Pass-Through and Welfare Effects of Regulations That Affect Product Attributes

Benjamin Leard*	Resources for the Future
Joshua Linn	University of Maryland and Resources for the Future
Katalin Springel	Georgetown University

January 2019

Abstract

A key finding in the literature is that the greater the pass-through of an input cost shock or tax to product prices, the larger the welfare loss to consumers. We show that the relationship between pass-through and welfare changes does not hold for a regulation that affects production costs and product attributes. An analytical model shows that the larger the willingness to pay (WTP) for the product attribute, the greater the pass-through but the smaller the welfare loss (or the larger the welfare gain) for consumers. We confirm this intuition in the context of passenger vehicle fuel economy standards using new estimates of consumer demand and an equilibrium model. Pass-through and welfare changes are positively correlated with WTP for fuel economy across demographic groups and manufacturers. Accounting for WTP breaks the direct link between pass-through and welfare changes identified in prior literature.

Keywords: incidence, willingness to pay, product attributes, fuel economy standards
JEL codes: H22, L50, L13, 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.

1 Introduction

A recent literature has examined pass-through and welfare effects of input cost shocks, taxes, and regulations that raise production costs in imperfectly competitive markets. Weyl and Fabinger (2013) show that in imperfectly competitive markets, the greater the pass-through of regulatory costs or taxes to product prices, the larger the welfare losses to consumers. Roughly speaking, price-inelastic product demand implies more pass-through, and larger welfare costs to consumers. Motivated by this theory, Ganapati et al. (2018) examine the pass-through of energy prices–as a proxy for a carbon price–to product prices among certain manufacturing industries. For petroleum refining, Muehlegger and Sweeney (2017) find low pass-through of idiosyncratic cost shocks and roughly full pass-through of aggregate shocks, suggesting that a global carbon price on oil would be passed through fully to petroleum prices.

Welfare effects of energy efficiency standards across demographic groups have become particularly contentious in the public debate and has received increasing attention in the literature. The Trump administration's proposal to weaken US fuel economy and greenhouse gas (GHG) standards is motivated partly out of concern for the possible adverse effects of the standards on low-income groups (EPA 2016; NHTSA 2018). Many vehicle manufacturers have claimed that they cannot fully pass through cost increases to consumers, and that tighter standards reduce their profits. Davis and Knittel (2016) and Levinson (2016) show that standards are likely to be more regressive than a carbon tax on fuels, although the conclusion depends on the use of the tax revenue (that is, whether a lump-sum rebate, a reduction in labor income taxes, and so on).

In this paper, we consider product market regulations that affect product attributes as well as production costs. Product market regulations often set standards for attributes of the product consumers value. For example, energy efficiency standards for refrigerators set minimum levels of energy efficiency that the products must attain. Typically, these regulations raise the cost of producing the product because manufacturers must innovate or adopt existing technology to meet the standards. These standards can also affect consumer demand if consumers value the energy efficiency.[1]

Because product attribute demand may vary across demographic groups, we address how attribute demand affects pass-through of regulatory costs and variation in welfare effects across demographic groups and firms. As a motivating example, consider an energy efficiency

[1]Product regulations can indirectly affect attributes of products that are not directly regulated. For example, Klier and Linn (2016) show that recently tightened fuel economy standards caused vehicle manufacturers to trade off horsepower and weight for fuel economy. We address this possibility in Section 6.

standard and suppose high-income groups are less price sensitive and have high willingness to pay (WTP) for a product attribute. The existing literature suggests that because high-income consumers are less price sensitive, pass-through rates of the standard's costs would be higher for products purchased by high-income consumers than for products purchased by low-income consumers. This would imply larger welfare losses for high-income consumers, and that the standards may be progressive. On the other hand, high-income consumers have greater WTP for the attribute, which would imply larger welfare increases and that the standards may be regressive. The objective of the paper is to disentangle these opposing forces, both in theory and in practice.

Section 2 extends the standard analysis of pass-through in an imperfectly competitive market to consider a case in which a regulation affects a product attribute that consumers value. In the standard analysis of a regulation that affects production costs but not product attributes, there is a direct relationship between pass-through and the welfare effects on consumers: the greater the pass-through, the larger the welfare loss. We use the analytical model to demonstrate that this relationship breaks down when the regulation affects product attributes. In this case, if the regulation causes an increase in the value of an attribute, consumer demand increases as demand curves shift out from the origin. The more the demand curve shifts (that is, the higher the WTP), the higher the pass-through. Moreover, the higher the WTP, the higher the welfare gain (or lower the welfare loss) for both consumers and manufacturers.[2] Hence, when regulation affects a product attribute, a greater rate of pass-through can be associated with a smaller welfare loss (or a greater welfare gain).

Having demonstrated this theoretical point, we use unique data and a new estimation strategy to examine the variation in welfare effects across demographic groups and firms of recently tightened fuel economy standards. The National Highway Traffic Safety Administration (NHTSA) has regulated fuel economy since the 1970s and the Environmental Protection Agency (EPA) has regulated GHG emissions since 2012. Since 2005, US fuel economy standards have been tightening after almost two decades in which they were fixed. However, these standards have been highly controversial, and in 2018 the agencies proposed weakening the standards largely because of the high expected costs (EPA 2018).

The main data set is built from survey responses of about 1.1 million new vehicle buyers between 2010 and 2015, and supplemented with data from the Consumer Expenditure Survey (CEX), IHS Automotive, the Bureau of Labor Statistics, and Wards Auto. The household

[2] Although this mechanism is present in papers that employ imperfect competition models to characterize the welfare effects of tighter standards (e.g., Klier and Linn (2012) and Jacobsen (2013)), those papers have not isolated this mechanism in explaining the welfare results. Our model also generalizes Houde (2018), who focuses on heterogeneous information across consumers.

survey data include demographics such as income, age, and urbanization, which we use to estimate heterogeneous preferences for vehicle attributes. Observed vehicle purchase patterns vary widely across groups.

We specify an equilibrium model in which consumers maximize utility by choosing a vehicle, and manufacturers maximize profits by choosing vehicle prices. The demand model has three distinguishing features. First, vehicles are defined at a highly disaggregated level, and the choice sets include about 1,000 unique vehicles. The choice sets are differentiated trims and power train configurations within a model, and conforms to the set of vehicles from which consumers choose in practice. The number of choices is several times greater than in most other vehicle demand models.[3]

Second, we estimate a distinct set of preference parameters for 20 demographic groups. This modeling approach is motivated by two considerations: it enables a transparent demonstration of the role of WTP in determining welfare changes; and the model can be estimated in two stages using a fixed effects regression and a straightforward generalized method of moments (GMM) estimator. Estimation is far simpler than a random preference coefficients model.

Third, the GMM estimator allows for the endogeneity of price, fuel economy, and performance (measured by the ratio of horsepower to weight; Leard et al. (2017)) using a novel set of instruments. We instrument for price using the physical dimensions of competing vehicles, which manufacturers take as pre-determined in the medium run (up to roughly 5-7 years). We identify WTP for fuel economy and performance by comparing vehicle prices and market shares for closely related vehicles.

The estimated preference parameters exhibit a substantial degree of heterogeneity across the demographic groups. Lower income groups tend to be more price responsive and have lower WTP for fuel economy and performance. For instance, members of the lowest income quintile have average WTP for fuel economy equal to about one-third of that of the highest income quintile. Younger households have lower WTP for fuel economy and urban households have higher WTP for fuel economy. The demand model performs well in predicting market shares, both in and out of sample.

Having estimated the preference parameters, we recover the marginal costs of each vehicle from the equilibrium first order condition for a vehicle's price that corresponds to the manufacturer's profit maximization problem. We model marginal costs as a function of

[3] Recent examples of vehicle demand models include Berry et al. (2004) (a choice set of 203 new vehicles), Train and Winston (2007) (200 new vehicles), Bento et al. (2009) (270 new and used vehicles), Whitefoot et al. (2017) (473 new vehicles).

the vehicle's fuel economy and other attributes, and estimate the effect of fuel economy on marginal costs.

We use the equilibrium model to demonstrate the role of WTP for fuel economy in determining the pass-through and private welfare effects of a mandatory fuel economy increase.[4] The fuel economy regulation raises the marginal costs of each vehicle according to the estimated marginal cost function. We simulate the equilibrium changes in vehicle prices and market shares and calculate changes in consumer welfare by demographic group and profits by manufacturer.

Variation of welfare effects across demographic groups and manufacturers is consistent with the analytical model. Holding fixed own-price elasticity of demand, pass-through is greater for demographic groups with higher WTP for fuel economy. Likewise, manufacturers that sell to consumers with high WTP have higher pass-through and a higher increase in profits than do other manufacturers.

The simulation results show that the connection between pass-through and welfare that the literature has emphasized does not hold when regulations affect product attributes that consumers value. In a hypothetical setting where consumers do not value fuel economy at all, we reproduce the standard result that higher pass-through implies larger welfare losses for consumers. However, when consumers value fuel economy based on our demand model estimates, we find that across demographic groups, pass-through is uncorrelated with welfare changes. This situation occurs because of two opposing factors. On the one hand, high-income groups are less price sensitive than low-income groups, which causes higher pass-through and welfare losses for high income groups. On the other hand, high-income groups also have high WTP for fuel economy, which causes high pass through and welfare increases. The two opposing effects roughly cancel one another, causing a near-zero correlation between pass-through and welfare changes.

Accounting for fuel economy WTP makes the standards regressive in the short run. High income households tend to gain more total surplus per household than low income households; this is true both in terms of absolute welfare changes and changes relative to income. These results also highight the importance of accounting for variation in WTP across demographic groups; ignoring this variation would cause us to conclude that the standards are progressive. Note that these distributional results are based on a short-run analysis because fuel economy and performance are exogenous, and we focus on welfare effects across new vehicle consumers

[4]For simplicity, we focus on changes in private welfare, including consumer and producer surplus, which represents most of the changes in social welfare due to fuel economy regulation (EPA 2018).

rather than all households. Accounting for the effects of standards on used vehicle consumers would likely strengthen the regressivity (Jacobsen 2013).

This paper contributes to the literature in several ways. First, we demonstrate that product manufacturers can pass through a higher proportion of regulatory costs if they sell products to consumers with higher WTP for the attribute affected by regulation. To our knowledge, the literature has not examined in theory the pass-through and welfare effects of regulations that affect product attributes. Second, we introduce a demand model that allows for extensive preference heterogeneity across demographic groups over a highly disaggregated choice set. Notwithstanding the disaggregated choice set and degree of modeled heterogeneity, the parameters are straightforward to estimate and the model performs well. Third, the model allows for endogeneity of prices, fuel economy, and performance while relaxing many of the identification assumptions in the literature. Whitefoot et al. (2017) allow for this endogeneity but use as instruments fuel type and other attributes that are likely to be correlated with the endogenous attributes. Klier and Linn (2012) and Leard et al. (2017) also use instruments that isolate medium-run variation in vehicle attributes, but they do not allow for observed preference heterogeneity across consumer groups. Finally, the simulated effects of tighter standards confirm the predictions of the analytical model, and show that in the short run the standards are regressive across new vehicle consumers.

2 Analytical Model of Pass-Through and Welfare

This section describes a general model of a regulation that affects a single product attribute. We derive closed-form expressions for pass-through and welfare changes for consumers and producers.

2.1 Setup of the model

We consider a static model of a monopolist that chooses price to maximize profits. Consumer demand for the product is denoted as $q(p, m)$, where p is the product's price and m is the level of a product attribute. We use m to indicate the attribute affected by the regulation, such as fuel economy. Consumer demand is decreasing in price and increasing in the level of the product attribute. The marginal cost of producing the product is $c(m)$, where $c(\cdot)$ is an increasing function of m.

The producer takes the level of the product attribute as exogenous and chooses p to maximize profits, π:

$$\pi = \max_p [p - c(m)]q(p, m). \tag{1}$$

The first order condition for price yields the well-known equation for equilibrium price:

$$p^* = c(m) - \frac{q(p, m)}{\partial q/\partial p}. \tag{2}$$

2.2 Welfare effects of an exogenous fuel economy increase

We now consider a regulation that requires m to increase. Differentiating the equilibrium price condition in Equation (2) with respect to m yields

$$\frac{dp^*}{dm} = \frac{dc}{dm} - \left[\left(\frac{\partial q}{\partial p}\frac{dp^*}{dm} + \frac{\partial q}{\partial m}\right)\frac{\partial p}{\partial q} + q(p, m)\frac{\partial^2 p}{\partial q^2}\frac{dp^*}{dm}\right]. \tag{3}$$

This expression simplifies to

$$\frac{dp^*}{dm} = \frac{\frac{dc}{dm} - \frac{\partial q/\partial m}{\partial q/\partial p}}{2 + q(p, m)\frac{\partial^2 p}{\partial q^2}}. \tag{4}$$

The second term in the numerator of Equation (4) represents the sensitivity of demand with respect to the product attribute relative to the sensitivity demand with respect to price. This term is equivalent to marginal WTP (MWTP) for the product attribute. To see this, differentiating demand in equilibrium $q(p^*(m), m) = q^*$ yields

$$\frac{\partial p}{\partial m} = -\frac{\partial q/\partial m}{\partial q/\partial p}. \tag{5}$$

The marginal change in price with respect to the product attribute along the demand schedule represents MWTP for the product attribute, denoted as $MWTP$. In words, $MWTP$ is the vertical shift in the demand curve, measured at the initial equilibrium quantity, caused by the change in the attribute. Making this substitution in (4) yields

$$\frac{dp^*}{dm} = \frac{\frac{dc}{dm} + MWTP}{2 + q(p, m)\frac{\partial^2 p}{\partial q^2}}. \tag{6}$$

The change in price depends on the marginal cost of increasing the level of the product attribute and the MWTP for the product attribute. The higher the MWTP, the higher the price increase. If the inverse demand function has no curvature ($\partial^2 p/\partial q^2 = 0$), then the change in price is the average of the marginal cost of increasing the level of the product attribute and the MWTP for the product attribute.

