduced-form estimated discontinuities for driving outcomes.

	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Total annual household VMT [mean = 31,540]	−689.7* (384.5)	−1,733.5*** (451.3)	112.0 (277.5)	−465.2 (312.8)	−1,548.7*** (407.3)	−1,566.7*** (402.1)
New vehicle miles traveled [mean = 13,750]	−652.2*** (151.4)	−736.0*** (175.2)	−439.3*** (107.7)	−588.9*** (121.9)	−546.9*** (159.4)	−618.3*** (157.1)
Household fuel consumption (gal) [mean=1646]	−62.79*** (21.31)	−105.42*** (24.80)	−29.90** (15.25)	−53.96*** (17.16)	−97.11*** (22.34)	−99.19*** (22.06)
Household VMT is unobserved [mean = 0.019]	−0.0003 (0.0024)	0.0012 (0.0028)	−0.0020 (0.0017)	−0.0007 (0.0019)	−0.0001 (0.0025)	−0.0004 (0.0025)
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	126,147	103,671	103,671	83,628	53,417	53,417

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$ 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 Eq. (2). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

that estimates in the second row are negative. These figures represent reductions ranging from 2 to 6 % across the specifications, and all of the estimates are statistically significant at the 5% level.

The results above show that the aggregate impact of exogenous increases in fuel economy coupled with downsizing does not increase VMT in our setting. We also estimate the effects separately for single- and multi-vehicle households. In Online Appendices D and E we show figures and corresponding regressions separately for single- and multi-vehicle households. Each household type was induced by the program to purchase a new car that was more fuel efficient than would have been purchased otherwise. And, for each type of household, the total household VMT and new car VMT (the two are identical for single-vehicle households) do not increase after being induced to purchase more fuel efficient and downsized vehicles. This suggests that “attributed-based adjustments” exist for both households with and without other vehicles in the household’s fleet that could serve as substitutes.

The results in this section imply that quasi-random changes in a household’s new vehicle purchase induced by the Cash for Clunkers program did not increase total annual household driving. We interpret the effect as being driven by the changes in fuel economy and associated vehicle characteristics. One might be concerned that a second factor is driving the result – our VMT measures for the barely eligible and ineligible could capture driving behavior during different periods of time. Recall that VMT data are generated from the new vehicle’s first smog check after purchase. If the barely eligible purchased substantially earlier than the barely ineligible households, then part of our estimated VMT effect could result from different macroeconomic conditions or gasoline prices during the time windows rather than from different vehicle fuel economy and associated characteristics.

However, we find no evidence that this alternative mechanism drives our VMT finding. First, we find that the barely eligible who purchased during the Pull-Forward Window purchased only around one month earlier than the barely ineligible, as shown in panel (a) of Fig. B.1.²² We confirm that the time window of VMT readings is similar

for the barely eligible and ineligible by looking at the average number of days between July 1, 2009 and the first odometer reading. The first odometer reading of barely eligible and barely ineligible households purchasing during the Pull Forward Window occur approximately one month apart as shown in panel (b) of Fig. B.1.

This implies that the time windows during which VMT is measured for the barely eligible and barely ineligible overlaps by all but one month on average. We next consider whether this (small) difference in the timing of driving is likely to be a contributing factor to our interpretation of the VMT finding. For context, aggregate driving patterns did not differ substantially over this time period. National trends in total VMT were quite flat from 2008 to 2012 (U.S. DOT, National Transportation Statistics, Table 1–36) in contrast to distinct upwards trends in driving from 1990 to 2007. We test for differences in two specific metrics that are likely to be determinants of household driving. First, we use the University of Michigan’s Consumer Sentiment Index (CSI) – a metric of how households view the economy’s strength – to measure macroeconomic effects that might be reflected in consumption decisions including driving. We use the CSI to test if consumer sentiment was substantially less for the barely eligible during the window used to create their annual VMT as compared to the barely ineligible during the window used to create their annual VMT. For each household purchasing during the Pull Forward Window, we calculate the average CSI for the calendar months that were included in the vehicle’s smog test odometer reading and call this “CSI during VMT window”.²³ If the timing of VMT readings creates systematic differences in consumer sentiment that is reflected in driving behavior, we should observe a discontinuity in our measure of “CSI during VMT window” at the eligibility cutoff. We find that the Consumer Sentiment Index during the VMT window was nearly identical for the barely eligible and barely ineligible, as seen in Fig. B.2.

