DISCUSSION PAPER

June 2016 ■ RFF DP 16-26

The Effect of Market Size on Fuel-Saving Technology Adoption in Passenger Vehicles

Thomas Klier, Joshua Linn, and Yichen Christy Zhou

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

The Effect of Market Size on Fuel-Saving Technology Adoption in Passenger Vehicles

Thomas Klier
Federal Reserve Bank of Chicago

Joshua Linn
Resources for the Future

Yichen Christy Zhou*
University of Maryland

June 2016

Abstract

Passenger vehicle fuel economy standards in the United States and many other countries require substantial new vehicle fuel economy improvements over the next decade. Although economic theory suggests that market size has a positive effect on fuel-saving technology adoption by vehicle manufacturers, there is no empirical evidence on the role of market size in passenger vehicle technology adoption. This paper documents a strong connection between market size, as measured by sales, and technology adoption. Using variation in consumer demographics and purchasing patterns to account for the endogeneity of market size, we find that a 10 percent increase in market size raises power train efficiency by 0.3 percentage points, as compared to a mean improvement rate of 1.4 percent per year between 1997 and 2013. Moreover, we find that fuel prices affect efficiency primarily via market size rather than other channels. Historically, fuel price and demographic-driven market size changes have had large effects on fuel-saving technology adoption. Fuel economy standards, a feebate, or fuel taxes tend to induce firms to adopt fuel-saving technology on their most efficient cars, thereby polarizing the fuel efficiency distribution of the new vehicle fleet.

Key words: passenger vehicles, technological change, market incentives, innovation incentives

JEL classification numbers: L62, Q4, Q55, O31

*Thomas Klier (Thomas.Klier@chi.frb.org) is a senior economist and research advisor at the Federal Reserve Bank of Chicago. Joshua Linn (linn@rff.org) is a senior fellow at Resources for the Future. Christy Zhou (zhou@econ.umd.edu) is a PhD candidate at the University of Maryland. Samuel Goldberg provided excellent research assistance. We thank Yasin Ozcan, Maureen Cropper, Andrew Sweeting, Roberton Williams, seminar participants at Georgia Tech and University of Maryland, and conference participants at the International Industrial Organization Conference.

1 Introduction

Improving vehicle fuel economy is a central part of international efforts to reduce the risks of climate change. Passenger vehicles account for about 15 percent of US greenhouse gas emissions and half of transportation sector emissions (IPCC, 2014). Fuel economy or carbon dioxide emissions standards are common across developed and developing countries. For example, US regulations require new vehicle fuel economy to roughly double between 2005 and 2025 to 54 miles per gallon (mpg).

Achieving the US fuel economy standards over the next decade requires substantial technology adoption (Knittel, 2012). Over the past several decades numerous fuel-saving technologies, such as advanced transmissions, have improved the efficiency of gasoline- and diesel-powered vehicles. Knittel (2012), Klier and Linn (2012), and Klier and Linn (2016) find that the adoption of fuel-saving technology over time has improved the efficiency by 1 to 2 percent annually since 1980.

This paper focuses on a manufacturer's decision to adopt fuel-saving technology across its individual vehicle models. Despite the importance of technology adoption in meeting future standards, the economic factors that drive adoption remain unclear. Klier and Linn (2016) show that the recent tightening of fuel economy standards has accelerated the market-wide rate of adoption of fuel-saving technology between 2000 and 2012. However, absent evidence on the major factors driving technology adoption, the literature on welfare effects of fuel economy standards (e.g., Jacobsen (2013) and Klier and Linn (2012)), as well as the regulatory agencies estimating costs and benefits of the standards, has assumed that technology adoption is driven by adoption costs and consumer willingness to pay for technology.

We depart from the traditional focus, and instead examine the role of market size in explaining technology adoption. According to theories of endogenous technological change, such as directed technological change (Acemoglu, 2002), market size may affect technology adoption, for example due to factors such as fixed costs of technology adoption. Carmakers typically redesign vehicles every five to seven years. Technology adoption entails fixed costs, such as spending resources to redesign and test the vehicle with a new technology before production begins. Fixed costs cause the average cost of technology adoption to decline with a vehicle's market size. A manufacturer can recover fixed adoption costs by charging markups over marginal costs that depend on willingness to pay for fuel costs and other vehicle performance characteristics (Berry et al., 1995). Comparing two hypothetical vehicles for which consumers have the same willingness to pay for fuel costs and other characteristics, this

theory predicts that at a specific point in time the manufacturer will adopt more technology for the vehicle with the larger market size.[1]

Whether market size affects passenger vehicle technology adoption is important for understanding the historical determinants of technology adoption. Furthermore, it has implications for consumer adoption of new technologies such as electric vehicles. Notwithstanding the media attention around electric vehicles and other alternative technologies, the internal combustion engine continues to dominate among passenger vehicles, accounting for about 97 percent of new vehicles in the United States in 2015. Given the market size advantage of gasoline-powered vehicles over alternative fuel technologies such as plug-in electric vehicles, a strong market size effect on technology adoption would imply that manufacturers will continue directing fuel-saving technology improvements to gasoline-powered vehicles (Acemoglu et al., 2012). Such improvements would increase the challenge that alternative fuel vehicles face in competing with gasoline-powered vehicles. Finally, a market size effect would also have implications for the welfare consequences of the various fuel consumption policies such as fuel economy standards, fuel or carbon taxes, and alternative fuel vehicle tax credits. The literature provides evidence of the importance of path dependence on passenger vehicle engine patenting (Acemoglu et al., 2016), and of the role of market size in pharmaceutical innovation (Acemoglu and Linn, 2004). Other studies (e.g., Newell et al. (1999)) document the effects of consumer demand on innovation and technology adoption in air conditioners and other industries. However, there is no evidence for the effects of market size or consumer demand on technology adoption among energy-intensive durable goods such as passenger vehicles.[2]

In this paper, we compare the effects on fuel-saving technology adoption of market size and the effects of fuel-cost-driven changes in consumer demand for technology. We measure the market size of a vehicle by its sales and focus on fuel-saving technology for the internal combustion engine, including gasoline- and diesel-powered engines, and associated transmissions. Using a novel empirical strategy to account for the endogeneity of market size, and using unique data on consumer demographics, consumer preferences for new vehicles, and vehicle-level characteristics, we show that market size has a statistically significant effect on fuel-saving technology adoption for individual vehicles, and that this effect can explain an economically important share of historical technology adoption across vehicles in the market.

[1] The US regulatory agencies include fixed costs in their cost-benefit analysis of fuel economy standards, but they do not consider the effect of market size on technology adoption.

[2] The international trade literature focuses on the the link between market size and productivity (Melitz and Ottaviano, 2008) or product choice (Mayer et al., 2014), but firm-specific productivity and the technology of each product is exogenous in these models.

Fuel costs also affect technology adoption but not independently of their effect on market size. Based on these results we consider several counterfactual scenarios that illustrate the economic and policy implications of the relationship between market size and technology adoption.

Specifically, we construct a unique data set that spans 1997 through 2013. It links US new vehicle sales and characteristics with consumer purchasing patterns by demographic groups. We begin by defining a vehicle power train's efficiency, which is distinct from its fuel economy (miles per gallon). For a given level of power train efficiency, a manufacturer can trade off fuel economy, horsepower, and weight, analogously to movement along a production possibilities frontier. By definition, when a manufacturer adds fuel-saving technology, it can increase fuel economy without affecting the other characteristics. We define the increase in efficiency that results from technology adoption as the increase in fuel economy that is feasible while holding fixed other vehicle characteristics. This definition accounts for the possibility that manufacturers adopt fuel-saving technology and use additional efficiency to boost horsepower or increase weight. We estimate the power train efficiency of each vehicle model by year similarly to Klier and Linn (2016).

The main empirical challenge is the endogeneity of a vehicle's market size. The endogeneity problem, which is common to nearly all empirical analysis of market-driven technological change, arises from both potential reverse causality and omitted variable bias. Adopting fuel-saving technologies may increase demand for a vehicle, causing sales to increase, and resulting in reverse causality. Furthermore, omitted demand or supply variables, such as a vehicle's acceleration, can be correlated with both market size and efficiency.

To address this challenge we construct an instrumental variable (IV) that takes advantage of variation in consumer demographics over time, combined with variation in purchasing behavior across consumer groups. For example, larger households tend to purchase more minivans than smaller households. The fact that the shares of large households in the United States has decreased over the sample has reduced demand for minivans relative to other market segments. To construct the instrument we use consumer preferences by demographic group that are measured at a specific point in time, combined with temporal variation in demographics. The validity of the instrument rests on the exogeneity of time series changes in demographics to the vehicles market. Acemoglu and Linn (2004) and DellaVigna and Pollet (2007) have similarly used demographic trends as exogenous determinants of market size in other industries.[3]

[3]Other papers have used pre-sample information to address endogeneity, such as Blundell et al. (1999).

We find that a one standard deviation increase in market size, which corresponds to about a 10 percent increase, raises a vehicle's efficiency by 0.3 percentage points. This estimate constitutes a large increase relative to the observed average annual efficiency increase of about 1.4 percentage points between 1997 and 2013. In addition to market size, we test whether a vehicle's efficiency responds to the efficiency of competing vehicles or to the manufacturer's stock of fuel-saving patents, finding some effects of competing vehicles. However, these effects are less precisely estimated than the primary effect of market size on efficiency. We also find that the main results are robust to alternative functional forms or constructions of the instrument.

We compare the effects of instrumented market size with the effects of fuel-cost-driven changes in willingness to pay on technology adoption. Fuel costs can affect technology adoption in two ways: first, by affecting consumers' willingness to pay for fuel-saving technology that raises fuel economy as studied in previous literature, and second, indirectly by affecting a vehicle's market size. Busse et al. (2013) and Allcott and Wozny (2014) demonstrate that high gasoline prices raise the market shares of vehicles with high fuel economy relative to vehicles with low fuel economy, and that consumers value the fuel savings offered by vehicles with high fuel economy. These findings suggest that an increase in fuel costs can affect technology adoption in two ways: first, by affecting market size, and second, by affecting consumers' willingness to pay for fuel-saving technology that raises fuel economy. In contrast to the strong effect of market size on technology adoption, after controlling for market size, we find that per-mile fuel costs do not have a statistically significant effect on technology adoption. However, fuel costs strongly predict market size, which is consistent with the literature. The results suggest that fuel costs affect technology adoption via market size, but not via willingness to pay for fuel-saving technology that raises fuel economy. This finding is consistent with Busse et al. (2013), who show that gasoline prices have large effects on market shares but relatively small effects on vehicle prices.

Our empirical findings are the basis for four sets of simulations that demonstrate their economic and policy implications. First, fuel-price-driven changes in market size have had large effects on technology adoption. We illustrate this point by simulating the effects of fuel prices on the cross-sectional distribution of efficiency. The 80 percent increase in real gasoline prices between 2003 and 2007 raised the market size of vehicles with high fuel economy relative to vehicles with low fuel economy. In turn, the changes in market size caused efficiency of the lowest fuel economy vehicles to be lower than if fuel prices had remained at the low 2003 levels. Likewise, efficiency of the highest fuel economy vehicles was higher in 2007 than if fuel prices had remained at 2003 levels.

Second, demographics have had a large effect on the efficiency distribution across models in different market segments. The overall shifts in demographics between 1980 and 2013 caused a shift in efficiency improvements away from light-duty trucks and toward cars. This effect occurred simultaneously with other demand- or supply-side effects on relative power train efficiencies of cars and light trucks, such as changes in gasoline prices.

Third, changes in market size for crossovers and sport utility vehicles (SUVs) have affected technology adoption. Between 2000 and 2004, per-model sales of crossovers increased sharply and per-model sales of SUVs decreased sharply. The increase in crossover market size raised crossover efficiency and the decrease in SUV market size reduced SUV efficiency.

The fourth simulation concerns policies that aim to reduce fuel consumption. Introducing a carbon price, fuel tax increase, feebate, or fuel economy standard will affect the vehicles market in different ways. Yet, these policies have the common feature that they increase the relative market size of vehicles with high fuel economy. For a carbon or gasoline tax increase, the market shares change because higher tax-inclusive fuel costs cause consumers to choose vehicles with higher average fuel economy (Klier and Linn, 2010; Li et al., 2014). In the case of fuel economy standards or a feebate (which jointly taxes and subsidizes vehicles according to their fuel economy), manufacturers respond to standards partly by reducing the relative price of vehicles with high fuel economy, raising their market size (Goldberg, 1995).

Such policy-induced changes in market size in turn affect the cross-sectional efficiency distribution. In our data, efficiency is positively correlated with fuel economy. These policies would strengthen this positive correlation by shifting sales to vehicles with high fuel economy and causing greater efficiency improvements for those vehicles than for vehicles with lower fuel economy. Because alternative fuel vehicles likely compete with high fuel economy gasoline-powered vehicles, the resulting increase in efficiency for vehicles with high fuel economy will present an even greater challenge for alternative technologies such as electric vehicles to gain market share. The market size effect also introduces perhaps an unintended effect of feebates or fuel taxes, which is to decrease the efficiency of vehicles with low fuel economy.

2 Data and Summary Statistics

2.1 Data

We assemble three data sets for the empirical analysis. The first data set includes vehicle characteristics and sales by model year and model version. This data set is constructed by merging vehicle characteristics by model year and model version with sales by model year,

model, and power type. The characteristics are from *Ward's Automotive Annual Yearbooks* from 1997 through 2013 (a change in reporting in 1997 prevents us from extending the sample to earlier years). A model year begins in September of the previous calendar year and ends in the current calendar year. A model version refers to a unique model, trim, body type, and fuel type, such as the two-door gasoline-powered Honda Accord coupe. In addition to these identifying characteristics, other vehicle characteristics include fuel economy, horsepower, torque, weight, transmission type, engine displacement, number of cylinders, and market segment (market segment is aggregated from the *Ward's* vehicle classes as in (Klier and Linn, 2016)). Our data exclude hybrid and electric vehicles, which account for 0.8 percent of sales between 1997 and 2013.[4]

The sales data are from *Ward's Automotive InfoBank*, which reports sales by month, model, and fuel type (gasoline, diesel fuel, and flex fuel, which refer to vehicles capable of using gasoline that contains a high percentage of ethanol). Because technology adoption depends on different factors for conventional internal combustion engines and hybrid electric vehicles, our analysis includes only gasoline and diesel fuel vehicles, as well flex-fuel vehicles; these vehicles account for about 97 percent of the US market in 2013 and 99 percent between 1997 and 2013. We aggregate sales by model year, model, and fuel type and merge with the characteristics data. We collect the real average state-level gasoline and diesel fuel prices by model year from the US Energy Information Administration, and merge the fuel prices to the sales and characteristics data.

The second data set contains vehicle purchases by demographic group and year. We use the 1995 National Personal Travel Survey (NPTS) and the 2001 and 2009 National Household Travel Survey (NHTS). The Department of Transportation makes the survey data available to the public; the three waves had similar scope and sampling methodologies, but the samples are considerably larger in the later years. We refer to all surveys as the NHTS for convenience. For each household in the multi-year sample, the survey collects information on demographics (age, income, etc.), vehicle holdings, and vehicle use. We only keep vehicles that were purchased new in the survey year.

