How Much Do Consumers Value Fuel Economy and Performance? Evidence from Technology Adoption

Benjamin Leard* Resources for the Future

Joshua Linn Resources for the Future

Yichen Christy Zhou Clemson University

December 2017

Abstract

During historical periods in which US fuel economy standards were unchanging, automakers increased performance but not fuel economy, contrasting with recent periods of tightening standards and rising fuel economy. This paper evaluates the welfare consequences of automakers forgoing performance increases to raise fuel economy as standards have tightened since 2012. Using a unique data set and a novel approach to account for fuel economy and performance endogeneity, we find undervaluation of fuel cost savings and high valuation of performance. Welfare costs of forgone performance approximately equal expected fuel savings benefits, suggesting approximately zero net private consumer benefit from tightened standards.

Key words: passenger vehicles, fuel economy standards, technology adoption, consumer welfare

JEL classification numbers: D12, L11, L62, Q41

*Benjamin Leard (leard@rff.org) is a fellow at Resources for the Future (RFF). Joshua Linn (linn@rff.org) is a senior fellow at RFF. Christy Zhou (yichen2@clemson.edu) is an assistant professor at Clemson University.

We thank Jim Sallee, Ashley Vissing, and seminar participants at the University of Maryland, Carnegie Mellon University, American University, Resources for the Future, EPA NCEE, and participants at the IIOC (Boston), AERE (Pittsburgh), Transportation, Economics, Energy, and the Environment TE3 Conference at Michigan, APPAM (Chicago), SRA (Washington), TRB (Washington), and AEA/ASSA (Philadelphia) for valuable comments. We thank Kevin Roth for sharing vehicle scrappage data. We are grateful to the Sloan Foundation for supporting the research.

1 Introduction

Motivated by climate and energy security concerns, the US Environmental Protection Agency (EPA) and National Highway Traffic Safety Administration (NHTSA) impose standards for passenger vehicle greenhouse gas emissions and fuel economy. The agencies project that the current standards will roughly double new vehicle fuel economy between 2011 and 2025, substantially reducing fuel consumption and greenhouse gas emissions.

In their benefit-cost analysis, EPA and NHTSA conclude that the standards create climate and energy security benefits (EPA 2012; EPA et al. 2016). In addition to these social benefits, the agencies argue that the standards create private welfare benefits because there is a market failure for fuel economy, which is often referred to as the *energy efficiency gap*: vehicle manufacturers and consumers fail to adopt technologies and increase fuel economy even when the value of the fuel savings exceeds the adoption costs. An extensive literature (e.g., NRC 2015) concludes that a gap exists by identifying numerous specific fuel-saving technologies, the value of whose fuel savings exceeds the adoption costs. The agencies argue that the standards increase consumer welfare by stimulating the adoption of fuel-saving technologies and correcting distortions from the market failure. In fact, the value of the fuel savings to consumers accounts for about 70 percent of the estimated benefits of the standards. According to the agencies' analysis, the standards would increase social welfare even without counting energy security and climate benefits.[1]

The energy efficiency gap literature has focused on whether new vehicle consumers undervalue fuel savings, meaning that they are willing to pay less for fuel savings than the present discounted value of the savings.[2] Undervaluation would be consistent with the energy efficiency gap because it would imply that manufacturers have insufficient incentive to adopt fuel-saving technology. Earlier studies yielded a wide range of results, from approximately zero valuation to substantial overvaluation (see literature reviews by Helfand and Wolverton 2009 and Greene 2010), but recent studies by Busse et al. (2013) and Allcott and Wozny (2014) have found full or nearly full valuation, implying that there is not an energy efficiency gap and that standards are unlikely to increase private consumer welfare.

[1] The literature has established that fuel or carbon taxes are more efficient than fuel economy or emissions standards at reducing energy security or climate market failures (e.g., Jacobsen 2013). However, because fuel or carbon taxes do not address the market failure associated with the energy efficiency gap (Jaffe and Stavins 1994), if the gap is large enough, standards could be more efficient than fuel and carbon taxes (Fischer 2010; Parry et al. 2007).

[2] A variety of factors could explain undervaluation, such as incomplete information about fuel economy (Gillingham et al. 2009) and sticky information about fuel prices (Allcott and Wozny 2014).

Economists and policy makers have focused on the energy efficiency gap under the presumption that if there is a gap, tighter standards would raise private consumer welfare. We argue that this inference is incorrect because it ignores the effects of tighter standards on vehicle performance. Klier and Linn (2016) and Reynaert (2015) document that tighter standards cause manufacturers to trade off performance for fuel economy, causing performance to increase less than if standards had not tightened. Therefore, the private welfare effects depend on the valuation of the forgone performance. However, for reasons we explain below, estimates of willingness to pay (WTP) for performance in the literature are likely to suffer from omitted variables bias. Moreover, estimates of WTP for fuel economy that account for the endogeneity of fuel economy rely on fuel price variation to identify WTP, and may not be relevant to regulatory-induced changes in fuel economy. We present new estimates of WTP for fuel economy and performance that address these issues. We find strong evidence that consumers undervalue fuel economy, suggesting the presence of an energy efficiency gap. Notwithstanding the undervaluation, once we account for changes in performance, we find that recent tighter standards have had approximately zero net effect on private consumer welfare.

Next, we describe the paper in more detail. Knittel (2011) and Klier and Linn (2012) argue that manufacturers can respond to tighter standards by trading off performance for fuel economy. Manufacturers can use fuel-saving technology to increase fuel economy or performance (such as towing capacity), for example, by retuning the engine so that the new vehicle has the same fuel economy and greater performance than the original vehicle (Klier and Linn 2012; Whitefoot et al. 2013; Zhou 2016).[3] As we show in Section 2, during time periods when fuel economy standards were not changing, manufacturers used fuel-saving technology to increase performance while maintaining fuel economy, improving vehicle efficiency by about 2 percent per year (Knittel 2011). During periods when the standards tightened, manufacturers chose to trade off performance for fuel economy.

Because of the technological trade-offs, the effects of tighter standards on private consumer welfare depend on changes in vehicle prices, fuel economy, and performance. In the absence of tighter standards, manufacturers adopt fuel-saving technology and boost performance. Tighter standards have two effects on vehicle attributes. First, tighter standards increase the incentive to adopt fuel-saving technology, raising the rate at which manufacturers add technology, as Klier and Linn (2016) demonstrate. This effect raises

[3]For example, between 1980 and 2014, Honda adopted a number of fuel-saving technologies to double the Civic's horsepower without changing its fuel economy. Certain technologies, such as turbochargers, improve performance and reduce fuel economy, whereas other technologies increase fuel economy without affecting performance. When adopting fuel-saving technologies, manufacturers can combine these technologies and retune the engine to achieve the desired combination of fuel economy and performance increases.

vehicle fuel economy and production costs, which may increase vehicle prices. Second, tighter standards cause manufacturers to trade off performance for fuel economy. Note that manufacturers may use these responses to increase fuel economy, in addition to reducing the relative prices of vehicles with high fuel economy to increase their market shares (Jacobsen 2013).

Given these manufacturer responses, undervaluation implies that marginally tightening standards raises private consumer welfare if two conditions hold. The first condition is that the marginal profit from increasing performance equals the marginal profit from increasing fuel economy. If this condition does not hold, in the absence of tighter standards manufacturers are at a corner solution such that all fuel-saving technology adoption is devoted to improving performance while maintaining the level of fuel economy that the standards require. Consequently, tightening standards reduces consumer welfare by inducing a trade-off from performance to fuel economy. A second condition is that marginal WTP for performance equals the cost of adopting fuel-saving technology adoption. A setting in which this condition does not hold would imply a "performance gap" for the adoption of fuel-saving technology that is analogous to the energy efficiency gap. Section 5.2 discusses these conditions in more detail.

Thus, the central questions regarding the effects of standards on private consumer welfare are whether consumers undervalue fuel economy and whether one of the two conditions hold. Whether one takes a structural or reduced-form approach to answering these questions, it is necessary to estimate consumer valuation of fuel economy and performance. As we argue next, nearly all existing WTP estimates are likely to be biased because they do not address a fundamental omitted variables problem. Consequently, we focus on obtaining unbiased estimates of consumer WTP for fuel economy and performance.

We make two improvements over the existing literature. First, most studies either have not estimated WTP for performance or have assumed that performance is uncorrelated with unobserved vehicle attributes.[4] Because vehicle manufacturers simultaneously choose fuel economy, performance, and other attributes, fuel economy and performance are likely to be correlated with other unobserved attributes (Klier and Linn 2012). Most earlier studies (e.g., Berry et al. 1995) that estimate WTP for performance assume that performance is exogenous, but a few recent papers, such as Whitefoot et al. (2013), instrument for performance. These recent studies primarily rely on variation from the vehicle's fuel type or drive type (e.g., 4-wheel-drive). However, because consumers directly value fuel type and drive type, and not just their effects on fuel economy and performance, the instruments are likely to be

[4] Recent papers that focus on consumer valuation of fuel economy, including Busse et al. (2013), Allcott and Wozny (2014), and Sallee et al. (2016), do not attempt to estimate WTP for performance.

correlated with unobserved vehicle attributes. For example, automakers may provide better (unobserved) technology packages for 4-wheel-drive vehicles than for 2-wheel-drive vehicles, causing biased estimates.

Second, our empirical analysis pertains directly to policies that affect fuel economy and performance in the medium to the long run. In contrast, recent estimates of consumer valuation of fuel economy (e.g., Busse et al. 2013; Allcott and Wozny 2014) are based on consumer responses to changes in fuel costs induced by fuel price variation, focusing on fuel price variation in the 1990s and early 2000s, when the stringency of fuel economy standards was unchanging. These studies are relevant to policies that directly affect fuel prices, such as carbon or fuel taxes, and they are relevant for the short run, in which attributes of market vehicles are held fixed. However, emissions or fuel economy standards cause fuel economy to increase over time without directly affecting fuel prices (Whitefoot et al. 2013; Reynaert 2015). Consumers could respond differently to fuel prices in the short run and fuel economy in the medium and long run for a variety of reasons, such as information or uncertainty about fuel prices and fuel economy (Metcalf and Hassett 1993; Dixit and Pindyck 1994).

We use a unique data set and a novel empirical strategy to account for the endogeneity of both fuel economy and performance, identifying WTP from changes in these attributes rather than from changes in fuel prices. Our data include 535,124 observations of new vehicles that were purchased or leased between the fourth quarter of 2009 and the third quarter of 2014. For each vehicle, we observe a vehicle transaction price, household demographics, and a vehicle identification number (VIN), which we use to assign extensive vehicle characteristics such as fuel economy, horsepower, torque, and weight. To compare our results with the recent literature, we adapt the empirical framework of Busse et al. (2013) to estimate average WTP for fuel economy and performance across all consumers in the market. We adopt two strategies to account for the endogeneity. First, we include vehicle fixed effects, defining vehicles at a highly disaggregated level, to control for cross-sectional correlations among fuel economy, performance, and unobserved vehicle attributes such as technology packages and safety features. Second, we use instrumental variables (IVs) constructed from EPA microdata on fuel-saving technology adoption. The instruments are indicators for the adoption of specific technologies in individual vehicle trims, and they are strong predictors of fuel economy and performance, reducing concerns about weak instruments bias. We report evidence supporting the underlying exclusion restrictions.[5]

[5] Klier and Linn (2016) report rough welfare estimates of the forgone performance, but the underlying WTP estimates are subject to shortcomings noted in the text. The estimation in this paper improves on our previous attempts to address endogeneity of fuel economy and performance (Klier and Linn 2012; Zhou 2016), by using actual transaction prices rather than manufacturer suggested retail prices, and by

We find that consumers undervalue fuel cost savings arising from higher fuel economy. The preferred estimates imply that consumers use a real discount rate of 12 percent to discount future fuel cost savings, compared to reported real market interest rates of 1.3 percent in our sample. The fact that the implicit discount rate exceeds the market rates suggests that consumers undervalue the fuel cost savings. An equivalent interpretation is that if we use market rates to discount future fuel cost savings, consumers pay 54 cents for $1 of discounted future fuel cost savings. In contrast, Busse et al. (2013) find full valuation and Allcott and Wozny (2014) estimate that consumers pay 76 cents for $1 of discounted fuel cost savings. We obtain similar undervaluation as in our baseline using our data and the methodology in Busse et al. (2013), suggesting that differences in sample period, rather than methodology, explain the discrepancies. The lower WTP for the most recent period is consistent with Leard, Linn, and McConnell (forthcoming), who show that new vehicle purchases responded differently in the late 1990s and early 2000s (when fuel prices were low or rising) than in the late 2000s and early 2010s (when fuel prices were high and volatile, and when fuel economy was increasing).

Consumers are willing to pay $94 for a 1 percent performance increase arising from fuel-saving technology adoption. This corresponds to a WTP of $1,100 for a 1-second reduction in the time needed to accelerate from rest to 60 miles per hour (0-to-60 time), which lies in the middle of the range of estimates in the literature (e.g., Whitefoot and Skerlos 2012; Greene et al. 2016). Comparing the ordinary least squares (OLS) and IV estimates, we conclude that failing to account for the endogeneity of fuel economy and performance would understate consumer valuation of fuel economy and performance.

The WTP estimates have three implications. First, combining our WTP estimates with estimates of the technological trade-offs between fuel economy and performance (Knittel 2011; Klier and Linn 2016), suggests that consumers are willing to pay about three times as much for a performance increase as for a fuel economy increase. This result is consistent with the observation (documented below) that during the 1990s and early 2000s, when vehicle standards were not tightening, manufacturers adopting fuel-saving technology used the technologies to increase performance rather than fuel economy.

Second, the estimates imply that, after accounting for the welfare costs of lower performance, recently tightened standards appear to have had approximately zero net effect on private consumer welfare. We consider a hypothetical tightening of the standards by 1 percent during our sample period. Using technology cost estimates from Leard et al. (2016) (which are based on EPA 2012), and estimated trade-offs between fuel economy and

relaxing assumptions on consumer demand and the exogeneity of power train attributes. Copeland (2014) demonstrates the importance of using transaction prices rather than retail prices.

performance, we find that tighter standards reduce consumer welfare by 0.4 percent per vehicle sold. This implication contrasts with the conclusion that one would obtain by following the conventional approach that considers only the estimated undervaluation and ignores performance changes. In that case, one would estimate that tighter standards raise consumer welfare by 0.6 percent per vehicle. These results therefore demonstrate the importance of including forgone performance in analyzing the welfare effects of the standards.

Above, we noted that in the presence of undervaluation, tighter standards would raise consumer welfare if two conditions hold. In practice, it appears that neither condition holds, as the marginal WTP for performance relative to fuel economy exceeds the technological trade-off between the two attributes, and the marginal WTP for performance exceeds the technology adoption cost.

The third implication regards the effect of fuel economy or greenhouse gas standards on consumer demand for new vehicles. A particularly contentious aspect of the existing standards is whether they reduce aggregate consumer demand for new vehicles, which the marketing literature refers to as consumer acceptance of new vehicles. This possibility is a manifestation of vintage differentiated regulation (Gruenspecht 1982; Stavins 2005)—that is, the fact that the regulations apply to new vehicles but not existing vehicles. This form of regulation raises the cost of purchasing a new vehicle compared with the cost of purchasing a second-hand vehicle, reducing aggregate new vehicle demand. Lower demand reduces manufacturer profits, and by delaying the replacement of older with newer vehicles, lower demand also reduces the overall fuel and greenhouse gas savings of the standards (Jacobsen and van Benthem 2015). We find that tightening standards by 1 percent reduces WTP for new vehicles by $236, or 0.8 percent.

The results illustrate the importance of estimating WTP for performance, and of accounting for the endogeneity of fuel economy and performance to estimate WTP. Our preferred estimates of fuel economy valuation contrast with other recent estimates, in that we find strong evidence of undervaluation. Yet, once we include the welfare costs of lower performance in the analysis, we find that tighter standards have had approximately no net effect on private consumer welfare, which contrasts with the conclusion that one would obtain by ignoring the costs of lower performance.

2 Data and Summary Statistics

2.1 Data

We assembled the main data set from several sources, the most important of which includes household survey data collected by MaritzCX. Based on vehicle registration information, MaritzCX contacts households that recently obtained new vehicles. The survey is administered online or by mail, with a 9 percent response rate. Our data include households that obtained new vehicles between October 2009 and September 2014. The final sample includes 535,124 observations, which represents about 1 percent of all new vehicles obtained during the five-year period.[6]

The survey includes questions about the new vehicle and household demographics. For each transaction in MaritzCX, we use the transaction price net of state taxes, prior to a trade-in, and without adjusting for trade-in credit.[7] As in many other recent vehicle market analyses (e.g., Busse et al. 2013; Copeland 2014), we use the transaction price, rather than the manufacturer suggested retail price (MSRP), to reflect the outcome of any price negotiation or unobserved incentives for the vehicle. In practice, we observe substantial differences between the MSRP and transaction price. Household demographic characteristics in the data include state of residence, household size and income, and the respondent's age, years of schooling, gender, marital status, and other characteristics.

The MaritzCX survey data include a vehicle identification number (VIN) for each observation. We use the VIN to define a unique model-variant for each vehicle, which is the combination of a vehicle's manufacturer, make, model name, trim/series, fuel type, drive type, displacement, and number of cylinders. For example, a unique model-variant is the Toyota Lexus HS250H Premium, with front-wheel drive and a gasoline-powered engine that has four cylinders and 2.4-liter displacement. Our definition of model-variant is similar to the definition of a unique vehicle used in recent studies (e.g., Allcott and Wozny 2014). Note that two versions of the same model-variant can have different body types, which we also observe in the data. The final sample contains 2,166 unique model-variants and about 250 observations per model-variant (Table 1). Time is indexed by model year and quarter, and the same model-variant may be observed in multiple time periods.

The VIN allows us to obtain an extensive set of vehicle attributes that are not found in the MaritzCX data. We supplement the MaritzCX data with the Chrome Automotive

[6]The raw data include 930,000 observations of new vehicle transactions. We drop 262,000 observations with missing transaction prices, and 126,000 observations with missing vehicle attributes or fuel-saving technologies. Because of the IV strategy, we exclude plug-in vehicles (both all-electrics and plug-in hybrids), which account for less than 1 percent of the MaritzCX sample. As explained in the text, we weight observations in the final data set to reflect non-random sampling, response rates, or missing data.

[7]These transaction price data are provided by survey respondents about a month after making a purchase. Given the short recall time and the high price associated with a new vehicle purchase relative to other consumer durable purchases, there is little risk of recall bias and these data are likely to accurately represent actual transaction prices. Some recent studies have used transaction prices reported by marketing companies such as J.D. Power. Unfortunately, based on personal correspondence, J.D. Power are not currently available for purchase by academic research teams.

Descriptive Service database, and use the VIN to obtain vehicle characteristics such as vehicle weight, horsepower, and torque.

