

# The Effects of Fuel Prices and Vehicle Sales on Fuel-Saving Technology Adoption in Passenger Vehicles

Thomas Klier\*      Federal Reserve Bank of Chicago

Joshua Linn      University of Maryland and Resources for the Future

Yichen Christy Zhou      Clemson University

January 2020

## Abstract

Although economic theory suggests that both sales and fuel costs affect technology adoption by vehicle manufacturers, there is very little empirical evidence on either effect. We document a strong connection between a vehicle's sales and its energy efficiency. Using a demographics-driven demand shifter to isolate demand-side changes in sales, we find that a one standard deviation increase in sales raises efficiency by 0.2 percent, compared with a mean improvement rate of 1.4 percent per year between 1997 and 2013. Higher fuel prices also increase technology adoption directly by increasing willingness to pay for fuel cost savings. The results have two implications: manufacturers will continue to focus technological improvements on top selling vehicles; and fuel taxes will have larger effects on technology adoption than fuel economy standards and feebates.

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

JEL classification numbers: L62, Q4, Q55, O31

---

\*Thomas Klier ([Thomas.Klier@chi.frb.org](mailto:Thomas.Klier@chi.frb.org)) is a senior economist and research advisor at the Federal Reserve Bank of Chicago. Joshua Linn ([linn@rff.org](mailto:linn@rff.org)) is a senior fellow at Resources for the Future (RFF). Christy Zhou ([yichen2@clemson.edu](mailto:yichen2@clemson.edu)) is an assistant professor at Clemson University. Samuel Goldberg provided excellent research assistance.

# 1 Introduction

Improving vehicle fuel economy is a central part of worldwide efforts to reduce the risks of climate change. In the United States, passenger vehicles account for about 15 percent of greenhouse gas emissions and half of transportation sector emissions (IPCC, 2014).

Meeting near- and long-term emissions targets requires substantial technology adoption (Knittel, 2012). An extensive literature (e.g., Goldberg (1995), Klier and Linn (2012), Jacobsen (2013), Roth (2015), Reynaert (2015)) has examined the welfare consequences of fuel economy and greenhouse gas standards for passenger vehicles. Klier and Linn (2016) find that these standards have increased the rate of technology adoption. However, the vast majority of this literature has assumed that technology adoption is exogenous and thus has missed a potentially important aspect of the standards.

The literature suggests that several factors may affect technology adoption. According to NRC (2015), variable costs of technology and resulting fuel savings affect adoption. Newell et al. 1999 document the effects of consumer demand and energy prices on innovation and technology adoption in air conditioners and other industries.<sup>1</sup> More broadly, the literature also suggests that a vehicle model's sales can affect adoption of technology through directed technical change (Acemoglu, 2002), fixed costs of technology adoption with variable markups (Berry et al., 1995), or learning by doing in vehicle production.

In this paper, we focus on a manufacturer's decision to adopt fuel-saving technology across its individual vehicle models. We show that fuel costs and vehicle demand substantially affect this decision, and we discuss implications for greenhouse gas emissions policy.

Motivated by the technology literature, we distinguish three demand-side drivers of fuel-saving technology adoption. The first channel is the *direct effect of fuel costs*. A vehicle's fuel costs increase with the price of fuel and decrease with fuel economy. An increase in fuel prices raises a consumer's willingness to pay for a given fuel economy improvement (Klier and Linn, 2012), and thus the profitability of adopting fuel-saving technology.

The second and third demand drivers operate via sales. As noted above, an increase in a vehicle's sales can cause technology adoption because of fixed costs of adoption or other reasons. Therefore, a demand shock that increases a vehicle's sales could also cause the adoption of fuel-saving technology. Busse et al. (2013) and Allcott and Wozny (2014) demonstrate that high gasoline prices raise the market shares of vehicles with high fuel economy. Thus, an increase in fuel prices can raise sales of vehicles with high fuel economy,

---

<sup>1</sup>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.

leading to more technology adoption for those vehicles. We label this demand driver the *indirect effect of fuel costs* because it operates via sales.

The distinction between the direct and indirect effects is that the direct effect holds sales fixed. For the consumers who would purchase a particular vehicle with low fuel prices, an increase in fuel prices raises their willingness to pay for fuel-saving technology, stimulating technology adoption. In addition, if that vehicle has high fuel economy, the increase in fuel prices results in additional consumers wanting to purchase the vehicle. The increase in sales creates the indirect effect. For example, an increase in gasoline prices raises fuel costs for all vehicles, such as the Toyota Prius. Higher fuel costs raise the incentive for Toyota to make the Prius more efficient, representing the direct effect. The gasoline price increase also raises the market share of the Prius, since it has relatively high fuel economy, increasing sales and further incentivizing technology adoption. Thus, for the Prius the direct and indirect effects both encourage more technology adoption. As we explain below, the two effects can work in the opposite direction for vehicles with low fuel economy, such as the Toyota Tundra (a pickup truck).

The third driver includes any other demand shock that increases sales, besides fuel prices. We denote this effect as the *direct effect of sales*. For example, the increase in demand for crossover vehicles during the 2000s and 2010s raised sales of those vehicles. Adding fuel-saving technology raises consumer willingness to pay for the vehicle, allowing the firm to raise the vehicle's price. The greater the vehicle's sales prior to adoption, the greater the revenue increase, making it more likely that manufacturers adopt technology when crossover demand is high than when it is low.

We use a novel identification strategy to estimate empirically the magnitudes of the three effects. We focus on fuel-saving technology for the internal combustion engine, including gasoline- and diesel-powered engines, and associated transmissions. We use unique data on consumer demographics (consumer preferences for new vehicles), and vehicle-level characteristics. Notwithstanding the media attention around electric vehicles and other alternative technologies, the internal combustion engine still accounted for about 99 percent of new vehicles sold in the United States in 2015 and about 98 percent in 2018 (authors' calculations).

We begin the analysis by defining a vehicle's powertrain efficiency, which is our measure of fuel-saving technology and is distinct from the vehicle's fuel economy (miles per gallon, mpg). For a given level of powertrain 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 shifts the frontier and can increase fuel economy without affecting other characteristics. We define the increase

in efficiency that results from technology adoption as the increase in fuel economy that is feasible holding other vehicle characteristics constant. 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 powertrain efficiency of each vehicle model by model year from 1997 to 2013 similarly to [Klier and Linn \(2016\)](#).

Firms generally choose efficiency by deciding which of a set of existing technologies to install. Consequently, the current or expected sales may affect the choice of efficiency. The main empirical challenge is the endogeneity of a vehicle's sales. 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. Improving a vehicle's efficiency may increase its demand, causing sales to increase and resulting in reverse causality. Furthermore, omitted supply variables, such as a vehicle's production costs, can be correlated with both its sales 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. The IV is a demand shifter that captures changes in demand for a particular vehicle, relative to demand for other vehicles, which arise from changes in consumer demographics over time. For example, larger households tend to purchase more minivans than smaller households. The decrease in the share of large households in the United States over the sample period has reduced demand for minivans relative to other market segments. To construct our instrument, we use consumer preferences by demographic group that are measured at a specific point in time, combined with temporal variation in demographics. [Acemoglu and Linn \(2004\)](#) and [DellaVigna and Pollet \(2007\)](#) have similarly used demographic trends as exogenous determinants of sales in the pharmaceuticals and toys markets (other papers, such as [Blundell et al. \(1999\)](#), have used pre-sample information to address endogeneity). The validity of the instrument rests on a) a positive correlation between the IV and actual sales; and b) the exclusion restriction, that the IV affects technology only via sales. We document a strong positive correlation between the IV and actual sales, and provide strong evidence supporting the exclusion restriction. Specifically, because consumer purchasing patterns are held fixed in constructing the instrument, changes in supply-side factors that affect purchasing patterns, such as changes of a vehicle's position in product space, do not affect the instrument. Also, we show that the instrument isolates variation in the potential demand for a vehicle that is uncorrelated with supply-side factors that affect sales, such as imperfect competition in non-price vehicle attributes. In keeping with the recent literature (e.g., [Acemoglu et al. \(2016\)](#)), we identify the effects of fuel costs on technology adoption assuming that fuel prices are exogenous to the market.

Sales positively affect a vehicle's efficiency. A one standard deviation increase in sales, which corresponds to about a 10 percent increase, raises a vehicle's efficiency by 0.2 percent. This estimate constitutes a substantial and statistically significant increase relative to the observed average annual efficiency increase of about 1.4 percent between 1997 and 2013. As an alternative interpretation of the estimated magnitude, the sales effect implies annual consumer benefits of \$100-400 million (2010 USD) for consumers who purchase high- rather than low-selling vehicles. In addition to sales, we test whether a vehicle's efficiency responds to the efficiency of competing vehicles or to the manufacturer's stock of efficiency-related patents. We find some effects of competing vehicles, but these effects are less precisely estimated than the primary effects of sales and fuel costs on efficiency. We also find that the main results are robust to alternative functional forms and constructions of the instrument.

Fuel prices affect technology adoption both directly and indirectly via sales. After controlling for sales, fuel prices have a positive and statistically significant effect on technology adoption. A fuel price increase causes average fuel economy to increase for two reasons: first, by increasing markets shares of vehicles with high fuel economy (Busse et al., 2013), and second, by inducing fuel-saving technology adoption. It turns out that the estimated effects are similar in magnitude, demonstrating the economic significance of the direct fuel price effect.

We illustrate the magnitudes of the sales and fuel cost effects using three sets of simulations. First, we compare the indirect and direct channels through which fuel costs affect efficiency, by focusing on the gasoline price increase that occurred between 2003 and 2007. The indirect effect of sales works as follows. The 80 percent increase in real gasoline prices in that period raised the sales of vehicles with high fuel economy relative to vehicles with low fuel economy. In turn, the changes in sales caused efficiency of the lowest-fuel-economy vehicles to be lower than they would have been 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. In contrast, the direct fuel cost effect works in the opposite direction, causing more technology adoption for low fuel economy vehicles. The increase in gasoline prices raises fuel costs disproportionately more for low fuel economy vehicles than high fuel economy vehicles. This change raises the willingness to pay for fuel cost savings more for low fuel economy vehicles than high fuel economy vehicles, causing more technology adoption for low fuel economy vehicles. In the simulations, the direct fuel cost effect was about twice as large as the indirect sales effect.

Second, we show that demographics affected the efficiency distribution across models in different market segments. The overall shifts in demographics between 1980 and 2013 caused a shift in cumulative efficiency improvements away from light-duty trucks and toward cars.

This effect occurred simultaneously with other demand- or supply-side effects on relative efficiencies of cars and light trucks, such as changes in gasoline prices and fuel economy standards, which affected consumer demand and manufacturer technology adoption.

Third, we find that changes in sales for crossovers and sport utility vehicles (SUVs) have affected technology adoption. Between 2001 and 2004, per-model sales of crossovers increased sharply and per-model sales of SUVs decreased sharply. The increase in crossover sales raised crossover efficiency and the decrease in SUV sales reduced SUV efficiency, relative to a counterfactual in which sales remained unchanged.

The empirical results have two main implications for policies aiming to improve passenger vehicle fuel economy and reduce greenhouse gas emissions. First, the strong sales effect implies that manufacturers will continue to improve the efficiency of vehicles with internal combustion engines, which currently dominate the market. Most analysts argue that vehicle electrification, combined with decarbonization of electricity generation, is necessary for substantially reducing transportation greenhouse gas emissions. The sales effect will increase consumer demand for internal combustion engines, and increase the challenges faced by alternative-fuel vehicles to gain market share, compared to a hypothetical scenario in which there is no sales effect for internal combustion engines. Existing welfare analyses of fuel economy and greenhouse gas standards (e.g., [Jacobsen \(2013\)](#)) do not account for this sales effect when characterizing the technology adoption caused by standards, and our results suggest that future analysis of the standards should do so.

