

NBER WORKING PAPER SERIES

WHO VALUES FUTURE ENERGY SAVINGS?
EVIDENCE FROM AMERICAN DRIVERS

Arik Levinson
Lutz Sager

Working Paper 28219
<http://www.nber.org/papers/w28219>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2020

The authors thank Georgetown University's McCourt School of Public Policy Massive Data Institute and the Georgetown Environmental Initiative for research funds, Sally Atwater, Spencer Banzhaf, Antonio Bento, Mary Ann Bronson, Meghan Busse, Raphael Calel, Ken Gillingham, and Shaun McRae for constructive suggestions, and Kevin Ankney and Juan Margitic for research assistance. Sager gratefully acknowledges financial support by the Grantham Foundation for the Protection of the Environment and the UK's Economic and Social Research Council. We are especially grateful to Ben Leard and Joshua Linn at Resources for the Future, for sharing with us calculations based on the MaritzCX data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w28219.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Arik Levinson and Lutz Sager. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Who Values Future Energy Savings? Evidence from American Drivers
Arik Levinson and Lutz Sager
NBER Working Paper No. 28219
December 2020
JEL No. Q48,Q58,R48

ABSTRACT

Regulators attest that tightened energy efficiency standards save consumers money. Efficient light bulbs, appliances, and vehicles cost more upfront but reduce energy expenses by more than enough to compensate. We use survey data on American cars and their drivers to examine whether individual drivers have indeed underinvested in fuel economy, given the gas prices they face and the miles they drive. We find that may be true, but only on average. Some drivers could likely have saved money by spending more upfront for efficient cars. But many others could have saved money purchasing less expensive, less fuel-efficient cars. In fact we find little correlation between individual drivers' annual fuel expenditures and their fuel economy choices: a driver's income, sex, age, and education are far more closely associated with their vehicle's fuel economy. We can rule out several explanations for the disconnect. Rich car purchasers do not seem to consider fuel expenses any more than poorer ones, undermining arguments that borrowing constraints prevent low-income consumers from investing in fuel efficiency. And the disconnect between fuel expenses and vehicle choice holds whether we examine anticipated or realized mileage, ruling out mistaken expectations about future driving as an explanation.

Arik Levinson
Department of Economics ICC 571
Georgetown University
3700 O St., NW
Washington, DC 20057
and NBER
Arik.Levinson@georgetown.edu

Lutz Sager
Georgetown University
37th & O St NW
Old North 408
Washington, DC 20057
lutz.sager@georgetown.edu

Who Values Future Energy Savings? Evidence from American Drivers

US regulators attest that energy efficiency standards—for light bulbs, appliances, buildings, and vehicles—save consumers money. The products themselves cost more to manufacture and purchase but quickly pay for themselves in energy savings. In the case of vehicles, the Corporate Average Fuel Economy (CAFE) standards for cars sold in 2012 through 2016 were expected to raise the average cost of a new car by \$950 but promised to save drivers enough gasoline to pay for that incremental cost in just the first three years of driving.¹ The CAFE standards for 2017 through 2025 were expected to add another \$1,800 to the cost of a new car but save more than \$5,000 in gasoline costs over its useful life.² And the March 2020 relaxation of those same standards is now projected to reduce vehicle costs by \$800 but add \$900 to \$1,200 in lifetime gas costs.³ Across administrations, US regulators have projected that even if external costs from pollution, congestion, and accidents are ignored, stricter fuel economy standards pay for themselves.

For that story to be true—obliging manufacturers to increase the fuel-efficiency of their products benefits their consumers—there must be some type of market failure. If efficient appliances or cars would save their owners money, buyers should be willing to pay more for them than their incremental manufacturing costs. Manufacturers in turn should be glad to produce and sell them at a profit. No regulation would be necessary. The fact that regulations are needed for manufacturers to sell products that buyers should prefer suggests that these markets for energy-using durables do not work efficiently.

Proffered explanations for that market failure include consumers' borrowing constraints, information asymmetries, and behavioral factors such as inattention, present

¹ In 2007 dollars. US EPA, “EPA and NHTSA Finalize Historic National Program to Reduce Greenhouse Gases and Improve Fuel Economy for Cars and Trucks,” EPA-420-F-10-014, April 2010. Link [here](#).

² In 2007 dollars. US EPA, “EPA and NHTSA Set Standards to Reduce Greenhouse Gases and Improve Fuel Economy for Model Years 2017-2025 Cars and Light Trucks,” EPA-420-F-12-051, August 2012. Link [here](#).

³ In 2018 dollars. US EPA “Final Regulatory Impact Analysis: The Safer Affordable Fuel-Efficient (SAFE) Vehicles Rule”, March 2020. Link [here](#).

bias, and mistaken estimation of future energy expenditures.⁴ All of these share a common empirical implication. Consumers would be better off paying higher upfront purchase prices in exchange for lower annual energy expenses. And evidence for those lost opportunities for savings should be apparent in data.

In what follows we look for evidence for those lost opportunities in data describing individual American drivers, their demographic characteristics, the cars they own, and how many miles they drive those cars. We use the US National Household Travel Survey (NHTS) to study drivers' actual mileage, effectively asking whether drivers' car choices turn out to have been mistakes, given how much they later use those cars. In follow-up comparisons, we examine used-car prices, effectively asking whether drivers would be better off exchanging their current cars for vehicles that are more (or less) efficient, and we examine time-of-purchase survey data to ask whether car buyers' fuel economy choices make sense, given the distances those buyers expect to be driving. In each case, we find nearly as many drivers who could have saved money by purchasing less expensive, less fuel-efficient vehicles as those who could be saving money in more expensive, more efficient ones. Drivers' characteristics such as income, sex, age, and education are far more closely associated with their vehicles' fuel economy than their fuel expenses.

Here's an example, a particular respondent to the 2017 NHTS. Albert (our pseudonym) was in his 40s, lived in the Southeast, and had annual household income between \$35,000 and \$50,000. He had a new(ish) gasoline-powered Toyota Camry, which he drove nearly 30,000 miles a year. If Albert had purchased a hybrid gas-electric version of the same Toyota Camry, he would have saved more than \$650 per year on gas.

Did Albert make a mistake? Perhaps. The hybrid cost about \$3,800 more than the least expensive gas-powered Camry, so he would have recouped the extra expenditure in a bit less than six years. That's longer than the three-year payback promised by the Environmental Protection Agency (EPA) and the National Highway Traffic Safety

⁴ See, for example, Allcott (2011), Allcott and Knittel (2018), Allcott and Wozny (2014), Ankney (2020), Busse et al. (2013), Leard et al. (2017), Sallee (2014), Sallee et al. (2016), and Gillingham et al. (2019).

Administration (NHTSA), but still respectable. Albert's failure to buy the hybrid might be just the type of mistake EPA and NHTSA used to justify the CAFE regulations.

Other car buyers appear to have made the opposite choice, by buying efficient cars when the energy savings don't compensate for the extra upfront costs. Betty (another pseudonym) drives the hybrid version of Albert's Toyota Camry. In 2017 she was in her 70s, lived in the West, and had household income between \$25,000 and \$35,000. She only drove 4,500 miles a year. Had she been driving the standard Camry, those miles would have cost her \$120 more per year. At that rate it will take her more than 30 years to recover the extra \$3,800 cost of her hybrid. Buying the hybrid could be seen as a mistake, at least from her own, self-interested, purely financial perspective.

Of course, these two car owners might have good reasons, unrelated to fuel expenditures, for preferring the gas or hybrid versions of their cars. Some drivers might be concerned about local pollution or climate change, or might want to display that concern to neighbors. Other drivers might like the quiet of the electric engine, or worry that it poses a danger to pedestrians or bicyclists. Or their commute distances or gas prices might have changed since they purchased their cars, and thus their vehicle decisions were financially optimal at the time of purchase but not with hindsight. These decisions are mistakes only in a personal, purely financial, ex post sense.

Those other external, nonfinancial, ex ante considerations are also omitted from the analyses used to justify the US fuel efficiency standards. By noting that the typical driver could save money by purchasing a more expensive but more fuel-efficient car, those regulatory analyses find that, from the perspective of private fuel expenses alone, American car buyers choose insufficiently efficient vehicles. If that's true, in the data we should expect to see more Alberts who could benefit from more fuel efficiency than Bettys who could have saved money by purchasing less fuel efficiency.

In the 2009 and 2017 rounds of the NHTS, there are 24,362 drivers of vehicles available in either a gas or hybrid version.⁵ For each driver, we can calculate the annual fuel cost difference between the two alternative versions of the car:

⁵ This total is limited to those NHTS observations with complete data on household demographics.

$$\text{Annual fuel cost difference} = \left(\frac{1}{\mu_n} - \frac{1}{\mu_e} \right) m_i \cdot p_g, \quad (1)$$

where μ_n is the fuel economy in miles per gallon (mpg) of the less efficient (gas) car, μ_e is the mpg for the more efficient (hybrid) car, m_i is miles driven per year, and p_g is the price of gasoline. Equation (1) differs for every driver depending on miles driven (m_i), gas prices paid (p_g), and the difference between the gas and hybrid fuel economies for the particular car (μ). For drivers of hybrid cars, (1) represents their annual savings, relative to driving the same number of miles in the gas version. For drivers of gas cars, (1) represents unrealized annual savings from not driving a hybrid.

Figure 1 plots two distributions of the annual fuel cost differences in (1). The outlined, unshaded bars plot the distribution for drivers of the 22,025 gas vehicles. Cases farther to the right in that distribution are more likely to represent (personal, purely financial, ex post) mistakes. They are gas car owners who miss out on the large annual savings of driving the hybrid version of their cars. The shaded bars plot the distribution of (1) for drivers of the 2,337 hybrids. Cars farther to the left are more likely to represent mistakes. Those drivers own hybrids but save little, either because they do not drive much or their gas prices are low. Accordingly, we should expect to see the shaded histogram (hybrids) to the right of the outlined histogram (gas powered), and that is exactly what Figure 1 shows.

Two points need to be made here. First, the shift of the hybrid distribution to the right of the gas-powered distribution could represent decisions by car purchasers based on their pre-purchase expected driving. People expecting to drive more miles could purchase the hybrids, although we show later, using different data collected at the time vehicles are purchased, that this does not seem to be the case. Alternatively, the shift in Figure 1 could represent post-purchase reactions by drivers to the relative differences in driving costs. If we randomly assigned cars to drivers, the recipients of hybrids would find driving less expensive and might drive more, realizing more savings.⁶ In other

⁶ If drivers choose cars based on other attributes, such as trunk space or perceived environmental awareness, and those are uncorrelated with demand for travel, then the car choice is equivalent to random assignment and the hybrid drivers will find driving less expensive, drive more, and realize more savings.

words, drivers facing higher costs could choose to buy hybrids, or hybrid owners could respond to the lower cost by driving more miles. Either way, hybrid owners will exhibit greater annual savings in Figure 1.

The post-purchase reaction to the different price of driving represents the “rebound” effect (Gillingham et al. 2016). Purchasers of efficient cars drive them more, making their investment in the efficiency appear smarter. Purchasers of less efficient cars drive them less, making them seem smart not to have invested in the costlier efficient version. Any analysis we do comparing drivers’ actual annual costs with costs that would be incurred driving the same distance in an alternative car will produce an underestimate of mistakes. Drivers’ post-purchase driving behavior will rationalize their vehicle choices.

The second point involves the complex set of assumptions necessary to figure out which drivers could be saving money in a different vehicle. To know whether an efficient hybrid car is worth the upfront investment, we need to compare the annual fuel savings in equation (1) with the annualized difference in the fixed costs of the two cars. That, in turn, depends on three things: (i) the difference in sales prices, net of any rebates or subsidies for buying an efficient car, (ii) the discount rate used to annualize the price difference, and (iii) the depreciation rate of the extra expenditure on the fuel efficient version.

Because that calculation is so complex and requires so many assumptions, we take a simple and intuitive approach. Whatever the values of those three variables, for each pair of car models there is some level of annual cost differences that would justify the investment in fuel economy. Think of it as a vertical line drawn in Figure 1. Any driver to the right of that line in a standard gas car has underinvested in fuel efficiency. Any driver to the left in a hybrid car has overinvested. We pick a value for that line for each car model, test the sensitivity of our findings to that value, and examine the demographics of drivers on either side.

To frame that empirical approach, we describe the intuition in a simple theoretical model in which cars differ only in their price and fuel economy. The model incorporates the trade-off between upfront efficiency costs and future gas savings. And the model helps us to develop two straightforward empirical tests of car buyer mistakes. Both tests

are designed to overcome the chief obstacle to assessing drivers' car choices: cars differ along many dimensions that are correlated with price and fuel efficiency, so any empirical test needs to control for car characteristics missing from the theoretical model. Our two empirical tests represent alternative approaches to that.

The first test controls for car characteristics correlated with price and fuel economy by examining only those car models, depicted in Figure 1, that can be purchased in either a gas or a hybrid version. We assume those pairs differ only in their fuel economy, as in the theory model, and we compare the fraction of drivers in the gas versions who would have been better off paying more for the hybrids with the fraction owning hybrids who could have saved money by purchasing the gas versions.

That gas-hybrid comparison depends on the strong assumption that the alternative versions do not differ in other respects, such as trunk space, acceleration, or maintenance costs. Because we know that assumption to be false in some cases, our second empirical approach uses all the cars in the 2009 and 2017 NHTS surveys and controls for other car characteristics statistically, rather than by comparing cars in gas-hybrid pairs. We estimate the average fuel economy premium—the incremental upfront cost of purchasing a vehicle with one extra mpg of efficiency—controlling statistically for other observed car characteristics correlated with efficiency and price. We then calculate each driver's actual annual driving costs, given the gas prices they face and their annual mileage, and what they would save annually in a slightly more fuel-efficient car with one extra mpg. If those annual savings exceed the fuel economy premium (annuitized appropriately), the driver would have been better off choosing a more efficient vehicle. If those savings are less than the fuel economy premium, the driver would have been better off paying less up front for a less efficient car. As with the gas-hybrid comparison, we calculate the proportion of drivers making both types of personal, purely financial, ex post mistakes.

Both empirical tests rely on assumptions. We assume we have controlled for other car characteristics, either in our gas-hybrid comparison or statistically. We follow the EPA and NHTSA regulatory analyses, which assume that cars last 14 years and discount future fuel savings at 3 and 7 percent.⁷ In the main part of the analyses, when we examine

⁷ For comparison, the real, inflation-adjusted, dealer-provided new-car loans studied by Busse et al. (2013) range from -0.9 to 9.0 percent.

actual, post-purchase fuel costs using the NHTS, we assume drivers paid the manufacturer's suggested retail price (MSRP) when they purchased their vehicles, adjusted for inflation. Later, we demonstrate the stability of our results when we consider each driver's annual decision to keep or sell a car at concurrent used-car market prices or, using a different survey, when we consider each car buyer's actual transaction cost at the time of purchase and expected annual mileage.

Our approach departs from existing work in important ways. We assess the actual, driver-specific fuel savings by individual drivers. Until now, with few exceptions, the evidence for whether drivers appropriately value future fuel savings has come from expected savings of typical drivers. Some researchers using that strategy have found the fuel economy premium to be less than the present discounted value of expected future savings, suggesting that on average, car buyers make mistakes by purchasing insufficiently fuel-efficient cars (Allcott and Wozny 2014; Gillingham et al. 2019; Grigolon et al. 2018). Others find the premium approximately equal to the discounted future savings, suggesting that on average, buyers are not mistakenly purchasing too little or too much fuel efficiency (Busse et al. 2013; Sallee et al. 2016).

