

Have US Fuel Economy and Greenhouse Gas Emissions Standards Improved Social Welfare?

Benjamin Leard* Resources for the Future

Joshua Linn University of Maryland and Resources for the Future

Katalin Springel Georgetown University

June 2020

Abstract

The US regulatory agencies predicted that the 2012-2016 passenger vehicle fuel economy standards would substantially increase social welfare. We provide the first comprehensive social welfare estimates of recent standards using a new equilibrium model that includes fixed and variable costs of raising fuel economy, manufacturer substitution between fuel economy and performance, and heterogeneous consumer preferences and manufacturer costs. The standards have reduced greenhouse gas emissions at a cost of \$6 per tonne of carbon, implying that standards have raised social welfare. Most fuel economy improvements have been achieved by trading off horsepower for fuel economy, rather than adjusting prices or adding fuel-saving technology.

Keywords: fuel economy standards, attribute trade-offs, medium run

JEL codes: L11, L50, L62, Q41

*Leard (leard@rff.org) is a fellow at Resources for the Future (RFF). Linn (linn@rff.org) is an associate professor at the University of Maryland and a senior fellow at RFF. Springel (Katalin.Springel@georgetown.edu) is an assistant professor at Georgetown University. We thank seminar participants at the Congressional Budget Office and the University of California at Davis, participants at the Association of Environmental and Resource Economists Summer Conference, and Ken Gillingham for helpful comments.

1 Introduction

Passenger vehicles are the main source of greenhouse gas (GHG) emissions from the transportation sector. Most countries that aim to reduce transportation sector emissions rely primarily on passenger vehicle fuel economy and GHG standards, and about half of all new passenger vehicles sold globally are subject to such standards.

The US standards, which have been tightening since 2005, have been particularly controversial. Starting in 2011, the Environmental Protection Agency (EPA) and National Highway Traffic Safety Administration (NHTSA) adopted standards that would have nearly doubled fuel economy by 2025. The agencies predicted that the standards would yield cumulative net social benefits of more than \$700 billion (2015\$; (EPA 2010; 2012)). However, in 2018 and 2020, EPA and NHTSA updated their analysis and substantially weakened the standards (EPA and NHTSA 2018).¹ Many vehicle manufacturers have supported weakening the standards, although not by as much as the agencies did; environmental groups have opposed the proposal vigorously because of the higher fuel consumption and GHG emissions.

The literature on fuel economy and GHG standards has highlighted several inefficiencies of the standards relative to a GHG emissions price: rebound (the increase in driving caused by lower fuel costs when standards tighten), delayed scrappage caused by vintage-differentiated regulation, and relying on tested rather than real-world emissions (Jacobsen 2013; Jacobsen and van Benthem 2015; Reynaert 2019). However, the regulatory agencies set standards based on welfare analysis of the standards that they make prior to setting the standards. The literature has not addressed the questions relevant to the agencies' decisions about setting future standards: have the recently tightened standards increased social welfare, and would weakening future standards increase social welfare? Among others, Klier and Linn (2016) suggest that the agencies' overstated the case for tighter standards by omitting certain types of manufacturer and consumer responses to standards, but the literature provides only rough estimates of the welfare effects of recent standards.

This paper reports the first retrospective welfare analysis of any fuel economy or GHG standards, focusing on the first phase of the US standards, which covered model years 2012–2016. We estimate welfare effects of the standards using a new computational model

¹In 2016, EPA and NHTSA reported analysis that largely supported their predictions from 2011 (EPA et al. 2016). Compared with the 2016 analysis, the 2018 analysis assumed higher technology costs based on updated modeling and data. The 2018 analysis concluded that tighter standards would increase traffic accidents and fatalities, which further supported weakening the standards, a finding that Bento et al. (2018) dispute.

that incorporates advances in modeling both the supply and the demand sides of the new-vehicles market.

In the supply component of the model, manufacturers choose the price, fuel economy, and horsepower of each vehicle in the market. In fact, fuel economy is chosen jointly with performance, and manufacturers can increase fuel economy by reducing performance. Most of the literature has assumed that other vehicle attributes are exogenous, and all of the literature has assumed that standards do not affect the rate of fuel-saving technology adoption. These assumptions contradict [Klier and Linn \(2016\)](#), who show that standards have increased the rate of technology adoption and reduced performance. We follow [Knittel \(2011\)](#) and [Klier and Linn \(2012\)](#) to estimate the trade-off between fuel economy and horsepower, and we allow for the possibility that standards affect the rate of technology adoption.

The supply component includes two main advances over the literature. First, we introduce a cost of adding fuel-saving technology that arises from redesigning and testing the vehicle. [Blonigen et al. \(2017\)](#) estimate that typical fixed costs of a redesign amount to around \$1 billion, which includes costs of adopting fuel-saving technology as well as other elements of a redesign, such as restyling the exterior. Adopting emissions-reducing technologies likely introduces fixed costs, as discussed in [National Research Council \(2015\)](#) and EPA and NHTSA rulemaking, such as [EPA \(2010\)](#).² EPA and NHTSA approximate these fixed costs by assuming a constant markup above marginal costs, but the economics literature has ignored the possibility that firms incur fixed costs when adopting technology.³ We introduce a redesign cost that increases with the amount of technology adoption but is fixed with respect to sales.

The second novel feature of the supply side is that we use the first-order conditions to a firm’s profit maximization to estimate all supply parameters using observed attribute choices. Specifically, we estimate the fixed costs from observed efficiency improvements over time, and we jointly estimate marginal costs of production and shadow costs of regulatory constraint from observed price and attribute choices. In contrast, most previous studies, such as [Klier and Linn \(2012\)](#), use National Academies of Sciences reports to estimate technology costs. We find that average fixed costs (that is, total fixed costs divided by sales) are roughly comparable to marginal costs of improving efficiency. Fixed costs rise quickly with efficiency

²By fixed costs, we mean technology costs that a manufacturer incurs that are independent of the number of vehicles that it sells.

³More precisely, EPA and NHTSA distinguish between direct costs of an emissions-reducing or fuel-saving technology, which include the increase in production costs caused by the adoption, and indirect costs, which include all other costs. Indirect costs may include both fixed and variable costs. In this paper, we distinguish between variable and fixed costs, rather than direct and indirect costs.

improvement. Estimated shadow prices are larger but have the same order of magnitude as the shadow price implied by observed credit prices (Leard and McConnell 2017).

The demand component of the model builds on Leard et al. (2019), allowing for additional heterogeneity in consumer preferences. In the model, consumers choose among a highly disaggregated choice set that roughly matches the vehicle definition in the agencies' analysis and increases by roughly an order of magnitude the number of choices compared with the literature. The disaggregation allows us to estimate more realistic preference parameters and substitution elasticities than models that include a more aggregated choice set, such as Berry et al. (1995) and Jacobsen (2013). Preferences are estimated from data that cover the study period. We find that consumers undervalue fuel economy; a typical consumer is willing to pay just less than 50 cents for \$1 of future fuel cost savings (i.e., a valuation ratio just below 0.5), with substantial heterogeneity across demographic groups around that mean. The typical consumer prefers horsepower to fuel economy.

Our welfare analysis allows for the possibility that consumers undervalue the fuel costs of vehicles they consider purchasing.⁴ Although Busse et al. (2013) conclude that consumer choices are consistent with rational behavior, Allcott and Wozny (2014) find some evidence of undervaluation, and Leard et al. (2017; 2019) and Gillingham et al. (2019) find substantial evidence of undervaluation, meaning that consumers are willing to pay less than \$1 for \$1 of future fuel cost savings. Thus, at a minimum, the recent literature leaves open the possibility that consumers systematically undervalue fuel economy. The regulatory agencies assume consumers undervalue fuel costs, but except for Allcott (2013), the literature has reported welfare estimates based on the assumption, either implicit or explicit, that consumers fully value fuel costs.⁵

We use the model to perform a static analysis of the standards by comparing steady states with different levels of the standards. The start year is model year 2011, which is prior to the first phase of the standards. The steady state is represented by the equilibrium five years later, which allows manufacturers time to redesign vehicles to meet the standards. We compare two levels of standards: maintaining the 2012 standards through model year

⁴If consumers systematically undervalue fuel costs, manufacturers would have insufficient incentive to adopt fuel-saving technology. Fuel economy or GHG standards could increase private welfare, and could be more cost effective than a carbon tax (Allcott and Greenstone 2012).

⁵NHTSA and EPA do not explicitly state their assumption that consumers undervalue fuel costs. The agencies implicitly assume consumer undervaluation, given their assumption about baseline technology adoption. The agencies assume that technologies that have a certain payback are adopted by manufacturers in the absence of fuel economy and GHG regulation. Other technologies are only adopted if it is cost-effective to do so because of regulation. The assumed payback period implies a substantial degree of undervaluation (Bento et al. 2018).

2016, versus tightening standards through 2016; the latter scenario requires 13 percent higher fuel economy than the former. The model relaxes all four assumptions mentioned above: horsepower is endogenous, standards can affect the amount of fuel-saving technology adoption, manufacturers incur fixed costs of technology adoption, and consumers may undervalue fuel cost savings.

We find that standards cost between \$6 and \$188 per tonne of carbon abated, where the range arises from the method used to compute consumer welfare. Net societal benefits range from −\$11 billion to \$4 billion per year. At one extreme, undervaluation could reflect hidden costs. For example, many consumers do not like stop-start ignitions, and undervaluation could reflect that disutility. At the other extreme, consumers could mistakenly undervalue fuel cost savings, causing them to be willing to pay just 50 cents for \$1 of future fuel cost savings. For a lower bound of consumer benefits, we use estimated preference parameters, and for the upper bound, we follow [Allcott \(2013\)](#) and [Train \(2015\)](#). The wide range of welfare estimates underscores the importance of allowing for the possibility that consumers mistakenly undervalue fuel cost savings.

We make two general points about the welfare results. First, the extensive amount of heterogeneity in the model yields accurate estimates of manufacturers’ compliance behavior as well as overall benefits and costs of the standards. Not only do consumers have heterogeneous preference parameters, but also manufacturers have heterogeneous technology adoption costs and costs of trading off performance for fuel economy. That is, the cost to manufacturers of trading off the attributes depends partly on consumer valuation of the attributes. For example, Fiat-Chrysler buyers typically have higher willingness to pay for performance than do Toyota buyers, making it more costly (in terms of forgone profits) for Fiat-Chrysler to trade off performance for fuel economy and comply with the standards.

Second, as we noted above, we relax several assumptions made in the literature, all of which have large welfare implications. Most of the literature has ignored changes in attributes other than fuel economy, has assumed zero fixed costs of technology adoption, and has ignored consumer internalities. Failing to account for attribute changes and fuel cost undervaluation would understate net welfare benefits, but failing to account for fixed costs would overstate net welfare benefits. Moreover, the analysis of the agencies shares the first two deficiencies (assuming horsepower is exogenous and failing to allow standards to affect the rate of technology adoption).⁶ An additional problem is that the agencies

⁶Agencies do include some technology adoption in simulations of scenarios that maintain constant standards. However, because the agencies assume performance does not change over time, their simulated adoption rates are far lower than rates reported in the literature (e.g., ([MacKenzie and Heywood 2015](#))).

do not use an internally consistent model of consumers’ and manufacturers’ behavior, in which consumer vehicle choices respond to vehicle fuel economy and price changes, and manufacturers anticipate those responses when choosing prices and technology. For example, Toyota buyers typically have higher willingness to pay for fuel economy than Chrysler buyers (Leard et al. 2019), which can affect the manufacturers’ choices of technology, performance, and attributes. Thus, existing estimates of the welfare effects of fuel economy and GHG standards include four main assumptions that are inconsistent with manufacturers’ and consumers’ behavior.

We find that the majority of fuel economy improvements are made by trading off horsepower for fuel economy, rather than by adding fuel-saving technology. Fixed costs of technology adoption amount to \$3 billion per year and thus account for a substantial portion of overall welfare changes. Finally, failing to account for technology adoption and performance improvements in the absence of tighter standards would increase net benefits by about \$11 billion per year. Thus, relaxing each of the assumptions has large welfare implications. Section 5 reports additional scenarios, such as assuming higher fuel prices and replacing footprint-based standards with uniform standards.

Returning to the critiques of the agencies’ ex ante analysis of the standards, we note that our estimates account for the decrease in gasoline prices that occurred in 2014, the possibility that the agencies underestimated technology costs, and trade-offs between performance and fuel economy. Our results demonstrate that, assuming consumer internalities for fuel economy choices, the first phase of the standards has positive net benefits, even accounting for these factors.