Second, as noted above, the literature has compared the welfare effects of policies that affect fuel prices (such as raising gasoline taxes) and policies that affect new vehicle fuel economy (such as standards or feebates). For a given emissions reduction, the relative costs of these policies depend on consumer and manufacturer responses. A comparison of these policies typically does not account for the effects of sales and fuel costs on technology adoption. Raising fuel prices differentially affects fuel costs of vehicles in the market and causes consumers to shift toward vehicles with higher fuel economy and away from vehicles with lower fuel economy. Therefore, the fuel price-based policies affect adoption via the fuel cost and sales effects, analogously to the 2003–2007 gasoline price increase. In contrast, starting from a particular market equilibrium, tightening fuel economy standards or introducing a feebate affects vehicle prices and sales, but does not affect fuel prices.

Such policy-induced changes in sales and fuel costs in turn affect the cross-sectional efficiency distribution. In our data, efficiency is positively correlated with fuel economy. A feebate would strengthen this positive correlation by shifting sales to vehicles with high fuel economy and causing greater cumulative efficiency improvements for those vehicles than for vehicles with lower fuel economy. Raising fuel taxes or introducing a carbon tax creates

opposing fuel cost and sales effects. Such policies do not strengthen the positive correlation as much as do vehicle-based policies. The fuel price- and vehicle-based policies therefore cause different changes in vehicle technology, and because consumer choices depend on technology, the sales and fuel cost effects imply differing welfare costs of the two types of policies.

## 2 Data and Summary Statistics

### 2.1 Data

We assemble three data sets for the empirical analysis. The first 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.<sup>2</sup> A model year begins in September of the previous calendar year and ends in August of 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. 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 \(2016\)](#)).<sup>3</sup>

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 refers 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 vehicles, as well as flex-fuel vehicles; these vehicles accounted 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 those data 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 with the sales and characteristics data.

---

<sup>2</sup>A change in reporting in 1997 prevents us from extending the sample to earlier years.

<sup>3</sup>The recent Volkswagen scandal raises some concerns about the accuracy of laboratory testing of vehicle emissions because the results are used to assess compliance with fuel economy and emissions standards. 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. In a few instances US laboratory tests have overstated 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 gas emissions. 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 and provides a disincentive for manufacturers to cheat on fuel economy ratings.

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

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.3 for definitions of the groups). The age and education groups are based on the attributes of the household head. 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 vehicle model name and demographic group. For example, we compute the average number of new Honda Civics 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 that we use for the NHTS.

## 2.2 Summary Statistics

In this subsection we present summary statistics of 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 sales on 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 entire sample. 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 waves and weight observations



using household survey weights. The figure indicates a substantial amount of variation in purchase behavior across groups. 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 lower-income households. Geographic variables are also correlated with purchase behavior: households in urban areas and the Northeast are much less likely to buy pickup trucks than other households. The demographic variables are correlated with one another; for example, households with high incomes tend to be well educated.

Figure 4 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 temporal variation with the variation in purchasing patterns across demographic groups illustrated in Figure 2 to construct the instrumental variable for sales. 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 sales. For example, after the initial introduction of the crossover segment in the late 1990s, the market share of crossovers increased from the late 1990s through the 2000s. This is consistent with the facts 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 vehicle market. Manufacturers continually redesign their vehicles, improving powertrain technology and other attributes that consumers value. Many vehicle models experience major redesigns at regular intervals, commonly every five to seven years. During a redesign, the manufacturer may make major changes to the powertrain, cabin, cargo space, or exterior. In between redesigns, manufacturers commonly make smaller changes to exterior design or to the powertrain, offering new options such as paint color or increasing the number of transmission speeds.

These alterations yield a process of steady technology adoption over time. Figure 3 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 U.S. Environmental Protection Agency (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 because manufacturers began replacing this technology with more advanced fuel injection technologies.

### 3 Empirical Strategy

In this section we review the literature on technology adoption, which motivates estimating a reduced-form relationship among vehicle efficiency, sales, and fuel costs. Then, we estimate efficiency of each vehicle model in the sample, and finally we derive the estimating equation and explain how the IV strategy addresses the identification issues.

#### 3.1 Possible Channels of the Sales and Fuel Cost Effects

There are several possible channels through which sales could affect technology adoption. First, manufacturers face fixed costs when improving a vehicle’s efficiency. Fixed costs may arise because of the need to redesign and test the vehicle before commencing full-scale production (Blonigen et al., 2013). These fixed costs are distinct from the additional production costs associated with adopting technologies. Additional production costs can arise because of the need for new components or because of greater complexity of the production process.<sup>4</sup> Consider a hypothetical manufacturer that sells two types of vehicles that are identical except that the first has larger expected sales than the second. Adopting technology to increase efficiency would require the same increase in fixed and production costs for the two vehicles. Because the first vehicle has larger expected sales, the first manufacturer can spread the fixed costs across a larger sales volume, making it more likely that the manufacturer adopts the technology for that model.

Learning by doing is a second possible channel linking sales and adoption. With learning by doing, the marginal costs of producing a vehicle decline with cumulative production. For example, Levitt et al. (2013) document large reductions in production costs and improvements in product quality following the introduction of a new vehicle at an automobile assembly plant. Learning by doing would imply greater cost reductions for vehicles with higher production volume and sales. Therefore, a manufacturer would be more likely to adopt technology for vehicles that are anticipated to have high sales. A third possible channel is directed technical change (e.g., (Acemoglu, 2002)), in which the profits accruing from innovation depend on a product’s market size and price. Innovation and efficiency improvements raise market size, which raises the returns to future innovation, and creates a positive feedback among market size, innovation, and technology adoption. In the directed technical change literature, market size corresponds roughly to sales.

These arguments suggest a positive effect of sales on technology adoption. However, other considerations would imply a weaker or perhaps negative relationship between sales

---

<sup>4</sup>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.

and technology adoption. For example, in the long run learning by doing could create lock-in. Suppose a manufacturer produces a high-selling vehicle and learning by doing has reduced its cost and price. If improving efficiency would require a new learning process, the initial cost of producing the more efficient vehicle could make it unprofitable to adopt the new technology. Moreover, competitive pressure could amplify or diminish each of these channels, depending on the distribution of consumer preferences and other factors.

In short, theory predicts the relationship between sales and technology adoption to be either positive or negative. Fixed costs of adoption and directed technical change suggest a positive relationship, whereas learning by doing could imply either a positive or a negative relationship. This ambiguity, and the absence of empirical evidence on these channels, motivates our reduced-form analysis that estimates the cumulative magnitude of these possible channels.

Fuel costs may affect technology adoption via two channels. First, consider a manufacturer that sells a particular type of vehicle. If the manufacturer raises the efficiency and fuel economy of the vehicle, consumers who would have purchased the vehicle without the fuel economy increase now have higher willingness to pay for the vehicle. We refer to this effect as the *direct effect of fuel costs*, and it captures the incentive to increase efficiency holding fixed the vehicle's sales. Second, the same efficiency increase may cause additional consumers to purchase the vehicle. We refer to this channel as the *indirect effect of fuel costs* via sales; this effect is indirect because it operates via vehicle demand.

### 3.2 Estimating Efficiency

The empirical objective is to estimate the effects of sales and fuel costs on efficiency. In this subsection we describe the construction of the dependent variable, which is efficiency.

The available data do not contain efficiency per se, but they 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 of vehicle characteristics. Efficiency is the amount of useful work or energy that the powertrain 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 powertrain 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 the fuel economy–horsepower–weight space. The frontier represents the maximum fuel economy that can be achieved given any particular level of horsepower and weight. That is, for a particular level of efficiency, as the manufacturer moves along the frontier, it can trade off fuel economy for weight and horsepower.

This framework yields a straightforward identification of efficiency improvements over time. 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. Specifically, we estimate an equation similar to [Klier and Linn \(2016\)](#):

$$\ln e_{jt} = \lambda_h \ln h_{jt} + \lambda_w \ln w_{jt} + \tau_{mt} + X_{jt}\delta + \varepsilon_{jt}, \quad (1)$$

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,  $\tau_{mt}$  is a set of interactions of model by model year,  $X_{jt}$  includes a vector of vehicle attributes,  $\varepsilon_{jt}$  is an error term, and the  $\lambda$ s and  $\delta$  are coefficients to be estimated. The coefficients on horsepower and weight govern the technological tradeoffs among fuel economy, horsepower, and weight. 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.<sup>5</sup> 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,  $\tau_{mt}$ , as the average efficiency of vehicles belonging to the model and sold in model year  $t$ . The difference between  $\tau_{mt}$  and  $\tau_{m(t-1)}$  is the change in efficiency of model  $m$  between model years  $t - 1$  and  $t$ . Equation (1) thus allows us to identify changes in efficiency over time, where efficiency is measured in units of log fuel economy.<sup>6</sup> We expect efficiency to increase over time as innovation reduces the cost of improving efficiency.

Before presenting the results from estimating equation (1), we briefly discuss identification and potential sources of bias. Equation (1) characterizes a technological relationship between

---

<sup>5</sup>These interactions are not meant to control for macroeconomic shocks or anything else; instead, they estimate efficiency improvements over time after controlling for changes in other variables that affect fuel economy (such as fuel type, horsepower, and weight). These interactions are defined at as disaggregated a level as possible because the estimated interactions become the dependent variable in equation (2).

<sup>6</sup>Equation (1) includes the assumption that the shape of the frontier does not change over time. Other papers that have estimated a frontier, such as [Knittel \(2012\)](#), have made the same assumption. Our results are robust to allowing the frontier to vary over time, such as by interacting the log of horsepower with a linear time trend. Also, the assumption that the shape does not change over time does not imply that fuel economy, horsepower, and weight increase proportionately over time. Instead, the manufacturer increases these attributes over time depending on other factors such as the tradeoffs and changes in consumer willingness to pay for the attributes.

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 (1), but in this context that would not bias the coefficients as long as the omitted variables affect fuel economy via horsepower, weight, or other included variables, 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. Note that efficiency improvements in the estimated  $\tau_{mt}$  may include aerodynamic improvements, because such improvements are not included in the other independent variables.<sup>7</sup>

Table 2 reports the main coefficient estimates from equation (1). Because fuel economy, horsepower, torque, and weight enter equation (1) in logs, the horsepower, torque, 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 35 percent higher fuel economy, and flex-fuel light trucks achieve about 24 percent lower fuel economy than gasoline-powered vehicles. The negative coefficient on flex-fuel vehicles reflects the lower energy content of ethanol compared with gasoline (that is, the coefficient estimate reflects the fact that EPA reports the fuel economy of the flex-fuel version; the coefficient does not reflect a difference in efficiency of flex-fuel and conventional engines). 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.<sup>8</sup>

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 powertrain 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 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 sales have a positive effect on efficiency-improving technology adoption.

---

<sup>7</sup>See Knittel (2012) and Klier and Linn (2016) for additional discussion of identification of equation (1).

<sup>8</sup>We have estimated versions of equation (1) that allow the horsepower and weight coefficients to vary across vehicles or over time. In general, we do not find such variation to be statistically significant.

### 3.3 Estimating the Effects of Sales and Fuel Costs on Efficiency

This subsection presents the strategy for estimating the effects of sales and fuel costs on efficiency. To motivate the estimating equation, we conceive of a market for new vehicles, which includes fixed numbers of consumers and manufacturers. Each consumer has a willingness to pay for a set of vehicle attributes, which includes things such as fuel costs and horsepower. To maximize profits, each manufacturer chooses the price, efficiency, horsepower, and other attributes. Fuel economy is determined by the manufacturer’s choices of horsepower and efficiency, according to equation (1). Each consumer chooses the utility-maximizing vehicle, given the attributes and prices of all vehicles in the market.

In equilibrium, the manufacturer’s choice of a vehicle’s efficiency depends on supply and demand-side factors.<sup>9</sup> For example, a reduction in technology adoption costs, perhaps due to learning-by-doing (see section 3.1) increases technology adoption and efficiency. In this model, fuel costs can affect efficiency in two ways. First, holding fixed the set of consumers who purchase a vehicle, an increase in their willingness to pay for fuel economy would increase equilibrium efficiency. Distinct from this effect is the possibility that a change in gasoline prices affects the equilibrium number of consumers who choose the vehicle. In particular, an increase in gasoline prices would raise equilibrium sales for vehicles with relatively high levels of fuel economy, which creates the indirect effect of fuel costs on efficiency that we discussed. Of course, the equilibrium sales response implies that sales are endogenous to fuel costs, and there may be reverse causality from efficiency to sales. We explain below how the IV strategy identifies the effect on efficiency of demand-driven changes in sales, addressing reverse causality. We also discuss how the IV addresses potential omitted variables bias, such as a potential correlation between a vehicle’s unobserved quality and its equilibrium sales.