Most prior analyses compare the fuel economy premium with the expected fuel savings at the time of purchase. At best, those expectations are based on miles driven by the average driver of a particular car model. As Bento et al. (2012) note, some individual drivers put many more miles on their cars each year than the average, and others drive less. The high-mileage drivers should be willing to pay more for fuel economy than low-mileage drivers, and that heterogeneity will bias the results toward suggesting that consumers undervalue future savings.

One recent exception that does examine individual drivers' vehicle choices is Banzhaf and Kasim (2019). They use the 2001 and 2009 versions of the NHTS to show that when gas prices increase, high-mileage drivers are more likely to choose efficient cars, and low-mileage drivers, less efficient cars. Like us, Banzhaf and Kasim also find a large mismatch between driving distances and car efficiency. According to their rough calculation, US gas consumption is 15 percent higher because of this mismatch. Without accounting for upfront vehicle costs or car characteristics other than mpg, though, there's no way to know whether that rematching would make individual drivers better off. We

take an entirely different approach, explicitly modeling the trade-off between upfront costs and fuel savings, controlling for a range of vehicle characteristics in addition to mpg.

Instead of comparing gas price elasticities by high- and low-mileage drivers, our approach compares the premium paid for fuel-efficient cars with the actual realized fuel savings of those cars' drivers, based on how far they drive each year and the gas prices they pay. This allows us to contribute to the literature in several ways. We estimate the fraction of drivers who underinvest and overinvest in fuel efficiency, retrospectively based on actual miles driven and, later in the paper, prospectively based on drivers' reported expected mileage at the time of purchase. We find lots of drivers who could have saved money by purchasing fuel-efficient cars. Those drivers undervalued future savings, as suggested by US regulators and by some of the previous work using average mileage. But we also find nearly as many drivers who could have saved money by purchasing *less* efficient cars, spending less to buy the cars but more each year to drive them. Similar to Banzhaf and Kasim, we find that even if the right *share* of drivers owns fuel-efficient cars, the wrong *types* of drivers own them. Of the people who have relatively efficient cars, many drive relatively few annual miles; of those with inefficient cars, many drive relatively more.

Because we use household survey data, we know details about the demographics of the drivers and can assess which particular groups are more likely to undervalue or overvalue future savings: male or female, rich or poor, educated or less educated, old or young. In general, we find that these demographic characteristics play a much larger role in predicting car fuel economy than fuel costs. We find that women and younger drivers are less likely to underinvest in fuel economy in general, but that men and older drivers are more likely to overinvest in hybrid cars relative to equivalent gas vehicles. Other groups that are less likely to underinvest are drivers living in urban rather than rural areas and drivers with graduate degrees. After controlling for those demographic differences, individual drivers' cost-saving potential has little influence on investments in fuel economy.

Before describing those analyses in detail, we recount the evidence on these issues so far, and then sketch a theory that frames our thinking as well as some comparative statics that demonstrate normative implications derived from the analyses.

I. Evidence to date comparing fuel economy premiums and expected annual savings

Our paper is not the first—by a large margin—to model the trade-off between upfront costs and future energy savings or to try to measure empirically consumers’ willingness to pay for future energy savings. Much of that research studies cars, in part because good data are available and in part because vehicle fuel economy has been a focus of regulatory attention. Allcott and Greenstone (2012) synthesize that research in a model that frames the main issues. In their model, relabeled here for consistency with our terminology, consumers choose among durable goods, such as cars, with different energy intensities, labeled μ . Think of cars with two different mpg ratings, efficient μ_e and inefficient μ_n , where $\mu_e > \mu_n$. Efficient cars cost more to purchase, all else being equal. Label that incremental upfront capital cost p_μ . Holding all other vehicle attributes constant, a driver will be better off buying the efficient car if that incremental cost is less than the discounted energy savings:

$$p_\mu < \gamma p_g m \left(\frac{1}{\mu_n} - \frac{1}{\mu_e} \right) F . \quad (2)$$

In equation (2), p_μ represents the price difference between the equivalent efficient and inefficient cars, m is consumption of the energy service (miles driven), p_g is the price of energy (gas), and future cost savings are collapsed into present values by factor F .⁸

The ad hoc parameter γ in equation (2) describes the weight consumers appear to assign to discounted future energy costs. If $\gamma=1$, consumers make rational trade-offs between current costs and future energy savings. If $\gamma<1$, some behavioral anomaly or market failure impedes consumers from purchasing cars or appliances that would save them money.

⁸ For example, applying NHTSA’s assumption that cars last 14 years, if future savings are discounted at 3 percent, F is 11.3. If savings are discounted at 7 percent, F is 8.7.

Estimating γ requires comparing cars that are identical in all respects other than efficiency. That difficulty—separately identifying preferences for fuel economy from other car attributes—is the big challenge. Many car characteristics, like size and power, are associated with more expensive, less fuel-efficient vehicles. As a result, efficient vehicles are not typically more expensive; they are cheaper. Only after controlling statistically for other observable car characteristics does the relationship between price and fuel economy become positive. But that leaves open concern about other car characteristics omitted from the analysis. Finding that buyers undervalue fuel savings could simply mean that car buyers have disutility from some unobserved characteristic of the efficient cars.

Five recent studies have taken creative steps to address the omitted variables problem.⁹ Allcott and Wozny (2014) use monthly data on new vehicle registrations in the United States. They test whether fuel-efficient cars command higher prices when gas prices rise. In other words, they identify $\hat{\gamma}$ using variation in p_g . Because their results are identified by price changes for specific car models, fixed effects can account for unobserved model characteristics. Allcott and Wozny estimate moderate undervaluation of future fuel savings ($\hat{\gamma} = 0.76$).

Sallee et al. (2016) examine used-car auctions in the United States. Cars with fewer miles on their odometers have more expected remaining useful years and therefore a higher payoff to fuel efficiency. When gas prices increase, the price of fuel-efficient used cars increases, and it increases more for used cars with lower odometer readings. Sallee et al. identify $\hat{\gamma}$ using joint variation in $F \cdot p_g$. Their estimates suggest that vehicle prices move approximately one-for-one with discounted future fuel savings, or that $\hat{\gamma} \approx 1$. As with Allcott and Wozny (2014), the identification is within model type, so unobserved differences between models pose less of a problem.

Grigolon et al. (2018) use European car sales data. They compare different variants of the same model, as in our gas-hybrid example. These variants differ in their suggested retail prices and fuel economies. Grigolon et al. exploit changes in market

⁹ Allcott and Greenstone (2012), Greene (2010), and Helfland and Wolverton (2011) contain thorough reviews of earlier papers.

shares of new-vehicle registrations of these variants. Their main estimate is that $\hat{\gamma} = 0.91$, implying that consumers only moderately undervalue discounted fuel savings.

Busse et al. (2013) take a different approach. They estimate the discount rate that would justify the price premium consumers pay for efficient cars. If that rate is implausibly large, it would suggest consumers undervalue future savings (i.e., that $\gamma < 1$). Busse et al. find implied discount rates in line with market interest rates for car loans, and therefore consumers do fully value future savings.

Gillingham et al. (2019) exploit the fact that in 2012 Hyundai and Kia corrected the reported fuel economy for 13 car models. Those cars' fuel economies had previously been exaggerated. The authors find little decline in willingness to pay for the cars, suggesting a value of γ in the range of 0.16–0.39. Their result amounts to a joint test of consumers' willingness to pay for future savings and the degree to which consumers get their information about savings from official reports or vehicle labels.

All of those papers calculate future savings based on the typical miles driven by the average driver. But Bento et al. (2012) remind us that the heterogeneity in mileage will result in underestimates of the value of future fuel savings ($\hat{\gamma}$). Grigolon et al. (2018) attempt to estimate the size of this bias by simulating the distribution of driving patterns based on actual mileage data from the United Kingdom in 2007. Their main result ($\hat{\gamma} = 0.91$) accounts for heterogeneity. But their analysis remains at the level of car models, not individual consumers, and their results are identified by fluctuations in market shares of variants of similar models.

Our approach differs from prior work in that we explicitly examine the realized, post-purchase mileage for each driver. We don't ask whether the price premium for a hybrid Toyota Camry is justified given the average annual mileage for all drivers, or even the average for Toyota Camry drivers. We examine the actual annual miles driven and gas prices paid for each driver. In the context of equation (2), our paper focuses on variation in individual mileage (m_i) across drivers, whereas much of the existing work focuses on variation in cost across products faced by the average driver (\bar{m}). Our approach is most similar to that of Banzhaf and Kasim (2019), who find imperfect matching between drivers who drive many miles and vehicles that are more fuel efficient. But imperfect matching does not in itself show that consumers make mistakes, because

the fuel-efficient cars cost more to purchase. Fuel savings have to be weighed against the extra cost of buying a more fuel-efficient vehicle. And both fuel economy and vehicle prices may correlate with vehicle specifications that influence vehicle choice. We directly account for those differences to see whether each individual driver would have been better off paying more up front for a more efficient car, or paying less for a less efficient car.¹⁰ To help frame the empirical analysis, we start with a simple model, following Levinson (2019).

II. A theoretical sketch

Consider a representative consumer or household, with utility over two goods: miles driven, m , and a numeraire, x . Consumers don't purchase miles directly but instead purchase cars and gasoline, g . Miles driven is the product of gasoline and the car's energy efficiency, μ , expressed as miles per gallon, or mpg: $m = \mu g$. Utility is then

$$U(m, x) = U(\mu g, x). \quad (3)$$

Driving requires two expenditures: gasoline g at price p_g and a vehicle at price P_v . For simplicity, assume the car price is annuitized, or equivalently that the car is leased and P_v represents the annual rental cost. Higher-efficiency cars with higher μ cost more to buy or lease. Call that premium p_μ as in equation (2). Here p_μ is the cost of one extra mpg of fuel economy. The full fixed cost of purchasing a car is $P_v + p_\mu \mu$, the base price plus an increment for extra fuel economy. A car owner's budget constraint is thus

$$Y = x + p_g g + P_v + p_\mu \mu. \quad (4)$$

This budget constraint can be represented by a plane in three dimensions: x , g , and μ . Those are the three purchases: the numeraire (x), gallons of gasoline (g), and fuel economy for the car (μ). But that's not the most instructive representation because miles m are in the utility function, not gasoline g .

¹⁰ Although it is not the focus of our paper, we do generate estimates of γ , the undervaluation parameter in equation (2). We estimate that, on average, drivers value \$1 of future fuel savings at approximately equal to \$0.10–0.50 in upfront costs (see Table 3). That is at the lower end relative to prior results, and it depends on the discount rate, whether we use cars' sticker prices or actual sales prices, and whether we examine anticipated or realized miles.

For a clearer view of the trade-offs, rewrite (4) replacing g with m/μ , and in a more familiar, graphable form with the numeraire on the left:

$$x = (Y - P_v) - p_g \left(\frac{m}{\mu} \right) - p_\mu \mu. \quad (5)$$

That's a budget surface in which two of the dimensions are goods in the utility function: the numeraire x and miles m . The third dimension is fuel economy, μ .

Maximizing utility in (3) with respect to the budget constraint (4), paired with a typical convexity assumption, results in two first-order conditions:

$$\begin{aligned} (i) \quad & U_m / U_x = p_g / \mu \\ (ii) \quad & U_m / U_x = p_\mu / g. \end{aligned} \quad (6)$$

The first condition in (6) indicates that the marginal rate of substitution between the two goods in the utility function, miles and the numeraire, should equal the cost of going one more mile by purchasing more gasoline. That cost is p_g / μ , the price of gas divided by the car's mpg. The second condition in (6) indicates that the same marginal rate of substitution should also equal the cost of going one more mile by purchasing a more efficient car. That cost is p_μ / g , the price of fuel economy divided by the number of gallons being used.

The intuition is simple. Drivers have two ways to travel an extra mile. They can buy more gasoline at price p_g for a car with any fixed fuel economy μ . That's the first condition in (6). For any given fuel economy μ , there's a linear trade-off between miles traveled (m) and the numeraire (x). Figure 2 plots that trade-off for two cars: a solid budget line for an inefficient car with mpg μ_1 and a dashed budget line for an efficient car with mpg μ_2 . The first condition in (6) is represented in Figure 2 by a tangency between an indifference curve (with slope U_m / U_x) and a budget line (with slope p_g / μ).

The second way that drivers can travel an extra mile is by purchasing a car with better fuel economy at price p_μ for a fixed amount of gasoline g . That represents a lowering of the vertical axis in Figure 2, as from μ_1 to μ_2 . To show that, Figure 3 sketches two different trade-offs between fuel economy and the numeraire, drawn for two

particular mileages, m_1 and m_2 . A driver can minimize the cost of driving m_1 miles by purchasing fuel economy μ until the cost of an extra unit of fuel economy p_μ equals the savings from that extra efficiency. The savings is just the derivative of the total cost of driving, mp_g/μ , with respect to μ . So in the optimum, the cost of that extra mpg, p_μ , equals those savings:

$$p_\mu = \frac{mp_g}{\mu^2}. \quad (7)$$

That is just another way of writing the combined first-order conditions in (6), replacing gallons g with m/μ . The optimum is depicted in Figure 3 as μ^* for a car owner driving m_1 miles per year. A car owner driving m_2 miles per year will have different trade-offs, represented by the dashed line with lower possible expenditures on the numeraire x and higher optimal fuel economy.¹¹

The two first-order conditions in (6), and their representations in Figure 2 and Figure 3, motivate our two empirical tests. Given two cars with efficiencies μ_1 and μ_2 , like the gas and hybrid Toyota Camrys in the introduction, drivers' choices will depend on their willingness to trade off miles traveled for other goods in the numeraire. Given the indifference curve depicted in Figure 2, the driver would be best served choosing the inefficient car and driving m^* miles. Given those same budget constraints, however, a driver with a flatter indifference curve might be better off choosing the efficient car and driving more than m^0 miles. Because we cannot observe utility, we cannot identify all of the cases where a driver would be better off in a more or less efficient vehicle. A different driver might be making a mistake choosing to drive m^* miles in the inefficient car.

We can, however, identify some cases that clearly cannot be optimal—they are personal, purely financial, ex post mistakes—while holding other vehicle attributes constant, even without knowing utility. In Figure 2 any car owner driving more than m^0

¹¹ Or, consider a third way of looking at the trade-offs. For any given expenditure on the numeraire, there's a quadratic relationship between miles (m) and mpg (μ). A household can spend more on energy efficiency to travel more miles, up to a certain point, holding the total budget and expenditures on x constant. Beyond that point, purchasing more efficiency doesn't pay off and results in fewer miles traveled. At the optimum

$$\mu^* = (Y - P_V - x) / 2p_\mu.$$

miles per year in a vehicle with efficiency μ_1 , choosing a point on the solid budget line to the right of m^0 , is making a mistake. Albert in our introduction may be in that position. He could drive that same number of miles at lower total cost, including the upfront price of the car, in the more efficient vehicle μ_2 . Similarly, anybody driving fewer than m^0 miles per year in a vehicle with efficiency μ_2 , the dashed budget line, is making a mistake. That may be Betty. She could have saved money by purchasing the less efficient car.