[4]The recent Volkswagen scandal raises some concerns about the accuracy of laboratory testing of vehicle emissions. Laboratory testing of US vehicle fuel economy typically overstates fuel economy by about 20 percent relative to the fuel economy values that appear on window stickers at new vehicle dealerships. There are a few instances of US laboratory tests overstating fuel economy by a substantially greater amount, but these events have affected a smaller number of vehicles than the Volkswagen event. Testing inaccuracies likely affect emissions of other pollutants—such as nitrogen oxides—more than fuel economy or greenhouse gases. The reason is that consumers can observe a vehicle's fuel economy (which is inversely proportional to its rate of greenhouse gas emissions) but they cannot directly observe emissions of other pollutants. This makes it easier to detect systematic cheating on fuel economy than on emissions of other pollutants, providing a disincentive for manufacturers to cheat on fuel economy ratings.

A demographic group is defined by a unique combination of age group, household income group, household size, education group, urbanization status, and geographic census division (see Appendix Table A.1 for definitions of the groups). The age and education groups are based on the attributes of the respondent. Other groups are based on the attributes of the household. Using sample weights for each of the three survey waves we compute the average number of new vehicles purchased per household by market segment and demographic group. For example, we compute the average number of new SUVs purchased by the group defined by households headed by a 35 to 54-year-old with 12 or more years of schooling, with annual household income of $75,000 to $100,000, containing two people, and located in an urban area in New England.

The third data set is constructed from the Current Population Survey (CPS), which is available at the National Bureau of Economic Research, from 1980 through 2013. We compute the number of households for each demographic group using the sample weights. We use the same six-dimension demographic groups as we use in the NHTS.

2.2 Summary Statistics

In this subsection we present summary statistics of vehicles market trends, consumer purchasing patterns, the evolution of consumer demographics over time, and manufacturer adoption of fuel-saving technology. Figure 1 shows total sales by market segment for model years 1997 through 2013, separately for cars and light trucks. The figure illustrates considerable variation in segment-level sales, such as the growth for crossovers that began in the late 1990s and the decline in sport utility vehicles (SUVs) that began shortly thereafter. This variation is useful in identifying the effect of market size on power train efficiency.

Table 1 shows average vehicle characteristics at various times in the sample. Average fuel economy was fairly flat through the mid-2000s, and increased at the end of the sample. Horsepower and weight increased over the sample. Vehicle torque, which represents light truck towing ability, followed a similar pattern.

Figure 2 summarizes the variation in vehicle purchasing patterns across demographic groups. To construct the figure we combine all three NHTS survey waves and weight observations by household survey weights. The figure indicates a substantial amount of variation in purchase behavior across groups. The variation is largely intuitive. For example, the age panel shows that younger households are more likely to buy small cars than older households, and wealthier households are more likely to buy crossovers and SUVs than small cars. Geographic variables are also correlated with purchase behavior; households in urban areas

and the Northeast are much less likely to buy pickup trucks than are other households. The demographic variables are correlated with one another; for example, households with high incomes tend to be well educated. As the next section explains, the IV accounts for this correlation.

Figure 3 shows changes in demographics over time from the CPS. Average age, education, and urbanization increased over time, whereas average household size decreased. As the next section explains, we combine this variation with the variation in purchasing patterns across demographic groups illustrated in Figure 2 to construct the instrumental variable for market size. The raw data support this approach by indicating that the time series changes in demographics, combined with heterogeneous purchasing patterns across demographic groups, are consistent with changes in market size. For example, the market share of crossovers increased from the late 1990s through the 2000s. This is consistent with the fact that older households are more likely to purchase crossovers than younger households, and that during the same time period the share of older households increased.

Finally, we present some background information about technology adoption in the US new vehicles market. Manufacturers continually redesign their vehicles, improving power train technology and other attributes that consumers demand. Many vehicles experience major redesigns at regular intervals, commonly every five to seven years. During a redesign, the manufacturer may make major changes to the power train, cabin, cargo, or exterior. In between redesigns, manufacturers commonly make smaller changes to exterior design or to the power train, offering new options such as paint color or increasing the number of transmission speeds.

These regularities yield a process of steady technology adoption over time. Figure 4 shows the share of vehicles in the market with the indicated fuel-saving engine or transmission technologies. The data cover 1986 through 2014, and are from the EPA Annual Fuel Economy Guides and Trend Reports. For many of these technologies, the figure suggests fairly typical patterns in the technology adoption literature, in which the penetration rate is very low initially, subsequently increases steeply, and then levels off—that is, an S-curve. Note that the penetration rate of multiport fuel injection decreases in the late 2000s, which is because many manufacturers had begun to replace this technology with other fuel injection technologies.

3 Empirical Strategy

In this section we motivate the reduced-form estimation equation using a simple model of technology adoption with fixed costs. Then, we estimate efficiency of each vehicle model in the sample, and finally we derive the estimating equation and explain the IV strategy.

3.1 Technology Adoption with Fixed Costs

We outline a simple model of technology adoption with fixed costs. In practice, manufacturers can improve efficiency by adding technologies such as those depicted in Figure 4. Adoption entails fixed costs because of the need to redesign and test the vehicle before commencing production of vehicles that include the new technologies. Many efficiency-improving technologies also increase marginal costs because adoption requires that new parts be installed in the power train. For example, NRC (2015) estimates that adding cylinder deactivation, which effectively shuts off a subset of a vehicle's cylinders when the vehicle is operating under a light load, increases production costs by \$118-\$133 per vehicle. To approximate these aspects of technology adoption, we assume that both marginal costs, c, and fixed costs, F, have positive first and second derivatives with respect to efficiency, τ.

The analysis focuses on a firm that sells one type of vehicle. The firm chooses the vehicle's price p and efficiency to maximize profits:

$$\pi(p, \tau) = \max_{p,\tau} \left(p - c(\tau)\right) q(\tau, p) - F(\tau)$$

where q is the quantity of vehicles demanded, which equals quantity supplied in equilibrium. Quantity demanded decreases with price and increases with efficiency (that is, $\frac{\partial q}{\partial p} < 0$ and $\frac{\partial q}{\partial \tau} > 0$).

This setup introduces the simplification that vehicle demand depends on the vehicle's efficiency. Most consumers likely care about fuel economy, horsepower, and other vehicle characteristics that affect efficiency, rather than efficiency per se. As we discuss in Section 3.2, for a vehicle with a particular level of efficiency there exist trade-offs among fuel economy, horsepower, and weight. Consequently, manufacturers will choose efficiency and then select fuel economy, horsepower, and characteristics that affect weight (such as electronic accessories). For simplicity we abstract from those subsequent decisions and instead focus on the choice of power train efficiency. We assume that quantity demanded increases with efficiency because a vehicle with higher efficiency allows the manufacturer to increase fuel economy without sacrificing other characteristics (or, likewise, to increase horsepower

or weight, without sacrificing other characteristics). Jointly modeling the manufacturer's choice of efficiency, horsepower, and fuel economy, as in Klier and Linn (2012), would not affect the main conclusions but would increase the complexity of the expressions below.

The first order conditions for price and efficiency are

$$\begin{aligned} \text{price:} \quad & (p-c)\frac{\partial q}{\partial p} + q = 0 \\ \text{efficiency:} \quad & (p-c)\frac{\partial q}{\partial \tau} - c'q - F' = 0 \end{aligned}$$

where $c'$and F' indicate the first derivatives with respect to efficiency. Combining the two first order conditions and rearranging them yields

$$F'(\tau) = q(-c' - \eta) \tag{1}$$

where $\eta = \frac{\partial q/\partial \tau}{\partial q/\partial p} < 0$, and η increases the more sensitive is quantity demanded to efficiency than to price. Provided that η does not decrease quickly with quantity (i.e., $\partial\eta/\partial q > -1$), and c' is relatively small in magnitude, equation (1) shows that equilibrium efficiency increases with vehicle sales. That is, vehicles with higher equilibrium market size have higher efficiency. In addition, η captures the willingness to pay for efficiency, and equation (1) shows that higher willingness to pay raises efficiency. Although we do not show it here, this conclusion is unaffected if the firm sells multiple vehicles (provided that fixed costs are vehicle-specific) or faces fuel economy standards. The positive relationship between sales and efficiency is therefore more general than the simple model implies.

3.2 Estimating Power Train Efficiency

The empirical objective is to estimate the effect of market size on power train efficiency. In this subsection we describe the construction of the dependent variable, which is power train efficiency.

We do not directly observe a vehicle's efficiency. The data do not contain efficiency per se, but they do include fuel economy and a number of other observable variables that affect efficiency, such as the number of engine cylinders.

We follow Knittel (2012) and Klier and Linn (2016) and estimate efficiency from the available data. We begin by defining a power train's efficiency as the amount of useful energy it produces per unit of fuel consumption. A vehicle's fuel economy (miles per gallon) is

distinct from its efficiency. Its fuel economy depends on the efficiency of its power train as well as characteristics such as horsepower, weight, and body type (which affects air resistance). As in Klier and Linn (2016), we conceive of an efficiency frontier defined in fuel economy-horsepower-weight space. The frontier represents the maximum fuel economy that can be achieved given any particular level of horsepower and weight, and holding efficiency fixed along the frontier. That is, for a particular level of efficiency, as it moves along the frontier, the manufacturer can trade off fuel economy for weight and horsepower.

This framework yields a straightforward identification of efficiency improvements over time. Specifically, we estimate the shape of the frontier using within-model variation in horsepower, weight, and fuel economy. As a baseline we assume that the shape of the frontier does not change over time. In that case, if we control for the effects of weight, horsepower, and other attributes on fuel economy, an increase in fuel economy is equivalent to an increase in efficiency. To implement this approach, we estimate an equation similar to Klier and Linn (2016):

$$\ln e_{jt} = \lambda_h \ln h_{jt} + \lambda_w \ln w_{jt} + \tau_{mt} + X_{jt}\delta + \varepsilon_{jt}, \tag{2}$$

where e_{jt} is the fuel economy of vehicle j in model year t, h_{jt} is horsepower for passenger cars (and torque for light-duty trucks), w_{jt} is weight, τ_{mt} is a set of interactions of model by model year, X_{jt} includes a vector of vehicle attributes, ε_{jt} is an error term, and the λs and δ are coefficients to be estimated. The coefficients on weight and horsepower capture trade-offs between these characteristics and fuel economy. We expect both coefficients to be negative. The controls in X_{jt} include fixed effects for whether the vehicle uses diesel fuel, whether the vehicle is flex-fuel capable, and whether the vehicle has a manual transmission, as well as fixed effects for the number of doors and the number of cylinders. Together, these variables allow for the fact that versions of a particular model sold in the same model year have different efficiency depending on fuel type and body type (as approximated by the number of doors). We estimate the equation separately for cars and light trucks to allow the coefficients to vary across the two classes.

We interpret the interactions of model by model year, τ_{mt}, as the average efficiency of vehicles belonging to the model, and sold in model year t. The difference between τ_{mt} and $\tau_{m(t-1)}$ is the change in efficiency of model m between model years $t-1$ and t. Equation (2) thus allows us to identify changes in efficiency over time, where efficiency is measured in units of log fuel economy.

Before presenting the results from estimating equation (2) we briefly discuss identification and potential sources of bias. Equation (2) characterizes a technical relationship between vehicle characteristics and fuel economy. It does not include certain vehicle attributes that consumers care about, such as seating comfort. Such attributes could be correlated with variables that are included in equation (2), but in this context that would not bias the coefficients as long as the omitted variables affect fuel economy via horsepower or weight, and not independently of the included variables. In other words, identification rests on the ability to include the variables that directly determine a vehicle's fuel economy. The high R-squared value reported below supports this estimation approach. See Klier and Linn (2016) for additional discussion of identification of equation (2).

Table 2 reports the main coefficient estimates from equation (2). Because fuel economy, horsepower, and weight enter equation (2) in logs, the horsepower and weight coefficients are elasticities. The coefficients on diesel fuel and flex fuel are the difference between log fuel economy of a vehicle that uses diesel fuel or is flex fuel capable and the log fuel economy of an otherwise comparable gasoline-powered vehicle. Diesel fuel vehicles achieve about 30 percent higher fuel economy, and flex-fuel light trucks achieve about 27 percent lower fuel economy than gasoline-powered vehicles. The negative coefficient on flex-fuel vehicles reflects the lower energy content of ethanol compared to gasoline. Overall, the estimates in Table 2 have the expected signs and are statistically significant at the 1 percent level. The magnitudes are similar to those reported in Klier and Linn (2016) for both cars and light trucks. The magnitudes of many of the coefficients are fairly similar across the car and light truck classes, which reflects a substantial degree of shared technology across the classes.

Because of the large number of estimated model by model year interactions, we aggregate across observations before reporting those estimates. Figure 5 plots the change in power train efficiency, averaged across cars and light trucks. The figure shows steady efficiency improvements for both vehicle classes. Table 3 shows the average change in efficiency by five-year time periods, separating models with sales above the median level of sales for the corresponding period and vehicles with sales below the median level of sales. Efficiency improvements are generally higher for the higher-selling models, which previews the main empirical finding that market size has a positive effect on efficiency-improving technology adoption.

3.3 Empirical Strategy for Estimating the Effect of Market Size on Efficiency

This section presents the strategy for estimating the effect of market size on efficiency. We assume a log-linear relationship between market size and efficiency that can be derived from equation (1) if we assume that fixed costs of technology adoption are iso-elastic (i.e., $F(\tau) \propto \tau^{\alpha}$, where $\alpha > 1$). Alternatively, a log-linear functional form can be derived from a model in which marginal costs decrease with production, because of learning by doing or scale economies in vehicle production.

The estimating equation is

$$\hat{\tau}_{mt} = \gamma_1 \ln Q_{mt} + \gamma_2 \bar{C}_{mt;s} + \phi_t + \phi_b + \phi_b \times t + \varepsilon_{mt} \tag{3}$$

where $\hat{\tau}_{mt}$ is power train efficiency estimated from equation (2), Q_{mt} is sales, $\bar{C}_{mt;s}$ is fuel costs per mile (dollars-per-mile), ϕ_t and ϕ_b are sets of year and brand fixed effects, $\phi_b \times t$ is the interaction of brand fixed effects with a linear time trend, and ε_{mt} is an error term. The two parameters of interest are γ_1 and γ_2, which are the effects of log sales and fuel costs on efficiency. Equation (1) implies a positive coefficient on log sales, which would indicate that manufacturers adopt more efficiency for high-selling vehicles. The variable $\bar{C}_{mt;s}$ is the weighted national average fuel price in year t and using weights constructed in period s, normalized by the vehicle's fuel economy (we explain the variable construction in more detail below). The coefficient on fuel costs is the effect of fuel costs on efficiency, holding market size fixed. We discuss the identification and interpretation of this coefficient at the end of the subsection. The year fixed effects control for aggregate demand or supply shocks and the brand fixed effects control for brand-level supply or demand shocks, such as consumer perceptions of brand quality. The interactions of the brand fixed effects with a linear time trend allow brand-specific demand and supply shocks to vary linearly over time. For example, the time trends control for changes in consumer preferences for brands as well as changes in brand quality.