In the empirical analysis, we use the ratio of horsepower to weight as a proxy for passenger car performance, and the ratio of torque to weight as a proxy for light truck performance. The performance definition follows previous studies that estimate vehicle demand, such as Berry et al. (1995), and we use different measures for cars and light trucks. Car consumers typically have stronger preference for acceleration (which is closely related to the ratio of horsepower to weight) than for towing ability, whereas light-truck consumers often have stronger preference for towing ability than acceleration (Knittel 2011). We note that several aspects of vehicle performance may affect consumer purchasing decisions, such as the time needed to accelerate from rest to 60 miles per hour, or the time needed to accelerate from 20 to 50 miles per hour (which is more relevant in certain situations such as merging onto a highway). In practice, these performance measures are highly correlated with one another. For example, the ratio of horsepower to weight accurately predicts 0-60 time (Greene et al. 2016; Linn 2016). The results are similar if we use the ratio of horsepower to weight for all vehicles rather than just for passenger cars.

We obtain fuel economy ratings (miles per gallon, mpg) and fuel-saving technology data from EPA.[8] The technology data include indicator variables for whether the vehicle has variable valve lift and timing, a turbocharger, a supercharger, gasoline direct injection, cylinder deactivation, continuously variable transmission, and other advanced transmissions. NRC (2015) concludes that each of these technologies raises a vehicle's fuel economy as well as production costs, holding fixed all other attributes including performance. For example, NRC (2015) estimates that cylinder deactivation, which effectively shuts off a subset of a vehicle's engine cylinders when the vehicle operates under a light load, raises fuel economy by as much as 5 percent, and raises production costs by \$118 to \$133 per vehicle. Because EPA data do not recognize potential differences in fuel economy across body types within a model-variant, we merge EPA data by vehicle model year and model-variant. Therefore, fuel economy and fuel-saving technologies can vary across model-variants but not within model-variants, and the definition of the model-variant preserves 99 percent of the EPA estimated fuel economy variation across new vehicles.[9]

[8] https://www3.epa.gov/fueleconomy/data.htm.

[9] We do not include fuel-saving technologies that were widely adopted at the beginning of the sample, such as variable valve timing, or technologies that consumers value directly (either negatively or positively), such as stop-start ignition. The EPA data include more detail on transmissions than Chrome. We average the technology variables across transmission type (automatic or manual), and for most observations in the final data set the technology variables are either zero or one, implying that the aggregation sacrifices little variation. Below we refer to the technology variables as indicator variables for convenience.

To correct for the non-random sampling of the MaritzCX survey, we obtained data on US national vehicle registrations from Information Handling Service Market (IHS Market). We observe the number of new vehicles registered by model year, model-variant, and body type for all vehicles registered each quarter in the United States from October 2009 through September 2014. We link the IHS to the MaritzCX data by vehicle model year, model-variant, body type, year and quarter of the transaction. As we show below, although the initial sample is not random, the weighted sample matches the distribution of new vehicle buyers from other data sources.

Monthly fuel prices come from the US Energy Information Administration (EIA). The data set includes the average monthly gasoline prices and diesel fuel prices by Petroleum Administration for Defense District (PADD), for each of four districts (Midwest, Gulf Coast, Mountain, and West Coast), and three subdistricts on the East Coast. When constructing the fuel cost variables described in the next section, we use gasoline prices for gasoline powered vehicles and flex-fuel vehicles, and diesel fuel prices for diesel fuel powered vehicles.[10] We deflate all transaction and fuel prices using the Consumer Price Index, and adjust them to 2010 US dollars.

We use measures of lifetime fuel costs in post-estimation calculations. Lifetime fuel costs are estimated from annual vehicle miles traveled (VMT) data from the 2009 National Household Travel Survey (NHTS), and proprietary data from R. L. Polk on annual scrappage rates from 2003-2014. Using the NHTS, we estimate average VMT by model year, income group, and vehicle class (cars or light trucks) following the methodology in Lu (2006). With the R. L. Polk data, we estimate a survival rate as a function of vehicle age following Lu (2006). The estimated schedules appear in Appendix Table B.5. We assume that vehicles have a maximum lifespan of 35 years for cars and 40 years for light trucks. Appendix A.1 explains the methodology for computing scrappage rates and VMT in more detail.

2.2 Summary statistics

We report summary statistics from the main data set, discussing vehicle attributes first and consumer demographics second. Panel A of Table 1 provides information about the distributions of certain vehicle characteristics. Observations are weighted by registrations, and the table indicates that most vehicles in the sample use gasoline rather than diesel fuel (recall that the sample excludes plug-in vehicles). Mean fuel economy is about 23.9 mpg, and the table indicates substantial variation in fuel economy and performance.

Figures 1 to 3 illustrate time series variation in several vehicle attributes and technologies. We plot registration-weighted model-year averages of vehicle attributes and

[10]Flex-fuel vehicles can use fuel that has a high ethanol content, but in practice few owners of flex-fuel vehicles use gasoline with ethanol content greater than 10 percent (Anderson and Sallee 2011).

technology adoption rates over time. The fuel economy standards for light trucks tightened throughout the period, and the standards for cars began tightening in model-year 2012. Figure 1 shows that average fuel economy increases after 2011. Horsepower, torque, and weight fluctuate over the same period.

Figure 2 reports statistics for engine and transmission attributes. Engine size, as measured by the number of cylinders or displacement, decreases over the sample period. Market shares of the three drive types are fairly stable over the time period. The market share of diesel fuel vehicles increases between model years 2010 and 2014 (the Volkswagen emissions scandal occurred after the end of the sample). The market shares of hybrids and flex-fuel vehicles decrease at the end of the sample. The latter may reflect the elimination of the flex-fuel vehicle credits that manufacturers could use to demonstrate compliance with the fuel economy standards (Anderson and Sallee 2011).

Figure 3 shows the market shares of fuel-saving technologies that we use to instrument for fuel economy and performance. In most cases the market shares increase over time, such as an increase in the gasoline direct injection market share from 9 to 56 percent. Most decreases in this figure arise from year-to-year changes in vehicle market shares rather than instances of manufacturers removing technologies from particular vehicles. Klier and Linn (2016) and Klier et al. (2017) suggest that tightening fuel economy standards as well as market factors such as fuel prices explain the technology adoption.

Figures 4 and 5 illustrate monthly variation in fuel prices and vehicle prices, with each dashed vertical line indicating the beginning of a calendar year. Although we do not use fuel prices to identify WTP for fuel economy, for context we summarize the fuel price variation during the sample. Panel A of Figure 4 shows that the sample includes periods of rising fuel prices (2009 through mid-2011) and volatile or declining fuel prices (mid-2011 through 2014). Panel B shows that regional prices are positively correlated with one another, and that prices in the West Coast and Midwest regions tend to be higher than in other regions. Regional price differences vary somewhat over time. Both Figures 4 and 5 indicate regular seasonal variation. Fuel prices tend to be higher in the summer than in other quarters, and vehicle prices tend to increase over the year, before decreasing at the end of the year.

Turning to consumer attributes, Panel A of Figure B.1 displays a histogram of the reported income distribution. The modal income is $75,000 to $100,000. Typical household income of vehicle buyers in our sample is higher than the typical US household income during this period, which reflects the fact that higher-income households are more likely than lower-income households to obtain new vehicles. The income distribution in our data is fairly close to the income distribution of new vehicle buyers as reported in the 2009 wave of the NHTS, which is a nationally representative survey. Panel B of Table 1 shows further information about the households in the sample, including average household size as well as the age, gender, urbanization, and marital status of the respondent.

Table B.1 reports information on the form of payment used to obtain the vehicle. About two-thirds of consumers finance their purchases, with an average nominal loan rate of 3.34 percent for about 5 years. About one quarter of consumers purchase their vehicles entirely via cash, and the remainder lease their vehicles.

Table 2 shows changes in vehicle fuel economy and horsepower since 1996 (we use data from Leard, Linn, and McConnell (forthcoming)). Recall that fuel economy standards for light trucks began increasing in 2005 and fuel economy for cars began increasing in 2012. The table shows that fuel economy increased much more quickly and horsepower increased much more slowly during periods when standards tightened; Klier and Linn (2016) demonstrate that the standards caused these changes. This evidence motivates our analysis of the effects of tightening standards on private consumer welfare, accounting for changes in fuel economy as well as performance.

3 Empirical Strategy

3.1 Empirical framework

Our empirical objective is to estimate consumer valuation for fuel economy and performance. We adapt the approach taken by Busse et al. (2013), which is to estimate separate reduced-form price and quantity regressions, and combine the results to estimate WTP. To illustrate this approach, we consider a hypothetical manufacturer that produces a single type of vehicle. For convenience, we conceive of a Bertrand model with heterogeneous products. (As we explain below, the empirical strategy does not depend on the underlying market structure.) We abstract from fuel economy and emissions standards for simplicity, and control for those standards in the empirical analysis as described below. The manufacturer faces a downward-sloping residual demand curve for that vehicle. We define the WTP for a fuel economy increase as the vertical shift of the demand curve caused by the fuel economy increase; WTP for a performance increase is defined similarly. The definition holds fixed all other attributes of the vehicle.

Figure 6 provides the intuition behind this approach. We describe the initial equilibrium using demand curve D_1 and marginal cost curve MC_1. The manufacturer chooses the price such that at the resulting quantity, Q_1, the marginal revenue curve (indicated by the downward sloping dashed line) intersects the marginal cost curve MC_1.

The figure illustrates a hypothetical situation in which the manufacturer adopts fuel-saving technology and increases the vehicle's fuel economy. The higher fuel economy reduces fuel costs, causing the demand curve to shift to D_2. The technology adoption increases marginal costs to MC_2, which results in the equilibrium price of P_2 and equilibrium quantity of Q_2.

The consumer WTP for the fuel economy increase corresponds to the vertical shift of the demand curve, which is equal to the sum $l_1 + l_2$. As explained in the next subsection, we use a regression of the vehicle's equilibrium price on its fuel costs to identify the first part of the sum, $l_1 \equiv P_2 - P_1$. We use a quantity regression to identify the equilibrium quantity effect $Q_2 - Q_1$. The term l_2 depends on the equilibrium quantity change, as well as the slope of the demand curve. Therefore, to estimate WTP for fuel economy, we estimate the effects of fuel economy on the equilibrium price and quantity, and calculate WTP by assuming a particular slope of the demand curve. As Busse et al. (2013) note, an alternative approach would be to estimate the demand curve directly, which would require certain assumptions on the structure of the demand at the outset. In contrast, the reduced-form approach requires only an assumption on the slope of the demand curve, which is made after estimating the two equations. An advantage of the reduced-form approach is that it facilitates accounting for the endogeneity of fuel economy and performance. Below, we show that the main conclusions are insensitive to the assumed demand elasticity.

3.2 Price regression

This subsection describes the estimation of the equilibrium relationship between a vehicle's transaction price, p_{ijt}, and its attributes, where the subscript indicates that household i obtained new passenger vehicle j in month t. The approach is similar to that taken in the hedonic literature (e.g., Rosen 1974). Specifically, we assume a log-log relationship between price and attributes:

$$\ln p_{ijt} = \alpha_f \ln fc_{ijt} + \alpha_p \ln perf_{jt} + X_{ijt}\delta + \varepsilon_{ijt} \tag{1}$$

where fc_{ijt} is the vehicle's fuel costs; $perf_{jt}$ is the vehicle's performance; X_{ijt} is a vector of variables described next; e_{jt} is an error term; and the αs and δ are coefficients to be estimated. The performance variable is the horsepower-to-weight ratio for cars and the torque-to-weight ratio for light trucks. The vector X_{ijt} includes PADD-month-fuel type fixed effects to account for aggregate and regional supply and demand shocks, as well as seasonality in fuel or vehicle prices (see Figures 4 and 5); state fixed effects to control for state-level demand or supply shocks; a model-year fixed effect to control for macroeconomic shocks and the demand for used vehicles; an indicator if the vehicle has flex-fuel capability; fixed effects of the number of transmission speeds, as well as the interactions of these variables with an indicator equal to one if the vehicle is a light truck; and controls for fuel economy regulatory stringency. At the end of the subsection, we explain the motivation for controlling for transmission speeds and flex-fuel capability.

The controls for fuel economy regulatory stringency incorporate two sources of stringency variation. First, under the current standards, a vehicle's fuel economy requirement depends on its size; manufacturers selling smaller vehicles must attain a higher overall level of fuel economy. Second, at the outset of the sample period, manufacturers varied in the difference between the level of fuel economy required by the standards and the level of fuel economy their vehicles actually attained (Jacobsen 2013). Stringency is measured as in Klier and Linn (2016), by computing the difference between the fleet level fuel economy a manufacturer must attain to meet the standards in model-year 2016 and the manufacturer's average fuel economy at the beginning of the sample. The stringency variable is interacted with model-year fixed effects, to allow for the possibility that regulatory pressure varies over time.

In equation (1) we separate fuel costs and performance from the other attributes because estimating separate consumer valuation of fuel costs and performance is the main focus of the paper. The fuel cost variable (measured in dollars per mile) is equal to the price of fuel in the month and the PADD region in which the vehicle is obtained, divided by the vehicle's fuel economy (mpg). Under the assumption that the expected real fuel price follows a random walk, which is consistent with Anderson et al. (2013), the ratio of the fuel price to fuel economy is proportional to the present discounted value of fuel costs over the lifetime of the vehicle (Busse et al. 2013). The PADD-month-fuel type interactions absorb the direct effect of fuel prices on fuel costs, because of which the coefficient α_f is identified by fuel economy variation.

Because the price, fuel costs, and performance variables enter equation (1) in logs, the coefficients represent elasticities. We expect the fuel cost coefficient to be negative because higher fuel costs raise the total cost of the vehicle over its lifetime, and we expect the performance coefficient to be positive. We interpret these estimates as the effect of fuel economy or performance on the average transaction price across all vehicles in the market.[11] The interpretation of the coefficients does not depend on the underlying demand or competitive structure of the market.

Note that we do not interpret the fuel cost and performance coefficients as being proportional to parameters in a consumer's utility function. The log-log functional form is not derived from an underlying utility function. Rather, we use the log-log functional form to approximate equilibrium relationships among vehicle prices and attributes; likewise, (Busse et al. 2013) use a functional form that approximates an equilibrium relationship rather than deriving the functional form from a utility function.

[11] We have estimated versions of equation (1) that allow the fuel cost and performance coefficients to vary across vehicles, such as by car or light truck. Although we find some evidence that the coefficients vary across vehicles, in many cases the differences are imprecisely estimated.

Although equation (1) yields a straightforward economic interpretation of the coefficients, the main identification concern is that the vehicle characteristics included in the regression may be correlated with omitted vehicle or household characteristics. For example, vehicles with high performance may include more comfortable seating or better entertainment devices than vehicles with lower performance. Although our data include an extensive set of characteristics, and more than the vehicle demand literature has typically used, we do not observe all vehicle characteristics that consumers value. For example, we observe seating material (cloth vs. leather), but overall seating comfort depends on other factors, such as lower back support, which our data do not include. OLS estimates of equation (1) would be biased if we fail to include all vehicle attributes that consumers value.

For expositional purposes we use the term *quality* to refer to the combined effect of all unobserved vehicle characteristics on the equilibrium price. The term includes seating comfort, entertainment devices, and anything else about the vehicle that consumers value but that is not included in equation (1). Quality also depends on consumer perceptions of the unobserved attributes. Using this definition, quality can vary across vehicles and within a vehicle over time. Obtaining unbiased estimates of WTP for fuel costs or performance therefore amounts to controlling for quality.

One approach to control for quality would be to include a full set of model-variant fixed effects—i.e., to adapt the approach taken in Busse et al. (2013) and several other recent studies of new vehicle demand. The fixed effects control for time-invariant vehicle quality, but they do not fully address the potential omitted variables bias because within-model-variant changes over time in fuel economy or performance may be correlated with changes in quality. Specifically, when a manufacturer redesigns a model-variant and alters its fuel economy or performance, it may change other vehicle quality attributes at the same time; the fixed effects do not control for such changes. Moreover, the fixed effects do not control for changes in consumer perceptions over time.

We could include interactions of model-variant fixed effects and model year, and identify the fuel cost coefficient by cross-sectional and time series variation in fuel prices. However, there would be two problems with this approach. The first is that the coefficient would be identified by fuel price variation rather than fuel economy variation. As we argued in the introduction, the consumer response to fuel economy is directly relevant to standards that affect fuel economy and not fuel prices, and consumers may respond differently to the two sources of variation in fuel costs. The second problem is that it is not possible to identify the performance coefficient because the model-variant by year interactions would be perfectly colinear with performance.

Given these considerations, we address potential omitted variables bias in equation (1) by adding vehicle model-variant fixed effects and instrumenting for fuel costs and performance. The estimating equation is

$$\ln p_{ijt} = \alpha_f \ln fc_{ijt} + \alpha_p \ln perf_{jt} + X_{ijt}\delta + \eta_j + \varepsilon_{ijt} \tag{2}$$

where η_j denotes a fixed effect for vehicle model-variant j. There is no fuel economy variation within a model-variant and model year, but fuel economy can vary across model-variants and within a model-variant over time. The fixed effects absorb the vehicle's fuel type and whether the power train is a hybrid, but they do not absorb the number of transmission speeds or whether the vehicle is flex-fuel capable. Consequently, we include those attributes in X_{ijt}. The instruments are seven indicators for the fuel-saving technologies shown in Figure 3: variable valve lift and timing, turbocharger, supercharger, gasoline direct injection, cylinder deactivation, continuously variable transmission, and other advanced transmissions. EPA (2014) and NRC (2015) identify these technologies as improving the efficiency of the engine or transmission. We further interact these instruments with an indicator equal to one if the vehicle is a light truck, which allows for the possibility that the technologies have different effects on fuel economy or performance across cars and light trucks (NRC 2015). Because of the model-variant fixed effects in equation (2), the first stage is identified by variation within a model-variant in fuel economy, performance, and technologies; that is, roughly speaking, by the time series variation illustrated in figures 1 and 3. The fact that fuel costs and performance enter equation (2) in logs is consistent with engineering assessments of the technologies that indicate that they affect fuel economy proportionately. That is, using the level of fuel costs rather than the log would be inconsistent with the technological relationships between the instruments and fuel economy.

Variation of the instruments arises from the tightening fuel economy and emissions standards, combined with the timing of vehicle redesigns. During the period of analysis, fuel economy standards tightened by about 4 to 5 percent per year after a long period in which they were unchanged. As Klier and Linn (2016) show, the tighter standards doubled the rate at which technologies were adopted, causing adoption to be more widespread across vehicles in the market than previously observed. Vehicles are typically redesigned in 4- to 6-year cycles, and manufacturers stagger the redesigns across vehicles. Because of the staggering, manufacturers do not adopt technologies simultaneously on all of their vehicles. Note that because we control for regulatory stringency, the first stage is identified by variation induced by the tightening standards interacting with staggered vehicle redesign.