Our first empirical test examines pairs of cars that come in both standard gas and hybrid gas-electric versions. We look for people with low annual gas expenses who have purchased the hybrids, and people with high annual expenses in the gas models. Our second empirical test focuses on the second first-order condition in (6) and on Figure 3. For each driver in our sample, we can calculate the annual gas savings from purchasing a slightly more efficient car, a model with one more mpg: mp_g / μ^2 . If that is higher than the cost of an extra mpg, the driver has underinvested in fuel economy and could be saving money in a more efficient car. If those annual savings are lower than the cost of an extra mpg, the driver has overinvested in fuel economy and would have been better off purchasing a less expensive, less efficient car.

To reiterate, the mistakes identified by both approaches are errors only in that the upfront cost of fuel efficiency does or does not exceed the associated future fuel savings. That is the calculation the Obama administration used to show that tightening fuel economy standards would save drivers money, and that the Trump administration used to show that relaxing the standards would cost drivers money.¹² The calculation ignores externalities such as pollution, other omitted vehicle attributes correlated with efficiency, and deviations from drivers' expected mileage. That's why we keep referring to the mistakes as personal, purely financial, and ex post. And so in the next section we assess which drivers in practice are driving more miles in cars that could be saving them money if they had spent more on fuel efficiency, and which are driving relatively few miles in cars that could have saved them money had they spent less on fuel efficiency.

¹² Bento et al. (2018) and Bento et al. (2019) discuss the assumptions underlying the findings by the consecutive administrations.

III. Identifying mistakes in practice: Comparing gas and hybrid cars

For post-purchase data on the annual miles people drive, and the cars in which they do that driving, we rely on the 2009 and 2017 waves of the National Household Travel Survey. It includes household demographics, the annual mileage in each vehicle, and the make, model, and model year of those vehicles. We match that information with WardsAuto data for each make and model year to get vehicle characteristics, including size, engine power, and the EPA estimated combined city-highway fuel economy.¹³

Table 1 presents some summary statistics. The first two columns contain data for only those vehicle models that come in both gas and hybrid versions. Drivers of hybrid cars get more miles per gallon of gas and travel more miles, as expected. The mileage difference translates to a 7 percent increase (relative to the midpoint of the two values). That's the rightward shift in the shaded, hybrid distribution in Figure 1.

That simple 7 percent difference illustrates our central finding. Compare it with the fuel economy difference in the second row of Table 1, which translates to a 30.4 percent drop in the cost of traveling each mile.¹⁴ The ratio of those two numbers, 0.23, is not an arc elasticity because it combines both the ex post response by hybrid owners to the lower cost of driving and the ex ante selection of hybrids by high-mileage drivers. By including selection, 0.23 overstates the true price elasticity. Labandeira et al. (2017) survey economic estimates of the true elasticity of driving with respect to the cost. Their meta-analysis has an average short-run price elasticity of demand for gas of 0.20 and an average long-run elasticity of 0.53. In other words, if we randomly assigned drivers to

¹³ The two waves of the NHTS have data on 294,409 households that own 776,731 cars and light trucks. We restrict our sample to those households in the NHTS for which we have information on miles driven and a complete set of covariates (income, age of household head, education, rural or urban). We further limit the sample to those vehicles that we successfully matched (by make, model, and year of production) to engineering characteristics (mpg, length, width, height, weight, liters, valves, horsepower, and revolutions per minute, rpm) from WardsAuto. This final sample contains 157,403 household-vehicle observations. A subsample of 24,362 comprises vehicle models that come with either a gas-powered or a hybrid engine, and for which both are present in the NHTS sample.

¹⁴ The cost of traveling one mile is p_g/μ . The price of gas cancels, and the percentage difference in the cost of driving a mile in columns 1 and 2 of Table 1 is $(1/\mu_e - 1/\mu_n)/((1/\mu_e + 1/\mu_n)/2)$.

gas and hybrid cars, the mileage difference we see in Table 1 between the alternative models, 7 percent, is at the low end of what we should expect to see based on those drivers' responses to their driving costs alone. In aggregate, selection of fuel-efficient cars by high-mileage drivers appears to play little or no role in the cars people drive.¹⁵

To examine the choices by individual drivers, we start with the simplest strategy that matches the introductory intuition and the first empirical test from Section II, comparing gas and hybrid versions of the same make and model, as in the first two columns of Table 1. For each driver with a car that comes in both versions, we calculate the annual value of the fuel cost difference between the gas and hybrid models. That cost difference, in equation (1), differs for every driver, depending on miles driven (m_i), gas prices (p_g), and the difference between the gas and hybrid fuel economies for the particular car. It is plotted in Figure 1 for two groups: the black-outlined bars for drivers of gas vehicles, and the shaded bars for hybrids.

As a way of comparing the two distributions in Figure 1, we note that for each gas-hybrid pair, there is some benchmark value of annual fuel savings above which driving a hybrid would make sense for personal financial reasons. Figure 4 sketches the distributions of drivers in Figure 1 as frequencies rather than densities, to emphasize that the hybrid market share is smaller. The benchmark cost difference is depicted as a starred vertical line. Drivers of standard gas cars spending more than the benchmark could have saved money by spending more upfront on the hybrid. The share of gas drivers in that position is $B/(A+B)$. Drivers of hybrids spending less than the benchmark could have saved money by purchasing the less expensive gas car. That share is $C/(C+D)$.

Finding the true benchmark annual cost difference is difficult. It requires knowing all the various parameters of equation (2), including the price premium for the hybrid, the cars' lifespans, and consumers' discount rates. Accordingly, as our main approach, we calculate the benchmark annual cost difference by assuming that the hybrid price premium for each model pair is the difference in the two cars' MSRPs. That may

¹⁵ Recent contributions have highlighted another margin for fuel savings, namely how a given number of miles are driven (Langer and McRae, 2017; Knittel and Tanaka, 2019). We focus only on the number of miles driven and gas prices paid, assuming constant fuel use per mile across drivers in the same car.

overstate the cost of fuel economy if consumers receive tax incentives or other rebates for purchasing the fuel-efficient versions, and that overstatement may be different for different model pairs. In a follow-up approach below, we will show that our main conclusions with regard to the demographics of drivers who make different fuel economy choices do not differ when we use the actual sales prices or used car prices in place of MSRPs.

Calculating mistakes using the assumptions behind the CAFE analyses

In 2012, EPA and NHTSA issued new fuel economy rules for cars to be produced in model years 2017 through 2025. That ruling was accompanied by a regulatory impact analysis (RIA) predicting that the new rules would add \$1,800 to the cost of an average new vehicle and save more than \$5,000 in gas costs over the life of the car.¹⁶ The EPA and NHTSA analyses hinge on vehicle depreciation rates and discount rates. They assume various lifespans for cars based on model-specific calculations. The average lifespan is 14 years, and we apply that to all the cars in our sample. EPA and NHTSA then apply either a 3 percent or a 7 percent discount rate to future fuel savings, following standard US government guidance.¹⁷

Our version of that analysis can be found in Table 2. The 2009 and 2017 NHTS data contain 24,362 cars that come in gas or hybrid versions, 2,337 of which are hybrids. We start by assuming those cars sell for their MSRP and last 14 years.¹⁸ If we discount future fuel savings at 3 percent, then 3,646 of the 22,025 owners of gas vehicles would be saving money in the hybrid version of their model pair. That's 17 percent, the fraction $B/(A+B)$ in Figure 4. Given their driving mileage and gas prices, their discounted future fuel savings would more than cover the price premium for the hybrid version. The aggregate statistics comport with prior results that car buyers undervalue savings: 4,120 drivers would be better off in hybrids, and only 2,337 are driving hybrids.

¹⁶ See footnote 2.

¹⁷ US Office of Management and Budget, Circular A-4, 68 FR 58366, October 9, 2003.

¹⁸ Technically, EPA and NHTSA's 14-year lifespan is for vehicles assumed to be identical in every respect other than fuel economy, but we know that gas and hybrid models differ in some respects. That is why in the next section we present an alternative approach controlling statistically for vehicle differences.

With the NHTS data on post-purchase driving by individual drivers, however, we can see not just the share of drivers in hybrids but which drivers are in hybrids. Of the 4,120 drivers in column 3 of Table 2 who would be better off in hybrids, only 12 percent (474) are actually driving hybrids. That's fraction $D/(B+D)$ in Figure 4. The other 88 percent are making the apparent mistake of underinvesting in fuel economy.

Again, these are mistakes only in an ex post financial sense, ignoring concerns about externalities due to pollution, traffic, or accidents, and assuming the gas and hybrid versions are identical. The drivers of gas cars who forgo savings from not buying the presumably identical hybrid version of their vehicles may simply favor attributes of the gas vehicles that differ from the hybrids, and the hybrid drivers who could have saved money in a standard gas version may prefer features of the hybrid.

If we discount future fuel savings at 7 percent, the aggregate statistics comport with prior results that car buyers fully value savings. There are 2,430 who would save money in hybrids, and nearly precisely that many, 2,337, actually drive hybrids. But again, our data show those are mostly the wrong 2,337 drivers. Of the 2,430 for whom a hybrid would save money, only 12 percent (286) actually drive hybrids. Increasing the discount rate reduces the number of drivers predicted to be better off in a hybrid but doesn't significantly change the share of those who actually do so, 12 percent.

Table 2 also reports the other type of mistake, overinvestment in fuel efficiency. Of the 20,242 drivers predicted to be saving money in gas cars, with future savings discounted at 3 percent, 18,379 are driving gas cars. The other 1,863, or 9 percent, are making an apparent mistake by driving hybrids: these low-expense drivers would be better off in less fuel-efficient, less expensive gas cars. Or, looked at another way, of the 2,337 drivers actually in hybrids, 1,863 (80 percent) would have saved money in gas cars. That's $C/(C+D)$ in Figure 4. They are making the apparent mistake of overinvesting in fuel economy.

And again, if we discount future fuel savings more, at 7 percent, the proportions are the same. More drivers would save money in the gas cars, 21,932. Of those, 91 percent do choose the gas versions. The other 2,051, or 9 percent, are making the apparent mistake of driving hybrids, overinvesting in fuel efficiency.

Why don't the shares of apparent mistakes change much, even though the discount rate substantially cuts the value of driving a hybrid? An obvious explanation is that drivers' gas-hybrid decisions are not based on potential fuel savings. Look again at Figure 1. The two distributions mostly overlap. Raising the discount rate from 3 to 7 percent shifts to the right the cutoff annual fuel savings that would make choosing the hybrid worthwhile. That shrinks the share of drivers for whom we consider choosing a hybrid to be optimal. But it also shrinks the share of drivers who have chosen to drive a hybrid by almost the same amount.

More concretely, if cars were randomly assigned to drivers in our sample, regardless of gas prices or annual mileage, then changing the discount rate would have no effect on observed apparent mistakes in car choice. That seems to be the case. Table 2 is just one more way of demonstrating the main takeaway from Figure 1: drivers' choices of fuel-efficient cars are unrelated to the cars' annual fuel costs.

It turns out, as we can show with these data, the association between drivers' demographic characteristics and their cars' fuel efficiency is far closer than the association between drivers' potential or realized fuel savings and their cars' fuel efficiency. Using the NHTS survey data, we can examine whether the failure to consider fuel savings differs across demographic characteristics of drivers.

Demographic characteristics and the choice between gas and hybrid cars

The natural place to begin examining demographic differences is with income. Figure 5 illustrates how the proportions of drivers who would likely be better off in different versions of their same vehicles differ by household income, given their annual mileage and the gas prices they pay.¹⁹ The figure uses the cutoff corresponding to 14 years of driving discounted at 3 percent, as in the left panel of Table 2. The upper line depicts the share of households who have annual expenses above the cutoff but are driving gas cars, $B/(B+D)$ in Figure 4. That is the mistake justifying fuel economy regulations: people don't invest enough in efficiency. Higher-income households are less likely to fit that pattern. The lower line in Figure 5 depicts the share of drivers with

¹⁹ The categories of household income in the 2009 and 2017 NHTS data are not directly comparable. We use the 2017 NHTS for this example.

annual expenses below the cutoff but who drive hybrids, $C/(A+C)$ in Figure 4. That's the opposite mistake: these people are likely to have overinvested in fuel efficiency. Higher-income households are more likely to have made that apparent mistake.

Hybrids make up only about 10 percent of the overall market share of these gas-hybrid pairs. What's interesting is not the shares themselves but how those shares differ by income. To normalize for the different market shares, we subtract from each income group's share the average market share for that vehicle type. Figure 6 plots versions of the two lines from Figure 5 adjusted for those market shares.²⁰ The left axis reports the difference between the share of the group likely to be making the particular mistake—under- or overinvesting in fuel economy—and the overall market share of that version of car. The differences in Figure 6 by income are based entirely on the propensity of different income groups to choose hybrid or gas cars, conditional on their annual mileage and gas prices, and factoring out aggregate market shares of vehicle types. Poorer households are relatively more likely to be missing out on annual fuel savings from driving hybrids, and richer households are more likely to have spent too much upfront for hybrids.

Figure 6 does not suggest that poor or rich households systematically make more mistakes. Rather, they make different types of mistakes. Poorer households are more likely to have chosen a gas car when a hybrid would be saving them money, and richer households are more likely to do the opposite.

In Figure 7 we conduct the same exercise for other household characteristics: sex, age, education, and rural versus urban. Unlike income, these other demographic variables are reported consistently across NHTS survey waves, so we include both 2009 and 2017 data here. We calculate the shares of apparent mistakes for each demographic group, normalized by the market shares for the two types of cars. The open circles represent the shares of drivers with high expenses who drive the gas versions of their vehicles. They could be saving money in a hybrid. The solid diamonds represent the opposite mistake, drivers with low annual expenses who drive the hybrid version. They could have saved money by purchasing the gas version.

²⁰ In the context of Figure 4, the top line in Figure 6 is $B(B+D)-(A+B)/(A+B+C+D)$. The bottom line is $C(A+C)-(C+D)/(A+B+C+D)$.

At the top of Figure 7, among drivers with high annual expenses, men are less likely than women to make the apparent mistake of choosing a gas car. But among drivers with low annual expenses, men are *more* likely than women to make the opposite apparent mistake, driving hybrids.

For drivers younger than 40, a higher proportion with high annual expenses drive gas cars even though they would be better off in hybrids. For drivers older than 40, that pattern is reversed.

Drivers who have not been to college look like the younger drivers. Fewer drive hybrids but would likely be better off in standard gas-powered cars, and more drive gas cars who would likely be better off in hybrids.

One final distinction in Figure 7 is worth noting. Rural and urban drivers appear less dissimilar in their propensity to make each type of apparent mistake. Our analysis controls for miles driven and gas prices, so those distinctions should not matter here. If the choices between hybrid and gas variants of the same model are determined by social or cultural differences, we would have expected to find more pronounced differences between rural and urban drivers.

The demographic analyses in Figure 6 and Figure 7 examine the sample according to single demographic characteristics, one at a time: income, gender, age, location, education. However, these characteristics covary. We cannot be certain, for example, whether differences between education groups are really due to differences across income levels, or vice versa. To examine how all five demographic characteristics together are associated with the fuel economy of the cars people drive, we estimate linear probability models of whether drivers have chosen the hybrid or gas versions of their car models.

Hybrid choice and driver demographics

Table 3 regresses the choice of each driver (1 for hybrid, 0 for gas) on the cumulative discounted value of future fuel savings (in \$1000s) from the hybrid relative to the gas version of the same model, given each driver's annual miles driven and fuel prices. The regressions also include the upfront incremental cost of the hybrid car. They control for all five demographic characteristics, vehicle characteristics that affect fuel economy, and a full set of make and year-by-type fixed effects. We also include state

fixed effects to account for regional differences, such as availability of hybrids or local tax rebates. The calculated fuel savings again mirror the assumptions from the CAFE analysis, summing annual savings over 14 years, discounted at 3 and 7 percent. Because income groups are not comparable between the 2009 and 2017 NHTS waves, we limit the analysis here to the larger 2017 survey.