Our modeling focuses on the new-vehicles market to address deficiencies in the existing models used to evaluate fuel economy standards. As noted above, the literature has identified three inefficiencies of standards compared with a carbon tax: rebound, vintage-differentiated regulation, and the use of tested rather than on-road emissions to assess compliance. Our welfare estimates account for rebound and on-road emissions using the same approach that the EPA and NHTSA use in their welfare analysis, which facilitates comparison between our results and theirs. Like their analysis, we do not include the effects of vintage-differentiated regulation. Moreover, we do not include certain benefits of higher fuel economy, such as reduced refueling time, because of the numerous additional assumptions that would be necessary. The results in Jacobsen and van Benthem (2015) and the agencies’ analysis suggest that including these margins would likely strengthen our finding that the standards have increased social welfare.

The three papers most closely related to ours are [Leard et al. \(2019\)](#), [Klier and Linn \(2012\)](#), and [Whitefoot et al. \(2017\)](#). The demand model builds on [Leard et al. \(2019\)](#), but their analysis is restricted to the short run, meaning that fuel economy and horsepower are exogenous. They do not consider the technology adoption that occurs over time when standards do not tighten, and they do not include fixed costs. [Klier and Linn \(2012\)](#) and [Whitefoot et al. \(2017\)](#) allow for trade-offs between fuel economy and horsepower, but they do not include fixed costs or technology adoption when standards do not tighten. Moreover, their consumer preferences are estimated using data from the mid-2000s, and we find evidence that consumer preferences have changed substantially over time.

More broadly, our estimation of supply parameters can be adapted for analyzing other product markets in which producers design and test their products, choosing multiple attributes that consumer value. For example, regulating the energy efficiency of home appliances could induce research and development in energy efficiency and affect multiple attributes of the products ([Houde and Spurlock 2015](#)).

2 Policy Context, Data, and Market Trends

2.1 Policy Context

The 2008 Energy Independence and Security Act requires NHTSA to set fuel economy standards through 2030. The 2009 Supreme Court ruling in *Massachusetts v. EPA* granted EPA the authority to regulate passenger vehicle’s GHG emissions under the Clean Air Act. In 2009, the Obama administration set requirements for the first phase of new joint fuel economy and greenhouse gas standards for passenger vehicles, requiring manufacturers to increase the fuel economy (and reduce tailpipe GHG emissions) of their vehicles sold during 2012–2016. Compared to standards in prior years, the new standards required a sharp increase in average fleetwide fuel economy of new vehicles sold, mandating an increase of about 10 miles per gallon above 2009 levels by 2016, or about 5 percent per year.⁷ In 2020, the Trump administration substantially weakened standards for the years 2021 through 2026. Figure 1 shows how the fuel economy standards (expressed in gallons of fuel consumption per 100 miles) evolved over time, with the dashed lines showing the recently weakened standards.

Along with setting the new federal standards, the federal agencies were required by law to perform a regulatory impact analysis (RIA). The RIA included an ex ante cost-benefit analysis of the 2012–2016 proposed standards. The RIA predicted that the standards would

⁷This increase of 10 mpg refers to the fuel economy used to assess compliance with the standards, which is based on a specified laboratory testing procedure. This procedure yields fuel economy ratings that are roughly 20–25 percent higher than the fuel economy ratings that appear on new vehicles’ window stickers.

raise private and social welfare, with fuel cost savings accounting for most of the overall welfare gains.

2.2 Data

The primary data source for our analysis is the MaritzCX New Vehicle Customer Survey. MaritzCX conducts a monthly survey of households that recently purchased or leased a new vehicle, with a response rate of about 9 percent. Our data include about 1.1 million new-vehicle transactions between 2010 and 2015, which represents about a 1 percent sample of new-vehicle buyers over this time period.

The MaritzCX survey asks detailed questions about the new vehicle and the demographics of the household. Survey respondents report sales or lease prices of their recently obtained vehicle, which include sale taxes but exclude any credit received from trade-ins.⁸ These transaction prices may be substantially different from the manufacturer’s suggested retail price (MSRP) frequently used by previous studies, reflecting the outcome of price negotiations between the survey respondent and salesperson or unobserved promotions offered by the dealer or the manufacturer. The advantage of using transaction prices rather than MSRP in our empirical analysis is that it reduces concerns about price measurement error, and the greater price variation helps identify the preference parameters. Demographic information for the survey includes the respondent’s age and years of schooling, household income (defined as total annual household income from all sources, including wage and investment, before taxes), state of residence, and urbanization (defined by population density).

The MaritzCX survey data also include the vehicle identification number (VIN) for each observation. This allows us to define vehicles at a highly disaggregated level that yields about 800 unique vehicle choices for consumers each year. Specifically, a vehicle is defined by a unique model year, make, model, trim or series, drive type, fuel type, body style, and number of cylinders.⁹ This level of aggregation better reflects the choice set available to consumers when they make their vehicle purchase or lease decision, compared with the approach that is commonly taken in the literature, which defines a vehicle at the model level. Furthermore, EPA and NHTSA define a unique vehicle in a similar fashion, which

⁸For about 10 percent of the transactions, we do not observe prices. We impute prices for these missing observations based on observed transaction prices of the most closely related vehicles in the data.

⁹Each model year is defined as September of the previous calendar year to August of the current calendar year. For example, MY 2015 is defined as September 2014 to August 2015. This definition of model year aligns closely with the typical vehicle production cycles observed in practice—that is, vehicle characteristics remain the same during each model year but may change across model years.

eases the comparison between the results of their welfare analysis of the fuel economy and GHG standards and ours. This highly differentiated choice set improves the identification of the preference parameters compared with prior work that relied on more aggregated choice sets. The level of disaggregation also helps us alleviate concerns regarding measurement error and the resulting bias in the estimated preference parameters.

We complement the survey data with a variety of other data. We use fuel economy ratings (miles per gallon) and energy efficiency technology data from EPA.¹⁰ For each vehicle’s fuel economy, we use combined city and highway fuel economy ratings. The fuel-saving technologies include engine technologies such as cylinder deactivation, turbochargers, gasoline direct injection, and valve timing and lift, as well as transmission technologies such as additional gears, continuously variable transmissions, and dual-clutch transmissions. For each vehicle, we merge vehicle characteristics data from Wards Automotive, which include MSRP, horsepower, size-related measures (such as width, height, and length), curb weight, and wheelbase.¹¹

We further supplement the main data set by adding new vehicle registration data obtained from IHS Automotive and new- and used-vehicle sales from the Consumer Expenditure Survey (CEX).¹² The IHS data include quarterly vehicle registrations for all US households, providing us with market-level information on vehicle demand. We obtain a count of used-vehicle purchases by year and demographic group from the CEX data to use in our consumer demand model, in which we define the outside option as the decision to purchase a used vehicle. We use the IHS Automotive and CEX data to correct for nonrandom sampling of the MaritzCX survey. In particular, we construct weights for the observations in the MaritzCX data to match the distribution of sales across vehicles in the IHS data and across demographic groups in the CEX data.¹³

We define 20 demographic groups: 5 income groups (constructed to approximately represent the income quintiles of the distribution of households that purchased a vehicle in the CEX data), 2 age groups (defined as the respondent being younger than the median age of 45 years in the data, or 45 and older), and an urbanization indicator (equal to one if the

¹⁰We link the EPA data to the MaritzCX survey data by vehicle with the exception of body style, which is not reported in the EPA data set. Nevertheless, this is unlikely to substantially reduce the available variation in fuel economy. Based on the Wards Automotive data, we find that our definition of a unique vehicle (excluding differentiation by body type) preserves 99 percent of variation in the EPA estimated fuel economy rating across all new vehicles.

¹¹For a small number of vehicle transactions with missing observations, we impute car attributes using data from Cars.com.

¹²We merge the IHS data to the MaritzCX survey data by vehicle and the CEX data by model year, make and vehicle class, which is an indicator equal to one if the vehicle is a light truck.

¹³See [Leard et al. \(2019\)](#) for more details on the weighting procedure.

household is located in an area with a population density above the median). Additionally, the demographic groups are defined to roughly match the count of households in the CEX data for each group. The detailed demographic information obtained from the MaritzCX survey allows us to differentiate vehicle purchases made across different demographic groups and estimate heterogeneous preferences for each vehicle attribute. Our choice of demographic groups is motivated by the fact that observed vehicle purchase patterns vary substantially within income groups and across age groups and urbanization rather than just across income groups.

Finally, fuel prices come from the US Energy Information Administration (EIA). We deflate all transaction and fuel prices using the Bureau of Labor Statistics Consumer Price Index and adjust them to 2015 US dollars.

2.3 Market Trends and Stringency of the Standards

In this section, we plot trends in the new-vehicle market during our sample period of 2010–2015. In Panel A of Figure 2, we plot sales-weighted average fuel economy of all new cars and light trucks sold in the United States. Fuel economy of cars and light trucks increases gradually, by about 4 and 3 miles per gallon, respectively (about a 15 percent increase relative to 2010 levels). In Panel B of Figure 2, we plot the sales-weighted performance for cars and light trucks as measured by the log of the ratio of horsepower to weight, which is a commonly used measure of performance because it is directly proportional to the time needed to accelerate the vehicle from rest to 60 miles per hour. The data show that performance changed by just a few percentage points during the sample period. Vehicle sales-weighted purchase prices denoted in 2015\$ are plotted separately for cars and light trucks in Panel C of Figure 2. Real purchase prices remained flat for both vehicle categories, with the average purchase price of cars and light trucks being about \$28,000 and \$36,000, respectively.

In Figure A.1, for every vehicle in the 2011 sample, we plot the vehicle’s fuel economy and its calculated requirement for 2012 and 2016, along with the vehicle’s footprint. Panel A shows the scatter plot for cars, and, Panel B, light trucks. In both panels, many 2011 vehicles already exceed their 2012 requirements. But by 2016, the standard is sufficiently stringent that most 2011 vehicles have fuel economy well below their requirements. Most of the vehicles exceeding their 2012 requirement have small footprints.

3 The Equilibrium Model

We model the equilibrium of the US market for new light-duty vehicles. This section presents the demand and supply sides of the market separately.

3.1 Demand

The general structure of our vehicle demand model follows [Leard et al. \(2019\)](#). We define each market as a model year, indicated by t . We assume that households decide to purchase a single new or used vehicle to maximize their private utility. Vehicle alternatives are indexed by $j = 1, 2, \dots, J$, and we denote the outside option to purchase a composite used vehicle as $j = 0$. We denote vehicle attributes by k and demographics of household by d . We also partition households into 20 mutually exclusive and collectively exhaustive groups, indexed by g , based on their demographics as defined in [Section 2](#): 5 income groups, 2 age groups, and an indicator for urbanization. We define the base demographic group as young and urban households that are in the lowest income category. The indirect utility u_{ijt} of household i from choosing vehicle j in market t is

$$u_{ijt} = v_{ijt} + \varepsilon_{ijt} = \sum_k \sum_g x_{jkt} h_{igt} \beta_{kg} + \sum_k x_{jkt} \bar{\beta}_k + \xi_{jt} + \nu_{gjt} + \varepsilon_{ijt}. \quad (1)$$

The term x_{jkt} stands for the value of characteristic k for vehicle j and model year t . We include six vehicle attributes: the transaction price, fuel economy, performance (defined as the logarithm of the ratio of horsepower to weight), footprint (defined as the product of the vehicle's wheelbase and width), fuel type, and drive type (the latter two extend the demand model in [Leard et al. \(2019\)](#) by allowing for additional heterogeneity). We add these interactions to address possible endogeneity concerns at the household group level. In particular, observed attributes such as fuel economy could be correlated with unobserved attributes, which would bias our estimates of willingness to pay (WTP) for fuel economy. Controlling for additional vehicle attributes, such as fuel type and drive type, reduces this bias. The term h_{igt} represents an indicator variable that equals one if household i is in demographic group g . The coefficient β_{kg} denotes the difference in the marginal utility of vehicle attribute k for households in demographic group g relative to households in the base demographic group, and the parameter $\bar{\beta}_k$ denotes the marginal utility of vehicle characteristic k for households in the base group. The term ξ_{jt} represents the national mean valuation of the unobserved vehicle attributes (such as vehicle color or cargo space) for vehicle j in market t , which is constant across demographic groups. The term ν_{gjt} represents group g unobserved mean utility for vehicle j in market t . We assume that the error term ε_{ijt} has a

type 1 extreme value distribution, such that the probability that household i chooses vehicle j in market t is

$$Pr_{ijt} = \frac{e^{v_{ijt}}}{\sum_l e^{v_{ilt}}}. \quad (2)$$

For each household demographic group, we normalize the outside good utility to zero. This yields a linear estimation equation relating market shares, household demographics, and vehicle characteristics:

$$\ln(s_{gjt}) - \ln(s_{g0t}) = \sum_k \sum_g x_{jkt} h_{igt} \beta_{kg} + \sum_k x_{jkt} \bar{\beta}_k + \xi_{jt} + \nu_{gjt} \quad (3)$$

where s_{gjt} and s_{g0t} are market shares for vehicle j and the outside option (defined as a used car purchase) for demographic group g in market t .