We begin by estimating the reduced-form relation between sales, fuel costs and efficiency. We assume a log-linear relationship between sales and efficiency that approximates nonlinear forms. The estimating equation is

$$\hat{\tau}_{mt} = \gamma_1 \ln Q_{mt} + \gamma_2 C_{mt} + \phi_t + \phi_{b(m)} + \phi_{b(m)} \times t + \varepsilon_{mt} \quad (2)$$

where  $\hat{\tau}_{mt}$  is efficiency estimated from equation (1) for model  $m$  in model year  $t$ ,  $Q_{mt}$  is sales,  $C_{mt}$  is fuel costs per mile (dollars-per-mile),  $\phi_t$  and  $\phi_{b(m)}$  are sets of year and make (i.e., brand) fixed effects,  $\phi_{b(m)} \times t$  is the interaction of make fixed effects with a linear time trend, and  $\varepsilon_{mt}$  is an error term. The two parameters of interest are  $\gamma_1$  and  $\gamma_2$ , which are the effects of log sales on efficiency (*sales effect*) and fuel costs on efficiency (*fuel cost effect*). As we

---

<sup>9</sup>See Roth (2015) for an example of modeling the choice of fuel economy and Zhou (2016) for an example of modeling the choices of technology adoption and vehicle performance.

discussed in section 3.1, theory implies that the coefficient on log sales could be positive or negative. Note that the sales variable is a proxy for the expected sales at the time that the manufacturer chooses efficiency; we return to the timing of the efficiency choice below. The variable  $C_{mt}$  is the cost of driving the vehicle one mile, which is proportional to the present discounted value of the vehicle’s fuel costs over its lifetime. We identify  $\gamma_2$  by variation in both fuel economy and fuel prices. The coefficient on fuel costs is the direct effect of fuel costs on efficiency, holding sales fixed. We discuss the fuel cost variable construction and identification at the end of the subsection. The year fixed effects control for aggregate demand or supply shocks, and the make fixed effects control for make-level supply or demand shocks, such as consumer perceptions of make-level quality and make-level productivity shocks. The make fixed effects also control for make-level economies of scope that may affect technology adoption, as well as make-specific pass through of technology adoption costs to vehicle prices. The interactions of the make fixed effects with a linear time trend allow these make-specific factors to vary linearly over time. For example, the time trends control for changes in consumer preferences for makes as well as changes in make quality.

Estimating equation (2) by ordinary least squares (OLS) is likely to yield biased estimates for two 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 efficiency. The make fixed effects and time trends control for make-level supply or demand shocks, but efficiency could be correlated with within-make 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 (2). That approach would yield an undesirable interpretation of  $\gamma$ , 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. For example, a firm may have time (i.e., the need to redesign the vehicle over a 5-year cycle) or borrowing constraints that prevent it from increasing efficiency for all models at the same time, in which case it could decide to focus on its best-selling models. Including model fixed effects would identify the choice of when to adopt technology for the model but not which model, and therefore might omit an important role of sales in technology adoption across vehicle models.

We instrument for sales and address both sources of bias. We choose to construct vehicle’s potential sales as our IV, which is a demand shifter that depends both on cross-sectional variation in consumer purchasing patterns at a point of time approximated by sales computed from the NHTS, and time series variation in demographics. The IV isolates changes in sales for an individual vehicle driven by changes in demographics. For example, compared with younger age groups, the older age groups tend to prefer large and luxury cars, and an increase in the population share of older age groups over time would increase demand for large and luxury cars relative to other vehicles. In standard models of imperfect competition, such as Bertrand with differentiated products, the increase in demand causes an increase in equilibrium sales. This implies a positive correlation between a vehicle’s endogenous sales and the IV.

More specifically, we define demographic group cell,  $g$ , by age, income, education, household size, urbanization, and census division (see Appendix Table A.3 for definitions of the groups). To measure purchasing behavior by group, we compute  $q_{mg;s}$  for each of the 2,628 cells as the number of vehicles of model  $m$  purchased per household by demographic group cell  $g$  in NHTS wave  $s$ . We use data from the NHTS years 1995, 2001, and 2009; each cell represents an average of about 28 households per year. The subscript  $s$  reflects the fact that  $q_{mg;s}$  varies across NHTS waves. We define time periods  $s$  based on the NHTS waves: 1997–2000, 2001–2008, and 2009–2013. 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, setting missing values equal to zero. The potential sales,  $\tilde{Q}_{mt;s}$ , is the product of the count of NHTS vehicle purchases per household and the CPS number of households, summed across demographic group cells:

$$\tilde{Q}_{mt;s} = \sum_g (q_{mg;s} \times w_{gt}) \quad (3)$$

Variation in the potential sales over time arises from variation in demographics over time, weighted by the NHTS quantities,  $q_{mg;s}$ . For each vehicle model, we subtract from  $\tilde{Q}_{mt;s}$  the average value for the corresponding model to obtain the instrument,  $\bar{Q}_{mt;s}$ . Subtracting the average value causes the variation of the IV to arise entirely from changes over time in potential sales caused by changes in demographics. Therefore, the IV does not include variation in cross-sectional purchasing patterns that may be correlated with unobserved vehicle attributes. Because  $q_{mg;s}$  does not vary within a period, within-period preference changes do not affect  $\bar{Q}_{mt;s}$ . Therefore, supply-side and demand-side factors that affect sales that are uncorrelated with demographics do not affect instrumented sales. For example, if a decrease in production costs causes vehicles to enter the market in a particular period, such



entry would not affect the instrument because  $q_{mg;s}$  does not respond to changes in supply conditions within a period. Likewise, fuel price-driven changes in vehicle attributes (other than fuel economy) would not affect the IV.

An ideal IV isolates cross sectional preference variation across demographic groups, interacted with time series variation in group size. The IV is based on the assumption that the group-specific population weight  $w_{gt}$  is uncorrelated with demand- and supply-side factors that affect technology adoption. Within each time period, we identify the first stage using cross sectional variation in purchase patterns across demographic groups, interacted with variation in demographic group sizes over time. Implicitly, this approach holds fixed consumer preferences within a time period, but allows preferences to vary across time periods. Temporal 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 4 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 unobserved demand and supply determinants 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.

The log potential sales,  $\ln \bar{Q}_{mt;s}$ , is the IV in the first stage for sales in equation (2)

$$\ln Q_{mt} = \beta_1 \ln \bar{Q}_{mt;s} + \beta_2 C_{mt} + \beta_3 I_{mt}^{imp} + \phi_t + \phi_{b(m)} + \phi_{b(m)} \times t + u_{mt} \quad (4)$$

where  $I_{mt}^{imp}$  is an indicator variable equal to one if the instrument is imputed using make-segment-year means.<sup>10</sup> The instrument yields unbiased estimates of the sales coefficient if it is strongly correlated with sales and is uncorrelated with the error term in equation (2). As we show below, the instrument is a strong predictor of log sales, reducing concerns about weak instruments bias and satisfying the first condition for the validity of the IV.

The second condition is that the IV affects efficiency only via sales and not via other channels. If either equilibrium purchases ( $q_{mg;s}$ ) or demographics ( $w_{gt}$ ) are correlated with unobserved demand or supply shocks, the IV estimates would be biased. In this subsection we provide several arguments for the validity of the IV strategy, and in the next section we provide empirical evidence supporting these arguments.

---

<sup>10</sup>Some vehicle models appear in the sales and characteristics data but not in the NHTS data. Most of these are low-selling models, such as the Chevy Aveo. In these cases, we impute the instruments using make-segment-year-level average NHTS weights.

The variable  $q_{mg;s}$  reflects equilibrium demand and supply conditions during the corresponding NHTS wave, and could be correlated with unobservables for several reasons. First, there could be unobserved preference shocks that persist across periods, and which affect  $q_{mg;s}$ . If the preference shocks also affect efficiency directly, the estimates would be biased. For example, if willingness to pay for fuel economy increases among consumers purchasing a particular vehicle, the manufacturer of that vehicle may adopt technology and increase fuel economy, raising sales.<sup>11</sup>

Second, changes in demographics ( $w_{gt}$ ) could be correlated with demand shocks over time. For example, the population aged during the sample period, and the older age groups' preference for large and luxury cars may also have changed. If manufacturers adjusted vehicle attributes in response to these preference changes, changes in  $w_{gt}$  over time would be correlated with the error term and the IV estimates would be biased.

Third, there could be persistent supply-side shocks that affect  $q_{mg;s}$  and efficiency. For example, innovation may reduce the cost of improving efficiency for one segment more than for others, causing  $q_{mg;s}$  to change over time and be correlated with unobserved supply-side factors affecting equilibrium efficiency.

In these three cases, changes in  $q_{mg;s}$  across periods would be correlated with omitted variables. We can address concerns related to persistent omitted demand and supply shocks by comparing results if we use a different NHTS wave to construct the IV. The weights constructed from the NHTS,  $q_{mg;s}$ , reflect preferences and vehicle attributes in period  $s$ . If changes in  $q_{mg;s}$  across periods for a particular demographic cell are correlated with changes in vehicle attributes or supply conditions, we would observe a correlation between demographics and changes in  $q_{mg;s}$  across NHTS waves. In that case the estimates of equation (2) would depend on whether we use all three NHTS waves, or just one survey wave. In section 4 we show that the results are unchanged if we use either the 1995 or 2009 survey wave rather than all three survey waves.

Demographics may also be correlated with unobserved variables indirectly due to product market competition. Suppose an increase in the share of elderly households raises demand for a particular vehicle. This demand shock could cause other manufacturers to change the attributes of the vehicles that compete with the first vehicle. In turn, such changes could cause the manufacturer of the first vehicle to change its efficiency. Therefore,  $\beta_1$  would be biased if changes in demographics affect technology adoption via sales and non-price competition in the product market. If the underlying changes in  $w_{gt}$  were persistent, they would also have affected  $q_{mg;s}$ . We show that using the single NHTS wave that our instrument

---

<sup>11</sup>Although advertising campaigns may affect preferences, advertising campaigns caused by variation in the instrument would not bias the estimates.

is unlikely to create bias for these reasons. Moreover, these results reduce concerns about omitted and persistent supply-side shocks that may be correlated with  $w_{gt}$ .

In short, the results using a single NHTS wave reduce concerns about bias caused by persistent omitted demand or supply shocks that may be correlated with the instrument. We address temporary demand and supply shocks below.

We interpret  $\beta_1$  as the effect on efficiency of a change in sales induced by a change in potential sales. In practice, changes in potential sales may affect equilibrium sales as well as vehicle prices. We do not control for vehicle prices in equation (2) because prices are likely to be correlated with unobserved demand or supply factors, and we lack suitable price instruments in the context in which technology and vehicle characteristics are endogenous (Klier and Linn, 2012). If the potential sales is a valid instrument, it is uncorrelated with unobserved supply or demand factors that affect vehicle price independently of sales, and omitting the vehicle's price would not cause spurious results. That is, omitting the price would affect only the interpretation of the log sales coefficient, as the effect of sales net of any potential sales or fuel cost-induced price changes. The econometric estimates represent the equilibrium efficiency change accounting for any coinciding price changes caused by the demand or fuel cost change (as a robustness check we report results including vehicle price as independent variable).

We note that because we do not observe a vehicle's efficiency, the dependent variable in equation (2) is generated from fixed effects in equation (1). The process generates prediction error for the dependent variable, which only affects the variance of our parameters in (2) but not the expected point estimates (Hausman, 2001). We address this concern by bootstrapping the standard errors. The imputation should not bias either the fuel cost or market size coefficients, even though fuel economy appears on the right hand side of equation (2). Supporting the latter claim is the fact that the market size coefficient is unaffected if we omit fuel costs from equation (2).

A final consideration 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, or on past sales in the presence of learning by doing. Because of the high persistence of sales across years, we use current sales in equation (2) as a proxy for lagged or expected sales. This introduces measurement error that could bias the estimated sales coefficient. In the empirical analysis we report estimates of equation (2) that use other measures of sales and assess the importance of potential measurement error.