Next, we analyze the welfare effects of an exogenous increase in the level of the attribute. Differentiating the profit function (1) with respect to the attribute and applying the envelope theorem yields

$$\frac{d\pi^*}{dm} = (p^* - c(m))\frac{\partial q}{\partial m} - \frac{dc}{dm} q(p^*, m). \tag{7}$$

We define the own-price elasticity of demand as $\epsilon_p = -\frac{\partial q}{\partial p}\frac{p^*}{q}$. Substituting $MWTP = -\frac{\partial q/\partial m}{\partial q/\partial p}$ into Equation (7) and re-arranging yields

$$\frac{d\pi^*}{dm} = q(p^*, m)\left[MWTP\frac{p^* - c(m)}{p^*}\epsilon_p - \frac{dc}{dm}\right]. \tag{8}$$

Equation (2) can be expressed as $\frac{p^* - c(m)}{p^*} = \frac{1}{\epsilon_p}$, and Equation (8) simplifies to

$$\frac{d\pi}{dm} = q(p, m)\left[MWTP - \frac{dc}{dm}\right]. \tag{9}$$

The change in profits is scaled by total sales of the product, q. The term within the brackets is the difference between the MWTP for the product attribute and the marginal cost of increasing the product attribute. The increase in profits is greater the larger is the MWTP for the product attribute.

To characterize the welfare effects of the attribute change on consumers, we first define equilibrium consumer surplus as

$$CS = \int_0^{q^*} [p(q, m) - p^*]\, dq, \tag{10}$$

where $p(q, m)$ represents inverse demand (or WTP) for the product. To determine the effect of a change in the product attribute on consumer surplus, we apply Leibnitz' rule for differentiation under the integral sign:

$$\frac{dCS}{dm} = (p(q^*, m) - p^*)\frac{\partial q^*}{\partial m} - 0 + \int_0^{q^*}\left[MWTP - \frac{dp^*}{dm}\right] dq. \tag{11}$$

The first term in Equation (11) cancels because the inverse demand function is evaluated at equilibrium demand, which equals the equilibrium price. Therefore

$$\frac{dCS}{dm} = \int_0^{q^*} \left[MWTP - \frac{dp^*}{dm} \right] dq. \tag{12}$$

Substituting Equation (4) into Equation (12) yields

$$\frac{dCS}{dm} = \int_0^{q^*} \left[MWTP - \frac{MWTP}{2 + q(p,m)\frac{\partial^2 p}{\partial q^2}} - \frac{\frac{dc}{dm}}{2 + q(p,m)\frac{\partial^2 p}{\partial q^2}} \right] dq. \tag{13}$$

For a sufficiently small curvature of the inverse demand function, the higher the MWTP for the product attribute, the larger the increase in consumer surplus. If we use the standard definition of pass-through, as the change in equilibrium price given an increase in marginal costs, we reach the following conclusions:

- A higher MWTP for the product attribute implies a higher pass-through.
- A higher MWTP for the product attribute implies a larger increase in profits.
- A higher MWTP for the product attribute implies a larger increase in consumer surplus.

Figure 1 illustrates these results, where the two panels represent separate markets for the product. Period 1 is prior to regulation, where the curve D_1 is consumer demand and MC_1 is the marginal cost of producing the product. Equilbrium prices and quantities are determined by the firm's profit maximization, and the period 1 equilibriums are the same in the two panels.

In period 2, a regulation in both markets raises the level of the product attribute, which causes the cost curves to shift up to MC_2. The difference between the two panels is that in Panel A consumers have lower MWTP for the product attribute than do consumers in Panel B. This difference is represented by the larger shift of the demand curve in Panel B than in Panel A.

The higher MWTP in Panel B causes a larger price increase than in Panel A. This highlights our first result that a higher MWTP for the product attribute implies a higher pass-through.

Profits are the rectangle bounded by marginal costs, the vertical line at equilibrium quantity, the horizontal line at the equilibrium price, and the vertical axis. Profits are the same in the two panels in period 1, but in period 2 profits are larger in Panel B than in Panel A; this is the second result above.

Consumer surplus is the triangle bounded by the demand curve, the horizontal dashed line at the equilibrium price, and the vertical axis. Consumer surplus is the same in the two panels in period 1. Consumer surplus increases by more in Panel B than in Panel A, which is the third result highlighted above.

The model includes a few simplifications. First, the firm is a single-product monopolist choosing price to maximize profits. The results are identical if we formulate the profit maximization problem over quantity rather than price. If the firm sells multiple products, it would consider cross-demand effects across its products when choosing the profit-maximizing price (or quantity). The expressions for pass-through and welfare changes would include these cross partial terms, and as long as the cross-partials are sufficiently small in magnitude, the three conclusions carry through. In the setting we consider below, of a mandatory fuel economy increase for the US passenger vehicles market, the cross partials are small in magnitude. A similar situation is likely to hold in other markets for consumer goods with highly differentiated products, such as home appliances.

A second simplification is that we have considered a monopolist rather than an oligopolist competing with other firms. If we model a market with differentiated products and Bertrand competition, the equilibrium quantity and price would depend on the attributes of all products in the market, rather than just the attribute of the product itself. If we consider a regulation that affects the attribute of a single product, the pass-through and welfare effects would be the same as those above. Alternatively, if the regulation affects attributes of all products, the pass-through and welfare changes would include cross partial derivatives of quantity with respect to the attributes of other products. As with the multi-product monopolist, the main results would carry through as long as the cross partials are sufficiently small.

Thus, the results generalize to a model with multi-product firms competing on price. For a firm, the average pass-through and profits change would depend on the average MWTP of consumers purchasing its products. For a group of consumers, such as a demographic group, the average pass-through and welfare change would depend on the MWTP for that group.

Finally, we have assumed that a consumer's expected utility from purchasing the product is the same as the realized utility. Consumer expectations may include systematic errors, which is sometimes referred to as an internality (Allcott et al. 2014; Allcott and Sunstein 2015). For example, consumers may enjoy more benefit from the product attribute than they expect. We discuss this possibility in the Appendix, and we show that in general, the main results carry through.[5]

3 Data and Summary Statistics

In the remainder of the paper, we use an equilibrium model of the new vehicles market to estimate the pass-through rate and private welfare effects of a fuel economy increase. This section describes the data and presents summary statistics.

3.1 Data

The data set consists of annual new vehicle sales by demographic group and vehicle, annual used vehicle sales by demographic group, and attributes and prices by vehicle. Individual survey responses to the MaritzCX New Vehicle Customer Survey (NVCS) data for the years 2010 through 2015 constitute the primary data source.

The NVCS data contain responses from households that recently obtained a new vehicle. The data include the purchase price of the vehicle (excluding trade-in and including sales tax) and demographics such as income, age, education, state of residence, and population density.

The survey also asks about the financing of the vehicle, down payment and loan terms (if any), and whether the vehicle was purchased or leased. The average response rate of the survey is 9 percent, and across the 6 survey years we have 1.1 million households, which represents about 1 percent of all vehicle buyers.

We supplement the MaritzCX data with data from EPA, Wards Automotive, IHS, CEX, and BLS. A vehicle is defined as a unique model year, make, model, trim, fuel type, drive type, body style, and the number of cylinders. For each vehicle and model year in the MaritzCX data, we merge EPA data on fuel economy and fuel-saving technologies as well as Wards data on manufacturer suggested retail price (MSRP), wheelbase, width, and horsepower.[6]

[5] We analyze a static model for simplicity. The literature on product market regulation has primarily used static models.

[6] The EPA fuel economy data are merged by vehicle except that we aggregate across body style because the EPA data do not include that variable. This likely sacrifices little fuel economy variation; based on Wards data, after controlling for the other vehicle identifiers, body style accounts for less than 1 percent of the remaining fuel economy variation. For each vehicle in the data, we merge combined city and highway fuel economy ratings that appear on new vehicle windows.

We use data from Cars.com to impute vehicle characteristics for the small number of missing observations.

We define five income groups that correspond roughly to quintiles of the distribution of new and used vehicle buyers in the CEX data, two age groups (above and below the median age of 45 years), and an indicator for urbanization that equals 1 for households living in areas above the median population density. The demographic groups are defined to roughly equate the number of households in the CEX data for each group. The CEX data also provide market shares for used vehicle purchases, which we aggregate to an outside option.

An advantage of the MaritzCX data is that they include reported transaction prices rather than MSRP. The transaction prices reflect any negotiation between the household and dealer as well as any dealer or manufacturer incentives. Much of the literature, such as Berry et al. (1995; 2004), uses MSRP. Unfortunately, the Maritz data do not include transaction prices for about 10 percent of the vehicles. The appendix describes the imputation of prices for missing observations.

A second advantage of the MaritzCX data is that they allow us to define a highly disaggregated choice set in the consumer demand model. The data include about 1,000 unique vehicles each year, which is several times larger than the number of unique choices that can be found in previous studies. For example, Berry et al. (1995) and Klier and Linn (2012) use the make and model to define a vehicle, which yields about 200 to 300 unique choices each year. Prior studies have aggregated the data in this manner either due to data or computational constraints. The level of vehicle aggregation in our data corresponds to the set of vehicles from which consumers choose, for instance including the choice between the base and luxury trims of a model, or between the 4-cylinder and 6-cylinder versions of a trim. The disaggregation reduces measurement error and the resulting bias in the estimated coefficients. The additional variation in our data, relative to more aggregated choice sets, helps us identify the preference parameters as we explain in Section 5.1. Moreover, the level of aggregation in our data aligns closely with the definition of a unique vehicle in the EPA and NHTSA analysis of GHG and fuel economy standards, facilitating comparison between our results and theirs.

We use the IHS Automotive and CEX data to weight observations in the MaritzCX data and account for potential variation in response rates across vehicles and demographic groups. The IHS data include registrations by year, quarter, and vehicle for all US households (that is, excluding fleet buyers). We use the CEX to compute the number of vehicles purchased by year, quarter, and demographic group. The appendix provides additional information about the CEX data and the procedure for weighting the MaritzCX observations to match the distributions of sales across vehicles according to IHS and across demographic groups according to CEX. In addition, because a used vehicle represents the outside option in the consumer demand model, we use the CEX data to construct a count of used vehicle purchases by year and demographic group.

Finally, we convert vehicle and fuel prices to 2015 dollars using the BLS Consumer Price Index. All dollar values reported in the paper are in 2015 dollars.

A few features of the data are worth highlighting. First, we use transaction prices rather than MSRP, which reduces measurement error and increases the variation available to identify preference parameters. Second, we construct purchases by demographic group, which allows us to estimate a unique set of preferences for each demographic group. Third, we use a highly disaggregated choice set, which reduces measurement error and helps identify the preference parameters as we explain in Section 5.1.

3.2 Summary statistics and background on fuel economy regulation

We provide summary statistics about vehicle purchase patterns and vehicle attributes. Figures 2 and 3 show extensive variation in sales-weighted mean vehicle attributes across demographic groups. Average purchase price varies by a factor of two across income groups. Rural households are more likely than urban households to purchase light trucks, and high income households are more likely to purchase plug-ins and hybrids. In the remainder of the paper, we use the log of the ratio of horsepower and weight as a proxy for the time needed for a vehicle to accelerate from rest to 60 miles per hour (Greene et al. 2018). The proxy is highly correlated with other potential measures of performance, such as the time needed to accelerate from 20 to 50 miles per hour (that is, for merging onto a highway). Footprint is the product of the vehicle's wheelbase (the distance between the two axles) and the width.[7] The figures show that low-income households tend to purchase vehicles with high fuel economy, low ratios of horsepower to weight, and small footprint. Used vehicles account for a larger share of total vehicle purchases for low income households than for high income households. In addition to the variation across income groups, there is substantial variation in mean vehicle attributes within income groups and across urbanization and age

[7] Throughout this paper, we compute the vehicle's footprint as the product of the wheelbase and width, because we do not observe the actual footprint as defined by EPA and NHTSA. This approximation likely introduces little measurement error.

groups. This variation motivates our choice of demographic groups; we observe substantially more heterogeneity using all three household characteristics, rather than solely income.

