Attribute Substitution in Household Vehicle Portfolios

James Archsmith
University of Maryland

Kenneth Gillingham
Yale

Christopher R. Knittel
MIT

David S. Rapson*
UC Davis

March 20, 2020

Abstract

Roughly three quarters of vehicles are purchased into multi-car households. We study whether households are willing to substitute attributes, such as fuel economy, across vehicles within their portfolio. We develop a novel strategy to separately identify idiosyncratic preferences for an attribute from these within-portfolio effects. Using the universe of household vehicle registration records in California over a six-year period, we find that two-car households exhibit strong substitution across vehicles when faced with an exogenous change to fuel intensity of a kept vehicle. This effect can erode a substantial portion of the benefit from major policies, such as Cash-for-Clunkers.

*Archsmith: Assistant Professor, Agricultural and Resource Economics, University of Maryland, College Park, archsmit@umd.edu. Gillingham: Associate Professor, Yale University and NBER, kenneth.gillingham@yale.edu. Knittel: George P. Shultz Professor of Applied Economics, Sloan School of Management, and Director of the Center for Energy and Environmental Policy Research, MIT, and NBER, knittel@mit.edu. Rapson: Associate Professor of Economics, UC Davis, dsrapson@ucdavis.edu. This paper has benefited from conversations with Mark Jacobsen, JF Houde, and seminar participants at Universidad Carlos III de Madrid, CESifo, ETH Zurich, University of Mannheim, European University Institute, London School of Economics, Toulouse School of Economics, Cornell, Fondazione Eni Enrico Mattei, the Energy Institute @ Haas Summer Energy Camp, AEA Meetings, IIOC, and NBER EEE Spring Meetings. We gratefully acknowledge financial support from the California Air Resources Board. The statements and conclusions presented are those of the authors and not necessarily those of the California Air Resources Board. The mention of commercial products, their source, or their use in connection with material reported herein is not to be construed as actual or implied endorsement of such products.

1 Introduction

In many settings, the sequential nature of decisions generates correlation between attributes of portfolio goods purchased over time (e.g., financial assets, household durables such as electronics, art or decor, clothing fashion, media subscriptions, higher education choices and more). One implication of this correlation is that the demand for products within a household will not be independent. If consumers view attributes across multiple goods as a bundle, policies that alter the attributes of one of these goods (e.g., energy efficiency standards in the context of energy-consuming durables) may set in motion a sequence of complicated substitution patterns that have important efficiency implications.

This paper examines attribute substitution by households in the context of multi-vehicle portfolios.[1] Our main motivation is to understand certain features of demand in the context of household vehicle *portfolio* choice and consequent implications for policies aimed to reduce gasoline demand and the accompanying external costs; but our contributions extend beyond this initial motivation. Identification of portfolio-level interdependence across vehicles is difficult because households with a strong preference over a particular product attribute will exhibit that preference in each sequential purchase. Complementarity in that attribute across vehicles would look similar to the researcher. This poses a challenge in identifying how changing the attributes of a household's existing vehicle affects the attributes of the next vehicle purchased.[2]

We make three main contributions. First, we develop a new identification strategy to separately identify the buyer's preferences for attributes from the substitutability or complementarity between attributes of the goods in a portfolio. Second, using a comprehensive dataset of vehicles and drivers in California, we demonstrate that household decisions in a vehicle portfolio are not independent as is generally assumed in durable goods demand estimation. Third, we demonstrate that household portfolio considerations have important implications for policies intended to alter the characteristics of vehicles purchased by multi-car households. The effect of policies intended to mitigate negative external costs from gasoline consumption—e.g. fuel economy standards and so-called "scrap-and-replace" programs—may be partly undone by the extensive and intensive margin substitution patterns exhibited by households.

Empirical models used to analyze the costs and benefits of fuel economy standards often capture

[1] The unconditional mean number of vehicles per household in the United States in 2014 was 2.09 (Oak Ridge National Laboratory), and in California roughly 75 percent of all vehicles are owned by households with two or more vehicles (California Department of Motor Vehicles).

[2] Researchers examining portfolio interactions in a number of settings face analogous identification challenges. There is a literature that suggests households engage in attribute substitution when it comes to children. Ben-Porath and Welch (1976) and Angrist and Evans (1998) show that households that have had two children of the same gender are more likely to have a third child, compared to two-children households endowed with one boy and one girl. This suggests that the attributes of the first two children, namely gender, affect the utility from having a third child. There is also evidence that households trade off, or substitute, the characteristics of occupations across spouses. For example, households may avoid having two occupations in the same sector as a way to reduce risk (see e.g., Udry et al. (1995) and Ellis (2000)). While intuitive, identification of such a phenomena in the broader labor market is difficult because matching costs may be lower within an occupation or sector; physicians tend to meet other physicians and not economists.

many key drivers of patterns of vehicle demand, but typically assume away interactions between preferences for multiple vehicles within a household. That is, these models of fuel economy standards assume that consumers choose only a single vehicle or, alternatively, that the choice of each vehicle in a household is independent of the others (Bento et al., 2009; Goulder, Jacobsen, and van Benthem, 2012; Jacobsen, 2013).

However, there are likely to be two sources of correlation in this choice. The first is that households may have particularly strong preferences for certain vehicle attributes, a feature that is captured in empirical models that allow for variation in the willingness to pay for vehicle attributes (e.g., Berry, Levinsohn, and Pakes (1995)). For example, the choice of fuel economy across vehicles within a household will be positively correlated for a household that particularly values horsepower. The second source of correlation is that household preferences may exhibit complementarities between portfolio goods, as suggested by Gentzkow (2007) for newspapers, Wakamori (2011) for Japanese vehicle purchasers, and Manski and Sherman (1980) for a small sample of U.S. vehicle purchasers.[3]

While we suspect that the second form of interdependence may not compromise the consistency of the parameter estimates of, for example, the mean and standard deviation of willingness to pay for attributes, it is likely to lead to biased predictions from policy counterfactuals related to fuel economy standards and gasoline taxes. For example, suppose a policy were to increase the chosen fuel economy of the newest vehicle for a given household at time t. When the household subsequently replaces the other vehicle at a later date, attribute substitution across vehicles (which Wakamori (2011) refers to as "complementarity" in the vehicle portfolio) could lead them to purchase a more fuel *intense* follow-on vehicle than they might have otherwise. Because fuel economy is correlated with other attributes, when there are attribute-based standards, such as the footprint-based standards in the United States (Gillingham, 2013; Ito and Sallee, 2014; Kellogg, 2017), this effect could further drive a wedge between the economic efficiency of fuel economy standards and Pigouvian gasoline taxes. Pigouvian gasoline taxes would still be economically efficient in the presence of attribute-substitution effects in portfolios, although these patterns may affect the counterfactual level of emissions reductions.

The ideal experiment to identify attribute substitution would randomly assign the "kept" vehicle attributes to households in the market for a vehicle, and then observe the relationship between an attribute, such as fuel economy, of this kept vehicle and that of the newly-acquired vehicle. Since this ideal experiment is obviously not feasible, our identification strategy must overcome two potential sources of endogeneity stemming from the non-random assignment of the kept vehicle. The first is the choice of which vehicle to replace. Since the household preference for particular features of a multi-car portfolio will directly inform the decision of which car to keep or drop, the attributes of the kept car are endogenous. The second challenge to causal identification is related to the

[3] Wakamori (2011) is the closest intellectual antecedent to this work. Wakamori focuses on the complementarity of a small automobile and a minivan in context of the Japanese car market. Manski and Sherman (1980) examine household demand for vehicles and assume complementarities between the vehicles.

presence of unobserved household preferences for vehicle attributes. Household fixed effects allow us to focus on within-household variation and can address time-invariant unobserved preferences, but there would still be a concern if preferences change over time. Time-varying preferences may produce a correlation between the desired attributes of the kept and newly-acquired vehicle, again implying that the attributes of the kept vehicle are endogenous.

To account for these potential sources of bias, we use an instrumental variables strategy, where we instrument for the household's kept vehicle using the gasoline price at the time of the purchase of the kept vehicle. A number of papers (Klier and Linn, 2012; Allcott and Wozny, 2014; Busse, Knittel, and Zettelmeyer, 2013; Gillingham, 2011) have shown that vehicle purchase behavior is influenced by contemporaneous gasoline prices. Given this literature and the finding in Anderson, Kellogg, and Sallee (2013) that consumers tend to use the current gasoline price in forming expectations of future prices, we would expect the fuel economy of the kept vehicle to be influenced by the gasoline price at the time of that purchase. We argue that this instrument for the kept-vehicle fuel economy satisfies the exclusion restriction because, after controlling for current gasoline price, past gasoline prices should not influence the choice of the new vehicle. This assumption rests on limited serial correlation in the residuals and assumes that consumers are using the contemporaneous gasoline price to form expectations of future gasoline prices.

Additional specifications to investigate the role of attribute substitution across vehicle types in a portfolio require additional instruments for identification. We develop instruments based on the idea that changes in the relative price of cars in a portfolio systematically affect the probability that the lowest fuel economy car is dropped. These changes occur in the gaps of time between portfolio decisions, rendering them exogenous with respect to the composition of the portfolio at the time when a previous replacement decision is made. Our preferred instrument is constructed from deviations from the expected change in relative vehicle prices since the time when the kept car was initially purchased. To the best of our knowledge, this instrument is new to the literature and could be deployed generally in durable goods settings in which a secondary market exists.

Our empirical analysis generates several important results. First, we find that there is dependence between vehicle purchases within a household portfolio. For a two-car household, decreasing the fuel intensity of the kept car in a replacement event induces households to demand higher fuel intensity (or a bundle of attributes that are linked with high fuel intensity) in the purchased car. The effects we estimate are equilibrium effects on the market, explicitly allowing for the entire bundle of vehicle attributes to change. We focus on fuel intensity as the measure of interest due to its high correlation with many other attributes (Knittel, 2011) and its particular policy relevance, as we readily acknowledge that households may be substituting an attribute that is correlated with fuel intensity, such as size or power, and not fuel intensity itself. Notably, we find similar attribute substitution for other policy-relevant attributes, such as the vehicle footprint (i.e., wheelbase x track length), which is policy-relevant under the current footprint-based fuel economy standards in the United States.

Second, an exogenous decrease in the fuel intensity of the kept car alters not only the fuel intensity of the newly purchased vehicle, but also how the two vehicles are used. When the kept vehicle has lower fuel intensity, the cost per mile of driving is lower, so the "rebound effect" implies that the kept vehicle will be driven more miles, some of which may come from the newly purchased vehicle.[4] Using data on household miles driven by both vehicles, we estimate that this rebound effect exhibited by drivers of the kept vehicle erodes roughly 40 percent of the potential fuel savings from an exogenous improvement in kept vehicle fuel intensity. Two additional forces are also at play. An increase in the fuel intensity of the newly-purchased vehicle will imply fewer miles traveled by this vehicle. This is a negative rebound brought about by attribute substitution and is just slightly smaller in magnitude than the effect on the kept vehicle. Furthermore, by changing the relative fuel intensity of the two vehicles within the household, miles will naturally flow away from the now more fuel-intense vehicle to the now less fuel-intense vehicle in a substitution of miles driven. We find, on net, these further reduce the savings from the exogenous decrease in fuel intensity, but the main channel remains through the higher fuel intensity of the newly-purchased vehicle.

Third, we find that changes in gasoline prices interact with household preferences in intuitive ways. As gasoline prices increase, the effect of decreasing the kept vehicle's fuel intensity becomes even stronger. In contrast, the probability of buying a car in the upper quartile of fuel intensity decreases as gasoline prices increase. These results imply that the cost of ignoring the attribute substitution effect when forming policy will scale with the gasoline price. For example, if fuel economy standards are implemented with the intention of reducing gasoline use, the forces of attribute substitution will render the standards less effective in periods of high gasoline prices.

To gauge the importance of attribute substitution for policy, we quantify the net effects of an exogenous decrease in the fuel intensity of the kept vehicle on gasoline consumption, accounting for the above effects.[5] An exogenous reduction in fuel intensity of the kept vehicle leads to an increase in fuel intensity of the next vehicle purchased, and these changes in (counterfactual) fuel intensity lead to subsequent adjustments along the intensive margin. In the two-car sample, decreasing the fuel intensity of the kept vehicle by 3.5 percent (as in Hoekstra, Puller, and West (2017)) results in a 2.1 percent increase in the fuel intensity of the purchased vehicle.[6] After accounting for the direct (kept car) and indirect (bought car) effects on VMT, on net only 32 to 33 percent of the total expected reduction in gasoline consumption from an exogenous reduction in fuel intensity of the kept vehicle is realized for households in this two-car sample, depending on whether we allow

[4]See Borenstein (2015) and Gillingham, Rapson, and Wagner (2016) for more on the rebound effect and its importance.

[5]For our exercise, we are agnostic as to what is causing this decrease in fuel intensity, but require that it is exogenous. This framing mirrors how the federal agencies model CAFE standards. The policy that most closely aligns with our thought experiment is The Car Allowance Rebate System, a 2009 federal U.S. program that is more popularly known as "Cash-for-Clunkers." Cash-for-Clunkers provided incentives to replace old, fuel-intense cars with new, fuel-efficient cars, and lasted for only a few months before disappearing. Alternatively, an exogenous decrease in fuel intensity could arise from tightening fuel economy standards, subsidizing fuel economy, or technological innovation.

[6]We also show that the attribute substitution effect operates strongly through attributes that are correlated with fuel economy, including vehicle footprint and weight. We do not attempt to separately identify these channels.

changes in the other attributes of the vehicle to affect VMT.

For most applications, we are interested in not just the gasoline consumption of two-car households, but in the gasoline consumption over the entire vehicle fleet. We thus extend our analysis to examine the effect of fuel intensity changes of the kept car in several different vehicle portfolio transitions (one-to-two, two-to-three, and a three-to-three replacement), using the same identification strategy.[7] When aggregating across the entire fleet, the direct and indirect effects of attribute substitution erodes 28 percent of the fuel savings from the fuel intensity decrease of the kept vehicle on net after accounting for all of the measurable factors. As a specific example, consider a 3.5 percent decrease in fuel intensity from the average vehicle in our sample. Given the gasoline consumption (540 annually per vehicle in our sample), this 3.5 percent fuel intensity decrease would directly lead to a 18.9 gallon decrease in annual fuel consumption. However, due to attribute substitution the next vehicle that the household purchases will be more fuel-intense than it otherwise would have been. This increase in fuel intensity of the newly-purchased vehicle reduces the fuel savings from our thought experiment to 15.01 gallons, holding usage of the two vehicles constant. But, we also find significant changes in usage patterns that further reduce the net fuel savings. The changes in usage patterns depend on whether we allow the attributes of the two vehicles, other than fuel economy, to endogenously adjust in response to the policy.[8] Alternatively, as is assumed in the government's cost-benefit analysis of Corporate Average Fuel Economy Standards, changes in fuel economy come from changes in fuel efficiency altering technologies, and therefore other attributes are held fixed. If we allow the other attributes adjust in response to the policy, accounting for all of the changes, the net savings of the exogenous decrease in fuel intensity falls from the naive estimate of 18.9 gallons to 13.6 gallons. If we, instead, assume that the other attributes remain the same, the net savings fall from 18.9 gallons to 10.93.

The remainder of the paper proceeds as follows. The next section describes the household vehicle choice problem and outlines a simple theoretical model (Section 2). We then describe our datasets, the restrictions that determine the sample used for our empirical tests, our identification strategy and empirical approach (Section 3). We next present our results and their economic importance (Sections 4 and 5). We conclude with a brief discussion of the implications for policymakers and empiricists (Section 6).

2 Context and Model

We begin by developing a simple economic framework of utility maximization in a setting where consumers enjoy a portfolio comprised of multiple units of a similar good. The model builds upon and shares elements from the frameworks presented in Gentzkow (2007) and Wakamori (2011),

[7] This extension relies on some assumptions about portfolio preferences, which we discuss in Section 5.

[8] In the case of the Cash for Clunkers, West et al. (2017) finds that while Cash for Clunkers increased fuel efficiency, the program also led to a decrease in size and performance. Given all of the changes in the bundle of attributes, they find no evidence that miles traveled increased in response to the increase in fuel efficiency. We find similarly small changes in total VMT when we all other attributes to adjust to the policy-induced increase in fuel efficiency.

but is focused on the policy-relevant question at hand: how a change in the attributes—and in particular the fuel economy—of a household's kept vehicle influences the choice of the next vehicle. Our focus on this question is based on how policies will change the attributes of vehicles, and we are interested in the ramifications this change in attributes will have in multi-vehicle households.[9]

While we will focus on the context of vehicles, our framework can easily be generalized to any product that has portfolio characteristics in which the choice of the good to be added to the portfolio is influenced by changes to the good currently in the portfolio. For example, this framework can also apply to other settings with sequential purchases, such as financial asset portfolios, some household durables (e.g., electronics, art, decor), clothing fashion, media subscriptions, higher education choices, and more.

Consider a consumer who currently possesses one vehicle and is purchasing a second. For simplicity, we model the household as an autonomous decision-maker (i.e., a single consumer). We begin with a random utility model in a standard discrete choice framework. Let the characteristics of a vehicle be given by the vector θ_V, where $V \in \{A, B, ...\}$ denotes distinct bundles of attributes, which we will call "types." For example, these can be defined as the class of vehicle (e.g., SUV or small car), or at a finer level such as at the make-model level.

Suppose the consumer is currently endowed with a product of type A, which is the result of a previous choice made by the consumer that was influenced by factors at that time, such as gasoline prices. The consumer is now deciding which subsequent product to purchase.[10] Because we are interested in the second purchase decision, we focus on the utility from the consumer's vehicle holdings at the time of the second purchase. We assume a consumer's utility function has two parts. The first is the standard component that captures utility directly from the characteristics of the new vehicle being chosen as well as the kept vehicle, as well as their prices. The second component captures the fact that the consumer can receive utility from having a portfolio of products with either similar or different characteristics. For example, suppose the subsequent vehicle purchased is of type B. In this case, we denote additional contribution to utility from having this particular portfolio of vehicles with kept vehicle A and subsequent vehicle B be given by Γ_{AB}.

There are several economic rationales for this 'portfolio' contribution to the utility. For example, the consumer may desire a 'commuter' car with high fuel economy and a larger 'utility' car for other types of trips that require carrying more passengers or materials, and thus there is added value from having a diverse portfolio for the diverse set of trips. Alternatively, there could be heterogeneous preferences within a multi-person household. One spouse may prefer an SUV, while the other prefers a Prius, so the combined household utility is maximized with a diverse portfolio. Or both spouses may prefer the same vehicle type, so that utility is maximized with a more similar portfolio. Another related possibility is that other goods the household owns (such as a boat) could

[9]Of course, standards also can change relative prices of vehicles of different fuel economies, but the focus here is on changes in attributes.

[10]We assume that the probability of not purchasing the subsequent product is not altered with a change in an attribute of product A. This allows us to simplify the model by ignoring the outside option.

be complements to different vehicle types, which could lead to either a preference for similar vehicles (so that all drivers in the household could tow the boat) or a preference for different vehicles (to have one more efficient vehicle and one towing vehicle). The key point is that Γ_{AB} could be positive or negative, depending on the preferences of the household and the characteristics of the vehicles.

The household's indirect utility derived from such a portfolio is given as:

$$u_{AB}(\theta_A, \theta_B, \Gamma_{AB}) = f(\theta_A) + f(\theta_B) + \Gamma_{AB} - \alpha(p_A + p_B) + \epsilon_{AB}, \tag{2.1}$$

where $f(.)$ is a function that maps product attributes into consumer utility, and p_V is the price of product V.

2.1 Implications for Product Choice

We seek to understand how an exogenous change in the attributes of the already-owned (kept) vehicle influences the choice of the attributes in the second vehicle. Thus, we extend the above framework by assuming that the consumer may choose between types B and C for the subsequent vehicle. The consumer chooses portfolio AB rather than AC if $u_{AB} > u_{AC}$. Thus, AB is chosen if:

$$f(\theta_B) - f(\theta_C) + \Gamma_{AB} - \Gamma_{AC} - \alpha(p_C - p_B) + \epsilon_{AB} - \epsilon_{AC} > 0. \tag{2.2}$$

This inequality indicates that the consumer will choose B as the second vehicle when the net utility from B dominates the net utility from C.

A set of policy-relevant comparative statics emerge from this framework. Consider what the model implies for the equilibrium portfolio choice probabilities. Conditional on purchasing a second vehicle, for the simple choice between B and C the choice probabilities are given as follows:

$$\begin{aligned} Pr_{AB} &= \int_{\mathbf{u}} I(u_{AB} > u_{AC}) dG(\mathbf{u}), \\ Pr_{AC} &= \int_{\mathbf{u}} I(u_{AC} > u_{AB}) dG(\mathbf{u}). \end{aligned} \tag{2.3}$$

Here $I(\cdot)$ is an indicator and $G(\cdot)$ is the distribution of utilities in the population.

We are interested in how the choice between portfolios AB and AC when θ_A changes, and in particular when the fuel economy attribute within the vector θ_A changes. Let the element of the vector θ_A that corresponds to any attribute l be denoted as θ_A^l. For the fuel economy attribute specifically, we denote the element of the vector as θ_A^m. This notation allows us to extend our framework further, by noting that vehicles B and C each also have a fuel economy attribute: θ_B^m and θ_C^m respectively. From a policy perspective, we are especially interested in how changes in θ_A, and especially θ_A^m, influence Pr_{AB} and Pr_{AC} based on the relative values of θ_B^m and θ_C^m.

There is a clear economic intuition for why a change in the fuel economy attribute θ_A^m may influence Pr_{AB} and Pr_{AC}. Consider that the derivative of (2.2) with respect to θ_A is just $\frac{\partial \Gamma_{AB}}{\partial \theta_A^m} - \frac{\partial \Gamma_{AC}}{\partial \theta_A^m}$ (this follows because only the portfolio terms are a function of θ_A^m). The sign of this derivative could be positive or negative depending on consumer preferences. For example, if the consumer has a higher fuel economy kept vehicle, she may desire more cargo space or acceleration (and thus lower fuel economy) in the subsequent purchase. By the same token, it is possible that if the kept vehicle has higher fuel economy, the consumer will demand more fuel economy in the subsequent purchase because of within-household bargaining (i.e., if one spouse gets to drive a higher fuel economy vehicle, the other may wish to do the same).

2.2 Attribute Substitutes and Complements

There is an obvious closeness in terminology between the substitutability and complementarity of goods and the attribute substitution/complementarity we introduce in this paper. However, these concepts are distinct and not nested. Here, we take a moment to formally define attribute substitution and attribute complementarity given the model laid out above. To generalize, consider any attribute l in attribute vector θ. Let B be the newly purchased good. As before, let A be the kept good and thus θ_B^l is the element in θ_B associated with attribute l of the bought good. We propose the following definition:

Definition 1. A consumer exhibits a preference for *attribute substitution* in a when $\frac{\partial \theta_B^l}{\partial \theta_A^l} < 0$ and *attribute complementarity* when $\frac{\partial \theta_B^l}{\partial \theta_A^l} > 0$.

Attribute substitution thus refers to an equilibrium change in the bought product's attribute when the endowed product's attribute changes. In the context of a two-car vehicle portfolio, this equilibrium outcome occurs due to a change in the willingness to pay for the attribute in the bought car from a change in the endowment of the attribute in the kept car. We focus on this equilibrium object because it is the most policy-relevant in our setting.