The IV strategy is valid if the instruments predict fuel economy and performance and are uncorrelated with the error term in equation (2). Failing to satisfy the first condition would raise concerns about weak instruments bias. However, the results reported in the next section indicate a strong correlation among the instruments, fuel economy, and performance, minimizing such concerns. Moreover, the results in the next section indicate that the values of fuel costs and performance predicted in the first stage are sufficiently uncorrelated with one another that we can identify the coefficients on fuel costs and performance in the second stage, equation (2).

The second condition is supported both by theoretical arguments that we present in this section and by empirical evidence that we present in the next section. First, we choose technology variables that consumers do not value per se (as opposed to the fuel economy or performance increase that they enable). If consumers valued the technologies, the technologies would violate the second condition because they would be correlated with the error term in equation (2). For this reason, we exclude technologies for which there are widespread reports of consumer dissatisfaction. For example, the Atkinson cycle gasoline-powered engine that Mitsubishi installed in some of its vehicles received negative reviews from consumers because it harmed performance or other vehicle attributes.[12] This feature of the IVs represents an improvement over other studies, such as Whitefoot et al. (2013), which have used power train characteristics as instruments because consumers likely value those characteristics directly, yielding biased WTP estimates.

Second, the fact that the standards roughly doubled the rate of technology adoption implies that manufacturers focused more on adopting technology during redesigns than they do typically. The source of technology variation is distinct from typical decisions about whether to install technology, when manufacturers may be more likely to redesign the vehicle to adopt technology as well as improve quality. For example, given time and resource constraints for redesigning vehicles, during our sample period a manufacturer is less likely to change vehicle quality in response to a demand shock than during prior periods in which standards were not tightening. Therefore, the tightening standards, combined with staggered redesigns, reduces the likelihood that the technology variables are correlated with quality.

Note that this consideration reduces concerns that household demographics, which equation (2) does not include, may be correlated with quality. For example, high-income households may have higher WTP for seating comfort. The fact that the standards drove fuel-saving technology adoption during the sample period reduces the likelihood that omitted demographics are correlated with quality; in the robustness analysis below, we show that the instruments are uncorrelated with demographics.

[12] There have been a few negative reports related to consumer perceptions of continuously variable transmission and cylinder deactivation. We prefer to include them because these technologies have been widely adopted (see Figure 3), and because the negative reports are scarce. In the robustness analysis we show that the coefficient estimates are similar if we omit these variables as instruments.

Third, manufacturers sometimes adopt fuel-saving technology in luxury vehicles before adopting it in other vehicles. This behavior would cause technology adoption to be correlated with unobserved quality at any point in time. For example, manufacturers may adopt technology first for luxury versions of a particular model, or they may adopt technology first for higher-end models prior to lower-end models (such as a Lexus sedan prior to a Toyota sedan). The model-variant fixed effects address cross-sectional and time-invariant correlations between quality and technology adoption. For example, the fixed effects control for situations in which a luxury vehicle has a fuel-saving technology throughout the sample period, whereas another vehicle does not have the technology during the period.

The main remaining concern is that manufacturers simultaneously change quality and adopt technology. We have argued that this is less likely to be the case during our sample than during historical periods of technology adoption. Moreover, Section 4.3 shows that the results are robust to adding several proxies for quality to equation (2).

3.3 Quantity regression

The empirical strategy for the quantity regression is similar to that for the price regression. We use the log of quarterly registrations as the dependent variable and estimate the equation at the household level:

$$\ln q_{jt} = \beta_f \ln fc_{ijt} + \beta_p \ln perf_{jt} + X_{ijt}\gamma + \xi_j + \nu_{ijt} \tag{3}$$

where the independent variables are the same as in equation (2). We use vehicle fixed effects and the same instruments to account for the endogeneity of fuel economy and performance. Note that vehicle fixed effects ξ_j are defined by trim, fuel type, drive type, and body type, to match the aggregation of the registration data. Because of the fixed effects, as with equation (2), in equation (3) the fuel cost coefficient is identified by variation in fuel economy rather than fuel prices.

The fact that the fuel cost and performance coefficients in equation (3) are identified by the same variation as the corresponding coefficients in equation (2) is an important aspect of our empirical strategy because it implies that the coefficients are identified by the same underlying consumer preferences and manufacturer supply responses. Consequently, we interpret the coefficients in both equations as the average equilibrium effects across vehicles in the market. In contrast, if we were using different estimation samples or empirical strategies for the two equations, one might be concerned that the coefficients represent averages across different sets of vehicles, in which case it would not be appropriate to combine the results to infer WTP for fuel economy and performance.

An important difference between interpreting the price and quantity regressions is that for the quantity regressions the signs of the fuel cost and performance coefficients are ambiguous. On the one hand, an increase in fuel economy (or performance) causes the demand curve to shift away from the origin, increasing equilibrium quantity (see Figure 6). This effect would cause a negative fuel cost coefficient and a positive performance coefficient. On the other hand, because the manufacturer adopts technology to raise fuel economy or performance, marginal costs increase, which reduces equilibrium quantity and pushes the coefficients in the opposite direction as the demand curve shift. The net equilibrium effect on quantity is ambiguous.

4 Estimation Results

4.1 Baseline estimates of willingness to pay for fuel cost savings and performance

Table 3 reports the main coefficient estimates. Column 1 shows the OLS estimates of equation (1) and the corresponding quantity regression, and column 2 includes model-variant fixed effects instead of the vehicle attributes that define the model-variant. We report the OLS results for comparison with our preferred IV estimates of equations (2) and (3) in column 3. The regressions include the independent variables indicated in the table notes, which control for demand and supply shocks at the regional, monthly, or state level, as well as for the stringency of fuel economy standards. The model-variant fixed effects in columns 2 and 3 control for model-variant-level unobservables that may be correlated with fuel costs or performance. Table B.2 reports the first stage estimates for fuel costs and performance.[13]

Because the transaction price, fuel costs, and performance enter equations (2) and (3) in logs, we interpret the fuel cost and performance coefficients as elasticities. Panel A reports the estimates of the price regression, equation (2). Comparing columns 1 and 2 shows that the model-variant fixed effects increase the magnitude of the fuel cost coefficient. Comparing columns 2 and 3, the OLS estimate of the fuel cost coefficient is -0.156, and the IV estimate is -0.354, both of which are negative and statistically significant at the one percent level. In both columns the fuel cost coefficient is identified by fuel economy variation because the other independent variables absorb the fuel price variation. The larger magnitude of the IV estimate suggests that time-varying quality is positively correlated with fuel costs (and negatively correlated with fuel economy), which biases the OLS estimate toward zero. The OLS estimate of the performance coefficient in column 2 is

[13] Some of the first stage coefficients have unexpected signs, which appears to be due to the high correlation among the instruments. Below we confirm the overall positive relationships among technology adoption, fuel economy, and performance.

negative, implying counterintuitively that in equilibrium consumers pay less for vehicles with better performance (i.e., those having a higher ratio of horsepower or torque to weigh). In contrast, the IV estimate of the performance coefficient is 0.203, which is positive and significant at the one percent level, suggesting that consumers are willing to pay for better performance. Comparing the OLS and IV estimates of the performance coefficient in columns 2 and 3 suggests that when model-variant fixed effects are included, unobserved quality is negatively correlated with performance. Thus, failing to account for the endogeneity of fuel costs and performance yields substantially biased estimates; adding model-variant fixed effects to the OLS equation in column 1 does not address the omitted variables bias.

Panel B reports the estimated coefficients from the quantity regression, equation (3). In column 3 the IV coefficient on fuel costs is -0.338 and the coefficient on performance is 0.371, both of which are statistically significant at the one percent level. Whereas Busse et al. (2013) find larger quantity than price responses, we find quantity and price responses of comparable magnitudes to one another. Below we discuss potential explanations for the differences between our results and theirs.

We briefly discuss the economic magnitudes of the estimated coefficients on fuel costs and performance. The baseline estimates in column 3 suggest that a 1 percent fuel economy increase (which reduces fuel costs by 1 percent) raises the equilibrium transaction price and quantity by about 0.3 percent. A 1 percent performance increase raises the transaction price by 0.2 percent and raises the quantity by 0.4 percent. To convert these estimates to WTP, we first compute the marginal equilibrium price effect (l_1 in Figure 6) using the price regression coefficients. Then we adjust for the quantity change (l_2 in Figure 6) using the the quantity regression coefficients and the assumed own-price elasticity of demand. For the baseline we assume an elasticity of -3, which lies in the middle of the range considered in Busse et al. (2013).[14]

Panel C converts the coefficient estimates to estimates of the WTP for a 1 percent fuel economy or performance increase. The baseline estimates suggest that consumers are willing to pay about \$133 for a 1 percent fuel economy increase and about \$94 for a one percent performance increase. The OLS estimates in column 1 are positive, as expected, but they are smaller than the IV estimates. The OLS estimates in column 2 yield a larger WTP for fuel economy than the preferred IV estimate, but an implausibly negative WTP for performance. For the IV estimates, Appendix Table B.4 reports estimates of l_1 and l_2; l_1

[14]Because the dependent variables are logs of price and quantity, to predict the levels of prices and quantities we would need to account for fact that the error term is log-normally distributed. However, because we are interested in changes in prices and quantities caused by attribute changes, the correction term cancels in these calculations, yielding unbiased WTP estimates.

explains 76 percent of the WTP for fuel economy and 62 percent of the WTP for performance. Using the estimated relationship between the ratio of horsepower to weight and 0-to-60 time from Greene et al. (2016), the performance coefficient estimate implies that consumers are willing to pay about $1,100 for a 1-second decrease in 0-to-60 time, which is similar to many estimates in the literature.[15]

4.2 Do consumers undervalue fuel cost savings?

In this section we use two measures of consumer valuation from the literature to interpret the magnitude of the fuel cost coefficients in column 3 of Table 3. The next section compares this magnitude with the performance estimate and draws implications for the energy efficiency gap.

The first measure is the valuation ratio, which is the amount the marginal consumer is willing to pay for a 1 percent fuel economy increase divided by the present discounted value of the associated future fuel cost savings. If the ratio equals one, the consumer fully values the fuel economy improvement; a value less than one implies undervaluation and a value greater than one implies overvaluation.

The amount the consumer pays for the fuel economy increase is reported in Panel C of Table 3, i.e., $133. For a vehicle purchased in year y, the present discounted value of future fuel costs is given by $PDV_{fc} = \sum_{\tau=y}^{y+T} \frac{\pi_\tau V_\tau f_\tau}{m(1+r)^\tau}$, where T is the maximum lifetime of the vehicle, π_τ is the probability that the vehicle is not retired before year τ (which is sometimes referred to as the survival probability rate), V_τ is the number of miles the vehicle is driven in year τ, f_τ is the real fuel price in year τ, m is the vehicle's fuel economy, and r is the real discount rate. See Section 2.1 for a summary of the methodology for estimating T, π_τ, and V_τ, and Appendix Sections A.1 and A.2 for details. The real discount rate r is computed using the observed average annual percentage rate (APR) adjusted by the average inflation rate. For consumers who lease or finance their purchases, the rate represents the opportunity cost of the monthly lease or loan payments. For consumers paying by cash, the rate represents the opportunity cost of investing the cash in other financial instruments (Allcott and Wozny 2014). In our sample, the average borrowing rate is about 3.3 percent and the average inflation rate is 2.0 percent, implying a 1.3 percent real borrowing rate. We

[15]In theory, households expecting to drive their vehicles intensively should have higher WTP for fuel economy than other households. We test this hypothesis using survey information about the household's expected annual miles traveled for the new vehicle. We compute the average mileage by household income group and vehicle type (car or light truck). We add to the baseline specification the interaction of this variable with log fuel costs. The interaction term has the expected positive sign (see Table B.9). The magnitude of the interaction coefficient implies relatively little variation across households. The estimated WTP for performance is similar to the baseline.

set household discount rates equal to this real borrowing rate.[16] Given the evidence reported in Anderson et al. (2013), we assume that real fuel prices follow a random walk, in which case the current price equals the expected real future price. We note that Allcott and Wozny (2014) and Sallee et al. (2016) directly estimate the valuation ratio, whereas we estimate the WTP and calculate the valuation ratio subsequently; inferences for consumer undervaluation do not depend on the approach. We choose this approach because it facilitates computation of multiple measures of consumer valuation that we can compare with the broader literature.

Panel A of Table 4 reports the valuation ratio results. The baseline calculation of the fuel cost savings is $249. For consistency with the WTP in Table 3, we weight the fuel cost savings $PDV_{fc,j}$ of each vehicle using the number of registrations. Combining our calculation of fuel cost savings, $249, with the WTP in Table 3 Panel C, we compute a valuation ratio of 54 percent, meaning that the marginal consumer pays 54 cents for $1 of present discounted fuel cost savings (where future fuel costs are discounted using the market rate). This valuation ratio is lower than the 76 percent reported in Allcott and Wozny (2014) and 100 percent in Sallee et al. (2016), but as we noted in the introduction, the broader literature has yielded a wide range of valuation ratios, from close to zero to much greater than 1.

Computing the valuation ratio requires a number of assumptions, and we report alternative calculations based on differing assumptions. Busse et al. (2013) evaluate the extent of consumer undervaluation using the same methodology from Lu (2006), but using older data than we use. If we use their data instead of ours, the present discounted value of fuel cost savings declines from $249 to $184. Using their data we obtain a valuation ratio of 73 percent, showing that the undervaluation is robust to the choice of data.

Table B.6 shows that the undervaluation is robust to other demand elasticities. The table also reports results using alternative real discount rates that have been used in the literature, of 5, 7, 10, and 12 percent. To put these alternative higher discount rates in context, a 7 percent real discount rate is about the national average interest rate for a 24-month personal loan, and 12 percent is close to the credit card real interest rate in our sample period.[17]A potential argument for using the credit card rate as the discount rate is that a substantial share of US households have credit card debt, and for these households the credit card rate would represent the marginal cost of borrowing. However, new vehicle buyers have

[16]Alternatively, for households paying cash and not taking out an auto loan, we could impute their discount rate using other market rates, such as the real rate of return of stocks or bonds. We prefer to use the APR because households that paid for their vehicle with cash could have taken out an auto loan that would have had a similar APR to the average APRs we observe. The decision not to take out a loan reveals that the APR is an upper bound to the opportunity cost of funds for these households. That is, if the opportunity cost of funds were higher than the APR, we would observe these households taking out auto loans and purchasing higher-yield investments. We evaluate the sensitivity of this assumption as a robustness check.

[17]Data on credit card interest rates are from the federal reserve: here.

higher income than typical households, and are less likely to have credit card debt than the typical household. In our sample, about 75 percent of survey respondents report having perfect credit with no late payments. For these households, it would be inappropriate to use credit card rates as the discount rate because the credit card rate does not represent the marginal cost of borrowing. Thus, the conclusion about undervaluation is robust to using discount rates that are appropriate for our sample. Moreover, we find undervaluation if, instead of assuming that fuel prices follow a random walk, we use projected fuel prices from the Energy Information Administration's Annual Energy Outlook. Thus, we consistently find undervaluation when we vary the survival probability, miles traveled, demand elasticity, discount rate, and fuel price projection.

We report a second valuation measure, which is the implicit discount rate. This is the discount rate that implies a valuation ratio equal to one. In other words, if a consumer uses the implicit discount rate to discount future fuel cost savings, the consumer would be willing to pay $134 (i.e., the amount reported in Panel C of Table 3) for a 1 percent fuel economy increase. An implicit discount rate equal to market borrowing rates would imply full valuation of fuel economy increases; a discount rate higher than market rates would imply undervaluation; and a discount rate below market rates would imply overvaluation. Panel B in Table 4 reports the baseline estimated implicit discount rate of 12 percent. This is much higher than the average reported real borrowing rate in our data, which is 1.3 percent, implying undervaluation of fuel economy improvements.

Our conclusion that the implicit discount rate exceeds market borrowing rates contrasts with Busse et al. (2013), who estimate implicit discount rates that are roughly equal to market borrowing rates. The second column in Table 4 shows that this difference does not arise from the fact that our baseline estimate is based on differing assumptions on vehicle miles traveled and survival probability. Using their assumptions yields a similar implicit discount rate to our baseline.

Another possible explanation for the difference between our results and theirs is that they identify consumer valuation from fuel cost variation induced by fuel price variation. If the consumer response to fuel price induced changes in fuel costs differs from the response to fuel economy induced changes in fuel costs, this could explain the discrepancy between our results and theirs.

However, our replication of their methodology using our data suggests otherwise (see Tables B.7 and B.8 for the estimation results). Table 5 shows that whereas Busse et al. (2013) report discount rates of -4.0 to 9.8 percent, using our data and their methodology we estimate higher discount rates of 2.1 to 25 percent (see Table B.6). Thus, we find consistent evidence of consumer undervaluation regardless of the estimation strategy or parameter assumptions.[18] This replication exercise also shows that differences between our functional form and theirs does not explain the differing results. Differences in the sample period could explain these results, if WTP depends on fuel prices (which were higher during our sample), on macroeconomic conditions (our sample includes the recovery from the 2008 to 2009 recession), or on other factors that differed between the two sample periods.[19]

4.3 Addressing potential sources of bias

As discussed in Section 3, the IV strategy would yield biased estimates if time-varying vehicle quality is correlated with the technology instruments, after controlling for average quality of each vehicle model-variant. This subsection provides evidence supporting the validity of the IV estimates.

If the instruments are correlated with quality, we would expect that the fuel economy and performance estimates would change if we add variables that are likely to be correlated with quality. We address this possibility in two ways, first by including variables that may directly measure vehicle quality, and second by including variables that may be indirectly correlated with quality. We begin by collecting variables from Chrome that are typically not included in vehicle demand models, and which may therefore reflect quality that is unobserved in these other studies. Specifically, in column 2 of Table 6 we add controls for the number of passengers, cubic feet of passenger volume, cubit feet of cargo volume, and a dummy for a moonroof or a sunroof. These variables are not observed for some of the observations in our data, which reduces the sample size. The coefficient estimates in the transaction price equation remain similar to the benchmark specification, while the fuel cost coefficient in the new registrations equation increases in magnitude. As a result, the implied willingness to pay for fuel economy is higher (as shown in Panel C of Table 2), suggesting a valuation ratio of 0.77. Although this ratio is higher than in the benchmark, the conclusion holds that consumers undervalue fuel economy. Moreover, the welfare conclusions in the next section are the same if we use these estimate rather than the baseline.

As an alternative measure of quality, we include consumer experience ratings reported in the MaritzCX survey. Respondents report ratings for a number of vehicle attributes, such as the vehicle's appearance and the quality of the sound system. We include 10 of these

[18]Given that Busse et al. (2013) rely on variation in fuel costs across models, the fact that we get similar results with both methods is consistent with consumers valuing fuel cost savings within model-variants in the same way that they value fuel cost savings across models or segments.