Figure 4 shows mean vehicle attributes by manufacturer. In this figure and in subsequent figures and tables, manufacturers are listed in order of declining total sales across the sample. The top 11 manufacturers collectively account for about 99 percent of the market. Vehicle attributes vary substantially across manufacturers. For example, BMW and Daimler sell vehicles with an average price about 40 percent higher than the average price of vehicles sold by other manufacturers. Toyota, Honda, and Hyundai sell vehicles with fuel economy about 25 percent higher than the vehicles sold by GM, Ford, and Fiat/Chrysler. There is also variation in the share of low income households and the share of urban households. This variation suggests that variation in consumer preferences across demographic groups could cause pass-through rates to vary across manufacturers. Although not shown in these figures, there is also substantial variation within a manufacturer and across make (for example, the Chevrolet and GMC makes sold by GM). For reference, Appendix Tables A.1, A.2, A.3, and A.4 show numerical values for the data reported in Figures 2, 3, and 4.

As we explain in Section 5.1, we identify WTP for performance using variation in performance across pairs of vehicles that have the same characteristics except the engine. We use the term engine twins to describe pairs of vehicles that have different engine configurations, are sold in the same market, and share a make, model, trim name, fuel type, drive type, and body style. Table 1 shows vehicle characteristics of the five most popular twins sold in 2015. The twin with the larger engine has higher performance, and tends to have a higher price, lower fuel economy and fewer sales than its smaller engine twin. For example, the version of the Ford F150 XL with the larger engine has more sales than the smaller engine version, suggesting that buyers of this vehicle value the performance.

Engine twins are common in the market. Figure 5 compares density functions of the attributes for all vehicles sold during 2010-2015 with density functions for engine twins. Although the right tails of the price and performance densities for the full vehicle sample are thicker, the shape of the densities is similar for all four attributes. The similarity of the means and density functions suggests that the engine twins are representative of vehicles sold in each market.

We next provide a brief background about the variation in fuel economy and GHG standards during our sample period. Fuel economy standards for light trucks increased throughout the sample, and standards for cars increased after 2011. EPA began setting GHG standards for cars and trucks in 2012. Between 2010 and 2015, fuel economy standards for light trucks increased by about 17 percent and fuel economy standards for cars increased by 32 percent.

Starting in 2012, for both cars and light trucks, the fuel economy and GHG requirements for each vehicle depend on its footprint, where the footprint is the area defined by the four wheels. Larger vehicles face lower fuel economy requirements than smaller vehicles, and cars face higher fuel economy requirements than light trucks. The GHG requirements are inversely related to the fuel economy requirements, so that larger vehicles and light trucks face higher GHG requirements than do smaller vehicles and cars. The overall GHG standard that each manufacturer faces is the sales-weighted average of the GHG requirements of its vehicles. The overall fuel economy standard that each manufacturer faces is the harmonic sales-weighted average of the fuel economy requirements. Because of the structure of the standards, manufacturers selling larger vehicles face lower fuel economy and higher GHG standards.

Figure 6 summarizes the stringency of the fuel economy standards in our sample period. Each x represents a unique vehicle in 2010, and the open circles and closed circles represent the fuel economy requirements for each vehicle in 2012 and 2015. For most vehicles, the fuel economy in 2010 lies well below the fuel economy requirement in 2012 and 2015, with a larger gap for light trucks than cars.

4 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.

4.1 Demand

We define each market as a model year, indicated by t, which represents the fourth quarter of the previous calendar year through the third quarter of the current calendar year. For example, model year 2012 begins in October of 2011 and ends in September of 2012. This definition of a model year is consistent with typical vehicle production cycles; vehicle attributes are constant during a model year but may change across model years.

Households maximize utility by choosing among a composite used vehicle and a set of new vehicles. Household i experiences utility u_{ijt} by choosing vehicle j in model year t, where $j = 1, 2, ..., J_t$ indexes new vehicles available in model year t and $j = 0$ is the used vehicle. We index vehicle attributes by k and household characteristics by d. We assign households to demographic groups indexed by g, and we define the 20 demographic groups that were discussed in the previous section: five income categories, two age categories, and two urban or rural categories. The base demographic group is defined as the lowest income group, young, and urban. Household utility is

$$u_{ijt} = v_{ijt} + \varepsilon_{ijt} = \sum_k \sum_g x_{jkt} h_{igt} \beta_{kg} + \sum_k x_{jkt} \bar{\beta}_k + \xi_{jt} + \varepsilon_{ijt}. \tag{14}$$

The term x_{jkt} represents vehicle j's value of attribute k in market t. The term h_{igt} is a dummy variable indicating whether household i is in demographic group g. Parameter β_{kg} is the difference between the marginal utility of vehicle attribute k for households in demographic group g, and the marginal utility of attribute k for the base demographic group. The term with a double summation measures observed heterogeneity across demographic groups in their marginal utilities. The coefficient $\bar{\beta}_k$ is the marginal utility for the base demographic group of attribute k. The second summation term on the right-hand side of Equation (14) is the utility of the vehicle for the base demographic group.[8] The term ξ_{jt} denotes the unobserved mean utility across demographic groups for vehicle j. The term ε_{ijt} is household i's unobserved utility for vehicle j that is unexplained by the observed vehicle attributes.

Under the standard assumption that the error term ε_{ijt} has a type 1 extreme value distribution, the probability that household i chooses vehicle j in market t is

$$Pr_{ijt} = \frac{e^{v_{ijt}}}{\sum_k e^{v_{ikt}}}. \tag{15}$$

For each household demographic group, we normalize the outside good utility to zero. Based on this assumption and Equation (15), in the Appendix we derive a linear equation linking observed market shares, product attributes and marginal utilities, and unobserved product attributes

$$\ln(s_{gjt}) - \ln(s_{g0t}) = \sum_k \sum_g x_{jkt} h_{igt} \beta_{kg} + \sum_k x_{jkt} \bar{\beta}_k + \xi_{jt}. \tag{16}$$

The left-hand side of Equation (16) is the difference between the log share of purchases of vehicle j in market t by demographic group g and the log share of purchases of the outside good in market t for demographic group g. The equation is linear in consumer heterogeneity

[8]It is common in the discrete choice literature to define the mean utility of the product and deviations from the mean utility, such as Berry et al. (1995). Instead, we define the utility of the base demographic group and differences between utility of other demographic groups and the base group. This definition is consistent with the two-stage estimation of the utility function parameters described in the next section, in which we first estimate the differences in parameters between each demographic group and the base group and subsequently estimate the parameters for the base group.

parameters as well as average preference parameters, which facilitates a simple estimation strategy described in Section 5.1.

We make several comments on the structure of the choice model that yields Equation (16). The model builds on the existing consumer demand literature and allows an extensive degree of preference heterogeneity across consumers. Each of the 20 demographic groups defined in the previous section has a unique price sensitivity and WTP for each of the attributes in Equation (16). The heterogeneity implies that aggregate substitution patterns are more plausible than those implied by a logit demand model without observed heterogeneity. For example, an increase in the price of all BMW vehicles has a larger effect on market shares of vehicles purchased by high-income consumers than vehicles purchased by low-income consumers.

According to the model, heterogeneity across households in market shares arises from variation in observable demographics. An alternative approach is to formulate a mixed logit model following Berry et al. (1995), Petrin (2002), Berry et al. (2004), or Train and Winston (2007). These models introduce consumer heterogeneity in the form of random taste variation across households that does not depend on the household's demographics.

We model heterogeneity based on observed demographics for four reasons. First, it enables a transparent interpretation of the results on pass-through and welfare, because there is a direct link between the estimated demand parameters and the variation across households in pass-through and welfare. Such a direct link would not be present in a mixed logit model.

The second reason for this modeling choice is that estimation of a mixed logit model would be computationally infeasible, given the number of households, the size of the choice set, and the number of markets that we are modeling. The household survey data contain 1.1 million households, five markets, and about one thousand vehicles in each choice set. In contrast, the prior literature using household data to estimate mixed logit demand models typically includes only a few hundred households, one market, and a few hundred choice set alternatives. Reducing the size of the household sample, number of markets, and number of choice alternatives for computational reasons would prevent us from identifying unbiased estimates of key parameter values and would mask some of the consumer responses to policy.

Third, mixed logit models may yield multiple sets of parameter estimates that imply a wide range of demand elasticities (Knittel and Metaxoglou 2014). This is because the computiational routines required to estimate these models do not guarantee a unique solution for the parameter values. Our model avoids this issue because our estimation is computationally fast and yields a unique set of parameter values.

Finally, in our model, the heterogeneity parameters are identified by variation across demographic groups in response to variation in vehicle attributes and prices across vehicles and markets. This contrasts with many vehicle demand models that estimate unobserved heterogeneity with random coefficients but without repeated choice microdata, such as Berry et al. (1995). In these models, heterogeneity parameters are identified by changes in choice sets across markets, making it difficult to determine whether the implied heterogeneity reflects the preference heterogeneity or something else (Ackerberg and Rysman 2005).

Note that our demand model imposes the assumptions that preference parameters do not vary across households that belong to the same demographic group. This contrasts with a random coefficients model, in which preference parameters vary randomly across households. Below, we show that the observed vehicle choices in our data are consistent with the assumption that preference parameters do not vary within a demographic group.

The logit structure imposes the assumption of independence of irrelevant alternatives within but not across demographic groups. Therefore, the cross-group heterogeneity allows for the possibility that a price increase for one vehicle can affect market shares of other vehicles disproportionately. For example, our model allows a price increase for all BMW vehicles to have a larger effect on market shares for Audi vehicles than for Chevrolet vehicles.

We model the outside option as the decision to buy a used vehicle. This decision effectively makes our choice model conditional on purchasing a new or used vehicle. As we explain in the Appendix, this modeling decision is a departure from most of the vehicle demand literature that either excludes an outside good altogether or treats the outside good as the decision to not select any vehicle at all. A benefit to including an aggregate used vehicle in the choice set is that we can examine the effect of fuel economy standards on the demand for used vehicles.[9]

4.2 Supply

The supply side is static, following Klier and Linn (2012) and Jacobsen (2013). Each manufacturer takes as exogenous the set of vehicles in each market and the non-price

[9]In principle, a change in demand for used vehicles may also affect prices of those vehicles (Jacobsen 2013). Because we do not model the used vehicle market explicitly, we do not estimate such price changes. Therefore, we interpret the welfare changes in our model as corresponding to the changes in welfare of consumers who purchase new vehicles with and without the policy change, as well as the difference in welfare between new and used vehicles for those consumers whose choice of a used vehicle is affected by the policy. The appendix discusses this interpretation in more detail.

attributes of those vehicles. Manufacturers compete by choosing the prices of their vehicles in a Bertrand-Nash imperfectly competitive market. Each manufacturer is subject to fuel economy standards, so that the harmonic mean of its car and truck fleet fuel economy must exceed a particular threshold. Although historically some manufacturers have elected to pay fines for noncompliance, during our sample period all manufacturers have complied. For simplicity, we assume that the GHG standards are harmonized with the fuel economy standards so that there is one set of binding standards. In addition, although the fuel economy standards include restrictions on credit trading across vehicles and manufacturers, for simplicity we assume that these restrictions are not binding.[10]

Each firm chooses vehicle prices and credit purchases to solve the maximization problem

$$\max_{\{p_{jt},x\}_{j\in J_{mt}}} \sum_{j\in J_{mt}} [p_{jt} - c_{jt}(m_{jt})]q_{jt} - p_x x \tag{17}$$

subject to

$$\sum_{j\in J_{mt}} \left(\frac{1}{m_{jt}} - \frac{1}{M_{jt}}\right) q_{jt} - x \leq 0. \tag{18}$$

Profits equal the product of vehicle sales (q) and the difference between vehicle price (p) and marginal costs (c), minus net costs of credit trading ($p_x x$). Marginal costs depend on the vehicle's fuel economy. The constraint (18) represents the fuel economy standard. The standard assigns a fuel economy target M_{jt} for every vehicle j in market t that depends on the vehicle's category (car or light truck) and footprint. Revenues and costs of credit transactions enter the objective function in the term $-p_x x$, where p_x is the credit price, and x is the number of credits purchased. A positive value of x represents a credit purchase, which relaxes the constraint. The credit price is endogenously determined by credit supply and demand.

We denote the Lagrange multiplier for constraint (18) by λ. The first-order conditions for price of vehicle k in market t and credit sales are

$$q_{kt} + \sum_{j\in J_{mt}} (p_{jt} - c_{jt})\frac{\partial q_{jt}}{\partial p_{kt}} - \lambda \sum_{j\in J_{mt}} \left(\frac{1}{m_{jt}} - \frac{1}{M_{jt}}\right) \frac{\partial q_{jt}}{\partial p_{kt}} = 0, \tag{19}$$

[10] Leard and McConnell (2017) identify differences in the crediting provisions across the fuel economy and GHG programs, such as over-crediting for plug-in vehicles. However, we abstract from such differences in the model.

$$p_x = \lambda. \tag{20}$$

These first order conditions are used in the marginal cost estimation described in the next section.