Estimating equation (3) by ordinary least squares (OLS) is likely to yield biased estimates for three main reasons. First, there would be reverse causality if increasing a vehicle's efficiency raises a vehicle's demand and, therefore, equilibrium sales. Second, sales may be correlated with unobserved supply or demand determinants of fuel-saving technology. The brand fixed effects and time trends control for brand-level supply or demand shocks, but efficiency could be correlated with within-brand variation in vehicle characteristics. For example, there is anecdotal evidence that manufacturers test efficiency-improving technologies

on luxury or performance vehicles before installing the technologies more broadly. This practice would cause sales and efficiency to be correlated with (omitted) characteristics such as seating quality or cabin space. Note that we could control flexibly for omitted model-level characteristics by including model fixed effects in equation (3). That approach would yield an undesirable interpretation of γ, however. The coefficient would be identified by within-model variation over time in sales and efficiency. In practice, manufacturers face choices not only about when to adopt technology for a particular model but also, given time and resource costs, about which of their models will receive improved technology at a particular time. Including model fixed effects would identify the former choice but not the latter, and therefore might omit an important role of market size in technology adoption across vehicle models.

A final source of bias is that technology adoption is a dynamic decision that includes fixed and irreversible costs. Efficiency may therefore depend on current sales as well as expected future sales. We use current sales in equation (3) as a proxy for expected sales, but this introduces measurement error that biases the estimated sales coefficient.

We instrument for sales and address all three sources of bias. The IV is the vehicle's potential market size, which depends on cross-sectional variation in consumer purchasing patterns and time series variation in demographics. We define demographic group cell, g, by age, income, education, household size, urbanization, and census division (see Appendix Table A.1 for definitions of the groups). To measure purchasing behavior by group, we compute $q_{mg;s}$ for each of the 2,700 cells as the number of vehicles of model m purchased per household by demographic group cell g in NHTS survey year s (recall that we use data from the NHTS survey years 1995, 2001, and 2009; each cell represents an average of about 20 households). To measure time-series variation in demographics, we compute the number of households in demographic group cell g in year t, w_{gt}, using CPS data. The potential market size is the product of NHTS vehicle purchases per household and CPS number of households, summed across demographic group cells:

$$\tilde{Q}_{mt;s} = \sum_{g} (q_{mg;s} \times w_{gt})$$

The subscript s in potential market size reflects the fact that $q_{mg;s}$ varies across NHTS survey waves. The IV is based on the assumption that w_{gt} is exogenous to the demand and supply of new vehicle technology. Changes in educational attainment, labor participation, and the US income distribution are driven by broad technological developments (such as information technology), the decrease in unionization, and other factors that are largely unrelated to

the new vehicles market (Black and Lynch, 2001; Bresnahan et al., 2002; Jorgenson, 2001; Johnson and Mieszkowski, 1970; Autor et al., 2008). Likewise, the overall increase in age depicted in Figure 3 arises from the aging of the baby boom generation. Household size, urbanization, and migration trends are similarly driven by changing preferences and other factors that are unrelated to demand and supply of new vehicle technology. The assumed exogeneity of these demographics follows assumptions made by Acemoglu and Linn (2004) and DellaVigna and Pollet (2007) for consumer demand in other industries.

We would be concerned about using the average per-household vehicle purchases, $q_{mg;s}$, however, because this variable may be correlated with demand and supply factors omitted from equation (3). For example, an increase in the efficiency of crossovers would increase per-household purchases of crossovers as measured in the NHTS, creating reverse causality between the dependent variable and the NHTS weights.

We make two refinements to the instrument to address this issue. First, we define time periods based on the NHTS survey waves: 1997-2000, 2001-2008, and 2009-2013. The variable $q_{mg;s}$ is measured at the beginning of each time period and does not vary across years within a period. Returning to the crossover example, the fact that $q_{mg;s}$ is constant within a time period reduces the likelihood that efficiency improvements that occur within the period are correlated with $q_{mg;s}$ because the variable is measured at the beginning of the period and does not change in response to efficiency improvements.

However, there could be unobserved and time-invariant characteristics of the vehicle that are correlated with efficiency and $q_{mg;s}$. For example, in a particular year a crossover with high efficiency could have higher sales per household than a crossover with lower efficiency. We use a second refinement to address this possibility, and demean $\tilde{Q}_{mt;s}$ by period and vehicle model. The demeaned value $\bar{Q}_{mt;s}$, is the instrument. This refinement eliminates correlation between the IV and unobserved model characteristics that are time invariant. Note that rather than demeaning the instrument we could add model by period interactions to equation (3), but this would affect the interpretation of the log market size coefficient as discussed above.

The demeaned potential market size, $\bar{Q}_{mt;s}$, is the IV in the first stage for market size in equation (3), which also contains fuel costs

$$\ln Q_{mt} = \beta_1 \ln \bar{Q}_{mt;s} + \beta_2 \bar{C}_{mt,s} + \beta_3 I_{mt}^{imp} + \phi_t + \phi_b + \phi_b \times t + u_{mt} \tag{4}$$

where I_{mt}^{imp} is an indicator variable equal to one if the instrument is imputed using brand-segment-year means.[5] As we show below, the instrument is a strong predictor of log sales, reducing concerns about weak instruments bias. Because the instrument is demeaned by time period, the identifying assumption is that within-period variation in demographics is uncorrelated with omitted demand and supply shocks. Supporting this assumption is the fact that demographic shifts are slow-moving and are driven by factors such as the aging of the baby boom generation and macroeconomic factors. Under this identifying assumption the IV strategy addresses reverse causality and omitted variables bias. The instrument is plausibly uncorrelated with measurement error in actual market size, addressing bias due to classical measurement error (we discuss non-classical measurement error in Section 4.2).

We interpret the coefficient on log sales in equation (3) as the effect on efficiency of a change in market size induced by a change in potential market size. In practice, changes in potential market size may affect equilibrium sales as well as vehicle prices. We do not control for vehicle prices in equation (3) because prices are likely to be correlated with unobserved demand or supply factors, and we lack suitable price instruments in this context in which technology and vehicle characteristics are endogenous (Klier and Linn, 2012). If the potential market size is a valid instrument, it is uncorrelated with unobserved supply or demand factors that affect vehicle price (and sales), and omitting the vehicle's price would not cause spurious results.

Next, we discuss the identification and interpretation of the vehicle's fuel cost per mile. The coefficient on fuel cost per mile is identified by fuel price and fuel economy variation across vehicles and over time. We use the contemporaneous fuel price under the assumption that price shocks are fully persistent (Busse et al., 2013). Previous research (Klier and Linn, 2010) has used the ratio of the national average fuel price to the vehicle's fuel economy to approximate per-mile fuel costs. Using this approach, fuel cost per mile varies because of time-series variation in fuel prices and cross-model variation in fuel economy. We slightly refine our previous approach and introduce additional variation by exploiting geographic variation in fuel prices and vehicle purchases. For example, fuel prices tend to be higher in the Northeast than the Midwest. Households purchase more small cars relative to pickup trucks in the Northeast than the Midwest, which causes the national average fuel price for households that purchase small cars to be higher than the national average fuel price for households that purchase pickup trucks. Formally, we compute a model-specific fuel price

[5] Some vehicle models appear in the sales and characteristics data but not in the NHTS data. Most of these are low-selling models. In these cases, we impute the instruments using brand-level average NHTS weights.

using NHTS data on vehicle purchases and EIA data on fuel prices by census division, d:

$$p_{mt;s} = \sum_{g} \left(p_{dt} \times q_{mg;s} \times w_{gt} \right) / \tilde{Q}_{mt}$$

We calculate the fuel cost per mile as the ratio of the model's fuel price to its fuel economy e_{m0}, which is measured in the first year the model is observed in the sample

$$\widetilde{C}_{mt;s} = \frac{p_{mt}}{e_{m0}}$$

In the cross section, $\widetilde{C}_{mt;s}$ is correlated with the vehicle's fuel economy (by construction), and may therefore be correlated with vehicle characteristics that are correlated with fuel economy, such as horsepower. Including $\widetilde{C}_{mt;s}$ in equation (3) as an independent variable would yield biased estimates because the variable would be correlated with the error term in equation (3). Similarly to the potential market size instrument, we eliminate this possibility by subtracting the mean fuel cost by model and period to obtain the demeaned independent variable $\bar{C}_{mt;s}$, which we include in equation (3). Note that demeaning this variable reduces concerns about the potential endogeneity of the NHTS weights in the fuel price.

The coefficient on fuel costs is the effect of fuel costs on efficiency after controlling for market size. Because the estimating equation includes year fixed effects, the coefficient is identified by within-year variation in fuel costs. We expect the coefficient to be positive because higher fuel costs raise the value of an efficiency improvement of a particular magnitude.

We interpret this coefficient as capturing the effect of consumer willingness to pay for fuel-saving technology that raises fuel economy. The coefficient on log sales is identified by variation in the IV as well as all of the other independent variables, which includes fuel costs. Therefore, fuel costs can affect efficiency via market size, as well as having a direct effect on efficiency.

4 Estimation Results

4.1 Main Results

Table 4 shows the main estimation results. Column 1 reports the OLS estimates of equation (3) for comparison with the IV estimates in column 2. In all columns of Table 4, the dependent variable is the efficiency estimated in Table 2 and observations are by model and model year from 1997 through 2013. To control for aggregate demand or supply

shocks as well as brand-specific shocks, the regression includes year fixed effects, brand fixed effects, and the interaction of a linear time trend with brand fixed effects. The table reports the estimated coefficient on log sales with the bootstrapped standard error in parentheses, clustered by brand (make) to allow for arbitrary correlation of the error term within brands and over time, and for the fact that the dependent variable is estimated in equation (2). The estimated coefficient on log sales is 0.006 and the estimate is statistically significant at the 5 percent level. The coefficient on fuel costs is negative, and the estimate is statistically significant at the 1 percent level.

The OLS estimates in column 1 are likely to be biased because of reverse causality, omitted variable bias, and measurement error (see Section 3.3). To address all three issues, we instrument for log sales using the log of demographics-driven potential market size, $\ln \bar{Q}_{mt;s}$. Column 2 in Panel B of Table 4 shows the results from the first stage. The instrument is a strong predictor of market size. The coefficient has the expected positive sign, and is statistically significant at the 1 percent level.

The magnitude of the IV estimate in Panel A is statistically and economically significant. Between 1997 and 2013, the average annual efficiency improvement is about 1.4 percent (see Figure 5). As shown in column 2, the estimated sales coefficient implies that a one standard deviation increase in log sales, or a 10 percent increase, raises efficiency by 0.3 percent. This estimate is substantially larger than the OLS estimate in column 1. It suggests that omitted variables or measurement error creates greater bias to the OLS estimate than does reverse causality (which biases the OLS estimate away from zero). In the following subsections we present a variety of additional estimation results and we refer to column 2 in Table 4 as our baseline estimate.

The coefficient on fuel costs in the second stage is negative, but the estimate is much smaller than in column 1 and it is not statistically significant. In addition to the lack of statistical significance, the results below in Section 5.1 imply that the OLS magnitude is economically small. Comparing the IV and OLS estimates suggests that fuel costs affect fuel-saving technology adoption primarily via market size. The negative and statistically significant negative coefficient on fuel costs in column 1 likely reflects the correlation between sales and fuel costs. The negative OLS coefficient may also reflect a negative correlation between fuel economy and other vehicle attributes such as horsepower. After controlling for fuel-cost-driven changes in market size in the first stage, fuel costs do not affect efficiency. Nevertheless, fuel costs affect market size, as the first stage coefficient on fuel costs indicates. The importance of market size in determining the adoption of fuel-saving technology, relative to the direct effect of fuel costs, is confirmed in column 3, in which we omit fuel costs from

the first and second stages. In this case, the coefficient on log sales is nearly identical to that in column 2.

4.2 Alternative Estimation Models

As discussed above, the IV addresses the main potential sources of endogeneity of the market size variable. In this subsection we show that the results are robust to adding further controls and alternative procedures for estimating efficiency and market size.

The baseline includes controls for brand-level demand or supply shocks, but there may also be segment-level shocks. For example, the increase in market shares of crossovers in the late 1990s and early 2000s, along with the decrease in market share of SUVs during the same period, could reflect a shift in consumer preferences toward smaller, car-like light trucks. Although we subtract mean preferences by model and period from the instruments, causing them to be orthogonal to segment-level shocks, there could be within-period preference changes correlated with the within-period demographic changes. Such a correlation would bias the IV estimate, but column 4 shows that the log sales coefficient is very similar if we add to the baseline the interactions of market segment fixed effects and a linear time trend.

Fuel economy standards have varied over the data sample in stringency and form. The standards were roughly constant in the 1990s and early 2000s, but began increasing for light trucks in 2005, and then for both cars and light trucks in 2011. Because of differences in fleet composition and market positioning, the standards impose varying degrees of pressure across manufacturers to improve fuel economy over time, which has affected the adoption of energy efficiency improving technology (Klier and Linn, 2016). The interactions of brand fixed effects with a linear time trend control imperfectly for the standards because the standards require manufacturers to achieve fleet-wide levels of fuel economy, and did not allow manufacturers to average fuel economy across classes (cars and light trucks). We have tried several semi-parametric approaches to controlling for fuel economy standards, in addition to using brand fixed effects and the interactions of brand fixed effects with a linear time trend. In particular, column 5 includes the interactions of brand fixed effects with a quadratic time trend. This controls for the nonlinear changes in the stringency of the standards over time across manufacturers. Historically the standards applied separately for cars and light trucks, but since 2011 manufacturers can average across their entire fleet (Leard and McConnell, 2016). To account for the differing regulatory pressure across cars and light trucks, column 6 includes triple interactions of brand fixed effects, vehicle class fixed effects (i.e., passenger car or light-duty truck), and a linear time trend. This controls for changes in stringency of the standards over time and across vehicle classes. The results are similar to our baseline in

column 2. Note that these specifications also address the potential bias caused by unobserved demand or supply shocks at the brand, market segment, or class level.

The results are also similar to the baseline if we control directly for the stringency of the standards as in Klier and Linn (2016). In Table 4 column 7, the coefficient on log sales is 0.032, with standard error 0.008, which is statistically significant at the 1 percent level and very similar to the baseline estimate. Finally, unobserved preference or cost shocks may be correlated with the shadow cost of the fuel economy standards, and therefore with incentives the standards create for technology adoption. To allow for this possibility in column 8 we add the interaction of the stringency variable with fuel costs (which depend on fuel economy and are therefore likely to be correlated with such shocks), yielding results that are similar to our baseline.