[19]Another commonly used measure of consumer valuation of fuel economy is the payback period. We follow the definition of payback period that EPA and NHTSA use, and compute the number of years from the time of purchase until the discounted stream of fuel savings equals the estimated WTP for a 1 percent fuel economy increase. Under our baseline assumption, the payback period for 1 percent fuel economy increase is 7 years. Under assumptions used in Busse et al. (2013), the payback period is 9 years.

attributes as covariates in column (3).[20] Although these measures are subjective, they are likely to be correlated with the consumer's perceived quality, and hence the transaction price. Identification rests on the assumption that the instrumented fuel economy and performance are uncorrelated with these quality measures, which is confirmed in column 3.

Recall that manufacturers typically make major redesigns of individual vehicles every 5-7 years; each redesign results in a new "generation" of the model. During a redesign, manufacturers are more likely to make major changes to the vehicle that could affect quality, compared to changes that are typically made between redesigns. This market regularity suggests that quality variation across generations may be more strongly correlated with the instruments, than quality variation within generations. If this is the case, interacting model-variant fixed effects with model generation fixed effects would cause WTP estimates to differ from the baseline. Columns 4 and 5 show that this is not the case. In column 4, we interact model-variant fixed effects with model generation fixed effects, and in column 5, we interact model-variant fixed effects with an indicator that equals to one if the model year represents a new generation. In each of these specifications, the implied valuations for fuel economy and performance (shown in Panel C of Table 2) are similar to those found in the benchmark model.

Next we turning to indirect proxies for quality in Table 7. First, we consider the example that vehicles may have (unobserved) automated safety features, such as blind spot detection. If manufacturers add automated safety features at the same time as adopting fuel-saving technology, quality would be correlated with the instruments. However, in this case quality would also be correlated with income and household size, as one expects households that have higher income or that include children to have higher WTP for automated safety features. Based on this reasoning, we add to the baseline IV specification of equations (2) and (3) six demographic controls: respondent's age, household size, male indicator, urban indicator, fixed effects for the respondent's education group (12 groups), and fixed effects for 23 household income groups. Note that the sample is smaller than the baseline because of missing demographics data. Column 2 of Table 7 reports the coefficient estimates when including these controls (column 1 repeats the baseline estimates for convenience), with Panel A reporting price regressions and Panel B reporting quantity regressions. The estimates are similar to the baseline. We estimate equations (2) and (3) with additional demographic controls in column 3, including the number of wage earners, number of children, an indicator equal to one if the respondent's spouse is employed, fixed effects for the respondent's race (6 categories), and fixed effects for the respondent's

[20] Each of these attributes is measured on a scale of 1 to 5. We represent these ratings as continuous variables.

occupation (20 categories). The additional demographics further reduce the sample size, but the coefficient estimates are similar to the baseline.

Quality may also vary geographically over time. Returning to the safety example, we may expect residents of the Northeast to have higher WTP for safety features because of the poor weather conditions in that region. The state fixed effects control for the average probability that the vehicles contain these features, but preferences or costs of the features may vary over time. If preference or cost changes are correlated with technology adoption, the IV estimates would be biased. In column 4 we include richer time fixed effects by interacting state fixed effects with model-year fixed effects, and interacting state fixed effects with month-of-year fixed effects. The coefficient estimates are similar to the baseline.

Above, we noted that there have been a few negative reports of consumer experiences with continuously variable transmissions and cylinder deactivation, particularly when these technologies first entered the market. If consumers value (either negatively or positively) these technologies for reasons other than their effects on fuel economy and performance, the IV estimates would be biased because the instruments would be correlated with quality. Column 5 shows that omitting these variables as instruments does not affect the point estimates, reducing such concerns.

If households face borrowing constraints, changes in financial market conditions could affect borrowing costs and the composition of households that choose to purchase a new vehicle. If WTP varies across households and the variation is correlated with borrowing costs, the WTP estimates could be biased. However, column 6 shows that controlling for financing arrangement and payment type does not affect the results, reducing this concern. Likewise, column 7 shows that the results are similar if we omit observations from 2009, when borrowing rates were relatively high following the economic recession.

As a final validation of the IV strategy, we report the reduced-form relationship between transaction prices and the fuel-saving technology instruments. Because the technologies can increase both fuel economy and performance, we expect a positive and monotonic relationship between a vehicle's price and the number of technologies it contains. In contrast, although we expect a positive correlation between the number of technologies and quality, the relationship between the number of technologies and quality is not necessarily monotonic. Therefore, if quality is correlated with the instruments, we may observe a non monotonic relationship between the number of fuel-saving technologies in a vehicle and its transaction price. We compute the number of technologies for each vehicle in the sample (we top-code the count at five because few observations contain more than five technologies). We regress the log of the transaction price on the same independent variables as in the baseline specification of equation (2), as well as fixed effects for the number of fuel-saving technologies. The top panel of Figure 7 plots the coefficients and 95 percent confidence intervals. The figure illustrates a positive and monotonic relationship between the transaction price and the technology count.

We estimate a second reduced-form regression of the transaction price on indicator variables for each technology. If the instruments are valid, each technology should increase the transaction price. However, if quality is positively correlated with some instruments and negatively correlated with others, we could observe negative correlations among transaction price and the latter technologies. The bottom panel of Figure 7 reports the estimated coefficients and confidence intervals. All coefficients are positive and most are statistically significant at the 5 percent level. Overall, both sets of reduced-form regressions support the IV strategy.

5 Implications

In this section we discuss the implications of our estimates for the effects of fuel economy and greenhouse gas emissions standards on consumer welfare. The approach is to consider small hypothetical changes in fuel economy and emissions standards, and to use the empirical estimates to infer the consumer welfare implications. For simplicity and consistency with EPA and NHTSA benefit-cost analysis, we focus on a representative consumer and assume that markets are imperfectly competitive with free entry and exit. Manufacturers pass to consumers cost changes, and profits are unaffected in these examples.

5.1 Comparing consumer valuation of fuel economy and performance

In this subsection we compare the magnitudes of the WTP for fuel economy and performance. Manufacturers can use fuel-saving technology, such as variable valve lift and timing, to increase fuel economy or performance. Historically, during periods of time in which the stringency of fuel economy standards was not changing, manufacturers have adopted fuel-saving technology and retuned engines to improve performance while maintaining fuel economy. Between 1990 and 2005 the standards did not change, and the market-wide average fuel economy was unchanged while the ratio of horsepower to weight increased by 33 percent (Klier and Linn 2012). We showed in Table 2 that when light truck standards began to tighten in 2005, the rate of horsepower improvements slowed while fuel economy began increasing. For cars, standards began to tighten in 2011, and we observe the same shift from horsepower to fuel economy improvements. Klier and Linn (2016) show that the tightening standards caused a shift to improving fuel economy and a shift away from improving other vehicle attributes. Because manufacturers typically use fuel-saving technology to raise performance when fuel economy standards are not tightening, these patterns suggest that consumers value performance more than fuel economy.

To assess whether our WTP estimates are consistent with these patterns, we combine the estimates with the estimated technological trade-off between fuel economy and performance from the literature. Our WTP estimates suggest that consumers would pay $133 for 1 percent fuel economy increase. Alternatively, suppose a manufacturer uses the same fuel-saving technology that would raise fuel economy by 1 percent, and increases performance rather than fuel economy. Knittel (2011) and Klier and Linn (2016) estimate technical trade-offs among fuel economy, horsepower, and other attributes. These estimates imply that, holding weight and marginal costs constant, rather than increasing fuel economy by 1 percent the manufacturer could increase performance by 3 to 6 percent (depending on market segment and the estimates from the two previous articles). Our WTP estimates suggest that consumers would pay about $394 for the performance increase, far exceeding the value of the fuel economy increase. Consumers would value vehicles more if automakers use fuel-saving technology to raise performance rather than fuel economy. Our estimates are therefore consistent with historical patterns of manufacturer attribute choices. The results suggest that the ratio of the marginal WTP for performance, relative to the marginal WTP for fuel economy, is about 0.7, which exceeds the technological trade-offs between the two attributes, which ranges from 0.17 to 0.33. This suggests that the passenger vehicle market is at a corner solution in the performance - fuel economy space, such that when fuel economy standards are unchanging over time, manufacturers use fuel-saving technology to increase performance and leave fuel economy unchanged.

5.2 How do tighter standards affect private consumer welfare?

In this subsection, we use our WTP estimates to assess the effect on private consumer welfare of tightening standards. Klier and Linn (2016) show that tighter standards cause manufacturers to adopt fuel-saving technology more quickly than they would have if standards had not tightened. The additional technology adoption raises fuel economy as well as vehicle production costs and vehicle prices. Klier and Linn (2016) show that the tighter standards cause manufacturers to trade off performance for fuel economy, despite the fact that consumers appear to have a high WTP for performance. This trade-off implies causes performance to be lower than if standards had not tightened.

Undervaluation implies that a marginal increase in the stringency of fuel economy standards raises private consumer welfare if two conditions hold. The first condition is that manufacturers equate the technological trade-off between fuel economy and performance to the ratio of the marginal WTP for performance to the marginal WTP fuel economy. The second condition is that manufacturers choose levels of performance for each vehicle such that the technology cost of increasing performance equals the marginal WTP for performance.[21] If either condition does not hold, tighter standards would reduce performance, which would cost consumers more than the benefit of higher fuel economy. The previous subsection suggests that the first condition does not hold, and in this subsection we show that the second condition does not hold, either.

To estimate the effects of tighter standards on private consumer welfare, this paper focuses on providing reliable estimates of consumer valuation. For technological trade-offs, as in the previous subsection we use the estimates from Klier and Linn (2016). We use technology adoption cost estimates from EPA (2012) and Leard et al. (2016).

For consistency with the marginal WTP estimates, we focus on the changes in vehicle attributes and prices caused by a 1 percent tightening of the standards in a single year. The estimates in Klier and Linn (2016) imply that, in response to a 1 percent fuel economy tightening, manufacturers adopted technology that increased vehicle efficiency and fuel economy by 0.12 percentage points more than they would have if the standards had not been tightened. Manufacturers trade off performance for fuel economy to attain the remaining 0.88 percentage points. Therefore, the total cost of the 1 percent fuel economy increase includes the cost of adopting the fuel-saving technology, as well as the welfare cost of the lower performance (i.e., relative to the counterfactual in which performance increases due to fuel-saving technology adoption). We compare these costs with the present discounted value of the fuel savings.

In Section 4.2, we reported that this fuel economy increase yields a present discounted value of fuel savings of $249. Based on technology cost estimates in EPA (2012), Leard et al. (2016) estimate that increasing fuel economy by 0.12 percent, while holding other attributes constant, raises costs by $11 per vehicle (this estimate includes the increase in marginal costs as well as average fixed costs).[22] Using the same assumptions as in the last subsection, the welfare cost of reducing performance to increase fuel economy by 0.88 percent is $347. Therefore, the tighter standards reduce private consumer welfare by $109 per vehicle, or 0.4 percent of the average transaction price in the sample. The negative estimate is robust to statistical uncertainty in Klier and Linn (2016) regarding the additional efficiency improvement; we have redone the calculations using the 95 percent confidence intervals from Klier and Linn (2016), which yields changes of private consumer welfare of -0.3 to -0.5

[21] This can be shown using the model in (Klier and Linn 2012) and applying the envelope theorem to a marginal tightening of the standards.

[22] Implicit in our analysis is the assumption that manufacturers comply with tighter fuel economy standards by adopting technology. In practice, they may also reduce the relative prices of vehicles with low fuel economy (Goldberg 1998), which would reduce the cost relative to our estimate. However, Klier and Linn (2012) suggest that this effect would be small in magnitude.

percent. Efficiency improvements would have to account for at least half of the fuel economy improvement for tighter standards to increase private consumer welfare.

We make two observations about this result. The first is that the estimate is much different from the estimate one would obtain by ignoring the costs of forgone performance. Contrary to recent evidence in the literature, in their benefit-cost analysis of the standards EPA and NHTSA assume that tighter standards do not cause manufacturers to trade off performance for fuel economy. Instead, to meet the 1 percent fuel economy increase required in this example, manufactures adopt sufficient fuel-saving technology to increase fuel economy by 1 percent. Using the same technology cost assumptions as in the preceding calculation, tighter standards raise vehicle prices by $91 per vehicle. Accounting for the value of the fuel savings, tightening standards by 1 percent would increase private consumer welfare by $158 per vehicle, or about 0.6 percent of the average transaction price.

Second, the consumer welfare effects depend on the effect of the standards on the rate of technology adoption. The more that standards increase this rate, the less manufacturers trade off performance for fuel economy, causing the standards to have less of a negative effect on consumer welfare. Our estimate of -$109 per vehicle is based on the estimated effect of standards on technology adoption from the post-2010 time period. Estimates from Klier and Linn (2016) for earlier periods indicate larger technology adoption effects of tighter standards. Those estimates imply that tightening standards by 1 percent changes consumer welfare by -$25 per vehicle, or 0.1 percent of average transaction price. The calculations imply negative consumer welfare effects and indicate some of the uncertainty around the point estimate of -$109. Overall, we conclude that tighter standards are unlikely to substantially improve consumer welfare, and our central estimate is that tighter standards have approximately zero effect.

These conclusions are subject to several caveats. The technology cost estimates are based on interpolations described in Leard et al. 2016. The reduction in consumer welfare refers to the private welfare of new vehicle consumers; it does not include the social benefits arising from improved energy security or climate—that is, the current standards may increase social welfare, even if standards do not noticeably increase private consumer welfare. Moreover, this conclusion does not account for potential induced innovation caused by tighter standards, market failures associated with insufficient market incentives for innovation (e.g., Fischer 2010; Porter and van der Linde 1995), market failures associated with imperfect competition (such as the possible underprovision of fuel economy), and interactions between the new and used vehicle markets (Jacobsen and van Benthem 2015). Finally, the conclusion does not account for transitional dynamics. Klier and Linn (2016) find that tighter standards increase the rate of technology adoption, implying that standards may trade off higher fuel economy in the near term for lower performance in the long term. Accounting for these effects would require a dynamic analysis of new vehicle standards, which remains for future research.

5.3 Tighter standards and consumer acceptance

A contentious issue regarding the fuel economy and greenhouse gas emissions standards is whether the standards reduce overall consumer demand for new vehicles. If the standards reduce demand, tighter standards could cause some consumers to forgo obtaining a new vehicle and instead obtain a used vehicle or continue using their existing vehicles longer than they would have. Lower demand would reduce the total number of new vehicles that manufacturers sell and their profits. In addition, lower demand would decrease the rate at which lower-emitting new vehicles replace higher-emitting existing vehicles, reducing equilibrium social welfare benefits of the standards.

We estimate the effects of tighter standards on consumer demand for a typical new vehicle—i.e., the marginal change in consumer surplus for the new vehicle—accounting for changes in vehicle prices, fuel economy, and performance. These calculations are identical to those used in the previous section, except that we use the WTP for fuel economy to value the fuel economy increase, rather than the discounted value of the fuel cost savings. This change is appropriate because consumers choose vehicles based on WTP rather than the discounted value of fuel savings. This measure is relevant to the effects of standards on consumer acceptance of new vehicles and aggregate vehicle demand.

Our estimates suggest that tighter standards reduce consumer demand in the short run. Specifically, tightening standards by 1 percent in our sample causes fuel economy to increase by the same amount, which increases WTP by $133. However, the same tightening of the standards raises vehicle prices by $11 and reduces WTP for performance by $347. Overall, consumer WTP for new vehicles, net of vehicle price, fuel economy, and performance changes, decreases by $227 per vehicle, or 0.8 percent of the average transaction price.

The result carries the same caveats as in the previous subsection. We leave for future work quantifying the welfare implications of this effect of fuel economy standards on total sales.

6 Conclusion

If an energy efficiency gap exists for passenger vehicles, new vehicle fuel economy or greenhouse gas emissions standards would increase private welfare of new vehicle consumers and producers. NHTSA and EPA argue that a gap exists and conclude that the benefits of the fuel savings from existing standards exceed the costs of achieving the standards; these benefits account for about 70 percent of the total benefits of the standards.

To draw welfare implications for standards, the literature has assessed whether there is an energy efficiency gap by asking whether consumers undervalue fuel economy. However, we argue that the literature has focused narrowly on consumer valuation of fuel economy and has not considered the welfare costs of forgone performance increases. Manufactures can use those fuel-saving technologies to increase either fuel economy or performance. There are certain fuel-saving technologies that manufacturers adopt regardless of whether standards tighten. If manufacturers use those technologies to increase performance if standards do not tighten, and if tighter standards cause manufacturers to use those technologies to increase fuel economy instead of performance, manufacturers forgo the opportunity to increase performance. The forgone performance reduces consumer welfare, opposing the positive consumer welfare effect of fuel savings caused by standards. As we explain, under certain conditions tighter standards could reduce private consumer welfare even in the presence of undervaluation.

We use a unique data set and novel identification strategy to estimate consumer valuation of fuel economy and performance. Consumers are willing to pay about 54 cents for $1 of discounted future fuel savings. This estimate is smaller than Busse et al. (2013) and Allcott and Wozny (2014), which likely reflects differences in sample period rather than methodology. The performance estimates imply that consumers pay about $94 for a 1 percent performance increase, which corresponds to $1,100 for a 1-second reduction in 0-to-60 time.

The estimated undervaluation of fuel economy would seem to suggest that tighter standards increase private consumer welfare. However, the estimated consumer valuation of performance is sufficiently large that the entire welfare cost of increasing fuel economy, including costs of adopting technology and reducing performance, approximately equals the value of the fuel savings. This conclusion is subject to the caveats we discuss in Section 5.2, and we note that standards may increase social welfare after accounting for the energy security and climate benefits.

Our WTP estimates suggest two puzzles related to technology adoption costs. First, the estimated WTP for a 1 percent performance increase ($394) exceeds the cost of adopting fuel-saving technology and increasing performance ($89), suggesting that manufacturers should adopt fuel-saving technology more quickly than they do. Second, the WTP for performance implies that manufacturers would avoid trading off performance for fuel economy because consumers value the performance so highly. Yet, the patterns in Table 2 as well as estimates in Klier and Linn (2016) suggest that manufacturers do make this trade-off when facing tighter standards. Future research can investigate whether hidden costs, consumer preference heterogeneity, or other factors explain these apparent puzzles.

Although fuel economy standards may not increase consumer welfare, other policies could improve consumer welfare by targeting the cause of the undervaluation. For example, if consumers lack information about fuel cost savings, and the lack of information causes them to undervalue savings, then improving information could increase consumer welfare. Future research could attempt to determine the cause of undervaluation and identify appropriate policies to correct market failures.

The results have implications for the effects of fuel economy and emissions standards on demand for new vehicles. Our estimates imply that tightening standards by 1 percent reduces consumer valuation by 0.8 percent per vehicle, although we suggest that these results should not be extrapolated far out of sample because they are based on marginal WTP. Future work could incorporate these effects in a comprehensive welfare analysis of the standards.