To demonstrate the range of outcomes for our analysis, we calculate two extremes. The first pairs the lower discount rate of 3 percent with a low estimate of the upfront investment cost, using half of the difference in MSRPs between the two vehicles in each model pair. That combination tilts the choice sharply in favor of choosing the hybrid: its upfront cost is low and future savings are highly valued. The other extreme pairs the higher discount rate of 7 percent with a high estimate of the upfront investment cost, twice the MSRP difference. That tilts the choice against the hybrid.

Start with the low option, favorable to the hybrid cars. Column 1 of Table 3 suggests that an additional \$1,000 in discounted fuel savings is associated with an increase in the probability of driving a hybrid by 0.38 percent. By comparison, a \$1,000 increase in the incremental cost of the hybrid is associated with a 4.3 percent reduction in the probability of driving a hybrid. If consumers fully valued future energy costs—that is, if $\gamma=1$ —the two coefficients would be equal. A dollar’s worth of discounted future fuel savings would have about the same association with car choices as a dollar’s worth of upfront costs. But our estimates in column 1 suggest significant undervaluation of future fuel savings relative to upfront investment cost: $\hat{\gamma}$ from equation (2) is about 0.09.

Demographic characteristics of drivers are strongly associated with the probability of driving a hybrid, even after controlling for fuel savings potential, upfront investment cost, and vehicle specifications. Drivers with higher incomes are more likely to own hybrids. This could be for many reasons, some nonfinancial. Perhaps high-income drivers have altruistic concern for the environment or want to signal that concern to neighbors (Sexton and Sexton, 2014). But other reasons rich drivers might own hybrids are financial. Perhaps they have less trouble affording or borrowing to buy the more expensive hybrids. Or perhaps they are more likely to do the math and get closer to choosing their optimal level of fuel economy.

Column 2 of Table 3 tests these financial explanations, that richer drivers are more likely purchase a hybrid when the gas savings justify the investment. The specification in column 2 includes an interaction between the cumulative fuel savings and an indicator for households with annual incomes above \$50,000. Although the coefficient on the interaction term is not quite statistically significant at 5 percent, its point estimate suggests that high-income drivers' propensity to own a hybrid is more responsive to fuel savings potential, though only to a small degree. Their estimated trade-off between future fuel savings and upfront investment costs is only marginally higher than for households with lower incomes ($\hat{\gamma}$ of about 0.10, compared with 0.05). That difference suggests that rich households are not significantly more likely to do the math, and that poor households are not more likely to be credit constrained. Higher-income drivers are on average more likely to own a hybrid, but evidently not for reasons of economic cost. Some other reason—altruism or environmental virtue signaling—must explain the income effect.

The regressions in Table 3 also show that some demographic traits are significantly more powerful predictors of observed vehicle choice than are the economic considerations that we would have expected to motivate the decision—fuel savings potential and upfront cost. Possessing a graduate degree or being older than 60 is associated with an increase in the probability of owning a hybrid that is seven or eight times larger than from a \$1,000 increase in fuel savings.

Columns 3 and 4 of Table 3 use the other extremes of our assumptions: a 7 percent discount rate and a price premium for hybrid vehicles equal to twice the MSRP difference. That makes purchasing a hybrid less financially worthwhile. The different assumptions have a predictable and mechanical effect on the regressions. We've discounted cumulative future fuel savings, so its coefficient grows, and we've quadrupled upfront costs, so its coefficient falls by three-fourths. But even with these exaggerated disincentives to drive a hybrid, fuel costs predict very little of the hybrid choice, and only half as much as upfront costs ($\hat{\gamma}=0.46$). Richer drivers are not statistically significantly more likely to take fuel costs into account, and having a graduate degree has 10 times the predictive power for hybrid ownership as \$1,000 of future fuel savings.

The results so far demonstrate that drivers' choices of fuel economy do not seem to be determined by future fuel savings. One possible explanation is that the gas-hybrid

model pairing is not the best way of controlling for other differences between less efficient and more efficient cars. Some hybrid vehicles clearly identify themselves as such, perhaps signaling the environmental concerns of their drivers. Battery life may limit the resale value of hybrid vehicles. Hybrids may have less interior room or cargo space. And the on-the-road performance of gas and hybrid vehicles may differ. For those reasons, in a second approach, based on the second condition in (6) and Figure 3, we expand the sample to include all cars, not just the models that come in gas and hybrid versions.

IV. Identifying mistakes in practice: A regression approach using all car models

Instead of controlling for other car characteristics with a simple gas-hybrid comparison, in this section we use the full sample of cars and a calculation of the incremental costs and benefits of purchasing a more efficient vehicle. A car owner i who drives m_i miles each year could save a little money on gasoline each year by purchasing a slightly more fuel-efficient car, with higher mpg (μ). Annual gas costs c_i are

$$c_i = \frac{m_i p_g}{\mu_i} . \quad (8)$$

And the savings from buying a more efficient car, given m_i , are

$$-\frac{\partial c_i}{\partial \mu_i} = \frac{m_i p_g}{\mu_i^2} . \quad (9)$$

The right-hand side of (9) is the same as the right-hand side of the first-order condition in equation (7). That incremental benefit from fuel economy is easy to calculate for each of the drivers in the data because we know how much they drive (m_i), what they pay for gas (p_g), and their cars' mpg (μ_i).

A rational, informed car buyer will purchase fuel economy until the present discounted savings in (9) are equal to the capital cost of purchasing a vehicle with an additional mpg, p_μ , holding all other attributes constant. So in equilibrium,

$$p_\mu = \left(\frac{m p_g}{\mu^2} \right) F \quad (10)$$

where F is the multiple that translates annual fuel savings into a present discounted value. Any car owner for whom the right-hand side of equation (10) is larger than the left-hand side will be making a (personal, financial, ex post) mistake and should have purchased a more efficient car, with higher mpg (μ). Any owner for whom the left side is larger than the right will be making the opposite mistake and should have purchased a less efficient, less expensive car.

The annual savings—the bracketed term in (10)—are simple to calculate using the NHTS data, but the other two terms, p_μ and F , are mostly unknown. As in our gas-hybrid comparison, for any particular values of p_μ and F there is some benchmark value for savings, mp_g / μ^2 , above which it would be worthwhile to pay for a more efficient car. All drivers with annual savings greater than that benchmark will be more likely to have underinvested in fuel efficiency, and all drivers with savings smaller than the benchmark will be more likely to have overinvested. We will test the robustness of our findings to the choice of the benchmark savings level, but we are primarily interested in the demographic characteristics of drivers with savings above or below that benchmark, and those turn out to be insensitive to the choice.

What is p_μ ? Before they were rolled back in March 2020, the CAFE standards were projected to raise fleetwide fuel economy from 35.5 mpg in model year 2016 to 54.5 in 2025, at an average cost of \$2,017 per vehicle in 2017 dollars.²¹ That is \$106 per mpg. The previous round of CAFE standards raised the fleetwide fuel economy from below 30 mpg to 35.5 mpg by 2016, at a projected cost of \$1,140 per vehicle, or \$207 per mpg.²²

To get a sense for what p_μ might be in the NHTS data, we regress car prices on car characteristics, using the MSRPs from the WardsAuto price data:

$$MSRP_{mk} = \alpha + \rho \cdot \mu_{mk} + \mathbf{bX}_{mk} + \psi_{mk} + \varepsilon_{mk} . \quad (11)$$

²¹ Converted from 2010 dollars used in the RIA using the Consumer Price Index. See footnote 2.

²² Converted from 2007 dollars. See footnote 1.

The price of trim t of model m and make k is a function of its fuel economy μ_{mk} as well as a range of other vehicle characteristics \mathbf{X}_{mk} . We also include either make or model fixed effects.

Table 4 presents estimates of (11) for the 14,789 cars in the WardsAuto data set. Column 1 has means and standard errors. Column 2 regresses price on fuel economy alone, without other characteristics, demonstrating the omitted variable problem. Cars with one extra mpg sell for \$1,420 less on average, not more. In column 3, we control for other car characteristics. In that specification, each extra mpg is associated with an extra \$197 in upfront vehicle costs. Column 4 adds make fixed effects, and column 5 adds model fixed effects. Our estimates for the marginal cost of fuel economy, p_μ , range between \$110 and \$340. This is similar to estimates from hedonic pricing methods used to estimate the contribution of fuel economy to vehicle prices and consumers' willingness to pay, though these vary widely (Greene et al. 2018).

For comparison, the mean annual savings from one extra mpg from equation (9) is \$56, and the median is \$42. Earned annually, discounted at 3 percent for 14 years, and summed, that's \$618. Discounted at 7 percent, it would be \$479. Both savings are larger than the coefficients on fuel economy in Table 4, consistent with the idea that car buyers undervalue future fuel savings, on average.

We start by using two estimates of p_μ , \$110 and \$340 from columns 4 and 5 of Table 4. Those are, respectively, slightly smaller and larger than the back-of-the-envelope calculations based on NHTSA's analysis of recent CAFE standard changes. But again, all we really care about is showing the demographic differences between drivers with savings larger and smaller than the benchmark, and as we show below, those demographic differences are not sensitive to the choice of p_μ .

The next step is to choose the discount and depreciation rates. We again follow the assumptions used in the EPA and NHTSA regulatory impact analyses, discounting future fuel savings for 14 years at 3 and 7 percent, which is equivalent to multiplying savings by $F=11.3$ and 8.75 in equation (10). We construct two cutoffs, one favorable to fuel economy investments, one less favorable. The first uses the low discount rate (3 percent) paired with the low price for fuel economy (\$110); the second uses the high discount rate (7 percent) paired with the high price for fuel economy (\$340). That leads to

two benchmark annual fuel savings level cutoffs, above which it would make sense for drivers to purchase one more mpg worth of energy efficiency: \$10 and \$39 per year.²³

We use these benchmark annual savings, \$10 and \$39, to classify all drivers as having purchased either too little fuel economy—because the marginal discounted lifetime savings from one more mpg outweigh the cost—or too much fuel economy. The first, \$10, means that any driver who would save more than \$10 per year in a vehicle with one more mpg of fuel economy has underinvested in fuel economy, and any driver who would pay less than \$10 more in fuel costs in a vehicle with one less mpg of fuel economy has overinvested. This classification is shown in Table 5.

When we discount future fuel savings at 3 percent and impose the low fuel economy price of \$110 per mpg, the analysis in Table 5 suggests that 142,924 (92 percent) of the 155,572 drivers in our sample bought cars with too little fuel economy. Their discounted lifetime savings from buying one more unit of mpg would outweigh the cost. The other 12,648 drivers (8 percent) overinvested in fuel savings. More than 10 times as many underinvest as overinvest. The fact that most drivers would benefit financially from investing more in fuel efficiency supports the argument in the benefit-cost analyses done for the US fuel economy rules. Those rules purportedly pay for themselves by requiring manufacturers to sell more cars that cost more but save fuel.

The bottom row of Table 5 shows not only the fraction of drivers who appear to underinvest in fuel efficiency, but also the magnitude of the average underinvestment. The 142,924 drivers who underinvested incur \$6,000 more in total lifetime average driving costs, relative to cars with the optimal level of fuel economy, whereas the 12,648 who overinvested lose only \$400.

Table 5 suggests that when we discount future fuel savings by 7 percent and impose the high fuel economy price of \$340 per mpg, the two types of mistakes are about equally likely: 53 percent underinvest and 47 percent overinvest. In other words, the \$39 cutoff is close to the \$42 median savings from an additional unit of mpg. And not only are the shares of mistakes more similar with the high discount rate and cost, the financial magnitudes of those mistakes are also closer. The average driver who underinvested

²³ For the first cutoff, $\$10 = \$115/11.3$. For the second, $\$39 = \$340/8.75$.

incurs an extra \$2,400 in total lifetime driving costs, whereas the average driver who overinvested loses about \$1,300.

Both the shares and the magnitudes of mistakes depend on discount rates and our estimates of p_{μ} . We are more interested in the types of people likely to under- or overinvest in fuel efficiency, which turn out to be relatively independent of those assumptions. Figure 8 begins to describe these distinctions, plotting the share deemed to have underinvested in fuel efficiency as a function of household income, again limited to the 2017 wave of the NHTS. It uses the high-cost combination: 7 percent discount rate and \$340 per mpg. Interestingly, this works in the opposite direction as our gas-hybrid comparison. Here the probability that a household underinvests in fuel efficiency *increases* with income, whereas the probability that a high-expense driver purchases a gas car *decreases* with income (Figure 6).

In this case, however, most of the distinction in Figure 8 results from the fact that richer households drive more miles. As there is little association between fuel savings potential and fuel economy choices, groups that drive more have higher fuel savings potential and are therefore more likely to appear to underinvest. In our gas-hybrid comparison, we controlled for driving by examining the share of high-expense drivers who drive the gas or hybrid cars, compared with those shares for low-expense drivers. We cannot do that in this case because we are classifying all drivers as likely to make one of the two mistakes, depending on whether their annual potential savings from an extra mpg are larger or smaller than our benchmark values.

To factor out these group differences in mileage, we construct a counterfactual version of Figure 8 where we assign to all drivers the median fuel economy (mpg) in the 2017 NHTS data. This erases differences in fuel economy choice between groups and instead identifies the shares of apparent mistakes based solely on group differences in miles driven and gas prices. We then subtract these counterfactual shares calculated with median fuel economies from the actual ones depicted in Figure 8. The remaining share differences are plotted in Figure 9. They are entirely due to differences in fuel economy choices between groups, after having controlled for differences in mileage and gas prices.

Figure 9 shows that the share of drivers who could benefit financially from an extra mpg, given their driving expenses, is larger for drivers with higher incomes.²⁴ Recall that the shares are a function of our assumed price of fuel economy (p_μ) and discount and depreciation rates. So we aren't focused on the level of mistakes per se, only on how they change with income. Higher-income households are distinctly more likely to choose cars that could be more cost-effective, even after factoring out their propensity to drive more miles. Presumably those households focus on other costly car features.

Figure 10 reports the results of the same exercise for other driver characteristics. As in Figure 7, we include both 2009 and the 2017 NHTS data. We calculate the share of households in each demographic group deemed likely to be driving cars with too little fuel efficiency, given their annual expenses and our assumptions about p_μ and depreciation and discount rates. We then calculate the counterfactual shares for each group as if every driver were assigned the median fuel economy. Then we report the difference.

Figure 10 shows that given their annual mileage and gas prices, men are more likely than women to be driving vehicles that are not optimal: they could be saving money in higher-mpg cars. Again, this is the opposite result from our gas-hybrid comparison, which found that low-expense men were more likely to be driving hybrids. People older than 60, with graduate degrees, and living in urban areas are less likely, given their mileage and gas prices, to be driving cars with too little fuel efficiency.