Table 5 reports estimates of equation (3) using alternative measures of efficiency as the dependent variable and alternative measures of market size as the key independent variable. In the baseline specification (repeated in column 1 for convenience), we estimate efficiency by model and model year using equation (2), implicitly assuming that efficiency is constant across versions of the same model and model year. This assumption is supported by the fact that versions of the same model, such as the Honda Accord, typically include engines produced on the same or a very similar production platform. However, because many technologies are installed at the engine platform rather than the model level and some models share an engine platform (Klier and Linn, 2012), platform-level market size could affect efficiency. To assess whether engine platform-level market size affects efficiency, columns 2 and 3 report estimates of equation (3) that are the same as the baseline, except for the estimation of the dependent variable. These specifications take advantage of highly detailed engine platform data, which allow us to identify the specific engine sold with each version. In column 2 we estimate efficiency in equation (2) by engine platform and model year rather than by model and model year, and use the estimated efficiency as the dependent variable in equation (3). The estimated coefficient on log sales is similar to the baseline. In column 3 we estimate efficiency by model and platform generation (such that a redesign of the engine constitutes a new generation).[6] The log sales coefficient is larger than the baseline. Thus, there is some evidence that the baseline understates the effect of market size on efficiency, but no evidence of a spurious estimate in the baseline specification.

As noted in Section 2.2, manufacturers typically make large changes to the power train or vehicle during major redesigns, and smaller changes between redesigns. The baseline

[6]Different models in the same year sold under the same brand could share a platform, as could one model in different years. It is also possible for models sold under a different brand to share a platform.

estimates include efficiency improvements that occur both within and across redesigns, but the relationship between log sales and efficiency may be different across redesigns from the relationship within redesigns. To allow for this possibility we define a change in model generation as occurring when the model is redesigned, and in column 4 we estimate efficiency by model generation and year (we collect model generation information from Automotive News). The estimated coefficient on log sales is similar to the baseline.

Because of regular production and redesign cycles in the vehicles market, efficiency may respond gradually to market size. Column 4 represents one approach to allowing for this possibility, by focusing on efficiency improvements across generations. Column 5 represents an alternative. In this case we use as the dependent variable the three-year moving average of efficiency. The estimate is close to the baseline.

Next, we consider possible sources of measurement error in the market size variable. Analogously to the three-year moving average for efficiency that was reported in column 5, we can use the three-year moving average of the model's sales to allow for the possibility that efficiency responds to the average market size over multi-year periods; the results are similar (not reported).

Given the time required to redesign and test a vehicle before beginning production of a new generation, there could be a lag between market size and adoption. The fact that shifts in demographics and potential market size can be forecast to some extent mitigates the lag between demographics-driven changes in market size and adoption, but there could nonetheless be a lag. We can consider this possibility empirically by replacing current log sales with the one-year lag of sales, and by replacing fuel costs with lagged fuel costs; column 6 reports results similar to the baseline.

Note that the potential market size variable is constructed from NHTS purchasing weights that vary across NHTS survey periods, but not within periods. We address potential correlation between these weights and unobserved vehicle attributes by demeaning the potential market size variable. However, this approach does not address possible measurement error in the weights due to the relatively small sample sizes from the 1995 and 2001 NHTS survey waves. To allow for this possibility, we construct the purchase weights using the 2009 NHTS survey, which has a much larger sample than the previous survey waves. The results are similar to the baseline.[7]

[7] Another source of measurement error for market size is that some vehicle models are produced on global platforms, and technology could respond to global market size for these models. However, even in such cases manufacturers commonly select engines and transmissions that are specific to the market, in which case the US market size would be most relevant to the chosen engine and transmission technologies for the vehicles

4.3 Additional Channels and Heterogeneous Technology Adoption

So far we have focused on the link between a vehicle's market size and its efficiency. In this subsection we consider possible indirect effects on efficiency, such as competitors' behavior. We report these results in Table 6. For comparison, column 1 repeats the baseline specification from column 2 of Table 4.

Manufacturers may respond to the efficiency of competing models (Fischer, 2010). Because consumers have heterogeneous preferences for efficiency and other vehicle attributes, the efficiency of competing models could have a positive or negative effect on a particular model's efficiency. On the one hand, if a manufacturer's competitors increase the efficiencies of their competing vehicles, that manufacturer might increase the efficiency of its vehicles to avoid losing customers to the other manufacturers' vehicles. On the other hand, instead of increasing efficiency in response to competitors' efficiency improvements, the manufacturer might not adopt efficiency or might even decrease efficiency and instead improve other attributes such as cabin size. These changes would attract customers with low valuation of efficiency and high valuation of cabin size. In column 2 of Table 6 we add to the baseline specification the mean efficiency of other vehicles in the same market segment and technology group, where we assign each manufacturer to one of three technology groups (Japanese, US, and other). This efficiency variable may be endogenous because of reverse causality and perhaps other reasons, and we instrument for it using the mean potential market size of the corresponding vehicles. Column 2 provides some evidence that efficiency of competing models has a positive effect on technology adoption; in column 2 the coefficient on competing vehicle efficiency is positive. The coefficient on competing efficiency suggests a more than one-for-one effect of competing vehicle efficiency, but the large standard errors reflect the limited variation of this variable.

Manufacturers could adopt efficiency-improving technology at the brand rather than the model level. This could occur if models share engine platforms or because of scale economies in redesigning vehicles. Table 5 addresses the former possibility but not the latter. We add to column 3 the mean frontier of other models sold under the same brand in the same market segment, and using mean potential market size of the corresponding models as an instrument. It is perhaps somewhat surprising that brand-level efficiency has a small and negative effect on efficiency, but this result could arise from the limited variation in this variable (as indicated by the large standard error).

sold in the United States. In addition, the United States represents about 20 percent of global sales, and therefore represents an important consideration in manufacturers' technology decisions.

Next, we consider the effect of the knowledge stock on technology adoption. To improve efficiency, manufacturers could adopt technologies that are already widely used in the market—either in their own vehicles or in those of competing manufacturers. Alternatively, they could innovate and adopt new technology. We construct a proxy for the effect of innovation and adoption of new technology by adding to equation (3) an estimate of a manufacturer's knowledge stock based on its historical patents. The variable is the cumulative number of fuel-saving patents for which a parent company has applied. The variable, which is sometimes referred to as the knowledge stock, is the sum of the depreciated patent stock from the previous period and the flow of patents in the current period (see Zhou (2016) for details on variable construction). Column 4 controls for knowledge stock between 1997 and 2010 and shows that knowledge stock has a positive effect on efficiency, but the estimate is not statistically significant.[8]

We have focused on the role of market size and fuel-cost-driven willingness to pay for technology that raises fuel economy. Consumer demand for other vehicle attributes, such as horsepower, may also affect the adoption of fuel-saving technology. If the IV strategy is valid, such omitted factors would not yield biased or spurious estimates of the market size effect. To demonstrate this point and to consider the role of other factors driving technology adoption, we add to the main regression the vehicle's price as a proxy for consumers' overall willingness to pay for the vehicle. Column 5 shows that adding the vehicle price does not affect the estimate of log sales, supporting the validity of the IV strategy. The price coefficient is positive, which suggests that vehicle demand affects technology adoption, but this coefficient is likely to be biased because of correlation with unobserved vehicle attributes.

Finally, we consider the possibility of heterogeneous effects of market size. In column 6, we interact market size with a dummy variable for light trucks, and instrument for this variable with the interaction between the potential market size and the light truck dummy. The point estimate of the market size effect is barely affected although it is not precisely estimated, and we cannot reject the hypothesis that market size affects technology adoption by the same amount for cars and light trucks. In column 7, we interact market size with a dummy variable equal to one for US-based manufacturers. Similarly to the light truck exercise we do not find strong evidence of heterogeneous effects, although the limited variation of the market size variable prevents strong conclusions regarding heterogeneity.[9]

[8]The stock of patent variable ends at 2010. OECD Triadic Patent Family (TPF) data are available up to 2015. However, it is common practice not to use the last four to five years of TPF data because of reporting lags from the US Patent and Trademark Office.

[9]We have also allowed for heterogeneity across market segments or across other types of manufacturers, yielding similar conclusions.

5 Implications

This section discusses four implications of the main estimates from Section 4. We quantify the effect of past variation in fuel prices on efficiency; the effect of changes in demographics on efficiency; and the effect of sales of crossovers and SUVs on efficiency; and we estimate the effect of a fuel tax, feebate, or fuel economy standards on the variation in efficiency across models sold in the market.

5.1 Effects of Gasoline Prices on Efficiency

In Section 4 we quantified the economic significance of the magnitude of the log market size coefficient by comparing the effect of a one standard deviation market size increase with the average efficiency improvement observed in the sample. To further illustrate the economic importance of this estimate, we compare efficiency levels across scenarios of low and high gasoline prices.

Between 2003 and 2007 the real price of gasoline increased almost 80 percent. Klier and Linn (2010) show that this price change increased sales of vehicles with high fuel economy at the expense of sales of vehicles with low fuel economy. The shifts in market shares increased sales-weighted average fuel economy by about 0.5 mpg.

We use equation (3) to estimate the effect of the resulting changes in market shares on efficiency, and report the results in Figure 6. We assign each model in the data to one of five fuel economy quintiles, based on each model's initial fuel economy when it first appears in the data. The first quintile includes vehicles with the lowest fuel economy, and the fifth quintile includes vehicles with the highest fuel economy. The colored bars show the average predicted efficiency increase for each quintile using the actual fuel prices between 2003 and 2007 and the baseline estimates of equation (3). The clear bars show the average predicted efficiency increase assuming fuel prices had remained at 2003 levels.

For this exercise, we hold the market-wide average efficiency increase equal in the two scenarios. To maintain consistency with equation (3), which includes year fixed effects that control for the market-wide average efficiency increase, we assume that fuel prices do not affect the market-wide average rates of efficiency changes. We do allow fuel prices to affect the cross-sectional distribution of sales, which in turn affects the cross-sectional distribution of efficiency. We use equations (3) and (4) to generate the counterfactual efficiency each year from 2003 to 2007. Because gasoline prices are lower in the counterfactual scenario, we expect counterfactual efficiency to be higher than predicted efficiency for the first quintile, which includes the lowest fuel economy vehicles.

In this counterfactual exercise, we focus on the effect of fuel prices on efficiency via market size, rather than the direct effect via consumer demand. Therefore, we adjust fuel prices for the first stage equation (4) but not the second stage equation (3). As a sensitivity check, Figure A.1 shows minimal differences if we also adjust fuel prices in the second stage.

Comparing the predicted and counterfactual cumulative efficiency improvements across quintiles, we observe that if fuel prices had remained at 2003 levels, efficiency would have improved 0.45 percentage points more for the lowest fuel economy quintile, and by 0.38 percentage points less for the highest fuel economy quintile. These effects are consistent with expectation and they are large relative to the predicted cumulative 6.8 percentage points efficiency improvements that actually occurred between 2003 and 2007.

5.2 Effects of Demographics on Efficiency

Next we analyze how demographics affect market size and the cross-sectional distribution of efficiency. As a counterfactual scenario, we suppose that all demographics remain unchanged at 1980 levels. The counterfactual reflects the changes in income, age, and other demographics over the 33 years between 1980 and 2013 (see Appendix Figure A.2). In Figure 7, we compare predicted changes in efficiency using actual demographic changes in the colored bars with counterfactual efficiency changes that would have occurred if demographics had remained fixed at 1980 levels in the clear bars.

Similarly to the fuel price exercise, we use equations (3) and (4) to predict counterfactual efficiency improvements each year from 1997 through 2013, again assuming that annual market-wide average efficiency improvements are unchanged from their observed levels. Figure 7 compares the predicted and counterfactual cumulative efficiency improvements from 1997 through 2013. The figure shows that if demographics had remained constant at 1980 levels, small cars would have been 0.48 percentage points less efficient, medium cars would have been 0.22 percentage points less efficient, SUVs would have been 0.54 percentage points more efficient, and pickup trucks would have been 0.68 percentage points less efficient, while other segments are less affected. For context, the overall change between 1997 and 2013 was 24 percentage points. Thus, demographics explain some of the cross-sectional variation in efficiency changes over the sample period.

5.3 Effects of Crossover and SUV Market Size on Efficiency

Figure 1 shows the large shifts in sales for crossovers and SUVs that occurred in the early 2000s. Those shifts reflect segment-level sales changes, and underlying model-level

sales changed in the same directions. Between 2000 and 2004, the average sales per model of crossovers increased by 44 percent and average sales per model of SUVs decreased by 25 percent (in contrast, the number of SUV models increased during this period). The empirical results suggest that these changes in market size caused efficiency to increase for crossovers and to decrease for SUVs, relative to a counterfactual in which market size had remained constant.

To quantify these effects, we analyze how cumulative efficiency would have been affected if the market size of crossovers and SUVs remains at 2000 levels through 2004.[10] On the left side of Figure 8 we compare the predicted and counterfactual efficiency of crossovers (CUVs). The colored bar shows the predicted cumulative efficiency improvement over 2000-2004. The clear bar shows the counterfactual efficiency bar holding market size of crossovers fixed at 2000 levels. In the counterfactual scenario, the lower market size of crossovers causes efficiency to be 0.52 percentage points lower. This is economically significant compared to a cumulative change of 3.6 percentage points for crossovers.

As shown in the right panel of Figure 8, the counterfactual scenario causes SUVs to be 0.38 percentage points higher than is predicted using the actual market size in 2004. This is substantial in magnitude compared to cumulative changes between 2000 and 2004 of 2.9 percentage points for SUVs. In short, comparing the simulations for fuel prices, demographics, crossovers, and SUVs, we observe that market size and fuel prices have had economically significant effects on the adoption of fuel-saving technology.

5.4 Effects of Taxes, Feebates, and Fuel Economy Standards on the Efficiency Distribution

Raising fuel taxes, introducing a carbon tax or feebate, or imposing fuel economy standards all increase sales of vehicles with high fuel economy at the expense of vehicles with low fuel economy. In this subsection we discuss the mechanism by which the policies affect market size and technology adoption, and use a hypothetical feebate to illustrate the implications of the market size effect.

A fuel tax increase or carbon price raises per-mile driving costs of all vehicles, but more so for vehicles with low fuel economy than vehicles with high fuel economy. The increase in

[10] Because we are interested in the effect of market size on crossover and SUV efficiency, our counterfactual represents a partial equilibrium outcome in which the market size of crossovers and SUVs does not affect total vehicle sales or total cumulative efficiency across the market. That is, total vehicle sales and average cumulative efficiency across all segments are identical in the predicted and counterfactual scenarios.

relative driving costs of the low fuel economy vehicles decreases their market size compared to those of high fuel economy vehicles.

The effect of fuel economy standards on vehicle sales arises from a different mechanism but nonetheless has the same qualitative effect on market size. If the standard applies to the sales-weighted (harmonic) mean fuel economy of a manufacturer's vehicles, as with the US standards and those in other regions, one compliance strategy available to manufacturers is to reduce the prices of vehicles with high fuel economy relative to vehicles with low fuel economy (Goldberg, 1995). The relative vehicle price change induces consumers to substitute from vehicles with low fuel economy to vehicles with high fuel economy. Consequently, standards cause the sales of low fuel economy vehicles to decrease relative to sales of high fuel economy vehicles. The shift in market size raises the manufacturer's sales-weighted average fuel economy.