A natural question that arises from this definition is how it relates to classic definitions of complements and substitutes. As is described in Samuelson (1974) and Newman (1994), there are many classic approaches to defining complements and substitutes. The basic intuition for all of these is that consumers change their equilibrium willingness to pay for a product when there is a change in the price of the other product or the quantity of the other product one is already endowed with. One common way this is operationalized is by defining two products as complements if the cross-price derivative of demand (usually Hicksian, but Marshallian under some definitions) is negative (substitutes if positive). Another common way is by defining two products as complements if the cross second-order partial derivative of the utility function with respect to quantities of both goods is positive (substitutes if negative).[11]

[11] There are several further definitions in the literature as well, including the Samuelson (1974) proposed approach of using von Neumann utility.

Our definition of attribute substitution is closest to the first definition of complements and substitutes. Under this definition, the focus is on changes in demand for a product in the context of relative changes in one particular attribute: price. In contrast, attribute substitution as defined in Equation 1 could occur in any attribute, including price. Moreover, attribute substitution focuses on the change in the product of interest's attribute, rather than demand for the entire bundle of attributes that make up a product. It is an equilibrium object, much in the way that whether two goods are complements or substitutes is an equilibrium object. But it should be apparent that while attribute substitution or complementarity may correlate with goods substitutability or complementarity, the concepts are non-nested. To see this, consider a circumstance where, on average, consumers exhibit a preference for attribute substitution over a particular non-price attribute. It is possible, indeed quite easy, to conceive of joint distributions of individual preferences over the remaining attributes (including price) that could support any number of patterns of goods complementarity or substitutability across individual vehicle pairs within the market. Just as the literature notes that complementarity between two products depends on the relation of those products to others in the market (Samuelson, 1974; Newman, 1994), our notion of attribute substitution also depends on the relation of the two products with others in the market.

3 Empirical Approach & Data

3.1 Baseline model

We use a panel dataset of household vehicle purchase decisions and portfolio holdings to quantify the effect that exogenous changes in attributes of a kept vehicle have on the choice of the next vehicle, and thus the overall composition of the portfolio. Thus, we use i to denote a household and t to denote the year of the subsequent vehicle purchase. Our focus is on fuel economy, and the units for fuel economy we use are the average gallons per mile (GPM) required for propulsion, which is also the fuel intensity.[12] Our empirical approach is directly motivated by the theory above.

To transition from our theory to the empirics, note that under standard monotonicity and continuity assumptions, the demand for the next purchased vehicle's fuel economy is an implicit function of all of the attributes of the kept vehicle, including the kept vehicle's fuel economy. Specifically, using the notation from above with B being the bought vehicle, under standard inversion assumptions (where we invert to recover a single parameter) we can retrieve the attribute demand function for fuel economy from the utility of a given portfolio is $u_{it,AB}(\theta_{it,A}, \theta_{it,B}, \Gamma_{it,AB})$:

$$\theta_{it,B}^{m} = u_{it,AB}^{-1}(\theta_{it,A}, \theta_{it,B}^{-m}, \Gamma_{it,AB}), \tag{3.1}$$

where $\theta_{it,B}^{-m}$ is defined as the vector of attributes of vehicle B not including fuel economy and $u_{it,AB}^{-1}$ is an inverse utility function (possibly a correspondence), which characterizes the attribute demand

[12] We prefer this definition to miles per gallon because total fuel consumption linearly scales with GPM and thus is the better measure (Larrick and Soll, 2008).

for fuel economy.

This insight directly leads to our empirical specification. We are interested in how consumers as a whole exhibit attribute substitution or complementarity in fuel economy, so we are interested in how changes to the fuel economy of the kept vehicle, θ_A^m, influence the choice of θ_B^m. Of course, other factors may influence the GPM of the newly bought vehicle, so in our empirical specification we also must include the gasoline price and a set of further covariates.

In our empirical specification, we linearize (3.1) and slightly modify the notation to apply the theory to the empirical setting. We model household i's bought vehicle GPM at time t, f_{it}^b, as the following function of the kept vehicle GPM, f_{it}^k:

$$f_{it}^b = \beta_0 + \beta_f f_{it}^k + \beta_g p_{it}^{gas} + \beta_{gf} p_{it}^{gas} \times f_{it}^k + \alpha_X X_{it}^k + \varepsilon_{it}, \tag{3.2}$$

where i's contemporaneous gas price in t is p_{it}^{gas}. X_{it} is a vector of kept vehicle attributes. These control for the asset value (via the estimated resale value), brand preferences (via kept vehicle manufacturer and class fixed effects), and replacement cycle effects (via kept vehicle age). Further, we account for the vehicle purchase type (indicator for new or used), seasonality and time-varying macroeconomic shocks using nonparametric time controls (year and month-of-year fixed effects), time-invariant household heterogeniety using household fixed effects, and other confounders that vary in both space and time, such as regional economic shocks, using demographic controls (monthly county-level unemployment).[13]

To investigate heterogeneity of effects across more or less fuel-intense vehicles, we also estimate a linear probability model where the dependent variable, $Pr\left(q(f_{it}^b) = s\right)$, equals one if f_{it}^b falls within the range of quartile $s \in \{1, med, 4\}$, keeping the rest of the specification as presented in (3.2). Following the theory, our coefficients of interest are β_f and β_{gf}. β_f tells us how changes in in the kept vehicle fuel economy affect the choice of the bought vehicle fuel economy (in terms of the theory, it is quantifying $\frac{\partial \theta_B^m}{\partial \theta_A^m}$). If β_f is negative, we have attribute substitution, while if it is positive, we have attribute complementarity. β_{gf} tells us how the magnitude of attribute substitution or complementarity may change with the current gasoline price.

3.1.1 Identification of the baseline model

The identification challenges in our setting can be most easily understood by considering the ideal experiment for answering our research question. Consider every two-car household that is about to exchange a vehicle for a different one. The ideal experiment would randomly assign one car to be the "kept" car, exogenously perturb the GPM of this kept car (f^k), and see how this exogenous change affects the household's observed choice of f^b, the GPM of the car purchased.[14]

[13] Macroeconomic shocks, seasonality, and changes in vehicle tastes over time should be absorbed in time fixed effects. Likewise, macroeconomic shocks at the time of the kept vehicle purchase should be absorbed by the aforementioned kept vehicle age fixed effects.

[14] Note that this basic approach of examining the effect of an exogenous change in GPM is exactly how the agencies in the United States perform their analysis of fuel economy standards.

The first identification challenge is that of unobserved preferences for fuel economy at the household level. To see this, consider a case in which we observe a cross-section of household vehicle pairs, where one is a newly-purchased second vehicle. If we estimated the model in (3.2) using only "between" variation, we may see that the purchased vehicle's GPM is increasing in the kept vehicle's GPM simply because households that have a fuel efficient kept vehicle may prefer more efficient vehicles in general. Such a preference for fuel efficiency for all vehicles in the portfolio is not the same as attribute substitution. To understand attribute substitution, we ideally would like to exogenously change the GPM of the kept vehicle A and observe its effect on the probability of choosing vehicle B versus C, but in the cross-section we cannot separate preferences over the level of attributes from the preferences for the mix of attributes.[15] Panel data help us overcome this challenge by observing repeated replacement choices by the same household. For example, a household repeatedly responding to an exogenous improvement in the fuel economy of the kept car by choosing a less fuel efficient second car allows us to separately identify preferences over level and gradient. Thus, household fixed effects are crucial for our identification.

However, even when utilizing repeated choices, there may be time-varying unobserved household characteristics that affect the vehicle choice decision (e.g., adding a household member or changing jobs). Household decision-making changing over time could bias an attempt to estimate the preference relationship between portfolio attributes because even with repeated choices the researcher would observe the jointly determined preference for attribute level and mix. We address this concern by relying on exogenous variation that perturbs the GPM of the kept vehicle. This instrumental variables approach may also address a wide variety of other potential confounders such as unobserved car attributes.

In our baseline specification, we instrument for the GPM in the kept vehicle (f^k) using the price of gasoline at the time of the kept vehicle purchase, $p^{gas}_{it_k}$, where t_k is the time of the kept vehicle purchase. This instrument is motivated by our theory. Recall that the attributes of the endowed vehicle A are the result of a previous choice made by the consumer, and are influenced by factors occurring at the time of the previous choice. Both theory and evidence (e.g., Busse, Knittel, and Zettelmeyer (2013), Klier and Linn (2010)) demonstrate that households consider future operating costs of the vehicle in their purchase decision. Changes in California gasoline prices several years prior are exogenous with respect to today's household choice, vary extensively over the time period of our data, and alter the expected lifecycle cost of vehicles according to each vehicle's GPM. Based on this logic, when gasoline prices are high at the time of the kept vehicle purchase, we would expect the household to purchase a more fuel efficient car than when gasoline prices are low (as also demonstrated in Busse, Knittel, and Zettelmeyer (2013)). The price of gasoline at the time of the kept car purchase thus provides exogenous variation in the potentially endogenous variable of interest, the GPM of the kept vehicle f^k.

[15] A cross-sectional analysis is analogous to taking the derivative of (2.2) with respect to a composite variable of θ_A and an unobserved preference variable.

Figure 1 illustrates the underlying relationship between gasoline price instrument ($p^{gas}_{it_k}$) and the endogenous kept vehicle fuel intensity (f^k). Many factors influence the consumer's purchase decision at the time of a (in this case the kept) vehicle purchase, so a plot of the raw data would reveal little detail of this relationship or first-stage power. Instead, we present this relationship after conditioning out household fixed effects and other covariates used in our primary regression specifications.

In these figures gray circles represent the mean value for 0.005 GPM bands of fuel intensity. The blue line is a kernel regression demonstrating the nonparametric relationship and the orange line is a linear regression with 95% confidence intervals shaded. There is a clear, downward sloping trend which agrees with the prior from economic theory. As gasoline prices increase, consumers, conditional on idiosyncratic household preferences, tend to purchase less fuel-intense vehicles.[16]

The $p^{gas}_{it_k}$ instrument can also be interacted with the current gasoline price to instrument for $p^{gas}_{it} \times f^k_{it}$. Thus, in our primary specification, we have a straightforward instrumental variables strategy to address endogeneity concerns. As such, while Figure 1 demonstrates the relationship between one instrument and one engougenous variable agrees with our prior, the true first stage – consisting of two instruments and two endogenous variables – is more complex.

3.2 Heterogeneous Effects

Our baseline specification assumes that the effect of a change in GPM of the kept vehicle on the fuel economy of the purchase vehicle is the same regardless of whether the household chose to keep its more efficient or less efficient vehicle. This is not required by our theory as consumer preferences can allow for considerable heterogeneity in the effects. To accommodate this possibility, we provide additional flexibility in our model of attribute substitution. This flexibility requires an additional instrument that provides exogenous variation in which vehicle the household keep. We discuss this instrument below.

In this specification, the dependent variable is still the GPM of the bought car itself (f^b_{it}), however, we expand on the specification in (3.2), by including an indicator variable for whether the high or low fuel intensity vehicle is kept from the original portfolio then interacting this indicator with GPM of kept car. For notational simplicity, denote the chosen vehicle by the following indicators:

$$\begin{aligned} \mathbb{1}^{k>d} &\equiv \mathbb{1}\{f^k > f^d\} \\ \mathbb{1}^{d\geq k} &\equiv \mathbb{1}\{f^d \geq f^k\} = \left(1 - \mathbb{1}^{k>d}\right) \end{aligned} \tag{3.3}$$

[16] One may also be interested in the reduced form relationship between $p^{gas}_{it_k}$ and f_b. We present graphical evidence of this relationship in C.1 of the Appendix.

This leads to the following specification:

$$
\begin{aligned}
f_{it}^{b} = \quad & \beta_0 + \beta_g p_{it}^{gas} + \mathbb{1}^{k>d} + \mathbb{1}_{it}^{k>d} \times \left(\beta_{fk} \times f_{it}^{k} + \beta_{gfk} p_{it}^{gas} \times f_{it}^{k}\right) + \\
& \mathbb{1}_{it}^{d \geq k} \times \left(\beta_{fd} \times f_{it}^{k} + \beta_{gfd} \times p_{it}^{gas} \times f_{it}^{k}\right) + \alpha_X X_{it}^{k} + \varepsilon_{it}
\end{aligned} \tag{3.4}
$$

with variables defined as in (3.2) and (3.3).

3.2.1 Identification of heterogeneous effects

This specification presents an additional challenge to identification. Just as the GPM of the kept vehicle (f_{it}^{k}) is likely endogenous, the relative position of the kept vehicle's GPM compared to other vehicles in the portfolio is likely also endogenous. This leads to an estimating equation with five endogenous variables: an indicator for observations where households replace the relatively efficient vehicle in the portfolio ($\mathbb{1}^{k>d}$), this indicator interacted with the endogenous kept vehicle GPM variables (f^k and $p^{gas} \times f^k$), and corresponding terms interacted with an indicator for when households replace the relatively more fuel-intense vehicle in the portfolio and the following matrix of endogenous variables:

$$
\mathbf{W}_{it} = \begin{bmatrix} \mathbb{1}_{it}^{k>d} & \mathbb{1}_{it}^{k>d} \times f_{it}^{k} & \mathbb{1}_{it}^{k>d} \times p_{it}^{gas} \times f_{it}^{k} & \mathbb{1}_{it}^{d \geq k} \times f_{it}^{k} & \mathbb{1}_{it}^{d \geq k} \times p_{it}^{gas} \times f_{it}^{k} \end{bmatrix}'.
$$

We address this challenge to identification in two steps. First, we include the gasoline price at the time the dropped vehicle was purchased ($p_{it_d}^{gas}$) as an instrument, just as before. This gasoline price is similarly correlated to the GPM of the dropped vehicle but exogenous with respect to the household's current choice. Second, we deploy a new instrument relying on unanticipated changes in the relative prices of each vehicle in the household's portfolio.

This instrumental variables strategy is new to the literature and can be generalized to portfolio durable goods whenever a secondary market exists (for example commercial aircraft as in Gavazza (2011) or textbooks as in Chevalier and Goolsbee (2009)). In our context, we would like to instrument for the choice of the kept vehicle and the GPM of the kept vehicle. A valid instrument will provide exogenous variation in the process that determines which of the household vehicles is kept and which is replaced. The exclusion restriction requires that the instrument affects the household's choice of f^b only indirectly, through the choice of which car to keep.

We instrument for a household's decision of which vehicle to drop using changes in the relative values of the vehicles in their portfolio. For exposition, let P_t^k and P_t^d be the average retail value of the kept and dropped cars, respectively, at the time when the car is dropped (t). Further, allow the change in price differences between time t and time 0 (when the kept car was purchased) to be $\Delta\Delta P^{kd} = (P_t^k - P_t^d) - (P_0^k - P_0^d)$. This instrument extracts the portion of the variation in the price difference-in-difference that occurs after the time of purchase (i.e., deviations from expectations about the trend of relative prices).

When constructing this instrument, we assume that households form expectations using lagged 5-year depreciation rates at the make-model-model year level, and project these into the future. Deviations from these projections are what we refer to as the "deviation from trend instrument," (or DfT) and we use differences in these deviations as an instrument. We find it difficult to come up with a violation of the exclusion restriction in this case. Any concern must posit a direct relationship between the instrument and GPM of the bought car that works outside of the relationship between the instrument and kept car attributes.

For both the kept and dropped vehicle we proxy the household's expectation of annual vehicle depreciation using an estimate of the depreciation of similar vehicles over the previous five years. For vehicle of make m, model year y, and value $V_{m,y,t}$ in year t, the expected depreciation is:[17]

$$\mathbf{E}[Dep_{m,y,t}] = \left(\prod_{s=1}^{5} \frac{V_{m,y-s+1,t-s+1} - V_{m,y-s,t-s}}{V_{m,y-s,t-s}} \right)^{\frac{1}{5}}. \tag{3.5}$$

We can then calculate the deviation from this expected depreciation rate for each car in the portfolio, and construct the DfT instrument. Assuming vehicle j has resale value $P_{j,t}$ in year t, this is:

$$\Delta\Delta V_{it}^{kd} = (P_{it}^{k} - \mathbf{E}[Dep_{it}^{k}] \cdot P_{i,t-1}^{k}) - (P_{it}^{d} - \mathbf{E}[Dep_{it}^{d}] \cdot P_{i,t-1}^{d}). \tag{3.6}$$

Figure 2 shows the first stage relationship between deviation from trend instrument and the probability the dropped vehicle was is the least valuable in the portfolio.[18] The first stage is clearly strong and possibly follows a nonlinear form. To address this potential nonlinearity, we will consider both a cubic function of $\Delta\Delta V_{it}^{kd}$, and a linear spline as candidate instruments.[19] Combined with the gasoline prices at the time of each vehicle purchase, we will refer to this set of instruments as V with the full set of first-stage equations:

$$\mathbf{W}_{it}^{w} = \mathbf{\Gamma}_0 + \mathbf{\Gamma}_V V_{it}^{wkd} + \mathbf{\Theta}\mathbf{X}_{it} + \mathbf{\Xi}_{it}^{w}. \tag{3.7}$$

The theoretical basis for this instrument may be most easily conveyed by recalling the model of adverse selection in Hendel and Lizzeri (1999). Consumers have heterogeneous preferences over attributes of new and used cars, and there is asymmetric information about the quality of goods transacted in the used market. In such a setting, an unexpected change in a state variable (e.g., an increase in the price of gasoline or the introduction of a new product) will catalyze a re-sorting of goods to consumers. Our empirical setting shares these characteristics. Information that enters

[17] As a more concrete example, for a household in year $t = 2005$ owning a 2002 Honda Civic, the expected depreciation is the geometric mean annual depreciation rate of 2001 Hondas in 2004, 2000 Hondas in 2003, 1999 Hondas in 2002, etc.

[18] The figure is limited to a range of $\Delta\Delta V_{it}^{kd} \in [-2500, 2500]$, which is 99.5% comprises 99.5% of the data.

[19] We select a cubic functional form as the simplest form that accounts for the nonlinearity while being an odd-order polynomial.

the market after a household forms its two-car portfolio may affect the value of each car, both to the present owners as well as (heterogeneously) to other market participants. Such information can include factors such as recalls of cars and the introduction of new models that put older models in a different light than expected. These factors generate exogenous variation in unexpected changes in relative prices that is correlated with the choice of which car to replace.

3.2.2 Estimation of the instrumental variables system for heterogeneous effects

While statistical power in the first stage is not a concern when estimating equation (3.2), we find that we have low first stage power when estimating Equation 3.7 using instruments V. The endogenous regressors are a system of interactions with both an exogenous regressor (the gasoline price at the time of bought vehicle purchase) and an endogenous binary indicator. As one might expect, this system is difficult to approximate using purely linear projections of the instruments in V.

To more closely approximate the hypothesized relationship between the endogenous variables and instruments, one may consider forming additional instruments that follow the functional form of these non-linear relationships by interacting V with the exogenous current gas price or using pairwise interactions from the row-wise Kronecker product of instruments, $V \otimes V$. This however, can lead quickly to a proliferation of instruments with the potential to greatly exacerbate any IV finite sample bias.[20] As an alternative, we follow Wooldridge (2010) and form a narrow set of instruments, approximating the functional form of the endogenous variables using interactions of projections from the space of exogenous variables.[21] This method is commonly deployed (e.g., Aizer and Currie (2017), Dorsch, Dunz, and Maarek (2015), Alcaraz, Chiquiar, and Salcedo (2012), and Michaels (2008)) when explanatory variables include interactions involving endogenous variables.[22]

We estimate the first-stage relationships for the uninteracted endogenous variables $\mathbb{1}^{k>d}$ and f_{it}^k and projections from the space of exogenous variables ($\widehat{\mathbb{1}^{k>d}}$ and $\widehat{f_{it}^k}$). We then compute four new instruments as interactions of these predictions:

$$\begin{aligned}
\widehat{\mathbb{1}_{it}^{k>d} \times f_{it}^k} &= \widehat{\mathbb{1}_{it}^{k>d}} \times \widehat{f_{it}^k} & \qquad \mathbb{1}_{it}^{k>d} \widehat{\times f_{it}^k} \times p_{it}^{gas} &= \widehat{\mathbb{1}_{it}^{k>d}} \times \widehat{f_{it}^k} \times p_{it}^{gas} \\
\widehat{\mathbb{1}_{it}^{d\geq k} \times f_{it}^k} &= (1 - \widehat{\mathbb{1}_{it}^{k>d}}) \times \widehat{f_{it}^k} & \qquad \mathbb{1}_{it}^{d\geq k} \widehat{\times f_{it}^k} \times p_{it}^{gas} &= (1 - \widehat{\mathbb{1}_{it}^{k>d}}) \times \widehat{f_{it}^k} \times p_{it}^{gas}.
\end{aligned} \tag{3.8}$$

We augment the vector of instruments (V_{it}^{wkd}) with these four additional instruments and estimate

[20] This specification utilizes 5 instruments. Simply forming all pairwise interactions and the interactions with gasoline prices would lead to 50 instruments.

[21] Alternatively, we construct an arbitrarily large space of candidate instruments and select a subset for the first stage using the IV-lasso procedure of Belloni et al. (2012). Results are shown in Section D.3 of the Appendix. This process generally leads to a weaker first stage, but broadly similar point estimates.

[22] Wooldridge (2010) Section 9.5.2 describes this method in detail and demonstrates that if the first stage instruments satisfy the exclusion restriction, the projections will as well. Aizer and Currie (2017); Dorsch, Dunz, and Maarek (2015); Michaels (2008) each instrument for an endogenous interaction term in a linear using the interaction of a covariate and projections from the first-stage regression. Alcaraz, Chiquiar, and Salcedo (2012) use the same procedure in the context of an instrumental variables probit.

the full system using GMM.

3.3 Data

The cornerstone of our dataset is the universe of California vehicle registration records that occurred from 2001-2007.[23] The DMV dataset includes every vehicle registered under the residential designation code (e.g., not commercial or government). In California every vehicle must be registered annually. Each record includes the 17-digit vehicle identification number (VIN) that uniquely identifies the vehicle, that year's registration date, the date when the vehicle was last sold, and various other information. A confidential version of the data includes registrant surnames and premise address. This information allows us to construct a household-level panel dataset of vehicle ownership in partnership with the California Air Resources Board.

Basic vehicle attributes (e.g., horsepower, weight, etc.) are available via a VIN decoder that we purchased from DataOne Software. We augment the decoder to include vehicle fuel intensity, which is available from the US Environmental Protection Agency. Vehicle-miles traveled (VMT) are available for each VIN whenever the vehicle is sold and upon receiving biannual Smog Check certification.[24] We thus have an average measure of miles traveled by each vehicle and, by extension, each household for each year in our sample. Our gasoline price data are from the Oil Price Information Service (OPIS) at the county-month level.