In some cases, the gas-hybrid comparison and the regression-based analysis of all cars lead to similar results. In particular, drivers with less education are more likely to be driving vehicles with too little fuel economy, after controlling for group-level differences in mileage and gas prices. But comparisons by gender, age, and region yield different results. In the gas-hybrid comparison in Figure 7, men are less likely than women to choose a gas version when they have high individual fuel savings potential. In the regression-based analysis involving all types of cars in Figure 10, men are significantly

²⁴ Group differences in Figure 10 are relative to shares calculated using group median mpg. Differences close to zero indicate a complete lack of sorting, while more negative numbers indicate better sorting of drivers in each group into appropriate fuel economy levels.

more likely to have underinvested in fuel economy. Similarly, drivers younger than 40 are more likely to own the gas cars in Figure 7, but less likely to underinvest in fuel economy in Figure 10. And in the gas-hybrid comparison in Figure 7, rural and urban drivers are surprisingly similar. Here, rural drivers appear much more likely to be driving cars with too-little fuel economy, given their mileage and gas prices.

Of course, the two analyses are based on very different samples. Owners of vehicles with both gas and hybrid versions represent a small subset of all drivers and cars. And in both samples and across methods, we see a common pattern: there are frequent mistakes in both directions among all demographic groups, as well as a weak relationship between fuel savings potential and observed fuel economy choice.

Regression-based approach to demographic analysis: All vehicles

As with the gas-hybrid comparison, this second approach so far has examined demographic characteristics one by one. To examine how all five demographic characteristics together are associated with the fuel economy of the cars people drive, we take a regression approach in Table 6. We restrict our analysis to the 110,231 drivers in the 2017 wave of the NHTS, within which income groups are comparable.

We start by calculating the optimal fuel economy for each driver, given annual mileage and gas prices, as represented in theory by equation (7) or Figure 3. Rewrite equation (10) to include the coefficient γ that captures the degree to which consumers value future fuel savings:

$$p_{\mu} = \gamma \left(\frac{mp_g}{\mu^2} \right) F . \tag{12}$$

Equation (12) is just the all-cars version of the gas-hybrid comparison in equation (2). Solving (12) for μ yields

$$\mu = \sqrt{\gamma \frac{mp_g}{p_{\mu}} F} . \tag{13}$$

The optimal choice of fuel economy, μ^* , is just the μ in (13) where $\gamma=1$.

The first column of Table 6 regresses each driver's car's actual fuel economy (μ) on the optimal fuel economy for each driver (μ^*), given miles driven and gas prices. It controls for all five demographic characteristics, vehicle characteristics that affect fuel

economy, and a full set of make, year-by-type, and state fixed effects.²⁵ We find a statistically significant but tiny positive association between μ and μ^* . For each additional mpg that drivers should purchase, based on mileage and gas prices, their cars have actual mpg that is only 0.003 higher. There is almost no economically meaningful relationship between drivers' optimal level of fuel economy and the observed fuel economy of their vehicles.

We can again interpret these regression results as evidence for the significance of future fuel savings relative to the upfront cost of fuel efficiency. The regressions in Table 6 estimate μ as a function of μ^* , so the coefficient on μ^* can be interpreted as an estimate of $\sqrt{\gamma}$ from equation (13). Or, $\hat{\gamma} = \hat{\beta}^2$. The $\hat{\beta}$ in column 1 of Table 6 is 0.003. If we square that, it shrinks to even less relevance, suggesting that drivers all but ignore fuel savings when choosing the energy efficiency of their vehicles. Or to be more precise, drivers appear to ignore their individual cost savings from fuel efficiency. Perhaps drivers choose cars based on other vehicle attributes, though we control for some of the most salient ones. Or maybe drivers act in a way akin to internalizing the cost savings for the average driver, but certainly not their own. Either way, we find close to no relationship between individual drivers' annual mileage and the fuel economy of their cars.

Meanwhile, demographic characteristics of drivers are clearly associated with observed fuel economy levels, even after controlling for optimal fuel economy levels and vehicle specifications. As in the gas-hybrid comparison in Table 3, drivers with higher incomes own cars that are more fuel efficient. This could be for reasons other than personal financial cost, such as altruistic concern for the environment or signaling environmental sentiment to neighbors. Or maybe the distinction has financial explanations, if rich drivers have less trouble affording or borrowing to buy the more expensive, more fuel-efficient vehicles. Or maybe they are more likely to do the calculations necessary to estimate their optimal level of fuel economy.

²⁵ The optimal fuel economy, μ^* , is just level of μ that solves the equation $p_\mu = F mp_g / \mu^2$, where F is the multiplication factor 11.3 corresponding to cumulative savings over 14 years, discounted at 3 percent and the price of fuel economy is \$110 per additional mpg.

Column 2 of Table 6 examines whether high-income households are more responsive to their optimal fuel economy. It includes in the regression an interaction between the optimal level of fuel economy (μ_i^*) and an indicator for households with annual incomes above \$50,000. The small and insignificant coefficient (0.0018) suggests that the observed fuel economy choices of high-income drivers are not much more responsive to the economically optimal level, if at all. Rich households are not significantly more likely to choose a vehicle's fuel economy based on financial calculations of mileage and gas prices. Higher-income drivers do on average choose higher levels of fuel economy, just as they are more likely to drive hybrids, but evidently not for reasons related to fuel savings. There is close to no meaningful association between the cost-minimizing level of fuel economy and observed choices for either rich or poor households. As with the gas-hybrid comparison, rich households are not more likely to do the math, and poor households are not more likely to be credit constrained.²⁶

V. Possible explanations for lack of attention to energy efficiency

Our analysis finds widespread apparent mistakes in both directions: many households don't drive sufficient miles but still own relatively fuel-efficient vehicles, and vice versa. What can explain this seeming absence of economic calculation? We have already seen, via the interaction terms in Table 3 and Table 6, that the responsiveness of fuel economy choices to economic incentives does not vary by income. That rules out credit constraints and other household characteristics correlated with income. Below we investigate further explanations: mistaken expectations about driving costs and systematic bias in vehicle prices.

²⁶ In the Appendix tables, we estimate a version of Table 6 that replaces the dependent variable with the difference between observed mpg and optimal mpg. That specification asks which households end up with a level of fuel economy that, given their driving, is further away from the economically optimal level. We estimate that drivers with higher incomes and more education, drivers under 40 years of age, male drivers, and drivers in rural areas own vehicles that have lower fuel economy relative to the level that would be economically optimal.

Possible differences between realized ex post mileage and ex ante expectations

Our main analyses in Sections III and IV rely on fuel expenses calculated using data on actual miles driven in the NHTS survey years, 2009 and 2017. But this may differ from how much drivers expected they would drive at the time they bought their cars. Although vehicle choices may seem to indicate mistakes ex post, they could have been in line with what car buyers expected. We repeat our analysis using alternative data for vehicle owners' expected annual mileage at the time of purchase, using data from the MaritzCX monthly surveys of vehicle buyers, conducted between 2010 and 2017.²⁷

The time-of-purchase surveys average a 9 percent response rate and represent approximately 1 percent of US car purchases. As with the NHTS data, we match those vehicles to characteristics from WardsAuto. The time-of-purchase sample differs a bit from the post-purchase NHTS data in Table 1. The time-of-purchase survey respondents are wealthier and more educated, more urban, and own larger, more powerful, and slightly less fuel-efficient cars. The vehicles in the time-of-purchase data are also newer—cars purchased between 2009 and 2017 rather than cars driven in 2009 and 2017.

Despite those differences, the time-of-purchase survey has two distinctions that make it a useful comparison with our post-purchase evidence. First, annual mileage comes from responses to a survey question about how many miles the car buyer intends to drive each year. That time-of-purchase expectation eliminates concerns that the NHTS calculations do not display a relationship between fuel economy and driving because car buyers inaccurately forecast their future driving. Second, the time-of-purchase data report the actual transacted sales price rather than the MSRP.

As a first step, we replicated Figure 1 using annual cost differences between hybrid and equivalent gas cars, based on expected miles at the time of purchase. The two distributions are indistinguishable, and so we have not reproduced the time-of-purchase version here.²⁸ Hybrid and gas purchasers expect to drive the same amount, supporting the finding that hybrids are not chosen for their expected fuel savings, and that the post-

²⁷ See Leard, Linn, and Zhou (2017), and Leard, Linn, and Springel (2019, 2020). Appendix Table 1 contains summary statistics for the MaritzCX data, the analog to Table 1. MaritzCX was acquired by another market research company, InMoment, in March 2020.

²⁸ See Appendix Figure 1.

purchase differences seen in Figure 1 are the result of rebound—drivers’ reactions to the hybrids’ low driving cost—not selection.

Second, in Panel A of Table 7 we have reproduced the evaluation of gas-hybrid choices from Table 2 but with time-of-purchase expected miles rather than post-purchase actual miles. The results are qualitatively similar, though a larger share of the cars are hybrids. Much of that difference is surely due to the growth in hybrid market share, but some may stem from the sample selection, for example if hybrid purchasers are more likely to respond to this survey.

Of 59,671 purchasers of cars that came in gas and hybrid versions, our calculations suggest that 23,111 would be better off choosing a hybrid, given their expected mileage, the sales price premium over the gas version, a 14-year life expectancy, and a 3 percent discount rate. Nearly as many do, 20,344. But as with the post-purchase driving evidence in Table 2, it’s mostly the wrong car buyers. Of the 23,111 who would better off purchasing a hybrid, only 8,924 do (39 percent). And as in Table 2, switching to a 7 percent discount rate dramatically drops the estimated number who would be better off choosing a hybrid, but it barely nudges the fraction of those who actually do (40 percent).

Table 7 also addresses a second concern about our main analyses in Sections III and IV: our use of MSRPs reported in WardsAuto to calculate the upfront premium paid for a more efficient vehicle. That MSRP may impart a bias if hybrid cars sell at a larger discount relative to MSRPs. The MaritxCX survey data contain the actual transacted sales prices. The qualitative results in Table 7 using sales prices are similar to those in Table 2 using MSRP—around 60 to 70 percent of drivers own gas vehicles, whether that appears optimal given their fuel expenses or not.

As a third step, in Table 8 we replicate Table 3 with the time-of-purchase data, estimating linear probability models of hybrid choice as a function of expected fuel savings, and using actual vehicle sales prices rather than MSRPs. The results are similar. The coefficients on fuel savings are slightly larger, and the coefficients on upfront costs smaller, so the implied γ is a little larger. But as before, other demographics play a larger role than fuel savings in determining hybrid choice. A graduate degree still has 10 times the predictive power as \$1,000 of savings. Being older than 60 has seven times the

predictive power, and earning more than \$150,000 instead of less than \$50,000 per year has 20 times the predictive power. Drivers' demographics, rather than expected fuel expenses, predict fuel economy choices.

Finally, in the last two columns of Table 8 we replicate Table 6 using the time-of-purchase data. Column 6 regresses actual fuel economy (μ) on our calculated optimal fuel economy (μ^*) for each driver, given expected annual mileage and local gas prices. As with the post-purchase NHTS data in Table 6, here we also find a statistically significant but tiny positive association between μ and μ^* . Each additional mpg that would be optimal is associated with actual mpg purchased that is only 0.007 higher. Squaring that coefficient to recover $\hat{\gamma}$, it shrinks to nothing, indicating that car buyers ignore their own expected future fuel savings when purchasing vehicles.

Again, demographic characteristics of car buyers have a much larger association with the fuel economy of the vehicles they purchase than their expected fuel savings. Columns 2 and 4 of Table 8 add the interaction with income. Rich households are not more likely to be calculating fuel savings, and poor households do not seem more credit constrained.

Possible difference in depreciation rates and used car prices

The main analyses in Sections III and IV explore whether, retrospectively, given their actual mileage in 2009 and 2017, drivers would have been better off had they purchased more efficient or less efficient cars. Above, we pose a different hypothetical and ask whether, prospectively, given their expected future mileage, car buyers would be better off by purchasing more efficient or less efficient cars. But there's a third way to frame the question, using data on used-car prices. Every driver who drives a lot of miles in an inefficient vehicle, like Albert in our introduction, has the option, in theory, to trade in his inefficient car for a more efficient one. The fuel savings would justify the upgrade. And every driver who owns an efficient vehicle but doesn't drive much, like Betty, could, in theory, downgrade her fuel economy and pocket the difference.

More generally, not all cars are purchased new and not all cars are held by the same owner for their entire lifespans. Any systematic difference in depreciation rates and

used-car prices between cars with different fuel economy levels could be responsible for at least some of the absence of sorting based on fuel savings potential that we find.

To frame the question this way, as a choice to keep or trade in one's current car, and to account for possible differences in depreciation rates, we complement our analysis with used-car prices for about 1 million vehicles listed in September 2017 on the website TrueCar.com.²⁹ We match those prices to each make, model, and year for vehicles in the 2017 NHTS. Only a relatively small subset of the relevant models had a matching listing on TrueCar.com, but even with that limited sample, we again see similar patterns. Panel B of Table 7 replicates Table 2 based on used-car prices and suggests two insights. First, more drivers would be better off in hybrids, presumably because the used hybrid price premium is smaller than the new one. Second, optimal and observed choices are again only weakly related. In other words, differences in depreciation and resale value cannot explain the observed lack of sorting into what looks like the cost-minimizing vehicle choice based on individual drivers' circumstances.

VI. Conclusion

Do drivers undervalue future fuel savings when they buy cars? Or do they appropriately balance the additional upfront cost of investing in fuel economy with the cumulative future fuel savings? Previous work in this area finds mixed evidence but suggests that vehicle prices and market shares respond to the cost savings from higher levels of fuel economy, at least partially and on average. We instead ask a related but somewhat different question: do individual drivers respond to their personal savings potential from higher levels of fuel economy based on their driving behavior and the fuel prices they face? We find very little evidence that they do.

Of the more than 150,000 car owners in the 2009 and 2017 NHTS survey waves, we find that a large portion have made what appear to be mistakes. Some own inefficient cars but drive them a lot, while others paid a premium for fuel-efficient cars that they drive very little. This disconnect between individual fuel economy and driving holds

²⁹ The used-car listings data from TrueCar.com are available under "CCO: Public Domain" license from the following website: <https://www.kaggle.com/jpayne/852k-used-car-listings>.

whether we look at ex post realized mileage or ex ante expected mileage at the time of purchase. It holds whether we use MSRPs or actual purchase prices when the cars were new, or used-car prices in 2017. And it holds whether we use the price premium for hybrid versions of car models that also come in a standard gas version or estimate the price premium for an extra unit of fuel economy, controlling for other car features statistically.

Our findings contain a puzzle: how can the market as a whole seemingly internalize the value of fuel economy while individual consumers mostly ignore it? One possible explanation might be that consumers respond to average incentives—they act as if they expect to drive the average number of miles—and ignore their individual deviation from the mean. Another explanation could be that fuel economy plays a less important role in vehicle choice than previously suggested.

Whatever the explanation, the observed lack of sorting into levels of fuel economy based on individual economic incentives suggests a significant, previously overlooked source of market inefficiency. Fuel policy discussions revolve around improving fleet-wide fuel economy levels, as with the CAFE standards. But we find many drivers already own cars for which the upfront premium they paid for fuel economy exceeds their annual savings. Any fleet-wide policy would make those drivers worse off.

Instead of a fleet-wide policy, more targeted policies might achieve the same fuel-conservation objectives and make more drivers better off. Consider two stylized alternatives: (1) increase every car's fuel economy by one mpg, or (2) keep all the cars the same but swap them among drivers so that the cars driven more than the median annual mileage have one mpg of additional fuel economy, and the cars driven less than the median annual mileage have one mpg less. In the 2017 NHTS sample, a fleet-wide increase in fuel economy of one mpg would save the average driver about \$50 per year. Merely swapping cars among drivers to increase high-mileage cars' efficiency by one mpg would save an average of \$30 per driver per year. In other words, better sorting at constant technologies achieves 60 percent of the savings from a fleet-wide fuel economy improvement. Banzhaf and Kasim (2019) show that higher fuel prices may trigger such sorting, but there may also be other policies to do so. Our analysis shows that this would have significant savings potential.