In this formulation, the fuel economy standard is equivalent to a government fee-bate that taxes vehicles with low fuel economy and subsidies vehicles with high fuel economy. The pivot point for each vehicle is $1/M_{jt}$ and vehicles are subsidized or taxed by the amount $p_x\left(\frac{1}{m_{jt}} - \frac{1}{M_{jt}}\right)$. The fee-bate and fuel economy standard are equivalent in the sense that they yield the same first order conditions for vehicle price and the same equilibriums. This equivalence is useful in the policy counterfactuals considered in Section 6.

5 Estimation

Estimation consists of three stages. First, we estimate differences in marginal utilities between each demographic group and the base demographic group (that is, β_{kg}), and simultaneously we estimate mean utilities for each vehicle. Second, we estimate the utility parameters of the base group, $\bar{\beta}_k$. Third, we estimate the marginal cost function, $c_{jt}(m_{jt})$. In this section, we first describe the estimation strategies for all three stages, and subsequently we present the estimation results.

5.1 Demand estimation strategy

5.1.1 First stage: Heterogeneous preference parameters

We estimate the marginal utilities for four vehicle characteristics: price, fuel costs (in dollars per mile), performance (defined as the log of horsepower divided by weight), and footprint (the product of wheelbase and track width). The set of vehicle characteristics is similar to the set used in many other demand papers, and it includes the characteristics that are most relevant to fuel economy standards.

As is common in the vehicle demand literature, we compute per-mile fuel costs as the price of gasoline in market t divided by the vehicle's fuel economy. This variable is proportional to the present discounted value of the vehicle's fuel costs if the current price of gasoline equals the expected real price of gasoline over the life of the vehicle. Based on the findings in Leard et al. (2017), we assume that WTP for a fuel economy increase is equal to WTP for an equivalent gasoline price decrease. Therefore, we use variation in fuel prices and fuel economy to identify the fuel cost coefficient, and we use the term WTP for fuel economy synonymously with WTP for a fuel cost reduction.

We define a vehicle by market fixed effect as $\delta_{jt} = \sum_k x_{jkt}\bar{\beta}_k + \xi_{jt}$. Inserting this definition in Equation (16) and allowing for measurement error in market shares yields the estimating equation

$$\ln(s_{gjt}) - \ln(s_{g0t}) = \sum_k \sum_g x_{jkt} h_{igt} \beta_{kg} + \delta_{jt} + \nu_{gjt}. \tag{21}$$

We estimate each of the vehicle market fixed effects as well as the 76 marginal utilities (19 demographic groups and 4 characteristics). Assuming that the measurement error in market shares is uncorrelated with vehicle attributes and fixed effects allows us to estimate Equation (16) by ordinary least squares (OLS).

5.1.2 Second stage: Preference parameters of the base group

In the second stage, we estimate the $\bar{\beta}_k$ terms in Equation (16). It is common in the literature to recover the mean marginal utilities for the vehicle characteristics in a second stage, in which the vehicle-market fixed effect is regressed on these attributes. By analogy, we could estimate the marginal utilities for the base demographic group by regressing the estimated fixed effects $\hat{\delta}_{jt}$ from Equation (21) on the four attributes that vary by vehicle and year.

However, the vehicle market fixed effects include both the observed attributes, x_{jkt} and the unobserved utility, ξ_{jt}. The unobserved utility includes vehicle attributes that are omitted from the utility function, such as cabin comfort. According to the first order condition for price, Equation (19), in the profit maximization problem, vehicle price is an implicit function of the observed and unobserved non-price attributes of other vehicles sold by the manufacturer as well as vehicle's sold by other manufacturers. For example, if a manufacturer redesigns the cabin of one of its vehicles to increase its comfort, the manufacturer may increase the price of the vehicle because the cabin improvement raises consumer demand for the vehicle. Because unobserved cabin comfort is correlated with the observed price, running an OLS regression of the vehicle market fixed effects on observed vehicle price and other attributes would likely yield biased results.

Berry et al. (1995) and many subsequent papers address this issue by instrumenting for a vehicle's price using the attributes of other vehicles, such as performance. The argument is based on the correlation between the price of a vehicle and the attributes of other vehicles in Equation (19). However, Klier and Linn (2012) and Leard et al. (2017) argue that this instrumental variables (IV) approach yields inconsistent estimates if one accounts for the fact that manufacturers choose the observed and unobserved attributes of the vehicles they sell. Consequently, the observed attribute of one vehicle (such as the Toyota Camry) is likely to be correlated with unobserved attributes of another vehicle (such as the Ford Focus). For example, if Ford makes the cabin of the Focus more comfortable, Toyota may simultaneously improve the cabin of the Camry as well as its performance. In this case, the performance of the Camry would be correlated with the (unobserved) cabin comfort of the Focus, making the performance of the Camry invalid as an instrument for the price of the Focus.

We introduce a new approach that builds on the traditional IV approach as well as Klier and Linn (2012) and Whitefoot et al. (2017). Specifically, we form three sets of moment conditions using three sources of plausibly exogenous variation. The first source of variation derives from the first order condition for vehicle price that was discussed in the preceding paragraphs. In practice, typically manufacturers change some vehicle attributes more often than they change other attributes. Attributes related to the power train, such as fuel economy, can be adjusted frequently by re-tuning the engine or replacing components of the engine and transmission. On the other hand, manufacturers change the physical dimensions of the vehicle much less frequently, typically only during major vehicle redesigns that occur every 5-7 years. Based on this regularity, we develop the first moment condition using the sales-weighted 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. Denoting the price and share parameter instruments by $\boldsymbol{z}_{jt}$, the moment condition for price is

$$G^1\left(\bar{\boldsymbol{\beta}}\right) = \sum_{t=1}^{T}\sum_{j=1}^{J_t}(\delta_{jt} - \bar{\boldsymbol{\beta}}\mathbf{x}_{jt})\boldsymbol{z}_{jt}. \tag{22}$$

We argue that this moment condition is valid during the period of time in which a vehicle's width, length, and height are fixed, or roughly 5-7 years. Product entry and exit generates variation within redesigns.

The second and third moment conditions are based on fuel economy and engine performance variation across closely related vehicles. The second source of variation exploits 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. Between 10 and 15 percent of all vehicles in our dataset are a twin, and these vehicles reflect the distributions of attributes for the full choice sets of vehicles. The twins help identify WTP for performance and fuel costs because all other physical attributes of the vehicles are the same, except for the engine size, which affects fuel economy and performance. We form the second moment condition by defining market by vehicle twin fixed effects, τ_{jt}:

$$G^2\left(\bar{\boldsymbol{\beta}}\right) = \sum_{t=1}^{T}\sum_{j=1}^{J_t} x'_{jt}(\delta_{jt} - \bar{\boldsymbol{\beta}}\mathbf{x}_{jt} - \tau_{jt}). \tag{23}$$

For the third moment condition, we define γ_j as the interaction of all of the attributes that define a vehicle except its market. In the data, two vehicles that have a common value of γ_j have the same make, model, trim/series, fuel type, drive type, and body style, as well as the same number of engine cylinders and liters. These vehicles may have different horsepower and fuel economy from one another, due to changes over time in the way a vehicle's engine is tuned or differences in the specific components in the engine or transmission. The moment condition is

$$G^3\left(\bar{\boldsymbol{\beta}}\right) = \sum_{t=1}^{T}\sum_{j=1}^{J_t} x'_{jt}(\delta_{jt} - \bar{\boldsymbol{\beta}}\mathbf{x}_{jt} - \gamma_j). \tag{24}$$

Whereas the second moment condition uses cross-sectional variation in performance and fuel economy across twins sold in the same market, the third moment condition uses time series variation in performance and fuel costs caused by the adoption of fuel-saving technology. This moment condition is similar to the identification strategy in Leard et al. (2017), with one key difference: we do not instrument for fuel economy and performance using engine technology adoption. We avoid instrumenting to exploit all of the variation over time in fuel economy and performance, including technology adoption and engine retuning that is independent of technology adoption. We found that this is necessary to have sufficient variation to identify the mean preference parameters for fuel economy and performance. An implicit assumption with our approach is that changes in fuel economy and performance over time within the same vehicle are uncorrelated with changes in other unobserved attributes that consumers value.

We estimate all second stage parameters in $\bar{\boldsymbol{\beta}}$ jointly using GMM. We estimate the parameters jointly instead of individually because the variation we exploit likely influences multiple endogenous variables. For instance, a vehicle's price may respond to its fuel economy, given empirical evidence by Busse et al. (2013); Langer and Miller (2013); Leard et al. (2017), among others. Furthermore, vehicle prices and fuel costs are correlated with the size of the engine, as illustrated in Table 1. The second and third moment conditions provide complementary variation that improves identification, compared to using one or the other. We stack the moment conditions $G^1(\cdot)$, $G^2(\cdot)$, and $G^3(\cdot)$ and use the two-step GMM estimator.

Thus, the identification strategy addresses the endogeneity of vehicle attributes and price caused by manufacturer choices between vehicle redesigns. The identifying assumptions are that the physical dimensions of other vehicles are uncorrelated with a vehicle's unobserved attributes, and that variations of fuel economy and horsepower within γ_{jt} and τ_{jt} are uncorrelated with the vehicle's unobserved attributes. Importantly, because γ_{jt} and τ_{jt} control for model and trim, the coefficient estimates are consistent even if manufacturers package the engine and transmission configuration with a particular trim. For example, many manufacturers offer a "sport" trim that includes a larger engine than the standard trim. Even if the sport trim differs in other unobserved dimensions from the standard trim, such as cabin features or exterior styling, γ_{jt} and τ_{jt} control for such differences.

Klier and Linn (2012) also address this source of endogeneity using proprietary information about engine attributes. Our method does not require engine data and could be used by researchers who do not have access to such data. Whitefoot et al. (2017) also instrument for endogenous vehicle attributes, but they use as IVs power train attributes that can vary between redesigns, potentially yielding inconsistent estimates.

5.2 Supply estimation strategy

Next, we estimate marginal production costs using the first-order conditions for vehicle prices and manufacturer net credit sales. Substituting Equation (20) into Equation (19) eliminates the Lagrange multiplier from the price first-order condition:

$$q_{kt} + \sum_{j \in J_{mt}} (p_{jt} - c_{jt}) \frac{\partial q_{jt}}{\partial p_{kt}} - p_x \sum_{j \in J_{mt}} \left(\frac{1}{m_{jt}} - \frac{1}{M_{jt}} \right) \frac{\partial q_{jt}}{\partial p_{kt}} = 0. \tag{25}$$

Given observed credit prices from Leard and McConnell (2017), we can use this first-order condition to compute the marginal cost for each vehicle and market, c_{jt}.

Because the simulations involve changes in fuel economy, we need to estimate the relationship between marginal costs and fuel economy. Following Leard and McConnell (2017), we specify a log-log model, and we control for other vehicle attributes:

$$lnc_{jt} = \delta lnm_{jt} + X_{jt}\eta + \sigma_j + \varsigma_t + \mu_{jt}. \tag{26}$$

The fuel economy coefficient, δ, is the elasticity of marginal costs to fuel economy. We estimate this equation by OLS, expecting a positive estimate of δ. The vector X_{jt} includes variables that may be correlated with fuel economy and that also affect marginal costs: the log of horsepower, the log of weight, and interactions of market fixed effects with body type fixed effects. The equation also includes vehicle fixed effects to absorb fixed variation across models, trims, fuel types, and body types.

5.3 Estimation results

This subsection presents the estimation results for the demand and supply sides of the model. Because the large number of preference coefficients are difficult to interpret, we focus on the implied own-price elasticities of demand and the WTP for fuel economy and performance. The fixed effects control for fixed variation across vehicles in marginal costs.

5.3.1 Demand estimates

We first report the estimation results and then assess the model's ability to reproduce observed consumer choices. Figure 7 illustrates the estimated own-price elasticity of demand; Appendix Table A.5 reports the corresponding numbers for reference. The figure shows that across groups, the average own-price elasticity is about -3.4, which is consistent with estimates in the literature (Berry et al. 1995; Train and Winston 2007). Moreover, the own-price elasticity of demand varies substantially across demographic groups. Low income groups tend to be more sensitive to prices than high income groups, ranging from about -5 for the lowest group to about -2 for the highest group. This variation across groups is similar to the variation across households estimated in mixed logit models (Train and Winston 2007). The similarity suggests that the demand model captures a substantial amount of preference heterogeneity across households. Income appears to explain most of the variation in own-price elasticities of demand, as the elasticities are similar within income groups and across urbanization and age groups.