References

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

Anderson, S. T., R. Kellogg, and J. M. Sallee (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management 66*(3), 383–403.

Anderson, S. T. and J. M. Sallee (2011, June). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review 101*(4), 1375–1409.

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

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.

Copeland, A. (2014). Intertemporal substitution and new car purchases. *The RAND Journal of Economics 45*(3), 624–644.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under Uncertainty.* Princeton University Press.

EPA (2012). Regulatory impact analysis: Final rulemaking for 2017–2015 light-duty vehicle greenhouse gas emission standard and corporate average fuel economy standards. (EPA-420-R-12-016 August 2012).

EPA (2014). Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2014. (EPA-420-R-14-023).

EPA, NHTSA, and CARB (2016). Draft technical assessment report: Midterm evaluation of light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards for model years 2022–2025. (EPA-420-D-16-900 July 2016).

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

Gillingham, K., R. G. Newell, and K. Palmer (2009). Energy efficiency economics and policy. *American Review of Resource Economics 1*, 597–620.

Goldberg, P. K. (1998, March). The effects of the corporate average fuel efficiency standards in the US. *Journal of Industrial Economics 46*(1), 1–33.

Greene, D. L. (2010). How consumers value fuel economy: A literature review. (EPA-420-R-10-008 March 2010).

Greene, D. L., A. Hossain, and R. Beach (2016). Consumer willingness to pay for vehicle attributes: What is the current state of knowledge? (RTI Work Assignment 4-11, RTI Project Number 0213244.004.011).

Gruenspecht, H. K. (1982). Differentiated regulation: The case of auto emissions standards. *The American Economic Review 72*(2), 328–331.

Helfand, G. and A. Wolverton (2009). Evaluting the consumer response to fuel economy: A review of the literature. (EPA Working Paper No. 09-04).

Jacobsen, M. R. (2013, May). Evaluating US fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy 5*(2), 148–87.

Jacobsen, M. R. and A. A. van Benthem (2015, March). Vehicle scrappage and gasoline policy. *American Economic Review 105*(3), 1312–38.

Jaffe, A. B. and R. N. Stavins (1994, May). The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics 16*, 91–122. A-13.

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 133*, 41–63.

Klier, T., J. Linn, and Y. C. Zhou (2017). The effects of fuel costs and market size on fuel-saving technology adoption: Direct and indirect effects. (RFF Discussion Paper No. 16-26-REV).

Knittel, C. (2011). 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 V. McConnell. Fuel prices, new vehicle fuel economy, and implications for attribute-based standards. *Journal of the Association of Environmental and Resource Economists,forthcoming*.

Leard, B., J. Linn, and V. McConnell (2016). How do low gas prices affect costs and benefits of US new vehicle fuel economy standards? (RFF Policy Brief No. 16-12).

Linn, J. (2016). Interactions between climate and local air pollution policies. (RFF Discussion Paper No. 16-51).

Lu, S. (2006). Vehicle survivability and travel mileage schedules. (Technical Report Working Paper DOT HS 809 952).

Metcalf, G. E. and K. A. Hassett (1993). Energy conservation investment: Do consumers discount the future correctly? *Energy Policy 21*(6), 710–716.

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

Parry, I. W. H., M. Walls, and W. Harrington (2007, June). Automobile externalities and policies. *Journal of Economic Literature 45*(2), 373–399.

Porter, M. E. and C. van der Linde (1995, December). Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives 9*(4), 97–118.

Reynaert, M. (2015). Abatement strategies and the cost of environmental regulations: Emission standards on the European car market. *Working Paper*.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy 82*(1), 34–55.

Sallee, J. M., S. E. West, and W. Fan (2016). Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics 135*, 61–73.

Stavins, R. N. (2005). The effects of vintage-differentiated environmental regulation. *Stanford Environmental Law Journal 25*(1).

Whitefoot, K. S., M. Fowlie, and S. J. Skerlos (2013). Compliance by design: Industry response to energy efficiency standards. *Working Paper*.

Whitefoot, K. S. and S. J. Skerlos (2012). Design incentives to increase vehicle size created from the u.s. footprint-based fuel economy standards. *Energy Policy 41*, 402–411. Modeling Transport (Energy) Demand and Policies.

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

Figures

Figure 1: **Fuel Economy, Weight, Horsepower, and Torque by Model Year, 2010–2014**

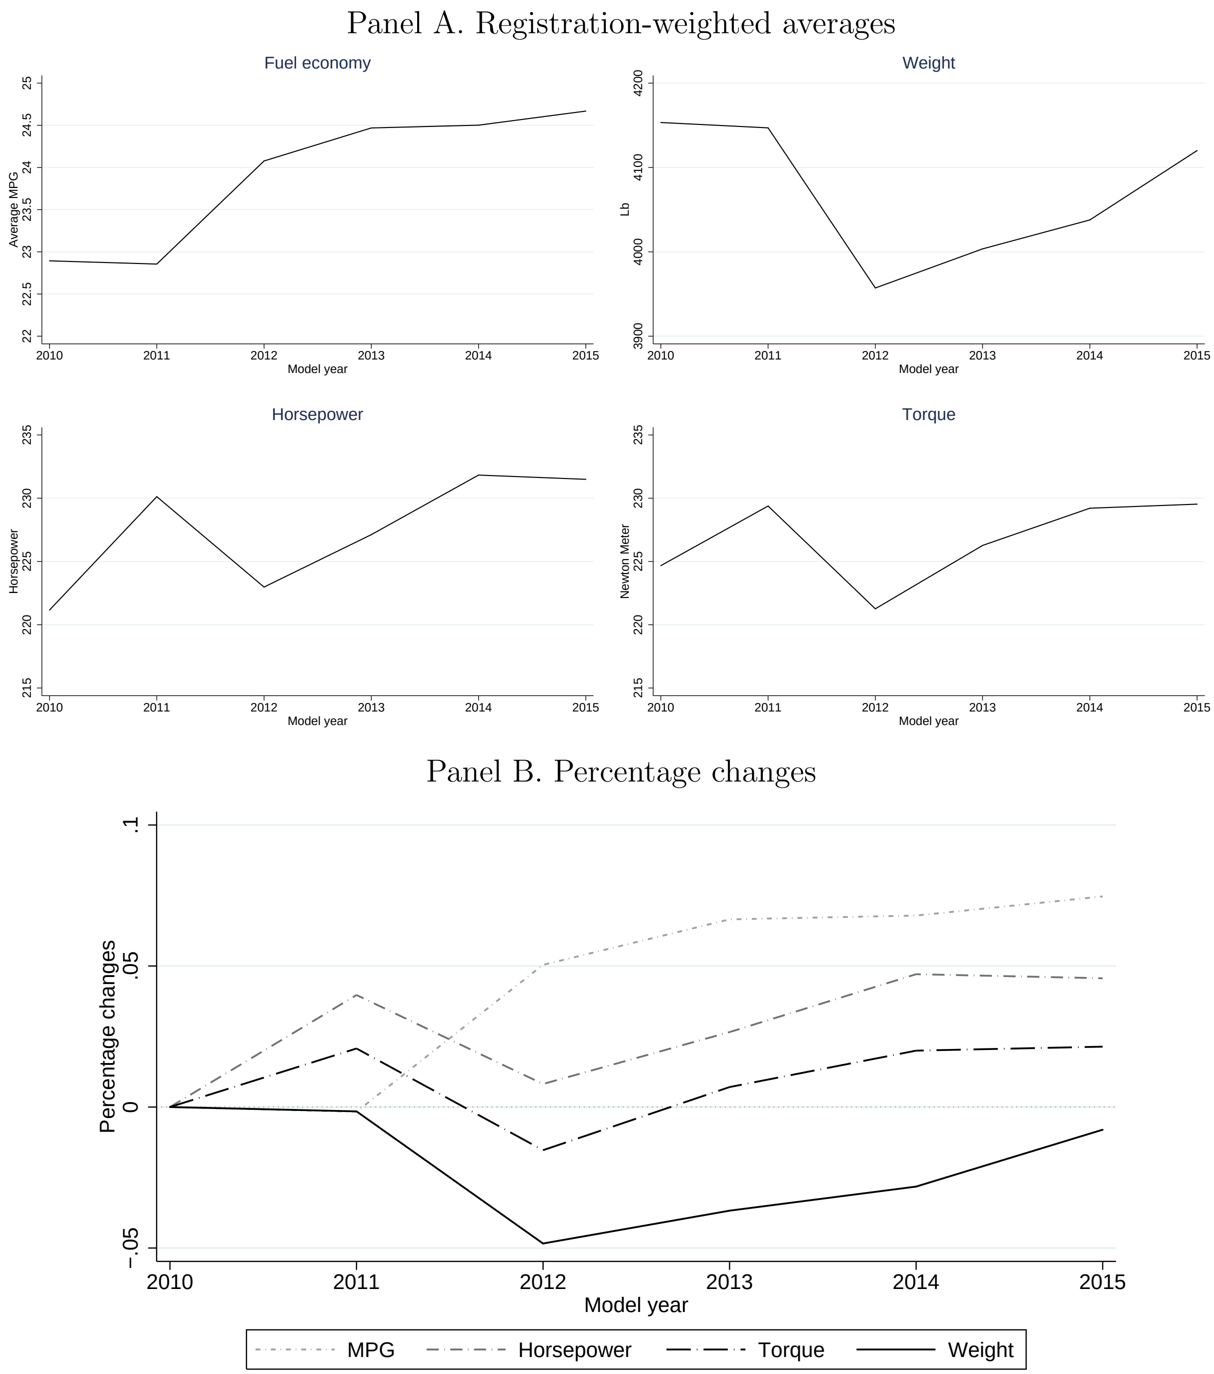

Notes: Panel A reports registration weighted average fuel economy, weight (in pounds, lb), horsepower, and torque (newton meters, nm) by model year. Panel B reports percent changes in these variables since the 2010 model year.

Figure 2: **Engine and Transmission Variables by Model Year, 2010–2014**

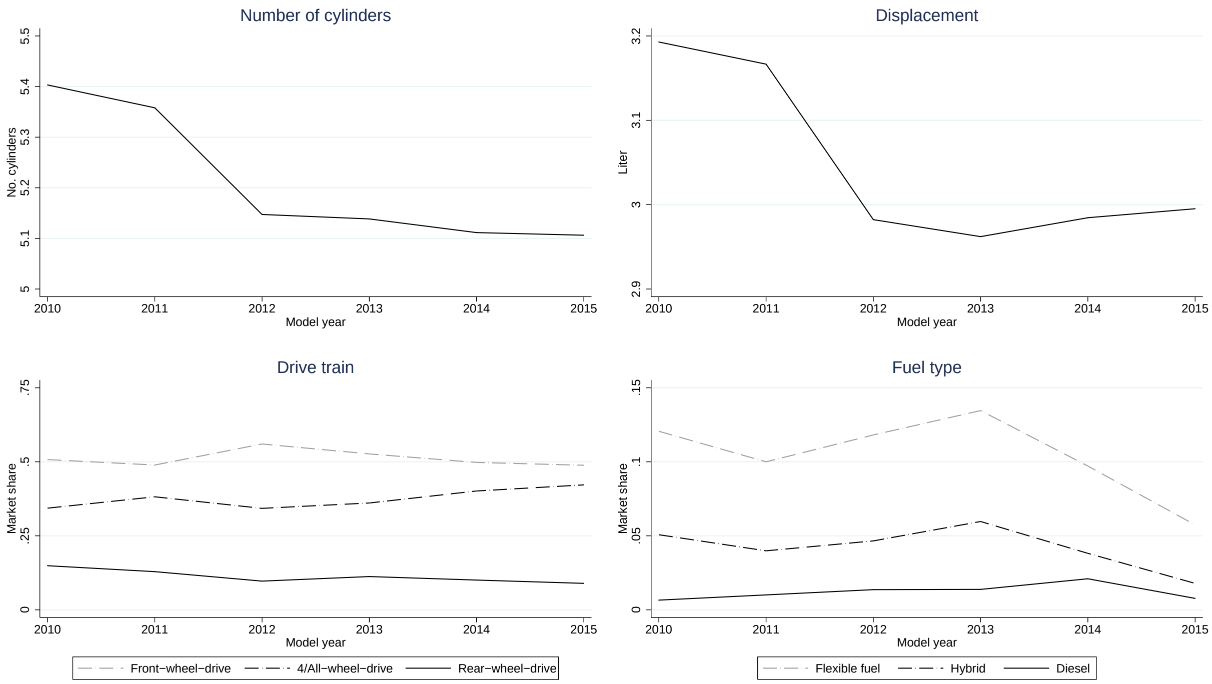

Note: The figure shows registration-weighted number of cylinders and engine displacement, as well as the market shares of drive train type and fuel type.

Figure 3: **Market Penetration of Selected Fuel-Saving Technologies, 2010–2014**

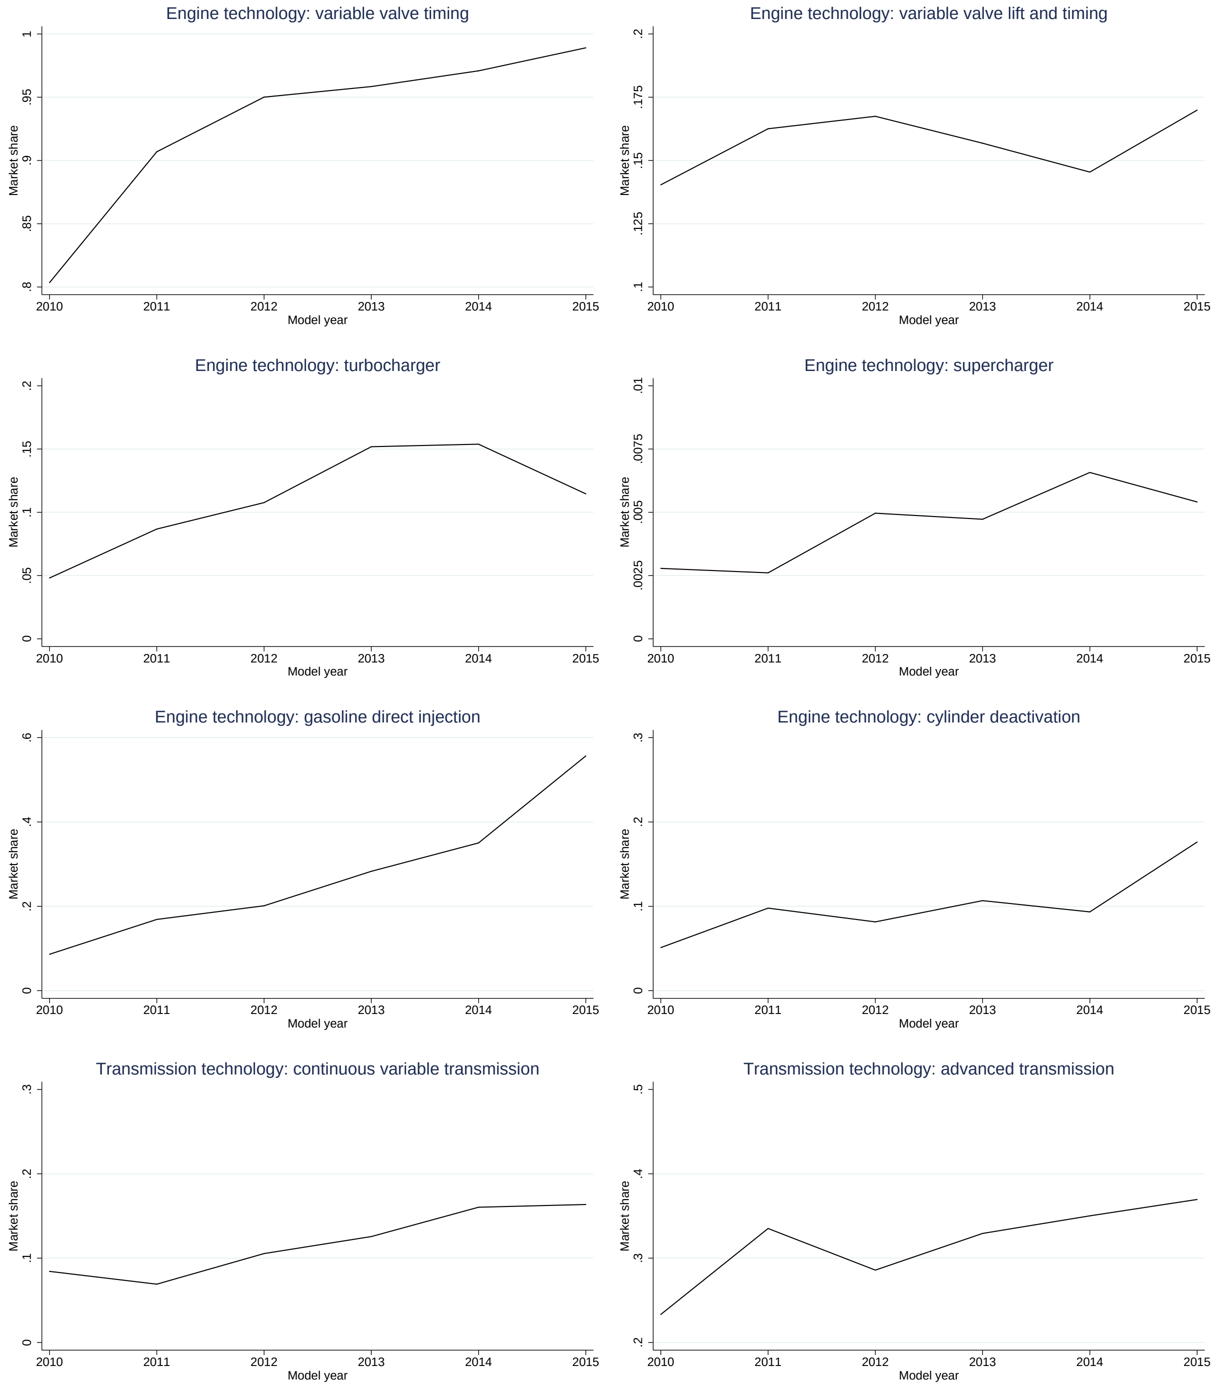

Note: The figure reports the the registration-weighted market shares of the engine and transmission variables used to construct the IVs.