The demand model allows for extensive preference heterogeneity across demographic groups. For each of the 20 demographic groups we estimate preference parameters for each of the six vehicle attributes. The demand model implies more plausible substitution patterns across vehicles relative to the simple logit demand model that restricts substitution between vehicles to being proportional to the vehicles' market shares, regardless of differences in attributes. We model consumer heterogeneity based on observed demographics. According to the model, for each demographic group, market shares drive substitution between vehicles. A benefit of this approach is that the model provides a direct link between the estimated demand parameters and the variation in welfare effects across households.¹⁴

3.2 Supply

Before describing the supply side of the model, we discuss briefly the three main supply-side assumptions that we relax. First, most of the literature has ignored the possibility that standards can affect other vehicle attributes besides fuel economy. [Goldberg \(1998\)](#) and [Jacobsen \(2013\)](#), among others, hold fixed other attributes, such as horsepower. However, as in [Knittel \(2011\)](#), one can conceive of a frontier defined by a level of fuel-saving technology. For a given size engine (such as one with four cylinders), the manufacturer

¹⁴Our approach avoids drawbacks of the random coefficients demand models explained in recent literature (e.g., ([Knittel and Metaxoglou 2014](#))), such as multiple local optima that can produce a range of potential elasticities and welfare effects. The model produces a unique set of preference parameters for each demographic group and can be estimated quickly. Furthermore, many papers in the literature include demand models in which changes in choice sets across markets identify heterogeneity parameters. This makes it difficult to determine whether the implied heterogeneity reflects the preference heterogeneity or something else ([Akerberg and Rysman 2005](#)). In contrast, in our demand model, heterogeneity parameters are identified by variation across demographic groups in response to variation in vehicle attributes across vehicles and markets.

can retune the engine or make small technology changes to reduce horsepower and raise fuel economy, moving along a frontier. [Klier and Linn \(2016\)](#) show that standards have affected performance (using horsepower as a proxy for performance), and [Ito and Saltee \(2018\)](#) suggest that standards could affect an even broader range of attributes. To our knowledge, only [Klier and Linn \(2012\)](#) and [Whitefoot et al. \(2017\)](#) allow performance to be endogenous in their models.

Second, the literature has not properly accounted for the industry dynamics of technology adoption. [Knittel \(2011\)](#) and [Klier and Linn \(2016\)](#) show that historically, during periods when standards were not tightening, such as the 1990s and early 2000s, manufacturers have adopted fuel-saving technologies and used the efficiency improvements to increase vehicle size and performance (i.e., horsepower or torque), while maintaining a constant level of fuel economy. In contrast, during periods when standards tighten, manufacturers adopt more technology and also shift along the frontier, trading off performance for fuel economy. This situation implies that tightening standards over time could increase fuel economy but decrease performance, relative to the equilibrium if standards remain fixed. Studies that assume exogenous horsepower cannot handle this possibility, and the two studies cited above that endogenize horsepower do not account for performance improvements if standards remain fixed.

Third, the literature does not include the fixed costs of technology adoption. Vehicle manufacturers redesign each vehicle every five to seven years, and as discussed above, the economics literature has largely ignored the possibility that firms incur fixed costs when adopting technology.

Following [Klier and Linn \(2012\)](#), we model the supply side in two periods. There is a set of vehicles in the market that have endowments of the same attributes that are included in the demand model. Each manufacturer maximizes profits of selling its vehicles.

In the first period, the manufacturer designs the vehicle and chooses a level of technology, $T_{d(j)}$, where the subscript indicates that the manufacturer chooses the same level of technology for each model, d . The firm has the opportunity to add technology for each of the models it sells.

The second period represents the five years that follow the design period, reflecting the typical cycle length. In the second period, the manufacturer chooses fuel economy, horsepower, and price. The level of technology corresponds directly to the energy efficiency of the vehicle, such that a higher level of technology allows the manufacturer to increase fuel economy without sacrificing horsepower. Technology is scaled such that increasing $T_{d(j)}$ by

0.01 units allows the manufacturer to increase fuel economy by 1 percent. Subsequently, having chosen a level of technology, the manufacturer can trade off fuel economy for performance. Thus, the representation of technology corresponds to a production possibilities frontier in horsepower–fuel economy space. Increasing the level of $T_{d(j)}$ causes the frontier to shift away from the origin, and after choosing $T_{d(j)}$ the manufacturer can move along the frontier as it trades off horsepower for fuel economy.

Thus, increasing technology from the endowed value allows the firm to provide higher fuel economy or horsepower (or both). Higher technology causes marginal costs to increase according to the function $c_j(T_{d(j)})$. Technology adoption also requires a fixed cost $F(T_{d(j)})$ to redesign and test the vehicle, where the fixed costs are incurred once for each model. Thus, when deciding whether to adopt technology, the firm faces a trade-off between higher revenue, which the technology enables, and the additional fixed and variable costs. The existence of the fixed costs implies that manufacturers adopt more technology for higher-selling vehicles, all else equal. This is consistent with the empirical evidence in [Klier and Linn \(2016\)](#).

We model the two periods of the manufacturer problem as static. The profit maximization problem is

$$\max_{\{p_j, m_j, h_j, T_{d(j)}\}_{j \in J_f}} \sum_{j \in J_f} [p_j - c_j(T_{d(j)})] q_j(p_j, m_j, h_j) - F(T_{d(j)}) \quad (4)$$

subject to

$$\sum_{j \in J_f} \left(\frac{1}{m_j} - \frac{1}{M_j} \right) q_j = 0 \quad (5)$$

and

$$\ln(m_j) = \ln(m_{0j}) + \delta^h \ln(h_j) + \delta^w \ln(w_j) + T_{d(j)}, \quad (6)$$

where marginal costs are given by

$$\ln(c_j) = \ln(c_{0j}) + \gamma T_{d(j)}. \quad (7)$$

Profits are the sum over vehicles of quantity multiplied by the difference between price and marginal costs ($c_j(T_{d(j)})$), less fixed costs ($F(T_{d(j)})$). Equation (5) is the regulatory constraint, which requires that the firm's sales-weighted fuel consumption rate (the reciprocal

of its fuel economy) equals the sales-weighted fuel consumption rate required by the standards. Note that the requirement is indexed by j because it depends on the vehicle’s class (car or truck) and footprint.¹⁵ The technology is chosen for each model, indexed by $m(j)$, during the design period. The fixed costs are incurred in the first period and must be recovered during the subsequent five years. Therefore, the sales q_j correspond to the sales over a five-year time period.

Equation (6) governs the trade-off between horsepower and fuel economy, given the vehicle’s exogenous weight and endogenous technology. This equation has the same form as that in [Klier and Linn \(2012\)](#), and is consistent with empirical analysis of the trade-off, such as in [Knittel \(2011\)](#).

The final equation links marginal costs to technology. Increasing technology by 0.01, which enables a 1 percent fuel economy increase according to Equation (6), increases marginal costs by γ percent.

In contrast to [Leard et al. \(2019\)](#), we do not model credit trading across manufacturers. During the first phase of the standards, few credits were traded ([Leard and McConnell 2017](#)), and only a handful of trades have been made. Credit trading likely has had modest effects on manufacturers’ decisions and compliance costs. Nonetheless, by leaving out this flexibility mechanism from our model, we expect our compliance cost estimates to be higher than those actually realized under the program.

3.3 Used Vehicle Market

The model includes the market for used vehicles in a reduced-form manner. Because new-vehicle fuel economy standards affect new-vehicle attributes and new and used vehicles are substitutes for one another, new-vehicle standards could affect demand for used vehicles ([Jacobsen 2013](#); [Bento et al. 2018](#); [Linn and Dou 2018](#); [Leard 2019](#)).

To account for this effect, we assume that used-vehicle supply is perfectly price-inelastic and that used-vehicle prices respond to used-vehicle demand changes caused by altering the standards. Therefore, if tighter standards reduce demand for new vehicles and increase demand for used vehicles, used-vehicle prices increase to restore equilibrium in the used-vehicle market. We assume an endogenous used-vehicle supply response and that used supply is completely inelastic. In practice, this involves adjusting the utility to the outside option to compute counterfactual equilibria.

¹⁵Other papers, such as [Jacobsen \(2013\)](#), include regulatory constraints that are nonlinear in vehicle sales, q_j . The linear constraint in this paper is a transformation of the constraint used in those papers.

The assumption of perfectly inelastic supply is broadly consistent with [Busse et al. \(2013\)](#), who report small effects of gasoline prices on equilibrium used-vehicle sales. Because households are both buyers and sellers of used vehicles, assuming perfectly inelastic supply likely has small welfare implications, despite the fact that the used-vehicle market includes several times more sales than does the new-vehicle market. [Jacobsen \(2013\)](#) finds that accounting for interactions between new and used vehicle markets has important implications for the distribution of welfare effects across demographic groups, since high-income households are more likely to buy new vehicles than are low-income households. For that reason, we are careful to interpret the estimated distributional effects of standards across demographic groups as pertaining only to new-vehicle buyers; the structure of the used-vehicle market likely has small effects on aggregate welfare estimates.

4 Estimation

Demand estimation follows the estimation strategy in [Leard et al. \(2019\)](#), with the exception that we add further heterogeneity in preferences for drive type and fuel type. The supply side uses observed attribute choices and first-order conditions to estimate the shadow costs of the regulation, marginal costs, and fixed costs.

4.1 Demand Estimation Strategy

We estimate household preference parameters using Equation (3) by following the estimation strategy in [Leard et al. \(2019\)](#). This strategy estimates the parameters in two stages. In the first stage, we regress the left-hand side of Equation (3) on interactions of vehicle characteristics with fixed effects for each of the 20 demographic groups, and we include vehicle by model year-specific fixed effects to absorb all mean utility terms. This yields estimates for the β_{kg} terms as well as mean utility fixed effects.

Assuming that the demographic-group specific utility from unobserved vehicle attributes is uncorrelated with utility from observed vehicle attributes allows us to estimate Equation (3) by including a fixed effect for each vehicle and market, δ_{jt} . This assumption is analogous to the assumption made in [Berry et al. \(1995\)](#) that consumer-specific utilities for unobserved product attributes are uncorrelated with utilities for observed product attributes. The fixed effects δ_{jt} control for unobserved vehicle attributes that vary across trims or engine configurations, such as seating material or the presence of a sunroof. Consequently,

the interaction terms are identified by preference variation across demographic groups, controlling for unobserved attribute variation across trims.¹⁶

In the second stage, we use the generalized method of moments (GMM) estimator from [Leard et al. \(2019\)](#) to estimate the mean utility preference parameters, $\bar{\beta}_k$. This approach is analogous to the standard approach in the literature (see [Goolsbee and Petrin \(2004\)](#) and [Train and Winston \(2007\)](#)) of recovering the mean marginal utilities for the vehicle characteristics in a second stage, in which the vehicle-market fixed effect is regressed on these attributes. Unlike most of the literature, the GMM estimation accounts for the endogeneity of vehicle prices, fuel economy, and horsepower by exploiting two main sources of variation. First, we instrument for vehicle prices using attributes that are adjusted only during major vehicle redesigns; specifically, the mean width, length, and height of vehicles sold by other manufacturers that belong to the same market segment, as well as the means of the same attributes of other vehicles sold by the same manufacturer but belonging to a different market segment. The use of these instruments is consistent with the supply side of the model, in which all vehicle attributes aside from prices, fuel economy, and horsepower are exogenous. The second source of variation arises from the fact that we observe many pairs of vehicles that are identical in all aspects with the exception that one has a higher performance engine than the other. That is, two such twins share a market, make, model, trim/series, fuel type, drive type, and body style but have different engine configurations. Consequently, twins typically have the same seating material and other features as one another, reducing concerns of omitted variables bias. See [Leard et al. \(2019\)](#) for additional details on the demand estimation.