Another factor – gasoline price variation over time – is also unlikely to drive our results. As shown in Fig. B.3, gasoline prices were actually relatively lower during the VMT measurement windows of the barely eligible households, which would work opposite our finding, but again the differences are small due to the small differences in the time windows. This analysis leads us to conclude that the primary difference between the treatment and control groups is the type of vehicle the household drove and not the time period during which the VMT was recorded.

²² For example, the figure shows that the average household with an 18 MPG clunker that purchased during the Pull Forward Window purchased its new vehicle 137 days after the start of the program (which corresponds to November 15, 2009). The average household with a 19 MPG clunker that purchased during the Pull Forward Window purchased 158 days after July 1 (which corresponds to December 6, 2009). This is driven by the fact that while some households were pulled forward from as far as 11 months in the future, the vast majority of households were pulled forward from the months during and just after the program. For this reason, while the “last” household was pulled forward from 11 months in the future, the average household was pulled forward by closer to one month. See Fig. 3 from our companion paper Hoekstra et al. (forthcoming) that estimates the months from which sales were pulled forward.

²³ For example, consider a household that purchased in July 2009 and had its first emissions test in June 2011. The “CSI during VMT window” is the average CSI for July 2009–June 2011.

Table 5
Robustness of estimated discontinuities to alternate time windows.

	Time window in months					
	9 months	10 months	11 (main)	12 months	13 months	14 months
New vehicle fuel economy (MPG)	0.8348*** (0.1300)	0.7604*** (0.1233)	0.7251*** (0.1177)	0.6099*** (0.1130)	0.5588*** (0.1092)	0.4980*** (0.1065)
Total annual household VMT	-1,716*** (439.7535)	-1,677*** (419.9054)	-1,567*** (402.1443)	-1,582*** (385.6548)	-1,496*** (371.5331)	-1,462*** (356.9656)
New vehicle miles traveled	-693.03*** (172.5030)	-621.21*** (165.3874)	-618.28*** (157.0661)	-618.3*** (151.2200)	-573.6*** (145.2261)	-559.93*** (140.8563)
Household fuel consumption (gal)	-116.68*** (24.0997)	-109.26*** (23.0204)	-99.19*** (22.0644)	-96.59*** (21.1822)	-90.93*** (20.4185)	-87.49*** (19.6219)
Household VMT is unobserved	-0.0008 (0.0027)	-0.0012 (0.0026)	-0.0004 (0.0025)	-0.0004 (0.0023)	-0.000005 (0.0022)	-0.0002 (0.0022)
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	44,617	48,882	53,417	57,720	61,782	66,526

*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 Eq. (2). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

4.4. Tests of the identification strategy

The identifying assumption of our empirical strategy is that all determinants of VMT in the year after purchase vary smoothly across the Cash for Clunkers eligibility cutoff among the households that purchased new vehicles from July 2009 to May 2010. In this section we test for a variety of potential threats to identification.

The identification assumption could be violated, for instance, if politicians endogenously selected the 18 MPG eligibility threshold based on types of vehicles that would qualify. We test this assumption in several ways. We use data from the 2009 National Household Travel Survey to test whether there are discontinuities in demographic characteristics among the population of vehicle owners. As shown in Fig. B.4, there is little compelling visual evidence of discontinuities in vehicle owner characteristics such as number of adults in household, number of weekly travel days, household income, living in an urban area, living in a single-family house, or race, which is consistent with our identifying assumption.

As described earlier, our analysis focuses on new car buyers, rather than all car owners. We do this in part because we otherwise cannot disentangle the effect of driving a more fuel-efficient vehicle from the effect of driving a new vehicle. As a result, our identifying assumption requires that for households that bought a new vehicle during the 11-month pull-forward window, all determinants of VMT in the year after purchase vary smoothly across the eligibility cutoff. We implement several tests where we focus on new car buyers using our DMV registration data.