Finally, a feebate refers to a system of taxes and rebates that jointly offers subsidies to vehicles with high fuel economy and imposes taxes on vehicles with low fuel economy. The taxes and rebates therefore mimic the pricing behavior of manufacturers facing fuel economy standards, and a feebate can be designed to achieve outcomes that are identical to those of a fuel economy standard (Roth, 2014).

Thus, taxes, feebates, and fuel economy standards increase sales of vehicles with high fuel economy relative to sales of vehicles with low fuel economy. Figure 6 shows that historical fuel price increases have affected the efficiency of vehicles in the market, and we would expect a carbon price, fuel tax increase, feebate, or standard to affect efficiency.

Rather than consider the effects of these policies on efficiency improvements for different fuel economy groups, as we did before, we further illustrate the implications of the market size effect by focusing on the policies' effects on the variation of efficiency across vehicles in the market. In principle, these policies could widen or narrow the distribution of efficiencies of vehicles in the market. If, in the absence of any policy, fuel economy is positively correlated with efficiency, the policies would widen the distribution. This is because the market size of vehicles with high fuel economy would increase, raising their efficiency, and the market size of vehicles with low fuel economy would decrease, reducing their efficiency. Both changes would strengthen the positive correlation between fuel economy and efficiency. If, on the other hand, fuel economy is negatively correlated with efficiency, the policies would narrow the efficiency distribution. In practice, the correlation is positive, 0.29, in which case we expect the policies to increase the variance of efficiency across vehicles in the market.

In this simulation we focus on a feebate (the other policies discussed above would have qualitatively similar effects). A feebate is determined by a pivot, which is the fuel economy level above which vehicles receive a subsidy and below which vehicles are taxed, and a rate of subsidy and taxation per unit of fuel economy. We set the pivot equal to the sales-weighted mean fuel economy in year t, e_t. For comparability with the fuel price counterfactual in Figure 6, the rate of taxation is chosen such that the sales-weighted average fuel economy increases by 0.5 mpg, which is the same fuel economy change as occurred in the counterfactual scenario considered in Figure 6. We consider a per-mile feebate rather than a lump-sum feebate, such that a model with fuel economy e_{jt} has a feebate of $(1/e_{jt} - 1/e_t) \times 1.53$. The counterfactual scenario includes a feebate for the years 2010-2013, and market conditions (e.g., fuel prices) are otherwise unchanged. We compute the predicted and counterfactual efficiency of each model in the sample for the years 2010 through 2013, and compute the cumulative predicted and counterfactual efficiencies for each model. As explained above, because of the feebate's effect on market size, we expect the feebate to increase the variance of efficiency across vehicles in the market.

Panel A of Figure 9 compares the cross-model distribution of predicted and counterfactual efficiencies, using the cumulative efficiency changes over 2010 through 2013. The solid line shows the estimated density function of the predicted efficiency, and the dashed line shows the estimated density function of the counterfactual efficiency had the feebate been in place. Consistent with expectations, the figure shows that the feebate increases the variance of efficiency across vehicles in the market.

Further illustrating the effects of the feebate on the distribution of efficiency in the market, Panel B of Figure 9 presents a scatter plot of efficiency and fuel economy for each model in the sample. The solid dots represent the predicted cumulative efficiencies of models sold in 2013 and the black circles are counterfactual cumulative efficiencies. Because the feebate reduces the market size of vehicles with fuel economy below the pivot, the counterfactual efficiency lies below the predicted efficiency for models with fuel economy below the pivot. In contrast, the feebate increases the market size of vehicles with fuel economy above the pivot, and causes counterfactual efficiency to lie above predicted efficiency for such vehicles. The lines in Panel B represent the fitted values of a linear regression of efficiency on fuel economy, which is estimated separately for the predicted and counterfactual data points. The counterfactual line is steeper than the predicted line, which indicates that the feebate strengthens the positive relationship between efficiency and fuel economy.

Previous welfare analysis of these policies has not considered their effects on market size and technology adoption, which have several implications for both the cost effectiveness

of the policies and their distributional effects. First, the feebate introduces a shadow cost of fuel economy, inducing manufacturers to increase the fuel economy of all vehicles. Our analysis illustrates an unintended consequence of the feebate, which is that because of the feebate's effect on market size, the feebate causes manufacturers to decrease the efficiency of their vehicles with fuel economy below the pivot. The market size effect lies behind this unintended effect, which could undermine the effectiveness of feebates or taxes at reducing fuel consumption.[11] Second, and following from the first, by affecting the distribution of efficiency across vehicles in the market, these policies result in consumers purchasing different vehicles than they would have if market size did not affect efficiency. While estimating the welfare consequences of these efficiency improvements lies outside the scope of this paper, the scenario considered here illustrates the effects of these policies on the distribution of efficiency across vehicles in the market. Third, alternative fuel technologies likely are closest substitutes to gasoline-powered vehicles with above average fuel economy. If a policy raises the efficiency of high fuel economy gasoline-powered vehicles, it would reduce consumer demand for alternative fuel vehicles. This implication for competition between alternative fuel technologies and conventional technologies represents an additional unintended consequence of these policies.

6 Conclusion

Current US fuel economy standards will dramatically increase the average fuel economy of new vehicles over at least the next decade. Despite the magnitude of this fuel economy change and the importance of technology adoption for meeting the standards, there is little empirical evidence on which factors determine a manufacturer's adoption of fuel-saving technology. This paper analyzes the effect of market size, as approximated by vehicle sales, on the adoption of efficiency-improving technology in the US new passenger vehicles market. We show that market size has a substantial effect on technology adoption and discuss implications for the evolution of technology adoption and fuel consumption policies.

We motivate the empirical analysis using a simple model of energy efficiency technology adoption, which shows that fixed costs of technology adoption cause adoption to depend on a vehicle's market size. The empirical analysis uses a unique data set that combines vehicle characteristics and sales with vehicle purchasing patterns by demographic group from 1997 to 2013. We address the endogeneity of market size by instrumenting for vehicle

[11] Although our analysis suggests that the market size effect undermines the effectiveness of fuel consumption policies, the market size effect does not necessarily undermine the economic efficiency of these policies. It may be economically efficient for the efficiency of the most efficient vehicles to improve the most.

sales using potential market size. Variation in potential market size arises from changes in demographics over time and cross-sectional heterogeneity in purchasing patterns across demographic groups. In the preferred specification we find that a 10 percent increase in sales (corresponding to about one standard deviation) increases efficiency by 0.3 percentage points, compared to a mean annual efficiency improvement of about 1.4 percentage points in the sample. Fuel costs affect efficiency via market size but not independently of market size. Acemoglu et al. (2016) find that high fuel prices increase clean technology innovation, and given our findings that fuel prices affect technology adoption primarily via market size, future research can address whether innovation responds directly to fuel prices or to fuel-price-driven changes in market size.

The results have four main implications. The first is that historical variation in fuel prices has had a substantial effect on the pattern of introducing new vehicle technology. Real fuel prices nearly doubled between 2003 and 2007. If gasoline prices had remained at 2003 levels, the efficiency of vehicles with low fuel economy would have increased 0.5 percentage points more between 2003 and 2007 than it did. Efficiency of the highest fuel economy vehicles would have increased by 0.4 percentage points less than it did.

Second, demographics have had a large effect on vehicle efficiency. Shifts in demographics increased efficiency of cars relative to light trucks. Third, shifts in market size of crossovers and SUVs have caused large changes in the efficiency of these vehicles. These three results imply that market size and fuel costs have had economically significant effects on technology adoption in the new vehicles market.

The final implication is that a fuel tax, feebate, or fuel economy standard increases the cross-sectional variation in efficiency by raising the efficiency of vehicles with high fuel economy relative to vehicles with low fuel economy. The existing literature on the welfare effects of these policies has not considered this effect, and future research could incorporate it. Accounting for the effects of market size and fuel costs on technology adoption would affect the estimated cost effectiveness and the distributional consequences of these policies, as well as the demand for alternative fuel vehicles such as plug-in electrics.

The empirical analysis does not identify the underlying reasons for why market size affects efficiency. The simple model we introduced to motivate the empirical analysis emphasizes the role of fixed costs in vehicle production, but other factors could also contribute to the positive relationship between market size and technology adoption, such as learning by doing, which may arise for reasons other than the presence of fixed costs. Future work may address

this question, which has important implications for the welfare effects of fuel consumption policies.

References

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies 69*(4), 781–809.

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The environment and directed technical change. *American Economic Review 102*(1), 131–166.

Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016). Transition to clean technology. *Journal of Political Economy 124*(1), 52–104.

Acemoglu, D. and J. Linn (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics 119*(3), 1049–1090.

Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics 96*(5), 779–795.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in US wage inequality: Revising the revisionists. *Review of economics and statistics 90*(2), 300–323.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society 63*(4), 841–890.

Black, S. E. and L. M. Lynch (2001). How to compete: the impact of workplace practices and information technology on productivity. *Review of Economics and statistics 83*(3), 434–445.

Blundell, R., R. Griffith, and J. Van Reenen (1999). Market share, market value and innovation in a panel of british manufacturing firms. *Review of Economic Studies 66*(3), 529–554.

Bresnahan, T., E. Brynjolfsson, and L. M. Hitt (2002). IT, workplace organization and the demand for skilled labor: A firm-level analysis. *Quarterly Journal of Economics 117*(1), 339–376.

Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013). Are consumers myopic? Evidence from new and used car purchases. *American Economic Review 103*(1), 220–256.

DellaVigna, S. and J. M. Pollet (2007). Demographics and industry returns. *American Economic Review 97*(5), 1667–1702.

Fischer, C. (2010). Imperfect competition, consumer behavior, and the provision of fuel efficiency in light-duty vehicles. (RFF Discussion Paper No. 10-60).

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the US automobile industry. *Econometrica: Journal of the Econometric Society*, 891–951.

IPCC (2014). *Climate Change 2014: Mitigation of Climate Change.*

Jacobsen, M. R. (2013). Evaluating US fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy 5*(2), 148–187.

Johnson, H. G. and P. Mieszkowski (1970). The effects of unionization on the distribution of income: A general equilibrium approach. *Quarterly Journal of Economics*, 539–561.

Jorgenson, D. W. (2001). Information technology and the US economy. *American Economic Review 91*(1), 1–32.

Klier, T. and J. Linn (2010). The price of gasoline and new vehicle fuel economy: Evidence from monthly sales data. *American Economic Journal: Economic Policy 2*(3), 134–153.

Klier, T. and J. Linn (2012). New-Vehicle characteristics and the cost of the Corporate Average Fuel Economy standard. *RAND Journal of Economics 43*(1), 186–213.

Klier, T. and J. Linn (2016). Technological change, vehicle characteristics and the opportunity costs of fuel economy standards. *Journal of Public Economics* (forthcoming).

Knittel, C. (2012). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review 101*(7), 3368–3399.

Leard, B. and V. McConnell (2016). New markets for pollution and energy efficiency. (RFF Discussion Paper No. 15-16).

Li, S., J. Linn, and E. Muehlegger (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy 4*(6), 302–342.

Mayer, T., M. Melitz, and G. Ottaviano (2014). Market size, competition, and the product mix of exporters. *American Economic Review 104*(2), 495.

Melitz, M. J. and G. I. Ottaviano (2008). Market size, trade, and productivity. *Review of Economic Studies 75*(1), 295–316.

Newell, R. G., A. B. Jaffe, and R. N. Stavins (1999, August). The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics 114*(3), 941–975.

NRC (2015). Cost, effectiveness, and deployment of fuel economy technologies for light-Duty vehicles, phase 2. *National Research Council Publication*.

Roth, K. (2014). The unintended consequences of uncoordinated regulation: Evidence from the transportation sector. *Working Paper*.

Zhou, Y. C. (2016). Knowledge capital, technology adoption, and environmental policies: Evidences from the US automobile industry. *Job Market Paper*.

Figures

Figure 1: **Vehicle Sales by Segment, 1997-2013**

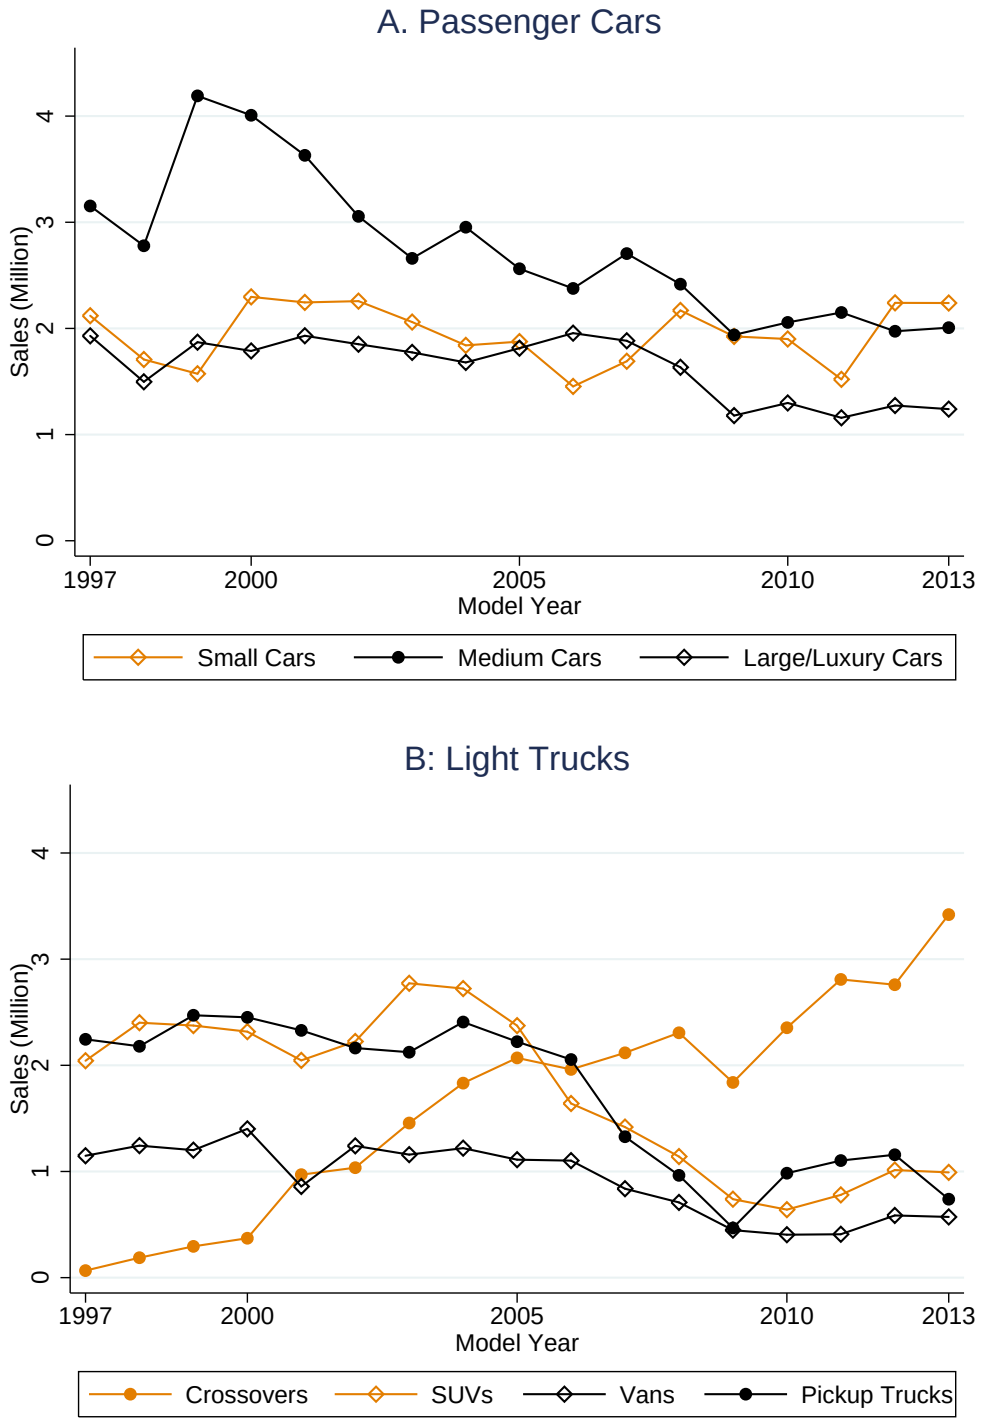

Notes: For each market segment, the figure plots the total model year sales. Panel A includes passenger cars and Panel B includes light-duty trucks.