4.2 Demand Estimation Results

We present the demand estimation results in a series of figures. In [Figure 3](#), we show the implied own-price elasticity of demand for each of the 20 demographic groups. The average own-price elasticity of demand is -4.02 and varies from about -5 to -2 across demographic groups. Lower-income households are much more price sensitive than higher-income households, a finding that is consistent with prior literature ([Leard et al. 2019](#);

¹⁶In the first stage, we allow for heterogeneity in preferences for four vehicle attributes: price, fuel economy, performance and footprint. We assume no heterogeneity in preferences for unobserved vehicle characteristics such as seating material. In the second stage, we address the endogeneity of vehicle price, fuel economy and performance, caused by the potential correlation between these characteristics and the unobserved attributes. This approach is consistent with exogeneity assumptions made in mixed logit specifications ([Berry et al. 1995; 2004; Goolsbee and Petrin 2004; Train and Winston 2007](#)). Moreover, because fixed effects control for mean sales by vehicle and year, first-stage preference parameters are identified by within-vehicle market shares and particularly low selling vehicles such as luxury sports cars do not affect the first-stage parameter estimates excessively.

Train and Winston 2007). In Panel A of Figure 4, we show the implied willingness to pay (WTP) for a 1 percent increase in fuel economy. This is computed based on the ratio of each demographic group’s fuel economy coefficient and its price coefficient. Lower-income households tend to have a lower WTP for fuel economy than higher-income households. This finding is consistent with Leard et al. (2019), although the figure shows a few exceptions, since lower-income rural groups tend to have a higher WTP than some higher-income urban groups.

In Panel B of Figure 4, we show WTP for horsepower. Household preferences for horsepower vary considerably across demographic groups, with some groups valuing horsepower gains very little and others valuing it considerably. The average WTP for a 1 percent increase in horsepower is \$41. This is lower than the estimates reported in Leard et al. (2017) and Leard et al. (2019), but it is within the range reported in Greene et al. (2018).

In Table 1, we show implied valuation ratios for fuel cost savings. The ratio equals the WTP for a 1 percent fuel economy increase divided by the present discounted value of the fuel cost savings caused by the fuel economy increase. The ratio represents how much households are willing to pay for a \$1 reduction in lifetime fuel costs. We use the same assumptions from Leard et al. (2019) on lifetime miles driven, private discount rates, and expectations of gasoline prices to compute the ratios. Full valuation that is consistent with a neoclassical model of decision-making predicts a valuation ratio of 1. Our estimates imply that most demographic groups undervalue fuel costs, with a mean ratio of 0.45. This value is similar to that reported in Leard et al. (2019), who estimate valuation ratios using similar data but a different identification strategy. The ratios vary from a low of about 0.20 to a high of 1.01, where lower-income households tend to value fuel costs less than higher-income households.

In Table 2, we show average WTP for fuel economy, valuation ratios, and WTP for horsepower, by manufacturer. These calculations are weighted averages, using purchase probabilities as weights, which vary by vehicle and manufacturer because of the heterogeneity in household preferences. WTP for fuel economy varies considerably across manufacturers. Daimler has a consumer base that has relatively high valuation of fuel economy, whereas Subaru sells vehicles to consumers that have relatively low valuation.

Figures A.2 and A.3 report a similar set of demand model validation exercises as those reported in Leard et al. (2019). Figure A.2 shows the predicted attribute means, by demographic group, against the observed sales-weighted averages. The predicted values fall

close to a 45-degree line, indicating that the model predicts accurately the attribute means by demographic group.

Next, we assess the model's ability to predict market shares out of sample. Panel A of Figure A.3 is a no-change forecast, which uses the observed brand-class market shares from the first market of our sample ($t = 2010$) to predict market shares from the final market of our sample ($t = 2015$; we aggregate market shares across vehicles to brand-class to account for vehicle entry and exit, which is exogenous in our model). Panel A plots the predicted market shares against the observed market shares in $t = 2015$. There is considerably less scatter if we use the demand model to predict 2015 market shares as in Panel B of Figure A.3, compared with the amount of scatter in Panel A. Thus, the demand model in Panel B outperforms the no-change forecast in Panel A.

4.3 Supply Estimation Strategy

We estimate the parameters of the supply model using first-order conditions of the manufacturer problem stated in Equation (4). The shadow price, trade-off between horsepower and fuel economy, and marginal costs are estimated via manufacturer pricing decisions and movement along the production possibility frontier in Equation (6). The fixed costs are identified by outward shifts of the frontier over time, and the assumption that the fixed costs rationalize observed technology choices.

Given the trade-off between horsepower and fuel economy in Equation (6), we have three sets of first-order conditions: price, fuel economy, and fuel-saving technology. For each vehicle j , these conditions are

$$\sum_{k \in J_f} \left[p_k - c_k - \lambda_f \left(\frac{1}{m_k} - \frac{1}{M_k} \right) \right] \frac{\partial q_k}{\partial p_j} + q_j = 0, \quad (8)$$

$$\sum_{k \in J_f} \left[p_k - c_k - \lambda_f \left(\frac{1}{m_k} - \frac{1}{M_k} \right) \right] \left(\frac{\partial q_k}{\partial m_j} + \frac{\partial q_k}{\partial h_j} \frac{\partial h_j}{\partial m_j} \right) + \lambda_f \frac{1}{(m_j)^2} q_j = 0, \quad (9)$$

$$\sum_{k \in J_f} \left[p_k - c_k - \lambda_f \left(\frac{1}{m_k} - \frac{1}{M_k} \right) \right] \left(\frac{\partial q_k}{\partial m_j} \frac{\partial m_j}{\partial T_{d(j)}} + \frac{\partial q_k}{\partial h_j} \frac{\partial h_j}{\partial T_{d(j)}} \right) - q_j \frac{\partial c_j}{\partial T_{d(j)}} - \frac{\partial F}{\partial T_{d(j)}} = 0. \quad (10)$$

The unknown parameters that we estimate are the J vehicle-specific marginal costs, c_k , the trade-off coefficient for horsepower and fuel economy in Equation (6) denoted by δ^h ,

manufacturer-specific shadow costs, λ_f , and the derivative of fixed costs with respect to technology. We estimate the parameters in a series of steps. In the first step, we initialize a guess for the marginal costs c_k . In the second step, given this guess, we rearrange first-order condition (9) such that the equation is linear in the manufacturer’s shadow costs, λ_f . We then estimate a regression to fit a value for λ_f that minimizes the squared distance of the linear equation. We do this separately for each set of manufacturer equations to obtain a manufacturer-specific shadow cost. In the third step, with these estimated shadow costs, we solve for marginal costs c_k in the first-order condition (8). These estimated marginal costs become our new guess of the marginal costs. We then repeat the second and third steps until the change in marginal costs across iterations is sufficiently small (within 0.01 percent).

The final steps involve substituting estimates for the marginal costs c_k and shadow costs λ_f in Equation (10) and estimating technology costs and fixed costs. We assume that each vehicle model has a unique fixed cost that varies linearly with the level of technology. We estimate these fixed-cost coefficients and the marginal-cost coefficient γ with a linear regression of Equation (10).

Finally, we estimate the function governing technology costs, $F(T_{d(j)})$. To simplify the estimation, we normalize each model’s technology endowment to zero and assume that fixed costs are a quadratic function of the technology change, $T_{d(j)}$. Under these assumptions, $\frac{\partial F}{\partial T_j} = 2\sigma T_{k(j)}$, where σ is a parameter to be estimated. Having estimated marginal costs c_k and shadow costs λ_f , Equation (10) is linear in technology and includes just the one unknown parameter σ . Similar to the procedure for estimating λ_f , we estimate a regression to fit a value for σ that minimizes the squared distance of the linear equation.

4.4 Supply Estimation Results

Tables 3 through 5 summarize the estimation results of the supply-side parameters. The first row of Table 3 reports each firm’s mean marginal costs. The variation across firms is intuitive; for example, Hyundai, which primarily sells low-price cars and crossovers, has lower marginal costs than Daimler, which sells luxury vehicles. The remaining rows in this table summarize the incentives that the standards create for the firms to raise fuel economy.

The average feebate is the sales-weighted mean of the quantity $\lambda_f \left(\frac{1}{m_k} - \frac{1}{M_k} \right)$ across each firm’s vehicles using observations from 2011 and the 2016 standards. Each vehicle’s feebate represents the incentives that the 2016 standards create to adjust prices of vehicles sold in 2011, and the average feebate is the incentive to increase the fuel economy of those vehicles (see Equations (8) through (10)). The feebate varies substantially across manufacturers. For

example, the average feebate is about 20 percent lower for Honda than for Fiat-Chrysler. Note that the cross-firm variation accounts for the footprint-based standards. Compared to uniform standards that do not depend on vehicle attributes, the footprint-based standards reduce cross-firm variation in the average feebate because firms selling larger vehicles are subject to lower fuel economy standards.

The final column of the table shows the fuel economy shadow cost, which is computed by applying the envelope theorem to the firm’s profit maximization problem; the shadow cost is equal to the derivative of the firm’s profits with respect to its fuel economy requirement. The shadow cost exhibits similar variation as the average feebate across firms.

The top panel of Table 4 shows the estimated trade-off between fuel economy and horsepower for cars and light trucks (i.e., the coefficient δ^h in Equation (6)). The coefficients are interpreted as elasticities and are larger in magnitude than those estimated in Klier and Linn (2016). We are using more recent data but otherwise the same methodology, which indicates that the trade-off between horsepower and fuel economy has become less severe over time.

The bottom panel of Table 4 reports the estimate of the efficiency coefficient in Equation (7). The coefficients are interpreted as elasticities, so the estimate in the first column means that increasing a car’s efficiency by 1 percent causes marginal costs to increase by 0.332 percent. The elasticity is about half as large for light trucks as it is for cars.

We obtain similar estimates using data from NHTSA’s technology model, which the agency uses to estimate benefits and costs of the regulations. In its model, firms minimize the cost of meeting the fuel economy regulations by choosing fuel-saving technologies for their individual vehicles. The model includes extensive detail on the costs and performance of the technologies, accounting for positive and negative interactions among technologies. The model output includes an estimate of each vehicle’s price and fuel economy. Because the agencies assume a constant markup over marginal costs, we can use their data to recover the efficiency coefficient. Doing so yields point estimates that lie within the range of the two estimates in the table. This similarity supports the validity of our approach to estimating the cost coefficient.

Table 5 provides context for the estimated fixed costs of technology adoption. Each column shows the costs of increasing fuel economy by the percentage indicated in the column heading. We compare the increase in marginal costs with the increase in average fixed costs—the ratio of the change in fixed costs to the vehicle’s sales. The marginal costs increase roughly linearly with the fuel economy improvement. However, average fixed costs

increase rapidly; they are typically lower than marginal cost increases for a 1 percent fuel economy increase, but are typically higher for 5 or 10 percent increases. This is reasonable because raising fuel economy by more than a few percentage points usually involves the simultaneous adoption of multiple fuel-saving technologies, which complicates the redesign process. Because the first phase of the standards required a 13 percent average fuel economy increase, the numbers in this table indicate that fixed costs play an important role in the overall welfare analysis.

5 Welfare Analysis

This section describes the two central scenarios that are simulated, reports results from those two scenarios, and presents results from additional scenarios.

5.1 Definition of Baseline and Central Policy Scenarios

In the baseline scenario, the standards are maintained at their 2012 levels for model years 2012–2016. Figure A.1 shows that the standards tightened by about 13 percent over these four years. We use 2012 rather than 2011, because 2011 was a transitional year and manufacturers had the option of complying with uniform or footprint-based standards.

The simulations are performed using the set of vehicles available in the market in 2015. We assume that the standards do not affect vehicle entry, exit, or any attributes other than price, fuel economy, and horsepower (all existing welfare analyses of the standards have made these assumptions implicitly).¹⁷ Under these assumptions, the baseline scenario and any other scenario include the same set of vehicles. Moreover, vehicle weight and the mean utility (ξ_{jt}) of each vehicle do not vary across scenarios. Therefore, if we define the set of vehicles in the simulations to include vehicles that were sold in 2011, we would have to impute weight and ξ_{jt} for each vehicle that enters the market after 2011. Because about 30 percent of models and 70 percent of vehicles that were sold in 2015 enter after 2011, this would require a substantial amount of imputation. We avoid having to impute ξ_{jt} by using data from the last year, 2015. For these vehicles, we use the observed fuel prices, vehicle characteristics, fuel economy requirements, and estimated marginal costs, fixed costs, and technology parameters estimated using the methodology described in the previous section.

¹⁷Wollmann (2018) argues that bankruptcy policy could affect product entry. In principle, fuel economy standards could also affect entry and exit, but for computational reasons we assume that entry and exit decisions are exogenous to the standards. Moreover, we assume that standards do not affect the frequency with which firms redesign their vehicles, which is necessary because of the length of our sample period. We are not aware of evidence that standards have affected the frequency of redesigns.