Next, we discuss the identification and interpretation of a vehicle's fuel cost per mile. Fuel costs can affect a vehicle's efficiency via two channels. The first is the indirect effect of fuel costs via sales. An increase in the fuel costs of a vehicle relative to the fuel costs of other

vehicles would reduce the demand for that vehicle and its sales; that is, fuel costs can act as a demand shifter analogously to the IV. The second is the direct effect of fuel costs. Higher fuel costs may raise consumers’ willingness to pay for technology that raises fuel economy, inducing technology adoption.

The first channel is captured in the first-stage equation (4). We aim to isolate the second channel in equation (2). 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, and that fuel prices are exogenous to the vehicle market (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, per-mile fuel costs vary because of time-series variation in fuel prices and cross-model variation in fuel economy. We 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 in the Midwest. Households purchase more small cars relative to pickup trucks in the Northeast than in 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. Similar to the construction of potential sales we compute a sales-weighted model-specific fuel price using NHTS data on vehicle purchases and U.S. Energy Information Administration (EIA) data on fuel prices  $p_{dt}$  by census division,  $d$ :

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

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:

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

In the cross section,  $\tilde{C}_{mt;s}$  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  $\tilde{C}_{mt;s}$  in equations (2) and (4) as an independent variable would yield biased estimates because the variable would be correlated with the error term in those equations. To address this concern, we first construct  $\tilde{C}_{mt;s}$  using initial-year fuel economy  $e_{m0}$  to eliminate temporal variation that can be correlated with unobserved demand and supply factors. Second, similarly to the potential sales instrument, we subtract average fuel costs from the variable to obtain the independent variable  $\bar{C}_{mt;s}$ , which we

include in equations (2) and (4). Subtracting average fuel costs eliminates cross-sectional variation arising from fuel economy or the NHTS weights, which could be correlated with unobserved demand or supply shocks.

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

Because fuel costs shift vehicle demand, we interpret this coefficient as capturing the effect of consumers' willingness to pay for fuel economy. The coefficient on log sales is identified by variation in the IV as well as fuel costs. Therefore, fuel costs can indirectly affect efficiency via sales. Fuel costs can also directly affect efficiency by raising willingness to pay for fuel economy, conditional on equilibrium sales.

## 4 Estimation Results

### 4.1 Main Results

Table 4 shows the main estimation results. Column 1 reports the OLS estimates of equation (2) for comparison with the preferred IV estimates in column 2. 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 make-specific shocks, the regression includes year fixed effects, make fixed effects, and the interaction of a linear time trend with make fixed effects. The table reports the estimated coefficient on log sales with the bootstrapped standard error in parentheses, clustered by make to allow for arbitrary correlation of the error term within makes and over time, and for the fact that the dependent variable is estimated in equation (1). The estimated coefficient on log sales is 0.008, and the estimate is statistically significant at the 1 percent level. The coefficient on fuel costs is positive, but it is not statistically significant.

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

The magnitude of the IV estimate in Panel A, 0.021, 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 0.2 percent. This estimate is substantially larger than the OLS estimate in column 1. This magnitude amounts to 13 percent of the annual rate of efficiency improvement from 2000 to 2012 (Klier and Linn, 2016). To interpret the efficiency using consumer fuel savings, this amounts to \$50 in savings per vehicle (2010 USD), using estimates from Leard et al. (2019b). As an alternative interpretation, the sales effect causes manufacturers to adopt additional technology for high-selling vehicles, compared to a counterfactual that does not include the sales effect. The additional technology for high-selling vehicles yields annual consumer benefits of roughly \$100-400 million, depending on whether manufacturers use the technology to raise fuel economy or horsepower.

The point estimate represents the equilibrium efficiency change accounting for unobserved price changes caused by the demand or fuel cost change. Reverse causality would bias the OLS estimate away from zero, whereas omitted variables could bias the OLS estimate in either direction. The fact that the IV estimate is larger than the OLS estimate suggests that omitted variables bias is the dominant source of bias. For example, vehicle entry could create omitted variables bias. Consider a crossover vehicle that enters the market in the early-2000s. As additional vehicles enter in subsequent years, sales of the initial vehicle may decline because of the increased competition. At the same time, the competition could induce the manufacturer of the initial vehicle to invest in efficiency improvements to attract consumers with high valuation of efficiency. 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 positive and statistically significant. A one standard deviation increase in fuel costs raises efficiency by 0.4 percent. In contrast, the coefficient on fuel costs is negative in the first stage, suggesting that the indirect effect via sales is opposite to the direct effect. Because fuel costs are an independent variable in the first stage, a potential concern is that the functional form in equation (2) may not properly distinguish between the two channels by which fuel costs affect efficiency. In fact, the baseline does appear to distinguish the two channels because the sales coefficient is similar in column 3 when we omit fuel costs from the first and second stages.

## 4.2 Possible Sources of Omitted Variables Bias

Our IV strategy rests on the assumption that the IV affects efficiency only via sales. This assumption implies that changes in demographics over time are uncorrelated with preference or supply changes over time. If this assumption does not hold, changes in purchasing patterns across survey waves would be correlated with changes in demographics, and the IV estimates

would be sensitive to the choice of which NHTS waves to use in constructing the IV. Table 5 repeats the baseline in column 1 and shows the estimates for two alternative specifications: using only the 1995 or the 2009 NHTS. The point estimates on sales and fuel costs are barely changed, providing strong support for the validity of the IV strategy.

As noted above, a potential concern about the baseline IV estimates is that changes in a vehicle’s sales may be correlated with changes in attributes of other vehicles. The IV strategy is premised on the notion that manufacturers improve vehicle efficiency in responses to changes in predicted sales driven by changes in demographics, i.e. that they position their product in response to demand (Petrin, 2002). This argument also implies that manufacturers may differentiate their products from competing products by changing non-price attributes (e.g., Fan (2013) and Fischer (2010)). In that case, a change in demand for a particular vehicle could cause competing manufacturers to change prices or non-price attributes of their vehicles, affecting the sales of the first vehicle. This situation would imply that the sales coefficient is biased because it would include the effect of supply conditions on sales and efficiency.

We present four arguments suggesting that the instrumented sales is uncorrelated with such changes in supply conditions. First, the preference weights used to construct the instrument depend on the supply conditions during the year in which the NHTS was implemented. If the instrumented sales were correlated with supply-side changes, the preference weights would be correlated with (unobserved) supply conditions. In that case, we would obtain different sales coefficient estimates using the 1995 or 2009 NHTS waves to construct the instrument, rather than all three waves. Table 5 showed that the results are similar under these alternative variable constructions, suggesting that instrumented sales is uncorrelated with unobserved supply conditions.

Second, we add to the baseline a control for product competition in non-price attributes. Akerberg and Rysman (2005) suggest two measures to control for unobserved product differentiation: the log number of products in a nest ( $\ln J_{nest}$ ), and the weighted distance between the product attributes of vehicle  $m$  and other products in the nest ( $R_{m|nest}$ ). The first measure controls for the number of competing products, and the second measure accounts for

the similarity of the competing products.<sup>12</sup> For both measures we use market segments to define the nests, noting that our results are robust to various definitions of nests.

Table 6 presents our results with two [Akerberg and Rysman \(2005\)](#) measures as controls in column 2 and 3 (column 1 repeats the baseline). Adding these controls does not affect the estimated sales effect. In the first stage, the control variables are correlated with a vehicle’s sales, as one would expect (first stage results are available upon request). The IV results imply that, although competition in product space affects sales, the instrumented sales from demographic changes excludes effects of competition on a vehicle’s efficiency.

Third, we control for the degree of competition directly by using the efficiency of competing vehicles. Consumers have heterogeneous preferences for efficiency and other vehicle attributes. The efficiency of competing models could therefore have a positive or negative effect on a particular model’s efficiency (see Section 3.1). In column 4 of Table 6 we add to the baseline specification the mean efficiency of vehicles sold under other makes in the same market segment and model year. The efficiency variable may be endogenous because of reverse causality and perhaps other reasons, and we instrument for it using the mean potential sales of the corresponding vehicles. In column 4 the sales coefficient is similar to the baseline, providing further evidence that the instrumented sales is uncorrelated with supply conditions.

Fourth, a related possibility is that the efficiency of other models sold under the same make affects a model’s efficiency. This could occur because of demand effects, such as a make-level demand shock, rather than the supply-side effects that were considered earlier. In column 5, we add to the baseline the mean powertrain efficiencies of other models sold under the same make in the same market segment, and we use mean potential sales of the corresponding models as an instrument. Make-level efficiency has a small and negative effect on efficiency, but the estimate is noisy; the sales coefficient is unaffected.

So far, we have focused on product space competition. Table 7 shows that the results are robust to adding controls for other supply or demand shocks. The baseline includes controls for make-level demand or supply shocks, but there may also be segment-level shocks. Column 2 shows that the results are similar if we include segment by make fixed effects. Segment-level shocks could also vary over time. For example, the increase in market shares of crossovers in

---

<sup>12</sup>[Akerberg and Rysman \(2005\)](#) propose this method to address misidentification and overidentification of the price elasticity in logit demand estimation due to congestion of the product space. Although we are not estimating a logit demand system, the same argument for these two measures applies in our context. In their paper  $R_{m|nest} = \sum_{k=1}^K \phi((X_m - X_k) \times cov(X)^{-1}(X_m - X_k))$ , where  $\phi$  is normal and  $K$  is the number of observed dimension of non-price attributes. In practice, we follow [Houde and Spurlock \(2015\)](#) and construct

$R_{m|nest} = \sqrt{\sum_{k=1}^K \left( \frac{x_{mk} - x_{nest,k}}{sd(x_{nest,k})} \right)^2}$  to account for the relative location of vehicle  $m$  in the product space using a vehicle segment as a nest.



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, carlike light trucks. A correlation between preference changes and demographics would bias the IV estimate, but column 3 shows that the sales and fuel cost coefficients are similar if we add to the baseline the interactions of market segment fixed effects and a linear time trend.

Fuel economy standards 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 technology (Klier and Linn, 2016).

Klier and Linn (2016) show that the stringency of the fuel economy standards was uncorrelated with fuel costs, addressing the concern that fuel economy standards bias the coefficient on fuel costs. However, this result leaves open the possibility that the stringency of the standards and sales may be correlated with each other. The interactions of make fixed effects with a linear time trend control imperfectly for the standards because during the sample period the stringency of the standards varied nonlinearly over time and within makes.

Because the shadow costs of fuel economy standards are unobserved, we have tried several semi-parametric approaches to controlling for the standards. In particular, column 4 includes the interactions of make fixed effects with a quadratic time trend. This controls for the nonlinear changes in the stringency of the standards over time and across manufacturers. Historically, the standards applied separately for cars and light trucks, but since 2011 manufacturers have been allowed to average across their entire fleet (Leard and McConnell, 2016). To account for the differing regulatory pressure across cars and light trucks, column 5 includes triple interactions of make 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. Note that these specifications also address the potential bias caused by unobserved demand or supply shocks at the make, 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). We define stringency as the difference between the level of the standard that a manufacturer faces at the end of the sample, and the average fuel economy of the manufacturer's vehicles at the beginning of the sample. We allow the effect of stringency on efficiency to vary over time, reflecting the fact that the standards were first tightened for light trucks in 2005 and for cars in 2011. Column 6 includes these controls, and

the coefficient on log sales is 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 7 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.

In addition, using the baseline definition of demographic groups to construct the sales instrument could include measurement error because of zero predicted sales for cells with few observations in the NHTS. We aggregate the baseline demographic groups from 2,836 cells to 431 cells, which substantially reduces the share of cells with zero predicted sales. Appendix Table A.4 shows that our results are robust to the aggregation of demographic cells.

An alternative to using sales by model and demographic group from the NHTS is to use a discrete choice model to predict sales. We use the model and parameter estimates from Leard et al. (2019a), which allows preference parameters to vary across demographic groups. However, to incorporate their estimates we must aggregate the demographic groups, which reduces the variation of the predicted sales in the first stage. Nonetheless, the point estimates are similar using the alternative IV (Appendix Table A.2).