Figure 4: **Monthly Fuel Prices, 2009–2014**

Panel A. National average monthly gasoline and diesel fuel prices

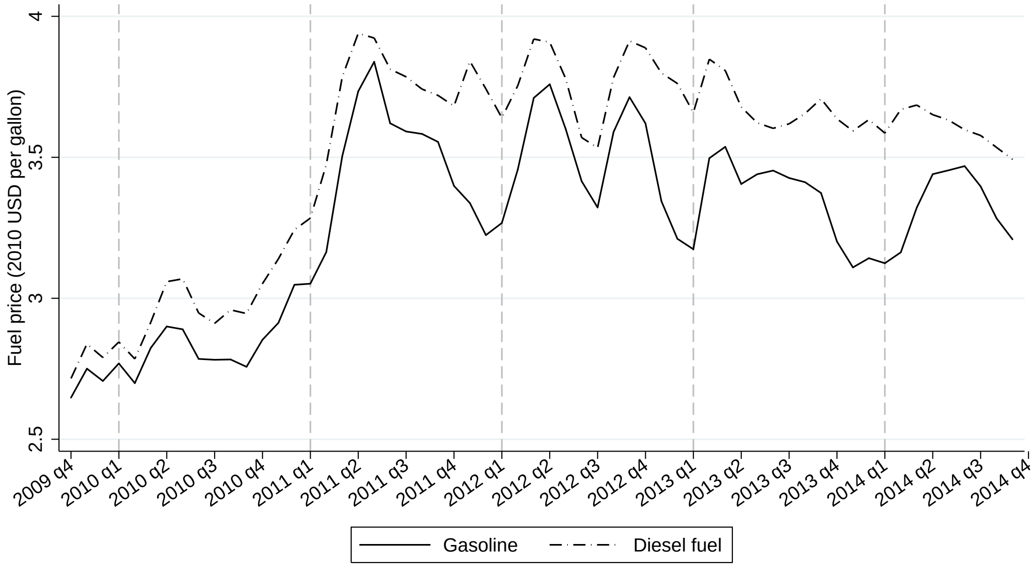

Panel B. Regional average monthly gasoline prices

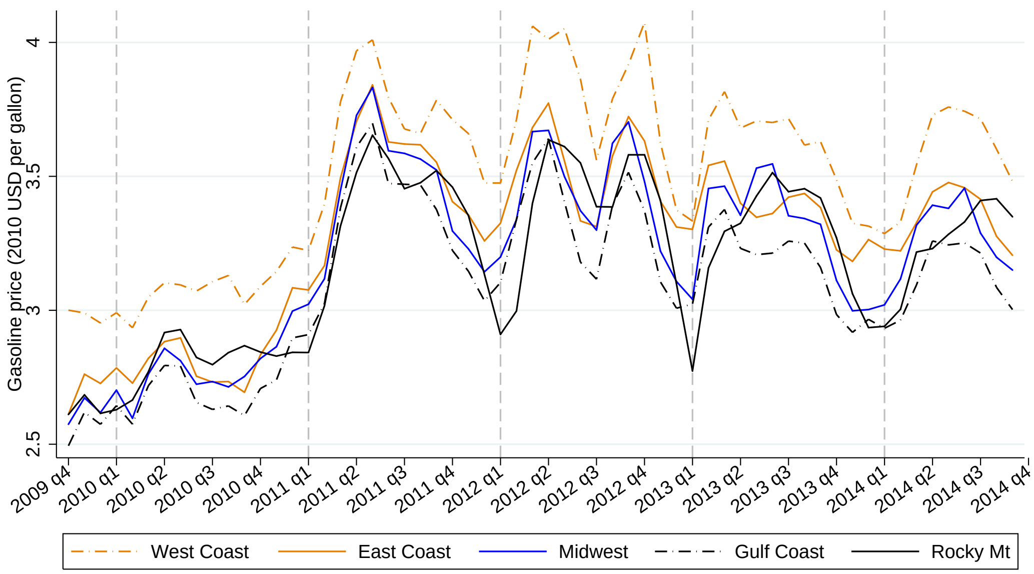

Notes: Panel A shows monthly average national gasoline and diesel fuel prices. Panel B shows monthly gasoline prices by petroleum administration for defense district. Dashed vertical lines indicate the beginning of calendar years.

Figure 5: **Vehicle Transaction Prices, 2009–2014**

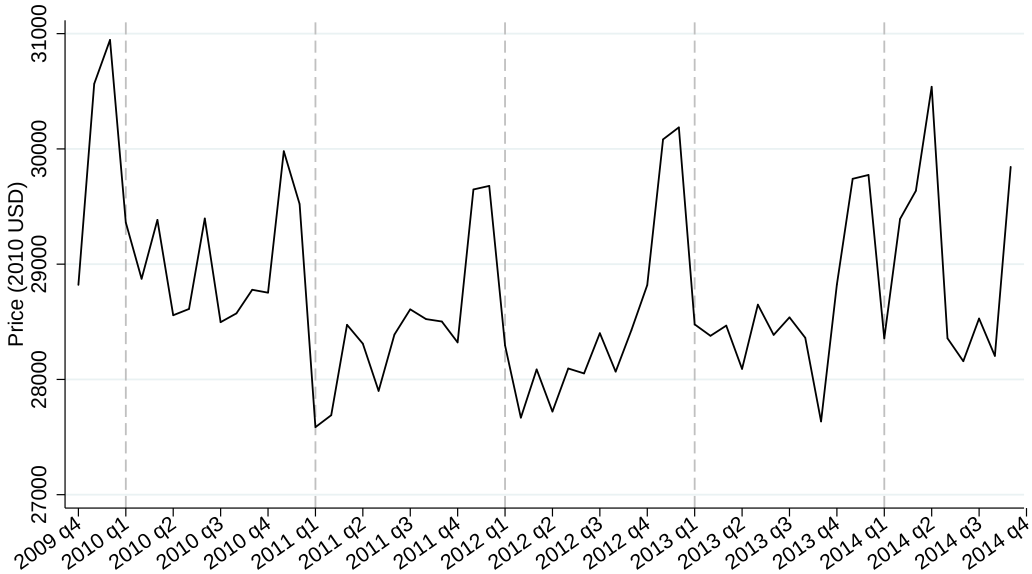

Notes: The figure shows the monthly registration-weighted average transaction prices, with dashed vertical lines indicating the beginning of calendar years.

Figure 6: **Effects of Fuel Economy Increase on Equilibrium Prices and Quantities**

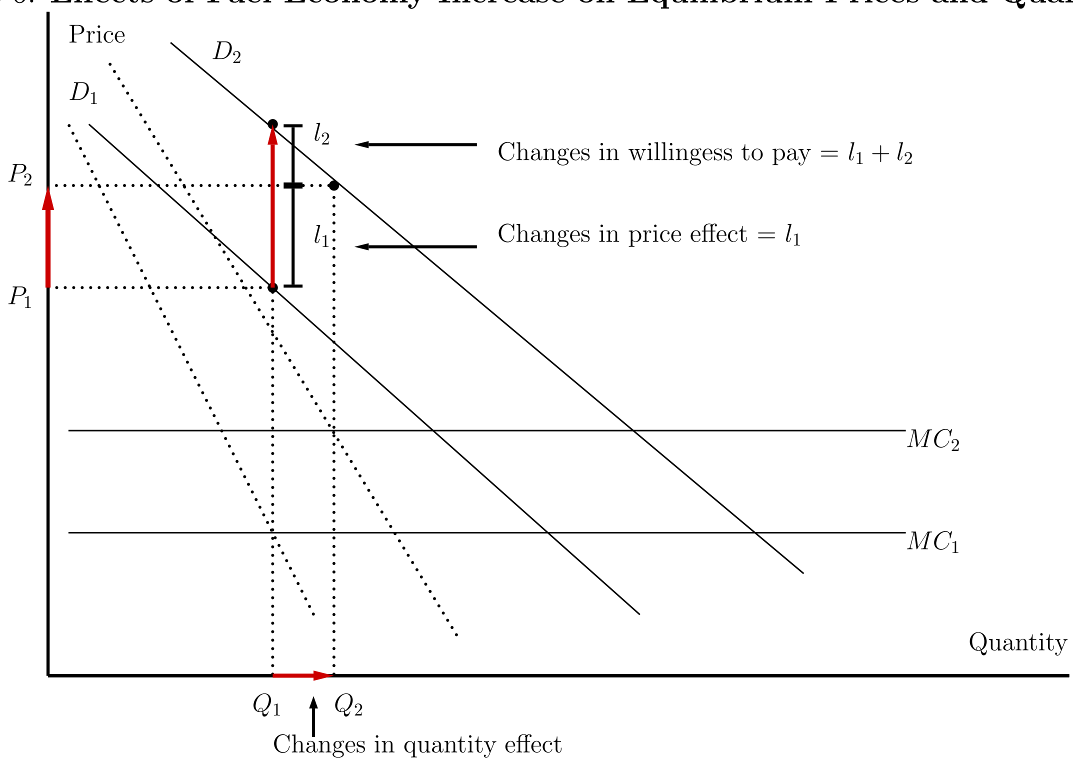

Figure 7: **Reduced-Form Relationships: Prices and Fuel-Saving Technologies**

Panel A. Number of fuel-saving technologies

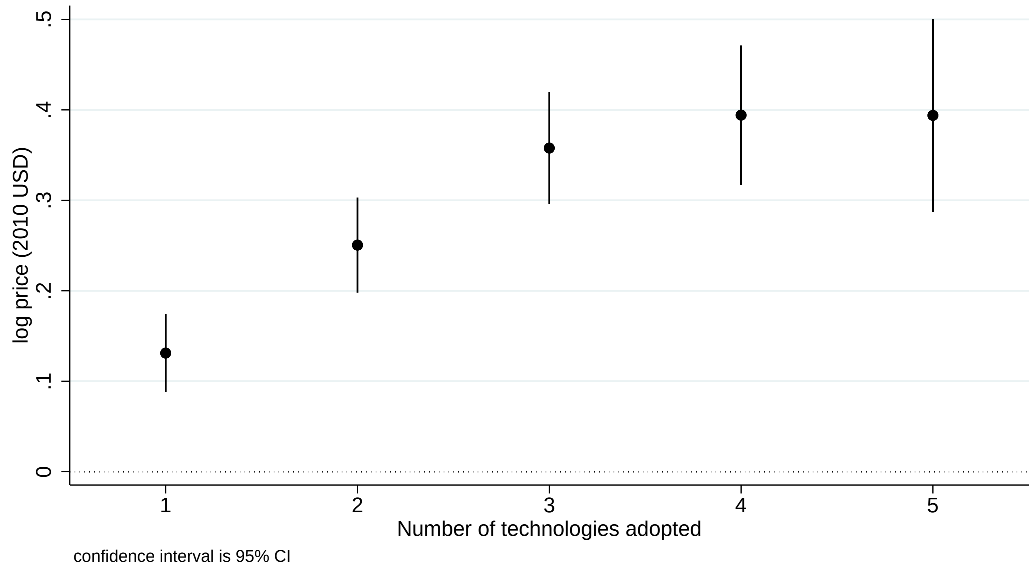

Panel B. Individual fuel-saving technologies

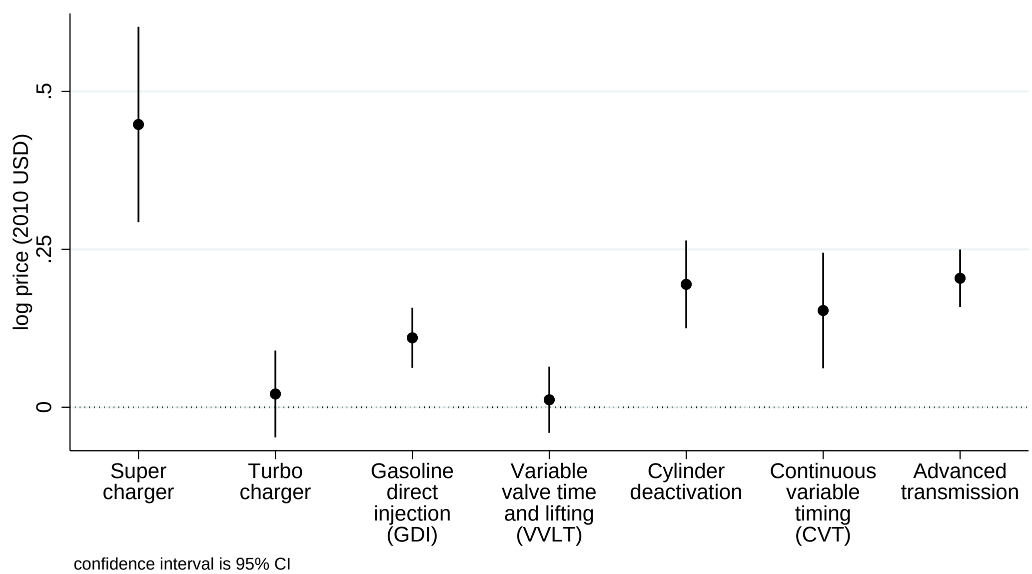

Notes: Panel A reports the coefficients on fixed effects for the number of fuel-saving technologies from a regression of log transaction price on the count fixed effects and the other independent variables from column 3 of Table 3. The number of technologies is top-coded at five because fewer than 1 percent of observations have more than five technologies. Panel B reports results from a similar regression, except that the count fixed effects are replaced by fixed effects for each technology. The vertical lines indicate 95 percent confidence intervals.

Tables

Table 1: **Summary Statistics**

	Mean	Std. dev.	Min.	Max.
Panel A. Price and vehicle characteristics				
Transaction price (2010 USD)	28,693	11,402	5,998	191,622
Fuel economy (miles/gallon)	23.9	6.6	12	50
Horsepower (hp)	226	78	70	662
Torque (newton meter, nm)	306	113	92	856
Weight (pounds, lb)	4,055	1,264	1,808	8,200
Engine displacement (liters)	3.0	1.2	1	8.4
Hybrid	0.05	0.21	0	1
Flex fuel	0.11	0.32	0	1
All-wheel/4-wheel-drive	0.37	0.48	0	1
Panel B. Demographics of respondent				
Household size	2.5	1.2	1	6
Age (years)	52.6	15.4	15	99
Male	0.61	0.49	0	1
Urban	0.55	0.50	0	1
Number of unique vehicle models				450
Number of unique vehicle trims				1,351
Number of unique vehicle model-variants				2,166
Number of observations				535,130

Notes: Panel A reports the registration-weighted average, standard deviation, minimum, and maximum of the variables indicated in the row headings. Engine displacement is the volume of the engine cylinders, in liters. Hybrid, and flex fuel are indicator variables for whether the vehicle has a hybrid power train, or is capable of using E85 fuel. All-wheel/4-wheel-drive is an indicator for whether the vehicle has all-wheel- or 4-wheel-drive. A unique model has a unique company name, manufacturer name, vehicle series name, and vehicle "nameplate" description. A unique trim is a unique model and a unique trim name. A unique model-variant is a trim with a unique combination of drive train specification (front-wheel-drive, rear-wheel-drive, or all/4-wheel-drive), fuel type (gasoline, diesel fuel, or other), displacement, and number of cylinders.

Table 2: **Annual Percent Growth of Vehicle Attributes by Time Period**

	Cars			Light trucks		
	Fuel economy	Horsepower	Weight	Fuel economy	Horsepower	Weight
1996–2000	-0.6	1.9	0.6	0.2	4.0	1.3
2001–2004	0.4	1.8	0.7	-0.6	4.7	3.2
2005–2011	0.2	1.2	0.4	1.0	1.0	-0.3
2012–2015	2.1	0.2	1.2	2.5	0.7	-0.9

Notes: The table reports annual percent growth rates for cars and light trucks by time period. The data are from Leard, Linn, and McConnell (forthcoming).

Table 3: **Willingness to Pay for Fuel Cost Savings and Performance**

	(1)	(2)	(3)
Estimated by	OLS	OLS	IV
Panel A. Dependent variable is log transaction price			
Log fuel cost (dollars/mile)	-0.113***	-0.156***	-0.354***
	(0.018)	(0.020)	(0.075)
Log performance (hp/lb or nm/lb)	0.068***	-0.230***	0.203***
	(0.014)	(0.020)	(0.074)
Model-variant fixed effect		Yes	Yes
Number of observations	457,525	535,124	535,124
RMSE	0.13	0.13	0.13
F-stat (fuel cost or fuel economy, excl var)	-	-	185.5
F-stat (performance, excl var)	-	-	243.4
Panel B. Dependent variable is log new registrations			
Log fuel cost (dollars/mile)	-1.651***	-0.636***	-0.338***
	(0.119)	(0.045)	(0.116)
Log performance (hp/lb or nm/lb)	-0.578***	-0.030	0.371***
	(0.061)	(0.028)	(0.083)
Model-variant fixed effect		Yes	Yes
Number of observations	457,525	535,124	535,124
RMSE	0.6	0.39	0.39
F-stat (fuel cost or fuel economy, excl var)	-	-	112.1
F-stat (performance, excl var)	-	-	150.1
Panel C. Willingness to pay (2010 USD)			
For 1 percent increases in			
• fuel economy	190.9	105.6	133.4
• performance	74.7	68.7	93.6

* $p<0.10$ ** $p<0.05$ *** $p<0.01$

Notes: Robust standard errors in parentheses, clustered by vehicle model-by-state. Performance for cars is the ratio of horsepower to weight and performance for trucks is the ratio of torque to weight. All specifications include as independent variables fixed effects for number of transmission speeds and a dummy variable for flex fuel capability, as well as the interactions of these variables with a dummy variable for light trucks. All specifications include fixed effects for state, model year, and PADD region-month-fuel type, as well as a lease dummy and a CAFE stringency variable interacted with model-year fixed effects (see text for details). In all price regressions, observations are weighted by the number of registrations, and all quantity regressions are not weighted. In column 1, regressions include trim fixed effects, displacement, weight, length, width, height, fuel tank volume, maximum number of passengers, and wheelbase. Column 1 in Panel A includes the number of cylinders and fixed effects for drive type and fuel type. Column 1 in Panel B includes fixed effects for body type, drive type, and fuel type. For column 2 and column 3, price regressions include model-variant fixed effects as defined in the Maritz data and Panel B includes model-variant fixed effects as defined in the IHS data. Column 1 and 2 are estimated by OLS and column 3 by IV. In column 3, log fuel costs and performance are instrumented using indicator technologies for the fuel-saving technologies from Figure 3, as well as the interactions of the indicator variables with a light truck indicator variable. First-stage results for price regressions are in Table B.7, and quantity regressions are in Table B.8. Panel C reports the change in WTP caused by a one percent increase in fuel economy or performance, assuming an own-price elasticity of demand equal to -3.

Table 4: **Valuation Ratios and Implicit Discount Rates**

	Our assumptions	Busse et al. (2013) assumptions
Panel A. Valuation ratio (percentage)		
Real discount rate = real reported APR 1.3 percent	53.6	72.7
Demand elasticity = 3		
Panel B. Implicit discount rate (percentage)		
Demand elasticity = 3	12.25	7.30

Notes: Panel A reports the valuation ratio, which is the ratio of the estimated WTP for a 1 percent fuel economy increase to the present discounted value of future fuel cost savings. Panel B reports the implicit discount rate, which is the discount rate that results in a valuation ratio of one. Both the valuation ratio and implicit discount rate are reported in percentages. The first column uses the baseline parameter assumptions and the second column uses the assumptions from Busse et al. (2013). See text for details on calculations and parameter assumptions.