Although financial incentives linked to fuel costs seem to play a minor role in drivers' vehicle choices, we identify important differences across demographic groups. Why this is important depends on how we interpret the apparent mistakes. If the choices reflect underlying unobserved preferences, then perhaps richer people have a higher willingness to pay for hybrids for some reason unrelated to fuel savings. And perhaps poorer drivers have higher discount rates, preferring to pay less up front in exchange for higher future costs. A policy that subsidizes hybrid cars would be an inframarginal transfer to rich drivers. A policy that penalizes gas-powered cars would impose a regressive burden on poorer drivers.

On the other hand, if the observed pattern is the result of some market failure—say, the liquidity constraints faced by poorer households—then the policy implication is different. A subsidy for efficient cars could enable poorer households to afford their upfront costs and save them future fuel expenses. Our analysis suggests that drivers' failure to consider fuel costs does not appear to be due either to financial constraints or to erroneous expectations about future driving.

References

- Allcott, Hunt. 2011. Consumers' perceptions and misperceptions of energy costs. *American Economic Review* 101(3): 98–104.
- Allcott, Hunt, and Christopher Knittel. 2018. Are consumers poorly informed about fuel economy? Evidence from two experiments *American Economic Journal: Economic Policy* 11(1): 1–37.
- Allcott, Hunt, and Michael Greenstone. 2012. Is there an energy efficiency gap? *Journal of Economic Perspectives* 26(1): 3–28.
- Allcott, Hunt, and Nathan Wozny. 2014. Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96(5): 779–95.
- Ankney, Kevin. 2020. Do credit constraints explain the energy efficiency Gap? Working paper. Georgetown University, Washington, DC.
- Banzhaf, H. Spencer, and M. Taha Kasim. 2019. Fuel consumption and gasoline prices: The role of assortative matching between households and automobiles. *Journal of Environmental Economics and Management* 95: 1-25.
- Bento, Antonio M., Shanjun Li, and Kevin Roth. 2012. Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias. *Economics Letters* 115(1): 44–48.
- Bento, Antonio M., Kenneth Gillingham, Mark R. Jacobsen, Christopher R. Knittel, Benjamin Leard, et al. 2018. Flawed analyses of U.S. auto fuel economy standards. *Science* 362(6419): 1119–21.
- Bento, Antonio, Mark Jacobsen, Christopher Knittel, and Arthur van Benthem. 2019. Estimating the costs and benefits of fuel-economy standards. Working paper 26309. National Bureau of Economic Research, Cambridge, MA.
- 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): 220–56.
- Gillingham, Kenneth, David Rapson, and Gernot Wagner. 2016. The rebound effect and energy efficiency policy. *Review of Environmental Economics & Policy* 10(1): 68–88.
- Gillingham, Kenneth, Sébastien Houde, and Arthur van Benthem. 2019. Consumer myopia in vehicle purchases: Evidence from a natural experiment. Working paper 25845. National Bureau of Economic Research, Cambridge, MA.
- Greene, David L. 2010. How consumers value fuel economy: A literature review. US Environmental Protection Agency, Oak Ridge, TN. March.
- Greene, David L., Hossain, Anushah, Hofmann, Julia, Helfand, Gloria, and Robert Beach. 2018. Consumer willingness to pay for vehicle attributes: What do we know? *Transportation Research Part A: Policy and Practice* 118: 258–79.

- Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2018. Consumer valuation of fuel costs and tax policy: Evidence from the European car market. *American Economic Journal: Economic Policy* 10(3): 193–225.
- Helfand, Gloria, and Ann Wolverton. 2011. Evaluating the consumer response to fuel economy: A review of the literature. *International Review of Environmental and Resource Economics* 5(2): 103–46.
- Labandeira, Xavier, José M. Labeaga, and Xiral López-Otero. 2017. A meta-analysis on the price elasticity of energy demand. *Energy Policy* 102: 549–68.
- Langer, Ashley, and Shaun McRae. 2017. Step on it: A new approach to improving vehicle fuel economy. Working paper. University of Arizona, Tucson.
- Leard, Benjamin, Joshua Linn, and Katalin Springel. 2019. Pass-through and welfare effects of regulations that affect product attributes. Resources for the Future, Washington, DC.
- . 2020. Have US Fuel Economy and Greenhouse Gas Emissions Standards Improved Social Welfare? Report. Resources for the Future, Washington, DC.
- Leard, Benjamin, Joshua Linn, and Christy Zhou. 2017. How much do consumers value fuel economy and performance? Evidence from technology adoption. Report. Resources for the Future, Washington, DC.
- Levinson, Arik. 2019. Energy efficiency standards are more regressive than energy taxes: Theory and evidence. *Journal of the Association of Environmental and Resource Economists*, 6(S1): S7-S36.
- Knittel, Christopher R., and Shinsuke Tanaka. 2019. Driving behavior and the price of gasoline: evidence from fueling-level micro data. Working paper 26488. National Bureau of Economic Research, Cambridge, MA.
- Sallee, James. 2014. Rational inattention and energy efficiency. *Journal of Law and Economics* 57(3): 781–820.
- Sallee, James, Sarah West, and Wei Fan. 2016. Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics* 134: 61–73.
- Sexton, Steven E., and Alison L. Sexton. 2014. Conspicuous conservation: The Prius Halo and willingness to pay for environmental bona fides. *Journal of Environmental Economics and Management* 67(3): 303–17.

Table 1. Summary statistics

	<u>Gas-hybrid pairs</u>		All cars (3)
	Gas (1)	Hybrid (2)	
Annual miles	12,350 (74.55)	13,251 (219.0)	12,428 (29.01)
Combined city-highway mpg	27.72 (0.0233)	37.67 (0.149)	25.72 (0.0154)
Household income (2017 \$)			
Less than \$25k	0.0850	0.0345	0.0891
\$25k–\$50k	0.186	0.102	0.181
\$50k–\$75k	0.199	0.129	0.187
\$75k–\$100k	0.165	0.165	0.161
\$100k–\$150k	0.217	0.271	0.216
More than \$150k	0.148	0.299	0.166
Age	52.73 (0.118)	54.58 (0.311)	53.22 (0.042)
Education:			
High school	0.169	0.0980	0.184
Some college	0.544	0.485	0.550
Graduate	0.242	0.380	0.235
Rural	0.213	0.199	0.245
Length	186.4 (0.0442)	186.0 (0.138)	186.7 (0.0312)
Width	71.63 (0.0136)	71.59 (0.0408)	72.60 (0.00868)
Height	59.80 (0.0325)	60.55 (0.113)	62.54 (0.0156)
Weight	3,279 (4.068)	3,630 (11.64)	3,527 (1.686)
Liters	2.417 (0.00355)	2.464 (0.0180)	2.776 (0.00217)
Valves	3.949 (0.00213)	3.914 (0.00837)	3.793 (0.00153)
Horsepower	171.9 (0.230)	162.4 (1.025)	195.6 (0.151)
Rpm	6,014 (1.753)	5,850 (4.756)	5,927 (1.280)
Observations	22,025	2,337	155,572

Sources: 2009 and 2017 NHTS, limited to households with cars and complete demographic information. Vehicle characteristics from Wardsauto.com. 24,362 gas-hybrid models in columns 1 and 2, and 155,572 observations for all cars in column 3. The income variables in columns 1–3 are only for the 2017 NHTS. Standard errors in parentheses.

Table 2. Two types of mistakes: Models with gas-hybrid choice

	Actual vehicle	Optimal vehicle (3% discount rate, 14-year life)		Optimal vehicle (7% discount rate, 14-year life)	
		Total (1)	Gas (2)	Hybrid (3)	Gas (4)
Total	24,362	20,242	4,120	21,932	2,430
Gas	22,025	18,379	3,646	19,881	2,144
(% of row)		(83%)	(17%)	(90%)	(10%)
(% of column)		(91%)	(88%)	(91%)	(88%)
Hybrids	2,337	1,863	474	2,051	286
(% of row)		(80%)	(20%)	(88%)	(12%)
(% of column)		(9%)	(12%)	(9%)	(12%)

Sources: 2009 and 2017 NHTS. Model years 2005–2017. All cars with complete data on prices and incomes and that come in gas and hybrid versions. The shaded boxes represent “mistakes”—drivers who would be financially better off in the alternative version of the same vehicle. Calculations assume a 14-year lifetime with discount rates of 3% and 7%. Note: The 14-year lifetime roughly replicates the CAFE 2017–2025 RIA. EPA and NHTSA use model year–specific lifetime estimates, which average 14 years.

Table 3. Characteristics of hybrid vehicle drivers

		NHTS data			
		“Low” cutoff (3% discount, ½ price)		“High” cutoff (7% discount, 2× price)	
Dependent variable = 1 if hybrid		(1)	(2)	(3)	(4)
Cumulative fuel savings (\$1,000)		0.0038*	0.0023*	0.0049*	0.0029*
		(0.0005)	(0.0010)	(0.0007)	(0.0013)
Fuel savings ×(Income>\$50,000)			0.0020 (0.0011)		0.0026 (0.0014)
Upfront investment cost (\$1,000)		-0.0429*	-0.0429*	-0.0107*	-0.0107*
		(0.0021)	(0.0021)	(0.0005)	(0.0005)
Income:	\$25k–\$50k	0.0140*	0.0144*		
		(0.0077)	(0.0070)		
	\$50k–\$75k	0.0170*	0.0108		
		(0.0069)	(0.0077)		
	\$75k–\$100k	0.0427*	0.0364*		
		(0.00727)	(0.0080)		
	\$100k–\$150k	0.0484*	0.0419*		
		(0.0070)	(0.0078)		
	Over \$150k	0.0841*	0.0776*		
		(0.0075)	(0.0084)		
Education:	Some college	0.0099*	0.0101*		
		(0.0048)	(0.0048)		
	Graduate	0.0332*	0.0334*		
		(0.0054)	(0.0054)		
Age:	40–60 years	0.0175*	0.0172*		
		(0.0043)	(0.0043)		
	Over 60 years	0.0287*	0.0283*		
		(0.0043)	(0.0043)		
Male		0.00292	0.00293		
		(0.0034)	(0.0034)		
Rural		0.0074	0.0074		
		(0.0043)	(0.0043)		
Includes car characteristics: length, width, height, weight, liters, valves, horsepower, and rpm					
Includes fixed effects by state, year-by-type, and make					
Implied $\hat{\gamma}$		0.09		0.46	
($\hat{\gamma}$ for income < \$50,000)			0.05		0.27
($\hat{\gamma}$ for income > \$50,000)			0.10		0.52
Observations		17,586	17,586	17,586	17,586
R-squared		0.37	0.37	0.37	0.37

Sources: See Table 1. Calculations for cumulative (and discounted) fuel savings assume a 14-year lifetime with discount rates of 3% and 7%. Sample using the 17,586 observations from the 2017 NHTS wave (because of different income classification in 2009 wave). * p<0.05.

Table 4. Fuel economy and vehicle prices

Dependent variable =	Means		Regressions		
	(1)	(2)	(3)	(4)	(5)
MSRP					
mpg μ	22.4 (5.5)	-1,420* (32)	197.0* (37.3)	339.7* (33.1)	109.9* (37.6)
Length (ins.)	192.4 (20.8)		-214.0* (9.0)	-99.3* (8.6)	17.7 (11.2)
Width (ins.)	73.7 (4.0)		-231.6* (58.4)	-214.1* (52.7)	237.4* (70.0)
Height (ins.)	64.7 (7.7)		-794.6* (31.0)	-235.5* (30.2)	-271.8* (64.7)
Weight (lbs.)	3,988 (855.5)		9.50* (0.42)	6.02* (0.39)	5.06* (0.43)
Liters	3.38 (1.22)		-3,767* (258)	-1,752* (241)	-2,436* (245)
Valves	3.64 (0.76)		271 (221)	-1,295* (224)	-1,898* (262)
Horsepower	248.7 (84.6)		227.5* (3.1)	170.2* (2.9)	158.8* (2.8)
Rpm	5,844 (603.8)		-4.12* (0.24)	-2.67* (0.23)	-2.86* (0.23)
Observations		14,789	14,789	14,789	14,789
R-squared		0.120	0.587	0.712	0.811
Make fixed effects		No	No	Yes	Yes
Model fixed effects		No	No	No	Yes

Sources: Authors' calculations from Wardsauto.com data. Prices are MSRP. Mpg is EPA's combined for city and highway driving. Each observation is a single make, model, and year combination.

Table 5. Two types of mistakes: All models

	Fuel economy level (3% discount rate, 14-year life, \$110 per mpg)			Fuel economy level (7% discount rate, 14-year life, \$340 per mpg)	
	Total (1)	Too much (2)	Too little (3)	Too much (4)	Too little (5)
NHTS 2009	45,291	2,717 (6%)	42,574 (94%)	17,599 (39%)	27,692 (61%)
NHTS 2017	110,281	9,931 (9%)	100,350 (91%)	55,959 (51%)	54,322 (49%)
Total NHTS	155,572	12,648 (8%)	142,924 (92%)	73,558 (47%)	82,014 (53%)
Avg. forgone savings (14 years)		\$417	\$6,008	\$1,294	\$2,383

Sources: See Table 1. Calculations assume a 14-year lifetime with discount rates of 3% and 7%. The cumulative (and discounted) monetary value of one more unit of mpg (according to equation 12) is compared with the marginal cost of mpg estimated in Table 4, columns 4 and 5 (\$110 and \$340). Forgone savings are the cumulative discounted cost difference over the 14-year vehicle lifetime relative to the cost-minimizing choice according to equation (13). The 14-year lifetime roughly replicates the CAFE 2017–2025 RIA. EPA and NHTSA use model-year-specific lifetime estimates, which average 14 years.

Table 6. Determinants of individual fuel economy choice

Dependent variable: mpg (μ)	(1)	(2)
Optimal mpg μ_i^*	0.0030*	0.0017
	(0.0005)	(0.0009)
(μ_i^*)		0.0018
×(Income > \$50k)		(0.0011)
Income: \$25k–\$50k	0.101*	0.103*
	(0.039)	(0.039)
\$50k–\$75k	0.199*	0.112
	(0.039)	(0.064)
\$75k–\$100k	0.336*	0.248*
	(0.041)	(0.065)
\$100k–\$150k	0.436*	0.347*
	(0.040)	(0.065)
Over \$150k	0.580*	0.490*
	(0.043)	(0.067)
Education: Some college	0.107*	0.107*
	(0.027)	(0.027)
Graduate	0.420*	0.420*
	(0.032)	(0.032)
Age: 40–60 years	0.0422	0.0410
	(0.025)	(0.025)
Over 60 years	0.094*	0.0922*
	(0.025)	(0.025)
Male	0.038	0.039
	(0.020)	(0.020)
Rural	0.0586*	0.0584*
	(0.0238)	(0.024)
Other car covariates: length, width, height, weight, liters, valves, horsepower		
Observations	110,231	110,231
R-squared	0.72	0.72

Sources: See Table 1. Calculations for optimal mpg assume a 14-year lifetime with a discount rate of 3%. The cumulative (and discounted) monetary value of one more unit of mpg (according to equation 12) is compared with the marginal cost of mpg estimated in Table 4, column 5 (\$110). Dependent variable is observed mpg. See Appendix Tables A2 and A3 for full results. * $p < 0.05$.