Recall that the model is static and includes two periods. During the first period, the regulator announces the level of the standards and the firm chooses a level of technology. Because models are typically redesigned every five years, we can conceive of period 1 as representing the year 2011, when the standards were determined, and period 2 as representing the steady state in 2016, which is five years later. In the baseline scenario, the standards correspond to the values of the standards that were actually applied to model year 2012, depicted by the orange marks in Figure A.1. Note that we do not model the transitional dynamics of compliance with the standards, which is the same as most of the literature (e.g., (Whitefoot et al. 2017)). However, unlike the rest of the literature, we interpret the steady state as occurring five years after the standards are announced. That allows us to account for technology adoption in a scenario in which standards do not tighten over time. Note that the EPA and NHTSA welfare analysis models transitional dynamics and assumes that standards do not change after 2016.¹⁸

The equilibrium for the baseline scenario is found iteratively.¹⁹ On average, firms have to increase fuel economy by about 1 mpg from 2011 levels to comply with the 2012 standards. We determine the initial conditions for the simulation by assuming that manufacturers use technology adoption and trade off horsepower for fuel economy in equal amounts, and that they do not adjust prices from observed 2011 levels. These assumptions determine initial conditions and not simulated outcomes, and they are consistent with Klier and Linn (2016) and Reynaert (2019). Starting from that initial condition, each firm chooses technology (T_{mt}), fuel economy, and price to maximize profits. Each firm’s decisions are the best response given choices of all other firms in the market, and we loop over firms and iterate across the entire market until simulated choices converge to within the chosen tolerance.²⁰

¹⁸This difference between our analysis and that of EPA and NHTSA probably has little effect on the welfare conclusions. Their analysis includes banking and borrowing of credits over time as well as credit trading across firms, as allowed under the regulations. These provisions reduce compliance costs, relative to regulations that do not include them. However, there have been few cross-firm credit trades (Leard and McConnell 2017), suggesting that this provision has not affected compliance costs substantially. Between 2012 and 2016, most firms complied on average, and a few firms overcomplied. Because there have been so few cross-firm trades, it is likely that the overcomplying firms are banking credits to use for their own compliance, in anticipation of tighter standards in future years. A welfare analysis of the 2012–2016 standards should not include this overcompliance, since it is caused by standards that apply in later years. For this reason, in the simulations we assume that all firms exactly comply with the standards.

¹⁹Convergence of the model is not guaranteed, and in fact for a minority of vehicles the model may cycle among two or more sets of price and attributes. These situations typically affect lower-selling vehicles and vehicles sold by small manufacturers.

²⁰Specifically, the tolerance is set such that each firm’s average fuel consumption rate is within 0.1 percent of its standard, and the maximum price change across iterations across vehicles in the market is 1 percent. For a few firms, the model does not converge after 50 iterations, in which case we use the simulated technology and attributes from the final iteration. Achieving convergence is especially challenging for Jaguar–Land

The central policy scenario is the same as the baseline, except that firms must achieve the 2016 standards rather than the 2012 standards. Welfare effects are calculated by comparing equilibriums in the baseline and central policy scenarios. The welfare analysis includes changes in consumer welfare, manufacturer profits, and the global value of avoided carbon dioxide emissions.

To facilitate comparison with the literature and the agencies' welfare analysis, we calculate consumer welfare in two separate ways. First, for consistency with the literature, we assume that consumers do not make mistakes, and that the estimated undervaluation represents a hidden cost of choosing a technology with high fuel economy. As noted in the Introduction, hidden costs may reflect consumers' perceptions of new technologies. We adopt the closed-form formula from [Small and Rosen \(1981\)](#) to compute consumer surplus changes. Second, for consistency with the agencies we follow [Allcott \(2013\)](#) and [Train \(2015\)](#) to adjust for consumer mistakes. For this calculation, we assume a 3 percent real discount rate for consistency with [EPA \(2010\)](#), and we use assumptions on vehicle scrappage and mileage from [Leard et al. \(2019\)](#).

For both sets of welfare calculations, we assume a rebound effect of 10 percent (which is the same as the assumption the agencies make for their analysis of the 2012--2016 standards). For the calculations that use preference parameters, we do not include the benefits from additional driving because the demand model implicitly incorporates these welfare changes. Households compare across vehicles the utility from owning and driving the vehicle. Increases in fuel economy lower household fuel costs but also let households drive more miles for the same fuel costs. These components are implicitly accounted for in the definition of utility. Because the calculations that correct for consumer mistakes do not include the benefits from additional driving, we use the conventional Harberger approximation similarly (EPA and NHTSA do the same).

For firms, we compute profits by including fixed costs. Finally, we compute the social value of avoided emissions using estimates from [EPA et al. \(2016\)](#) of the social cost of carbon dioxide emissions, using the same 3 percent discount rate that we use for the fuel savings calculations.

Note that we do not include refueling time costs and external costs of fine particulate matter, accidents, noise, congestion, and energy security. Estimating these costs and benefits requires many assumptions that lie well outside the scope of the paper. EPA and NHTSA

Rover, which sells a small number of high-end vehicles and faces particularly stringent standards, but this company accounts for less than 0.1 percent of total sales in the market.

estimate combined net welfare benefits of \$7 billion for these categories. Therefore, including them would strengthen our findings that the standards have increased social welfare.

5.2 Results for Baseline and Central Policy Scenarios

The simulated mean attributes for the baseline and central policy scenarios appear in the row headings of Table 6. The fractional technology improvement in the baseline scenario implies an average annual efficiency improvement of about 2.5 percent, which is roughly consistent with historical averages. The tighter standards in the central policy scenario increase the incentive to adopt technology, causing more technology adoption than in the baseline scenario. This result is consistent with empirical analysis in Klier and Linn (2016).

Sales-weighted average fuel economy is about 14 percent higher under the 2016 standards, which corresponds to a 13 percent decline in average fuel consumption rate required by the standards.²¹ Performance (as measured by the log of the ratio of horsepower to weight) is about 0.37 log points lower in the central policy scenario than in the baseline scenario, and it is within a few percentage points of the observed 2011 and 2015 values (see Panel B of Figure 2). Thus, the simulations imply that in the baseline, horsepower is about 30 percent higher than the observed 2011 value, and that in the central policy, horsepower is about equal to the observed 2011 value. In other words, tighter standards do not decrease horsepower over time in absolute terms, but they decrease horsepower relative to the counterfactual of weaker standards.

Mean vehicle prices are slightly higher in the central policy scenario. The similarity reflects two opposing effects on vehicle prices, which roughly cancel each other. On the one hand, the standards induce technology adoption, which raises marginal costs and prices. On the other hand, the standards reduce consumer demand for new vehicles (which Table 7 shows), reducing equilibrium vehicle prices.

The magnitudes of the technology, horsepower, and fuel economy changes indicate that manufacturers comply with the standards largely by trading off horsepower for fuel economy. The difference in technology across the two scenarios means that the additional technology adoption reduced the average fuel consumption rate by about 4 percent. The remaining 9 percent reduction is caused by trading off horsepower for fuel economy. The main explanation for this pattern is that adopting technology gets increasingly expensive, as Table 5 shows. More specifically, manufacturers have three options for complying with tighter

²¹Because the fuel consumption rate is the reciprocal of fuel economy, the percentage change of the sales-weighted average fuel economy does not equal precisely the negative of the sales-weighted percentage change in the fuel consumption rate, across the two scenarios.

standards: adjusting vehicle prices, adding technology, and trading off fuel economy for horsepower. In equilibrium, manufacturers equate the marginal costs of the three options. Starting from the 2011 equilibrium, it turns out that the cost of trading off horsepower for fuel economy is roughly linear, based on the estimated preference parameters and trade-off coefficient in Equation (6). However, average fixed costs increase faster than linearly with the fuel economy increase, making it increasingly expensive to comply by adopting technology. Consequently, in equilibrium, firms rely more on attribute trade-offs than technology adoption for compliance. In contrast, it turns out that the cost of adjusting vehicle prices for compliance increases quickly with the price change, and firms make only small price adjustments, usually within a few percentage points.

The estimated welfare effects of the 2016 standards appear in Table 7. The first row of the table shows the difference in consumer surplus between the central policy and the baseline scenarios using the two computations described in the previous subsection. Using the estimated preference parameters to compute consumer surplus in the first column, we find that consumer surplus decreases by about \$9.6 billion. The 2016 standards reduce consumer surplus because the decrease in welfare from lower performance outweighs the increase in welfare from higher fuel economy (Table 6 shows that vehicle prices are approximately the same in the two scenarios). The value of the fuel cost savings is about \$30 billion (not shown in the table), and the mean valuation ratio of 0.5 implies that consumers value only about half of those savings. Consequently, when we correct for consumer mistakes in the second column, consumer surplus increases by about \$15 billion.

The second row shows the change in producer welfare, as measured by profits. The 2016 standards reduce profits by about \$5.8 billion, which includes two effects. First, tighter standards cause firms to adopt more technology, which raises marginal costs, but manufacturers are unable to pass through the full cost increase to consumers.²² The resulting lower markup reduces profits, which contrasts with the EPA and NHTSA benefit-cost analysis, which assumes full pass-through of compliance costs to vehicle prices. Second, the additional technology adoption shown in Table 6 translates to an additional \$3 billion per year across the market.

The bottom two rows report the net social benefits and the private welfare cost per ton of carbon dioxide abated. Using the estimated preference parameters to evaluate consumer

²²This result differs from Leard et al. (2019), who show that manufacturers can pass through the full cost increase. However, in their model, horsepower is exogenous, and firms cannot trade off horsepower for fuel economy. In our analysis, because firms reduce horsepower, demand for new vehicles falls, preventing manufacturers from passing through the full cost increase. This difference demonstrates the importance of endogenizing horsepower for accurate welfare estimates.

surplus in the first column, we find that tighter standards reduce social welfare by \$11.4 billion. This cost exceeds the social value of the avoided emissions. The implied cost per metric ton of carbon dioxide emissions abated, which is calculated as the ratio of the private welfare cost and the tons of abated, equals \$187 per metric ton and exceeds the estimated social cost of carbon dioxide.

However, the standards increase social welfare if we include the full value of the fuel cost savings. In this case, the tighter standards reduce private welfare by \$470 million, which is outweighed by the value of the avoided carbon dioxide emissions. As a result, social welfare increases by \$3.6 billion. The cost per metric ton of carbon dioxide emissions abated equals \$6, which is less than the estimated social cost of carbon dioxide including global damages.²³

Table A.1 reports attribute and welfare changes by demographic group. Figure 4 shows that low-income groups have lower WTP for fuel economy and horsepower than do high-income groups. It turns out that manufacturers find it less costly, in forgone profits, to trade off horsepower for fuel economy for vehicles that tend to be purchased by low-income consumers. Consequently, those consumers experience larger fuel economy increases and horsepower decreases than do high-income groups. When using the full value of the fuel cost savings to compute welfare changes in the right-most column, we find that welfare increases for low-income groups and decreases for high-income groups; thus, the standards are progressive. This result contrasts with Leard et al. (2019), who find that standards are regressive if horsepower is held exogenous. The difference arises from the fact that high-income groups have higher demand for horsepower than do low-income groups.

The effects of the standards on manufacturers appear in Table A.2. The standards reduce per vehicle profits for Ford and Chrysler. Compared with other manufacturers, Ford and Chrysler have higher estimated fixed costs of technology adoption and face larger costs of trading off horsepower for fuel economy because their customers have higher WTP for horsepower than do other customers (see Table 2). The tighter standards reduce consumer demand for vehicles sold by these manufacturers, which increases demand for vehicles sold by other manufacturers. Consequently, profits for some manufacturers increase, including Toyota and Honda.

²³Recall that these estimates do not include savings of refueling costs and certain other costs and benefits that EPA and NHTSA include in their analysis. Including those terms would cause us to estimate that the 2012–2016 standards caused welfare changes of –\$4 billion to \$11 billion, depending on the computation of fuel cost savings.

5.3 Other Results

Table 8 shows results from additional scenarios, described in the column headings. Each column reports the differences in attribute changes between the policy scenario and the corresponding baseline, with the first column repeating the results from the central policy and baseline scenarios for comparison purposes.

The Introduction noted that fuel prices dropped after the agencies finalized the 2012–2016 standards and performed their welfare analysis. Lower fuel prices reduce consumer demand for fuel economy, making it more costly for manufacturers to comply with standards.²⁴ Lower fuel prices also reduce the value of the fuel cost savings, lowering the benefits. The average gasoline price in 2015 was about \$1 per gallon lower than the average price in 2011, corresponding to a 30 percent decline. In column 2 we resimulate the baseline and policy scenarios using observed 2011 fuel prices (i.e., corresponding to the time at which the agencies finalized the standards) rather than the observed 2015 fuel prices. With the higher fuel prices, net social benefits are substantially higher than with lower fuel prices.