### 4.3 Alternative Measures of Efficiency and Sales

In this subsection we show that the results are similar using alternative measures of efficiency or sales. In the baseline specification (repeated in column 1 of Table 8 for convenience), we estimate efficiency by model and model year using equation (1), 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 sales could affect efficiency. To assess whether engine platform-level sales affect efficiency, columns 2 and 3 report estimates of equation (2) 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 (1) by engine platform and model year rather than by model and model year, and we use the estimated efficiency as the dependent variable in equation (2). 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

platform constitutes a new platform generation).<sup>13</sup> The log sales coefficient is similar to the baseline. Thus, we find similar elasticities of efficiency to sales across different levels of aggregation of the dependent variable.

As noted in Section 2.2, manufacturers typically make large changes to the powertrain or vehicle during major redesigns, and smaller changes between redesigns. The baseline 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 experiences a major redesign (model redesigns do not always coincide with engine redesigns because many models share engine platforms). In column 4 we estimate efficiency and sales by model generation and year (we collect model generation information from Automotive News, as in Blonigen et al. (2013)). The estimated coefficient on log sales is similar to the baseline.<sup>14</sup>

Next, we consider possible sources of measurement error (classical or nonclassical) in the sales variable. Because of regular production and redesign cycles in the vehicles market, efficiency may respond gradually to sales. 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. Analogously to the three-year moving average efficiency 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 average sales over multi year periods; the results are similar to the baseline (not reported).

Using the 2009 NHTS rather than the 1995 (or 2001) NHTS increases the average number of households per cell from 19 (or 27) to 38 (using the 2009 NHTS also increases the number of cells). We note that the specification using the 2009 NHTS addresses the possibility that the 1995 and 2001 NHTS waves have greater measurement error than the 2009 survey wave because of their smaller sample size.

Column 6 reports results using lagged sales, an exercise motivated by a few considerations. First, learning by doing suggests that lagged sales may be correlated with efficiency. Second, efficiency may lag behind sales because of the time required to redesign and test a vehicle

---

<sup>13</sup>Different models in the same year sold under the same make could share a platform, as could one model in different years. It is also possible for models sold under a different make to share a platform.

<sup>14</sup>Alternatively, we also tried to include sales of other models from a manufacturer, instrumented by the potential sales of other models sold by the same firm. Unfortunately, the two predicted variables from the first stage (sales of the model and sales of other models) are highly correlated with one another, and it is not possible to separately identify the effects. The high correlation arises from the fact that manufacturers tend to specialize to some extent. For example, Toyota tends to sell fairly inexpensive vehicles with high fuel economy, and predicted sales of one Toyota model are highly correlated with predicted sales of other models.

before beginning production. The fact that shifts in demographics and potential sales can be forecast to some extent mitigates the lag between demographics-driven changes in sales and adoption, but there could nonetheless be a lag. We can consider these possibilities empirically by replacing current log sales with the one-year lag of log sales and by replacing fuel costs with lagged fuel costs. Column 6 shows that the results are similar using the lags (the high correlation between lagged and current sales prevents us from using a distributed lag model that includes both variables in the same regression).

Finally, the baseline allows sales to affect efficiency at the model level. We showed that sales have a similar effect at the platform level as at the model level. Because of economies of scale or scope, aggregate sales, such as make by segment, may affect efficiency. Using make by segment level efficiency and sales, we find results similar to the baseline (column 7).<sup>15</sup>

#### 4.4 Additional Channels and Heterogeneous Technology Adoption

So far we have focused on the link between a vehicle’s sales and its efficiency. In this subsection we consider possible indirect effects on efficiency and possible heterogeneity across vehicles in the effect of sales on efficiency. We report these results in Table 9, repeating the baseline in column 1 for comparison. We do not find strong evidence of indirect effects or of heterogeneity, but this may reflect the limited variation in the variables we use to assess these possibilities.

First, 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 (2) an estimate of a manufacturer’s knowledge stock based on its historical patents. The variable is the cumulative number of efficiency-related 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 2 controls for knowledge stock between 1997 and

---

<sup>15</sup>Another source of measurement error for sales is that some vehicle models are produced on global platforms, and technology could respond to global sales 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 sales would be most relevant to the chosen engine and transmission technologies for the vehicles 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 for vehicles sold in the United States and other markets.

2010 and shows that knowledge stock has a positive effect on efficiency, but the estimate is not statistically significant.<sup>16</sup>

We have focused on the role of sales 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 efficiency. If the IV strategy is valid, such omitted factors would not yield biased or spurious estimates of the sales 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 (i.e., the manufacturer’s suggested retail price) as a proxy for consumers’ overall willingness to pay for the vehicle. Column 3 shows that adding the vehicle price does not affect the estimate of log sales, supporting the exogeneity of the IV to omitted demand and supply shocks. 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 supply shocks.

Finally, we consider the possibility of heterogeneous effects of sales. In section 4.3 we considered heterogeneous effects for new and continuing vehicle models and found that lag sales have a similar effect on efficiency as contemporaneous sales. In addition, we also consider possible heterogeneous effects across market segments or firms. In column 4, we interact sales with a dummy variable for light trucks, and we instrument for this variable with the interaction between the potential sales and the light truck dummy. The point estimate on sales is barely affected although it is not precisely estimated, and we cannot reject the hypothesis that sales affect technology adoption by the same amount for cars and light trucks. Note that the light truck estimate suggests a weaker effect of sales on efficiency for cars than light trucks.

In column 5, we interact sales with a dummy variable equal to one for US-based manufacturers. As in the light truck exercise we do not find strong evidence of heterogeneous effects, although the limited variation of the sales variable prevents strong conclusions regarding heterogeneity. In column 6, we allow the effect of market size to depend on whether market size is increasing or decreasing by interacting log sales with dummies that indicate if log sales increases. We also interact our instrument, predicted sales, with a dummy that indicates if the predicted sales increase. We do not find evidence that the effect of sales depends on whether sales are increasing or decreasing.<sup>17</sup>

---

<sup>16</sup>The patent stock variable ends at 2010. OECD Triadic Patent Family (TPF) data are available up to 2015. However, it is common practice to omit the last four to five years of TPF data because of reporting lags from the US Patent and Trademark Office.

<sup>17</sup>We have also allowed for heterogeneity across market segments or across other types of manufacturers, yielding similar conclusions.

## 5 Implications

### 5.1 Effects of Gasoline Prices on Efficiency

In Section 4 we quantified the economic significance of the sales and fuel cost coefficients by comparing the effect of a one standard deviation change in the independent variables. To further illustrate the economic importance of these estimates, we compare efficiency levels across scenarios of low and high gasoline prices. Note that consumers benefit from higher efficiency because manufacturers can use an efficiency increase to raise fuel economy, horsepower, or both according to equation (1). Because the manufacturer choices of fuel economy and horsepower lie outside the scope of the empirical analysis, we focus on efficiency in the simulations in this section.

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 1.1 mpg. Therefore, the gasoline price increase affected both fuel costs and sales.

We isolate these two channels, first considering the sales channel. We assign each model in the data to one of five fuel economy quintiles, based on each model's fuel economy in the first year it appears in the data. The first quintile consists of vehicles with the lowest fuel economy, and the fifth quintile consists of vehicles with the highest fuel economy. To maintain consistency with equation (2), which includes year fixed effects that control for the market-wide average efficiency increase, we hold fixed the market-wide average efficiency change across the scenarios. In the simulations, fuel prices affect the cross-sectional distribution of sales, which in turn affects the cross-sectional distribution of efficiency. We use equations (2) 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 consists of the lowest-fuel-economy vehicles. To isolate the sales effect, we adjust fuel prices for the first stage equation (4) but not the second stage equation (2).

In Panel A of Figure 6 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 (2). The clear bars show the average predicted efficiency increase assuming fuel prices had remained at 2003 levels. 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 by 0.51 percent more for the lowest-fuel-economy quintile and by 0.47 percent less for the highest-fuel-economy quintile. These effects are

consistent with expectation and they are large relative to the predicted cumulative 6.4 percent efficiency of improvement that actually occurred between 2003 and 2007. Thus, via the sales effect the increase in gasoline prices caused manufacturers to improve the efficiency of vehicles with high fuel economy.

In Panel B of Figure 6 we isolate the direct effect of fuel costs on efficiency. For this counterfactual we hold sales fixed. We use the fuel cost coefficient in the second stage and gasoline price change to predict the cumulative efficiency change for each model.

A gasoline price increase raises fuel costs for all vehicles, but by more for vehicles with low fuel economy. Therefore, efficiency should increase more for vehicles with low fuel economy because of the fuel cost effect, and we expect the fuel cost effect to work in the opposite direction as the sales effect. Panel B shows this to be the case. Comparing the two simulations suggests that the fuel cost effect is larger in magnitude than the sales effect. Therefore, accounting for both effects implies that an increase in gasoline prices causes greater efficiency improvements for vehicles with low fuel economy. The results demonstrate the importance of distinguishing between the sales and fuel cost effects.

## 5.2 Effects of Crossover and SUV Sales 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 2001 and 2004, the average sales per crossover model 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 sales increased efficiency for crossovers and decreased efficiency for SUVs, relative to a counterfactual in which sales had remained at 2001 levels.

To quantify these effects, we estimate the cumulative efficiency changes between 2001 and 2004 that would have occurred if crossover sales had not changed over this period.<sup>18</sup> On the left side of Figure 7 we compare the predicted and counterfactual efficiency of crossovers. The colored bar shows the predicted cumulative efficiency improvement over 2001–2004. The clear bar shows the counterfactual cumulative efficiency change holding sales of crossovers fixed at 2001 levels. In the counterfactual scenario, the lower sales of crossovers causes efficiency to be 0.09 percent lower.

As shown in the right panel of Figure 7, the counterfactual scenario causes efficiency of SUVs to be 0.17 percent higher than is predicted using the actual sales in 2004. This is

---

<sup>18</sup>Because we are interested in the effect of sales on crossover and SUV efficiency, our counterfactual represents a partial equilibrium outcome in which the sales 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.

substantial compared to the 2.9 percent cumulative change between 2001 and 2004 for SUVs. In short, comparing the simulations for fuel prices, demographics, crossovers, and SUVs, we observe that sales and fuel prices have had economically significant effects on efficiency.

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

Next, we discuss the policy implications of the estimated sales and fuel cost coefficients in equation (2). The typical sales of gasoline-powered vehicles are an order of magnitude larger than the typical sales of alternative-fuel vehicles. The positive sales effect implies that manufacturers will continue to improve the efficiency of gasoline-powered vehicles, which currently dominate the market. Unless the sales effect of alternative-fuel vehicles is multiple orders of magnitude larger than that for gasoline-powered vehicles, the sales effect increases the challenge of alternative-fuel vehicles to compete with gasoline-powered vehicles, and reduces the effectiveness of policies that directly subsidize alternative-fuel vehicles, relative to a hypothetical in which there were no sales effect.

The second policy implication concerns the effects of fuel price and vehicle-based policies aiming to reduce passenger vehicle fuel consumption and greenhouse gas emissions. We focus on the immediate incentives of the policies on technology choices; characterizing the long-run effects would require an equilibrium model and lies outside the scope of the paper. Our discussion of fuel price-based policies includes fuel taxes or a carbon tax imposed on fuels, and the discussion of vehicle-based policies includes fuel economy standards, greenhouse gas emissions rate standards, and feebates, as all of these policies can affect vehicle prices and sales.

If the standard applies to the mean fuel economy of a manufacturer's vehicles, as with the US standards and those in other regions, manufacturers can 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 sales shift raises the manufacturer's sales-weighted average fuel economy, helping the manufacturer achieve the standard. The sales shift also creates incentives for technology adoption, which arises from the changes in vehicle prices. A vehicle tax works in similar ways as to fuel economy standards (e.g. Grigolon, Reynaert, and Verboven, (forthcoming); Huse and Koptyug (2017)). It would therefore have similar effects on technology adoption via the sales effect.