Our first test examines whether the households that purchased during the pull-forward window were different in the year prior to the Cash for Clunkers program. Results are shown in Fig. 5 with corresponding regression estimates in Table 6. The first row of the table shows the discontinuity in the fuel economy of the household's vehicle fleet *excluding the clunker*.²⁴ Estimates indicate that if anything, barely eligible households may have a preference for slightly lower MPG vehicles, which suggests that our treatment effect estimates presented earlier may somewhat understate the increase in new vehicle fuel economy due to the program. However, the estimated difference is relatively small as compared to differences in MPG induced by Cash for Clunkers.

We also test if total household VMT in the year prior to Cash for Clunkers differed across the eligibility threshold. As seen in panel (b)

of Fig. 5, the prior year's driving is quite noisy for households with clunker MPG very close to 18. As a result, corresponding estimates in Table 6 vary substantially across bandwidths and specifications, making it difficult to precisely estimate the pre-treatment VMT difference. For example, while the estimated discontinuity in VMT is statistically significant -1443 when using a bandwidth of 4 and a quadratic fit, changing to a linear fit results in a marginally significant estimate of (positive) 435. The lack of a robust discontinuity is consistent with the raw data shown in panel (b) of Fig. 5, which shows no visually compelling evidence that barely eligible households drove more or fewer miles than barely ineligible households.

As another identification test, we consider the possibility that there is some general underlying difference between new car buyers with "clunkers" on either side of the 18 MPG threshold. Here, we test whether new-car buying households just below the cutoff *always* tend to buy vehicles that are smaller and more efficient. To test for this possibility, we analyze the purchase and driving behavior of households in 2008 – the year prior to Cash for Clunkers. Fig. B.5 and Table B.2 show new vehicle characteristic discontinuity estimates for households that purchased a new vehicle in calendar year 2008 as a function of the household's "clunker" in 2008.²⁵ Visually there are no discontinuities in new vehicle characteristics in Fig. B.5. While some of the formal estimates in Table B.2 are statistically significant, they are sensitive to specification and are all economically small. For example, only three of the six estimates for curb weight are statistically significant at the 5 % level, and even the largest discontinuity (-56.8 pounds) is less than one third the size of the (robustly) estimated effect in Table 3.

Results for driving outcomes of 2008 purchasers are shown in Fig. B.6, with corresponding regression estimates shown in Table B.3. As with the 2009 purchasers, VMT is noisy for households with clunker MPG close to 18, so discontinuity estimates are sensitive to bandwidth and functional form. For example, the quadratic specification with a bandwidth of 4 shown in row (1) of column (2) results in a statistically significant estimate of -1279, while a linear specification with the same bandwidth results in a statistically significant estimate of (positive) 1150. Overall, however, it is clear from the raw data shown in panel (a) of Fig. B.6 that households with clunker MPG below 18 do not tend to drive less than households with clunker MPG

²⁴ If we included the fuel economy of the clunker – which defines the running variable – then mechanically the relationship would be smooth through the 18 MPG clunker threshold.

²⁵ To define the household's clunker in 2008, we use a similar approach to define the clunker, except that we use the oldest vehicle in a household's vehicle stock as of December 2007.




Fig. 5. Identification check: Discontinuities in pre-treatment characteristics for purchases during July 2009–May 2010.

above 19, which suggests that we are not biased towards finding lower VMT among the barely eligible.²⁶