In each year households are characterized by the starting and ending number of vehicles in their portfolio. In year t a household's starting portfolio size (N^s) is the number of vehicles registered in that year. If the household also registers exactly N^s vehicles in year $t+1$ or $t+2$, then the ending portfolio size (N^e) in year t is N^s. If the number of vehicles registered in years $t+1$ and $t+2$ are identical to each other, but not equal to N^s then the ending portfolio size is the number of vehicles registered in the later years.[25]

Table 1 shows the distribution of household portfolio transitions. Rows indicate the number of cars in year t, and columns indicate the number of cars in $t+1$. The table represents all possible household transitions. The large mass on diagonals indicates that many households do not increase or decrease the number of cars that they register from year to year. A careful interpretation of "0" is necessary: households owning 0 cars are unobserved in our dataset, so transitions from 0 occur when a Californian household without a car in t registers one in $t+1$, or with observational-equivalence, a household moves to California from another state. Similarly, transitions to 0 occur

[23] We thank the California Department of Motor Vehicles (DMV) for making these data available for research.

[24] New vehicles are not required to have a smog check until six years after registration, subsequently it is a biennial check. Hybrid electric vehicles and electric vehicles are also exempt.

[25] We examine one and two years in the future as a household that may register more cars in one year than they ever owned simultaneously. For example, consider a household that owns two cars in year t. In year $t+1$ they re-register both previously owned vehicles and the registrations expire. Then, toward the end of the year, they sell one vehicle and replace it with a new one (which requires registration of the new vehicle). This household has registered three unique vehicles in year $t+1$ but only ever owned two at any given time. In year $t+2$, barring the purchase of yet another new vehicle, the household would return to registering two vehicles.

either when a household sells all of its registered cars, if it exits the data via a move to another state or a dissolution of the household.

The key regressions that follow are estimated using a sample of two-car households that replace one of their cars, a sample which we call "2x2 replacement households."[26] While other transitions are certainly interesting, and we examine them later, two-car replacement households provide the cleanest experiment. Households increasing the number of cars in their portfolio are likely to be experiencing an unobserved development that increases their demand for transportation (e.g., having a baby). Furthermore, it is unclear how to characterize the channels through which the consumers may have preferences for attribute substitution when there are multiple kept cars. Do these households substitute attributes based on the highest-VMT kept car, or the newest? Or is a higher dimensional analysis required?

Given that no clear answer exists to these questions, for our primary analysis we choose the transparent path of focusing on the replacement decisions of two-car households, consistent with the simple theory model presented above (we also perform additional analyses examining other transitions). Moving forward with our 2x2 replacement sample is valid when considering small deviations from an interior consumer choice optimum. For such deviations, the probability of a different transition (e.g., not purchasing a second vehicle or purchasing a third vehicle) is not affected by a small change in an attribute of the kept vehicle.

Table 2 shows summary statistics for all 2x2 replacement households, including segmentation based on the fuel economy of the bought car. Households that purchase relatively fuel efficient vehicles (gallons per mile quartile 1) tend to keep relatively fuel efficient cars as well. The converse is true for households buying fuel inefficient vehicles, suggesting that households may have an overall preference for either high or low fuel economy cars.

Some of the analyses that follow use the quartile of fuel intensity to describe bought and sold cars. The GPM cutoffs are presented in Table 3, along with their corresponding fuel economy analogs in miles-per-gallon (MPG) for reference. Figures 3a - 3b present histograms of the number of transactions per household under various sample restrictions. It reveals that we are left with approximately 235,000 households for our primary instrumental variables specification that includes household fixed effects.

[26]We define a household as replacing one vehicle if the starting (in year t) and ending (in year $t+1$ or $t+2$) portfolios differ by one vehicle. The household may conduct multiple vehicle transactions, as long as one of the two vehicles appears in both the starting and ending portfolios. We do not consider households where both vehicles in the two-vehicle portfolio change as the relative timing of each purchase becomes important for defining the portfolio at the time of each vehicle's purchase.

4 Results

This section presents our primary estimation results.[27] We first demonstrate the importance of the instrumental variables approach and inclusion of household fixed effects, both of which qualitatively and quantitatively alter key coefficient estimates. We then present the marginal effects of kept car GPM on bought car GPM, which reveal household preferences for attribute substitution. We also investigate heterogeneity in these marginal effects by the relative level of the retained vehicle's GPM in the portfolio. Motivated by the correlation between GPM and other vehicle attributes, we also examine the relationship between kept car GPM and footprint, engine displacement, and weight of the bought car. These results provide context for the discussion of policy implications that follows.

4.1 Effect of Kept GPM on Bought GPM

Table 4 presents the baseline regression results from the full sample, and separately for new and used car purchases. This first set of results is based on estimating the model in equation (3.2). Our object of interest, the effect of kept vehicle fuel intensity on bought vehicle fuel intensity, is a function of the coefficients on the kept vehicle fuel intensity (GPM^K) and its interaction with the gasoline price ($GPM^K \times P^G$). Recall that beyond these parameters of interest our primary regression specification includes linear controls for the current gasoline price, the bluebook resale value of the kept vehicle, and the current month's county unemployment rate plus fixed effects for each household, whether the purchased vehicle was new or used, the vehicle class and manufacturer of the kept vehicle, the age of the kept vehicle, year and month of year. For ease of interpretation and to focus on potential attribute substitution, in Table 5 we additionally compute the marginal effect of kept vehicle GPM on the bought vehicle GPM for each specification at gasoline prices of \$2, \$3, and \$4 per gallon.

In specifications failing to account for unobserved heterogeneity across households or the endogeneity of kept vehicle fuel intensity (column 1), the results suggest that households will tend to replace their dropped car with one that is qualitatively similar in GPM to the kept car. For example, the coefficient in the first row of Table 5 Column 1 shows that an increase in the fuel intensity of the kept vehicle is correlated with an increase in the fuel intensity of the bought vehicle. This result suggests that when looking across households, we see higher fuel intensity in the kept car (regardless of which one is dropped) being correlated with higher fuel intensity in the bought car. This is consistent with different households having different preferences for fuel sippers or gas guzzlers in general (presumably due to correlated attributes including power, comfort, safety, etc).

Columns 2 and 3 take different approaches toward accounting for challenges to identifying the causal effect of kept vehicle attributes on the attributes of the follow-on purchase. Column 2

[27] Estimates presented in this section rely in part on user-contributed software in Baum, Schaffer, and Stillman (2002) and Correia (2014).

instruments for the fuel intensity of the kept vehicle, but ignores time-invariant household-level preferences for the level of attributes in the portfolio. Column 3 takes an alternate approach, controlling for unobserved household preferences of the level of portfolio attributes using household fixed effects. In each case, the estimated effects have the opposite sign from the OLS estimates, again consistent with households having heterogeneous preferences for the level of fuel economy in their portfolio.

Finally, our preferred results in column 4 combines these approaches. These results show an increase in fuel intensity of the kept car *decreases* the fuel intensity of the bought car. Negative coefficients reflect a household's revealed desire to buy a car with more (less) GPM as the GPM of its kept car decreases (increases). All of the marginal effects have a negative sign, and are statistically significant in the full and used vehicle sample.[28] Following our theory, we interpret this result as causal evidence of attribute substitution for the GPM attribute.

It is clear from this progression of specifications that it is important to account for both endogeneity and unobserved household heterogeneity. Specifications without household fixed effects primarily rely on variation across households and do not reflect the thought experiment described earlier, which relies on within-household variation. In many cases, the inclusion of household fixed effects flips the sign of the estimated coefficient, indicating that within versus across variation may lead to important differences in interpretation. Deploying instruments has the overall effect of magnifying coefficient estimates. This not surprising in a setting in which many factors enter into the household vehicle purchase decision, including those outside of our channels of interest.

The extent of bias associated with OLS underscores the importance of valid instruments. For each regression using instrumental variables, In Table 4 we present the Kleibergen-Paap rk F statistic as an assessment of weak identification. The statistics associated with each of our baseline regressions in Table 4 offer reassurance that the instruments are indeed strong, with the possible exception of the new vehicle sample. In Section D.1 of the Appendix, we apply further common diagnostics of first stage power and reject the null of weak instruments whenever a formal statistical test is available.[29]

Further, comparing the marginal effects within each specification allows us to investigate how gasoline prices impact attribute substitution. As the gasoline price increases, the magnitude of the attribute substitution effect likewise increases. This may be due to the relative importance of fuel intensity (or attributes correlated with fuel intensity) in household decisions when gasoline prices are higher.

[28]Identification of household fixed effects is driven only by households with multiple transactions during our sample period. To investigate potential sample selection induced by relying on multiple transaction households, in Section D.4 of the Appendix, we repeat the OLS and IV estimates using only transactions which contribute to identification in the HHEFIV specification. The parameter estimates from these restricted regressions are nearly identical to the unrestricted regressions reported here.

[29]Further, in Section 3.2 we identify additional candidate for a specification with additional endogenous variables. In Section D.2 we utilize these additional instruments in tests of overidentifying restrictions. These additional instruments have power and in each case we fail to reject the null that our instruments satisfy the exclusion restriction.

The overall story is clear: households incorporate portfolio considerations in their vehicle purchase decisions and have a preference for attribute substitution in GPM. That is, if we were to exogenously increase the fuel intensity of the kept car, households would buy a second car that has attributes associated with lower fuel intensity. This implies that there is an explicit dependency between the kept car and bought car – a dependency that is rarely discussed in previous work.

4.1.1 Robustness Tests

One may be concerned that household income or wealth effects stemming from the vehicle purchase decision or changes in gasoline prices may influence parameter estimates. In Appendix Section D we present alternative estimates to address two of these concerns. First, it is possible high gasoline prices at the time of vehicle purchase may induce a household to purchase a less expensive, but more fuel efficient vehicle. If gasoline prices fall in the future, the household's follow-on vehicle purchase may be a more expensive and less fuel efficient vehicle in response to the relaxed budget constraint. This effect would be most salient when the vehicle was recently purchased, so we reestimate our primary specification excluding observations where the dropped vehicle is less than three years old at the time of replacement and find similar parameter estimates.

Second, it is possible the sale or scrap value of the dropped vehicle may have an income effect on the choice of attributes for the follow-on purchase. As a test of this income channel, we include the price paid by the household for their dropped vehicle at the time of purchase as a covariate. While this is a potentially endogenous control variable, one would expect its inclusion to alter our parameter estimates if this income channel is biasing our results. Again, results of this regression are similar to our primary specification.[30]

4.2 Heterogeneous effects

Attribute substitution may arise because households see each of their vehicles as serving different household needs. For example, a household may value having a small, efficient vehicle for commuting and a larger, less efficient vehicle for family trips. If this is the case, the effect of kept vehicle attributes on the follow-on vehicle may vary depending on which vehicle a household is replacing. To investigate these potentially heterogeneous effects, we allow the marginal effect of kept vehicle attributes to differ depending on whether the household retained the least or most fuel intense vehicle in their portfolio. As described in Section 3.2 this specification introduces three new endogenous variables and requires additional instruments to identify the system of first stage equations.

[30] We additionally considered whether the time elapsed between the purchase of the kept and bought vehicle explain some of our results. Conditioning our primary specification on fixed effects for the year the kept vehicle was purchased absorbs much of the variation in the kept vehicle gasoline price, which leads to a very weak first stage and no precision in the second stage estimates. Alternatively, including fixed effects for the number of years elapsed since the kept vehicle was purchased leads to a weaker first stage than the primary specification and marginal effects of kept vehicle fuel intensity that are larger in magnitude than our primary specification.

Table 6 presents the marginal effect of kept vehicle fuel intensity from these regressions on the full sample of vehicle transactions. The effect of kept vehicle fuel intensity is allowed to vary depending on whether households make the (endogenous) decision to keep the less ($f^d \geq f^k$) or more ($f^k > f^d$) fuel-intense vehicle in the portfolio. Column 1 presents the HHFEIV marginal effects from Table 5 for reference. Columns 2 and 3 present marginal effects using a cubic function of the "price deviations from trend" instrument described in Section 3.2. Columns 4 and 5 show estimates utilizing a linear spline of that same instrument.

Each specification indicates that an increase in fuel intensity of the kept car *decreases* the fuel intensity of the bought car. For households that dropped the more fuel-intense car (*i.e.*, columns 2 and 4), this implies that the household responds to an exogenous decrease in fuel intensity of the already more-efficient kept car by acquiring a more fuel intense car. This is consistent with substitution across attributes: if the kept car is made less fuel intense, the household prefers to substitute fuel intensity for other attributes in the bought car, increasing the fuel intensity of the bought car.[31]

For households that kept the more fuel-intense car, the results in columns 3 and 5 imply that the household responds to an exogenous decrease in the fuel intensity of the less-efficient kept car by acquiring an slightly more fuel intense bought car. This again implies a preference for attribute substitution: if the kept car is made more efficient, the household again prefers to substitute fuel economy for other attributes in the bought car.[32]

In general, the marginal effect of kept vehicle fuel intensity on the bought vehicle's fuel intensity is larger in magnitude when the kept vehicle is the more fuel intense vehicle in the portfolio. However, we are unable to reject the null that attribute substitution effects are identical when replacing either the most or least fuel-intense vehicle at all but the highest of gasoline prices observed in our sample.

Figures 4a and 4b display the marginal effects from a linear probability model of f^k_{it} on the probability of buying a car in the highest (red and least fuel efficient) or lowest (blue and most fuel efficient) GPM quartile. The qualitative story remains the same. Over most of the gasoline price range, increasing the GPM of the kept car increases (decreases) the probability of purchasing a car in the least (most) fuel intense quartile. This finding is qualitatively similar across new and used cars.

For households that keep their less fuel-intense vehicle, the magnitude of the attribute substitution effect increases with the gasoline price. This may be due to the relative importance of fuel intensity (or attributes correlated with fuel intensity) in household decisions when gasoline prices are higher. Thus far, all of the qualitative results hold similarly for both new and used car purchase

[31] In the case of the kept vehicle being less fuel intense, this result is also consistent with households diversifying their portfolio, for it suggests that if the more fuel intense kept car has an even further decrease in fuel intensity, households would respond by increasing the fuel intensity of the more fuel intense bought car.

[32] While above one could see the result as consistent with a household diversifying their portfolio, this is a case where we clearly see attribute substitution and *not* diversification.

instances.

4.3 The Role of Non-GPM Attributes

Our empirical specification intentionally omits many kept-car vehicle attributes from the set of control variables. This allows us to interpret the portfolio effect in what we believe is the most policy-relevant way: allowing other vehicle attributes to change along with GPM. We now directly explore the effect of changing kept car GPM on three vehicle attributes: footprint, engine displacement, and weight. We chose these attributes because they appear with good coverage in our dataset and because of their economic relevance. *Ceteris paribus*, increases in weight, power, and size increase fuel intensity, and thus it is likely that the results we have discussed thus far are (at least in part) operating via these attribute channels.

Table 7 displays marginal effects from specifications that are analogous to column 4 of Table 5, except with the alternative attribute of the bought car as the dependent variable.[33] For each attribute, a negative estimate can be interpreted as households demanding less of that bought car attribute as f_{it}^k increases. All point estimates are negative, providing evidence in favor of the hypothesis that the GPM portfolio effect that we observe is operating at least in part through portfolio preferences over other attributes.

The importance of the effect on vehicle footprint has direct policy relevance. Fuel economy standards in many countries worldwide are vehicle attribute-based (e.g., footprint in the US and weight in several countries), whereby larger or heavier cars receive a less stringent fuel economy requirement. To the extent that the portfolio effect manifests through preferences for vehicle size, there will be consequences for the realized effectiveness of fuel economy standards relative to expectations. We now turn to an exercise designed to illustrate the implications for policy.

5 Policy Implications

In this section we quantify the strength of the forces we uncover in Section 4. We do this through two thought experiments, each of which reflects a common policy intended to shift drivers into more fuel efficient cars. First, we measure the net effect of an increase in the fuel economy of the kept vehicle in a manner consistent with a "Cash for Clunkers" program. This exercise allows for the fuel economy of the bought vehicle to change in ways consistent with the results in Section 4. The second thought experiment uses empirical estimates on the welfare costs associated with Corporate Average Fuel Economy (CAFE) standards to measure the added welfare costs associated that result from attribute substitution.[34]

[33] We also condition on the same attribute of the kept vehicle, which we assume to be exogenous.

[34] Given the large number of households in our sample, estimating household-level treatment effects is infeasible. Therefore, we estimate the average treatment effect across all households. Under the assumption that households respond similarly in their choice of vehicle fuel intensity to changes in expected future operating costs and current capital costs, we are extracting the average treatment effect of interest for these policy simulations.

It is important to note that we do not provide a full *counterfactual* analysis, but rather rely on our estimated reduced-form relationships to provide insight into the policy implications. In the presence of either of these policies, prices and the availability of vehicles are likely to respond; indeed, both of the policies discussed above operate through changing equilibrium prices. Our reduced-form relationships are functions of observed equilibrium prices and vehicle availability; thus, any policy-induced supply responses would alter these reduced-form relationships. Our thought experiments are best viewed as measuring the importance of incorporating portfolio effects within structural models that seek to generate true counterfactuals of policies similar to cash for clunker programs and fuel economy standards. It is possible that the response of firms acts to undo a significant portion of our reduced-form effects, but this is unlikely to be costless to the firm and consumers.

Our first thought experiment investigates the net effect on gasoline consumption of decreasing the fuel intensity of a household's initial ("kept") vehicle through a limited-duration program causing a one-time reduction in the fuel intensity of vehicles purchased, analogous to the Cash for Clunkers program. Such policies are surprisingly common. In addition to the well-known U.S. federal Car Allowance Rebate System program in 2009, several states operate similar programs and they are also widely used in Europe and elsewhere internationally. For example, California and Texas currently have programs similar to the defunct federal Cash for Clunkers program.[35] Because incentives might exist for used-car purchases, we also include used vehicles in the thought experiment. We report the results for used cars in the appendix.

A cash for clunkers program will set in motion a number of forces. We show below that this can have a dramatic effect on the net fuel savings. Our estimates in Section 4 imply that, given a decrease in the fuel intensity of the kept vehicle, the next vehicle purchased by the household will be more fuel intense. However, there is likely to be additional changes in behavior. In particular, the exogenous reduction in the fuel intensity of the kept vehicle may also lead to changes in usage patterns across the household's two vehicles. For one, we might expect to see a rebound effect: decreasing the fuel intensity of the kept vehicle reduces the marginal cost of driving, leading to more miles traveled within the household (the "rebound effect"). We might also expect to see a change in the usage across vehicles in the household given that the *relative* fuel intensities of the two vehicles has changed. Furthermore, this shifting of mileage will be exacerbated by the fact that the newly purchased vehicle becomes even more fuel intense due to attribute-substitution. To account for these changes in usage patterns, in the next sub-section we augment our empirical results on attribute substitution with estimates on how changes in fuel intensity affect a household's total vehicle miles traveled, as well as how these miles are divided across the two vehicles within the household.

[35] California's program is called the California Vehicle Retirement Program. See https://www.cashforclunkers.org/california-cash-for-clunkers-program/.

5.1 Household Usage Substitution Across Vehicle Portfolio

Our household vehicle data include (roughly) biennial odometer readings. We use these data to estimate how usage responds to changes in the relative per-mile costs of vehicles within a household. The details of this empirical exercise are provided in Appendix A. In brief, we exploit two sources of variation in vehicle operating costs: variation in gasoline prices over time while holding the vehicle portfolio fixed and changes in operating costs resulting from changes in the fuel intensity of vehicles in the portfolio. For each vehicle $i \in \{1, 2\}$ in a two-vehicle portfolio, we compute the fuel cost in dollars per mile DPM_i as the price of gasoline, in dollars per gallon, times the fuel intensity, in gallons per mile. As these are two-vehicle portfolios, attributes of the other vehicle included in the regression are subscripted j.

We construct a yearly panel of two-vehicle households. For each vehicle i in year t, we compute the mean annual VMT (VMT_{it}) as miles driven between the closest preceding (at time $\underline{t}$) and upcoming (at time $\bar{t}$) odometer measurements for that vehicle.[36] We estimate the impact of operating costs on VMT using the following specification:

$$\log(VMT_{h,i,t}) = \beta_i DPM_{h,i,t} + \beta_j DPM_{h,j,t} + \Xi_h + \Theta_{h,t} + \Psi(\underline{t}, \bar{t}) + A^i(i) + A^j(j) + \varepsilon_{h,i,t}, \quad (5.1)$$

where Ξ_h are household fixed effects, $\Theta_{h,t}$ are fixed effects for the county of residence of household h in year t, $\Psi_{h,\bar{t}}$ are fixed effects controlling for seasonality in driving,[37] $A(.)$ are controls for vehicle attributes,[38] and $\varepsilon_{h,i,t}$ is an idiosyncratic error which may have arbitrary correlation within households. Following the identification strategy outlined in Section 3.1.1, we instrument for operating costs using the gasoline price at the time the vehicle was purchased multiplied by the current gasoline price. We elect a log linear specification to allow for gasoline prices to have heterogeneous impacts in the demand for VMT across vehicles in the household's portfolio depending on the fuel intensity of each vehicle.[39]

Estimates from Equation (5.1) are shown in Table 8. The top panel shows the impact of vehicle operating costs on VMT of the more fuel-intense vehicle and the second set show the impacts for the relatively fuel-efficient vehicle. In either case, an increase in the cost of driving (DPM) of one vehicle introduces an incentive to shift VMT from that vehicle to the other vehicle in the household

[36] We obtain odometer readings through DMV records each time a vehicle is transacted and at the time of biannual smog checks for vehicles six years and older.

[37] Demand for VMT follows seasonal patterns and odometer readings do not necessarily occur at the same time each year for a given vehicle. In fact, one may be concerned that the timing of odometer readings may be correlated with demand for VMT. We deploy two sets of controls to account for seasonality in the VMT measurement. First, in a simpler specification, we include fixed effects for the quarter-of-year of the upcoming odometer reading. In our preferred specification, we interact these fixed effects with counts of each quarter-of-year elapsed since the previous odometer reading.

[38] All specifications include nonparametric controls for the age of both vehicles in the portfolio and indicators for leased vehicles. Additional attribute controls in our preferred specification include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

[39] In a log-log specification, the log of the operating cost is the sum of the log of fuel intensity and the log of the gasoline price. This would imply a change in the gas price would have an identical impact on the allocation of VMT across vehicles regardless of the VMT of each vehicle in the portfolio.

portfolio. We focus here on the final two columns, (5) and (6), which include the richest sets of controls. Beyond operating costs, these specifications include household fixed effects, county of residence fixed effects, seasonality fixed effects (captured by the quarter the odometer was read). Column (5) additionally includes controls for vehicle age, attributes, and leases. As an alternative, Column (6) excludes controls for vehicle attributes. We discuss each column as well as robustness in the appendix.

It is important to note, columns (5) and (6) represent fundamentally different thought experiments in how households may respond to changes in vehicle operating costs. Column (5), which includes controls for vehicle attributes, replicates the EPA's approach to simulating CAFE effects, where technological advances reduce vehicle fuel intensity without altering underlying characteristics. Column (6), which excludes vehicle attribute controls, more closely follows a policy, such as Cash-for-Clunkers, where households are nudged into purchasing less fuel intense vehicles with corresponding changes in attributes associated with the change in fuel intensity, holding technology fixed.

The results on usage shifting are consistent with intuition. Increasing the cost per mile of a given vehicle in the household reduces the number of miles that particular vehicle is driven, but increases the mileage of the other vehicle. For example, for the average vehicle in our sample of households (i.e., a gallons per mile of 0.052) and at a gasoline price of \$2 per gallon, the estimates in Column (5) imply that increasing the dollars per mile of vehicle 1, the more fuel intensive vehicle in the household, by 10% (i.e., a change of $0.052 \cdot 2$) decreases the number of miles driven by vehicle 1 by 3.84% ($10\% \cdot -3.697 \cdot 0.052 \cdot 2$). The estimates in Panel 2 suggest that a large portion of these miles will be shifted to vehicle 2. In particular, that same 10% change in vehicle 1's cost per mile *increases* vehicle 2's miles driven by 2.24% ($10\% \cdot 2.152 \cdot 0.0502 \cdot 3$).