Table 5: **Implicit Discount Rates Using Busse et al. (2013) Methodology**

Assumed demand elasticity	Implicit discount rate	
	Results reported in Busse et al. (2013)	Our results using Busse et al. (2013) methodology
-2	-4.0	2.1
-3	1.0	9.8
-4	5.5	17.6
-5	9.8	25.3

Notes: The implicit discount rate is computed by comparing vehicles in the fourth fuel economy quartile (highest fuel economy) with vehicles in the first fuel economy quartile (lowest fuel economy) assuming the own-price demand elasticities indicated in each row. Busse et al. (2013) results are repeated from their Table 9 column "NHTSA VMT and NHTSA PSR" and rows "Q1 versus Q4". To produce our results using their methodology, we estimate a price regression in Table B.7 (column 4) and quantity regression in Table B.8. We convert our estimates to implicit discount rates using the spreadsheet provided by Busse et al. (2013).

Table 6: **Directly Control for Vehicle Quality**

	(1) Baseline	(2)	(3)	(4)	(5)
Panel A. Dependent variable is log transaction price					
Log fuel cost	-0.354***	-0.385***	-0.312***	-0.172***	-0.297***
	(0.075)	(0.078)	(0.054)	(0.055)	(0.056)
Log performance	0.203***	0.280***	0.205***	0.335***	0.255***
	(0.074)	(0.041)	(0.048)	(0.039)	(0.045)
Control for vehicle quality		Quality attributes	Consumer experience ratings	Model-variant FE interacted with model generation FE	Model-variant FE interacted with generation change dummy
Number of observations	535,124	410,770	454,660	535,124	535,124
RMSE	0.13	0.13	0.13	0.39	0.13
F-stat (fuel cost)	185.5	163.5	174.6	110.0	163.1
F-stat (performance)	243.4	102.1	216.0	206.3	272.0
Panel B. Dependent variable is log new registrations					
Log fuel cost	-0.338***	-0.860***	-0.319***	-0.722***	-0.258**
	(0.116)	(0.149)	(0.115)	(0.154)	(0.117)
Log performance	0.371***	0.353***	0.320***	0.298***	0.362***
	(0.083)	(0.099)	(0.084)	(0.074)	(0.078)
Control for vehicle quality		Quality attributes	Consumer experience ratings	Model-variant FE interacted with model generation FE	Model-variant FE interacted with generation change dummy
Number of observations	535,124	410,770	454,660	535,124	535,124
RMSE	0.39	0.40	0.39	0.39	0.39
F-stat (fuel cost)	112.1	104.5	110.8	110.0	118.4
F-stat (performance)	150.1	229.4	141.9	206.3	210.4
Panel C. Willingness to pay (2010 USD)					
For 1 percent increases in					
• fuel economy	133.4	192.5	119.7	118.2	109.6
• performance	93.6	113.8	89.2	124.3	107.5

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

Notes: Robust standard errors in parentheses, clustered by vehicle model by state. Column 1 repeats the baseline in Table 3. Column 2 includes additional characteristics from the Chrome dataset to capture vehicle quality: number of passengers, passenger volume (cubic ft), cargo volume (cubic ft), , and "moonroof" or "sunroof" dummy variables. Column 3 adds controls of consumers' experience rating in the MaritzCX survey on a scale of 1 to 5: overall appearance; usefulness for carrying passengers; performance of entertainment system; exterior styling and workmanship; overall front room; interior material including seating and interior styling; quietness inside the vehicle; well equipped to prevent theft and vandalism; and exterior workmanship and attention to detail. In column 4, we further interact model-variant fixed effects with model generation fixed effects. In column 5, we further interact model-variant fixed effects with an indicator if the model is a new generation in the observed model year.

Table 7: **Including Proxies for Vehicle Quality and Other Sources of Bias**

	(1) Baseline	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Dependent variable is log transaction price							
Log fuel cost	-0.354***	-0.351***	-0.352***	-0.356***	-0.387***	-0.333***	-0.336***
	(0.075)	(0.055)	(0.056)	(0.054)	(0.083)	(0.055)	(0.054)
Log performance	0.203***	0.221***	0.228***	0.207***	0.200***	0.215***	0.210***
	(0.074)	(0.048)	(0.050)	(0.046)	(0.050)	(0.045)	(0.047)
Control for vehicle quality		Demo-graphic	Demo-graphic	Richer time controls	Drop CVT, deactiv--ation		
Finance control						Yes	
Drop 2009							Yes
Num. of obs.	535,124	497,867	450,515	535,124	535,124	515,994	507,461
RMSE	0.13	0.13	0.13	0.13	0.13	0.13	0.13
F-stat (fuel cost)	185.5	182.3	181.0	186.3	68.4	187.9	182.8
F-stat (perform.)	243.4	239.9	233.4	247.2	290.8	229.6	229.7
Panel B. Dependent variable is log new registrations							
Log fuel cost	-0.338***	-0.348***	-0.334***	-0.325***	-0.055	-0.339***	-0.363***
	(0.116)	(0.116)	(0.118)	(0.037)	(0.142)	(0.116)	(0.102)
Log performance	0.371***	0.363***	0.345***	0.356***	0.505***	0.371***	0.184**
	(0.083)	(0.084)	(0.083)	(0.022)	(0.136)	(0.083)	(0.075)
Control for vehicle quality		Demo-graphic	Demo-graphic	Richer time controls	Drop CVT, cylinder deactiv--ation		
Finance control						Yes	
Drop 2009							Yes
Num. of obs.	535,124	497,867	450,515	535,124	535,124	515,994	507,461
RMSE	0.39	0.40	0.40	0.39	0.40	0.39	0.38
F-stat (fuel cost)	112.1	111.2	109.2	112.5	77.9	112.9	112.3
F-stat (perform.)	150.1	147.6	143.0	149.5	127.3	149.8	138.2
Panel C. Willingness to pay (2010 USD)							
For 1 percent increases in							
• fuel economy	133.4	133.7	132.6	132.9	116.0	127.7	130.8
• performance	93.6	97.2	98.1	93.2	105.4	96.9	77.6

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

Notes: Robust standard errors in parentheses, clustered by vehicle model by state. Column 1 repeats the baseline in Table 3. Column 2 adds to column 1 six demographic controls: respondent's age, household size, indicator for male, urbanization indicator, 12 respondent education group fixed effects, and 23 household income group fixed effects. Column 3 adds to column 2 five additional demographic controls: number of wage earners, number of children, indicator equaling one if the respondent's spouse is employed, six respondent race fixed effects, and 20 respondent occupation fixed effects. Column 4 includes state by model-year fixed effects and state by month-of-year fixed effects. In column 5, we drop continuously variable transmission, cylinder deactivation, and their interactions with truck as instruments. In column 6, we include fixed effects for financing source (arrange own financing, finance via dealership, or do not finance) and fixed effects for payment type (automaker's loan/lease, bank loan/lease, friend/relative, cash, credit union loan, another finance company loan/lease, or other). In column 7, we drop observations if a transaction took place in 2009.

Appendix for Online Publication

A Vehicle Miles Traveled and Survival Rate

A.1 Vehicle Miles Traveled Schedules

We estimate vehicle miles traveled (VMT) over the lifetime of each vehicle by building on the models presented in Lu (2006). The data source we use to estimate VMT schedules is the 2009 National Household Travel Survey (NHTS). We use the publicly available data files on vehicle and household information, which contain 309,163 individual vehicles held by 150,147 surveyed households. We estimate the relationship between VMT and two variables: vehicle age and household income. We include household income as a covariate to account for the effect that the recession had on driving. We follow Lu (2006) in specifying a cubic relationship between VMT and vehicle age, where vehicle age is measured in years. We take a semi parametric approach in specifying the relationship between VMT and household income. We create 13 bins of household income, which correspond to bins present in both the NHTS and Maritz survey data, and we aggregate bins where necessary to make the bins consistent between the surveys. Furthermore, we convert income bins from the NHTS to 2014$ corresponding to bins in the 2014 wave of the Maritz survey data. We do this to be able to apply our estimated VMT model to households in the Maritz data, which we convert all incomes to 2014$. We estimate a separate intercept for each income group by regressing VMT on a fixed effect for each group. We also interact these fixed effects with a linear age variable to capture differences in VMT across income groups for different vehicle vintages. The interaction effects model the possibility that household driving intensity over the lifetime of a vehicle varies by income. Following Lu (2006), we estimate separate VMT models for cars and light trucks. We aggregate vehicle/household level observations to vehicle age by household income bin averages, giving us a total of 869 and 785 observations for the car and light truck specifications, respectively. The estimates for both models appear in Appendix Table B.11.

The estimates are plausible and most are statistically significant. For both vehicle classes, VMT increases with household income. The vehicle age/household income interaction terms are mostly negative and significant and are decreasing in household income. This implies that the marginal reduction in VMT from a vehicle aging by one year is larger for high-income households. This seems plausible given the preferences that high income households have for driving new vehicles more frequently by substituting miles away from their older vehicles to their newer vehicles. Conversely, low-income households tend to keep vehicles longer and drive them more when they are older. This relationship is apparent by plotting VMT schedules as a function of vehicle age for high- and low-income groups. Appendix Figure A.1 illustrates this effect for cars and light trucks, respectively. To account for the effect of fuel prices on VMT, we adjust the estimated VMT schedules by the change in national average fuel prices between the period of the 2009 NHTS (March 2008 to April 2009) and each year of the Maritz sample, assuming an elasticity of VMT to fuel prices of -0.1.

Figure A.1: **Estimated Vehicle Miles Traveled by Vehicle Age and Household Income**

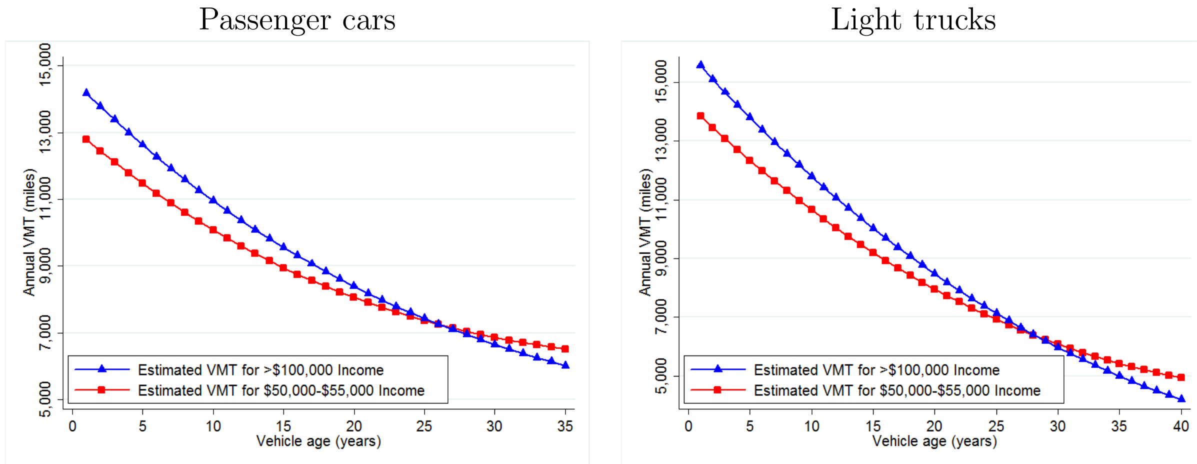

A.2 Vehicle Survival Schedules

We update the vehicle survival schedules in Lu (2006) using R. L. Polk data on vehicle registrations from 2002 to 2014. The R.L. Polk data are disaggregated by vehicle class (e.g., car and light truck), vehicle age, and year, where registrations are recorded for each vehicle age up to age 14. We drop observations with age 1 due to the increase in some vehicle counts from vehicle ages 1 and 2 across consecutive years, which would imply survival rates above 1. We estimate the following model:

$$age_{it} = \gamma_0 + \gamma_1 \ln(-\ln(1 - rate_{it}))$$

The variable is the survival rate of vehicles of age in year and is computed as the number of registered vehicles of age in year divided by the number of registered vehicles of age in year. Inverting the above equation yields a model that is comparable to the coefficient estimates in Lu (2006):

$$rate_{it} = 1 - \exp(-\exp(-\gamma_0/\gamma_1 + age_{it}/\gamma_1))$$

Defining $A = -\gamma_0/\gamma_1$ and $B = 1/\gamma_1$, Appendix Table B.12 presents estimates comparable to Lu (2006).

Appendix Figure A.2 plots the survival schedules for cars and light trucks, respectively. The figure illustrates that cars and light trucks are lasting longer than they have been historically. This is consistent with Lu (2006), who documents longer survival schedules than earlier time periods. The figures also highlight the importance of using more recent data for estimating vehicle survival schedules, as the newer data suggest greater VMT–and hence greater expected fuel costs–over vehicle lifetimes.

Figure A.2: **Vehicle Survivability Schedule**

Passenger cars Light trucks

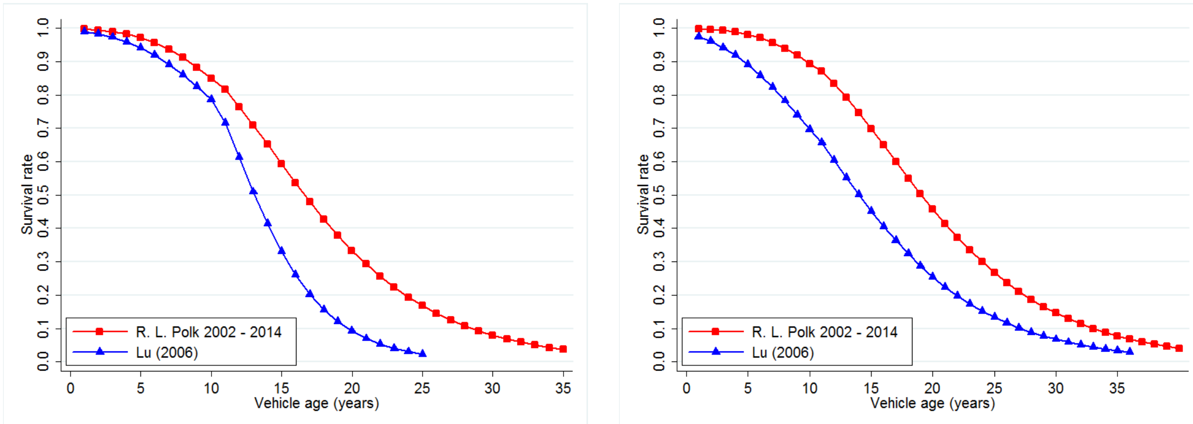

B Additional Summary Statistics, First-stage Results, and Robust Results

Figure B.1: **Distributions of Income and Education**

Panel A. Household income

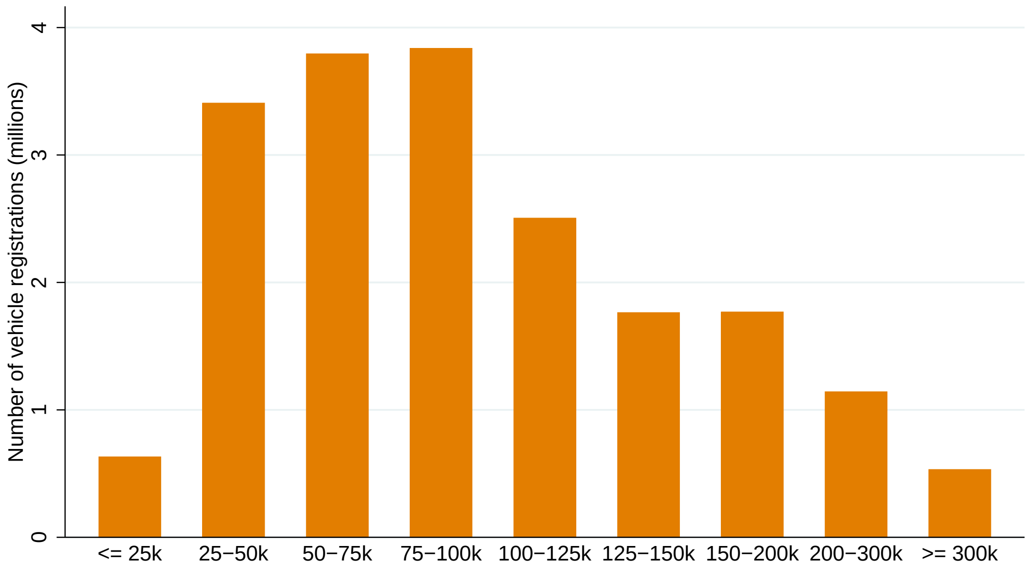

Panel B. Education

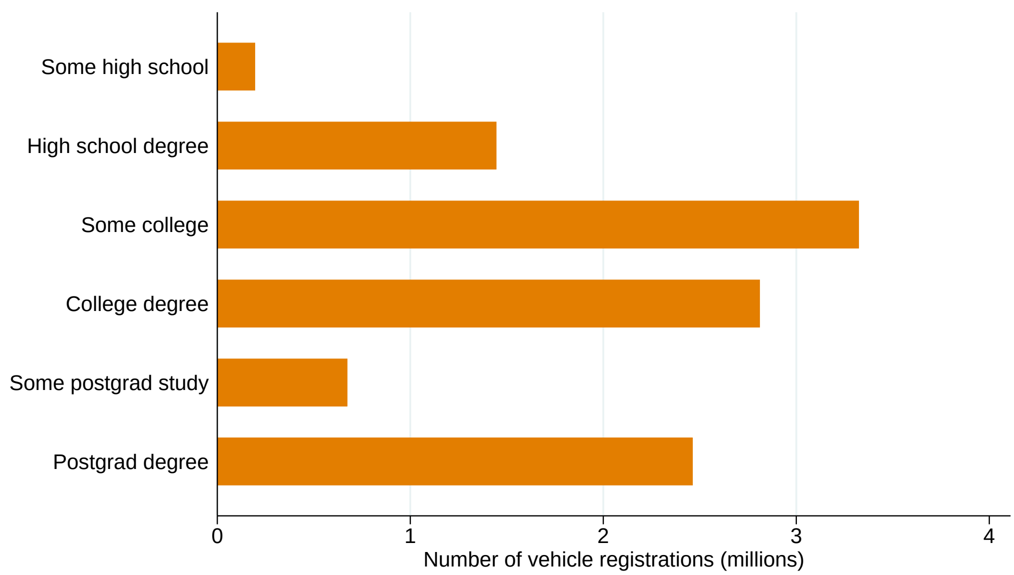

Table B.1: **Summary Statistics on Financing and Purchase Terms, 2009–2014**

Payment method	Share of vehicles (%)	Annual percentage rate (%)	Length (months)	Monthly payment (USD)	Down payment (USD)
Panel A. Purchased					
1. Financed	63.7	3.34	59.6	471	2,884
2. Cash	23.6	NA	NA	NA	NA
Panel B. Leased	12.7	NA	37.0	423	9,417

Notes: Annual percentage rate, length of the loan or lease, and payment information are weighted by registrations.