4.5. External validity

Finally, we consider the extent to which the households that we study are representative of the new car buying population of the U.S. Our estimated LATE applies to a particular subpopulation – households in Texas MSAs who owned a used car with fuel economy near 18 MPG (depending on bandwidth used in the RD specification) and purchased a new car during the 11-month period July 2009–May 2010. We assess external validity by using the 2009 National Household Transportation Survey which contains household and driving characteristics. We cannot directly link our DMV registration data to the NHTS, so we choose NHTS households that resemble our sample. Specifically, we focus on NHTS households who lived in an MSA in 2009 that owned a “clunker” within 2 MPG of the eligibility cutoff and that recently purchased a new vehicle. Comparisons are shown in Table B.1. The first row reports household characteristics from NHTS that capture the types of households in our sample – households in Texas MSAs who (b) own a “clunker” within 2 MPG of the cutoff. In contrast, the second two rows show households outside of Texas, and enable us to address the two primary issues of external validity – whether urban Texas is representative of urban areas in the U.S. and whether households owning used vehicles near

the 18 MPG cutoff are representative of car owners in general. The second row addresses the first issue of whether urban Texas is representative of urban areas in the U.S. Results indicate that households are quite similar in the number of drivers and the number of weekly travel days. However, Texas households drive 4% more miles and own vehicles that are 1.6% less fuel efficient. The third row addresses the second issue of external validity by showing all households in non-Texas MSAs, regardless of clunker MPG. As compared to those households, Texas households with clunker MPG near the eligibility cutoff have slightly fewer drivers but slightly more weekly travel days. In addition, our households have somewhat fewer vehicles and lower MPG, but drive a similar number of miles. Overall, we conclude based on Table B.1 that while there are some small differences between the households in our sample and households away from the eligibility threshold and outside of Texas, by and large the new car buyers we study are similar to those in other urban areas in the U.S.

5. Conclusion

A critical energy policy question is whether increases in fuel economy will increase miles driven and thus partially mitigate gains from efficiency improvements, which will exacerbate externalities associated with driving and gasoline consumption. To our knowledge, this is the first paper to address this question using quasi-random variation in a household's fuel economy. We show that while households who were barely eligible for the subsidy purchased significantly more fuel efficient and downsized vehicles, they did not respond by driving more miles. As a result, in our setting we find that there is no evidence of a rebound effect that offsets the reduction in fuel consumption.

²⁶ In addition, we note that even if one differences out the estimates in row (1) of Table B.3 from the corresponding estimates in row (1) of Table 4, there is still no evidence of a positive rebound effect on VMT, as the most positive adjusted estimate would be -220 miles.

Table 6
Identification check: Discontinuities in pre-treatment characteristics for purchases during July 2009–May 2010.

	Estimated discontinuity					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-clunker household MPG	−0.2689*** (0.0804)	−0.5331*** (0.0944)	0.0736 (0.0581)	−0.1680** (0.0657)	−0.3130*** (0.0863)	−0.2981*** (0.0859)
Total annual household VMT	−633.35* (366.72)	−1,443.32*** (430.19)	435.21* (264.54)	−67.56 (300.57)	−1,530.13*** (402.00)	−1,609.94*** (397.22)
Annual fuel consumption (gal)	−31.050 (21.028)	−62.914** (24.465)	7.243 (15.043)	−9.026 (16.984)	−73.634*** (22.613)	−78.765*** (22.360)
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	126,147	103,671	103,671	83,628	53,417	53,417

*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 Eq. (2). Columns vary the bandwidth and included control terms. Standard errors are reported in parentheses.

This paper has important implications for policies that target gasoline consumption with fuel economy standards. If future fuel economy standards require households to downsize vehicles, then the standard rebound effect is likely to be mitigated due to attribute-based adjustments in driving. Using rebound estimates that hold vehicle characteristics constant can overstate the driving response to fuel economy standards. The “policy-relevant rebound effect” includes not only how households respond to a lower price-per-mile but also the response to less desirable vehicle attributes.

Thus, the critical issue for policymakers is to assess the extent to which technological innovation will relax the tradeoff between fuel economy and desirable attributes without substantially increasing vehicle price. Historically there has been a tradeoff between improvements in fuel economy and characteristics such as horsepower, size, and weight. The question is whether this tradeoff will be strong in the future, in which case policymakers need to account for the attribute-based adjustments when making assumptions about rebound. On the other hand, if fuel economy can be increased without large attribute sacrifices, then these adjustments are likely to be small. That being said, it is important to note that if the tradeoff is only relaxed at high prices that are outside the range of most household budgets, then many households may still face the tradeoff and choose to downsize. It is beyond the scope of this paper to assess the likely technological tradeoff in the future. But the policy upshot is clear – the assumed rebound effect that should be built into fuel economy standards needs to account for a vehicle-attribute response in driving that incorporates a realistic assessment of the expected future technological tradeoff.