Considering Column (6), which excludes controls for vehicle attributes, the VMT elasticity of operating costs is nearly an order of magnitude smaller. A 10% increase in the fuel intensity of the more fuel-intense vehicle leads to a 0.58% ($10\% \cdot -0.557 \cdot 0.052 \cdot 2$) decrease in VMT of that vehicle. In contrast to specifications including attribute controls, the cross-vehicle operating cost effect on VMT is also negative – i.e., increasing the operating cost of the more fuel intense vehicle by 10% decreases VMT of the other vehicle in the portifolio by 0.82% ($10\% \cdot -0.789 \cdot 0.052 \cdot 2$). Effects are similar, but smaller, when considering the operating costs of the less fuel intense vehicle in the portfolio.[40] Next, we discuss implications that policy-makers may glean from these estimates.

5.2 Implications for Cash for Clunker Programs

With the estimates of usage in hand, we can calculate what our reduced-form relationships imply for the a decrease in the kept vehicle's fuel intensity through a one-time program such a Cash For Clunkers. We only require an assumed change in the kept vehicle's fuel economy. Hoekstra, Puller,

[40] As we discuss explore in detail in the next section, this result is consistent with households deriving utility from attributes such as horsepower or interior volume that are positively correlated with fuel intensity.

and West (2017) find that the 2009 Federal policy reduced fuel intensity by 3.5% for households taking advantage of the program.[41] We adopt this as our assumed change in the fuel intensity of the kept vehicle and restrict our sample to households that have replaced a vehicle. If a household has replaced more than one vehicle, we treat each replacement as an independent event.[42] For each household we observe the fuel economy of the kept and replacement vehicles, as well as the average annual miles driven for each vehicle. This allows us to generate baseline gasoline consumption for each vehicle.

We describe our procedure in the context of a two-vehicle household as follows:

1. We begin by reducing the fuel intensity of the kept vehicle by 3.5%. For example, suppose the kept vehicle's observed fuel intensity is 5.0 gallons per 100 miles (20 MPG), we decrease the fuel intensity to 4.83 gallons per 100 miles.

2. Using specification (4) from subtables (b) and (c) in Table 5, we increase the fuel intensity of follow-on vehicle. We estimate this effect using the in-sample mean gasoline price of $2. For example, suppose replaced vehicle was new and its observed fuel intensity was 4 gallons per 100 miles (25 MPG). Our simulated fuel intensity of the replaced vehicle would be: $4.0 + 0.38 \cdot 0.18 = 4.07$ gallons per 100 miles—an increase of 0.41%.

3. Using the own-effect coefficient in specification (6) in Table 8, we calculate the direct effect on miles driven on both vehicles. Continuing with our previous example, the kept vehicle is more fuel intense, corresponding to vehicle "1" in Table 8. In this case the direct effect on the kept vehicle's usage would be an increase of: $3.5 \cdot 0.577 \cdot 0.1 = 0.202\%$. The direct effect for the replacement vehicle, vehicle "2", would be an increase in miles driven of: $0.410 \cdot 0.422 \cdot 0.08 = 0.013\%$.[43]

4. Using the cross-effect coefficient in specification (6) in Table 8, we calculate the indirect effect on miles driven on both vehicles. The indirect effect on hypothetical kept vehicle would be to increase miles driven by: $0.41 \cdot 0.789 \cdot 0.08 = 0.026\%$. The indirect effect on the replacement vehicle would be to increase miles driven by: $3.5 \cdot 0.345 \cdot 0.1 = 0.121\%$.

5. We sum the direct and indirect effects to get a final change in miles driven. For our example, the miles driven of the kept vehicle would increase by: $0.202 + 0.121 = 0.323\%$. The miles driven of the replacement vehicle would decrease by: $0.013 + 0.026 = 0.040\%$.

6. We sum up the effects from each observed replacement in the data.

[41] Hoekstra, Puller, and West (2017) find Cash for Clunkers improved fuel economy between 0.65 and 0.81 MPG over the counterfactual purchase in the absence of the program. The mean vehicle kept by a two-car household in our sample has a fuel economy of 19.19 MPG, hence, 0.7 MPG increase translates to a 3.5% decrease in GPM.

[42] That is we do compound the effects of simulated changes in fuel intensity over follow-on replacements.

[43] The VMT regressions in Table 8 use a log-linear specification, so the estimated VMT elasticity of operating costs is the parameter on operating costs times the operating cost in dollars per mile. Within our sample, the mean fuel intensity is 0.052 GPM (19.19 MPG) and the mean gasoline price is almost exactly $2/gal, leading to a mean operating cost of approximately $0.1/mile.

Table 9 Panel (a) breaks down the effect into the components discussed above for the entire sample. We present the population average for all two-car households across cases where the dropped vehicle is the more or less fuel intense vehicle in the starting portfolio and across new and used vehicle purchases. The first row reports the baseline annual gasoline consumption for the kept (initial) and purchased vehicles (follow-on).[44]

The second row reports the first effect of our thought experiment. By construction the fuel consumption of the initial vehicle falls by 3.5%. Given the attribute substitution effect, this will *increase* the fuel consumption of the follow-on bought vehicle by an average of 12.5 gallons; this is over 60% of the fuel savings from the 3.5% improvement of the initial vehicle. The next two rows report the impact from changes in usage. The effect in the third row (Direct Effect on VMT) comes directly from the 3.5% decrease in the fuel intensity of the initial vehicle, but holds the fuel intensity of the follow-on vehicle constant. With the added miles on the initial vehicle, we see the miles driven by the follow-on vehicle decreasing (holding constant the intensity of the initial vehicle). For the two average vehicle types in our sample, this increases fuel savings because miles traveled are shifting to the initial vehicle. The net effect could theoretically increase or decrease total household VMT because the 3.5% decrease in fuel intensity of the initial vehicle also leads to a net increase in miles driven.[45] However, under attribute substitution, the relative fuel efficiency of the initial and follow-on vehicles changes as well. The fourth row (Indirect Effect on VMT) calculates the impact of the additional vehicle-usage shifting that comes from the fact that the relative fuel efficiency of the two vehicles will be changed. This shifts additional miles from the follow-on bought vehicle to the initial vehicle and reduces the fuel savings from our thought experiment.

Once all of the forces are considered, the fuel savings from our thought experiment fall from the naive estimate of 19.8 gallons to 6.4 gallons. Of this change, the attribute substitution effect plays a leading role, with 12.9 gallons or 96% of the erosion of the savings coming about from effects that are explicitly due to attribute substitution (note of course that even the direct effect on the initial vehicle is calculated in equilibrium along with all of the other changes). The dramatic reduction in savings is quite startling and may have unfortunate implications for the effectiveness of policies that explicitly regulate fuel economy, especially if such policies are short term in nature. Specifically, these calculations are most relevant to a fuel economy standard that increases, but then plateaus (an extremely common situation when looking historically). These forces would also be relevant for other policies mandating changes in fuel intensity, the effects are especially relevant

[44] Given technical progress and the positive trend in fuel prices in our data, the fuel consumption of the bought vehicle is lower than for the initial vehicle.

[45] In households where the kept vehicle is much less fuel efficient than the follow-on bought vehicle, this direct effect could decrease the fuel savings from our thought experiment. The reader might wonder whether our VMT specification should include the possibility of a knife-edge around the point where the kept vehicle's fuel efficiency equals the follow-on vehicle's fuel efficiency. A complete analysis of VMT shifting is beyond the scope of this paper but the topic of current research. An argument against such a knife-edge is that comfort and fuel economy tend to be inversely related. Therefore even if the kept vehicle is less efficient compared to the follow-on bought vehicle, a marginal decrease in the fuel intensity of the kept vehicle will lead to more miles traveled because it is, on average, the more comfortable vehicle within the household.

for such policies as Cash for Clunkers, which provided a one-time subsidy with the aim of reducing new vehicle fuel intensity. Our estimates suggest that 68% of the initial fuel savings from Cash for Clunkers would have been eroded from attribute substitution and rebound.

The thought experiment we consider here is a policy analogous to Cash for Clunkers, where some incentive nudges households into purchasing a less fuel intense vehicle, with the corresponding change in attributes, holding technology fixed. That is, efficiency gains come through reductions in attributes negatively correlated with fuel economy, such as size and horsepower. Alternatively, in Table 9 Panel (b), we perform a similar simulation, decreasing the fuel intensity of the initial purchase vehicle and allowing the effects to propagate to the follow-on purchase and driving decisions but holding the attributes of each vehicle fixed.[46] This simulation is analogous to the EPA's analysis of CAFE standards, where fuel intensity reductions are assumed to arise from technological advances, with no impact on other vehicle attributes.

The overall impact of a reduction in fuel intensity of the initial vehicle is quite similar to what would be expected under a program like Cash for Clunkers, however the channels through which fuel consumption changes are quite different. When holding attributes fixed, household driving behavior is much more responsive to operating costs, resulting in an annual increase in gasoline consumption of the initial vehicle of 10.64 gallons. The follow-on vehicle, now being less efficient in comparison, shows fuel consumption reduced by 9.9 gallons per year. The total effect of VMT changes nearly cancels out, leading to a net effect on fuel consumption very similar to a program like Cash for Clunkers, where technology is held fixed.

This difference in channels is unsurprising given recent work such as West et al. (2017), Hoekstra, Puller, and West (2017), and Leard, Linn, and Zhou (2017). These papers find that consumers value attributes, such as horsepower, that are positively correlated with fuel intensity. Computing the VMT response to operating costs holding attributes fixed, consistent with technological change reducing fuel intensity, focuses only on this single channel and households exhibit large VMT responses. When other attributes are not held constant, however, the VMT consumption reflects both value from reduced operating costs, but lost utility from other attributes that are correlated with fuel intensity, attenuating the response to changes in operating costs.

We additionally extend this thought experiment to all households holding three or fewer vehicle in their portfolio. This comprises approximately 85% of all households in California during the period of our sample. These results require estimates of attribute substitution and portfolio VMT effects for additional household types beyond the 2x2 portfolio, details of which are available in the Appendix. Table 10 panel (a) shows the population average effect from an exogenous decrease of 3.5% in the fuel intensity of one vehicle in a household's portfolio, through the one vehicle replacement event. The estimated effects are the average across all households with three or fewer vehicles.

[46]We hold vehicle attributes fixed by conditioning the household VMT response to operating costs on vehicle weight, footprint, and horsepower. These results are presented in Table 8 column (5).

Here, exogenously reducing the fuel intensity of a single vehicle in the household's portfolio has a mean direct effect of reducing gasoline consumption of 18.9 gallons per year. Households respond to this change in the attributes of their vehicle portfolio by purchasing a follow-on vehicle that is more fuel intense than they would have otherwise, increasing gasoline consumption by 3.9 gallons per year at the mean. This change in vehicle attributes alters the household's demand for VMT and the allocation across vehicles, increasing the mean gasoline consumption of the initial vehicles by 1.08 gallons per year and decreasing consumption by the follow-on vehicle 0.34 gallons per year.[47] On net, the exogenous reduction of fuel intensity of the kept vehicle reduces mean annual household gasoline consumption 13.6 gallons, about 72% of the naive estimate of 18.9 gallons.

These results provide a measure of the cost per ton of carbon dioxide emissions avoided under the Cash for Clunkers program. Assuming a typical vehicle life of 12 years, the results in Table 10 predict a total savings of 163.1 gallons of gasoline, avoiding 1.45 metric tons of carbon dioxide emissions. The average Cash for Clunkers subsidy in Hoekstra, Puller, and West (2017) was \$4210, leading to a cost of \$2904 per ton of avoided emissions. The bulk of this subsidy, however, is a transfer between the government and the household. Assuming a marginal cost of public funds of 0.3, the welfare cost of Cash for Clunkers was \$871 per ton of avoided carbon dioxide emissions.[48]

One could argue that these estimates are conservative. The used car market, which is not covered by cash for clunkers, is another channel through which attribute substitution may manifest. Increases in the fuel economy of initial kept vehicles due to an increased standard will increased demand for used fuel-inefficient vehicles. The increase in demand will lead used gas guzzlers to be more valuable, and thus more slowly retire from the fleet (similar to the effect documented in Jacobsen and van Benthem (2015)). Furthermore, it cannot realistically be argued that fuel economy of the bought vehicle will be constrained by the standards themselves. When stringency of the standard is a function of other attributes, such as size or weight (as is the case for fuel economy standards in every developed country in the world), inducing a change in demand for these attributes will directly affect the effective stringency of the standard.

5.3 Implications for the Welfare Effects of CAFE Standards

Our second thought experiment focuses on the welfare costs of CAFE standards. It is simpler in nature and relies on the empirical estimates of the welfare costs of CAFE in Jacobsen (2013). The thought experiment is a sustained change in the average fuel economy required under CAFE of 1 MPG. Therefore, the fuel economy of both the initial (kept) and follow-on (bought) vehicles are

[47] This small net change in VMT is consistent with the results of West et al. (2017) which find that Cash for Clunkers increased fuel efficiency of purchased vehicles, but those vehicles were small and less powerful. They find no evidence that, accounting for changes in other attributes, household VMT increased.

[48] As before, in Table 10 panel (b), we simulate the same exogeneous reduction in fuel intensity across all households holding 3 vehicles or less, holding attributes fixed when estimating the VMT response to changes in operating costs. The pattern is similar the 2x2 portfolio; reduced operating costs cause households to drive the initial vehicle more, attribute substitution leads to a more fuel intense follow-on vehicle, which households then drive less. The total reduction in fuel consumed is approximately 42% smaller than the naive estimate based on the fuel intensity alone.

forced to increase by 1 MPG. Jacobsen (2013) calculates the equivalent variation from such a 1 MPG change in CAFE, but his calculations do not include the additional welfare costs that would operate through attribute substitution. The welfare costs accounting for attribute substitution will be larger because the *desired* fuel economy of the follow-on vehicle will now be lower. Therefore forcing the follow-on vehicle's fuel economy to increase by 1 MPG will have a larger welfare consequence.

We can put numbers to this effect. Table 8 in Jacobsen (2013) implies that the average equivalent variation across all households from a 1 MPG increase in CAFE standards is $264 (in year 10). Our results suggest that the welfare costs for the bought vehicle will be larger because the desired fuel economy for the average bought vehicle in our data *decreases* by 0.40 MPG due to the increase in the initial vehicle's fuel economy. Therefore, a sustained 1 MPG increase in fleet fuel economy will be a 1.40 MPG increase in the average "desired" fuel economy of follow-on vehicles, again noting that these are not true counterfactual estimates. This would in turn increase the welfare costs by 40% to $368.

6 Conclusions

Much in the same way that *products* can be complements or substitutes, households may view *attributes* of one product as substitutable with or complementary to attributes of another. This observation is potentially relevant to understanding consumer decisions relating to a broad set of goods ranging from financial asset portfolios, household durables and clothing fashion to media subscriptions, higher education, and more. We develop an identification strategy to separate household preferences for level effects from attribute substitution within the household portfolio. Aspects of the methodology may be generally applied to many of the aforementioned household goods. This paper focuses on the vehicle market due to both the particular suitability of the choice setting for identifying attribute substitution and the potential implications for transportation market regulations.

The effects of a number of polices applied to the vehicle market depend crucially on consumer choice patterns. Empirical estimates of vehicle choice typically assume that the vehicle choices within a household are made independently. We provide evidence that this assumption does not hold. Using panel data on the portfolio of vehicles within a household and an instrumental variables approach, we find evidence that households exhibit a preference for attribute substitution. Exogenous decreases in the fuel intensity of the kept car increase the fuel intensity of the purchased car. We show this using both a continuous measure of fuel intensity, as well as by estimating the probability a household purchases a vehicle in the upper and lower quartiles of the fuel intensity distribution. An decrease in the fuel intensity of the kept car reduces the probability the household purchases a car in the lower quartile of gallons per mile, while such an increase reduces the probability the household buys a car in the upper quartile.

We also find that gasoline prices affect the preference for attribute substitution in intuitive ways.

As gasoline prices increase, the effect of the fuel intensity of the kept vehicle on the probability of buying a car in the lower quartile of fuel intensity becomes even more positive. In contrast, as gasoline prices decrease, the effect of fuel intensity of the kept vehicle on the probability of buying a car in the upper quartile of fuel intensity becomes even more negative. These effects manifest through substitution across vehicle attributes beyond fuel economy. We find evidence households view attributes such as vehicle size (footprint), weight, and horsepower as substitutes across their portfolio of vehicles.

These results have substantial economic importance for the understanding of major policies to improve the fuel economy of the light duty vehicle fleet, such as subsidies for more fuel-efficient vehicles (e.g., Cash-for-Clunkers) and fuel economy standards that increase and then level off. We use our results to estimate the net effect of a one-time exogenous decrease in fuel intensity of the kept vehicle and find that the attribute substitution effect can erode as much as 60% of the fuel savings from the decrease in fuel intensity. Moreover, our results suggest that this erosion of the savings is likely to be especially problematic under attribute-based standards, such as the current footprint-based standard in the United States and other countries in the world.

While this research setting considers household purchases in the vehicle market, these results highlight the challenges in design or evaluation of any policy intending to alter consumer choices over a portfolio of goods. When households view the attributes of those goods as substitutes, any shift in the attributes of one good will shift the unconstrained optimal choice for attributes of other goods in the portfolio in the opposite direction. This will lead to subsequent purchases which potentially erode or magnify effects of the policy or, if the policy also prevents the unconstrained optimal purchase, increases the true welfare costs of the policy.

References

Aizer, Anna and Janet Currie. 2017. "Lead and Juvenile Delinquency: New Evidence from Linked Birth, School and Juvenile Detention Records." Working Paper 23392, National Bureau of Economic Research. URL http://www.nber.org/papers/w23392.

Alcaraz, Carlo, Daniel Chiquiar, and Alejandrina Salcedo. 2012. "Remittances, schooling, and child labor in Mexico." *Journal of Development Economics* 97 (1):156 – 165. URL http://www.sciencedirect.com/science/article/pii/S0304387810001318.

Allcott, Hunt and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Review of Economics and Statistics* 96 (5):779–795.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe About Future Gasoline Prices?" *Journal of Environmental Economics and Management* 66 (3):383–403.

Angrist, Joshua D. and William N. Evans. 1998. "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size." *The American Economic Review* 88 (3):450–477. URL http://www.jstor.org/stable/116844.

Baum, Christopher F, Mark E Schaffer, and Steven Stillman. 2002. "IVREG2: Stata module for extended instrumental variables/2SLS and GMM estimation." URL https://ideas.repec.org/c/boc/bocode/s425401.html.

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen. 2012. "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain." *Econometrica* 80 (6):2369–2429.

Ben-Porath, Yoram and Finis Welch. 1976. "Do Sex Preferences Really Matter?" *Quarterly Journal of Economics* 90 (2):285–307.

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *American Economic Review* 99 (3):667–699.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 64 (4):841–890.

Borenstein, Severin. 2015. "A Microeconomic Framework for Evaluating Energy Efficiency Rebound and Some Implications." *Energy Journal* 36:1–21.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1):220–256.

Chevalier, Judith and Austan Goolsbee. 2009. "Are Durable Goods Consumers Forward-Looking? Evidence from College Textbooks." *Quarterly Journal of Economics* 124 (4):1853–1884.

Correia, Sergio. 2014. "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects." URL `https://ideas.repec.org/c/boc/bocode/s457874.html`.

Dorsch, Michael T., Karl Dunz, and Paul Maarek. 2015. "Macro shocks and costly political action in non-democracies." *Public Choice* 162 (3):381–404. URL `https://doi.org/10.1007/s11127-015-0239-x`.

Ellis, Frank. 2000. "The Determinants of Rural Livelihood Diversification in Developing Countries." *Journal of Agricultural Economics* 51 (2):289–302.

Gavazza, Alessandro. 2011. "The Role of Trading Frictions in Real Asset Markets." *American Economic Review* 101 (4):1106–1143.

Gentzkow, Matthew. 2007. "Valuing New Goods in a Model with Complementarity: Online Newspapers." *American Economic Review* 97:713–744.

Gillingham, Kenneth. 2011. "How Do Consumers Respond to Gasoline Price Shocks? Heterogeneity in Vehicle Choice and Driving Behavior." Manuscript: Yale University.

———. 2013. "The Economics of Fuel Economy Standards Versus Feebates." *NEPI Working Paper* .

Gillingham, Kenneth, David Rapson, and Gernot Wagner. 2016. "The Rebound Effect and Energy Efficiency Policy." *Review of Environmental Economics and Policy* 10:68–88.

Goulder, Lawrence H., Mark R. Jacobsen, and Arthur A. van Benthem. 2012. "Unintended Consequences from Nested State and Federal Regulations: The Case of Pavley Greenhouse-Gas-per-Mile Limits." *Journal of Environmental Economics and Management* 63:187–207.

Hendel, Igal and Alessandro Lizzeri. 1999. "Adverse Selection in Durable Goods Markets." *American Economic Review* 89 (5):1097–1115.

Hoekstra, Mark, Steven L. Puller, and Jeremy West. 2017. "Cash for corollas: When stimulus reduces spending." *American Economic Journal: Applied Economics* 9 (3):1–35.

Ito, Koichiro and James Sallee. 2014. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." *NBER Working Paper 20500* .

Jacobsen, Mark R. 2013. "Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity." *American Economic Journal: Economic Policy* 5 (2):148–187.

Jacobsen, Mark R. and Arthur A. van Benthem. 2015. "Vehicle Scrappage and Gasoline Policy." *American Economic Review* 105 (3):1312–1338.

Kellogg, Ryan. 2017. "Gasoline Price Uncertainty and the Design of Fuel Economy Standards." *NBER Working Paper 23024* .

Klier, Thomas and Joshua Linn. 2010. "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data." *American Economic Journal: Economic Policy* 2 (3):134–153.

———. 2012. "New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standards." *RAND Journal of EConomics* 43 (1):186–213.

Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review* 101 (7):3368–3399.

Larrick, Richard and Jack Soll. 2008. "The MPG Illusion." *Science* 320 (5883):1593–1594.

Leard, Benjamin, Joshua Linn, and Y Zhou. 2017. "How much do consumers value fuel economy and performance."

Manski, Charles F. and Leonard Sherman. 1980. "An Empirical Analysis of Household Choice Among Motor Vehicles." *Transportation Research A* 14 (5-6):349–366.

Michaels, Guy. 2008. "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System." *The Review of Economics and Statistics* 90 (4):683–701. URL https://doi.org/10.1162/rest.90.4.683.

Newman, Peter. 1994. "Substitutes and complements." *New Palgrave Dictionary of Economics* .

Samuelson, Paul. 1974. "Complementarity: An essay on the 40th anniversary of the hicks-allen revoluation in demand theory." *Journal of Economic Literature* 12 (4):1255–1289.

Stock, James H and Motohiro Yogo. 2002. "Testing for Weak Instruments in Linear IV Regression." *NBER Technical Working Paper* (T0284).