Figure 8 shows the WTP for fuel economy and performance implied by the estimated utility function coefficients; Appendix Table A.5 reports the corresponding numbers for reference. The figure plots the WTP for a 1 percent increase in fuel economy or horsepower for each demographic group. The mean WTP for fuel economy across demographic groups is about $68. This estimate is roughly half of that reported in Leard et al. (2017), who also use MaritzCX data but identify WTP for fuel economy from fuel economy changes over time caused by manufacturer adoption of fuel-saving technology. The difference arises partly from the fact that we estimate a larger own-price elasticity of demand than they assume in their

WTP calculation. Using their demand elasticity would increase WTP for fuel economy by about 20 percent.[11]

Overall, low income groups have lower WTP for fuel economy than high income groups. There is also substantial variation across age and urbanization groups within income groups; for example, the young age group typically has lower WTP. We believe that the vehicle demand literature has not previously quantified the variation in WTP across demographic groups.

To provide an economic interpretation of the fuel economy WTP estimates, Table 2 reports valuation ratios as in Leard et al. (2017). The valuation ratio is the ratio of the WTP for a 1 percent fuel economy increase to the present discounted value of the fuel savings arising from the fuel economy increase. A valuation ratio of 1 would imply full valuation. The Appendix describes the calculations in detail. Importantly, the calculations account for the fact that higher income groups typically drive their vehicles more miles and have lower discount rates because they have lower borrowing costs; see Appendix Tables A.1 and A.2 for variation in borrowing costs across demographic groups.

The valuation ratios rise with income, which is consistent with the correlation between income and WTP shown in Figure 2. However, note that the valuation ratio varies somewhat less across income groups than does WTP, which is because lower income groups tend to drive their vehicles fewer miles and have higher discount rates. In other words, lower income groups have lower WTP partly because they have higher borrowing rates and drive fewer miles, but these factors only partly explain the estimated WTP variation.

Panel B of Figure 8 illustrates the WTP for performance by demographic group. The mean across demographic groups is about $83, which is similar to the mean WTP reported in Leard et al. (2017). Consistent with intuition, higher income groups have higher WTP for performance. The younger age group typically has lower WTP than the older age group, and urban households often have higher WTP than rural households. As with fuel economy, we believe this is the first quantitative assessment of variation in WTP for performance across demographic groups.

Figures 9 and 10 assess the model's ability to reproduce observed consumer choices. We begin by using the demand model and observed vehicle attributes to predict market shares of

[11] The differing identification strategy may also play a role, as consumers may respond differently to variation in fuel economy and performance over time caused by technology adoption, as used in their paper, than to other sources of variation that we use. Unfortunately, we cannot use their identification strategy to estimate the demand parameters, because the data do not contain sufficient variation to identify price and fuel cost coefficients from technology adoption.

each vehicle and market, according to Equation (15). We use the predicted market shares to compute the market share-weighted average of each attribute, by demographic group. Figure 9 plots the predicted attribute means against the observed sales-weighted averages. The fact that the predicted values fall close to a 45-degree line indicates that the model accurately predicts means of the attributes across demographic groups.

Figure 10 provides additional evidence on the model's performance. There is a substantial amount of vehicle entry and exit in the data, which complicates efforts to evaluate the model's out-of-sample prediction power, because entry and exit are exogenous to the model. Because there is little brand entry and exit in our data, for this analysis we aggregate vehicles by brand and class. Panel A is a no-change forecast; we predict market shares in the final market of our sample, $t = 2015$, using the observed brand-class market share from the first market of our sample, $t = 2010$. Panel A plots the predicted market shares against the observed market shares in $t = 2015$. The brand-class market shares are fairly stable over time, as the predicted market shares are strongly correlated with the observed market shares. However, there is a fair bit of scatter in Panel A, and in fact there is considerably less scatter if we use the demand model to predict 2015 market shares as in Panel B. Thus, the demand model in Panel B outperforms the no-change forecast in Panel A.

As we noted in Section 5.1, in our model, for a particular demographic group the utility function parameters do not vary across households or the vehicles they choose. To validate our assumption, we re-estimate the model on a subset of our data. For each vehicle and market, we randomly sample 10 of the 20 demographic groups, and re-estimate the model using the chosen sub sample. Then, we use the estimated coefficients to compute the market shares out of sample. If the demand model assumptions were not valid, we would predict poorly the out-of-sample market shares. Panel C shows that this is not the case; there is a strong correlation between the predicted and observed market shares, and in fact the correlation is nearly as strong as when we use the full sample to estimate the demand parameters as in Panel B.

5.3.2 Supply estimates

We use the first order condition for vehicle price, Equation (25), to compute each vehicle's marginal costs, c_{jt}. In that equation all variables are observed in the data or are computed from the demand estimates, except for the price of compliance credits, p_x. Leard and McConnell (2017) report compliance credit prices for the years 2012 through 2014. We use these credit prices for each year in our sample. For the 2015 market, we assume a credit price equal to the 2014 price reported in their paper. As the first order condition illustrates, the credit price represents the shadow cost on the fuel consumption rate created by the standards (the fuel consumption rate is the reciprocal of fuel economy). Unfortunately, because cross-firm credit transactions were not allowed prior to 2012, we cannot observe the shadow cost for 2010 and 2011. For that reason, we do not compute c_{jt} for 2010 and 2011.

Table 3 shows the estimated fuel economy coefficient in Equation (26), which is the elasticity of the vehicle's marginal costs to its fuel economy. Each column and panel reports the results of a separate regression. Observations are by vehicle and market, and the dependent variable is the marginal costs computed from Equation (25). In addition to log fuel economy, all regressions include log horsepower, log weight, vehicle fixed effects, and the interaction terms described in the bottom row of the table. Column 1, which we consider to be the baseline specification, shows an elasticity of marginal costs to fuel economy of about 0.25 for cars and 0.18 for light trucks. These estimates are fairly stable across the remaining columns, which include additional controls that may be correlated with marginal costs and fuel economy.

As we describe in the Appendix, our estimates are similar to those we obtain from NHTSA estimates of technology costs. We prefer to use the estimates from Table 3 rather than the estimates from the NHTSA data because the estimates in Table 3 are internally consistent with the other estimated parameters. The welfare results in the next section are similar if we use the NHTSA-based estimates instead.

6 Welfare Results

This section reports the results of simulating a hypothetical increase in fuel economy mandated by regulation. We compare changes in consumer welfare across demographic groups and changes in profits across manufacturers. The main conclusions confirm the results of the analytical model from Section 2: All else equal, pass-through is positively correlated with WTP for fuel economy; manufacturers selling to consumers with higher WTP experience larger profits increases; and consumers with higher WTP experience larger welfare gains.

6.1 Primary model specification and counterfactual

This subsection describes the setup of the baseline and policy scenario and reports the results.

6.1.1 Definition of the baseline and primary policy counterfactual

The baseline and policy scenarios conform to the economic environment in which the demand and supply parameters are estimated. We focus on a single market, choosing the year 2012 because that was the year in which the fuel economy and GHG standards began tightening for both cars and light trucks. It is also the first year in which we observe the the credit prices and are able to compute marginal costs.

Given the vehicle prices and attributes observed in 2012, we use Equation (15) to compute market shares. We use credit prices observed om 2012 to compute marginal costs. We use these values as well as marginal costs to compute mean utilities for each demographic group and profits for each manufacturer. Because of the equivalence between a fuel economy standard and a fee-bate, we interpret the baseline equilibrium as arising from a fee-bate equal to the observed credit price.

The primary policy counterfactual consists of an exogenous 1 percent fuel economy increase for all vehicles. This scenario corresponds to a fuel economy standard that does not allow for credit trading across vehicles. The scenario is motivated by the fact that manufacturers have relied heavily on meeting GHG standards in the US and Europe (Klier and Linn (2016) and Reynaert (2017)).[12] In the counterfactual policy scenario, each vehicle's fuel economy increases by 1 percent from its 2012 level. We treat fuel economy as exogenous and uniform across vehicles to isolate the effects of cross-household heterogeneity in WTP.

More specifically, we take advantage of the equivalence between a fuel economy standard and fee-bate. Each vehicle's fuel economy increases by an exogenous 1 percent. In addition to the fuel economy increase, each vehicle sold in 2012 is subject to the same fee-bate as in the baseline. We compare the equilibrium with the observed standards (i.e., fee-bate, the baseline scenario) and the equilibrium with the fee-bate plus fuel economy improvement (i.e., the policy scenario). The differences in outcomes across the two scenarios represent the effect of the exogenous fuel economy increase.

The higher fuel economy raises each vehicle's marginal costs according to Equation (26). The Appendix explains the algorithm that we use to solve for the profit-maximizing vehicle prices of each manufacturer. Given these prices, we use Equation (15) to compute market shares, and then compute the changes in consumer welfare (that is, equivalent variation)

[12] Alternatively, we could model a tightening of the fee-bate. In that case, fuel economy would be endogenous. We prefer the main scenario because it isolates the effects of technology adoption and conforms to the case considered in the analytical model, in contrast to the fee-bate, which would also incentivize manufacturers to adjust vehicle prices.

by demographic group and changes in profits by manufacturer, relative to the baseline equilibrium.

6.1.2 Results

Table 4 and Figure 11 show the effects of the mandated higher fuel economy on consumer welfare and manufacturer profits. Table 4 shows that the fuel economy mandate raises total consumer welfare by about \$172 million and raises profits by about \$133 million. The fact that both consumers and manufacturers are better off with the higher fuel economy arises from the fact that the WTP for a 1 percent fuel economy increase (\$68) exceeds the estimated marginal cost increase (\$48).[13]

Figure 11 indicates that these benefits are not uniformly distributed across demographic groups and manufacturers. High-income urban households benefit substantially, whereas lower income groups benefit less. Typically, urban households benefit more than rural households. Some households, especially lower-income rural households, experience welfare losses. The figure shows absolute welfare changes, and if we normalize the welfare changes by mean income in each group we conclude that the standards are regressive, meaning that low-income households benefit less relative to their income than do high-income households.

Changes in profits also vary across manufacturers. The US-based manufacturers (GM, Ford, and Fiat-Chrysler) benefit substantially, as does Daimler. The three largest Japanese manufacturers (Toyota, Nissan, and Honda) benefit somewhat less than those manufacturers, and several of the manufacturers experience welfare losses. For reference, Appendix Tables A.6 and A.7 report the numerical values for each demographic group and manufacturer.

The analytical model in Section 2 indicates that WTP for fuel economy can explain the variation in pass-through and welfare effects across demographic groups and manufacturers. The pass-through rate for a demographic group is the average pass-through for the vehicles purchased by the demographic group, weighted by the vehicle's share of purchases in total purchases by the demographic group. According to the model, pass-through rates should be higher for demographic groups with higher WTP for fuel economy than for other demographic groups. Panel A displays a positive correlation between WTP for fuel economy and the pass-through rate across demographic results, which is consistent with the model's result. Moreover, the model suggests that the curvature of the demand curve should be positively

[13] This result would appear to imply that there exists a market failure for fuel economy, and that tighter fuel economy or GHG standards would increase private welfare. However, Leard et al. (2017) show that this inference is not correct because it does not account for changes in performance caused by the standards. Below, by including the trade-off, we report results that confirm their conclusion.

correlated with pass-through, and in fact Panel B shows a positive correlation between the own-price elasticity of demand and the pass-through rate.[14]

Figure 13 shows the correlations among welfare effects, WTP, and own-price elasticity of demand for demographic groups and manufacturers. Panel A shows a weak correlation between WTP and the welfare change by demographic group. Panel B shows a weak correlation between the own-price elasticity and the welfare change. The correlations do not change materially if one excludes the outliers to the far right-hand side of the diagrams.

The weak correlations would appear to be surprising, given the conclusions from the analytical model. Panel C provides intuition for the weak correlations, showing a strong correlation between WTP and the own-price elasticity of demand. The strong correlation creates opposing effects on consumer welfare that roughly offset one another. On the one hand, demographic groups with high WTP experience larger welfare gains than groups with lower WTP. On the other hand, demographic groups with more inelastic demand should experience smaller welfare gains than groups with more elastic demand. As Panel C shows, it turns out that the demographic groups with high WTP also tend to have more inelastic demand, which causes the two effects to oppose one another. The magnitudes roughly cancel one another causing welfare changes to be uncorrelated with pass-through. However, we find that the mandatory fuel economy increase is regressive, in that high-income groups benefit more than low-income groups.

Turning to manufacturers, according to the model in Section 2, manufacturers selling to consumers with higher WTP should experience higher pass-through rates than other manufacturers. Panel C of Figure 12 shows a positive correlation between the average WTP of each manufacturer's consumers and the pass-through rate. Interestingly, Panel D shows approximately no correlation between the own-price elasticity of demand and the pass-through rate for manufacturers.

Panel D of Figure 13 shows a strong positive correlation between the WTP for fuel economy and the change in profits per vehicle. In Panel E, there is a weak correlation between the own-price elasticity of demand and the change in profits, and Panel F shows a weak correlation between WTP and own-price elasticity. Thus, for manufacturers, WTP appears to be more important than the own-price elasticity of demand in explaining the

[14] According to the model in Section 2, the second derivative of demand with respect to price is negatively correlated with the pass-through rate. In the figures and tables, we report own-price elasticities of demand rather than second derivatives, because the elasticities are more intuitive and are widely reported in the literature. In practice, the own-price elasticities are strongly correlated with the second derivatives; a larger own-price elasticity (in magnitude) implies a larger second derivative (in magnitude).

variation across manufacturers in changes in profits per vehicle. The positive correlation in Panel D is consistent with the conclusions from the analytical model.