Appendix A. Appendix – Data

A.1. Defining a household’s fleet

The Texas Department of Motor Vehicles (DMV) provided us with confidential access to all Texas vehicle registrations for the years spanning our study. From these records, we attribute individual vehicles to households as follows. First, we used ESRI’s ArcMAP software to geocode 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. 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.²⁷ After determining pairs of vehicles belonging to the same household, we chain these connections to allocate the population of vehicles to households for each date included in our data.

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

A.2. Calculating household VMT

We calculate vehicle miles traveled for each unique vehicle (VIN17) using three sources of odometer readings. Primarily, we use data from annual vehicle emissions tests/safety inspections conducted in the seventeen EPA non-attainment counties in Texas, which were provided to us by the Texas Commission on Environmental Quality (TCEQ) for January 1, 2004 through August 20, 2012. In these counties, Texas law requires personal vehicles to undergo emissions testing annually beginning at the vehicle’s second year.²⁸ New residents are allowed thirty days to obtain a vehicle emissions test. We augment these odometer readings with data from the Texas

²⁷ We use a dynamic Levenshtein distance metric to match last names. First, we trim each of the two last name fields to fifteen letters. Then, we match them pairwise using a Levenshtein critical value of 0.34. The most common entry errors for names in the database are omitted letters (an L-distance of one) and transposed letters (an L-distance of two). For a six letter last name, an L-distance of two requires a critical value of 0.34 to correct. A nine letter last name is allowed three transformations under this critical value.

²⁸ The annual emissions inspection requirement is waived for vehicles older than twenty-four years. More information on Texas emissions testing requirements is provided by the Texas Department of Public Safety at <http://www.txdps.state.tx.us/InternetForms/Forms/VI-51.pdf>

DMV database, which reports the odometer value for each vehicle transaction involving a Texas buyer. Finally, for a fairly small set of vehicles we append odometer readings reported to the U.S. DOT for vehicles scrapped in the Car Allowance Rebate System (CARS, or “Cash for Clunkers”).²⁹

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

We use these odometer readings for each VIN to calculate the “daily” vehicle miles traveled on every observed date of each vehicle’s lifetime. Then, we aggregate these daily VMT measures for each VIN to a longer calendar period of interest. For our analysis of post-treatment VMT, we use the period from June 1, 2010 through May 31, 2011, the first full calendar year subsequent to the pull-forward purchase window. In the event that a household sold a vehicle during this time period, we only attribute miles to the household that were driven during the household’s ownership of the vehicle, which is straightforward to do as we observe the odometer reading at each sale date. In the rare case that the final observed odometer reading for a VIN occurs prior to the end of this time window (while the household’s ownership of the vehicle continues) we extrapolate the most recent daily VMT throughout the remainder of the analysis period.³⁰

We estimate the quantity of gasoline consumed by each household in each vehicle by dividing the total VMT driven in that vehicle in the analysis period by the vehicle’s EPA rating for combined fuel economy. As with VMT, we only attribute fuel consumption to the household that occurred during the household’s ownership of the vehicle.

Fig. A.1 shows example patterns of odometer readings for three hypothetical new vehicles purchased on July 1, 2009. The analysis period from June 1, 2010 through May 31, 2011 is indicated by the vertical lines and odometer readings by the solid circles.

Vehicle A traveled 27,500 miles before its first emissions test on July 1, 2011. Dividing these VMT by the elapsed 730 days yields 37.671 miles per day. Thus, we estimate Vehicle A traveled 13,750 miles during the one-year analysis period.

Vehicle B traveled 19,514 miles before its first (and only) emissions test on December 1, 2010. Dividing these VMT by the elapsed 518 days yields 37.672 miles per day. Thus, we estimate Vehicle B




Fig. A.1. Example odometer reading timelines.

also traveled 13,750 miles during the one-year analysis period. Note the extrapolation that is present in (only) these types of cases.