Udry, Christopher, John Hoddinott, Harold Alderman, and Lawrence Haddad. 1995. "Gender Differentials in Farm Productivity: Implications for Household Efficiency and Agricultural Policy." *Food Policy* 20 (5):407–423.

Wakamori, Naoki. 2011. "Portfolio Considerations in Differentiated Product Purchases: An Application to the Japanese Automobile Market." *Bank of Canada Working Paper 2011-27* .

West, Jeremy, Mark Hoekstra, Jonathan Meer, and Steven L. Puller. 2017. "Vehicle miles (not) traveled: Fuel economy requirements, vehicle characteristics, and household driving." *Journal of*

Public Economics 145:65 – 81. URL http://www.sciencedirect.com/science/article/pii/S0047272716301463.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data (Hardcover)*. MIT Press, 3 ed.

Table 1: Number of Unique Households by Portfolio Size

Start	End Portfolio Size			
Portfolio Size	1	2	3	4+
1	7,262,111	1,360,594	187,558	75,150
2	1,172,278	4,632,425	839,546	259,098
3	168,745	849,703	2,169,948	675,040
4+	35,810	141,618	381,226	1,489,926

Each cell represents the count of unique households from 2001 to 2007 observed to have the starting portfolio size shown in each row and the ending portfolio size shown in the column. These counts provide a measure of the number of households providing identifying variation in each portfolio cell. A single household may appear in multiple cells if their portfolio changes over time but is counted at most once in each cell. For example, two-car household that replaces one car every year would add one to the count of the (2,2) cell. If instead, that household adds a third vehicle in 2004 and returns to a two-car portfolio in 2006 it would add one to the count of the (2,2) cell, one to the count of the (2,3) cell, one to the (3,3) cell, and one to the count of the (3,2) cell. Each household may have zero, one, or multiple vehicle transactions during this time period.

Table 2: Summary Statistics for Continuous Variables - 2x2 Replacement Households

	All Households	Bought GPM Qtile 1	Bought GPM Qtile 2 or 3	Bought GPM Qtile 4
Kept Vehicle GPM	0.0522	0.0507	0.0523	0.0533
	(0.0106)	(0.0106)	(0.0105)	(0.0108)
Bought Vehicle GPM	0.0516	0.0388	0.0503	0.0664
	(0.0108)	(0.0037)	(0.0036)	(0.0063)
Dropped Vehicle GPM	0.0511	0.0478	0.0507	0.0549
	(0.0103)	(0.0098)	(0.0094)	(0.0112)
Gasoline Price at Bought Purchase (US$)	2.380	2.434	2.377	2.335
	(0.747)	(0.763)	(0.745)	(0.733)
Gas Price at Kept Vehicle Purchase (US$)	2.064	2.105	2.056	2.041
	(0.539)	(0.558)	(0.534)	(0.527)
(Kept - Sold) Value DfT (US$)	4.483	22.629	2.567	-9.056
	(848.059)	(731.634)	(850.055)	(942.603)
Kept Vehicle Age (yr)	7.320	7.466	7.450	6.919
	(5.924)	(5.962)	(5.944)	(5.828)
Dropped Vehicle Age (yr)	9.948	10.651	9.990	9.187
	(5.899)	(5.849)	(5.847)	(5.957)
Kept vehicle value (US$)	9,905	9,082	9,953	10,626
	(8,352)	(7,295)	(8,457)	(9,024)
Bought Vehicle Value (US$)	11,283	7,468	11,819	13,911
	(9,162)	(5,267)	(9,999)	(9,273)
Dropped Vehicle Value (US$)	7,794	6,149	7,957	9,072
	(7,871)	(5,970)	(8,126)	(8,671)
N Transactions	2,004,312	491,010	1,003,044	510,258
N Households	1,452,896	392,168	768,517	413,367

Summary statistics of continuous variables for 2x2 replacement households. Standard deviations shown in parentheses.

Table 3: Distribution of observed fuel economy

Percentile	Gallons per Mile (GPM)	Miles Per Gallon (MPG)
25th Percetile	0.045	22.0
Median	0.052	19.3
75th Percentile	0.059	17.0

Percentiles of observed fuel intensity and corresponding fuel economy from the 2x2 replacement sample.

Table 4: Regression Estimates

(a) All Vehicle Transactions

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
GPM^K	0.091	-2.551	-0.115	-0.830
	(0.004)***	(0.509)***	(0.013)***	(0.300)***
$GPM^K \times P^G$	-0.004	0.000	-0.139	-0.186
	(0.002)**	(0.004)	(0.005)***	(0.012)***
N Non-singleton	1,171,976	1,169,006	509,664	508,407
Kleibergen-Paap rk F	.	18.72	.	17.77

(b) New Vehicle Transactions

New				
GPM^K	0.086	-2.266	-0.145	-0.106
	(0.006)***	(0.683)***	(0.023)***	(0.499)
$GPM^K \times P^G$	0.007	0.009	-0.121	-0.134
	(0.003)***	(0.006)	(0.010)***	(0.035)***
N Non-singleton	520,181	519,397	205,788	205,492
Kleibergen-Paap rk F	.	8.53	.	3.14

(c) Used Vehicle Transactions

Used				
GPM^K	0.086	-0.568	-0.110	-0.470
	(0.006)***	(0.363)	(0.021)***	(0.262)*
$GPM^K \times P^G$	-0.007	-0.006	-0.137	-0.166
	(0.002)***	(0.003)*	(0.008)***	(0.010)***
N Non-singleton	651,795	649,609	260,624	259,960
Kleibergen-Paap rk F	.	12.54	.	24.67

Regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. IV and HHFEIV specificationsuse the gasoline price at the time of kept vehicle purchase and its interaction with the current gasoline price as described in 3.1.1 as instruments for endogenous regressors. Weak instrument diagnostics robust to heteroskdasticity and clustering with household prodvided by the Kleinbergen-Paap rk F. Additional weak instrument tests and diagnostics shown in Table 20 of the Appendix.

Table 5: Marginal Effect of bought vehicle fuel intensity on kept vehicle fuel intensity

(a) All Vehicle Transactions

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
$P^G = \$2.00$	0.082	-2.551	-0.392	-1.202
	(0.002)***	(0.509)***	(0.007)***	(0.319)***
$P^G = \$3.00$	0.077	-2.551	-0.531	-1.389
	(0.002)***	(0.509)***	(0.008)***	(0.329)***
$P^G = \$4.00$	0.073	-2.551	-0.669	-1.575
	(0.004)***	(0.509)***	(0.012)***	(0.339)***

(b) New Vehicle Transactions

New				
$P^G = \$2.00$	0.101	-2.249	-0.387	-0.375
	(0.003)***	(0.687)***	(0.012)***	(0.566)
$P^G = \$3.00$	0.108	-2.240	-0.509	-0.509
	(0.004)***	(0.690)***	(0.015)***	(0.599)
$P^G = \$4.00$	0.115	-2.231	-0.630	-0.643
	(0.005)***	(0.692)***	(0.022)***	(0.633)

(c) Used Vehicle Transactions

Used				
$P^G = \$2.00$	0.072	-0.580	-0.384	-0.801
	(0.003)***	(0.366)	(0.010)***	(0.267)***
$P^G = \$3.00$	0.064	-0.586	-0.521	-0.967
	(0.003)***	(0.367)	(0.012)***	(0.271)***
$P^G = \$4.00$	0.057	-0.593	-0.658	-1.133
	(0.005)***	(0.369)	(0.018)***	(0.274)***

Marginal effect of kept vehicle GPM on bought vehicle GPM from a regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. IV and HHFEIV specifications use the gasoline price at the time of kept vehicle purchase and its interaction with the current gasoline price as described in 3.1.1 as instruments for endogenous regressors.

Table 6: Marginal Effect of Kept Vehicle GPM on Purchased Vehicle GPM by Relative Portfolio Position

	No Interaction	Value DiD - Cubic		Value DiD - Spline	
		$f^d \geq f^k$	$f^k > f^d$	$f^d \geq f^k$	$f^k > f^d$
	(1)	(2)	(3)	(4)	(5)
$P^G = \$2.00$	-1.202	-1.212	-0.974	-1.194	-0.944
	(0.319)***	(0.181)***	(0.118)***	(0.188)***	(0.121)***
$P^G = \$3.00$	-1.389	-1.609	-1.221	-1.574	-1.177
	(0.329)***	(0.192)***	(0.126)***	(0.198)***	(0.127)***
$P^G = \$4.00$	-1.575	-2.005	-1.468	-1.954	-1.411
	(0.339)***	(0.205)***	(0.136)***	(0.210)***	(0.136)***
N Non-Singleton	508,407	303,772		303,772	
Kleibergen-Paap rk F	17.77	31.25		35.19	
Instruments:					
p^k_{gas}	Y	Y		Y	
$\Delta\Delta P^{kd}$	N	Cubic		Linear Spline	

Marginal effects of kept vehicle GPM from a regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. Columns 2/3 and 4/5 each show results from a single regression. $f^d \geq f^k$ ($f^k > f^d$) show marginal effects when the dropped vehicle was the most (least) fuel intense vehicle in the portfolio prior to the purchase. Columns 2 through 5 instrument for endogenous regressors using gas price at the time of kept vehicle purchase, gas price at the time of dropped vehicle purchase, the either a cubic (Columns 2 and 3) or a linear spline with knots at ±500 (Columns 4 and 5) of "Price deviations from trend" instruments, and projections from the space of exogenous variables described in Section 3.2 as instruments for endogenous regressors. Weak instrument diagnostics robust to heteroskdasticity and clustering with household prodvided by the Kleinbergen-Paap rk F.

Table 7: Bought Vehicle Attributes - Kept GPM and Attribute Marginal Effects

(a) All Vehicles

	Footprint (1)	Displacement (2)	Curb Weight (3)
$P^G = \$2.00$	-56,171	-6,286	-44,635
	(23,949)**	(2,957)**	(10,976)***
$P^G = \$3.00$	-66,240	-7,150	-53,970
	(25,069)***	(3,009)**	(11,460)***
$P^G = \$4.00$	-76,309	-8,014	-63,306
	(26,208)***	(3,062)***	(11,957)***
Kept Vehicle Attribute	-0.226	0.179	0.007
	(0.057)***	(0.269)	(0.116)
N Non-singleton	512,423	513,924	507,435
Kleinbergen-Paap rk F	48.35	16.40	120.49
Bought Vehicle Attribute	Footprint	Engine Disp.	Curb wt.
Units	ft ^{2}	L	tons

Marginal effects of kept vehicle GPM from a regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. All specifications instrument for endogenous kept vehicle GPM using the gasoline price at the time of kept vehicle purchase and its interaction with the current gasoline price. Weak instrument diagnostics robust to heteroskdasticity and clustering with household prodvided by the Kleinbergen-Paap rk F.

Table 7: Bought Vehicle Attributes - Kept GPM and Attribute Marginal Effects (cont.)

(b) New Vehicles

	Footprint (1)	Displacement (2)	Curb Weight (3)
$P^G = \$2.00$	-90,481	-1,395	-52,707
	(31,263)***	(3,716)	(12,772)***
$P^G = \$3.00$	-100,066	-1,961	-60,289
	(33,547)***	(3,790)	(13,486)***
$P^G = \$4.00$	-109,651	-2,527	-67,870
	(35,876)***	(3,866)	(14,243)***
Kept Vehicle Attribute	-0.068	-0.257	0.032
	(0.077)	(0.325)	(0.127)
N Non-singleton	210,138	210,347	208,829
Kleinbergen-Paap rk F	20.39	6.82	76.25
Bought Vehicle Attribute	Footprint	Engine Disp.	Curb wt.
Units	ft ^{2}	L	tons

(c) Used Vehicles

	Footprint (1)	Displacement (2)	Curb Weight (3)
$P^G = \$2.00$	1,902	-3,069	-11,383
	(27,654)	(2,949)	(14,084)
$P^G = \$3.00$	-6,429	-3,901	-19,749
	(28,412)	(3,000)	(14,539)
$P^G = \$4.00$	-14,760	-4,732	-28,115
	(29,213)	(3,051)	(15,017)*
Kept Vehicle Attribute	-0.380	-0.110	-0.307
	(0.061)***	(0.270)	(0.150)**
N Non-singleton	259,154	260,318	255,936
Kleinbergen-Paap rk F	41.66	16.00	70.88
Bought Vehicle Attribute	Footprint	Engine Disp.	Curb wt.
Units	ft ^{2}	L	tons

Marginal effects of kept vehicle GPM from a regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. All specifications instrument for endogenous kept vehicle GPM using the gasoline price at the time of kept vehicle purchase and its interaction with the current gasoline price. Weak instrument diagnostics robust to heteroskdasticity and clustering with household prodvided by the Kleinbergen-Paap rk F.

Table 8: Regression of Log VMT on Fuel Cost Per Mile

Outcome		(1)	(2)	(3)	(4)	(5)	(6)
$\log(VMT_1)$	DPM_1	-0.557	-3.705	-3.709	-3.718	-3.697	-0.557
		(0.057)***	(0.113)***	(0.113)***	(0.114)***	(0.114)***	(0.057)***
	DPM_2	-0.806	2.433	2.435	2.457	2.435	-0.789
		(0.065)***	(0.139)***	(0.139)***	(0.141)***	(0.141)***	(0.066)***
$\log(VMT_2)$	DPM_2	-0.420	-3.926	-3.925	-3.968	-3.952	-0.422
		(0.071)***	(0.145)***	(0.145)***	(0.146)***	(0.146)***	(0.072)***
	DPM_1	-0.363	2.105	2.100	2.157	2.152	-0.345
		(0.058)***	(0.116)***	(0.116)***	(0.117)***	(0.117)***	(0.059)***
N		2,942,024	2,942,024	2,942,024	2,903,315	2,903,315	2,903,315
N Households		854,299	854,299	854,299	845,121	845,121	845,121
Household FE		Y	Y	Y	Y	Y	Y
County FE		N	N	Y	N	Y	Y
Seasonality FE		N	N	N	Y	Y	Y
Attribute Controls		N	Y	Y	Y	Y	N

Regression of vehicle log VMT on the operating cost level for each vehicle and covariates. Reported coefficients are semi-elasticites; the elasticity can be approximated multiplying each coefficient by the vehicle operating cost, approximately 0.05 GPM + \$2.00/gal = 0.1 $\frac{\$}{mi}$ mean in-sample. Variables subscripted with 1 denote the more fuel intense vehicle and 2 denotes the less fuel intense vehicle. Vehicle cost per mile (DPM) instrumented using gasoline price at the time the vehicle was purchased and its interaction with current gasoline prices. All regressions include household fixed effects, nonparametric controls for the age of both vehicles in the portfolio, county-level unemployment, and indicators for leased vehicles. Standard errors clustered by household shown in parentheses. Seasonality fixed effects account for seasonal patterns in driving behavior and consist of the quarter of year of the most recent VMT measurement interacted with counts of each quarter type since the previous VMT measurement. Attribute controls include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

Table 9: Effect of a Policy-Induced Decrease in Fuel Intensity on Vehicle Fuel Consumption, All 2x2 Households

(a) Cash-for-Clunkers Thought Experiment (No Attribute Controls)

	Initial Vehicle	Follow-on Vehicle	Portfolio Total
Base Fuel Consumption (gal/yr)	561.44	594.64	1,156.08
Direct Effect (gal/yr)	-19.76	12.49	-7.27
	[-100.00%]	[63.21%]	[-36.79%]
Direct AS Effect on VMT (gal/yr)	1.00	-0.63	0.37
	[5.09%]	[-3.19%]	[1.90%]
Indirect AS Effect on VMT (gal/yr)	-0.69	1.23	0.54
	[-3.49%]	[6.24%]	[2.75%]
Total Effect (gal/yr)	-19.44	13.09	-6.35
	[-98.40%]	[66.26%]	[-32.15%]

(b) CAFE Thought Experiment (With Attribute Controls)

	Initial Vehicle	Follow-on Vehicle	Portfolio Total
Base Fuel Consumption (gal/yr)	561.44	594.64	1,156.08
Direct Effect (gal/yr)	-19.76	12.49	-7.27
	[-100.00%]	[63.21%]	[-36.79%]
Direct AS Effect on VMT (gal/yr)	7.85	-4.92	2.92
	[39.71%]	[-24.91%]	[14.80%]
Indirect AS Effect on VMT (gal/yr)	2.79	-4.98	-2.20
	[14.10%]	[-25.22%]	[-11.12%]
Total Effect (gal/yr)	-9.12	2.58	-6.54
	[-46.18%]	[13.08%]	[-33.11%]

Average effect of a 3.5% decrease in fuel intensity of a vehicle through the purchase of the next vehicle across all 2x2 households. Assumes an average 0.7 MPG improvement from Hoekstra, Puller, and West (2017) across a fleet average 19.19 MPG. Panel (a) models the VMT response without attribute controls (Column 6 of Table 8). Panel (b) models the VMT response controlling for vehicle attributes (Column 5 of Table 8). Direct Effect is the effect from the exogeneous shock to the fuel intensity of the kept vehicle and the resulting change in fuel intensity of the follow-on purchase. Direct Effect on VMT is the own-vehicle effect in fuel consumption due to the change in operating costs changing VMT. Indirect Effect is the effect of cross-vehicle substitution of VMT. Base fuel consumption and vehicle VMT are the sample mean for two car households. VMT effect assume a gasoline price of $2 per gallon. Each effect size as a percentage of the direct effect to kept vehicles shown in brackets.

Table 10: Effect of a Policy-Induced Decrease in Fuel Intensity of Vehicle Fuel Consumption, Three or Fewer Vehicles

(a) Cash-for-Clunkers Thought Experiment (No Attribute Controls)

	Initial Vehicle(s)	Follow-on Vehicle	Portfolio Total
Base Fuel Consumption (gal/yr)	764.42	570.49	1,334.91
Direct Effect (gal/yr)	-18.88	3.87	-15.01
	[-100.00%]	[20.48%]	[-79.52%]
Direct AS Effect on VMT (gal/yr)	1.18	-0.20	0.98
	[6.26%]	[-1.04%]	[5.22%]
Indirect AS Effect on VMT (gal/yr)	-0.10	0.54	0.44
	[-0.53%]	[2.84%]	[2.31%]
Total Effect (gal/yr)	-17.79	4.20	-13.59
	[-94.27%]	[22.27%]	[-72.00%]

(b) CAFE Thought Experiment (With Attribute Controls)

	Initial Vehicle(s)	Follow-on Vehicle	Portfolio Total
Base Fuel Consumption (gal/yr)	764.42	570.49	1,334.91
Direct Effect (gal/yr)	-18.88	3.87	-15.01
	[-100.00%]	[20.48%]	[-79.52%]
Direct AS Effect on VMT (gal/yr)	7.38	-1.54	5.85
	[39.12%]	[-8.14%]	[30.99%]
Indirect AS Effect on VMT (gal/yr)	0.40	-2.17	-1.77
	[2.10%]	[-11.49%]	[-9.39%]
Total Effect (gal/yr)	-11.09	0.16	-10.93
	[-58.78%]	[0.85%]	[-57.93%]

Average effect of a 3.5% decrease in fuel intensity of a vehicle through the purchase of the next vehicle across all households. Assumes an average 0.7 MPG improvement from Hoekstra, Puller, and West (2017) across a fleet average 19.19 MPG. When households keep more than one vehicle in the portfolio, the most valuable vehicle decreases fuel-intensity. Population mean effect over 1x1, 1x2, 1x3, 2x1, 2x2, 2x3, 3x1,3x2, and 3x3 households. These comprise over 85% of the population of households. Panel (a) models the VMT response without attribute controls (Column 6 of Table 8). Panel (b) models the VMT response controlling for vehicle attributes (Column 5 of Table 8). Direct effect is the exogenous change in fuel consumption (for "Initial Vehicles") or the households optimal response in fuel intensity of the purchased vehicle (for "Follow-on Vehicle"). Direct Effect on VMT is the own-vehicle effect in fuel consumption due to the change in operating costs changing VMT. Indirect Effect is the effect of cross-vehicle substitution of VMT. Base fuel consumption and vehicle VMT are the sample mean for two car households. VMT effect assume a gasoline price of $2 per gallon. Each effect size as a percentage of the direct effect to kept vehicles shown in brackets.

Figure 1: First stage relationship: Gas price at time of kept car purchase

(a) New Vehicle Purchases **(b)** Used Vehicle Purchases

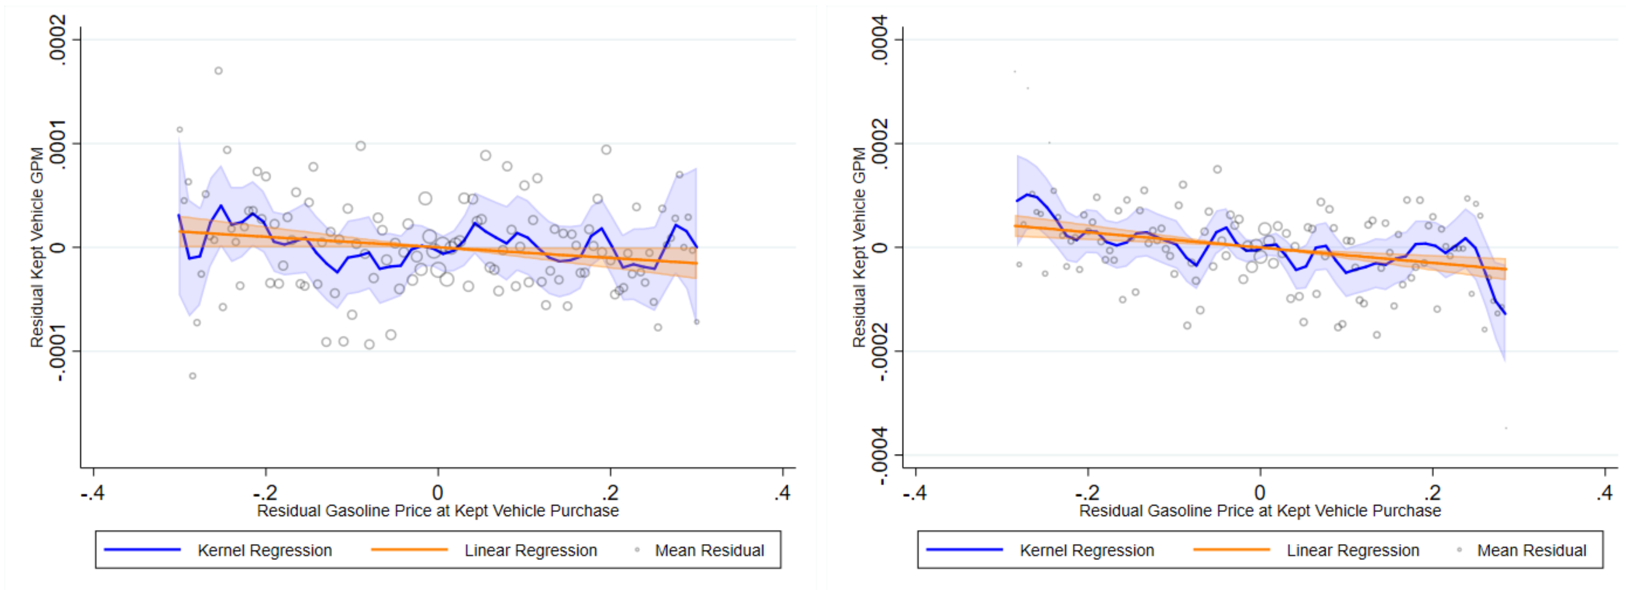

Plot of the first stage relationship between gasoline price at the time of kept vehicle purchase and the fuel economy (in GPM) of that vehicle. Both variables are partialed of all other regression covariates (Household, vehicle class, manufacturer, age, year, month, and new purchase fixed effects, and continuous controls for the current gasoline price, and the unemployment rate). Graphs are limited to the 5th through 95th percentiles of residual kept vehicle gasoline price. Excludes observations where the household fixed effect perfectly predicts the outcome. Blue line is a kernel regression with Epanechnikov kernel and bandwidth 0.01. The orange line is the linear relationship estimated using OLS. Shaded regions are 95% confidence intervals from these estimates. Gray circles are mean residuals for each 0.005 in kept vehicle GPM. The size of each circle is proportional to the number of observations used to compute the mean residual.