Table 7. Time-of-purchase and used-car price analysis: Models with gas-hybrid choice

	Actual vehicle Total (1)	Optimal vehicle (3% discount rate, 14-year life)		Optimal vehicle (7% discount rate, 14-year life)	
		Gas (2)	Hybrid (3)	Gas (4)	Hybrid (5)
<i>Panel A: Time-of-purchase analysis (MaritzCX)</i>					
Total	59,671	36,260	23,111	42,963	16,078
Gas (% of row) (% of column)	39,327	24,840 (63%) (69%)	14,187 (36%) (61%)	29,251 (74%) (68%)	10,076 (26%) (60%)
Hybrids (% of row) (% of column)	20,344	11,420 (56%) (31%)	8,924 (44%) (39%)	13,712 (67%) (32%)	6,632 (33%) (40%)
<i>Panel B: Used price analysis (NHTS + TrueCar.com)</i>					
Total	3,254	921	2,333	1,162	2,092
Gas (% of row) (% of column)	2,999	874 (29%) (95%)	2,125 (71%) (91%)	1,098 (37%) (94%)	1,901 (63%) (91%)
Hybrids (% of row) (% of column)	255	47 (18%) (5%)	208 (82%) (9%)	64 (25%) (6%)	191 (75%) (9%)


Sources: Panel A: MaritzCX survey years 2010–2017, vehicle characteristics from Wardsauto.com. Panel B: NHTS 2017 survey, limited sample with available used-car prices from TrueCar.com. All cars with complete data on prices and incomes and that come in gas and hybrid versions. The shaded boxes represent “mistakes”—drivers who would be financially better off in the alternative version of the same vehicle. Calculations assume a 14-year lifetime with discount rates of 3% and 7%. Note: The 14-year lifetime roughly replicates the CAFE 2017–2025 RIA. EPA and NHTSA use model-year-specific lifetime estimates, which average 14 years.

Table 8. Time-of-purchase fuel economy choice

	Dependent variable = 1 if hybrid		Dependent variable: mpg	
	“Low” cutoff (3% discount, ½ price) (1)	“High” cutoff (7% discount, 2× price) (2)	(3)	(4)
Cumulative fuel savings (\$1000)	0.00446* (0.00124)	0.00576* (0.00160)		
Fuel savings ×(Income>\$50,000)	-0.00268* (0.00127)	-0.00346* (0.00164)		
Optimal mpg μ_i^* (μ_i^*) ×(Income > \$50k)			0.00715* (0.000271)	0.00726* (0.000719) -0.000120 (0.000758)
Upfront investment cost (\$1000)	-0.0180* (0.00138)	-0.00451* (0.000345)		
Income:				
\$25k–\$50k	0.00783 (0.0150)		0.161* (0.0607)	0.162* (0.0607)
\$50k–\$75k	0.0401* (0.0158)		0.381* (0.0605)	0.388* (0.0748)
\$75k–\$100k	0.0609* (0.0158)	<i>coefficients same as column (1)</i>	0.490* (0.0606)	0.497* (0.0750)
\$100k–\$150k	0.0800* (0.0158)		0.603* (0.0603)	0.610* (0.0748)
Over \$150k	0.102* (0.0158)		0.793* (0.0606)	0.800* (0.0751)
Education:				
Some college	0.0160* (0.00339)		0.142* (0.0136)	0.142* (0.0136)
Graduate	0.0432* (0.00370)		0.463* (0.0152)	0.463* (0.0152)
Age:				
40–60 years	0.0128* (0.00286)		0.118* (0.0114)	0.118* (0.0114)
Over 60 years	0.0322* (0.00298)		0.294* (0.0122)	0.294* (0.0122)
Male	0.0175* (0.00216)		0.0134 (0.00932)	0.0134 (0.00932)
Rural	-0.00652* (0.00317)		-0.0310* (0.0126)	-0.0310* (0.0126)
Includes car characteristics: length, width, height, weight, liters, valves, horsepower, and rpm				
Includes fixed effects by state, year-by-type, and make				
Observations	59,671	59,671	607,024	607,024
R-squared	0.736	0.736	0.752	0.752

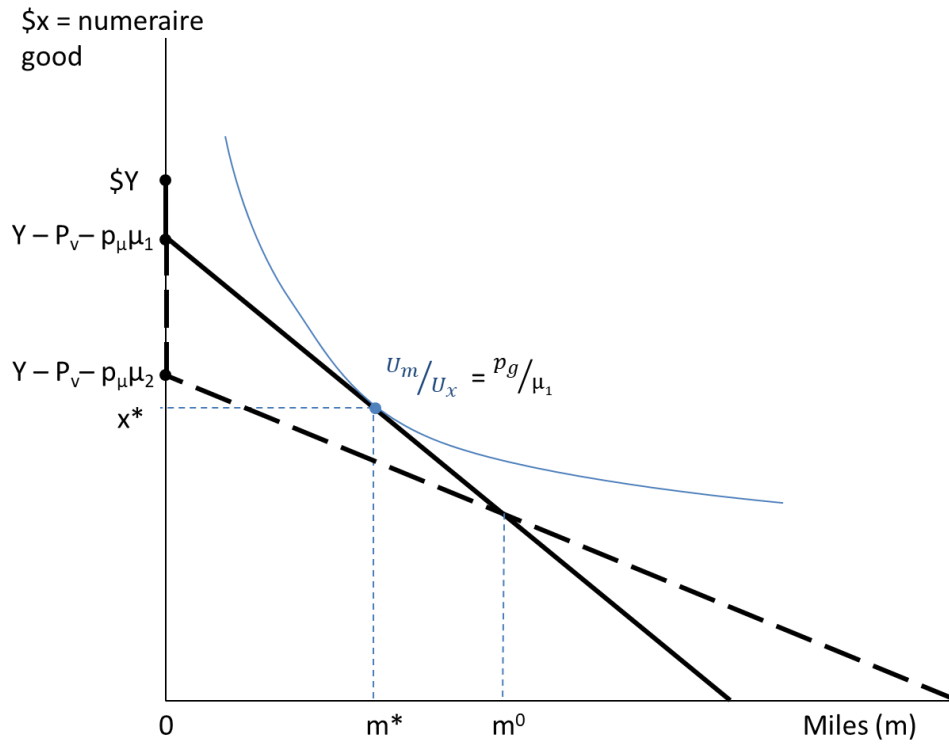
Sources: MaritzCX survey years 2010–2017, vehicle characteristics from Wardsauto.com. Calculations for cumulative (and discounted) fuel savings in columns 1 and 2 assume a 14-year lifetime with discount rates of 3% and 7%. Calculations for optimal mpg in columns 3 and 4 assume a 14-year lifetime with a discount rate of 3%. The cumulative (and discounted) monetary value of one more unit of mpg (according to equation 12) is compared with the marginal cost of mpg estimated in Table 4, column 5 (\$110). See Appendix Tables A2 and A4 for full results. * p<0.05.

Figure 1. Annual fuel cost differences, gas versus hybrid



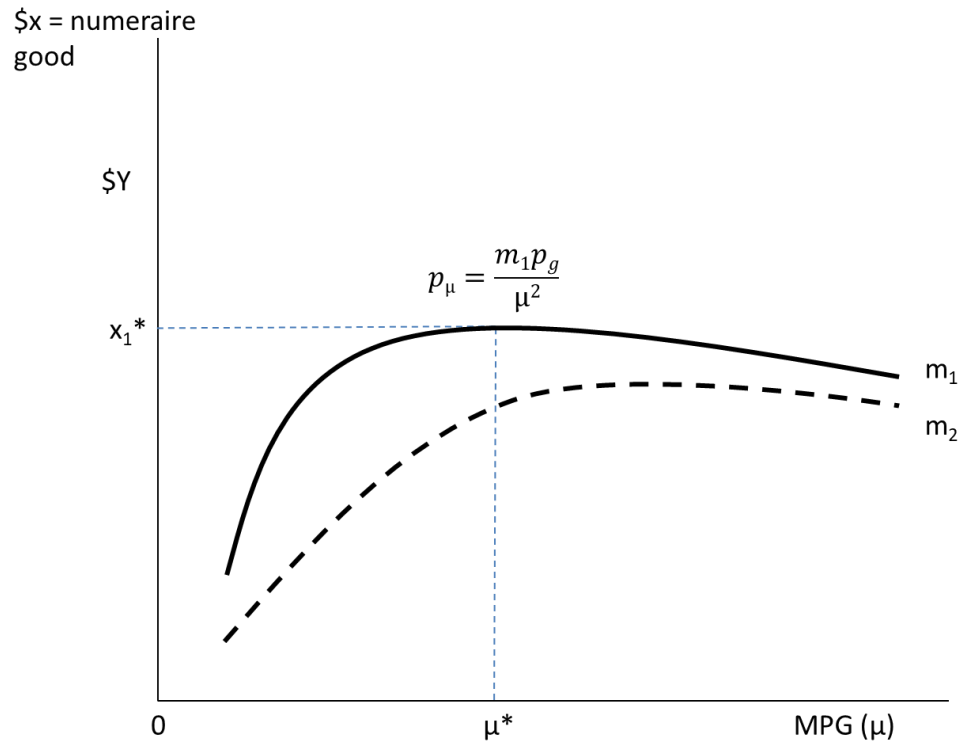
Notes: For each vehicle in the 2009 and 2017 NHTS that comes in both gas and hybrid versions, we calculate the annual fuel cost difference between the two versions, given mileage and gas prices. The shaded bars represent the cost savings for hybrid drivers relative to driving a gas version of the same car; the outlined bars represent the extra annual costs borne by gas drivers.

Figure 2. One view of the budget constraint




Notes: Illustration of the budget trade-off between miles driven (m) and other consumption (x), comparing a less efficient vehicle purchased at a lower price (solid budget line) and a more efficient vehicle purchased at a higher price (dashed budget line).

Figure 3. A different view of the budget constraint




Notes: Graph illustrates the optimal choice of fuel economy (MPG) that minimizes fuel expenditure and thus maximizes consumption of the numeraire (x), holding miles driven (m) constant. Comparison between a driver with low mileage (solid budget line) and one with higher mileage (dashed budget line).

Figure 4. Benchmark annual savings, gas versus hybrid



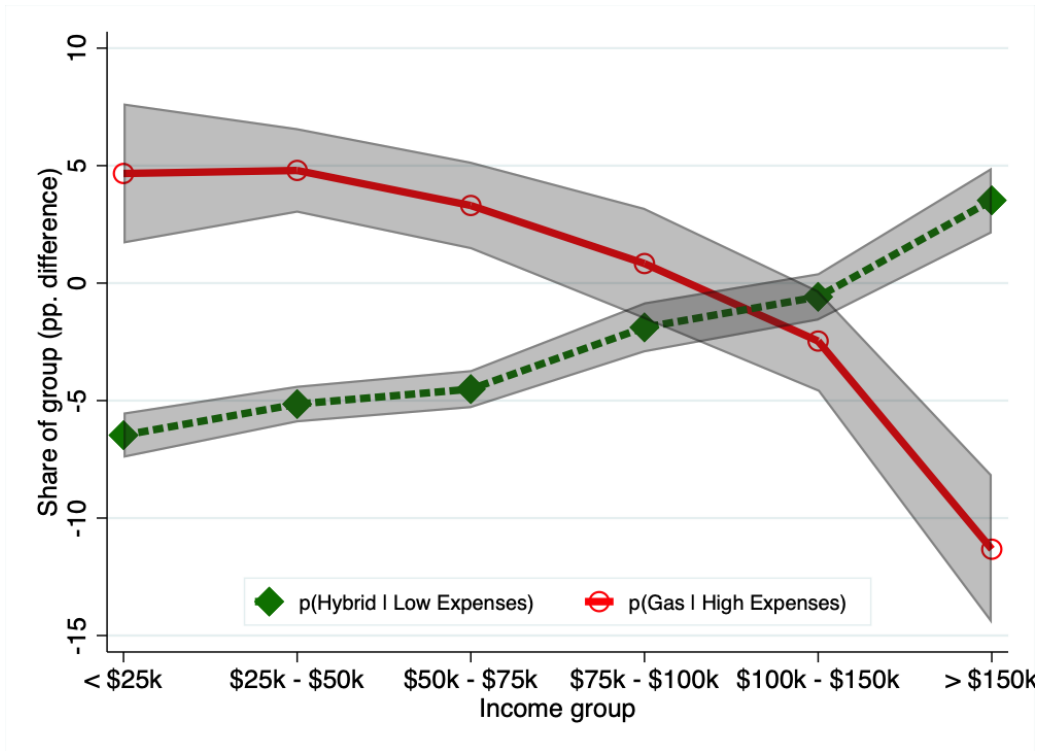
Notes: Graphs illustrate a cutoff value in counterfactual annual fuel cost savings for switching from a gas to a hybrid version of the same car model. These savings are forgone by owners of gas vehicles (left) and realized by owners of hybrid vehicles (right). Hybrid owners saving less than the cutoff (area C) would have saved money by buying the gas version; owners of gas vehicles with savings exceeding the cutoff (area B) would have saved money by buying the hybrid.

Figure 5. Car-buying mistakes, by income




Notes: Based on the 2017 NHTS, a 7% discount rate, and a price of efficiency (p_{μ}) of \$340. The top line plots the share of high-expense drivers, who would be financially better off in hybrids but are driving gas cars. The bottom line plots the share of low-expense drivers in hybrids. Shaded areas represent 95% confidence intervals of the shares.

Figure 6. Car-buying mistakes, by income and adjusted for market shares




Notes: A version of Figure 5, adjusted for the relative difference in market shares between gas and hybrid models. The downward-sloping line plots the share of high-expense drivers in gas cars, net of the overall market share of gas cars. The upward-sloping line plots the share of low-expense drivers in hybrids, net of the market share of hybrids. Shaded areas represent 95% confidence intervals of the shares.

Figure 7. Car-buying mistakes, by other demographics




Notes: Data from 2009 and 2017 NHTS. The open circles plot the share of each demographic with high expenses who would be financially better off in hybrids but are driving gas cars, minus the overall market share of gas cars. The solid diamonds plot the share of each group with low expenses driving hybrids, minus the market share of hybrids. Lines represent 95% confidence intervals of the shares.