6.2 Alternative model calibrations and scenarios

This subsection reports results from alternative model calibrations that provide additional context for the results in the previous subsection.

6.2.1 Zero willingness to pay for fuel economy

The analytical model emphasizes the role of WTP for fuel economy in determining pass-through rates and welfare effects of a fuel economy increase. In the model, WTP affects pass-through and welfare, whereas the previous literature has emphasized the role of the shape of the demand curve in determining pass-through and welfare effects. In fact, the simulation results for consumers in Figure 13 indicated that the own-price elasticity may explain some of the variation in welfare effects across demographic groups.

To assess the role of the own-price elasticity of demand in the welfare results and relate our results to the literature, we re-estimate the demand model imposing the constraint that the coefficient on fuel costs equals zero for all demographic groups. This constraint means that consumers do not place any valuation on fuel economy. More generally, this constraint corresponds to a situation in which a regulation affects a product attribute that consumers do not value; in that case, the standard conclusions would apply, that the shape of the demand curves governs pass-through and welfare changes. Comparing this version of the model with the full version is instructive because the comparison illustrates the role of the WTP for fuel economy in determining the welfare results.

Panel A compares the estimated own-price elasticity of demand for this version of the model with the full version, for each demographic group. Although the magnitude of the estimates for the zero-WTP model are smaller than for the full model, the pattern across demographics groups is identical. This similarity suggests that any differences in the welfare effects in this version of the model and the full version arise from setting the WTP equal to zero, rather than from changes in the estimated demand elasticities.

Table 4 shows the aggregate welfare effects of the simulated 1 percent fuel economy increase, for the version of the model with zero WTP for fuel economy. Raising fuel economy by 1 percent reduces total consumer welfare and manufacturer profits. This result is not surprising. Unlike in the full model with non-zero WTP, with zero WTP manufacturers cannot pass along as much of the cost increase to consumers because the consumers do not value the fuel economy increase, and the lower pass-through rate means that manufacturer profits decline. Because consumers do not value the fuel economy increase, they do not benefit from the higher fuel economy and their welfare declines because of higher vehicle prices.

Panel B of Figure 13 plots the change in consumer welfare against the own-price elasticity of demand for each demographic group. There is a strong negative correlation between the welfare change and the own-price elasticity of demand, indicating that consumers with more price-inelastic demand experience larger welfare decreases. This result is consistent with the standard theory on pass-through and welfare. This result also supports the argument we made above when discussing the lack of correlation between welfare changes and own-price elasticity of demand for the full version of the model. We argued that the effects of WTP and own-price elasticity offset one another, and this argument is confirmed by the fact that we find a negative correlation when we impose the restriction of zero WTP for fuel economy.

6.2.2 Alternative policy scenarios

Here, we discuss two alternative definitions of the policy scenarios, which yield similar conclusions as the primary scenario considered previously. First, we consider a policy that raises each manufacturer's fuel economy requirement proportionally, rather than raising each vehicle's fuel economy by a uniform 1 percent. The proportional assumption is similar to the NHTSA analysis of fuel economy standards.

Recall that each manufacturer faces a fuel economy requirement that depends on the footprint and class of its vehicles; manufacturers selling more light trucks and larger vehicles face a lower fuel economy requirement than others. In the proportional scenario, each vehicle's fuel economy increases in proportion to the manufacturer's fuel economy requirement in 2012, so that fuel economy increases by more (in percentage terms) for vehicles sold by manufacturers with higher requirements than for other manufacturers. For example, if a manufacturer faces a requirement that is 10 percent lower than the requirement of another manufacturer, the counterfactual raises fuel economy of vehicles sold by the first manufacturer by 10 percent less than that for the second manufacturer. This situation would correspond to a policy that raises fuel economy more quickly for smaller vehicles than for other vehicles, and corresponds roughly to the actual structure of the standards in recent years.

Appendix Table A.8 compares the changes in profits per vehicle by manufacturer for the primary and proportional scenarios. The changes in profits are similar across the two scenarios, as manufacturers that experience large profits increases in the uniform scenario are likely to experience large profits increases in the proportional scenario. The differences in fuel economy changes explain the differences in profits changes across the two scenarios. For example, GM experiences a larger fuel economy increase in the proportional scenario than the primary scenario, which causes its profits to increase by more for the proportional scenario. The table also shows that market shares by manufacturer are similar to one another in the two scenarios.

Second, we allow for the possibility that requiring higher fuel economy induces trade-offs between fuel economy and horsepower. As noted above, during historical periods in which the stringency of fuel economy standards did not change, manufacturers steadily improved the energy efficiency of power trains, which allowed the manufacturers to increase performance without reducing fuel economy. However, Klier and Linn (2016) conclude that increasing the stringency of the standards causes manufacturers to increase horsepower less quickly than they would have. These results imply that tighter standards induce a trade-off between fuel economy and horsepower, in that horsepower is lower and fuel economy is higher when standards are tighter than when they are weaker. A straightforward extension of the analytical model in which the regulation affects multiple product attributes would show that pass-through and welfare changes would depend on the WTP for all attributes affected by the regulation.

We construct a scenario that approximates this trade-off. Specifically, we assume that half of the required 1 percent fuel economy increase in the primary scenario is achieved by trading off fuel economy and performance. We compare the observed equilibrium in 2012 with a counterfactual equilibrium that includes lower performance and 1 percent higher fuel economy. The performance change is computed from the estimated trade-off between fuel economy and performance, where we have updated the trade-off estimation in Klier and Linn (2016) using more recent data. We expect welfare increases to be lower in this counterfactual than in the primary one, because this counterfactual accounts for the fact that consumers forgo performance increases that they would have enjoyed had the standards not been tightened.[15]

Figure 15 shows that all consumer groups and manufacturers experience welfare losses in the scenario with trade-offs. In addition, the ranking of the welfare changes across groups or firms is different in the two scenarios. Variation in WTP for performance across demographic groups and manufacturers explains the change in ranking. For example, urban households

[15]Modeling such a trade-off explicitly requires a dynamic treatment of the fuel economy standards. However, such an analysis would require us to model technology adoption and the trade-offs between fuel economy and horsepower, which lies outside the scope of the paper. As noted above, we treat fuel economy as exogenous to focus on the role of WTP in the welfare effects.

in the highest income group experience the largest welfare gains without trade-offs, but they experience large welfare losses with trade-offs. The result follows from the fact that these demographic groups have relatively high WTP for performance, as seen in Figure 8. Relative to household income, the standards are regressive when performance is held fixed and are roughly neutral when we include the trade-offs.

Table 4 shows that the scenario with trade-offs reduces consumer welfare and manufacturer profits, in contrast to the increases observed in the uniform scenario that does not include the trade-offs. This result suggests that modeling the trade-offs is important for welfare analysis of the standards, which we leave for future work. Thus, the main results underscore the importance of WTP in determining the welfare effects of regulations that affect product attributes.

7 Conclusions

In this paper, we analyze regulations that affect product attributes and show analytically that the pass-through of regulatory costs and private welfare effects depend on consumer WTP for the attribute. The higher the WTP for the attribute, the greater the pass-through rate, the greater the change in firm profits, and the greater the change in consumer surplus. To the best of our knowledge, these are new results in the literature on pass-through and welfare effects of policies in imperfectly competitive markets.

We evaluate the relevance of these theoretical results in the context of fuel economy standards for light-duty vehicles, which affect fuel economy and possibly other vehicle attributes. We build an equilibrium model of the new vehicles market that includes a highly disaggregated vehicle choice set and allows preferences to vary across demographic groups. We use the model to simulate the effects of tightening fuel economy standards.

The results confirm the intuition provided by the analytical model. Demographic groups with higher WTP have higher average pass-through rates and higher welfare increases than do other groups. Likewise, across manufacturers, pass-through rates and profits increases are positively correlated with the WTP of their consumers. Moreover, accounting for WTP breaks the link between pass-through and welfare changes that the previous literature has emphasized. If we ignore WTP, high pass-through implies a large welfare loss. However, once we account for WTP, pass-through is uncorrelated with welfare changes across demographic groups.

The empirical findings depend on consumer valuation of fuel costs and performance. We estimate that consumers undervalue increases in fuel economy and highly value performance, a result that is consistent with Leard et al. (2017). Our finding that consumers undervalue fuel economy gains contrasts with other recent reduced-form estimates that use earlier data and gasoline price variation for identification (for example, Busse et al. (2013)). These contrasting findings motivate further analysis of consumer valuation of fuel economy.

Our analysis has several limitations that future work could address. First, we assume that fuel economy and performance changes are exogenous. We find that the standards are regressive when we hold performance fixed and that they are roughly neutral when we account for trade-offs between performance and fuel economy. Extending our model to endogenize these attributes is therefore an important direction for future work.

Second, we do not explicitly model the effects of fuel economy standards on prices of used vehicles. Although we do model the substitution between new and used vehicles, used vehicle price adjustments can have important welfare implications (Jacobsen 2013). Future work could model explicitly the used vehicle market.

Our theoretical and empirical framework can be applied to other settings besides assessing the passenger vehicle fuel economy standards. Other policies affect the passenger vehicle market and vehicle attributes, such as the zero emission vehicle program, which mandates sales of plug-in and fuel cell vehicles in California and other states. More broadly, a wide range of regulations affect product attributes that consumers value, such as appliance standards. Future work could examine whether variation in WTP across demographic groups plays an important role in the welfare effects of such regulations.

References

Ackerberg, 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., S. Mullainathan, and D. Taubinsky (2014). Energy policy with externalities and internalities. *Journal of Public Economics 112*, 72–88.

Allcott, H. and C. Sunstein (2015). Regulating internalities. *Journal of Policy Analysis and Management 34*(2), 698–705.

Bento, A., L. Goulder, R. V. Hafen, and M. Jacobsen (2009). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review 99*(3), 1–37.

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.

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.

Davis, L. W. and C. R. Knittel (2016). Are fuel economy standards regressive? *NBER Working Paper No. 22925*.

EPA (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. Technical report, Environmental Protection Agency.

EPA (2018). Preliminary regulatory impact analysis for the safer affordable fuel-efficient (safe) vehicles rule for model year 2021 - 2026 passenger cars and light trucks. *Environmental Protection Agency and National Highway Traffic Safety Administration*.

Ganapati, S., J. Shapiro, and R. Walker (2018). The incidence of carbon taxes in us manufacturing: Lessons from energy cost pass-through. *NBER Working Paper No. 22281*.

Greene, D., A. Hossain, J. Hofmann, G. Helfand, and R. Beach (2018). Consumer willingess to pay for vehicle attributes: What do we know? *Transportation Research Part A: Policy and Practice 118*, 258–279.

Houde, S. (2018). How consumers respond to product certification and the value of energy information. *The RAND Journal of Economics 49*(2), 453–477.

Jacobsen, M. (2013). Evaluating U.S. fuel economy standards in a model with producer and consumer heterogeneity. *American Economic Journal: Economic Policy 5*(2).

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. R. and K. Metaxoglou (2014). Estimation of random-coefficient demand models: Two empiricists' perspective. *The Review of Economics and Statistics 96*(1), 34–59.

Langer, A. and N. Miller (2013). Automakers' short-run responses to changing gasoline prices. *Review of Economics and Statistics 95*(4), 1198–1211.

Leard, B., J. Linn, and V. McConnell (2017). Fuel prices, new vehicle fuel economy, and implications for attribute-based standards. *Journal of the Association of Environmental and Resource Economists 4*(3), 659–700.

Leard, B., J. Linn, and Y. Zhou (2017). How much do consumers value fuel economy and performance? Evidence from technology adoption. *Resources for the Future Report*, 1–37.

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.

Levinson, A. (2016). Energy efficiency standards are more regressive than energy taxes: Theory and evidence. *NBER Working Paper No. 22956*.

Muehlegger, E. and R. Sweeney (2017). Pass-through of input cost shocks under imperfect competition: Evidence from the us fracking boom. *NBER Working Paper 24025*.

NHTSA (2018). Preliminary regulatory impact analysis of the safer affordable fuel-efficient (safe) vehicles rule for model year 2021-2026 passenger cars and light trucks. Technical report, National Highway Traffic Safety Administration.

Petrin, A. (2002). Quantifying the benefit of new products: the case of minivan. *Journal of Political Economy 110*, 705–729.

Reynaert, M. (2017). Abatement strategies and the cost of environmental regulation: Emission standards on the european car market. *Working paper*.

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.

Weyl, E. and M. Fabinger (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy 121*(3), 528–583.

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.