A feebate would also create a sales effect for technology adoption. A feebate refers to a combination of taxes and rebates: taxes on vehicles with low fuel economy and rebates on



vehicles with high fuel economy. Note that prior to the efficiency response to the feebate, the policy does not affect fuel costs. The taxes and rebates mimic the pricing behavior of manufacturers facing fuel economy standards and affect sales in a similar manner as does a fuel economy standard. In fact, a feebate or standard can be designed to have identical effects on the sales of each new vehicle, absent additional policies that interact with these policies (Roth, 2015). Consequently, the policies create the same sales effect on technology adoption.

Fuel price-based policies differ from vehicle-based policies in their effects on technology adoption. The literature has focused on differences between these two classes of policies that arise because of differences in driving incentives (i.e., the rebound effect) and the fact that standards or an equivalent feebate cover only new and not existing vehicles, which introduces the inefficiencies associated with vintage differentiated regulation (Stavins (2006) and Jacobsen and van Benthem (2015)). In this paper we identify a further distinction between the two classes of policies. We compare the immediate effects of these policies on technology adoption holding each vehicle's fuel economy fixed. In this setting, all policies induce the sales effect. Fuel taxes or a carbon tax differ from fuel economy standards or a feebate in that they immediately affect fuel costs.<sup>19</sup>

We conduct two simulations to illustrate this difference. In the following analysis we focus on a feebate and fuel tax as representative of vehicle- and fuel-based policies. Rather than consider the effects of feebates and fuel taxes on efficiency improvements for different fuel economy groups, as we did before, we further illustrate the implications of the sales and fuel cost effects 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 because of the sales effect. This is because the sales of vehicles with high fuel economy would increase, raising their efficiency, and the sales 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 due to the sales effect. The widening of the distribution would be greater for vehicle-based than

---

<sup>19</sup>In the long run these policies also induce changes in technology and fuel economy; we focus on the short-run fuel cost and sales effects to isolate the differences in the policies that have not been considered elsewhere in the literature.

fuel-based policies because the latter also include the fuel cost effect, which works in the opposite direction as the sales effect, as explained in section 5.1.

In the first simulation we implement the feebate as a price on fuel economy. The feebate is implemented as a fee or rebate on vehicle purchase, which depends on the vehicle's fuel economy. The feebate is defined by a pivot, which is the fuel economy level above which vehicles receive a subsidy and below which vehicles are taxed; the subsidies and taxes are added to the purchase price of the vehicles. 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 1.1 mpg, which is the same fuel economy change as that which occurred in the counterfactual scenario considered in Figure 6. We calibrate the feebate so that a model with fuel economy  $e_{jt}$  has a feebate of  $(1/e_{jt} - 1/e_t) \times 1.1$  dollars per mile. The counterfactual scenario includes a feebate for the years 2010–2013, and market conditions (e.g., fuel prices) are otherwise unchanged. As the feebate affects technology adoption and fuel economy over time, the feebate would also create a direct fuel cost effect. Here, we consider only the immediate incentive for technology adoption caused by the indirect effect via sales, and we do not trace out the choice of fuel economy over time.<sup>20</sup> 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. The feebate is implemented over four years for comparability with the gasoline price scenario in section 5.1. Because of the feebate's effect on sales, we expect the feebate to increase the variance of efficiency across vehicles in the market.

Panel A of Figure 8 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 sales of vehicles with fuel economy below the pivot, the counterfactual efficiency lies below the predicted efficiency for models with fuel economy below the pivot. For these vehicles the cumulative efficiency from 2010 to 2013 would have been smaller had the feebate been in place. In contrast, the feebate increases the sales of vehicles with fuel economy above the pivot and causes counterfactual efficiency to lie above predicted efficiency for such vehicles. The lines in Panel A represent the fitted values of a linear regression of cumulative 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

---

<sup>20</sup>To be consistent with the counterfactual in Figure 6 and with the fuel tax counterfactual in Panel B of 8, we conduct the simulations using a per-mile tax (or subsidy) equivalent to the lump-sum fees (or rebates).

economy. Figure A.2 provides an alternative view of the effect of the policy by showing this effect by illustrating the widening of the efficiency distribution caused by the feebate.

To compare with the feebate we consider a fuel tax of \$1.14 per gallon, which increases the sales-weighted average fuel economy by the same 1.1 miles per gallon. The fuel tax affects both fuel costs and sales. Panel B of Figure 8 shows that because these two effects oppose one another, the fuel tax does not strengthen the positive relationship between efficiency and fuel economy nearly as much as does the feebate.

These simulations illustrate the differing effects of vehicle and fuel-based policies on technology adoption, and have two implications for the welfare costs of the policies. First, these differences affect consumer choices among vehicles, and therefore affect the welfare costs of achieving a particular emissions reduction. Although estimating the welfare consequences 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.

The second implication pertains to the competitiveness of alternative-fuel vehicles. Survey data from recent vehicle buyers suggests that individuals who obtain alternative-fuel vehicles typically choose between those vehicles and gasoline-powered vehicles that have relatively high fuel economy. By widening the efficiency distribution, the vehicle-based policies would therefore increase the challenges of alternative-fuel vehicles in competing with gasoline-powered vehicles, and by more than the fuel-based policies.

## 6 Conclusion

Fuel economy standards across the world, including the U.S., are set to increase.<sup>21</sup> Technology adoption features importantly in meeting tighter standards, yet there is little empirical evidence on which factors determine a manufacturer's choice of efficiency. This paper analyzes the effects of a vehicle's sales and fuel costs on the model-level adoption of efficiency-improving technology in the US new passenger vehicle market. We show that sales and fuel costs have substantial effects on technology adoption and discuss implications for fuel consumption policies.

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 sales by instrumenting for sales using potential sales as a demand shifter. Variation in potential sales arises from changes in demographics over time and cross-sectional heterogeneity in purchasing patterns across demographic groups. In the preferred specification a 10 percent increase in sales (corresponding to about one standard deviation)

---

<sup>21</sup>The current efforts to weaken the federal fuel economy and greenhouse gas standards is likely to slow down, but not reverse this trend in the U.S.

increases efficiency by 0.2 percent, compared with a mean annual efficiency improvement of about 1.4 percent in the sample. Fuel costs affect efficiency via sales and also independently of sales. [Acemoglu et al. \(2016\)](#) find that high fuel prices increase patents on alternative-fuel technologies, and our findings suggest that efficiency improvements may respond directly to fuel costs or indirectly via sales.

In several ways, we quantify the economic importance of the sales and fuel cost effects on efficiency. First, historical variation in fuel prices has had a substantial effect on efficiency. Real fuel prices nearly doubled between 2003 and 2007, which affected the relative sales of vehicles according to their fuel economy. Had gasoline prices remained at 2003 levels, the efficiency of vehicles in the lowest fuel economy group would have increased 0.5 percent more between 2003 and 2007 than it did. Efficiency of the vehicles in the highest fuel economy group would have increased by 0.4 percent less than it did. The direct effect of fuel costs on technology adoption worked in the opposite direction, more than offsetting the sales effect.

Second, shifts in sales of crossovers and SUVs have caused large changes in the efficiency of these vehicles. These two results imply that sales and fuel costs have had economically significant effects on technology adoption.

Finally, we discuss two policy implications of the empirical findings. First, the sales effect implies that manufacturers in the US market will continue to improve the efficiency of their best-selling gasoline-powered vehicles. This pattern of technology adoption decreases the competitiveness of low-selling gasoline-powered vehicles and alternative-fuel vehicles. We suggest that future welfare analysis of passenger vehicle fuel economy and greenhouse gas standards account for this sales effect.

Second, we show that the policies commonly discussed for reducing vehicle fuel consumption and greenhouse gas emissions have differing effects on the efficiency distribution of vehicles in the market. Because of the sales effect, feebates, a fuel economy standard, or an emissions rate standard would widen the distribution, further harming the competitiveness of conventional hybrid, plug-in hybrid vehicles and other alternative-fuel vehicles assuming the market size effect is quantitatively similar in the new market segment. Fuel or carbon taxes do not have this effect because they induce a fuel cost effect that opposes the sales effect. The policies differentially affect technology adoption and consumer choice, changing the welfare costs of achieving a particular emissions reduction. Moreover, the vehicle-based policies have a greater effect on the efficiency of vehicles with high fuel economy, against which alternative-fuel vehicles typically compete.

The simulations in this paper hold fixed the set of vehicles in the market, but entry and exit of new vehicles may be an additional channel by which sales and fuel costs affect technology adoption. Future work could endogenize entry and exit decisions.

The empirical analysis does not identify the underlying reasons why sales affects efficiency. As we discuss in Section 3.1, a range of factors could explain a positive effect of sales on technology adoption. Future work may distinguish among these possibilities, which would have 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., 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.
- Ackerberg, D. and M. Rysman (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics* 36(4), 771–788.
- 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.
- Blonigen, B. A., C. R. Knittel, and A. Soderbery (2013, April). Keeping it fresh: Strategic product redesigns and welfare. (NBER Working Paper No.18997).
- 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.
- Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review* 103(5), 1598–1628.
- 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.
- Grigolon, L., M. Reynaert, and F. Verboven. Consumer valuation of fuel costs and tax policy: Evidence from the european car market. *American Economic Journal: Economic Policy*, forthcoming.
- Hausman, J. (2001). Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *The Journal of Economic Perspectives* 15(4), 57–67.
- Houde, S. and A. Spurlock (2015). Do energy efficiency standards improve quality? Evidence from a revealed preference approach. *Working Paper*.
- Huse, C. and N. Koptuyg (2017). Taxes vs. standards as policy instruments: Evidence from the auto market. (Working Paper).

- 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.
- Jacobsen, M. R. and A. A. van Benthem (2015, March). Vehicle scrappage and gasoline policy. *American Economic Review* 105(3), 1312–38.
- 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., J. Linn, and K. Springel (2019a). Pass-through and welfare effects of regulations that affect product attributes. (RFF Working Paper No. 19-07).
- Leard, B., J. Linn, and Y. C. Zhou (2019b). How much do consumers value fuel economy and performance? Evidence from technology adoption. Technical Report RFF Policy Report, Resource for the Future.
- Leard, B. and V. McConnell (2016). New markets for pollution and energy efficiency. (RFF Discussion Paper No. 15-16).
- Levitt, S. D., J. A. List, and C. Syverson (2013). Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy* 121(4), 643–681.
- 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*.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* (110), 705–729.
- Reynaert, M. (2015). Abatement strategies and the cost of environmental regulations: Emission standards on the European car market. *Working Paper*.
- Roth, K. (2015). The unintended consequences of uncoordinated regulation: Evidence from the transportation sector. *Working Paper*.
- Stavins, R. N. (2006). Vintage-differentiated environmental regulation. *Stan. Envtl. LJ* 25, 29.
- Zhou, Y. C. (2016). Knowledge capital, technology adoption, and environmental policies: Evidences from the US automobile industry. *Job Market Paper*.