Vehicle C traveled 15,510 miles before its first (of two) emissions test on December 1, 2010. Dividing these VMT by the elapsed 518 days yields 29.942 miles per day during the first half of the analysis period. Vehicle C then traveled 16,571 miles before its second emissions test on December 1, 2011. Dividing these VMT by the elapsed 365 days yields 45.4 miles per day during the second half of the analysis period. Thus, we estimate Vehicle C also traveled 13,750 miles during the one-year analysis period.

Note that in all three examples the VMT during the analysis period is 13,750 miles, despite the different nature of the odometer readings.

Appendix B. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jpubecon.2016.09.009>.

References

- Allcott, H., Sweeney, R., 2015. Can retailers inform consumers about energy costs? Evidence from a field experiment. January. Working paper.
- Allcott, H., Wozny, N., 2014. Gasoline prices, fuel economy, and the energy paradox. December. *Rev. Econ. Stat.* 96 (10), 779–795.
- Anderson, M., Auffhammer, M., 2014. Pounds that kill: the external costs of vehicle weight. *Rev. Econ. Stat.* 81 (2), 535–571.
- Anderson, S.T., Parry, I.W.H., Sallee, J.M., Fischer, C., Winter 2011. Automobile fuel economy standards: impacts, efficiency, and alternatives. *Rev. Environ. Econ. Policy* 5 (1), 89–108.
- Anderson, S.T., Sallee, J.M., 2016. Designing policies to make cars greener: a review of the literature. *Ann. Rev. Resour. Econ.* forthcoming.
- Austin, D., 2008. Effects of gasoline prices on driving behavior and vehicle markets. January. Congressional Office Report 2883.
- Bento, A.M., Goulder, L.H., Jacobsen, M.R., von Haefen, R.H., 2009. Distributional and efficiency impacts of U.S. gasoline taxes. *June. Am. Econ. Rev.* 99 (3), 667–699.
- Borenstein, S., 2015. A microeconomic framework for evaluating energy efficiency rebound and some implications. *Energy J.* 36 (1), 1–21.
- Busse, M.R., Knittel, C.R., Silva-Risso, J., Zettelmeyer, F., 2012. Did “Cash for Clunkers” deliver? The consumer effects of the Car Allowance Rebate System. November. Working paper.
- Calonico, S., Cattaneo, M.D., Titiunik, R., 2015. Robust nonparametric confidence intervals for regression-discontinuity designs. November. *Econometrica* 82 (6), 2295–2326.
- Copeland, A., Kahn, J., 2013. The production impact of “Cash for Clunkers”: implications for stabilization policy. January. *Econ. Inq.* 51 (1), 288–303.
- Currie, J., Walker, R., 2011. Traffic congestion and infant health: evidence from E-Zpass. January. *Am. Econ. J. Appl. Econ.* 3 (1), 65–90.
- Energy Information Administration, U.S., 2014. Light-duty vehicle energy demand: demographics and travel behavior. Issues in Focus
- Environmental Protection Agency, 2015. Inventory of U.S. greenhouse gas emissions and sinks: 1990–2013. February. Draft Edition

²⁹ The CARS data are available from the www.nhtsa.gov National Highway Safety Traffic Safety Administration.

³⁰ We restrict such extrapolation to at most one full calendar year. In light of non-compliance, households moving out of emissions testing counties, and other factors precluding odometer observations, we view this as a reasonable trade-off. The overall fraction of VMT determined using such extrapolation is quite small.