The stated objective of the footprint-based standards was to reduce incentives for manufacturers to comply by selling smaller vehicles. NHTSA argued that smaller vehicles are less safe in multivehicle collisions. Flat standards, in which a vehicle’s fuel economy requirement depends on its class but not any other attribute, incentivizes firms to reduce prices of smaller vehicles, which have higher fuel economy than do larger vehicles. The footprint-based standards reduce this incentive to lower prices for small vehicles. Column 3 of Table 8 shows that with uniform standards, the 2016 standards cause a larger decrease in average footprint than in column 1. However, the magnitude of this difference is small in percentage terms, reflecting the fact that manufacturers rely little on vehicle price changes to increase average fuel economy. Quantifying the welfare implications of this difference lies outside the scope of the paper, since it requires modeling the interactions between new vehicles and existing vehicles in multivehicle collisions.

Coinciding with the increase in stringency that occurred during the 2012–2016 phase of the standards was a change in structure: allowing manufacturers to comply with an average standard across their car and light truck fleets, rather than with separate standards for

²⁴More specifically, lower fuel prices increase the cost of trading off horsepower for fuel economy and decrease the ability of the manufacturer to pass through cost increases to consumers. The adoption of the footprint-based standards mitigates a third reason why lower gas prices may increase costs. With uniform standards, a drop in gasoline prices causes consumers to substitute to vehicles with lower fuel economy, making it more challenging for the manufacturer to comply with the standards. With footprint-based standards, however, vehicles with lower fuel economy tend to be larger and therefore have lower fuel economy requirements. Leard et al. (2017) show that the footprint-based standards reduce but do not eliminate the importance of this effect.

each class, and the replacement of flat standards with footprint-based standards. Allowing manufacturers to comply with a single standard should reduce welfare costs of the standards because manufacturers can take advantage of compliance cost variation across classes. That is, if it is less costly for a manufacturer to comply with a standard for one class than for the other, a single standard across the classes allows the manufacturer to overcomply for the low-cost class and undercomply with the high-cost class, reducing total costs. Column 4 of Table 8 shows that this change in the standards had little effect on overall welfare costs.

Finally, we show the implications of relaxing two assumptions commonly made in the literature and by the regulatory agencies: that horsepower is fixed in all scenarios, and that manufacturers do not increase horsepower in the baseline scenario. If we assume that horsepower is exogenous in column 6, consumer surplus is substantially higher than in column 1 because consumers do not suffer the welfare losses of lower horsepower. Manufacturers' costs are also higher, largely because fixed costs roughly double. Columns 1 and 6 demonstrate the importance of endogenizing horsepower in the model and including fixed costs: consumer benefits are \$14 billion higher, and profits are \$3 billion lower if horsepower is exogenous and welfare analysis ignores performance increases in the baseline. Overall, these two assumptions overstate net benefits by about \$11 billion.

6 Conclusions

The US market for new passenger vehicles is dynamic and complex. Firms spend billions of dollars per year designing and producing highly differentiated vehicles; these investments steadily improve vehicles over time.

This situation presents a challenge to analyzing the welfare effects of fuel economy and GHG standards. Despite the numerous studies of consumer demand for new vehicles, the literature on the standards has largely ignored three key supply-side factors: trade-offs between performance and fuel economy, technology adoption that would occur in the absence of the standards, and fixed costs of adopting technology. Moreover, the literature has ignored the possibility that consumers systematically make mistakes when evaluating differences in fuel costs across vehicles.

In this paper, we analyze the first phase of combined fuel economy and GHG standards, which covered model years 2012–2016, using a new equilibrium model of the passenger vehicles market that addresses these supply and behavioral factors. In the demand component of the model, consumers choose among a highly differentiated set of vehicles. Consumer preferences for price and attributes vary across demographic groups. Using

data from a household survey of recent vehicle buyers, we estimate substantial preference heterogeneity across demographic groups. On average, consumers undervalue fuel cost savings and prefer horsepower to fuel economy.

In the supply component of the model, manufacturers choose vehicle prices, fuel economy, horsepower, and fuel-saving technology, where technology adoption increases fixed costs as well as marginal costs. We estimate all cost parameters using first-order conditions to a firm’s profit maximization problem and observed attribute choices.

We estimate benefits and costs of the standards by comparing two scenarios that differ according to the level of standards that apply in model year 2016. In the first case, we assume that standards are maintained at their 2012 levels for model years 2012–2016. In the second, we use the actual 2016 standards for model year 2016. The comparison accounts for technology adoption and horsepower changes that occur in the 2012 scenario, differing from EPA and NHTSA welfare analysis as well as estimates in the literature in which horsepower is endogenous (i.e., (Klier and Linn 2012) and (Whitefoot et al. 2017)).

After accounting for the full value of the fuel cost savings, the 2012–2016 standards reduced GHG emissions at a cost of \$6 per metric ton of carbon dioxide, implying that the standards increased social welfare. The largely unexpected decline in fuel prices that began in 2014 substantially reduces the net benefits of the standards. The shift from uniform to footprint-based standards had a small effect on overall welfare and average footprint.

As noted in the Introduction, our model focuses on the new-vehicle market to address the shortcomings in the literature. We follow EPA (2012) to incorporate the effects of standards on vehicle utilization. Accounting for scrappage would not overturn the main conclusion that the standards have increased social welfare.

The net benefits of the standards are substantially lower than the agencies predicted. This is partly due to the drop in fuel prices that occurred after they performed their welfare analysis. Failing to account for technology adoption and performance improvements that would have occurred if standards had not tightened appears to explain much of the difference between our results and theirs. Our results suggest that tightening standards beyond 2016 levels also increases social welfare because our estimated compliance costs are less than the social cost of carbon dioxide. During our period of analysis, plug-in vehicles accounted for less than 1 percent of the market, and for simplicity our model allows for limited preference heterogeneity across demographic groups for plug-in vehicles. However, the market share of these vehicles was above 2 percent in 2019 and is expected to continue rising. Therefore,

analyzing welfare effects of future standards requires a more detailed treatment of plug-in vehicles, which we leave for future work.

References

- Akerberg, D. A. and M. Rysman (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics* 36(4), 1–19.
- Allcott, H. (2013). The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3), 30–66.
- Allcott, H. and M. Greenstone (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives* 14(5), 3–28.
- Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96, 779–795.
- Bento, A., K. Roth, and Y. Zuo (2018). Vehicle lifetime trends and scrappage behavior in the US used car market. *The Energy Journal* 39(1).
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy* 112(1), 68–105.
- Blonigen, B. A., C. R. Knittel, and A. Soderbery (2017). Keeping it fresh: Strategic product redesigns and welfare. *International Journal of Industrial Organization* 53, 170–214.
- Busse, M., C. Knittel, and F. Zettelmeyer (2013). Are consumers myopic? Evidence from new and used car purchases. *American Economic Review* 103(1), 220–256.
- EPA (2010). Regulatory Impact Analysis: Final Rulemaking to Establish Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. *EPA-420-R-10-009 April 2010*.
- EPA (2012). Regulatory Impact Analysis: Final Rulemaking for 2017–2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. *EPA-420-R-12-016 August 2012*.
- EPA, CARB, and NHTSA (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*.
- EPA and NHTSA (2018). Preliminary Regulatory Impact Analysis for The Safer Affordable Fuel-Efficient (SAFE) Vehicles Rule for Model Year 2021–2026 Passenger Cars and Light Trucks.
- Gillingham, K., S. Houde, and A. Van Benthem (2019). Consumer myopia in vehicle purchases: Evidence from a natural experiment. *NBER Working Paper No. 25845. Cambridge, MA*:


National Bureau of Economic Research.

- Goldberg, P. K. (1998). The effects of the corporate average fuel efficiency standards in the US. *The Journal of Industrial Economics* 46(1), 1–33.
- Goolsbee, A. and A. Petrin (2004). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica* 72(2), 351–381.
- Greene, D., A. Hossain, J. Hofmann, G. Helfand, and R. Beach (2018). Consumer willingness to pay for vehicle attributes: What do we know? *Transportation Research Part A: Policy and Practice* 118, 258–279.
- Houde, S. and C. A. Spurlock (2015). Do energy efficiency standards improve quality? Evidence from a revealed preference approach.
- Ito, K. and J. M. Sallee (2018). The economics of attribute-based regulation: Theory and evidence from fuel economy standards. *Review of Economics and Statistics* 100(2), 319–336.
- Jacobsen, M. (2013). Evaluating U.S. fuel economy standards in a model with producer and consumer heterogeneity. *American Economic Journal: Economic Policy* 5(2).
- Jacobsen, M. and A. van Benthem (2015). Vehicle scrappage and gasoline policy. *American Economic Review* 105(3), 1312–1338.
- 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). The effect of vehicle fuel economy standards on technology adoption. *Journal of Public Economics* 133, 41–63.
- Knittel, C. (2011). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review* 101(7), 3368–3399.
- Knittel, C. R. and K. Metaxoglou (2014). Estimation of random-coefficient demand models: Two empiricists’ perspective. *Review of Economics and Statistics* 96(1), 34–59.
- Leard, B. (2019). Estimating consumer substitution between new and used passenger vehicles. *Working Paper 19-02. Washington, DC: Resources for the Future.*
- Leard, B., J. Linn, and K. Springel (2019). Pass-through and welfare effects of regulations that affect product attributes. *Working Paper 19-07. Washington, DC: Resources for the Future.*
- Leard, B., J. Linn, and Y. Zhou (2017). How much do consumers value fuel economy and performance? Evidence from technology adoption. *Report. Washington, DC: Resources for the Future.*
- Leard, B. and V. McConnell (2017). New markets for pollution and energy efficiency: Credit trading under automobile greenhouse gas and fuel economy standards. *Review of Environmental Economics and Policy* 11(2), 207–226.
- Linn, J. and X. Dou (2018). How do US passenger vehicle fuel economy standards affect purchases of new and used vehicles? *Working Paper. Washington, DC: Resources for the Future.*

- MacKenzie, D. and J. B. Heywood (2015). Quantifying efficiency technology improvements in us cars from 1975–2009. *Applied Energy* 157, 918–928.
- National Research Council (2015). *Cost, effectiveness, and deployment of fuel economy technologies for light-duty vehicles*. Washington DC: National Academies Press.
- Reynaert, M. (2019). Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market. *Working paper*.
- Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica* 49(1), 105.
- Train, K. (2015). Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples. *Journal of Choice Modelling* 16, 15–22.
- Train, K. E. and C. Winston (2007). Vehicle choice behavior and the declining market share of U.S. automakers. *International Economic Review* 48, 1469–1496.
- Whitefoot, K. S., M. L. Fowlie, and S. J. Skerlos (2017). Compliance by design: Influence of acceleration tradeoffs on CO₂ emissions and costs of fuel economy and greenhouse gas regulations. *Environmental Science and Technology* 51(18), 10307–10315.
- Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review* 108(6), 1364–1406.

Figures


Figure 1: US Fuel Economy and Greenhouse Gas Standards, Historical and Projected



Notes: The figure shows the actual and projected fuel economy standards in terms of fuel consumption (gallons per 100 miles) for passenger cars and light trucks, respectively.


Source: NHTSA, CAFE Final Rules, <http://www.nhtsa.gov/fuel-economy>

Figure 2: Trends in Vehicle Attributes, 2010–2015




Notes: The figure shows the sales-weighted average attribute of all new cars and light trucks sold in the United States between 2010 and 2015. Performance is the log of the ratio of horsepower to weight.

Figure 3: Own-Price Elasticity of Demand by Demographic Group



Notes: Each bar shows the own-price elasticity of demand for the indicated demographic group. The estimates are computed from the estimated demand model coefficients, and all estimates are weighted across vehicles and markets using predicted market shares as weights.

Figure 4: Willingness to Pay for Fuel Economy and Performance by Demographic Group



Notes: Panel A reports the WTP for a 1 percent fuel economy increase, and Panel B reports the WTP for a 1 percent performance increase. Each bar shows the estimate for the indicated demographic group. The estimates are computed from the estimated demand model coefficients, and all estimates are weighted across vehicles and markets using predicted market shares as weights.

Tables

Table 1: Estimated Valuation Ratios by Demographic Group

| Panel A: Rural | | |
|----------------|----------|-----------|
| Income | Age < 45 | Age >= 45 |
| < 44k | 0.30 | 0.20 |
| 44k – 91k | 0.36 | 0.35 |
| 91k – 123k | 0.60 | 0.36 |
| 123k – 185k | 0.62 | 0.52 |
| > 185k | 0.74 | 0.99 |
| Panel B: Urban | | |
| Income | Age < 45 | Age >= 45 |
| < 44k | 0.22 | 0.227 |
| 44k – 91k | 0.35 | 0.33 |
| 91k – 123k | 0.24 | 0.42 |
| 123k – 185k | 0.37 | 0.57 |
| > 185k | 0.83 | 1.01 |

Notes: The table reports the valuation ratio for each demographic group using the WTP estimates from Figure 4.