Figures

Figure 1: The Effect of MWTP for an Attribute on Pass-Through and Welfare

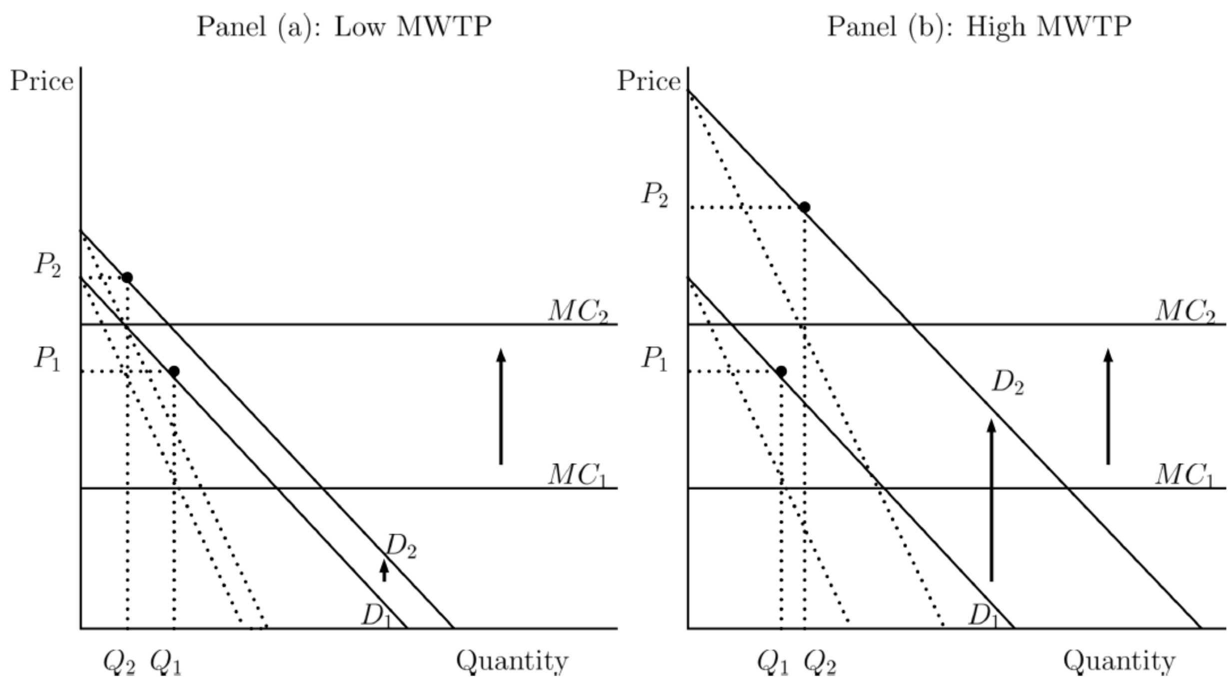

Figure 2: Means of Vehicle Attributes by Demographic Group

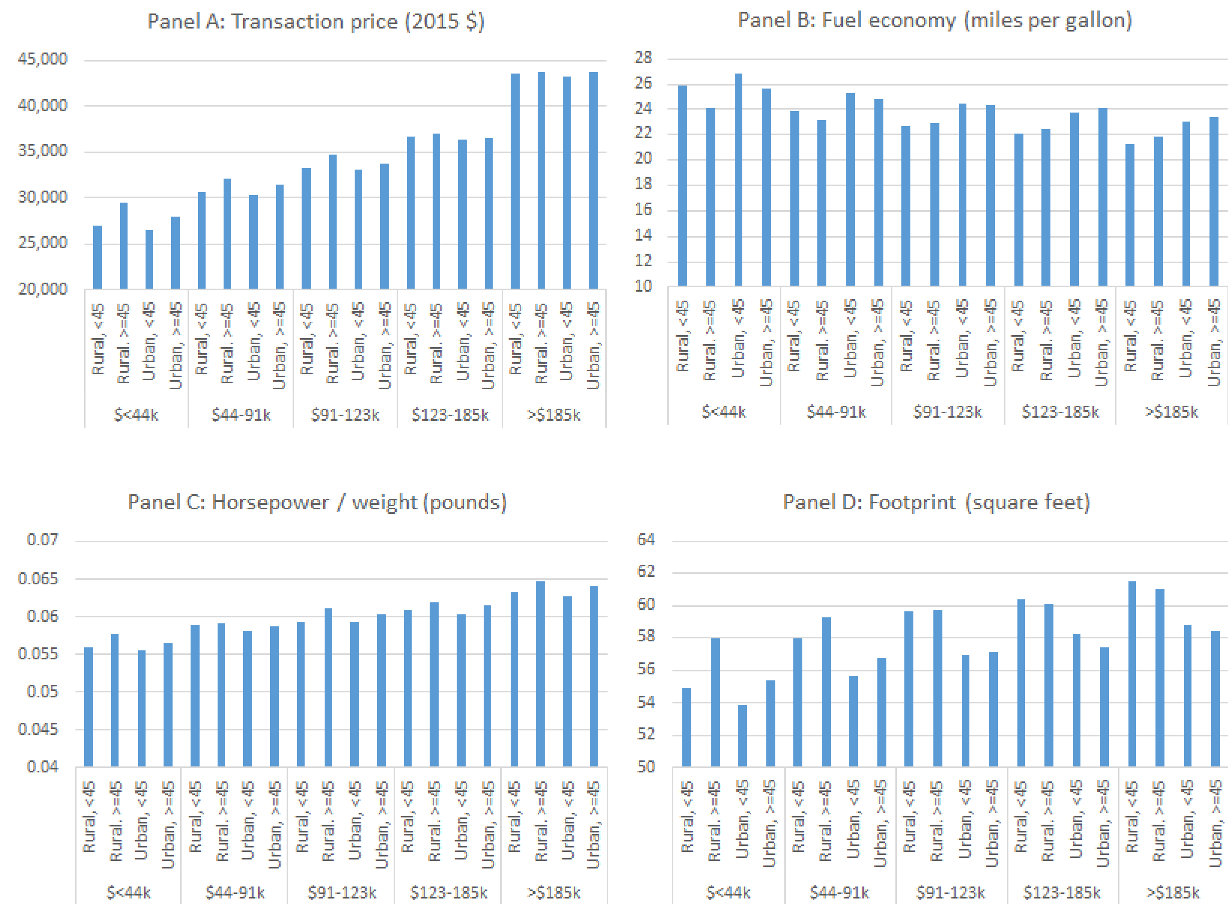

Notes: The figure shows the sales-weighted mean attribute for each demographic group. The sample includes all vehicles purchased between 2010 and 2015.

Figure 3: Shares of Light Trucks, Hybrids, Plug-ins, and Used Vehicles by Demographic Group

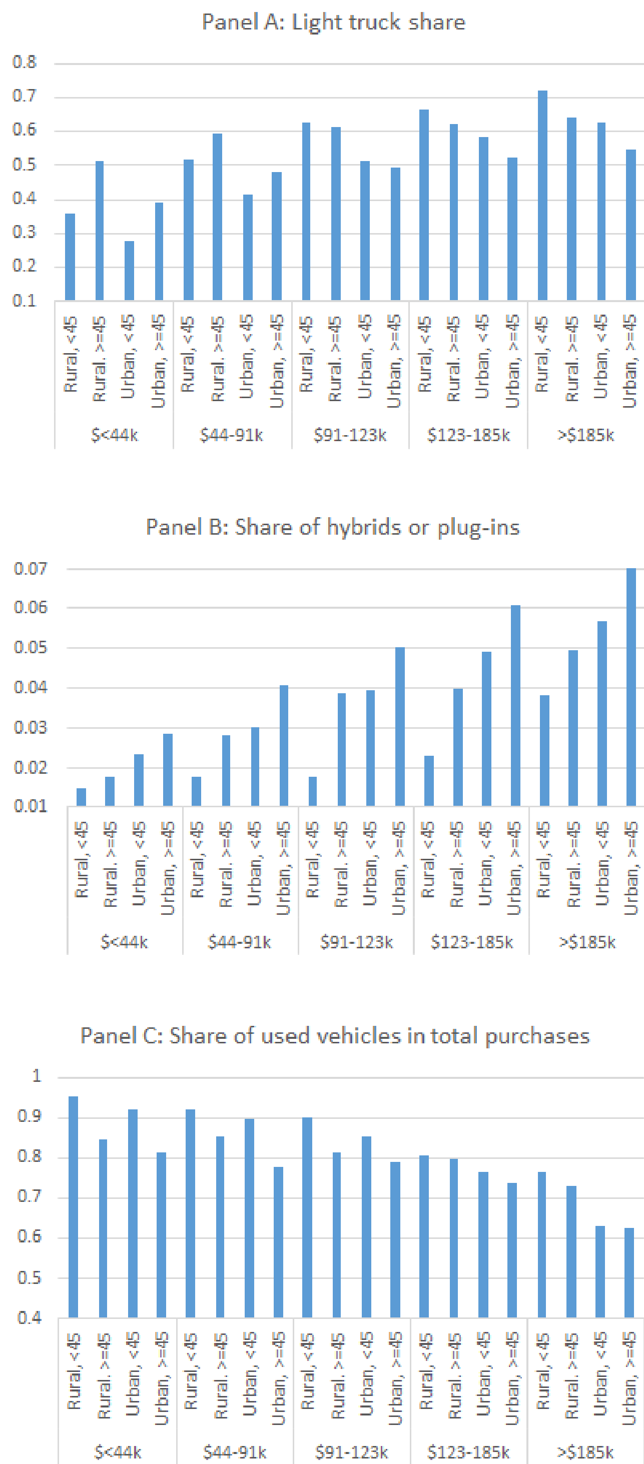

Notes: The figure shows the sales-weighted market share for each demographic group. The sample includes all vehicles purchased between 2010 and 2015.

Figure 4: Means of Vehicle and Household Attributes by Firm

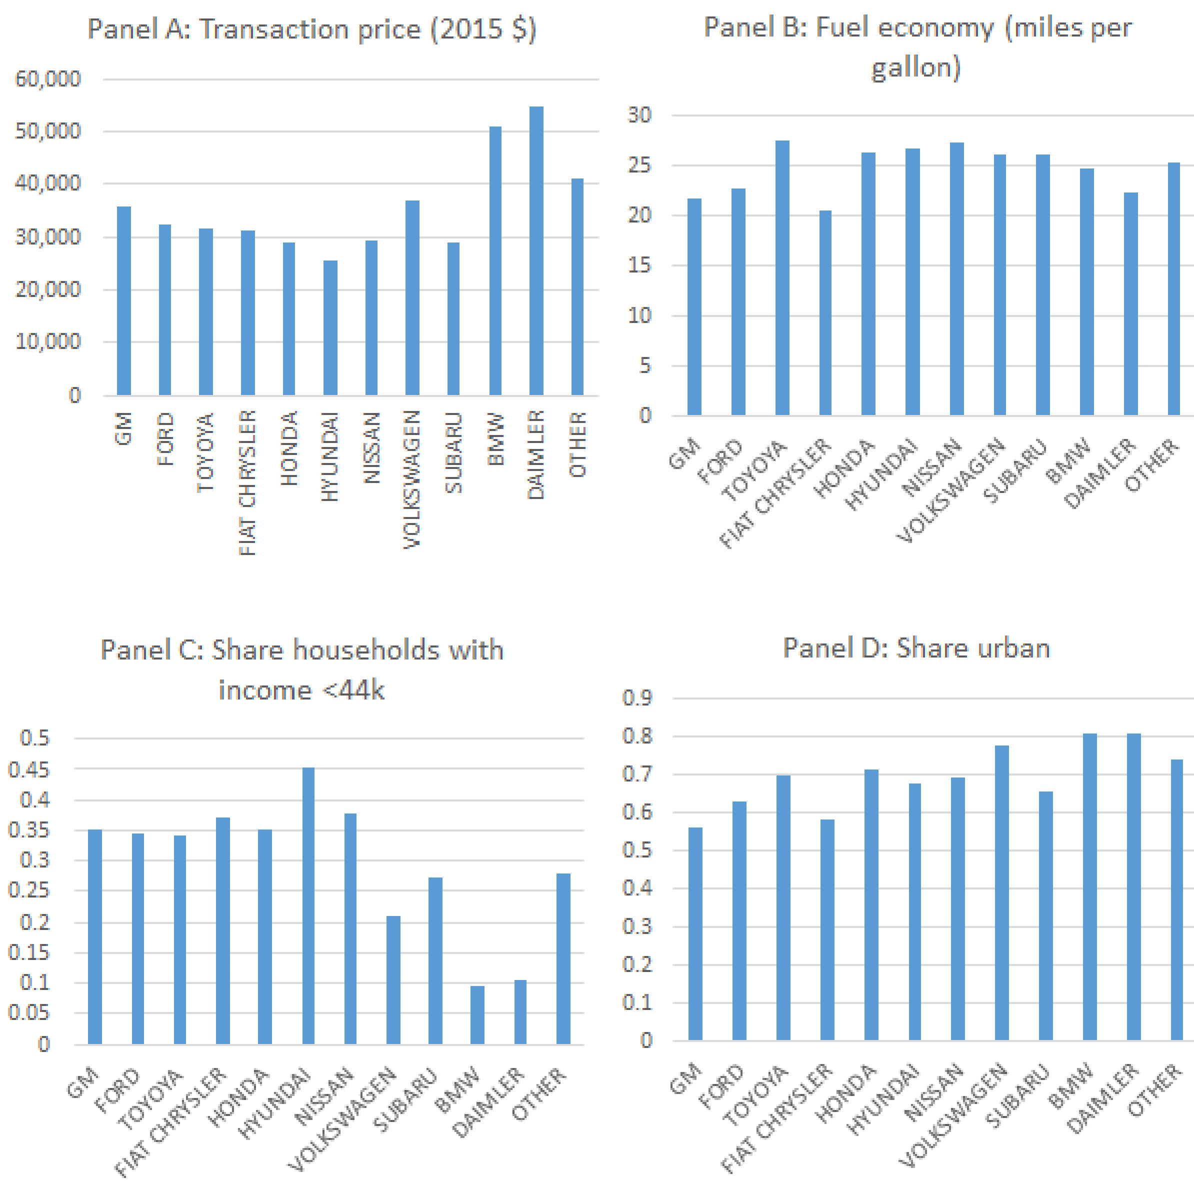

Notes: The figure shows the purchases-weighted mean attribute or market share for each manufacturer. The sample includes all vehicles purchased between 2010 and 2015.