Figure 8. Probability of having too little fuel economy, by income



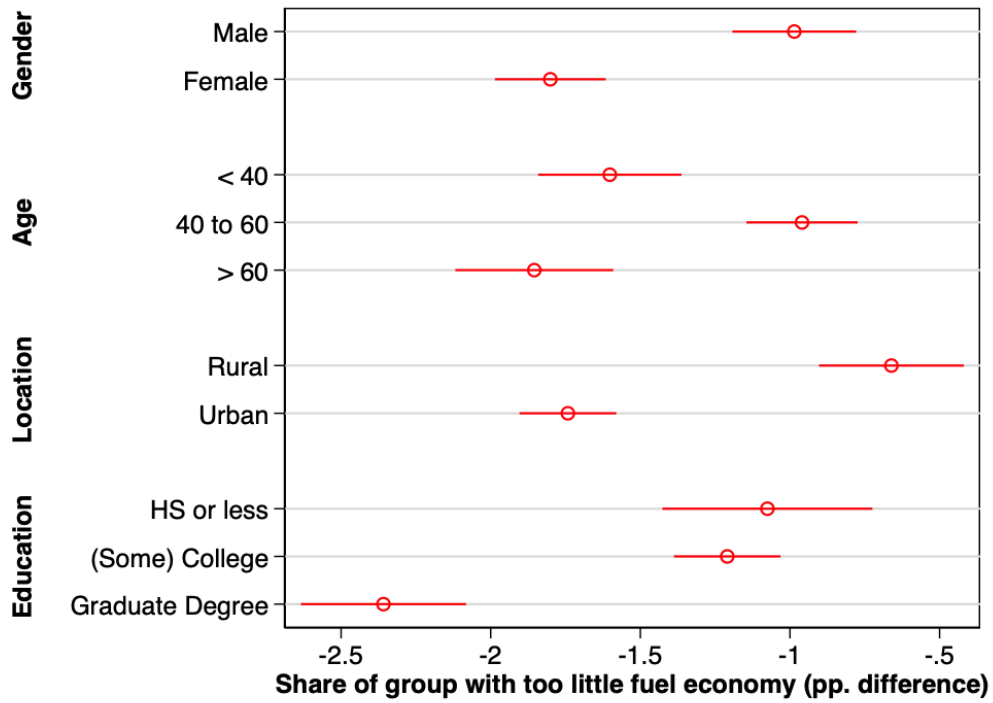
Notes: Data from 2017 NHTS. Each point represents the share of each income group that could be financially better off by driving a car with one more mpg of fuel economy. Assumes a 7% discount rate and a price of efficiency (p_{μ}) of \$340. Shaded areas represent 95% confidence intervals of the shares.

Figure 9. Probability of having too little fuel economy, accounting for mileage and gas price



Notes: A version of Figure 8. Actual shares of each group deemed to have underinvested in fuel economy, minus what that share would be if all drivers were assigned the median fuel economy. Those counterfactual shares depend only on mileage and gas prices, so the probabilities reported here are versions of those in Figure 8, adjusted for those income-group differences. Shaded area represents 95% confidence intervals of the shares.

Figure 10. Probability of having too little fuel economy, accounting for mileage and gas price



Notes: Data from 2009 and 2017 NHTS. Actual share of each demographic group deemed to be driving cars with too little fuel economy, minus that same share calculated under the counterfactual assumption that each driver's car had the median mpg for each wave of the NHTS. Assumes a 7% discount rate and a price of efficiency (p_{μ}) of \$340. Lines represent 95% confidence intervals of the shares.

Appendix

Appendix Table A1. Summary statistics for time-of-purchase data

	<u>Gas-hybrid pairs</u>		All cars (3)
	Gas (1)	Hybrid (2)	
Annual miles	12,515 (32.45)	13,201 (46.69)	12,189 (8.196)
Combined city-highway mpg	24.41 (0.0280)	36.65 (0.0453)	23.97 (0.00838)
Household income			
Less than \$25k	0.006	0.002	0.005
\$25k–\$50k	0.122	0.0542	0.108
\$50k–\$75k	0.160	0.105	0.143
\$75k–\$100k	0.163	0.131	0.149
\$100k–\$150k	0.268	0.282	0.265
Over \$150k	0.282	0.426	0.330
Age	53.97 (0.0782)	56.33 (0.0990)	53.12 (0.0192)
Education:			
High school	0.141	0.060	0.122
Some college	0.005	0.0018	0.005
Graduate	0.282	0.462	0.303
Rural	0.143	0.0973	0.143
Length	193.4 (0.0709)	188.3 (0.0451)	189.5 (0.0225)
Width	73.63 (0.0166)	72.63 (0.0118)	73.63 (0.00490)
Height	62.66 (0.0354)	61.08 (0.0347)	64.11 (0.00905)
Weight	3,786 (4.035)	3,818 (3.760)	3,861 (1.061)
Liters	3.134 (0.00615)	2.539 (0.00467)	3.146 (0.00161)
Valves	3.997 (0.000413)	4 (0)	3.908 (0.000537)
Horsepower	233.0 (0.338)	169.7 (0.295)	243.5 (0.111)
Rpm	5,985 (2.614)	5,828 (1.752)	5,919 (0.686)
Observations	39,327	20,344	607,024

Sources: MaritzCX survey years 2010–2017, vehicle characteristics from Wardsauto.com. Limited to households with cars and complete nonmissing demographic information. Vehicle characteristics from WardsAuto. 59,671 gas-hybrid models in columns 1 and 2, and 607,024 observations for all cars in column 3. Standard errors in parentheses.

Appendix Table A2. Alternative columns for Table 3, with model-specific fixed effects

Dependent variable = 1 if hybrid	NHTS postpurchase data		MartizCX time-of-purchase data	
	“Low” cutoff	“High” cutoff	“Low” cutoff	“High” cutoff
	(1)	(2)	(3)	(4)
Cumulative fuel savings (\$1000)	0.0020* (0.0005)	0.0026* (0.0006)	0.00261* (0.000362)	0.00337* (0.000467)
Upfront investment cost (\$1000)	-0.0487* (0.0025)	-0.0122* (0.0006)	-0.0132* (0.00189)	-0.00331* (0.000474)
Income: \$25k–\$50k	0.0092 (0.0066)		0.0232 (0.0139)	
\$50k–\$75k	0.0111 (0.0066)		0.0382* (0.0138)	
\$75k–\$100k	0.0317* (0.0068)		0.0531* (0.0138)	
\$100k–\$150k	0.0373* (0.0066)	<i>coefficients same as</i>	0.0699* (0.0137)	<i>coefficients same as</i>
Over \$150k	0.0634* (0.0071)	<i>column (1)</i>	0.0858* (0.0138)	<i>column (3)</i>
Education: Some college	0.0104* (0.0045)		0.0160* (0.00314)	
Graduate	0.0333* (0.0052)		0.0383* (0.00342)	
Age: 40–60 years	0.0165* (0.0041)		0.00470 (0.00265)	
Over 60 years	0.0210* (0.0041)		0.0175* (0.00277)	
Male	0.0032 (0.0032)		0.0132* (0.00200)	
Rural	0.0049 (0.0041)		-0.00628* (0.00293)	
Includes car characteristics: length, width, height, weight, liters, valves, horsepower, and rpm Includes fixed effects by state, year-by-type, and model				
Implied $\hat{\rho}$	0.04	0.21	0.56	1.02
Observations	17,586	17,586	59,671	59,671
R-squared	0.43	0.43	0.775	0.775

Sources: See Table 1 and Appendix Table A1. Calculations for cumulative (and discounted) fuel savings assume a 14-year lifetime with a discount rate of 3% (“Low”) and 7% (“High”). * p<0.05.

Appendix Table A3. Full version of Table 6, using NHTS data

		Actual mpg: μ		Actual – optimal mpg: $\mu - \mu^*$
		(1)	(2)	(3)
Optimal mpg μ_i^*		0.0030*	0.0017	
		(0.0005)	(0.0009)	
$(\mu_i^*) \times (\text{Income} > \$50k)$			0.0018	
			(0.0011)	
Income:	\$25k–\$50k	0.101*	0.103*	-1.304*
		(0.039)	(0.039)	(0.240)
	\$50k–\$75k	0.199*	0.112	-1.860*
		(0.039)	(0.064)	(0.242)
	\$75k–\$100k	0.336*	0.248*	-2.540*
	(0.041)	(0.065)	(0.251)	
	\$100k–\$150k	0.436*	0.347*	-2.94*
		(0.040)	(0.065)	(0.24)
	Over \$150k	0.580*	0.490*	-2.60*
		(0.043)	(0.067)	(0.26)
Education:	Some college	0.107*	0.107*	-1.82*
		(0.027)	(0.027)	(0.16)
	Graduate	0.420*	0.420*	-2.41*
		(0.032)	(0.032)	(0.19)
Age:	40–60 years	0.0422	0.0410	1.47*
		(0.025)	(0.025)	(0.16)
	Over 60 years	0.094*	0.0922*	8.80*
		(0.025)	(0.025)	(0.15)
Male		0.038	0.039	-1.37*
		(0.020)	(0.020)	(0.12)
Rural		0.0586*	0.0584*	-3.09*
		(0.0238)	(0.024)	(0.15)
Length (ins.)		-0.0127*	-0.0127*	-0.121*
		(0.0016)	(0.0016)	(0.010)
Width (ins.)		-0.263*	-0.263*	-0.228*
		(0.007)	(0.007)	(0.048)
Height (ins.)		-0.182*	-0.182*	-0.413*
		(0.004)	(0.004)	(0.026)
Weight (lbs.)		0.00106*	0.00106*	0.000736*
		(0.00004)	(0.00004)	(0.000260)
Liters		-2.46*	-2.46*	-1.49*
		(0.04)	(0.04)	(0.27)
Valves		-0.0061	-0.0061	-0.302
		(0.0276)	(0.0275)	(0.169)
Horsepower		-0.0199*	-0.0199*	-0.00952*
		(0.00057)	(0.00057)	(0.00347)
Rpm		-0.00179*	-0.00179*	-0.00200*
		(0.00003)	(0.00003)	(0.00019)
Observations		110,231	110,231	110,231
R-squared		0.72	0.72	0.13

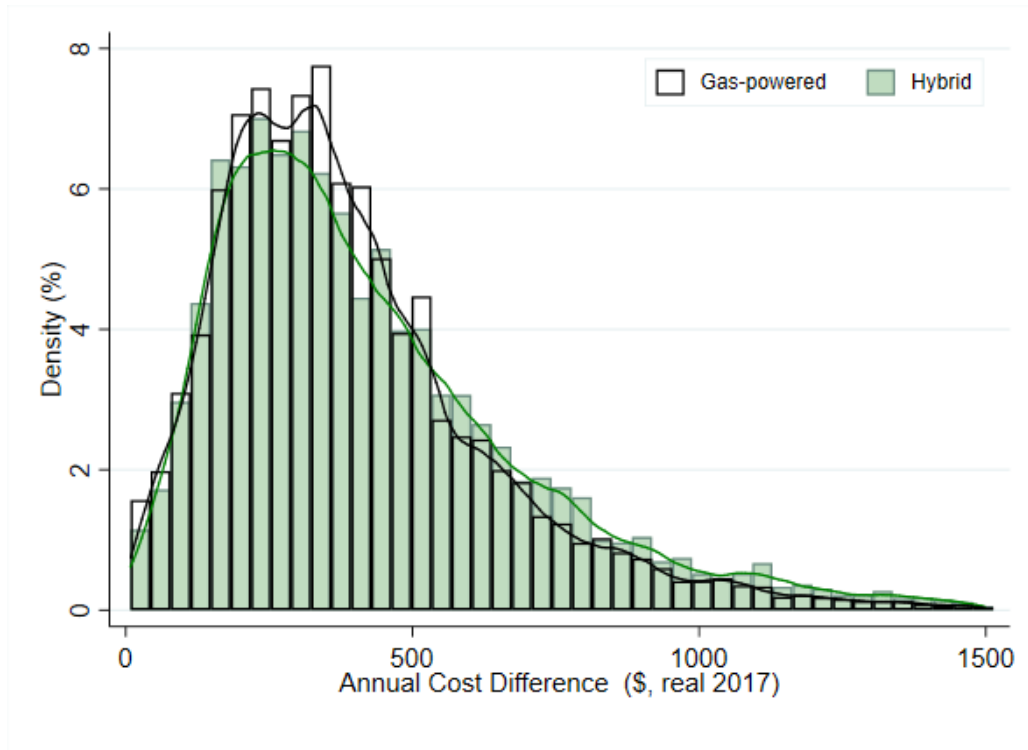
Sources: See Table 1. Calculations for optimal mpg assume a 14-year lifetime with a discount rate of 3%. The cumulative (and discounted) monetary value of one more unit of mpg (according to equation 12) is compared with the marginal cost of mpg estimated in Table 4, Column 5 (\$110). Dependent variable is observed mpg in columns (1) and (2) and difference between observed mpg and optimal mpg in column (3). * $p < 0.05$

Appendix Table A4. Full version of Table 6, using MartizCX data

		Actual mpg: μ		Actual – optimal mpg: $\mu - \mu^*$
		(1)	(2)	(3)
Optimal mpg μ_i^*		0.00715*	0.00726*	
		(0.000271)	(0.000719)	
$(\mu_i^*) \times (\text{Income} > \$50\text{k})$			-0.000120	
			(0.000758)	
Income:	\$25k–\$50k	0.161*	0.162*	-1.943*
		(0.0607)	(0.0607)	(0.292)
	\$50k–\$75k	0.381*	0.388*	-3.432*
		(0.0605)	(0.0748)	(0.291)
	\$75k–\$100k	0.490*	0.497*	-4.477*
		(0.0606)	(0.0750)	(0.291)
	\$100k–\$150k	0.603*	0.610*	-5.641*
		(0.0603)	(0.0748)	(0.290)
	Over \$150k	0.793*	0.800*	-6.803*
		(0.0606)	(0.0751)	(0.291)
Education:	Some college	0.142*	0.142*	-1.645*
		(0.0136)	(0.0136)	(0.0652)
	Graduate	0.463*	0.463*	-1.066*
		(0.0152)	(0.0152)	(0.0728)
Age:	40–60 years	0.118*	0.118*	1.651*
		(0.0114)	(0.0114)	(0.0546)
	Over 60 years	0.294*	0.294*	6.174*
		(0.0122)	(0.0122)	(0.0579)
Male		0.0134	0.0134	-1.239*
		(0.00932)	(0.00932)	(0.0447)
Rural		-0.0310*	-0.0310*	-4.582*
		(0.0126)	(0.0126)	(0.0603)
Length (ins.)		0.0282*	0.0282*	-0.115*
		(0.000682)	(0.000682)	(0.00327)
Width (ins.)		-0.0450*	-0.0450*	0.0232
		(0.00273)	(0.00273)	(0.0131)
Height (ins.)		-0.204*	-0.204*	-0.533*
		(0.00200)	(0.00200)	(0.00958)
Weight (lbs.)		-0.000136*	-0.000136*	-0.000928*
		(2.02e-05)	(2.02e-05)	(9.70e-05)
Liters		-0.683*	-0.683*	0.368*
		(0.0104)	(0.0104)	(0.0498)
Valves		-1.148*	-1.148*	-0.757*
		(0.0135)	(0.0135)	(0.0647)
Horsepower		-0.0407*	-0.0407*	-0.0183*
		(0.000132)	(0.000132)	(0.000633)
Rpm		-0.000779*	-0.000779*	-0.000989*
		(1.01e-05)	(1.01e-05)	(4.86e-05)
Observations		607,024	607,024	607,024
R-squared		0.752	0.752	0.294

Sources: See Appendix Table 1. Calculations for optimal mpg assume a 14-year lifetime with discount rates of 3 percent. The cumulative (and discounted) monetary value of one more unit of mpg (according to equation 12) is compared to the marginal cost of mpg estimated in Table 4, Column 6 (\$110). Dependent variable is observed mpg in columns (1) and (2) and difference between observed mpg and optimal mpg in column (3). * $p < 0.05$.

Appendix Figure A1. Time-of-purchase evidence for annual fuel costs, gas versus hybrid



Notes: Replicates Figure 1 but uses time-of-purchase data from the 2010–2017 MaritzCX survey years. For each vehicle that comes in both gas and hybrid versions, we calculate the annual fuel cost difference between the two models, given expected mileage and gas prices at the time of purchase. The shaded bars represent the cost savings for hybrid drivers relative to driving a gas version of the same car; the outlined bars represent the extra annual costs borne by gas drivers.