Table 2: Own-Price Elasticity and Willingness to Pay by Manufacturer

| Firm | Own-price elasticity of demand | WTP for 1 percent fuel economy increase | Estimated valuation ratio | WTP for 1 percent horsepower increase |
|---------------|--------------------------------|---|---------------------------|---------------------------------------|
| GM | –3.86 | 156.17 | 0.63 | 49.68 |
| Ford | –3.84 | 123.14 | 0.49 | 43.47 |
| Toyota | –3.61 | 122.39 | 0.49 | 46.86 |
| Honda | –3.76 | 100.76 | 0.40 | 38.97 |
| Hyundai | –3.73 | 75.85 | 0.30 | 33.87 |
| Fiat-Chrysler | –3.80 | 134.26 | 0.54 | 42.57 |
| Nissan | –3.79 | 108.89 | 0.44 | 40.48 |
| Volkswagen | –3.86 | 132.51 | 0.53 | 50.36 |
| BMW | –3.94 | 190.66 | 0.77 | 64.98 |
| Subaru | –3.70 | 99.91 | 0.40 | 38.46 |
| Daimler | –3.92 | 227.66 | 0.91 | 68.86 |
| Other | –4.01 | 199.61 | 0.80 | 56.94 |

Notes: The table reports the average own-price elasticity of demand, willingness to pay for fuel economy, valuation ratios, and willingness to pay for horsepower by firm. All values are a weighted average using predicted sales for each vehicle and demographic group as weights.

Table 3: Estimated Marginal Costs and Feebate by Manufacturer

| Firm | Average marginal costs (2015\$/vehicle) | Average feebate (2015\$/vehicle) | Standard deviation of feebate (2015\$/vehicle) | Fuel economy shadow cost (\$ per mpg per vehicle) |
|---------------|--|-------------------------------------|---|---|
| GM | 26,864 | 722 | 491 | -144 |
| Ford | 24,524 | 687 | 632 | -121 |
| Toyota | 25,266 | 691 | 952 | -204 |
| Honda | 22,856 | 634 | 547 | -177 |
| Hyundai | 19,924 | 475 | 422 | -173 |
| Fiat-Chrysler | 23,897 | 802 | 383 | -119 |
| Nissan | 23,202 | 718 | 808 | -163 |
| Volkswagen | 29,181 | 765 | 633 | -184 |
| BMW | 41,165 | 1741 | 1134 | -261 |
| Subaru | 22,789 | 1093 | 458 | -161 |
| Daimler | 42,555 | 1887 | 980 | -203 |
| Other | 31,673 | 1237 | 253 | -87 |

Notes: The table reports the estimated marginal costs and feebate by firm. All values are a weighted average using predicted sales for each vehicle and demographic group as weights.

Table 4: Estimated Fuel Economy-Horsepower Trade-off and Marginal Cost Function

| | Dependent variable is log fuel economy | |
|----------------------------------|--|-------------------|
| | Cars | Trucks |
| Log horsepower | -0.239 (0.054) | -0.218 (0.047) |
| N | 2,508 | 2,901 |
| R-squared | 0.996 | 0.990 |
| | Dependent variable is log marginal costs | |
| | Cars | Trucks |
| Efficiency | 0.332 (0.082) | 0.173 (0.075) |
| N | 2,508 | 2,901 |
| R-squared | 0.988 | 0.978 |
| Change in marginal costs per mpg | 266 (66) | 234 (102) |

Table 5: Changes in Marginal Costs and Average Fixed Costs from Raising Fuel Economy

| | Percentage efficiency increase | | |
|------------------------------|--------------------------------|--------------------------------|---------------------------------|
| | 1 | 5 | 10 |
| Marginal costs (2015\$) | 60.76 (30.71) {53.34} | 299.48 (151.65) {263.04} | 588.65 (298.80) {516.13} |
| Average fixed costs (2015\$) | 5.58 (24.27) {1.77} | 139.47 (606.84) {44.15} | 557.87 (2427.34) {176.60} |

Notes: Each cell reports the mean with standard deviation in parentheses and the median in brackets.

Table 6: Simulated Efficiency, Vehicle Attributes, and Transaction Price

| | 2012 standards (baseline) | 2016 standards (central policy) |
|--|------------------------------|------------------------------------|
| Efficiency (fractional improvement since 2011) | 0.125 | 0.161 |
| Fuel economy (miles per gallon) | 23.82 | 27.19 |
| Log (horsepower/weight) | -2.50 | -2.87 |
| Transaction price (2015\$) | 33,580 | 33,810 |
| Footprint (square feet) | 49.71 | 49.70 |

Notes: The table reports average efficiency, vehicle attributes, and transaction prices for the two scenarios described in the row headings. In column 1, the standards are fixed at 2012 levels from model years 2011 through 2016. In column 2, the standards increase to 2016 levels.

Table 7: Estimated Welfare Changes: Difference between 2016 and 2012 Standards

| | Welfare changes (billion 2015\$) | |
|---------------------------------------|--|---|
| | Preference parameter for fuel cost savings | Present discount value of fuel cost savings |
| Consumers | −9.65 | 5.29 |
| Firms | −5.76 | −5.76 |
| Carbon emissions | 4.05 | 4.05 |
| Total social welfare | −11.35 | 3.59 |
| Private welfare cost per tonne abated | 187 | 6 |

Notes: The table reports the estimated costs and benefits of the 2016 standards. Welfare changes are disaggregated by consumers and producers. We compute changes in consumer surplus in two ways. In column 1, we use estimated preference parameters. Column 2 augments consumer surplus changes with the difference between the present value of fuel cost savings and the valuation from column 1. Welfare changes for manufacturers are the change in firm profits. Welfare changes due to changes in carbon emissions are calculated by multiplying the change in lifetime emissions of each vehicle by the social cost of carbon. Total social welfare is the sum of private welfare changes and change in social welfare due to reduced carbon emissions. The implied cost per metric ton of carbon emissions abated is the ratio of the private welfare cost and the metric tons of carbon emissions abated.

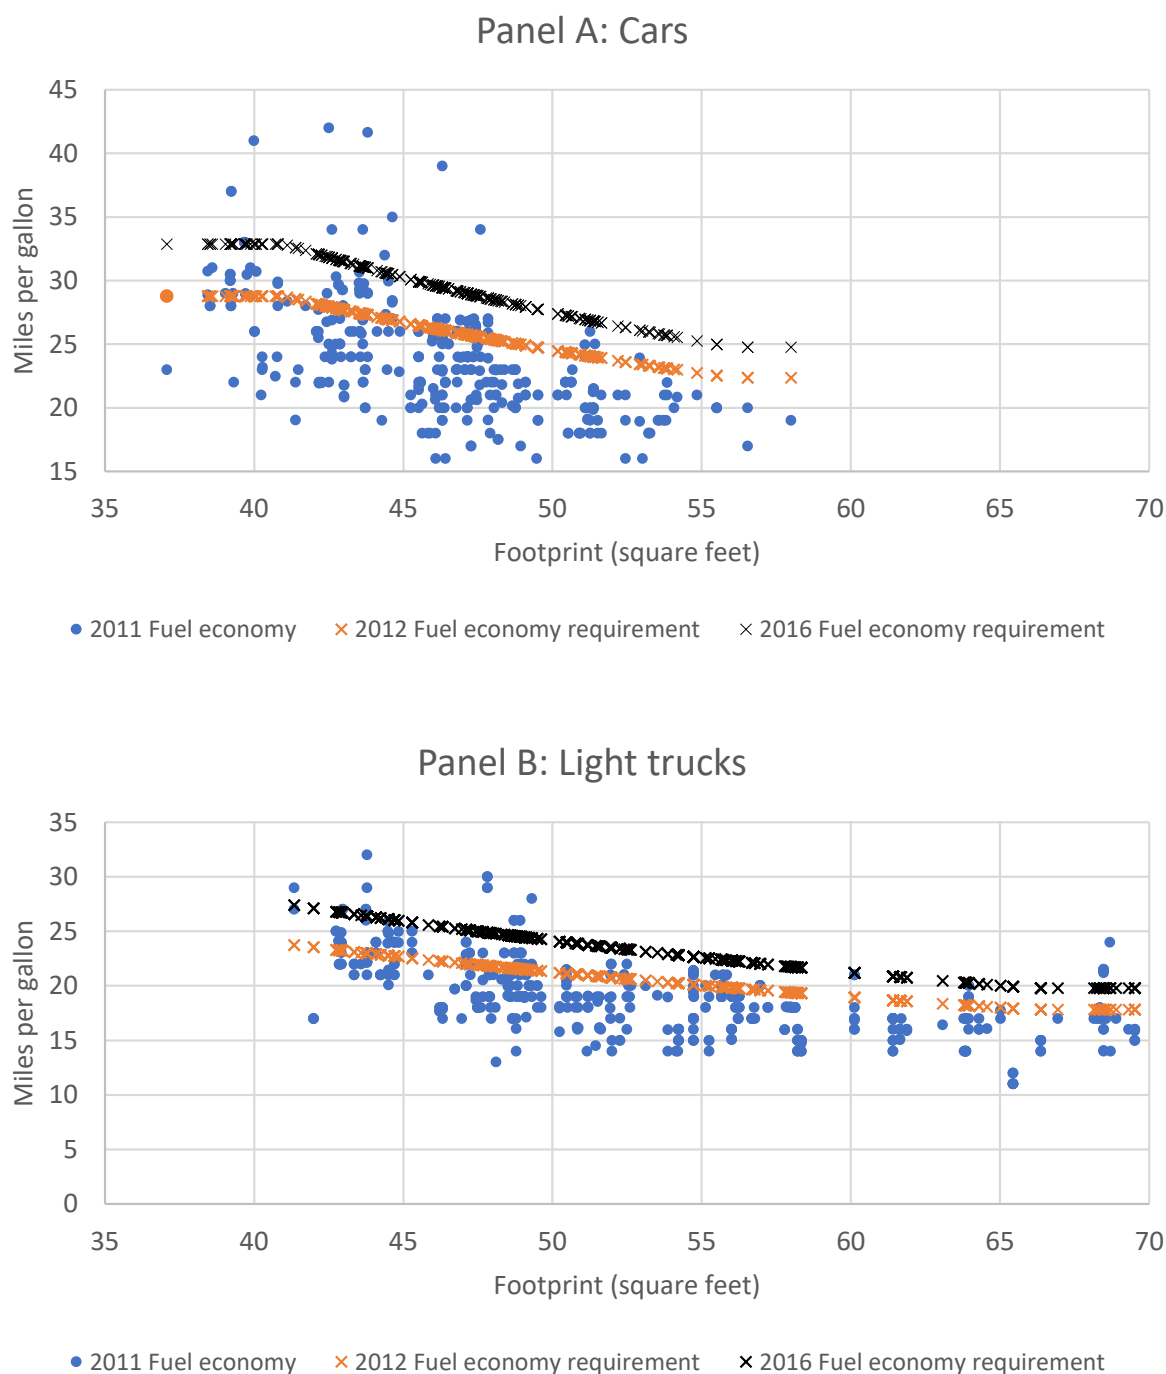
Table 8: Alternative Assumptions on Standards, Fuel Prices and Used Vehicle Supply

| | Central policy | 2011 fuel prices | Flat standards | No car/truck averaging | EPA/NHTSA fuel price and SCC assumptions | Horsepower exogenous and static baseline |
|--|-------------------|---------------------|-------------------|------------------------------|---|---|
| Panel A: Attribute changes (2016 vs. 2012 standards) | | | | | | |
| Fuel economy (mpg) | 3.36 | 3.22 | 3.44 | 3.12 | 3.22 | 2.38 |
| Log (hp/wt) | -0.36 | -0.37 | -0.35 | -0.34 | -0.37 | 0.00 |
| Footprint (sq ft) | -0.01 | 0.02 | -0.24 | 0.01 | 0.02 | 0.70 |
| Price | 231 | 200 | 69 | 215 | 200 | 1,049 |
| Panel B: Welfare changes (2016 vs. 2012 standards, billion 2015\$) | | | | | | |
| Consumer (WTP) | -9.65 | -6.28 | -9.40 | -9.40 | -6.28 | 6.06 |
| Consumers (PDV) | 5.29 | 14.38 | 5.63 | 5.38 | 6.76 | 19.73 |
| Firms | -5.76 | -5.43 | -6.15 | -6.04 | -5.43 | -9.01 |
| Carbon emissions | 4.05 | 4.05 | 4.07 | 4.04 | 1.94 | 3.93 |

Notes: The table reports results from the scenarios described in the column headings. Each column reports the differences in attribute and welfare changes between the policy scenario and the corresponding baseline. The first column repeats the results from the central policy and baseline scenarios. SCC = social cost of carbon; WTP = willingness to pay; PDV = present discounted value.