Figure 2: Instrumental variables first stage relationship - Price deviation from trend DiD IV

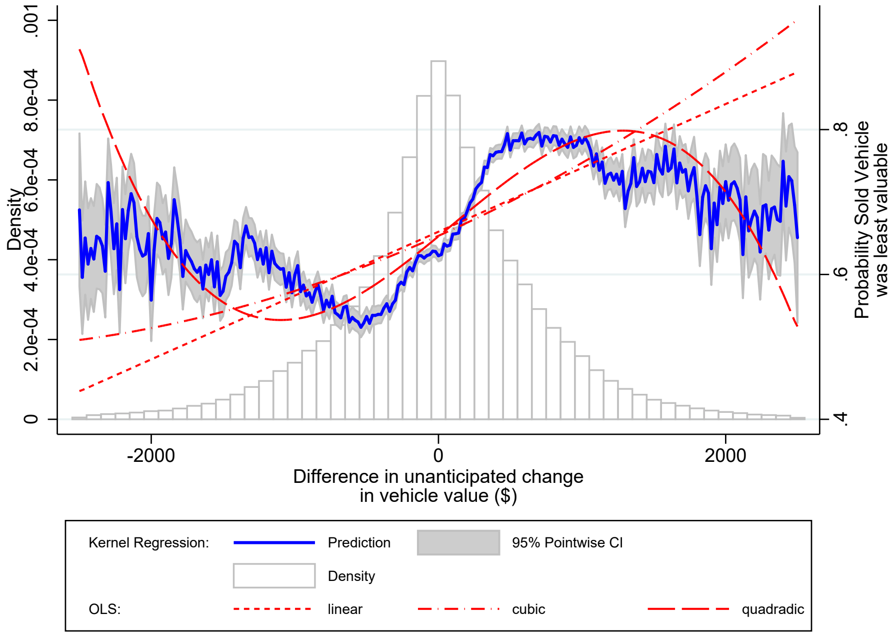

Plot of the first stage relationship between the price deviations from trend instrument (DfT) described in section 3.2 and probability a household drops the lowest value vehicle. Both variables are partialed of all other regression covariates (Household, vehicle class, manufacturer, age, year, month, and new purchase fixed effects, and continuous controls for the current gasoline price, and the unemployment rate). Graphs are limited to the 5th through 95th percentiles of the DfT instrument. Excludes observations where the household fixed effect perfectly predicts the outcome. Probabilities conditional on a vehicle purchase (new or used) estimated within $1,000 bins and shown in blue with binomial 95% confidence intervals shown in gray. Red lines represent a linear, quadradic, and cubic relationship between the instrument and the outcome.

Figure 3: Number of Transactions per 2x2 Replacement Household

(a) Full Sample

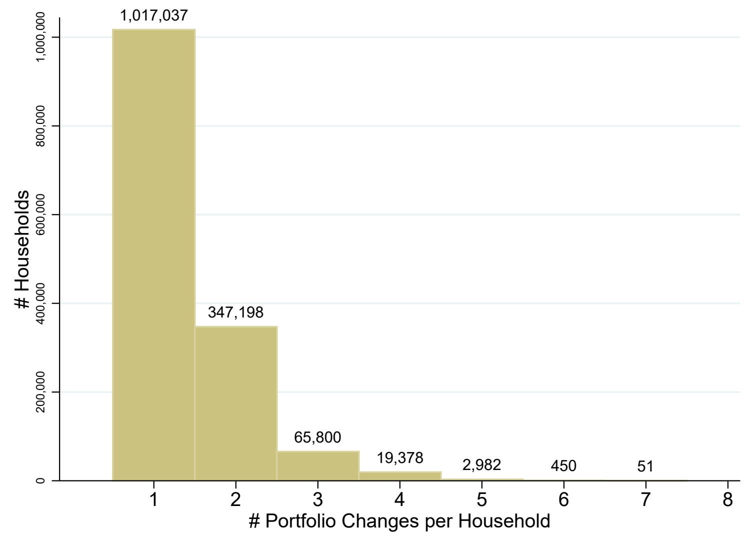

(b) IV Sample

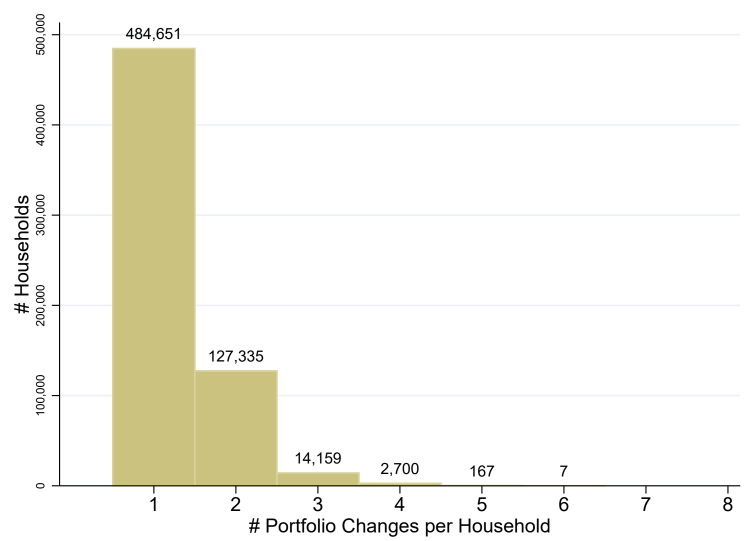

Distribution of the total number of observed vehicle transactions for each household from 2001 to 2007 for (a) the full sample of 2x2 replacement households and (b) households for which the data support deploying our IVs. In specifications including household fixed effects the fixed effect perfectly predicts the decision of a household if it only engages in one transaction. Other model parameters are identified by households engaging in multiple transactions from 2001 to 2007.

Figure 4: Marginal Effect of Kept Vehicle GPM on Bought Vehicle GPM

(a) New Vehicle Purchases

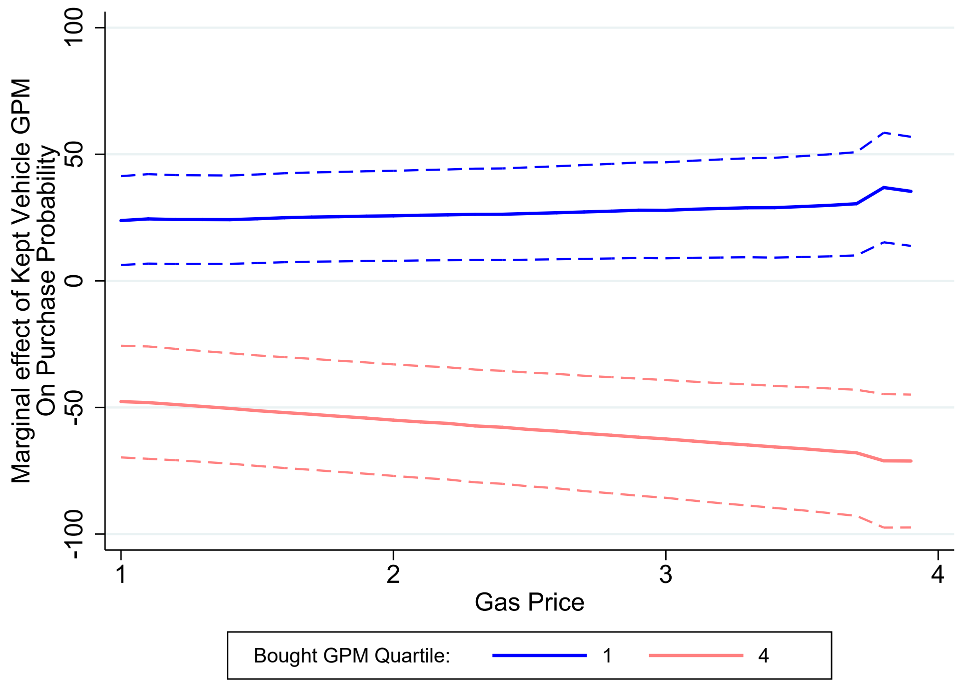

(b) Used Vehicle Purchases

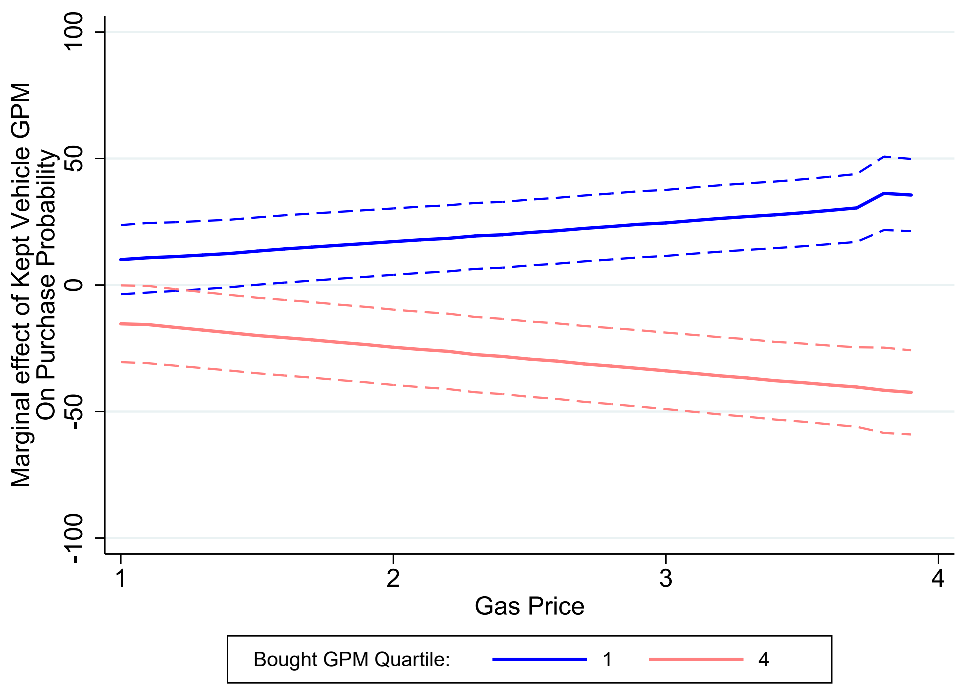

Population average marginal of the kept vehicle GPM on the probability model a household purchases a vehicle in the 1st (blue) or 4th (red) quartile of the GPM distribution for used vehicle purchases. Estimated as a linear probability model using the specification shown in Equation (3.2). 95% confidence intervals robust to heteroskedastiticy and clustered by household shown in dashed lines.

A Substituting Vehicle-Miles Traveled

When the cost of transportation services changes differentially across vehicles in the portfolio, households may adjust their usage on the intensive margin by substituting towards the less expensive vehicle. We investigate the relationship between vehicle operating cost per mile and the household's allocation of VMT across the vehicles in its portfolio. We exploit two sources of variation in vehicle operating costs: variation in gasoline prices over time while holding the vehicle portfolio fixed and changes in operating costs resulting from changes in the fuel intensity of vehicles in the portfolio. For each vehicle $i \in \{1, 2, \ldots, J\}$ in a J-vehicle portfolio, we compute the fuel cost in dollars per mile DPM_i as the price of gasoline, in dollars per gallon, times the fuel intensity, in gallons per mile. As these are J-vehicle portfolios, attributes and operating costs for each vehicle in the portifolio are indexed by j. Own-vehicle fuel economy effects are cases where $i = j$, with other cases representing attributes and cross-portfolio effects.

We construct a yearly panel of J-vehicle households. For each vehicle i in year t, we compute the mean annual VMT (VMT_{it}) as miles driven between the closest preceding (at time $\underline{t}$) and upcoming (at time $\bar{t}$) odometer measurements for that vehicle.[49] We estimate the impact of operating costs on VMT using Equation A.1,

$$VMT_{h,i,t} = \sum_{j=1}^{J} \beta_j DPM_{h,j,t} + \Xi_h + \Theta_{h,t} + \Psi(\underline{t}, \bar{t}) + \sum_{j=1}^{J} A^j(j) + \varepsilon_{h,i,t} \tag{A.1}$$

where Ξ_h are household fixed effects, $\Theta_{h,t}$ are fixed effects for the county of residence of household h in year t, $\Psi_{h,\bar{t}}$ are fixed effects controlling for seasonality in driving,[50] $A()$ are controls for vehicle attributes,[51] and $\varepsilon_{h,i,t}$ is an idiosyncratic error which may have arbitrary correlation within households.

Estimating the causal impact of operating costs on VMT consumption using a household panel presents challenges to identification. A household's vehicle operating costs are the product of the current gasoline price, which we assume to be exogenous, and the fuel intensity of the household's vehicles. Over time the change in a household's preference for vehicles with low fuel intensity, or other attributes correlated with low fuel intensity, may be correlated with changes in demand for VMT over time. For example, a positive household income shock may increase demand for vehicle

[49] We obtain odometer readings through DMV records each time a vehicle is transacted and at the time of biannual smog checks for vehicles six years and older.

[50] Demand for VMT follows seasonal patterns and odometer readings do not necessarily occur at the same time each year for a given vehicle. In fact, one may be concerned that the timing of odometer readings may be correlated with demand for VMT. We deploy two sets of controls to account for seasonality in the VMT measurement. First, in a simpler specification, we include fixed effects for the quarter-of-year of the upcoming odometer reading. In our preferred specification, we interact these fixed effects with counts of each quarter-of-year elapsed since the previous odometer reading.

[51] All specifications include nonparametric controls for the age of both vehicles in the portfolio and indicators for leased vehicles. Additional attribute controls in our preferred specification include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

horsepower (increasing fuel intensity) but also increase demand for driving in general. Similar to previous regressions, we instrument for operating cost per mile using the gasoline price at the time the vehicle was purchased and its interaction with current gasoline prices. Finally, as households may shift VMT between the more and less fuel-intense vehicles in response to changes in operating costs, we estimate separate regressions for the more (subscript 1) and less (subscript 2) fuel intense vehicle in the portfolio.

Estimates from Equation (A.1) for two-vehicle households are shown in Table 8. Similar estimates for one and three-vehicle households are shown in Tables 11 and 12, respectively. The sets of coefficients at the top of each section show the impact of vehicle operating costs on VMT of the more fuel-intense vehicle and the second set show the impacts for the relatively fuel-efficient vehicle. In every case, increasing the operating cost of a vehicle has the expected effect; the household responds by decreasing VMT of that vehicle. For multi-vehicle households, an increase in the cost of driving (DPM) of one vehicle introduces an incentive to shift VMT from that vehicle to the other vehicle(s) in the household portfolio.

Table 11: Regression of Log VMT on Fuel Cost Per Mile, 1-Vehicle Households

Outcome		(1)	(2)	(3)	(4)	(5)
$\log(VMT_1)$	DPM_1	-1.272	-1.931	-1.931	-1.918	-1.908
		(0.020)***	(0.025)***	(0.025)***	(0.025)***	(0.026)***
N		7,731,172	7,731,172	7,731,172	7,677,739	7,677,739
N Households		2,083,201	2,083,201	2,083,201	2,071,206	2,071,206
Household FE		Y	Y	Y	Y	Y
County FE		N	N	Y	N	Y
Seasonality FE		N	N	N	Y	Y
Attribute Controls		N	Y	Y	Y	Y

Regression of vehicle log VMT on covariates. Vehicle cost per mile (DPM) instrumented using gasoline price at the time the vehicle was purchased and its interaction with current gasoline prices. All regressions include household fixed effects, nonparametric controls for the age of both vehicles in the portfolio, county-level unemployment, and indicators for leased vehicles. Standard errors clustered by household shown in parentheses. Seasonality fixed effects account for seasonal patterns in driving behavior and consist of the quarter of year of the most recent VMT measurement interacted with counts of each quarter type since the previous VMT measurement. Attribute controls include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

Table 12: Regression of Log VMT on Fuel Cost Per Mile, 3-Vehicle Households

Outcome		(1)	(2)	(3)	(4)	(5)
$\log(VMT_1)$	DPM_1	-0.730	-4.425	-4.415	-4.444	-4.421
		(0.173)***	(0.341)***	(0.340)***	(0.347)***	(0.347)***
	DPM_2	-0.249	2.027	2.015	2.029	2.046
		(0.196)	(0.360)***	(0.360)***	(0.367)***	(0.367)***
	DPM_3	-0.398	1.679	1.674	1.692	1.667
		(0.195)**	(0.408)***	(0.408)***	(0.413)***	(0.414)***
$\log(VMT_2)$	DPM_2	-0.038	-4.179	-4.184	-4.147	-4.165
		(0.243)	(0.427)***	(0.427)***	(0.440)***	(0.440)***
	DPM_1	-0.887	1.271	1.271	1.284	1.298
		(0.178)***	(0.326)***	(0.326)***	(0.333)***	(0.334)***
	DPM_3	-0.161	1.368	1.371	1.253	1.262
		(0.224)	(0.482)***	(0.482)***	(0.493)**	(0.494)**
$\log(VMT_3)$	DPM_3	-1.118	-3.993	-4.004	-3.985	-3.975
		(0.235)***	(0.472)***	(0.472)***	(0.484)***	(0.483)***
	DPM_1	0.392	1.380	1.386	1.422	1.404
		(0.164)**	(0.296)***	(0.296)***	(0.304)***	(0.304)***
	DPM_2	-0.527	0.739	0.736	0.712	0.735
		(0.207)**	(0.368)**	(0.368)**	(0.380)*	(0.379)*
N		468,033	468,033	468,033	455,690	455,690
N Households		165,085	165,085	165,085	161,135	161,135
Household FE		Y	Y	Y	Y	Y
County FE		N	N	Y	N	Y
Seasonality FE		N	N	N	Y	Y
Attribute Controls		N	Y	Y	Y	Y

Regression of vehicle log VMT on covariates. Variables subscripted with 1 to 3 in order of decreasing fuel intensity. Vehicle cost per mile (DPM) instrumented using gasoline price at the time the vehicle was purchased and its interaction with current gasoline prices. All regressions include household fixed effects, nonparametric controls for the age of both vehicles in the portfolio, county-level unemployment, and indicators for leased vehicles. Standard errors clustered by household shown in parentheses. Seasonality fixed effects account for seasonal patterns in driving behavior and consist of the quarter of year of the most recent VMT measurement interacted with counts of each quarter type since the previous VMT measurement. Attribute controls include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

The specification in column (1) includes only operating costs, household fixed effects, and controls for vehicle age and leases. Here an increase in the operating cost of one vehicle appears to decrease the VMT of both vehicles in the portfolio. This may be reasonable in the presence of large income effects, but we suspect that it may be driven by omitted vehicle attributes. Column (2) adds controls for vehicle attributes to the previous specification. Here, all estimated parameters are of larger magnitude and the sign on the cross-effect for both vehicles becomes positive. That is, an increase in the operating cost of vehicle A reduces VMT of vehicle A but increases the VMT of vehicle B. These effects are similar in magnitude across the more and less fuel-intense vehicles in the portfolio, but the difference between the direct and indirect effects are larger for the relatively fuel-efficient vehicle, implying that households shift VMT demand from less to more fuel efficient vehicles in the face of increasing operating costs.

These results are robust to inclusion of alternative controls. The next three columns include additional fixed effects to account for various forms of unobserved heterogeneity. Column (3) adds fixed effects for the county of residence of household h in year t. Column (4) adds indicators for the quarter-of-year of the upcoming odometer reading to the specification in column (2). Column (5) is our preferred specification and includes county fixed effects from column (3) and provides robust controls for seasonality in driving patterns. It does so by using fixed effects for quarter-of-year of the upcoming odometer reading interacted with counts of each quarter-of-year elapsed since the previous odometer reading. This will flexibly control for determinants of VMT such as the number of summer seasons (high VMT) that elapsed between odometer readings. Coefficient estimates are quite similar cross Columns (2) to (5), showing our results are robust to the specific set of included fixed effects.

B Effects for Additional Portfolio Types

Table 10 shows the average effect of decreasing the fuel intensity of a single vehicle in a household's portfolio by 10%. We consider all households holding three or fewer vehicles in their portfolio. The portion of households by starting and ending portfolio sizes are shown in Table 13. Three or fewer vehicle households comprise 86.0% of the households in California during our sample period.

Table 13: Portion of households by portfolio type

Start	End Portfolio Size			
Portfolio Size	1	2	3	4+
1	0.335	0.063	0.009	0.003
2	0.054	0.213	0.039	0.012
3	0.008	0.039	0.100	0.031
4+	0.002	0.007	0.018	0.069

Observed portion of household-year observations by starting and ending portfolio size.

Computing the effect of changes in vehicle fuel economy requires additional assumptions and empirical estimates of household behavior. In cases where households have more than one kept vehicle, the highest value vehicle experiences the decrease in fuel intensity. For cases where households have fewer vehicles in the ending portfolio than their starting portfolio, there are no attribute substitution effects because the household is not purchasing a new vehicle. Attribute substitution effects for other sizes are computed using methods analogous to those for 2x2 portfolios.The marginal effects of kept vehicle fuel intensity on purchased vehicle fuel intensity are shown in Table 14. Estimates for the 1x2 and 2x3 portfolio are very noisy due to small sample size. For the purposes of this simulation, we assume no portfolio effects for the 1x2 and 2x3 portfolios. Finally, there are no effects on fuel consumption for 1x1 household as they are replacing their only vehicle and we assume the vehicle purchase is unconstrained.

Table 14: Marginal effect of kept vehicle GPM on bough vehicle GPM for other portfolio sizes

	2x2 (1)	1x2 (2)	2x3 (3)	3x3 (4)
$P^G = \$2.00$	-1.178	1.976	-6.257	-1.759
	(0.320)***	(4.422)	(23.565)	(0.898)*
$P^G = \$3.00$	-1.362	1.797	-6.551	-1.916
	(0.329)***	(4.420)	(24.152)	(0.864)**
$P^G = \$4.00$	-1.547	1.617	-6.844	-2.073
	(0.339)***	(4.418)	(24.743)	(0.834)**
N Non-singleton	511,243	164,897	29,586	124,529
Starting Port. Size	2	1	2	3
Ending Port. Size	2	2	3	3

Marginal effects of kept vehicle GPM from a regression of the continuous bought vehicle GPM on covariates. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. The endogenous kept vehicle fuel ecnonomy and its interaction with gasoline prices instrumented with the gasoline price at the time the kept vehicle was purchased and its interaction with the current gasoline price.

The effect of portfolio fuel intensity of VMT is estimated on all households with the same ending portfolio size, using the specification presented in Equation A.1. Estimates for two vehicle portfolios are in Table 8 of the paper and one and three-vehicle portfolios are shown in Tables 11 and 12 of the Appendix, respectively.