Table B.2: **First-Stage Coefficient Estimates from Baseline Price Specification**

Dependent variable	Log fuel cost		Log performance	
Supercharger	0.013**	(0.006)	0.156***	(0.003)
Turbocharger	-0.006**	(0.003)	0.086***	(0.027)
Gasoline direct injection	-0.055***	(0.007)	0.070***	(0.004)
Var. valve lift and timing	0.023***	(0.005)	0.001	(0.002)
Cylinder deactivation	0.033***	(0.006)	0.006***	(0.002)
Cont. variable transmission	-0.126***	(0.004)	-0.035***	(0.006)
Advanced transmission	-0.024***	(0.004)	-0.011***	(0.004)
Supercharger × truck	-0.002	(0.007)	-0.177***	(0.019)
Turbocharger× truck	-0.029***	(0.007)	0.110***	(0.031)
Gasoline direct inject. × truck	0.056***	(0.009)	-0.042***	(0.005)
Var. valve lift and timing × truck	-0.088***	(0.006)	0.021***	(0.004)
Cylinder deactivation × truck	-0.015**	(0.006)	-0.014***	(0.002)
Cont. variable transmission × truck	0.026***	(0.007)	0.047***	(0.006)
Advanced transmission × truck	-0.019***	(0.005)	0.002	(0.005)
Num. of observations	535,124		535,124	
F-stat (1st stg excl var.)	185.5		243.4	

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

Notes: Robust standard errors in parentheses, clustered by vehicle model and state. The table reports the first stage coefficient estimates for the baseline specification from column 3 of Table 3, Panel A. The bottom row reports the F-statistic on the test that the instruments are jointly equal to zero.

Table B.3: **First Stage Coefficient Estimates from Baseline Quantity Specification**

Dependent variable	Log fuel cost		Log performance	
Supercharger	0.053***	(0.014)	0.270***	(0.021)
Turbocharger	-0.081***	(0.006)	-0.033***	(0.012)
Gasoline direct injection	0.016***	(0.005)	0.103***	(0.009)
Var. valve lift and timing	-0.033***	(0.008)	0.006	(0.009)
Cylinder deactivation	0.109***	(0.007)	0.216***	(0.011)
Cont. variable transmission	-0.096***	(0.009)	-0.056***	(0.011)
Advanced transmission	0.007*	(0.004)	-0.014	(0.008)
Supercharger × truck	-0.066***	(0.021)	-0.098***	(0.023)
Turbocharger × truck	-0.020*	(0.010)	0.149***	(0.015)
Gasoline direct inject. × truck	-0.004	(0.008)	-0.093***	(0.012)
Var. valve lift and timing × truck	0.040***	(0.010)	0.014	(0.012)
Cylinder deactivation × truck	-0.076***	(0.008)	-0.102***	(0.014)
Cont. variable transmission × truck	0.071***	(0.015)	0.024*	(0.013)
Advanced transmission × truck	-0.008***	(0.001)	0.005***	(0.001)
Num. of observations	535,124		535,124	
F-stat (1st stg excl var.)	112.1		150.1	

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

Notes: Robust standard errors in parentheses, clustered by vehicle model and state. The table reports the first stage coefficient estimates for the baseline specification from column 3 of Table 3, Panel B. The bottom row reports the F statistic on the test that the instruments are jointly equal to zero.

Table B.4: **Composition of Willingness to Pay for Fuel Cost Savings and Performance**

Willingness to pay (2010 USD) for 1 percent increases in	Fuel economy (1)	Performance (2)
Panel A. WTP (Baseline)		
• price effect l_1	101.2	58.2
	[98.5, 104.0]	[56.1, 60.3]
• quantity effect l_2, assuming elasticity = 3	32.2	35.4
	[28.3, 35.9]	[31.3, 39.2]
• overall equilibrium effect, assuming elasticity = 3	133.4	93.6
Panel B. Average alternative elasticity		
• overall equilibrium effect, assuming elasticity = 2	149.5	111.3
• overall equilibrium effect, assuming elasticity = 4	125.4	84.7
• overall equilibrium effect, assuming elasticity = 5	120.5	79.4

Notes: For equilibrium price effect l_1 and additional price from quantity effect l_2, we report 95% confidence interval in parentheses using delta method.

Table B.5: **Assumptions for Implicit Discount Rate Calculations**

Vehicle age (years)	Our assumptions: VMT cars	Our assumptions: VMT trucks	Our assumptions: Survival rate cars	Our assumptions: Survival rate trucks	Assumptions of Busse et al. (2013): VMT cars	Assumptions of Busse et al. (2013): VMT trucks	Assumptions of Busse et al. (2013): Survival rate cars	Assumptions of Busse et al. (2013): Survival rate trucks
1	13,379	14,821	0.9972	0.9982	14,231	16,085	0.9900	0.9741
2	12,963	14,334	0.9944	0.9964	13,961	15,782	0.9831	0.9603
3	12,563	13,864	0.9897	0.9933	13,669	15,442	0.9731	0.9420
4	12,179	13,409	0.9823	0.9885	13,357	15,069	0.9593	0.9190
5	11,810	12,969	0.9714	0.9813	13,028	14,667	0.9413	0.8913
6	11,456	12,545	0.9564	0.9711	12,683	14,239	0.9188	0.8590
7	11,117	12,136	0.9367	0.9574	12,325	13,790	0.8918	0.8226
8	10,792	11,742	0.9122	0.9399	11,956	13,323	0.8604	0.7827
9	10,482	11,363	0.8828	0.9184	11,578	12,844	0.8252	0.7401
10	10,185	10,997	0.8488	0.8927	11,193	12,356	0.7866	0.6956
11	9,902	10,646	0.8168	0.8724	10,804	11,863	0.7170	0.6501
12	9,633	10,309	0.7650	0.8345	10,413	11,369	0.6125	0.6040
13	9,376	9,,985	0.7093	0.7922	10,022	10,879	0.5094	0.5517
14	9,131	9675	0.6515	0.7466	9,633	10,396	0.4142	0.5009
15	8,900	9,377	0.5932	0.6986	9,249	9,924	0.3308	0.4522
16	8,680	9,093	0.5357	0.6493	8,871	9,468	0.2604	0.4062
17	8,471	8,821	0.4804	0.5996	8,502	9,032	0.2028	0.3633
18	8,274	8,,561	0.4280	0.5505	8,144	8,619	0.1565	0.3236
19	8,088	8314	0.3791	0.5027	7,799	8,234	0.1200	0.2873
20	7,913	8,078	0.3341	0.4568	7,469	7,881	0.0916	0.2542
21	7,748	7,854	0.2931	0.4133	7,157	7,565	0.0696	0.2244
22	7,593	7,642	0.2562	0.3724	6,866	7,288	0.0527	0.1975
23	7,448	7,440	0.2231	0.3343	6,596	7,055	0.0399	0.1735
24	7,312	7,250	0.1938	0.2992	6,350	6,871	0.0301	0.1522
25	7,186	7,070	0.1679	0.2670	6,131	6,739	0.0227	0.1332
26	7,068	6,,900	0.1451	0.2377		6,663		0.1165
27	6,959	6,740	0.1252	0.2111		6,648		0.1017
28	6,857	6,591	0.1079	0.1871		6,648		0.0887
29	6,764	6,451	0.0928	0.1655		6,648		0.0773
30	6,678	6,320	0.0797	0.1462		6,648		0.0673
31	6,600	6,199	0.0684	0.1290		6,648		0.0586
32	6,528	6,086	0.0587	0.1137		6,648		0.0509
33	6,463	5,982	0.0503	0.1001		6,648		0.0443
34	6,404	5,887	0.0431	0.0880		6,648		0.0385
35	6,352	5,800	0.0369	0.0773		6,648		0.0334
36		5,720		0.0679		6,648		0.0290
37		5,648		0.0596				
38		5,584		0.0522				
39		5,527		0.0458				
40		5,477		0.0401				

Notes: The table reports the estimated vehicle miles traveled (VMT) and survival probability for cars and light trucks by vehicle age. Our estimates are from the 2009 wave of the National Household Travel Survey following the methodology of Lu (2006). The four columns on the right of the table show the assumptions from Busse et al. (2013).

Table B.6: **Alternative Assumptions for Computing Valuation Ratios and Implicit Discount Rates**

	Our assumptions of VMT and survival probability	Assumptions of Busse et al. (2013)
Panel A. Valuation ratio (percentage)		
A.1 Alternative demand elasticity		
A.1.1 Real discount rate = 1.3 percent, demand elasticity = 2	60.0	81.4
A.1.2 Real discount rate = 1.3 percent, demand elasticity = 3 **(base)**	53.6	73.0
A.1.3 Real discount rate = 1.3 percent, demand elasticity = 4	50.3	68.3
A.1.4 Real discount rate = 1.3 percent, demand elasticity = 5	48.4	65.6
A.2 Alternative real discount rate		
A.2.1 Real discount rate = 1.3 percent, demand elasticity = 3 **(base)**	53.6	73.0
A.2.2 Real discount rate = 5 percent, demand elasticity = 3	69.1	89.5
A.2.3 Real discount rate = 7 percent, demand elasticity = 3	77.7	98.6
A.2.4 Real discount rate = 10 percent, demand elasticity = 3	90.4	112.4
A.2.4 Real discount rate = 12 percent, demand elasticity = 3	98.9	121.5
A.3 Alternative future gasoline price assumptions		
A.3.1 Gasoline price follows random walk **(base)**	53.6	72.7
A.3.1 Gasoline price follow EIA AEO projection	57.2	77.6
Panel B. Implicit discount rate (percentage)		
Alternative demand elasticity		
B.1 Real discount rate = 1.3 percent, demand elasticity = 2	9.72	4.95
B.2 Real discount rate = 1.3 percent, demand elasticity = 3 **(base)**	12.25	7.30
B.3 Real discount rate = 1.3 percent, demand elasticity = 4	13.79	8.71
B.4 Real discount rate = 1.3 percent, demand elasticity = 5	14.83	9.63

Notes: The table reports valuation ratios in Panel A and implicit discount rates in Panel B, in percentages. The calculations use the same assumptions as in Table 4, except as indicated in the column and row headings.

Table B.7: **Price Regression Using Busse et al. (2013) Methodology**

Dependent variable: price	(1)	(2)	(3)	(4)
Gas prices × MPG quartile 1 (least efficient)	-142.052***	-149.354***	-104.193***	-112.484***
	(25.341)	(25.611)	(23.813)	(24.062)
Gas prices × MPG quartile 2	-22.614*	-25.443**	-20.104*	-24.213**
	(11.584)	(11.171)	(11.102)	(10.967)
Gas prices × MPG quartile 3	-40.029**	-40.828**	-37.303**	-38.539**
	(15.435)	(17.662)	(16.854)	(18.531)
Gas prices × MPG quartile 4 (most efficient)	25.754	31.412*	6.596	12.342
	(16.824)	(18.694)	(18.303)	(20.767)
State FE	Yes	Yes		
Model-year FE	Yes	Yes		
Month-of-year FE	Yes	Yes		
State × year FE			Yes	Yes
State × month-of-year FE			Yes	Yes
Include demographics		Yes		Yes
Number of observations	535,130	457,324	535,130	457,324
R-squared	0.90	0.90	0.90	0.90
Differences in WTP of Q1 versus Q4	$167	$180	$110	$135
* p<0.10 ** p<0.05 *** p<0.01				

Notes: Standard errors in parentheses, clustered by trim. The specifications are similar to Busse et al. (2013). The dependent variable is the transaction price, and the reported independent variables are interactions of the fuel price with fixed effects for the vehicle's fuel economy quartile. Observations are weighted by registrations, and regressions include model-variant fixed effects as well as the fixed effects indicated at the bottom of the table.

Table B.8: **Quantity Regressions Using Busse et al. (2013) Methodology**

Dependent variable: quantity	Coef.	SE	Average new cars registered per month per state (100)	Percentage change
Gas prices × MPG quartile 1 (least efficient)	-6.353***	(1.928)	87.99	17.41
Gas prices × MPG quartile 2	-3.479*	(2.057)	96.62	20.47
Gas prices × MPG quartile 3	8.848***	(2.489)	109.73	24.27
Gas prices × MPG quartile 4 (most efficient)	25.442***	(5.668)	122.57	30.84
Number of observations				12,182
R-squared				0.87
* p<0.10 ** p<0.05 *** p<0.01				

Notes: Standard errors in parentheses, robust to heteroskedasticity. The regression follows the Busse et al. (2013) methodology reported in their Tables 6 and 7. The dependent variable is the registrations by fuel economy quartile, state, and month. The regression reported in this table includes interactions of state fixed effects and transaction year fixed effects, interactions of state fixed effects and month of year fixed effects, and fuel economy quartile fixed effects. Observations are weighted by registrations.

Table B.9: **Baseline WTP by Expected Vehicle Miles Traveled (VMT)**

	(1) Baseline	(2)
Panel A. Dependent variable is log transaction price		
Log fuel cost	-0.354***	2.894***
	(0.075)	(0.785)
Expected VMT (in 1 million miles)		-35.329***
		(8.564)
Log fuel cost × expected VMT		-17.508***
		(4.251)
Log performance	0.203***	0.176***
	(0.074)	(0.050)
Number of observations	535,124	450,635
RMSE	0.13	0.14
F-stat (fuel cost)	185.5	185.5
F-stat (fuel cost by VMT)		188.6
F-stat (performance)	243.4	243.4
Panel B. Dependent variable is log new registrations		
Log fuel cost	-0.338***	-13.131***
	(0.116)	(3.695)
Expected VMT (in 1 million miles)		133.401***
		(38.981)
Log fuel cost × expected VMT		67.245***
		(19.642)
Log performance	0.371***	0.498***
	(0.083)	(0.086)
Number of observations	535,124	450,635
RMSE	0.39	0.43
F-stat (fuel cost)	112.1	112.1
F-stat (fuel cost by VMT)		
F-stat (performance)	150.1	150.1
Panel C. Willingness to pay (2010 USD)		
For 1 percent increases in		
• fuel economy	133.4	
at average VMT at 0.19 million miles		156.3
with one s.d. of VMT at 0.01 million miles		[141.3, 172.0]
• performance	93.6	97.9
* p<0.10 ** p<0.05 *** p<0.01		

Notes: Robust standard errors in parentheses, clustered by vehicle model by state. Column 1 repeats the baseline in Table 3. In column 2, we include expected lifetime VMT as an exogenous variable and its interaction with fuel costs as an endogenous variable. The lifetime VMT depends on household income group and broad market segment (car or truck). We construct it from survival data and annual VMT data as described in Section A.1.

Table B.10: **Alternative Measure for Performance**

Dependent variable: log price or quantity	(1) Baseline	(2)
Panel A. Price regression estimates		
Log fuel cost	-0.354***	-0.334***
	(0.075)	(0.111)
Log performance (hp/lb, or nm/lb)	0.203***	
	(0.074)	
Log performance (hp/lb)		0.217*
		(0.123)
Number of observations	535,124	535,130
RMSE	0.13	0.13
F-stat (fuel cost)	185.5	19.1
F-stat (performance)	243.4	98.0
Panel B. Quantity regression estimates		
Log fuel cost	-0.338***	-0.580***
	(0.116)	(0.038)
Log performance (hp/lb, or nm/lb)	0.371***	
	(0.083)	
Log performance (hp/lb)		0.589***
		(0.026)
Number of observations	535,124	535,130
RMSE	0.39	0.40
F-stat (fuel cost)	112.1	1540.2
F-stat (performance)	150.1	2047.9
* p<0.10 ** p<0.05 *** p<0.01.		

Notes: Standard errors in parentheses, clustered by trim. Column 1 repeats the baseline. Column 2 use horsepower-to-weight ratio for all vehicles. Column 2 uses torque-to-weight ratio for all vehicles.

Table B.11: **Estimates for Predicting Vehicle Miles Traveled**

Dep. var.: vehicle miles traveled	(1)		(2)	
Variables	Cars		Light truck	
Vehicle age	-298.5***	(16.87)	-341.0***	(21.61)
Vehicle age squared	6.493***	(0.582)	5.013***	(0.839)
Vehicle age cubed	-0.0391***	(0.00698)	-0.0152	(0.0110)
Household income $20,000-$25,000	-206.2	(258.2)	-538.1*	(322.5)
Household income $25,000-$30,000	810.5***	(252.5)	-258.2	(319.3)
Household income $30,000-$35,000	557.0**	(232.4)	37.95	(284.9)
Household income $35,000-$40,000	1,607***	(262.1)	710.3**	(328.7)
Household income $40,000-$45,000	1,099***	(225.5)	953.9***	(277.1)
Household income $45,000-$50,000	2,132***	(257.6)	1,651***	(327.5)
Household income $50,000-$55,000	2,096***	(227.5)	1,331***	(276.1)
Household income $55,000-$65,000	2,608***	(207.6)	1,883***	(261.6)
Household income $65,000-$75,000	2,878***	(216.3)	1,988***	(262.4)
Household income $75,000-$85,000	3,061***	(213.3)	2,311***	(262.8)
Household income $85,000-$100,000	3,647***	(201.8)	2,828***	(249.5)
Household income >$100,000	3,526***	(182.8)	3,098***	(231.3)
Vehicle age x household income $20,000-$25,000	21.99	(19.35)	27.94	(22.45)
Vehicle age x household income $25,000-$30,000	-45.81**	(18.53)	7.359	(21.73)
Vehicle age x household income $30,000-$35,000	-12.81	(17.57)	-13.43	(19.01)
Vehicle age x household income $35,000-$40,000	-50.82**	(20.03)	-21.02	(24.13)
Vehicle age x household income $40,000-$45,000	-25.37	(16.73)	-52.54***	(19.57)
Vehicle age x household income $45,000-$50,000	-80.87***	(20.01)	-65.82***	(25.24)
Vehicle age x household income $50,000-$55,000	-71.09***	(17.42)	-68.50***	(17.42)
Vehicle age x household income $55,000-$65,000	-86.40***	(15.27)	-82.73***	(19.13)
Vehicle age x household income $65,000-$75,000	-88.93***	(16.78)	-88.46***	(19.75)
Vehicle age x household income $75,000-$85,000	-94.87***	(16.25)	-91.95***	(20.13)
Vehicle age x household income $85,000-$100,000	-119.1***	(15.16)	-111.6***	(18.92)
Vehicle age x household income >$100,000	-125.9***	(13.64)	-131.2***	(16.74)
Constant	11,069***	(177.7)	12,937***	(228.6)
Observations	869		785	
R-squared	0.893		0.905	

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

Table B.12: **Estimates for Survival Rate**

	(1) Cars		(2) Light truck	
	Age ≤ 10	Age > 10	Age ≤ 10	Age > 10
$A = -\gamma_0/\gamma_1$	1.90	2.28	1.96	2.21
$B = 1/\gamma_1$	-0.13	-0.16	-0.12	-0.14