# Figures

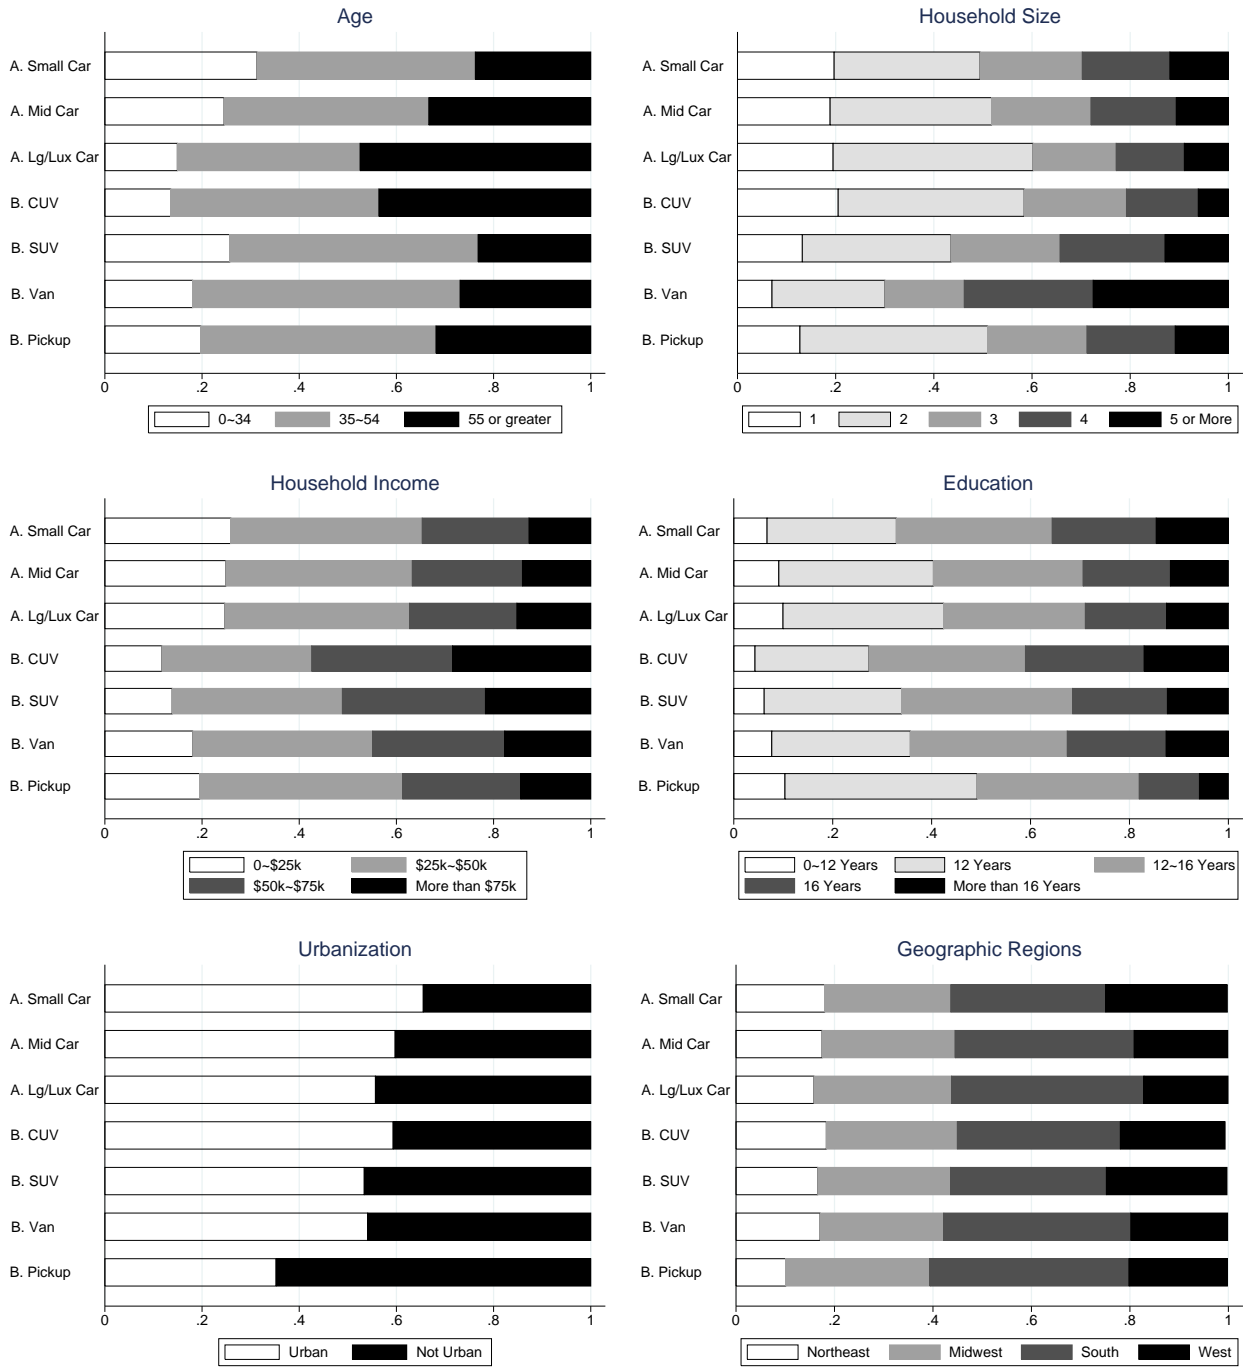
Figure 1: Total Vehicle Sales by Segment, 1997–2013



Notes: For each market segment, the figure plots the total model year sales.

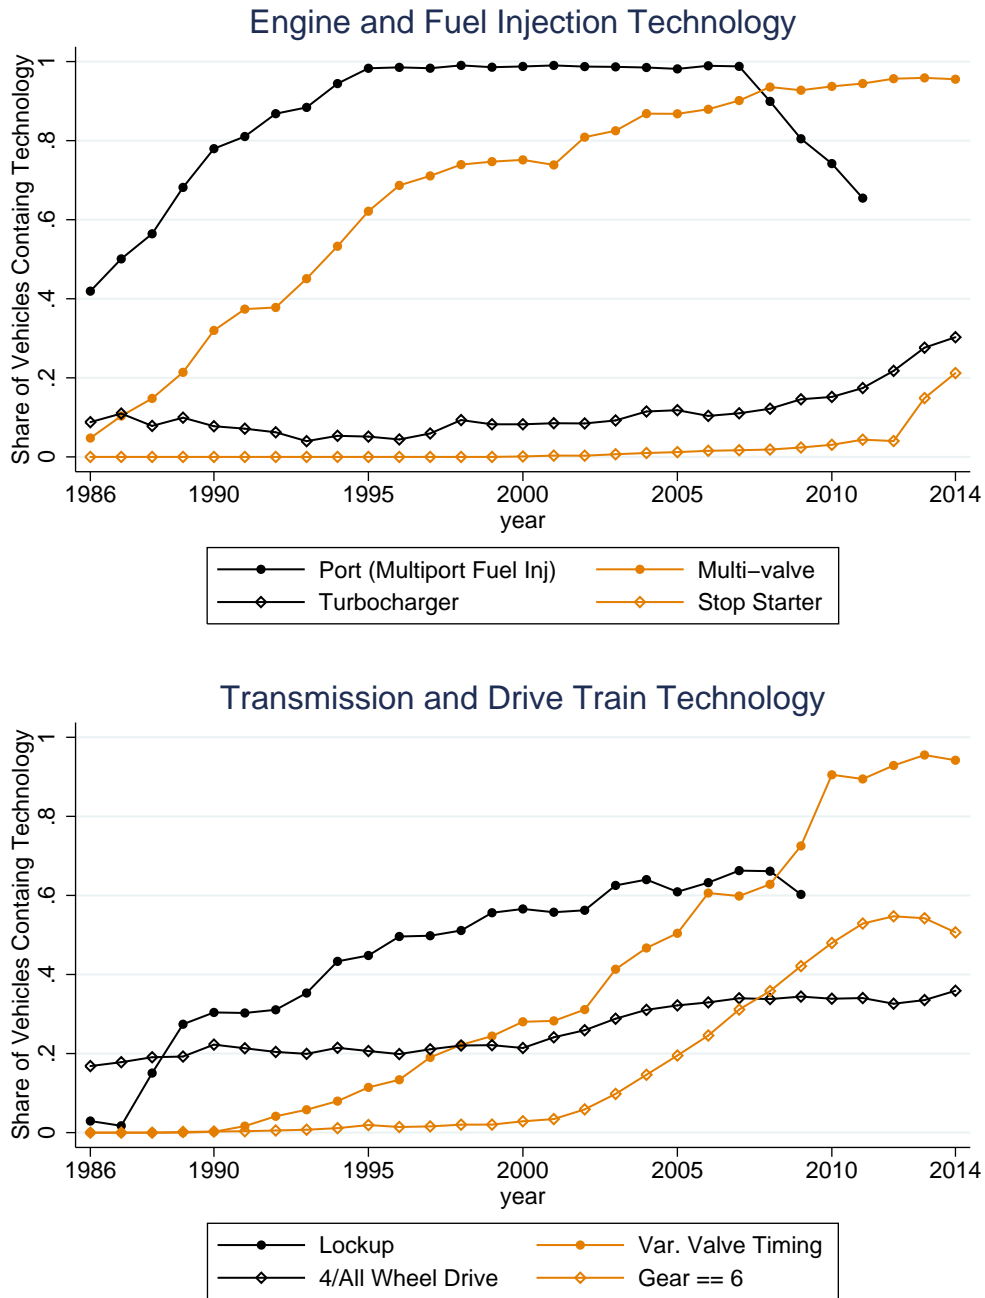


Figure 2: Vehicle Purchase Patterns by Demographic Group




*Notes:* The figure is constructed using the NHTS 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: 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 4: 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 5: Estimated Powertrain Efficiency, 1997–2013




Notes: The figure plots the mean estimated efficiency across passenger cars and light trucks estimated from equation (1). 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 (2), 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 (2) and (4) to predict the efficiency change for each observation between 2003 and 2007. Panel A simulates the isolated sales effect and Panel B simulates the isolated fuel cost effect.

Figure 7: **Effect of Sales on Efficiency, 2001–2004**



*Notes:* The colored bars show the mean cumulative predicted efficiency increase between 2001 and 2004 for crossovers (left panel) and SUVs (right panel). Predicted values are obtained from the estimation of equation (2) 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 2001 levels and uses equation (2) to predict the counterfactual efficiency change for each crossover and SUV between 2001 and 2004.

Figure 8: Correlation of Fuel Economy and Efficiency



*Notes:* For each observation in equation (2), the frontier is predicted using the estimates reported in column 2 of Table 4. Panel A is a scatter plot of efficiency and fuel economy for each model in 2013. 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. The counterfactual efficiency of each vehicle is computed from the sales caused by introducing a feebate of  $(1/e_{jt} - 1/e_t) \times 1.1$  \$/gal, where  $e_{jt}$  is the fuel economy of model  $j$  in model year  $t$  and  $e_t$  is the harmonic mean of fuel economy in model year  $t$ . In Panel B, the counterfactual efficiency of each vehicle is computed from both its direct effect and its effect from the sales, caused by introducing a fuel tax at 1.14 \$/gal.

# Tables

Table 1: Average Vehicle Characteristics, 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	24.8	201	317	3746	6.2
2005	24.7	233	345	4035	6.3
2010	26.0	263	369	4223	6.2
2013	28.3	278	381	4234	6.1

*Notes:* 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*** (0.037)	-0.424*** (0.034)
Diesel	0.336*** (0.017)	0.260*** (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 (1), 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**

Time period	High-Selling Vehicles			Low-Selling Vehicles		
	Efficiency in		Cumulative	Efficiency in		Cumulative
	starting year	ending year	change by period	starting year	ending year	change by period
1997–2000	0	0.017	0.017	0	0.012	0.012
2001–2005	0.019	0.070	0.051	0.006	0.057	0.051
2006–2009	0.068	0.140	0.072	0.078	0.135	0.057
2010–2013	0.156	0.239	0.083	0.161	0.234	0.074

*Notes:* Efficiency is estimated by model, market segment, and model year in equation (1), 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 Sales and Fuel Costs on Efficiency**

	(1)	(2)	(3)
Estimated by	OLS	IV, Baseline	IV
<b>Panel A.</b> Dependent variable: Efficiency			
Log sales	0.008*** (0.001)	0.021*** (0.004)	0.025*** (0.006)
Fuel costs	0.112 (0.101)	0.256*** (0.091)	
<b>Panel B.</b> First-stage estimate. Dependent variable: Log sales			
Potential sales (log)		0.138*** (0.029)	0.121*** (0.030)
Fuel costs		-20.543*** (1.776)	
If potential sales is imputed		-0.651*** (0.055)	-0.564*** (0.053)
Make fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Make fixed effects×linear time trend	Yes	Yes	Yes
Observations	2,740	2,740	2,740
RMSE	0.06	0.06	0.06
F (1st stage excl. var.)	NA	83.69	46.69

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

*Notes:* The table reports coefficient estimates from equation (2), with bootstrapped standard errors in parentheses, clustered by make (i.e., brand). Observations are by model and model year. Column 1 is estimated by ordinary least squares (OLS). Columns 2 and 3 are estimated by instrumental variables, using potential sales 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 make fixed effects, year fixed effects, and make fixed effects interacted with a linear time trend.