C Additional Detail of Empirical Estimates

C.1 Reduced form relationship

The reduced form relationship between the the gasoline price instrument and our outcome variable of interest, f^b, is presented in Figures 5a and 5b. Many factors influence a consumer's choice of vehicle attributes, including f^b, so a plot of the raw data reveals little about the underlying relationship between our variables of interest. Instead, we present the variables after partialing out other covariates. The x-axis and y-axis are the residuals retrieved from regressing $p^{gas}_{it_k}$ and f^b, respectively, on covariates.

A clear relationship emerges, indicating a negative correlation between the gasoline price at the time of purchase of the kept car and the GPM of the bought car. In theory, the relationship between $p^{gas}_{it_k}$ and f^b could be positive or negative depending on the correlation between f^k and f^b. This figure provides some of the first suggestive evidence that attributes across cars in the portfolio are negatively correlated (i.e., attribute substitution). When gasoline prices were higher at the time the household purchased the kept vehicle (accordingly leading to lower GPM of the kept vehicle), households prefer to buy a vehicle with higher GPM.

Figure 5: Reduced form relationship: Gas price at time of kept car purchase

(a) New Vehicle Purchases

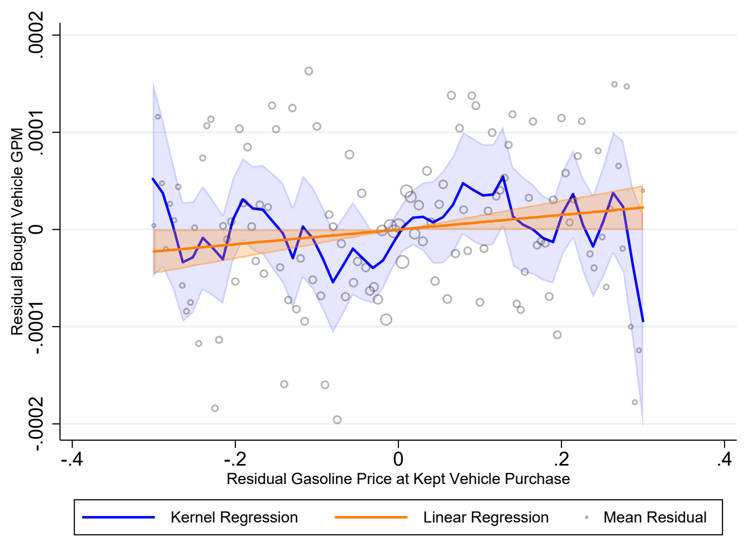

(b) Used Vehicle Purchases

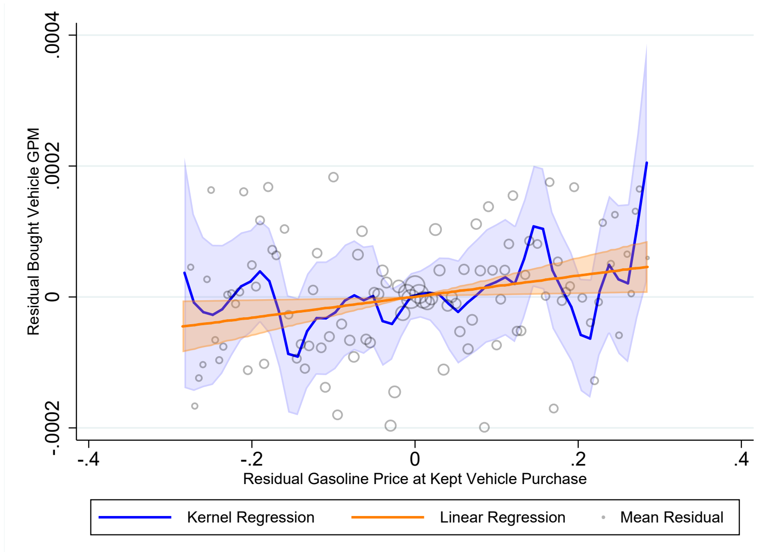

Plot of the reduced form relationship between gasoline price at the time of kept vehicle purchase and the fuel economy (in GPM) of the purchased vehicle. Both variables are partialed of all other regression covariates (Household, vehicle class, manufacturer, age, year, month, and new purchase fixed effects, and continuous controls for the current gasoline price, and the unemployment rate). Graphs are limited to the 5th through 95th percentiles of residual kept vehicle gasoline price. Excludes observations where the household fixed effect perfectly predicts the outcome. Blue line is a kernel regression with Epanechnikov kernel and bandwidth 0.01. The orange line is the linear relationship estimated using OLS. Shaded regions are 95% confidence intervals from these estimates. Gray circles are mean residuals for each 0.005 in kept vehicle GPM. The size of each circle is proportional to the number of observations used to compute the mean residual.

C.2 Heterogeneous Effects by Relative Fuel Intensity

Table 6 presents estimates of our parameter of interest, the marginal effect of kept vehicle GPM on bought vehicle GPM from a variety of specifications. For completeness, Table 15 shows the underlying parameter estimates for the kept vehicle GPM, current gasoline price, the interaction of kept vehicle GPM and current gasoline price and an indicator set to one if the kept vehicle was the most fuel-intense in the portfolio (where applicable).

Table 15: Regression of Kept Vehicle GPM on Purchased Vehicle GPM by Relative Portfolio Position

	No Interaction	Value DiD - Cubic		Value DiD - Spline	
		$f^d \geq f^k$	$f^k > f^d$	$f^d \geq f^k$	$f^k > f^d$
	(1)	(2)	(3)	(4)	(5)
GPM^K	-0.830	-0.419	-0.060	-0.434	-0.042
	(0.300)***	(0.171)**	(0.070)	(0.178)**	(0.071)
P^G	0.008	0.015		0.014	
	(0.001)***	(0.001)***		(0.001)***	
$GPM^K \times P^G$	-0.186	-0.397	0.149	-0.380	0.146
	(0.012)***	(0.026)***	(0.009)***	(0.025)***	(0.009)***
$\mathbf{1}\left[GPM^K > GPM^D\right]$		-0.004		-0.003	
		(0.003)		(0.003)	
N Non-Singleton	508,407	303,772		303,772	
Kleibergen-Paap rk F	17.77	31.25		35.19	
Instruments:					
p^k_{gas}	Y	Y		Y	
$\Delta\Delta P^{kd}$	N	Cubic		Linear Spline	

Parameter estimates from a regression of continuous bought vehicle GPM. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. Columns 2/3 and 4/5 each show results from a single regression. $f^d \geq f^k$ ($f^k > f^d$) show marginal effects when the dropped vehicle was the most (least) fuel intense vehicle in the portfolio prior to the purchase. Columns 2 through 5 instrument for endogenous regressors using gas price at the time of kept vehicle purchase, gas price at the time of dropped vehicle purchase, the either a cubic (Columns 2 and 3) or a linear spline with knots at ± 500 (Columns 4 and 5) of "Price deviations from trend" instruments, and projections from the space of exogenous variables described in Section 3.2 as instruments for endogenous regressors.

D Additional Tests of Robustness

The empirical results described in Section 4 are robust to a range of alternative specifications. Table 16 presents an object of interest, the marginal effect of the kept vehicle GPM on purchased vehicle GPM for each specification. For comparison, Column (1) repeats the primary specification using household fixed effects and the gasoline price plus deviation from trend instruments. Column (2) limits the sample to households whose dropped vehicle is at least three years old at the time of the new vehicle purchase. Column (3) includes the price paid for the dropped vehicle when it was purchased as an additional control. In each case, the marginal effects in each alternative specification are broadly similar to those in the primary specification.

Table 16: Marginal effect of gasoline price, Base Model

	Pref. Spec (1)	$\geq$ 3 yr (2)	Dropped Price (3)
$P^G = \$2.00$	-1.202	-1.203	-1.592
	(0.319)***	(0.324)***	(0.460)***
$P^G = \$3.00$	-1.389	-1.410	-1.792
	(0.329)***	(0.335)***	(0.474)***
$P^G = \$4.00$	-1.575	-1.617	-1.992
	(0.339)***	(0.347)***	(0.488)***
N Non-singleton	508,407	457,431	415,700

Marginal effect of the kept vehicle GPM on bought vehicle GPM. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. "$\geq$ 3 yr" excludes observations where the dropped vehicle has been held by the household for less than 3 years at the time of new vehicle purchase. "Dropped Price" includes the price of the dropped vehicle at the time of purchase, as reported to the DMV, as a covariate.

Table 17: Marginal effect of operating costs

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
DPM^K	0.086	-2.691	-0.295	-1.407
	(0.002)***	(0.538)***	(0.006)***	(0.531)***
N Non-singleton	1,171,976	1,169,006	509,664	508,407
Kleibergen-Paap rk F	.	35.41	.	13.30

Marginal effect of the kept vehicle operating costs on bought vehicle operating costs in dollars per mile. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively.

Next, we consider an alternative specification where, as opposed to fuel intensity, operating costs are considered the outcome and independent variable. Table 17 shows the marginal effect of kept vehicle operating costs in dollars per mile ($DPM^K = P^G \times f^K$) on the bought vehicle operating cost ($DPM^B = P^G \times f^B$). This specification combines two distinct sources of variation in operating costs, fuel intensity and gasoline prices, but leads to a simpler system containing only a single endogenous variable (DPM^K). In specification using instrumental variables, we instrument for DPM^K using the the gasoline price from the time of kept vehicle purchase time the current gasoline price. The estimated marginal effect of operating costs is similar to the marginal effect of kept vehicle fuel intensity from the primary specification.

As an alternative to household fixed effects, we proxy for household-level heterogeneity using an estimate of household-level VMT demand. This approach offers potential advantages over using household fixed effects. Since we observe vehicle VMT every two years (for vehicles that are required to undergo smog checks) we can more precisely estimate household VMT demand than

Table 18: Marginal effect of kept vehicle GPM

	HHFEIV (1)	OLS (2)	IV (3)	OLS^3 (4)	IV^3 (5)
$P^G = \$2.00$	-1.202	0.027	-1.282	0.027	-1.286
	(0.319)***	(0.004)***	(0.352)***	(0.004)***	(0.356)***
$P^G = \$3.00$	-1.389	0.011	-1.294	0.012	-1.298
	(0.329)***	(0.006)**	(0.351)***	(0.006)**	(0.354)***
$P^G = \$4.00$	-1.575	-0.004	-1.306	-0.003	-1.309
	(0.339)***	(0.008)	(0.349)***	(0.008)	(0.353)***
Kleibergen-Paap rk F	17.77	.	19.77	.	19.39
HH Control	FE	VMT	VMT	VMT^3	VMT^3

Marginal effect of the kept vehicle GPM on bought vehicle GPM in dollars per mile. Standard errors robust to arbitrary heteroskedasticity clustered by household shown in parentheses. *,**,*** denote significance at the 10%, 5%, and 1% levels, respectively. Column 1 repeats the preferred specification with household fixed effects. Columns 2 to 5 replace household fixed effects with predicted household VMT demand. Column 2 estimates the specification using OLS (treating kept vehicle GPM as exogenous). Column 3 instruments using the same IV strategy as the HHFEIV specification. Columns 3 and 4 repeat these OLS and IV estimators controlling for household-level heterogeneity using cubic orthogonal polynomials of predicted VMT.

vehicle purchase behaviors. If VMT demand is the primary driver of household preferences for the level of fuel intensity across the portfolio, then including predicted VMT demand as a covariate would obviate the need for household fixed effects.

Our analysis proceeds in two steps. First, we estimate a model of household VMT demand conditional on the fuel intensity of each vehicle in the portfolio, household, and time fixed effects. We instrument for fuel intensity using the gasoline price at the time each vehicle was purchased. This model is then used to predict a VMT demand for each household. We then repeat our primary specification including predicted VMT (or powers of VMT) as a control. The results are presented in Table 18.

The OLS specifications fail to account for the endogenous kept vehicle GPM and much like other specifications, estimate marginal effects close to zero. The IV specifications, however, estimate marginal effects similar to the HHFEIV specification, suggesting that household VMT demand is a major contributing factor (or proxy for) household preferences for the level of portfolio fuel intensity.

D.1 Weak Identification Tests

We subject our primary estimates to an array of tests for weak and under identification of the first stage regressions. The results of these tests are shown in Table 20. The Anderson-Rubin test rejects the null that the endogenous parameter vector is jointly equal to zero. Unlike a typical Wald test, the Anderson-Rubin test has no power when the instruments are weak and thus, we can conclude

Table 19: Instrumental Variable Diagnostics - HHFEIV Specification

		All (1)	New (2)	Used (3)
Weak Identification				
Anderson-Rubin F		341.62	58.25	153.29
		[0.0000]	[0.0000]	[0.0000]
Cragg-Donald F		29.43	5.33	43.97
Kleibergen-Paap rk F		17.77	3.14	24.67
Sanderson-Windmeijer Partial F	GPM^K	35.55	6.28	49.40
		[0.0000]	[0.0122]	[0.0000]
	$GPM^K \times P^G$	36.09	6.30	63.34
		[0.0000]	[0.0121]	[0.0000]
Under Identification				
Kleinbergen-Paap rk LM		35.53	6.27	49.29
		[0.0000]	[0.0123]	[0.0000]
Sanderson-Windmeijer χ^2	GPM^K	35.56	6.28	49.41
		[0.0000]	[0.0122]	[0.0000]
	$GPM^K \times P^G$	36.10	6.30	63.36
		[0.0000]	[0.0121]	[0.0000]

Statistics from common instrumental variables diagnostics. Hypothesis test p-values shown in square brackets, where appropriate. The Anderson-Rubin test is a weak-instrument-robust test that the coefficients on endogenous variables are jointly equal to zero. The Cragg-Donald F and Kleinbergen-Paap rk F are statistics measuring the power of instruments in the first stage regressions. The Sanderson-Windmeijer Partial F tests the null of weak instruments in each first stage regression, conditional on variation in the other first stage regressions. The Kleinbergen-Paap rk LM tests the null of an underidentified first stage and the Sanderson-Windmeijer χ^2 test each first stage regression for underidentification, conditional on variation in the other first stage equations.

the endogenous variables are not equal to zero, even if the instruments are weak.

The Cragg-Donald F and Kleibergen-Paap rk F are measures of the power of the excluded instruments in the system of first stage equations. Larger values indicate a more powerful first stage and less potential for weak instrument bias. There is little theoretical guidance for the appropriate level of these statistics, but both the full sample and used vehicle sample are much larger than the Stock and Yogo (2002) critical values for low bias.

The Sanderson-Windmeijer Partial F measure the first stage power of each first stage equation, conditional on variation in the other equations. In each case, we reject the null that the first stage equations are weak.

Finally the Kleibergen-Paap rk LM and Sanderson-Windmeijer χ^2 tests measure the under identification – that the rank condition for the system of first stage equations holds. In each case, we reject the null hypothesis that the first stage is under identified.

D.2 Overidentification Tests

Our primary instrumental variables specification is just-identified, instrumenting for the kept vehicle GPM and its interaction with gasoline price using the gasoline price at the time of kept vehicle purchase and its interaction with gasoline price. As described in Section 3.2, estimating separate effects for cases when the kept vehicle is more or less fuel intense than the dropped vehicle requires additional instruments. We can also use these additional instruments as a test of overidentifying restrictions in our primary specification. In Table 20, we reestimate our preferred HHFEIV specification including additional instruments in each column. For clarity, only the marginal effect of kept vehicle GPM is shown. Adding instruments strengthens (in the case of the Value DiD instrument) or weakens (gasoline price at the time of dropped vehicle purchase) the first stage. However, the estimated marginal effects are broadly similar across specification, with the possible exception at a gasoline price of $4/gal, which has poor support in the data.

Importantly, in each case, Hansen's J test fails to reject the null of a violation of the exclusion restriction. This suggests none of the suggested instruments violate the exclusion restriction.

Table 20: Overidentification and Weak Instrument Tests

	Base	Sold GP	Value DiD	All
	(1)	(2)	(3)	(4)
Estimated Marginal Effects				
$P^G = \$2.00$	-1.202	-1.384	-1.264	-1.251
	(0.3193)***	(0.3765)***	(0.1614)***	(0.1534)***
$P^G = \$3.00$	-1.389	-1.232	-1.160	-1.149
	(0.3291)***	(0.3188)***	(0.1832)***	(0.1695)***
$P^G = \$4.00$	-1.575	-1.080	-1.056	-1.046
	(0.3390)***	(0.5414)**	(0.2365)***	(0.2187)***
Overidentifying Restrictions				
Hansen's J Statistic		0.354	2.449	6.274
		(0.8377)	(0.2940)	(0.1796)
Weak Identification				
Anderson-Rubin F	341.621	6.552	20.400	15.704
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Cragg-Donald F	29.433	9.897	84.191	62.000
Kleibergen-Paap rk F	17.773	4.873	28.468	21.197
Instruments:				
Gas price at time of Kept Purchase	Y	Y	Y	Y
Gas price as time of Dropped Purchase	N	Y	N	Y
Difference in Unanticipated Change in Value	N	N	Y	Y

Estimated marginal effect of kept vehicle fuel intensity (in GPM) on bought vehicle fuel intensity across a range of gasoline prices for over-identified IV regressions. For marginal effects, standard errors clustered by household shown in parentheses. Over-identification and weak instrument tests show the p-value in parentheses where appropriate. The first column repeats the household fixed effects plus IV specification from Table 5 using the gasoline price at the time the kept vehicle was purchased and its interaction with the current gasoline price as instruments for the two endogenous regressors. Column 2 adds the gasoline price at the time the dropped vehicle was purchased and its interaction with the current gasoline price as an instrument. Column 3 adds the difference in unanticipated value change instrument from Section 3.2 and its interaction with the current gasoline price. Finally, column 4 includes all aforementioned instruments.

D.3 Lasso-based instrument selection

The IV estimates presented in Tables 4 and 5 use economic reasoning and an examination of the first stage relationships to choose a first stage specification. As an alternative, one could define a pool of candidate instruments and use a lasso-based selection algorithm described in Belloni et al.

(2012) to choose first stage specifications.

We apply this method with caution. Importantly, Belloni et al. demonstrate estimating instrumental variables regressions leads to asymptotically consistent estimates when the set of candidate instruments is sparse – the true parameter matrix of the first stage on the candidate instrument pool contains mostly zeros. We are not convinced this is the case in our setting. The true first stage relationship may depend on many high order interactions between our instruments. Our preferred specification limits the number of interactions to minimize bias from including many instruments.

In Table 21 we show the estimated marginal effect of the kept vehicle fuel intensity on the bought vehicle fuel intensity, using instruments described in Section 3.2. In the first row, we repeat the specification from Table 5. In subsequent rows we show post-lasso IV estimates using 1-way to 4-way interactions between the instruments and a single interaction with the current gasoline price as the candidate instrument pool. In general, the first stage is substantially weaker using the post-lasso IV method is weaker and the point estimates are biased in the direction of the OLS estimates, consistent with a weaker first stage.

Table 22 shows similar estimates, replacing the cubic in the Value DiD instrument with a spline with knots at $\pm 500, \pm 1000, \pm 1500, \pm 2000$. Again the results are quite similar to the primary specification but with a weaker first stage.

Table 21: Comparison of selected IV specification with post-lasso IV

IV Spec	Subpopulation	$P^G = 2$	$P^G = 3$	$P^G = 4$
Primary Specification	$\mathbb{1}^{d \geq k}$	-1.162***	-1.570***	-1.978***
		(0.161)	(0.170)	(0.182)
$F^{rk} = 20.36$	$\mathbb{1}^{k>d}$	-0.957***	-1.217***	-1.477***
		(0.105)	(0.115)	(0.128)
1-way Interactions	$\mathbb{1}^{d \geq k}$	-1.259**	-1.149*	-1.040
		(0.525)	(0.631)	(0.845)
$F^{rk} = 3.06$	$\mathbb{1}^{k>d}$	-0.985***	-0.873***	-0.760*
		(0.127)	(0.244)	(0.431)
2-way Interactions	$\mathbb{1}^{d \geq k}$	-1.239***	-1.253**	-1.266*
		(0.451)	(0.506)	(0.673)
$F^{rk} = 3.62$	$\mathbb{1}^{k>d}$	-0.959***	-0.926***	-0.893**
		(0.121)	(0.228)	(0.393)
3-way Interactions	$\mathbb{1}^{d \geq k}$	-1.230***	-1.223**	-1.217*
		(0.455)	(0.504)	(0.659)
$F^{rk} = 3.27$	$\mathbb{1}^{k>d}$	-0.964***	-0.921***	-0.877**
		(0.119)	(0.218)	(0.375)
4-way Interactions	$\mathbb{1}^{d \geq k}$	-1.203***	-1.259**	-1.315**
		(0.455)	(0.495)	(0.636)
$F^{rk} = 2.90$	$\mathbb{1}^{k>d}$	-0.963***	-0.962***	-0.961***
		(0.116)	(0.206)	(0.356)

Estimated marginal effect of kept vehicle fuel intensity (in GPM) on bought vehicle fuel intensity across a range of gasoline prices. Marginal effects are allow to vary across the endogenous decision to drop the more ($\mathbb{1}^{d \geq k}$) or less ($\mathbb{1}^{k>d}$) fuel intense vehicle. Regressions include five endogenous covariates. Standard errors shown in parentheses. The first section repeats Table 5 and shows instruments selected using economic reasoning and an examination of the first stage relationships. Each subsequent section uses the post-lasso first stage selection procedure described in Belloni et al. (2012), choosing instruments from the set of row wise Kronecker products of candidate instruments described in Section 3.2 and the current gasoline price.

Table 22: Comparison of selected IV with post-lasso IV using splines

IV Spec	Subpopulation	$P^G = 2$	$P^G = 3$	$P^G = 4$
Primary Specification	$\mathbb{1}^{d \geq k}$	-1.162***	-1.570***	-1.978***
		(0.161)	(0.170)	(0.182)
$F^{rk} = 20.36$	$\mathbb{1}^{k>d}$	-0.957***	-1.217***	-1.477***
		(0.105)	(0.115)	(0.128)
1-way Interactions	$\mathbb{1}^{d \geq k}$	-1.007*	-1.316**	-1.624**
		(0.530)	(0.588)	(0.665)
$F^{rk} = 2.88$	$\mathbb{1}^{k>d}$	-0.966***	-1.139***	-1.311***
		(0.129)	(0.166)	(0.222)
2-way Interactions	$\mathbb{1}^{d \geq k}$	-1.180**	-1.456**	-1.731**
		(0.562)	(0.638)	(0.735)
$F^{rk} = 2.59$	$\mathbb{1}^{k>d}$	-0.977***	-1.128***	-1.278***
		(0.129)	(0.173)	(0.239)
3-way Interactions	$\mathbb{1}^{d \geq k}$	-0.961*	-1.256**	-1.551**
		(0.496)	(0.550)	(0.624)
$F^{rk} = 2.70$	$\mathbb{1}^{k>d}$	-0.951***	-1.119***	-1.287***
		(0.121)	(0.154)	(0.208)
4-way Interactions	$\mathbb{1}^{d \geq k}$	-0.843*	-1.107**	-1.370**
		(0.468)	(0.511)	(0.574)
$F^{rk} = 2.93$	$\mathbb{1}^{k>d}$	-0.939***	-1.093***	-1.247***
		(0.118)	(0.148)	(0.200)

Estimated marginal effect of kept vehicle fuel intensity (in GPM) on bought vehicle fuel intensity across a range of gasoline prices. Marginal effects are allow to vary across the endogenous decision to drop the more ($\mathbb{1}^{d \geq k}$) or less ($\mathbb{1}^{k>d}$) fuel intense vehicle. Regressions include five endogenous covariates. Standard errors shown in parentheses. The first section repeats Table 5 and shows instruments selected using economic reasoning and an examination of the first stage relationships. Each subsequent section uses the post-lasso first stage selection procedure described in Belloni et al. (2012), choosing instruments from the set of row wise Kronecker products of candidate instruments described in Section 3.2, replacing the "Price deviations from trend" instrument with a linear spline with knots at $\pm 500, \pm 1000, \pm 1500, \pm 2000$ and the current gasoline price.