Figure 2: **Vehicle Purchase Patterns by Demographic Group**

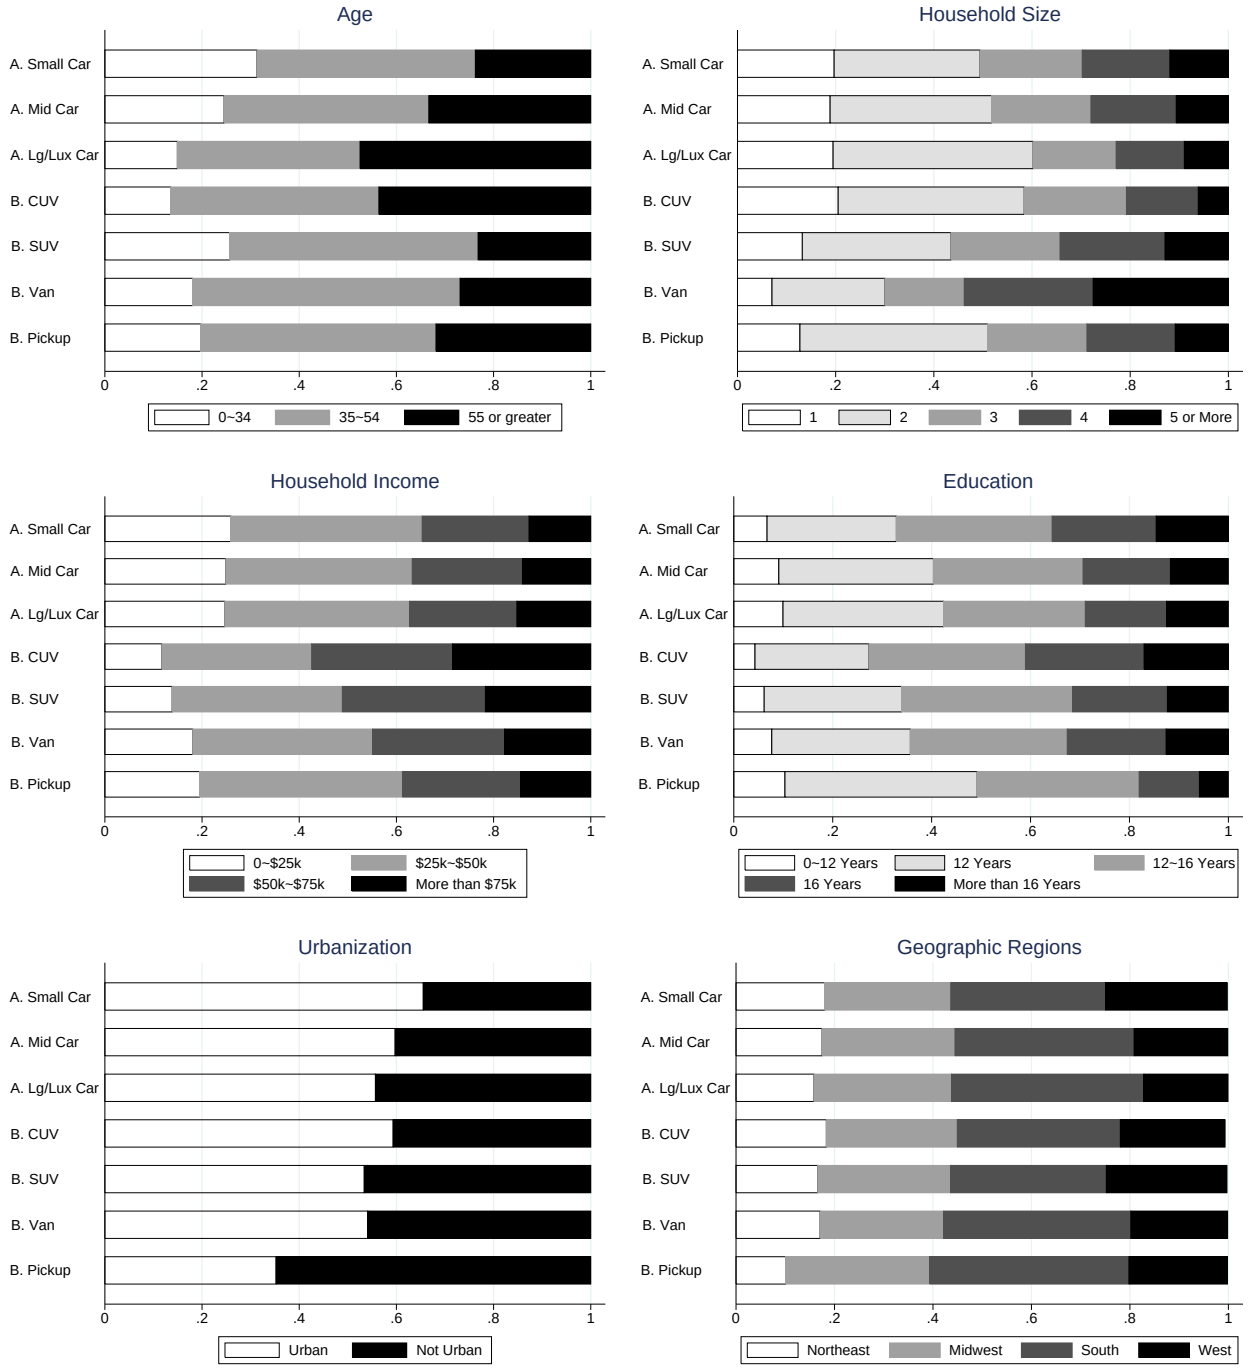

Notes: The figure is constructed using the NHTS survey data from the 1995, 2001, and 2009 survey waves. Each panel illustrates purchasing patterns for the indicated demographic variable. For households purchasing vehicles in a particular market segment, we compute the share of those households belonging to each category of the demographic variable, using the NHTS household survey weights. For example, among the households that purchase small cars, 64% of them live in urban areas.

Figure 3: **Changes in Demographics, 1997-2015**

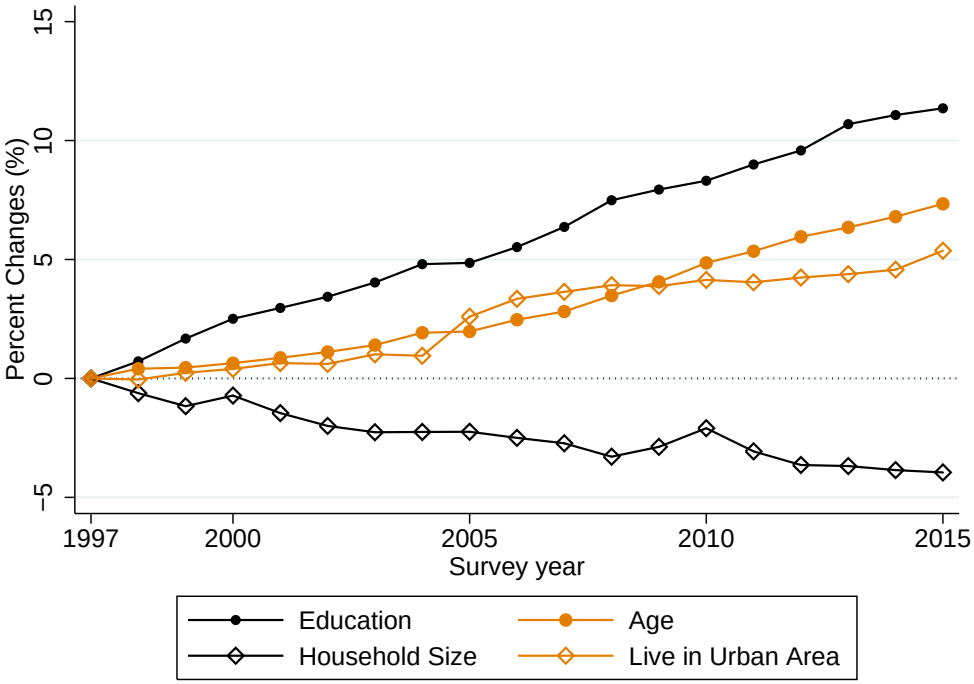

Notes: Using household survey weights from the CPS, we compute the weighted average of each demographic variable by year. The figure plots the percentage change since 1997 of each variable.

Figure 4: **Market Penetration of Selected Fuel-Saving Technologies, 1986-2014**

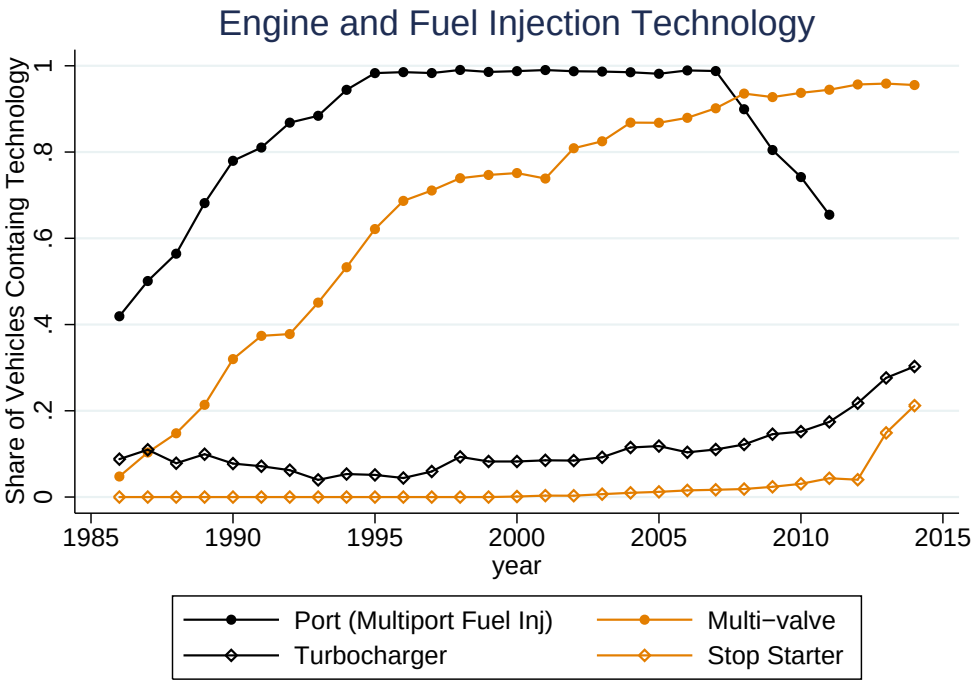

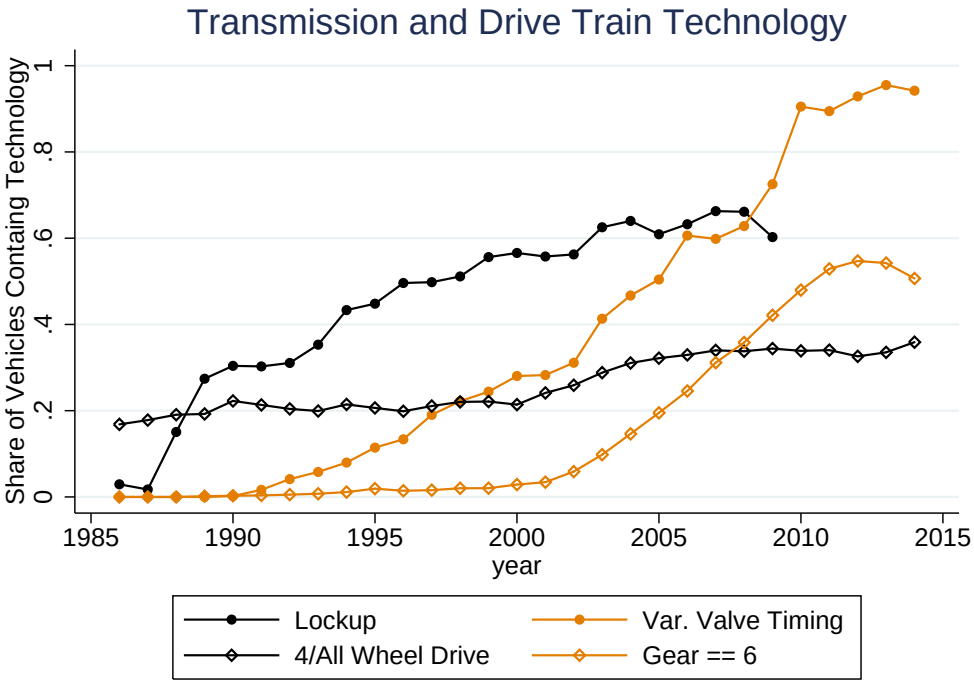

Notes: The figure is constructed from the EPA Fuel Economy Guide and EPA Fuel Economy Trends data. Technology penetration rates are the unweighted average across all vehicles in the corresponding model year.