Table 5: **Alternative NHTS Waves**

Estimated by	(1) Baseline, All NHTS	(2) 1995 NHTS	(3) 2009 NHTS
<b>Panel A.</b> Dependent variable: Efficiency			
Log sales	0.021*** (0.004)	0.020*** (0.004)	0.021*** (0.004)
Fuel costs	0.256*** (0.091)	0.269*** (0.086)	0.264** (0.127)
<b>Panel B.</b> First-stage estimate. Dependent variable: Log sales			
Potential sales (log)	0.138*** (0.029)	0.106*** (0.036)	0.072** (0.028)
Fuel costs	-20.543*** (1.776)	-18.445*** (1.787)	-18.422*** (1.716)
If potential sales is imputed	-0.651*** (0.055)	-0.749*** (0.057)	-0.552*** (0.054)
Make fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Make fixed effects×linear time trend	Yes	Yes	Yes
Observations	2,740	2,791	2,736
RMSE	0.06	0.06	0.06
F (1st stage excl. var.)	83.69	84.27	61.77

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

*Notes:* The table reports coefficient estimates from equation (2), with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4, in which the potential sales is constructed using the 1995, 2001 and 2009 NHTS waves. In column 2, we construct the predicted sales using only the 1995 NHTS wave. In Column 3, we construct the potential sales using only the 2009 NHTS wave. The number of observations differs across the three columns because different models were purchased in each of the three NHTS waves.

Table 6: **Imperfect Competition Channels**

Dependent variable: Efficiency	(1)	(2)	(3)	(4)	(5)
	Baseline				
Log sales	0.021*** (0.004)	0.021*** (0.006)	0.021*** (0.005)	0.035** (0.014)	0.019** (0.007)
Fuel costs	0.256*** (0.091)	0.264** (0.124)	0.269*** (0.092)	-0.105 (0.317)	0.204 (0.150)
AS (2005) $\ln J_{nest}$ number of product		-0.002 (0.004)	-0.004 (0.004)		
AS (2005) $R_{j nest}$ attribute distance		0.001 (0.002)	0.000 (0.003)		
Nest is defined by		segment model year	segment		
Efficiency of competing models				-0.105 (0.317)	-0.829 (1.719)
Competing models are from				other makes, same segment and year	same make, segment, and year
Make fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Make fixed effect $\times$ linear time trend	Yes	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740	2,740
RMSE	0.06	0.06	0.06	0.07	0.06
F (1st stg. excl. var.)	83.7	46.9	50.2	55.8	55.5

\*  $p < 0.10$  \*\*  $p < 0.05$  \*\*\*  $p < 0.01$ .

*Notes:* The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4. Column 2 include two measures to account for unobserved competition and congestion based on [Akerberg and Rysman \(2005\)](#): the number of products within a nest, and the distance of product  $j$  to other product in a nest. In column 2, a nest is defined by a segment and model year. In column 3, a nest is defined by a segment. Column 4 includes the average efficiency by technology group as an independent variable. This variable is instrumented using the corresponding average potential sales of those models. Column 5 includes the average efficiency of other models sold under the same make in the same market segment, using the average potential sales as an instrument. All regressions are estimated by instrumental variables using potential sales as an instrument, as in Table 4.

Table 7: Controlling for Other Potential Sources of Omitted Variable Bias

Dependent variable: Efficiency	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Baseline						
Log sales	0.021*** (0.004)	0.021*** (0.004)	0.020*** (0.005)	0.019*** (0.004)	0.033*** (0.005)	0.021*** (0.004)	0.023*** (0.005)
Fuel costs	0.256*** (0.091)	0.314*** (0.100)	0.233** (0.105)	0.136 (0.103)	0.006 (0.121)	0.251*** (0.090)	0.479*** (0.127)
CAFE Stringency						0.027 (0.036)	0.074** (0.031)
Fuel costs×CAFE stringency							-2.878*** (0.796)
Make fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Make fixed effects×linear time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Segment fixed effects		Yes	Yes				
Segment fixed effects×make fixed effects		Yes					
Segment fixed effects×linear time trend			Yes				
Make fixed effects×quadratic time trend				Yes			
Make fixed effects×linear time trend×truck class					Yes		
Observations	2,740	2,740	2,740	2,740	2,740	2,740	2,740
RMSE	0.06	0.05	0.06	0.06	0.06	0.06	0.06
F (1st stage excl. var.)	83.69	113.14	82.16	87.79	62.86	82.50	74.58

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

*Notes:* The table reports coefficient estimates from equation (2), with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4. Column 2 includes a set of segment fixed effects and their interaction with make fixed effects. Column 3 includes a set of segment fixed effects interacted with a linear time trend. Column 4 includes make fixed effects interacted with a quadratic time trend. Column 5 includes the triple interaction of make fixed effects by light truck class by linear time trend. Column 6 includes the fuel economy stringency variable described in [Klier and Linn \(2016\)](#), and column 7 includes the interaction of this variable with fuel costs.

Table 8: **Alternative Methods for Estimating Efficiency and Sales**

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Efficiency							
Efficiency estimated by:	Model by year (baseline)	Platform by year	Model by platform generation	Model by model generation	Model by year (3-yr moving average)	Lagged log sales	Make by segment by year
Log sales	0.021*** (0.004)	0.023*** (0.005)	0.024** (0.009)	0.021*** (0.008)	0.025*** (0.005)	0.026*** (0.005)	0.019** (0.009)
Fuel costs	0.256*** (0.091)	0.567*** (0.166)	0.468** (0.221)	0.579** (0.244)	0.525*** (0.129)	0.352*** (0.123)	0.021 (0.100)
Make fixed effects	Yes		Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Make fixed effects ×linear time trend	Yes		Yes	Yes	Yes	Yes	Yes
Company fixed effects		Yes					
Company fixed effects ×linear time trend		Yes					
Observations	2,740	1,956	532	538	2,096	2,396	2,600
RMSE	0.06	0.06	0.06	0.06	0.06	0.06	0.05
F (1st stage excl. var.)	83.69	65.90	16.49	20.73	50.10	65.97	22.57

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

*Notes:* The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4. 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 sales rather than contemporaneous sales, as well as lagged fuel cost, potential sales, and impute dummy. In column 7 efficiency is estimated by make, segment, and model year. 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 sales as an instrument, as in Table 4.

Table 9: **Additional Channels and Heterogeneity**


Dependent variable: Efficiency	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline			Truck	US firm	Nonlinear
Log sales	0.021*** (0.004)	0.028** (0.011)	0.021*** (0.005)	0.030*** (0.007)	0.022** (0.009)	
Log sales $\times$ {1 = sales increases}						0.023*** (0.007)
Log sales $\times$ {1 = sales declines}						0.028* (0.017)
Fuel costs	0.256*** (0.091)	-0.237 (0.150)	0.241** (0.102)	0.178 (0.138)	0.241* (0.128)	0.210 (0.146)
Knowledge stock		0.003 (0.009)				
Log price			0.032*** (0.009)			
Log sales $\times$ truck				-0.038 (0.024)		
Log sales $\times$ US firm					-0.005 (0.028)	
Make fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Make fixed effect $\times$ linear time trend	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,740	2,318	2,740	2,740	2,740	2,740
RMSE	0.06	0.06	0.06	0.07	0.06	0.07

\*  $p < 0.10$  \*\*  $p < 0.05$  \*\*\*  $p < 0.01$ .

*Notes:* The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4. Column 2 includes the manufacturer’s knowledge stock, which is the cumulative number of efficiency-related patents that a parent company has applied for. Column 4 includes the log of the vehicle’s price as an independent variable. Columns 1 and 3-5 include observations from 1997 to 2013 and column 2 includes observations from 1997 to 2010. Column 4 includes a dummy for light truck, the interaction of sales with a dummy variable for light trucks, and the corresponding instrument. Column 5 includes a dummy for US-based manufacturers, the interaction of sales with a dummy for US-based manufacturers, and the corresponding instrument. All regressions are estimated by instrumental variables using potential sales as an instrument, as in Table 4. In column 6, we interact log sales with dummies that indicate if sales of a model increases or drops from the previous year. We interact our predicted sales IV with dummies that indicate if the predicted sales of a model increases or drops from the previous year.


# Appendix

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



*Notes:* For each observation in equation (2), 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 (2) and (4) to predict the efficiency change for each observation between 2003 and 2007.

Figure A.2: **Effect of Feebate on Efficiency**



*Notes:* For each observation in equation (2), the frontier is predicted using the estimates reported in column 2 of Table 4. The counterfactual efficiency of each vehicle is computed from the sales 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 model year  $t$  and  $e_t$  is the harmonic mean of fuel economy in model year  $t$ . The above figure shows the estimated density functions of cumulative predicted and counterfactual efficiencies over the period 2010 through 2013.

Table A.1: Robustness Using Alternative Group Definition

Estimated by	(1) Baseline	(2) Alternative Groups
<b>Panel A.</b> Dependent variable: Efficiency		
Log sales	0.021*** (0.004)	0.020*** (0.004)
Fuel costs	0.256*** (0.091)	0.253** (0.113)
<b>Panel B.</b> First-stage estimate. Dependent variable: Log sales		
Potential sales (log)	0.138*** (0.029)	0.114*** (0.028)
Fuel costs	-20.543*** (1.776)	-20.997*** (1.828)
If potential sales is imputed	-0.651*** (0.055)	-0.880*** (0.057)
Make fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Make fixed effects×linear time trend	Yes	Yes
Observations	2,740	2,740
F (1st stage excl. var.)	83.69	94.15

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

*Notes:* The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by make. Column 1 repeats the baseline from Table 4. Column 2 use alternative group definition in Appendix Table A.4. The alternative group definition reduce the number of cells from 2,628 to 431 and increase the number of surveyed households per cell to 178.



Table A.2: Robustness Using Alternative Instrumental Variable

Estimated by	(1) Baseline	(2)	(3) Alternative IV	(4) Alternative IV
<b>Panel A.</b> Dependent variable: Efficiency				
Log sales	0.021*** (0.004)	0.025*** (0.006)	0.024 (0.075)	0.018 (0.031)
Fuel costs	0.256*** (0.091)		0.807 (1.811)	
<b>Panel B.</b> First-stage estimate. Dependent variable: Log sales				
Potential sales (log)	0.138*** (0.029)	0.121*** (0.030)	0.029* (0.015)	0.064*** (0.017)
Fuel costs	-20.543*** (1.776)		-28.054*** (2.495)	
If potential sales is imputed	-0.651*** (0.055)	-0.564*** (0.053)		
Make fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Make fixed effects×linear time trend	Yes	Yes	Yes	Yes
Observations	2,740	2,740	1,761	1,761
F (1st stage excl. var.)	83.69	46.69	4.5	7.2

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

*Notes:* The table reports coefficient estimates with bootstrapped standard errors in parentheses, clustered by make. Column 1 and 2 repeat the baselines from Table 4. Column 3 and 4 use alternative instrumental variable using p parameters of price, cost per mile, horsepower-to-weight, and footprint from Leard et al. (2019a) for 3 income groups (that overlap with 5 income groups in Leard et al. (2019a)), 2 urban groups, and 2 age groups. We construct alternative log potential sales by predicting log sales from their demand model.

Table A.3: **Definitions of Demographic Groups**  
Panel A: Group Definition

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 urban	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						2,628

Panel B: Group Information

	All NHTS waves	NHTS 1995	NHTS 2001	NHTS 2009
Total number of groups	2,628	2,196	2,153	2,473
Number of households per cell	28	19	27	38

Table A.4: **Definitions of Alternative Demographic Groups**  
 Panel A: Group Definition

Group number	Age (years)	Household income (thousand nominal dollars)	Education (years)	Household Size	Urban	Census division
1	0–34	0–25	0–12	1 or 2	urban	New England Middle Atlantic
2	35–54	25–60	12+	3 or 4	not urban	Pacific
3	55+	60+		5+		East North Central West North Central Mountain
4						East South Central West South Central South Atlantic
No. of Groups	3	3	2	3	2	4
Total number of groups						431

Panel B: Group Information

	All NHTS waves	NHTS 1995	NHTS 2001	NHTS 2009
Total number of groups	431	420	423	428
Number of households per cell	178	154	265	229