- Fischer, C., Harrington, W., Parry, I.W.H., 2007. Should automobile fuel economy standards be tightened? December. *Energy J.* 28 (4), 1–29.
- Fullerton, D., Gan, L., 2005. Cost-effective policies to reduce vehicle emissions. May. *Am. Econ. Rev. Pap. Proc.* 95 (2), 300–304.
- Gillingham, K., 2012. Selection on anticipated driving and the consumer response to changing gasoline prices. Yale University Working paper.
- Gillingham, K., Kotchen, M.J., Rapson, D.S., Wagner, G., 2013. Energy policy: the rebound effect is overplayed. January. *Nature* 493, 475–476.
- Gillingham, K., Rapson, D., Wagner, G., 2013b. The rebound effect and energy efficiency policy. Working paper.
- Goldberg, P.K., 1998. The effects of the Corporate Average Fuel Efficiency standards in the U.S. March. *J. Ind. Econ.* 46 (1), 1–33.
- Greening, L.A., Greene, D.L., Difiglio, C., 2000. Energy efficiency and consumption, the rebound effect: a survey. June. *Energy Policy* 28 (6–7), 389–401.
- Hoekstra, M., Puller, S.L., West, J., 2016. Cash for Corollas: when stimulus reduces spending. *Am. Econ. J. Appl. Econ.* forthcoming.
- Hymel, K.M., Small, K.A., 2013. The rebound effect for automobile travel: Asymmetric response to price changes and other quirks of the 2000s. University of California at Irvine Working Paper.
- Interagency Working Group, 2013. Technical update on the social cost of carbon for regulatory impact analysis - under executive order 12866. May. Interagency Working Group on the Social Cost of Carbon
- Jacobsen, M.R., 2013. Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. May. *Am. Econ. J. Econ. Pol.* 5 (2), 148–187.
- Khazzoum, D.J., 1980. Implications of mandated efficiency in standards for household appliances. *Energy J.* 1 (4), 21–40.
- Klier, T., Linn, J., 2012. New-vehicle characteristics and the cost of the Corporate Average Fuel Economy standard. Spring. *RAND J. Econ.* 43 (1), 186–213.
- Knittel, C.R., 2009. The implied cost of carbon dioxide under the Cash for Clunkers program. August. Center for the Study of Energy Markets Working Paper.
- Knittel, C.R., 2011. Automobiles on steroids: product attribute trade-offs and technological progress in the automobile sector. December. *Am. Econ. Rev.* 101 (7), 3368–3399.
- Knittel, C.R., 2013. Transportation fuels policy since the OPEC embargo: paved with good intentions. May. *Am. Econ. Rev. Pap. Proc.* 103 (3), 344–349.
- Knittel, C.R., Miller, D., Sanders, N.J., 2011. Caution, drivers! Children present. Traffic, pollution, and infant health. NBER Working Paper No. 17222.
- Knittel, C.R., Sandler, R., 2011. Cleaning the bathwater with the baby: the health co-benefits of carbon pricing in transportation. November. NBER Working Paper No. 17390.
- Li, S., Linn, J., Spiller, E., 2013. Evaluating Cash-for-Clunkers: program effects on auto sales and the environment. March. *J. Environ. Econ. Manag.* 65 (2), 165–193.
- Manning, F., Winston, C., 1985. A dynamic empirical analysis of household vehicle ownership and utilization. Summer. *RAND J. Econ.* 16 (2), 215–236.
- McConnell, V., 2013. The new CAFE standards: Are they enough on their own? May. RFF Discussion Paper 13-14.
- Mian, A., Sufi, A., 2012. The effects of fiscal stimulus: evidence from the 2009 “Cash for Clunkers” program. August. *Q. J. Econ.* 127 (3), 1107–1142.
- National Research Council, 2013. *Transitions to Alternative Vehicles and Fuels*. The National Academies Press.
- NHTSA, 2012. Corporate average fuel economy for MY 2017-MY 2025 passenger cars and light trucks. U.S. Department of Transportation: National Highway Traffic Safety Administration
- Portney, P.R., Parry, I.W., Gruenspect, H.K., Harrington, W., 2003. The economics of fuel economy standards. Fall. *J. Econ. Perspect.* 17 (4), 203–217.
- Small, K.A., van Dender, K., 2007. Fuel efficiency and motor vehicle travel: the declining rebound effect. *Energy J.* 28 (1), 25–51.
- West, S.E., 2004. Distributional effects of alternative vehicle pollution control policies. March. *J. Public Econ.* 88 (3–4), 735–757.