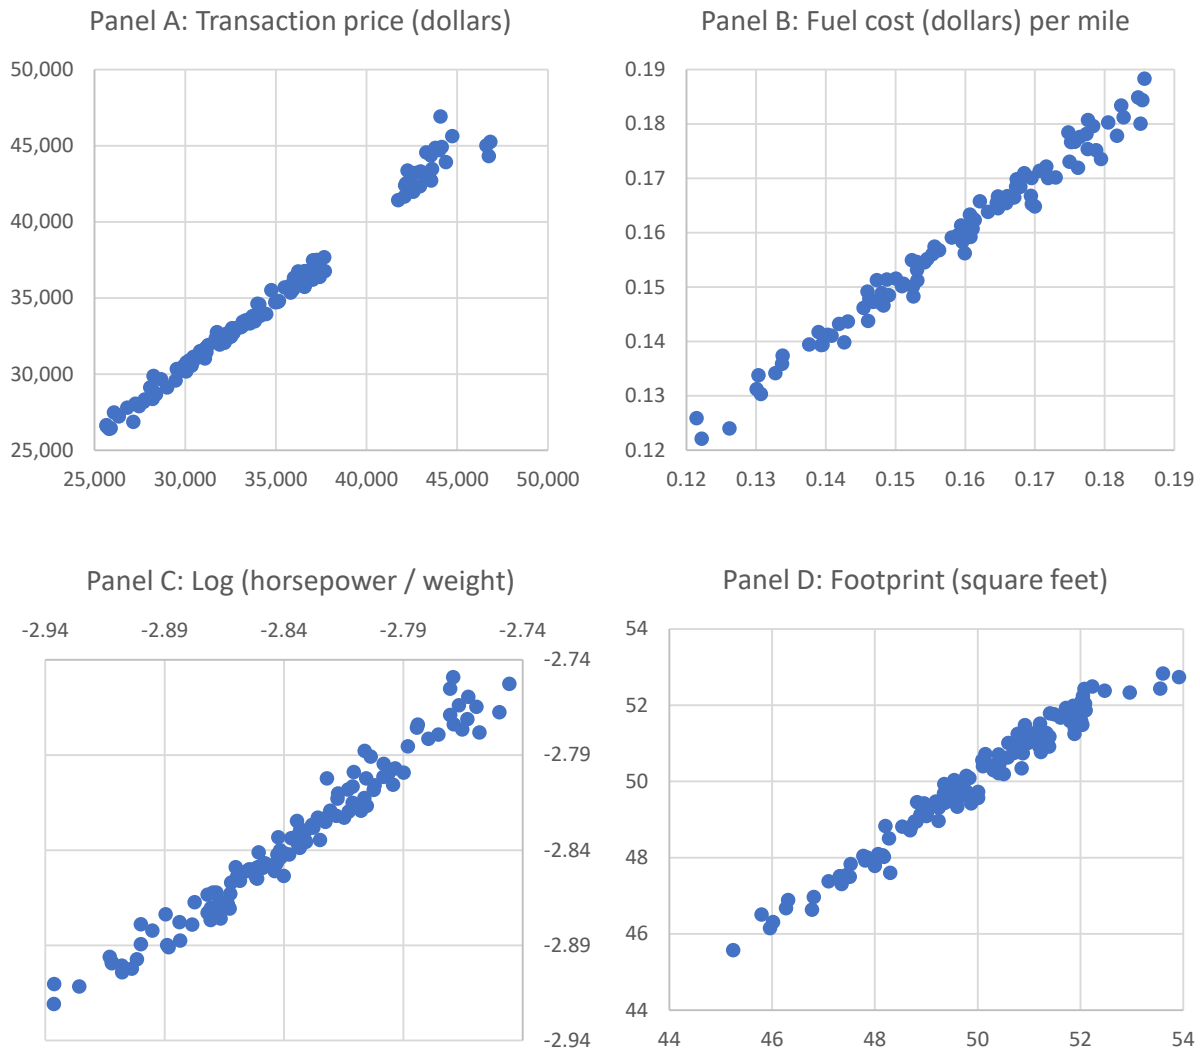
Appendix Figures

Figure A.1: Fuel Economy for Vehicles Sold in 2011 and Requirements in 2012 and 2016



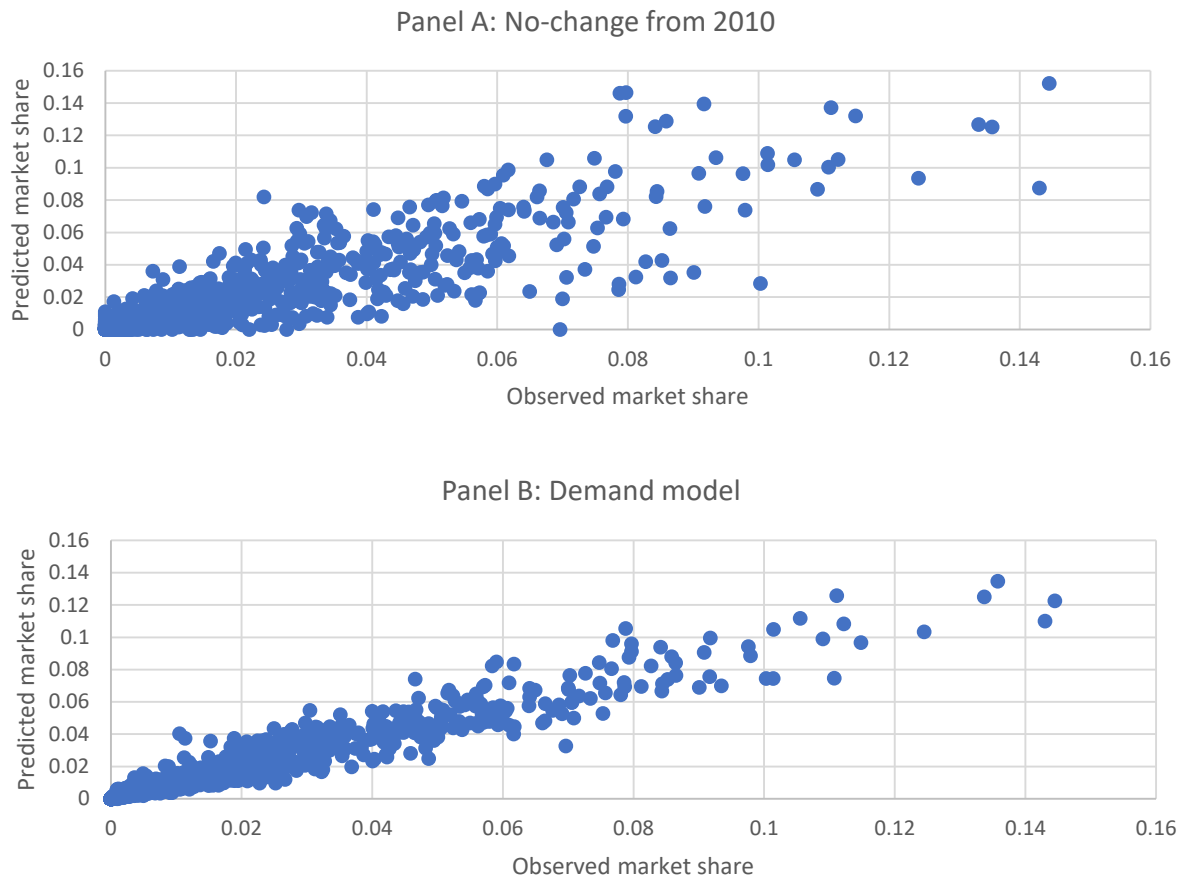
Notes: Each circle and x represent a unique vehicle sold in 2011. Each blue circle is the vehicle's fuel economy, in miles per gallon, plotted against its footprint, in square feet, where the footprint is computed as the product of the wheelbase and width. Each orange x shows the vehicle's fuel economy requirement in 2012, and each black x shows the fuel economy requirement in 2016.

Figure A.2: Comparison of Predicted and Observed Attributes by Demographic Group and Year



Notes: For each demographic group, we compute the predicted mean attribute indicated in the panel title using the vehicle market shares predicted by the model. The figure plots the predicted mean against the observed sales-weighted mean.

Figure A.3: Comparison of Predicted and Observed 2015 Market Shares by Demographic Group, Brand, and Class: No-change vs. Demand Model



Notes: Vehicles are aggregated by brand and class. The figure plots the predicted against observed market share by aggregated vehicle and demographic group. In Panel A, the prediction is equal to the observed market share in 2010. In Panel B, the prediction is made using the demand model.

Appendix Tables

Table A.1: Effects of Standards on Vehicle Prices and Attributes and Consumer Surplus

| Income group | Age group | Urban? | Average transaction price (2015\$) | Average fuel economy (mpg) | Average log (horsepower /weight) | Average consumer surplus using preference parameters | Average consumer surplus using fuel cost savings |
|--------------|-----------|--------|------------------------------------|----------------------------|----------------------------------|--|--|
| < 44k | < 45 | No | 154 | 3.93 | -0.43 | -42 | 60 |
| < 44k | < 45 | Yes | 31 | 3.92 | -0.43 | -34 | 99 |
| < 44k | >=45 | No | -126 | 3.59 | -0.42 | -13 | 144 |
| < 44k | >=45 | Yes | -109 | 3.57 | -0.37 | 4 | 129 |
| 44k – 91k | < 45 | No | 7 | 3.57 | -0.43 | -44 | 101 |
| 44k – 91k | < 45 | Yes | 89 | 3.67 | -0.39 | -51 | 79 |
| 44k – 91k | >=45 | No | 127 | 3.30 | -0.39 | -94 | 46 |
| 44k – 91k | >=45 | Yes | 219 | 3.50 | -0.36 | -165 | -8 |
| 91k – 123k | < 45 | No | 0 | 3.23 | -0.42 | -18 | 140 |
| 91k – 123k | < 45 | Yes | 217 | 3.53 | -0.38 | -118 | 71 |
| 91k – 123k | >=45 | No | 431 | 3.03 | -0.35 | -273 | -49 |
| 91k – 123k | >=45 | Yes | 396 | 3.42 | -0.34 | -171 | -47 |
| 123k – 185k | < 45 | No | 295 | 2.97 | -0.37 | -160 | 124 |
| 123k – 185k | < 45 | Yes | 305 | 3.11 | -0.34 | -123 | 70 |
| 123k – 185k | >=45 | No | 536 | 2.82 | -0.31 | -226 | -45 |
| 123k – 185k | >=45 | Yes | 825 | 2.89 | -0.28 | -286 | -165 |
| > 185k | < 45 | No | 174 | 2.57 | -0.33 | -1 | 155 |
| > 185k | < 45 | Yes | 489 | 2.62 | -0.28 | -85 | 31 |
| > 185k | >=45 | No | 933 | 2.37 | -0.25 | -305 | -226 |
| > 185k | >=45 | Yes | 1,060 | 2.48 | -0.22 | -360 | -353 |

Notes: The table reports changes in transaction prices, vehicle attributes, and consumer surplus by demographic group, comparing the baseline and central policy scenarios. See Section 5.2 for description of the two ways we define change in consumer surplus.

Table A.2: Effects of Standards on Vehicle Prices and Attributes and Profits, by Firm

| Firm | Average transaction price (2015\$) | Average fuel economy (mpg) | Average log (horsepower /weight) | Profits per vehicle (2015\$) |
|---------------|------------------------------------|----------------------------|----------------------------------|------------------------------|
| GM | 251 | 3.00 | −0.41 | 17 |
| Ford | 123 | 3.08 | −0.35 | −11 |
| Toyota | 260 | 3.18 | −0.33 | 20 |
| Honda | 302 | 3.51 | −0.29 | 32 |
| Hyundai | 148 | 3.71 | −0.46 | 1 |
| Fiat-Chrysler | −28 | 3.57 | −0.45 | −29 |
| Nissan | 340 | 4.30 | −0.31 | 16 |
| Volkswagen | 546 | 3.70 | −0.28 | 37 |
| BMW | 250 | 3.16 | −0.40 | −39 |
| Subaru | 452 | 3.02 | −0.26 | 31 |
| Daimler | −94 | 3.49 | −0.35 | −110 |
| Other | 1,487 | 3.08 | −0.18 | −10,315 |

Notes: The table reports the changes in transaction prices, vehicle attributes, and profits by manufacturer, comparing the baseline and central policy scenarios.