Figure 5: **Estimated Power Train Efficiency, 1997-2013**

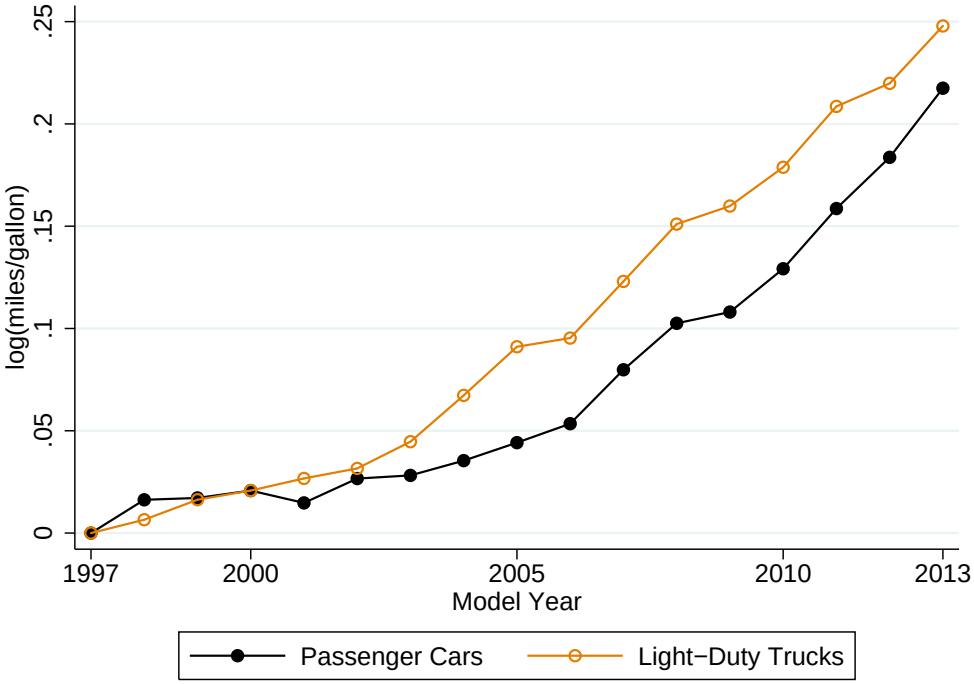

Notes: The figure plots the mean estimated efficiency across passenger cars and light trucks estimated from equation (2). To construct this figure, efficiency is normalized to zero for all observations in 1997.

Figure 6: **Effect of 2003-2007 Gasoline Price Increase on Efficiency**

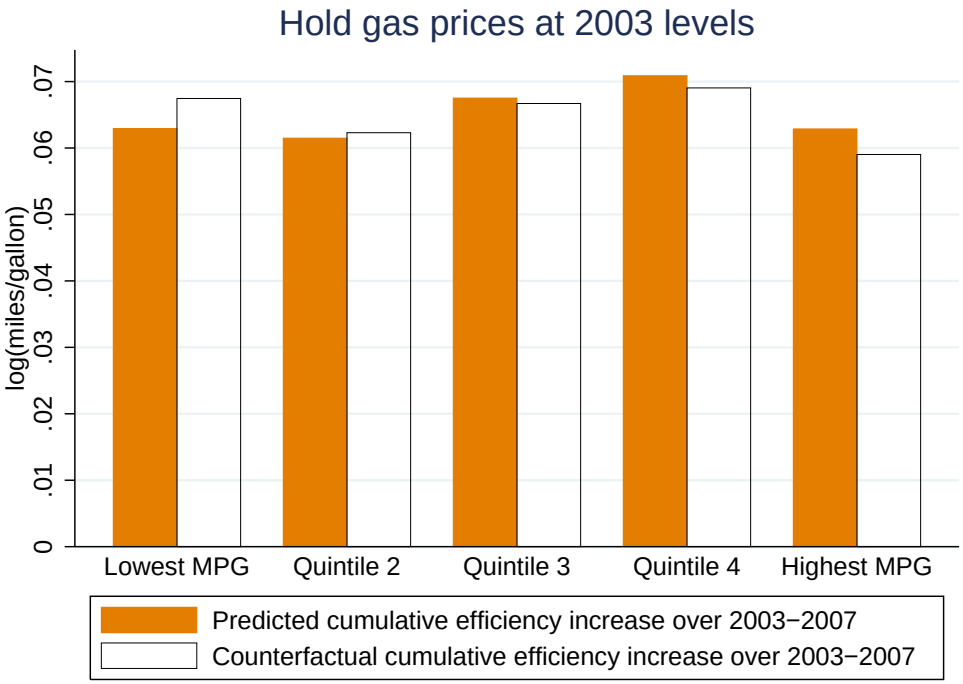

Notes: For each observation in equation (3), the frontier is predicted using the estimates reported in column 2 of Table 4. All observations are assigned to a fuel economy quintile based on the fuel economy distribution across observations between 2003 and 2007, using each vehicle model's initial fuel economy when the model enters the market. The predicted frontier in each colored bar is the mean cumulative predicted efficiency change between 2003 and 2007 for each quintile. The clear bars show the cumulative counterfactual efficiency change by quintile. Counterfactual efficiency changes are computed by holding fixed fuel prices at 2003 levels and using equations (3) and (4) to predict the efficiency change for each observation between 2003 and 2007.

Figure 7: **Effect of Demographics on Efficiency, 1980-2013**

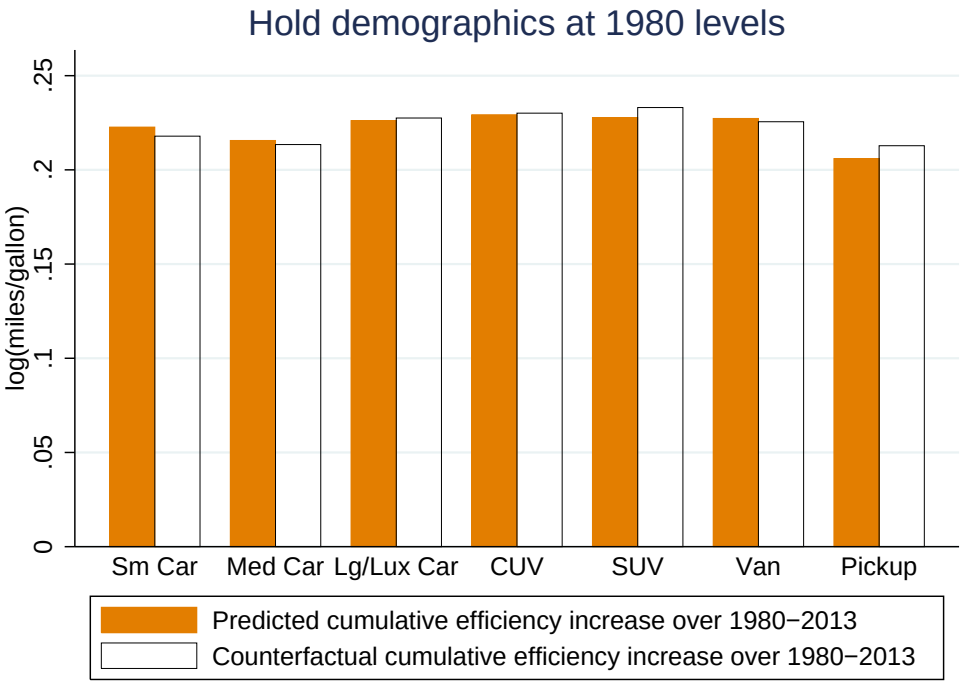

Notes: The colored bars show the mean cumulative predicted efficiency increase between 1997 and 2013, for each market segment. Predicted values are obtained from the estimation of equation (3) reported in column 2 of Table 4. The clear bars show the cumulative counterfactual efficiency change by market segment. The counterfactual holds fixed demographics at 1980 levels and uses equations (3) and (4) to predict the counterfactual efficiency change for each observation between 1997 and 2013.

Figure 8: **Effect of Sales on Efficiency, 2000-2004**

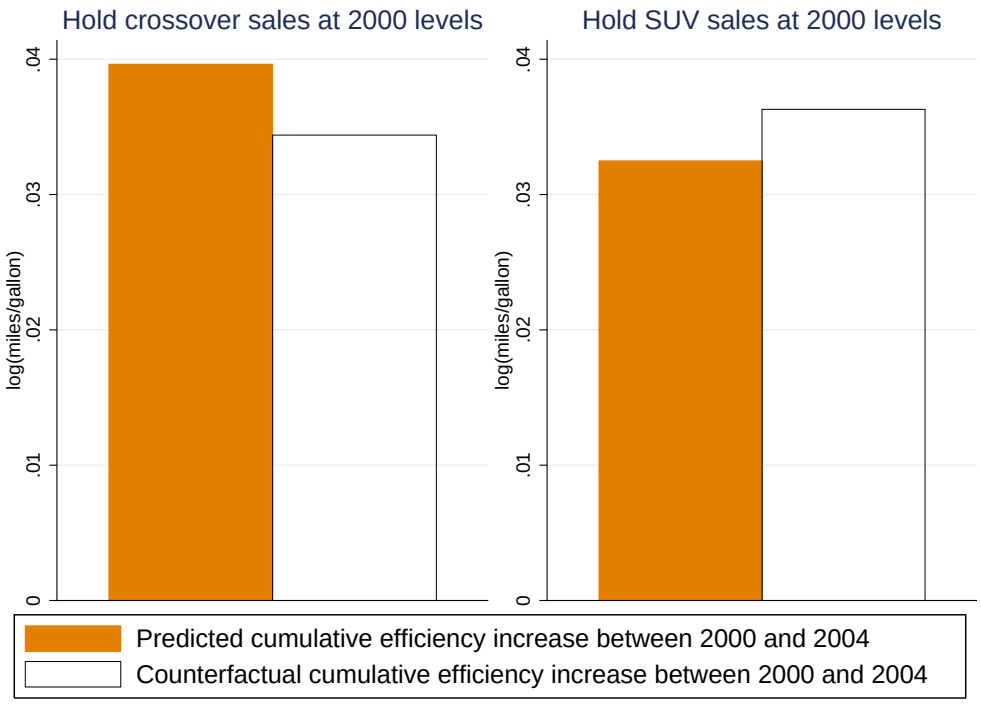

Notes: The colored bars show the mean cumulative predicted efficiency increase between 2000 and 2004 for crossovers (left panel) and SUVs (right panel). Predicted values are obtained from the estimation of equation (3) reported in column 2 of Table 4. The clear bars show the cumulative counterfactual efficiency changes for crossovers and SUVs. The counterfactual holds fixed crossover sales at 2000 levels and uses equation (3) to predict the counterfactual efficiency change for each crossover and SUV between 2000 and 2004.

Figure 9: **Effect of Feebate on Efficiency**

A. Distribution of Efficiency

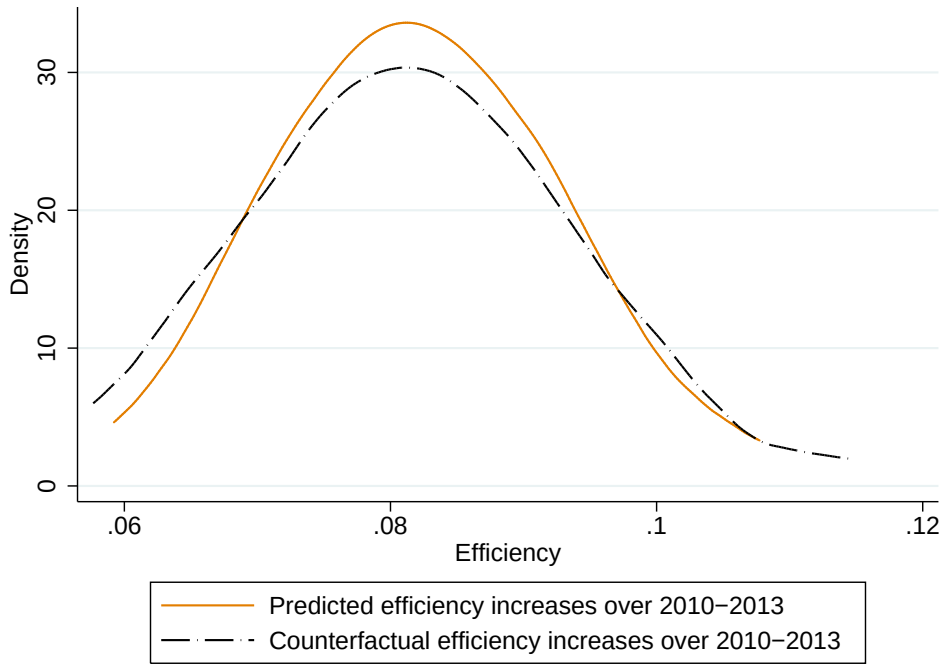

B. Correlation of Fuel Economy and Efficiency

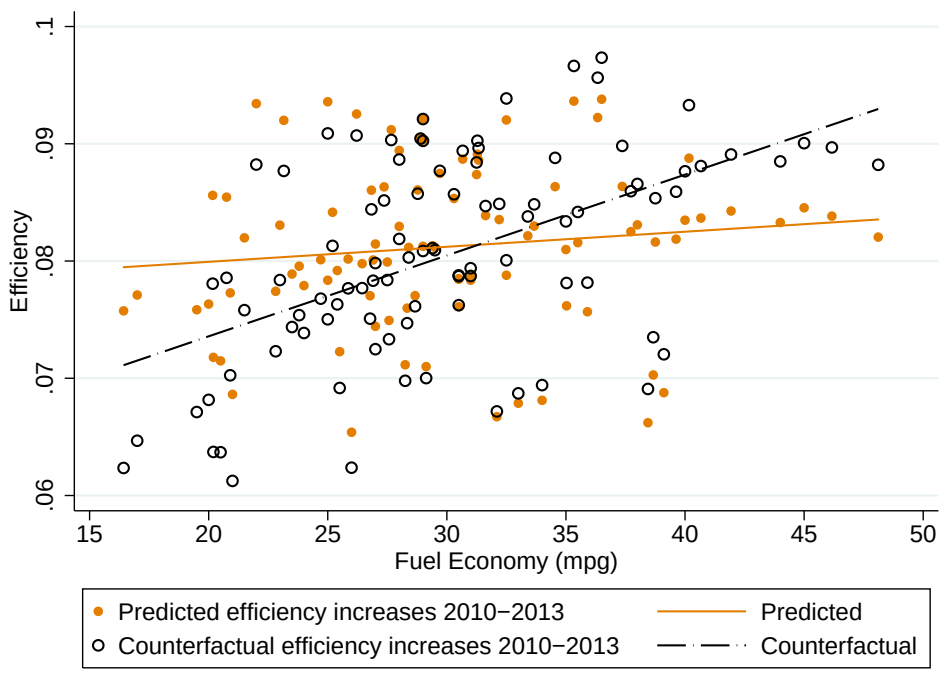

Notes: For each observation in equation (3), the frontier is predicted using the estimates reported in column 2 of Table 4. The counterfactual efficiency of each vehicle is computed from the market size caused by introducing a feebate of $(1/e_{jt} - 1/e_t) \times 1.53$, where e_{jt} is the fuel economy of model j in year t and e_t is the harmonic mean of fuel economy in year t. Panel A shows the estimated density functions of cumulative predicted and counterfactual efficiencies over the period 2010 through 2013. Panel B is a scatter plot of efficiency and fuel economy for each model in the sample. The solid dots represent cumulative predicted efficiency and the circles represent cumulative counterfactual efficiency. The two lines are the linear prediction of efficiency on fuel economy.