D.4 Sample Restrictions

The primary estimation strategy relies on household fixed effects to control for household-level idiosyncratic preferences over fuel economy in the purchased vehicle. Parameters in this model are identified by households engaging in at least two transactions during the sample period and couple potentially lead to a selected sample. To investigate this concern, we repeat our estimation strategy, limiting to the sample of households with multiple transactions during the sample period.

The top panel of Table 23 repeats estimates from Table 5, which rely on all available observations. Panel (b) restimates the OLS, IV, and HHFE models using only observations from the HHFEIV estimator with identified fixed effects. The estimates are broadly similar across the sample restriction, suggesting selection is not a major driver of the differences between the OLS and HHFEIV estimates.

E Regressions Restricted to the HHFEIV Sample

Table 23: Linear Model of Purchased Vehicle Fuel Intensity

(a) Full Sample

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
GPM^K	0.091	-2.551	-0.115	-0.830
	(0.004)***	(0.509)***	(0.013)***	(0.300)***
$GPM^K \times P^G$	-0.004	0.000	-0.139	-0.186
	(0.002)**	(0.004)	(0.005)***	(0.012)***
N Non-singleton	1,171,976	1,169,006	509,664	508,407
Kleibergen-Paap rk F	.	18.72	.	17.77

(b) Household fixed effects sample

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
GPM^K	0.077	-2.592	-0.116	-0.830
	(0.008)***	(1.297)**	(0.013)***	(0.300)***
$GPM^K \times P^G$	0.001	-0.008	-0.138	-0.186
	(0.003)	(0.011)	(0.005)***	(0.012)***
N Non-singleton	508,407	508,407	508,407	508,407
Kleibergen-Paap rk F	.	2.84	.	17.77

Regression of bought vehicle GPM on covariates. Standard errors robust to heteroskedasticity and clustered at the household level shown in parentheses. Panel (a) repeats specifications show in Table 5. Panel (b) repeats each specification, limiting to households making at least two transactions during the sample period.

E.1 Alternate Functional Form

As an alternative to the linear relationship between gasoline prices, kept vehicle GPM and the bought vehicle GPM, we reestimate the OLS, IV, HHFE, and HHFEIV models under a log-log relationship. In these regressions, we replace gasoline prices and vehicle GPM with their logs. In this log-log specification, the log interaction between current gasoline prices and kept vehicle GPM is collinear by the log of gasoline price and the log of kept vehicle GPM. Consequently, there is only one parameter of interest, the coefficient on the log of kept vehicle GPM. We instrument for this endogenous variable using only the log gasoline price at the time the kept vehicle was purchased.

Table 24 show the estimated marginal effects from the linear models in Panel (a) and the estimated kept-bought vehicle GPM elasticity in Panel (b). Since the kept and bought vehicle GPM are of similar scale, the estimated elasticities should be of similar scale to the marginal effects from the linear specification. Comparing Panels (a) and (b) there is little difference between estimated effects using a linear versus a log-log specification.

Table 24: Linear Model of Purchased Vehicle Fuel Intensity

(a) Linear Relationship

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
$P^G = \$2.00$	0.082	-2.551	-0.392	-1.202
	(0.002)***	(0.509)***	(0.007)***	(0.319)***
$P^G = \$3.00$	0.077	-2.551	-0.531	-1.389
	(0.002)***	(0.509)***	(0.008)***	(0.329)***
$P^G = \$4.00$	0.073	-2.551	-0.669	-1.575
	(0.004)***	(0.509)***	(0.012)***	(0.339)***

(b) Log-Log Relationship

All	OLS (1)	IV (2)	HHFE (3)	HHFEIV (4)
$\log(GPM^K)$	0.087	-1.875	-0.402	-1.291
	(0.002)***	(0.421)***	(0.007)***	(0.397)***
N Non-singleton	1,171,976	1,169,006	509,664	508,407
Kleibergen-Paap rk F	.	37.86	.	24.62

Regression of bought vehicle GPM on covariates. Standard errors robust to heteroskedasticity and clustered at the household level shown in parentheses. Panel (a) repeats marginal effects from linear specifications shown in Table 5. Panel (b) replaces gasoline prices and vehicle GPM with their logs.

Table 25: Log-log Regression of VMT on Fuel Intensity and Fuel Cost, 2-Vehicle Households

Outcome		(1)	(2)	(3)	(4)
$\log(VMT_1)$	$\log(P^G)$	-0.317	-0.321	-0.321	-0.324
		(0.007)***	(0.007)***	(0.007)***	(0.007)***
	$\log(GPM_1)$	-5.788	-5.821	-5.852	-5.937
		(0.216)***	(0.217)***	(0.216)***	(0.221)***
	$\log(GPM_2)$	3.882	3.871	3.922	3.939
		(0.147)***	(0.148)***	(0.147)***	(0.150)***
$\log(VMT_2)$	$\log(P^G)$	-0.253	-0.256	-0.260	-0.262
		(0.006)***	(0.006)***	(0.006)***	(0.006)***
	$\log(GPM_2)$	-3.387	-3.401	-3.521	-3.556
		(0.143)***	(0.143)***	(0.144)***	(0.146)***
	$\log(GPM_1)$	2.124	2.104	2.154	2.134
		(0.167)***	(0.167)***	(0.169)***	(0.171)***
N		3,015,388	3,015,388	2,946,343	2,946,343
N Households		877,978	877,978	858,607	858,607
Household FE		Y	Y	Y	Y
County FE		N	Y	N	Y
Seasonality FE		N	N	Y	Y
Attribute Controls		Y	Y	Y	Y

Log-log Regression of VMT on fuel intensity of each vehicle in the portfolio and current gasoline price (P^G) covariates. Log vehicle fuel intensity (GPM) instrumented using gasoline price at the time the vehicle was purchased. All regressions include household fixed effects, nonparametric controls for the age of both vehicles in the portfolio, county-level unemployment, and indicators for leased vehicles. Standard errors clustered by household shown in parentheses. Seasonality fixed effects account for seasonal patterns in driving behavior and consist of the quarter of year of the most recent VMT measurement interacted with counts of each quarter type since the previous VMT measurement. Attribute controls include indicators for vehicle class and continuous measures of vehicle curb weight, wheelbase, vehicle width, and engine displacement.

F Additional Summary Statistics

The following appendix provides additional summary statistics describing.

F.1 Vehicle Class Transitions in Purchase Decisions

Tables 26 and 27 show the probability a household purchases a vehicle of the class shown in the column header when they drop (Table 26) or keep (Table 27) a vehicle shown in the row header. These summary statistics show households persistence in their preferences for vehicle class. Households dropping an SUV are much more likley to replace it with an SUV, regardless of the class of the kept vehicle.

Table 26: Vehicle Class Purchase Decisions

(a) Kept vehicle class: Car

Dropped Vehicle Class	Bought Vehicle Class SUV	Crossover	Car	Luxury	Truck
SUV	0.316	0.108	0.289	0.097	0.190
Crossover	0.112	0.370	0.288	0.097	0.134
Car	0.093	0.070	0.571	0.108	0.158
Luxury	0.099	0.057	0.391	0.272	0.181
Truck	0.138	0.066	0.300	0.072	0.424

(b) Kept vehicle class: Truck

Dropped Vehicle Class	Bought Vehicle Class SUV	Crossover	Car	Luxury	Truck
SUV	0.351	0.091	0.281	0.098	0.179
Crossover	0.145	0.350	0.278	0.105	0.122
Car	0.113	0.076	0.546	0.105	0.159
Luxury	0.115	0.061	0.361	0.296	0.166
Truck	0.158	0.053	0.327	0.081	0.381

(c) Kept vehicle class: SUV

Dropped Vehicle Class	Bought Vehicle Class SUV	Crossover	Car	Luxury	Truck
SUV	0.302	0.079	0.271	0.159	0.189
Crossover	0.155	0.306	0.245	0.166	0.128
Car	0.120	0.063	0.483	0.165	0.168
Luxury	0.126	0.064	0.272	0.396	0.142
Truck	0.147	0.045	0.287	0.102	0.419

Each cell shows the portion of households purchasing a vehicle of the class in the column header when they drop a vehicle in the row header. Table (a) shows households where the kept vehicle class was a car, Table (b) shows households keeping trucks, and Table (c) shows households keeping an SUV.

Table 27: Vehicle Class Purchase Decisions

(a) Dropped vehicle class: Car

Kept	Bought Vehicle Class				
Vehicle Class	SUV	Crossover	Car	Luxury	Truck
SUV	0.120	0.063	0.483	0.165	0.168
Crossover	0.074	0.088	0.529	0.171	0.137
Car	0.093	0.070	0.571	0.108	0.158
Luxury	0.102	0.081	0.497	0.190	0.131
Truck	0.113	0.076	0.546	0.105	0.159

(b) Dropped vehicle class: Truck

Kept	Bought Vehicle Class				
Vehicle Class	SUV	Crossover	Car	Luxury	Truck
SUV	0.147	0.045	0.287	0.102	0.419
Crossover	0.111	0.078	0.288	0.088	0.435
Car	0.138	0.066	0.300	0.072	0.424
Luxury	0.151	0.074	0.234	0.123	0.417
Truck	0.158	0.053	0.327	0.081	0.381

(c) Dropped vehicle class: SUV

Kept	Bought Vehicle Class				
Vehicle Class	SUV	Crossover	Car	Luxury	Truck
SUV	0.302	0.079	0.271	0.159	0.189
Crossover	0.261	0.135	0.273	0.152	0.180
Car	0.316	0.108	0.289	0.097	0.190
Luxury	0.414	0.142	0.175	0.144	0.125
Truck	0.351	0.091	0.281	0.098	0.179

Each cell shows the portion of households purchasing a vehicle of the class in the column header when they keep a vehicle in the row header. Table (a) shows households where the dropped vehicle class was a car, Table (b) shows households dropping trucks, and Table (c) shows households dropping an SUV.

F.2 Fuel Intensity Transitions in Purchase Decisions

Investigating the quartile of fuel intensity shows a similar pattern. Tables 28 and 29 show the portion of households purchasing a vehicle in the GPM quartile show in the column header when they drop (Table 28) or keep (Table 29) a vehicle in the GPM quartile shown in the row header. Again, households appear to have persistent preferences for portfolio types. Households dropping a vehicle in a given GPM quartile a more likely to purchase a replacement in that same quartile. Further, the correlation is the largest when the kept vehicle is also in that same GPM quartile.

Table 28: Vehicle Fuel Intensity Purchase Decisions

(a) Kept vehicle GPM Quartile: 1

Dropped Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.438	0.258	0.167	0.137
Q2	0.298	0.330	0.197	0.174
Q3	0.250	0.260	0.278	0.212
Q4	0.208	0.208	0.214	0.370

(b) Kept vehicle GPM Quartile: 2

Dropped Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.384	0.278	0.183	0.154
Q2	0.235	0.369	0.217	0.179
Q3	0.191	0.267	0.322	0.220
Q4	0.154	0.196	0.228	0.423

(c) Kept vehicle GPM Quartile: 3

Dropped Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.370	0.283	0.186	0.161
Q2	0.228	0.357	0.227	0.188
Q3	0.178	0.266	0.332	0.224
Q4	0.147	0.193	0.230	0.430

(d) Kept vehicle GPM Quartile: 4

Dropped Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.341	0.280	0.187	0.193
Q2	0.215	0.342	0.222	0.221
Q3	0.181	0.253	0.306	0.260
Q4	0.163	0.193	0.214	0.431

Each cell shows the portion of households purchasing a vehicle in the GPM quartile shown in the column header when they drop a vehicle with the GPM quartile shown in the row header. Table (a),(b),(c), and (d) shows households keeping vehicles in the 1st, 2nd, 3rd, and 4th GPM quartiles, respectively.

Table 29: Vehicle Fuel Intensity Purchase Decisions

(a) Dropped vehicle GPM Quartile: 1

Kept Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.438	0.258	0.167	0.137
Q2	0.384	0.278	0.183	0.154
Q3	0.370	0.283	0.186	0.161
Q4	0.341	0.280	0.187	0.193

(b) Dropped vehicle GPM Quartile: 2

Kept Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.298	0.330	0.197	0.174
Q2	0.235	0.369	0.217	0.179
Q3	0.228	0.357	0.227	0.188
Q4	0.215	0.342	0.222	0.221

(c) Dropped vehicle GPM Quartile: 3

Kept Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.250	0.260	0.278	0.212
Q2	0.191	0.267	0.322	0.220
Q3	0.178	0.266	0.332	0.224
Q4	0.181	0.253	0.306	0.260

(d) Dropped vehicle GPM Quartile: 4

Kept Vehicle GPM Quartile	Bought Vehicle GPM Quartile Q1	Q2	Q3	Q4
Q1	0.208	0.208	0.214	0.370
Q2	0.154	0.196	0.228	0.423
Q3	0.147	0.193	0.230	0.430
Q4	0.163	0.193	0.214	0.431

Each cell shows the portion of households purchasing a vehicle in the GPM quartile shown in the column header when they keep a vehicle with the GPM quartile shown in the row header. Table (a),(b),(c), and (d) shows households dropping vehicles in the 1st, 2nd, 3rd, and 4th GPM quartiles, respectively.

F.3 Vehicle Replacement Cycles

The following histograms summarize the time elapsed between purchases of vehicles in each household in the 2x2 replacement sample's portfolios. Figure 6 shows the distribution of time elapsed between the purchase of the "bought" vehicle and the vehicle it replaced. It is important to note that households replacing the same vehicle twice in the same year are systematically eliminated from our sample as they are very likely to have overlapping active registrations of four vehicles, in which case we are unable to determine if they were, for some period of time, an actual three car household. The median replacement cycle is 2.33 years and is slightly shorter for used vehicles on average. However, the mean difference between households buying new or used vehicles is only 0.06 years. Looking at the distribution, this is driven mass for households purchasing a new vehicle being concentrated around the 3-year mark, consistent with typical leasing cycles.

Figure 7 considers the time between the purchase of the "bought" vehicle and the date of purchase of the vehicle the household"kept" through that purchase. Here systematic attrition of households from the sample is not an issue, but still very few households replace both vehicles in their portfolio in the same year. The median time between the purchase of the bought and kept vehicle is 2.25 years (for both households purchasing new and used vehicles.) The duration has a longer right tail when the replacement vehicle is used. However, the vast majority of households have a kept vehicle less than 6 years when replacing the other vehicle.

Figure 6: Duration from dropped vehicle purchase to bought vehicle purchase

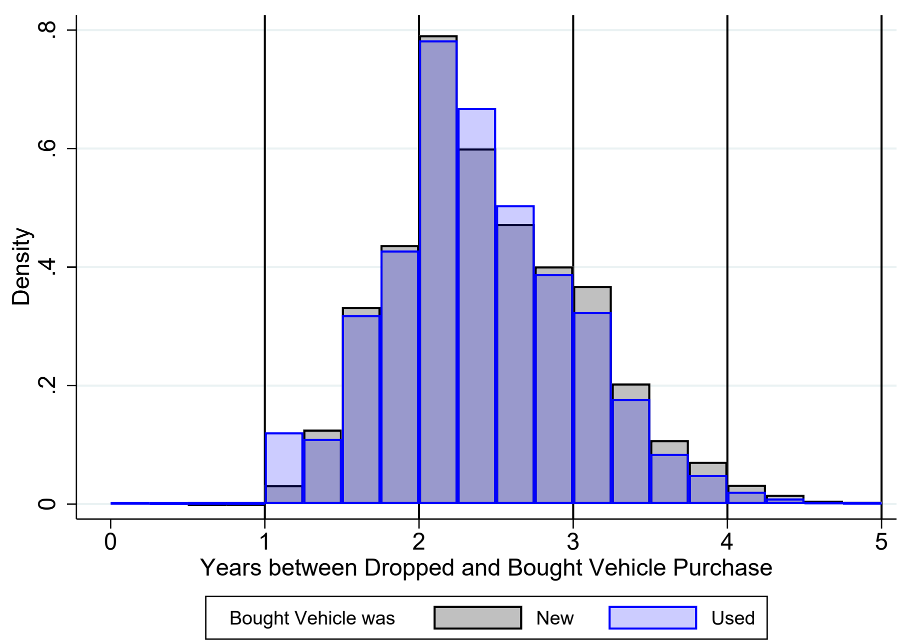

Histogram of the time elapsed between a household's purchase of a vehicle and the year in which the vehicle it replaced was purchased. Households purchasing a new vehicle are shown in gray and a used vehicle in blue. Households limited to the 2x2 replacement sample.

Figure 7: Duration from kept vehicle purchase to bought vehicle purchase

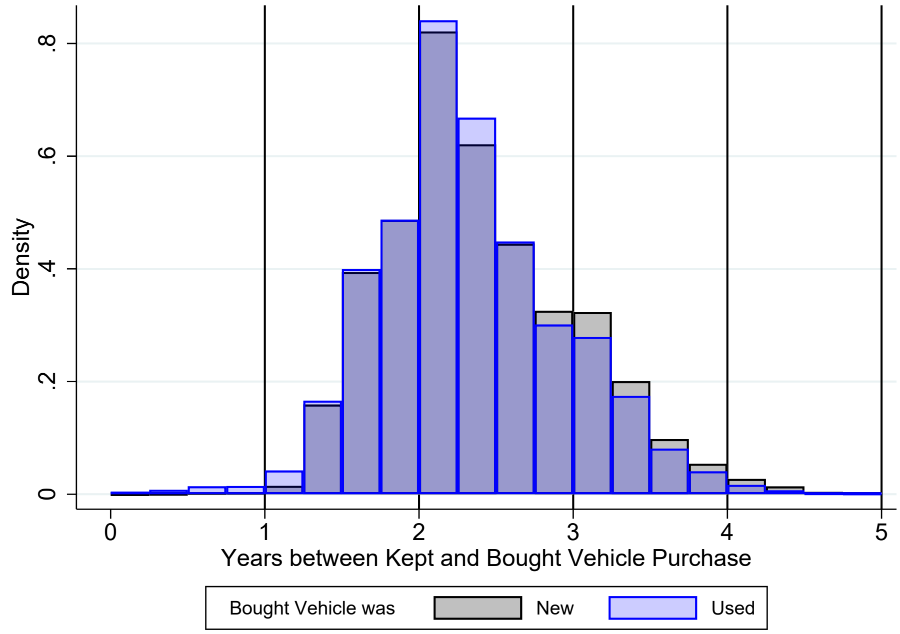

Histogram of the time elapsed between a household's purchase of a vehicle and the year in which the kept vehicle was purchased. Households purchasing a new vehicle are shown in gray and a used vehicle in blue. Households limited to the 2x2 replacement sample.

F.4 Residual Variation in Gasoline Price

The primary regression specification uses gasoline price at the time the kept vehicle was purchased as an instrument for its endogenous fuel intensity. Figure 8 shows histograms and the associated kernel density plot using the Epanechnikov kernel for the gasoline price at the time of kept vehicle purchase partialed of all other covariates and fixed effects in the primary regression specification.

Figure 8: Duration from kept vehicle purchase to bought vehicle purchase

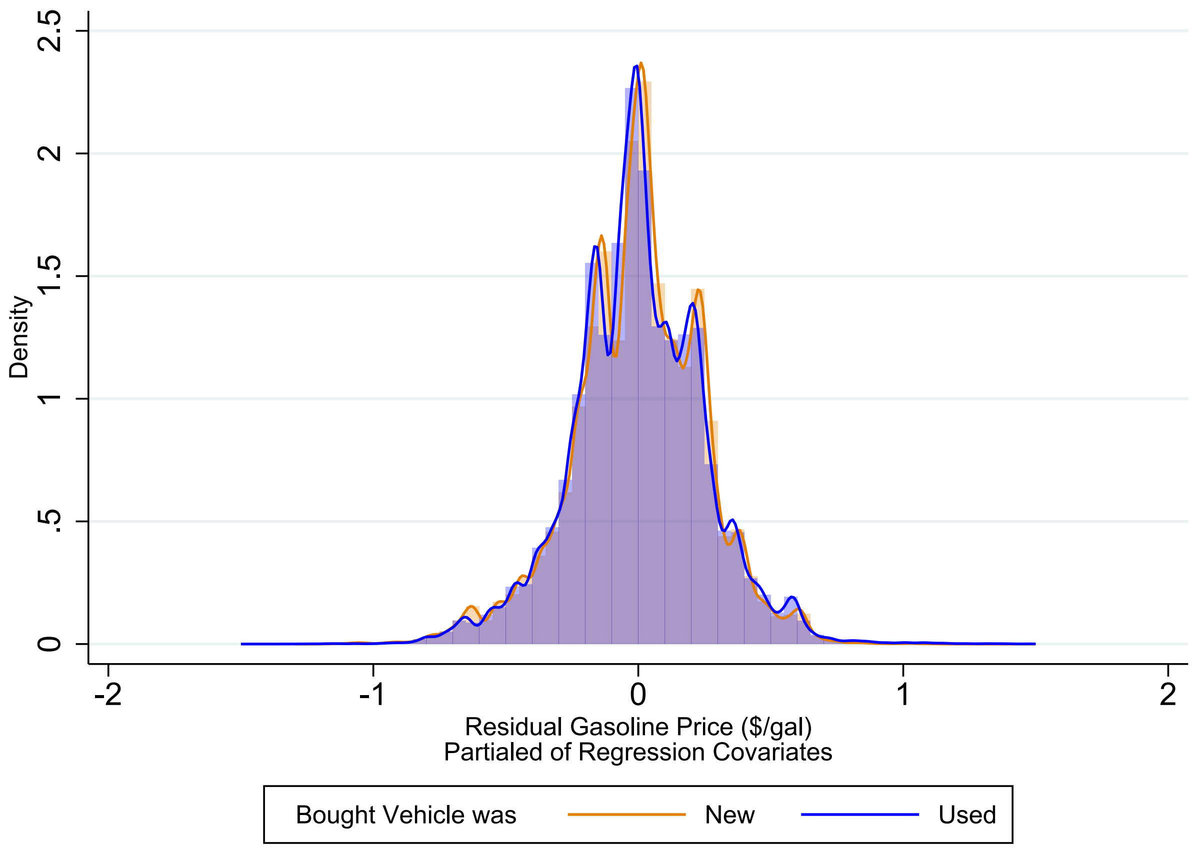

Histogram and kernel density of the gasoline price at the time of kept vehicle purchase partialed of all covariates and fixed effects included in the primary regression specification. Households purchasing a new vehicle are shown in orange and a used vehicle in blue. Households limited to the 2x2 replacement sample. Kernel density computed using the Epanechnikov and automatically generated bandwidths.