Tables

Table 1: **Average Vehicle Characteristics over 1997-2013**

Model Year	Fuel economy (miles per gallon)	Horsepower	Torque (newton-meters)	Weight (pounds)	Number of cylinders
1997	25.4	184	301	3607	6.0
2000	24.9	201	317	3746	6.2
2005	24.9	232	344	4028	6.3
2010	26.3	262	368	4230	6.2
2013	29.1	275	377	4226	6.0

Note: The table reports the sales-weighted average of fuel economy (in miles per gallon), horsepower, torque (maximum torque in newton-meters), weight (in pounds), and number of cylinders for the indicated years.

Table 2: **Estimated Trade-offs Between Fuel Economy and Other Characteristics**

Dependent variable: Log fuel economy	Passenger Cars	Light-Duty Trucks
Log horsepower	-0.224***	-
	(0.014)	-
Log torque	-	-0.157***
	-	(0.016)
Log weight	-0.317***	-0.424***
	(0.037)	(0.034)
Diesel	0.336***	0.260***
	(0.017)	(0.015)
Manual transmission	0.008	-0.004
	(0.004)	(0.004)
Flex fuel	-	-0.272***
	-	(0.012)
Observations	8676	15836
R-squared	0.95	0.93

* p<0.10 ** p<0.05 *** p<0.01.

Notes: The table reports coefficient estimates from equation (2), with standard errors in parentheses, clustered by model and model year. Observations are by model year and model version. The sample in column 1 includes passenger cars and the sample in column 2 includes light-duty trucks. In addition to the reported coefficients, the regressions include model by model year interactions, fixed effects for the number of cylinders, and fixed effects for the number of doors, similarly to Klier and Linn (2016).

Table 3: **Estimated Efficiency for High- and Low-Selling Vehicles**

	High-Selling Vehicles			Low-Selling Vehicles		
	Efficiency in		Cumulative	Efficiency in		Cumulative
Time Period	starting year	ending year	change	starting year	ending year	change
1997-2000	0	0.017	0.017	0	0.0126	0.012
2001-2005	0.019	0.069	0.050	0.006	0.057	0.050
2006-2009	0.067	0.140	0.072	0.078	0.134	0.056
2010-2013	0.156	0.239	0.083	0.160	0.232	0.072

Notes: Efficiency is estimated by model, market segment, and model year in equation (2), using the specification reported in Table 2. Models are assigned one of two categories depending on whether their sales are above the median sales in the initial year of the indicated time period. The table reports the mean estimated efficiency across the two groups and time periods in the first and last years of each period, as well as the cumulative change in mean efficiency over the time period.

Table 4: **Estimation Results: Effect of Market Size and Fuel Costs on Efficiency**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Dependent variable: Efficiency								
Log sales	0.006***	0.030***	0.030***	0.034***	0.029***	0.025***	0.032***	0.030***
	(0.001)	(0.007)	(0.007)	(0.005)	(0.007)	(0.007)	(0.008)	(0.008)
Fuel costs	-0.410***	-0.168		0.201	-0.178*	-0.153	-0.178	-0.306***
	(0.052)	(0.110)		(0.162)	(0.099)	(0.144)	(0.127)	(0.115)
CAFE Stringency							0.048	-0.126
							(0.033)	(0.087)
Fuel costs×CAFE Stringency								1.306**
								(0.648)
Panel B. First Stage Estimate. Dependent variable: Log sales								
Potential market size (log)		0.139***	0.112***	0.161***	0.117***	0.179***	0.115***	0.159***
		(0.037)	(0.039)	(0.037)	(0.040)	(0.036)	(0.039)	(0.042)
Fuel costs		-12.017***		-14.133***	-8.800***	-16.544***	-10.759***	-10.527***
		(1.155)		(1.812)	(1.164)	(1.570)	(1.277)	(1.717)
If market size is imputed		-0.391***	-0.385***	-0.389***	-0.506***	-0.521***	-0.398***	-0.432***
		(0.053)	(0.056)	(0.053)	(0.053)	(0.053)	(0.053)	(0.058)
Estimated by:	OLS	IV, Baseline	IV	IV	IV	IV	IV	IV
Brand fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Brand fixed effects × linear time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Segment fixed effects				Yes				
Segment fixed effects× time trend				Yes				
Brand fixed effects× quadratic time trend					Yes			
Brand fixed effects× linear time trend× truck class						Yes		
Observations	2722	2722	2727	2722	2722	2722	2722	2727
RMSE	0.07	0.07	0.07	0.06	0.06	0.06	0.07	0.07
F (1st stage excl. var.)	NA	32.69	31.76	29.73	31.96	43.42	27.40	25.96

* p<0.10 ** p<0.05 *** p<0.01.

Notes: The table reports coefficient estimates from equation (3), with bootstrapped standard errors in parentheses, clustered by brand (make). Observations are by model and model year. Column 1 is estimated by ordinary least squares (OLS) and columns 2-8 are estimated by instrumental variables, using potential market size and the imputation dummy as instruments according to equation (4). The bottom of the table reports the F statistics of a joint test of the significance of the excluded variables. All regressions include brand fixed effects, year fixed effects, and brand fixed effects interacted with a linear time trend. Column 4 includes a set of segment fixed effects with a linear time trend. Column 5 includes brand fixed effects interacted with a quadratic time trend. Column 6 includes the triple interaction of brand fixed effects by light-duty truck class by linear time trend. Column 7 includes the fuel economy stringency variable described in Klier and Linn (2016), and column 8 includes the interaction of this variable with fuel costs.

Table 5: **Alternative Methods for Estimating Efficiency**

Efficiency estimated by:	(1) Model by year (baseline)	(2) Platform by year	(3) Model by platform generation	(4) Model by model generation	(5) Model by year (3-yr moving average)	(6) Lagged log sales	(7) NHTS 2009 Round Only
Log sales	0.030***	0.041***	0.039**	0.035***	0.040***	0.040***	0.025***
	(0.007)	(0.009)	(0.018)	(0.013)	(0.010)	(0.009)	(0.010)
Fuel costs	-0.168	-0.033	-0.554***	-0.560***	0.083	0.036	0.062
	(0.110)	(0.144)	(0.171)	(0.200)	(0.141)	(0.138)	(0.099)
Brand fixed effects	Yes		Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Brand fixed effects × linear time trend	Yes		Yes	Yes	Yes	Yes	Yes
Company fixed effects		Yes					
Company fixed effects × linear time trend		Yes					
Observations	2722	1953	526	541	2084	2384	2727
RMSE	0.07	0.07	0.06	0.07	0.06	0.07	0.06
F (1st stage excl. var.)	32.69	16.23	7.74	11.64	20.32	27.39	25.33

* p<0.10, ** p<0.05, *** p<0.01.

Notes: The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by brand (make). In column 2 efficiency is estimated by platform and model year. In column 3 efficiency is estimated by model and platform generation. In column 4 efficiency is estimated by model generation and model year. In column 5 efficiency is estimated by model and model year, as in the baseline, but the dependent variable is the three-year moving average of efficiency. Column 6 includes the one-year lag of market size rather than contemporaneous market size, as well as lagged fuel cost, potential market size, and impute dummy. Column 7 uses potential market size constructed using only the 2009 NHTS survey round. In all columns, the independent variables are aggregated to match the aggregation of the dependent variable. All regressions are estimated by instrumental variables using potential market size as an instrument, as in Table 4.

Table 6: **Additional Channels and Heterogeneity**

	(1) Baseline	(2)	(3)	(4)	(5)	(5) Truck	(7) US firm
Log sales	0.030***	0.034	0.023*	0.028*	0.031***	0.026	0.035*
	(0.007)	(0.023)	(0.011)	(0.011)	(0.007)	(0.022)	(0.012)
Fuel costs	-0.168	1.156	-0.239	-0.237	-1.026***	0.408	-0.157
	(0.110)	(1.305)	(0.153)	(0.150)	(0.075)	(2.574)	(0.141)
Efficiency of		3.527					
competing models		(3.058)					
Efficiency of			-1.722				
models by same brand and segment			(1.233)				
Knowledge stock				0.003			
				(0.009)			
Log price					0.089***		
					(0.012)		
Log sales×truck						-0.004	
						(0.018)	
Log sales×US firm							-0.051
							(0.075)
Brand fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Brand fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
× linear time trend							
Observations	2722	2722	2722	2318	2722	2727	2727
RMSE	0.07	0.06	0.06	0.06	0.06	0.06	0.06
F (1st stg. excl. var.)	32.69	25.17	22.92	20.91	35.32	22.07	21.90

* p<0.10 ** p<0.05 *** p<0.01.

Notes: The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by brand (make). Column 1 repeats the baseline regression from column 2 of Table 4. We assign each manufacturer to one of three technology groups: Japanese, US, and other. Column 2 includes the average efficiency by technology group as an independent variable. This variable is instrumented using the corresponding average potential market size of those models. Column 3 includes the average efficiency of other models sold under the same brand in the same market segment, using the average potential market size as an instrument. Column 4 includes the manufacturer's knowledge stock, which is the cumulative number of fuel-saving patents that a parent company has applied for. Column 5 includes the log of the vehicle's price as an independent variable. Columns 1-3 and 5-7 include observations from 1997 to 2013 and column 4 includes observations from 1997 to 2010. All regressions are estimated by instrumental variables using potential market size as an instrument. Columns 2 and 3 use the potential market size of the corresponding vehicles to instrument for competing models and models sold under the same brand. Column 6 includes the interaction of market size with a dummy variable for light trucks, and the corresponding instrument. Column 7 includes the interaction of market size with a dummy for US-based manufacturers, and the corresponding instrument. All regressions are estimated by instrumental variables using potential market size as an instrument, as in Table 4.

Appendix

Figure A.1: **Effect of 2003-2007 Gasoline Price Increase on Efficiency**

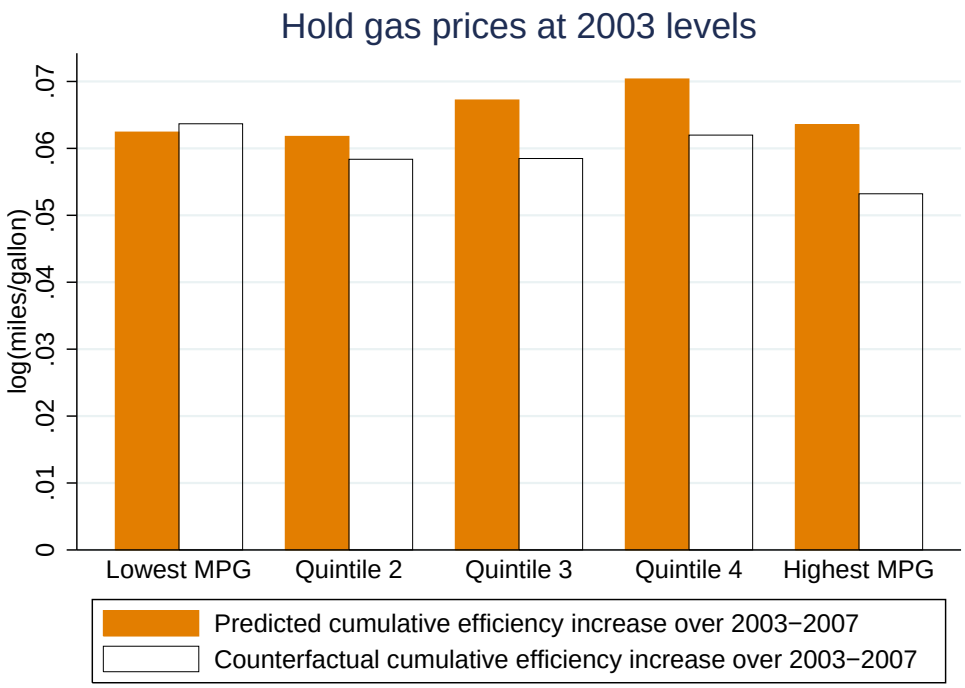

Notes: For each observation in equation (3), the frontier is predicted using the estimates reported in column 3 of Table 4. All observations are assigned to a fuel economy quintile based on the fuel economy distribution across observations between 2003 and 2007, using each vehicle model's initial fuel economy when the model enters the market. The predicted frontier in each colored bar is the mean cumulative predicted efficiency change between 2003 and 2007 for each quintile. The clear bars show the cumulative counterfactual efficiency change by quintile. Counterfactual efficiency changes are computed by holding fixed fuel prices at 2003 levels and using equations (3) and (4) to predict the efficiency change for each observation between 2003 and 2007.

Figure A.2: **Changes in Demographics, 1980-2015**

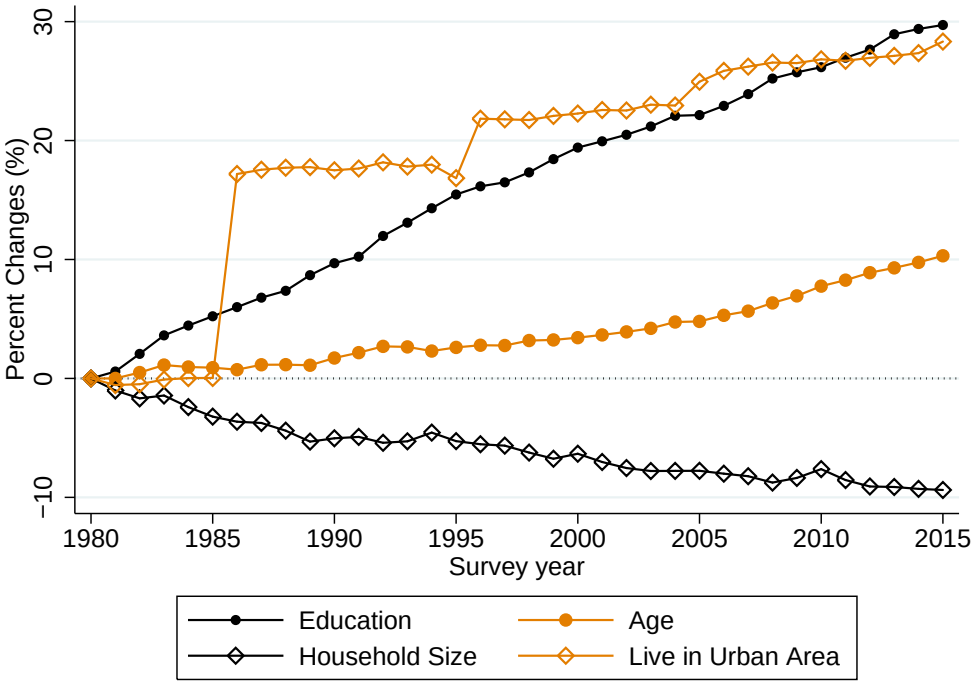

Notes: Using household survey weights from the CPS, we compute the weighted average of each demographic variable by year. The figure plots the percentage change since 1980 of each variable.

Table A.1: **Definitions of Demographic Groups**

Group Number	Age (years)	Household Income (thousand nominal dollars)	Education (years)	Household Size	Urban	Census Division
1	0-34	0-25	0-12	1	urban	New England
2	35-54	25-50	12+	2	not urbanized	Middle Atlantic
3	55+	50-75		3		East North Central
4		75-100		4		West North Central
5		100+		5+		South Atlantic
6						East South Central
7						West South Central
8						Mountain
9						Pacific
No. of Groups	3	9	2	5	2	9
Total Number of Groups:						2700