Figure 5: Vehicle Attribute Densities: Engine Twins and the Full Sample

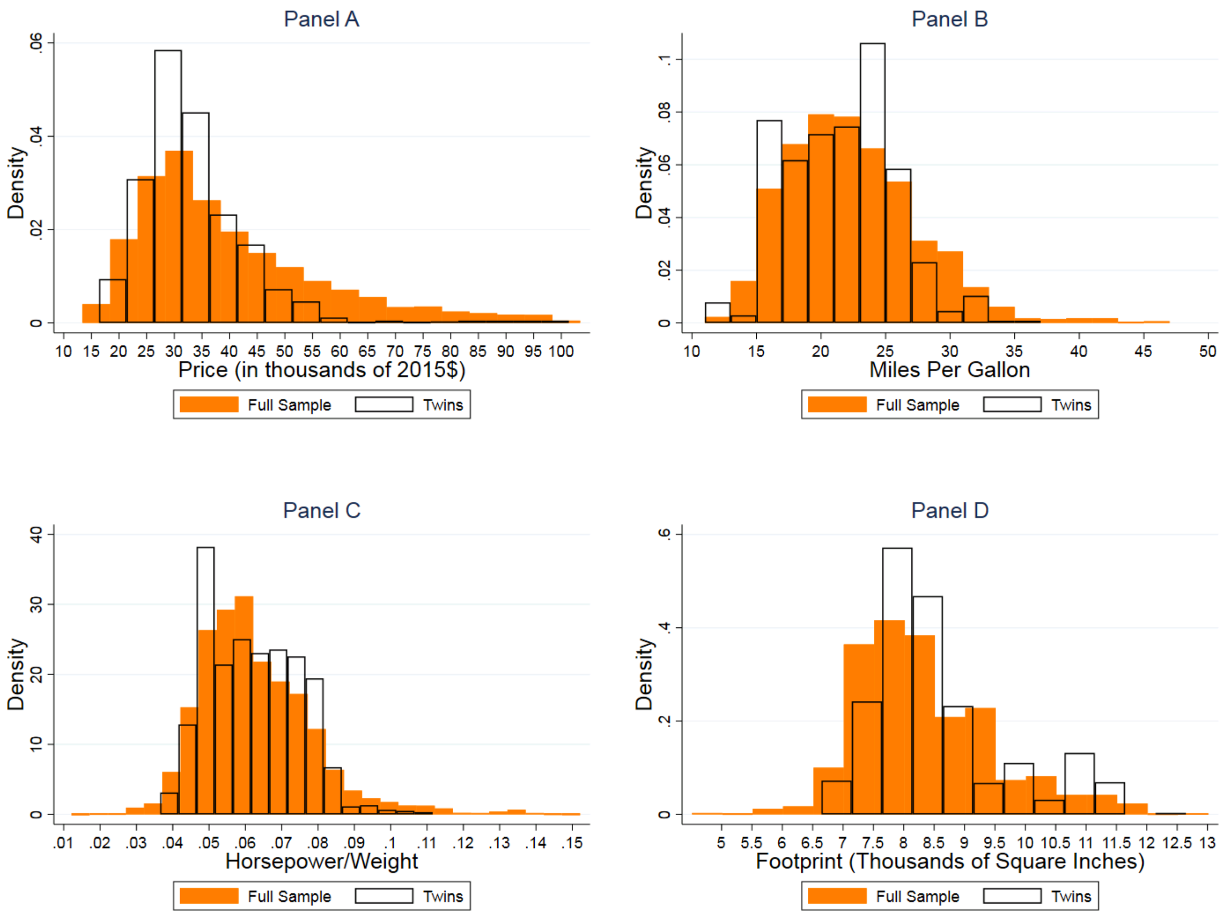

Notes: The figure shows the unweighted densities of vehicle attributes for the full sample of vehicles and engine twins sold between 2010 and 2015. Footprint is the product of the vehicle's wheelbase and width.

Figure 6: Fuel Economy for Vehicles Sold in 2010 and Standards in 2012 and 2015

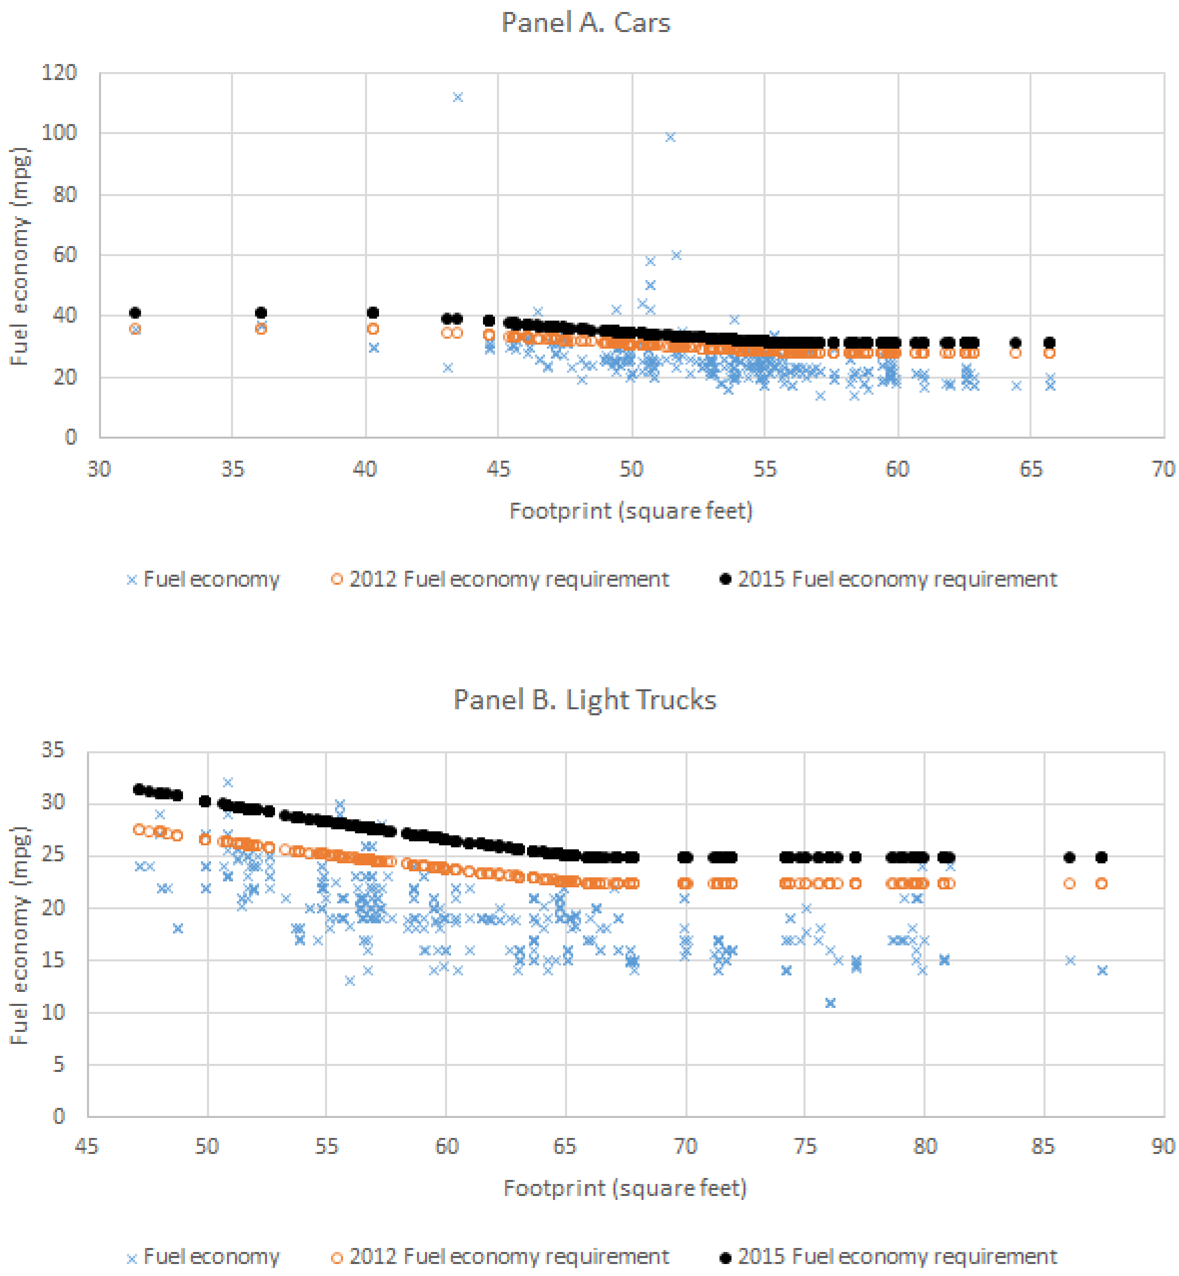

Notes: Each x and circle in the figure represents a unique vehicle sold in 2010. The x plots the vehicle's fuel economy, in miles per gallon, against its footprint, in square feet, where the footprint is computed as the product of the wheelbase and width. The open circles show the vehicle's fuel economy requirement in 2012 and the filled circles show the fuel economy requirement in 2015.

Figure 7: Own-Price Elasticity of Demand by Demographic Group

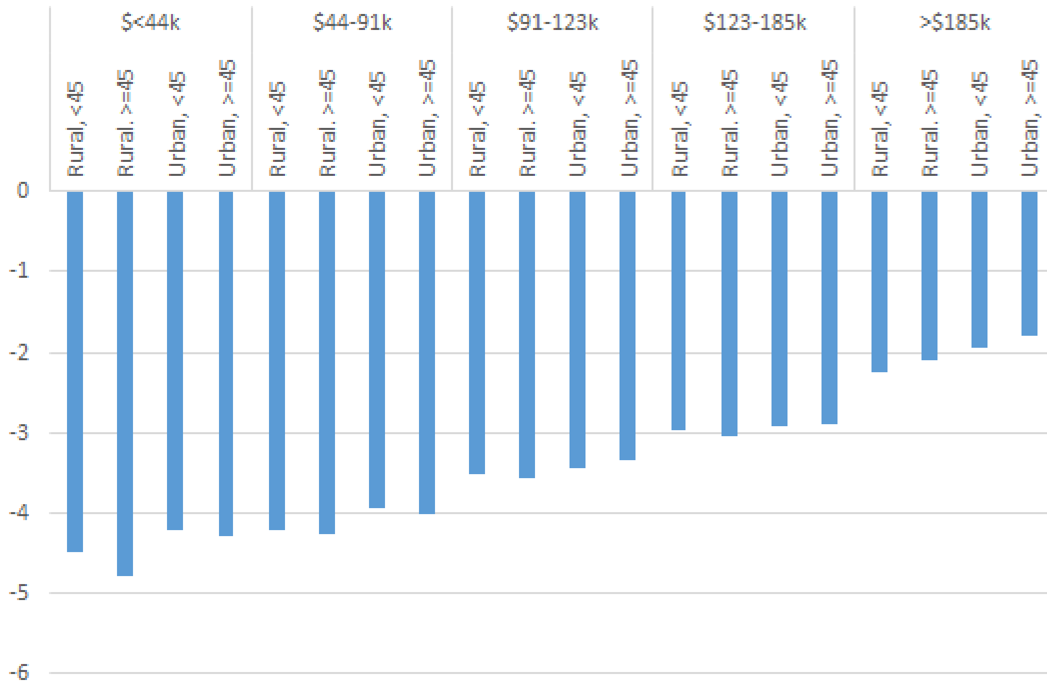

Notes: Each bar shows the own-price elasticity of demand for the indicated demographic group. The estimates are computed from the estimated demand model coefficients, and all estimates are weighted across vehicles and markets using predicted market shares as weights.

Figure 8: Willingness to Pay for Fuel Economy and Performance by Demographic Group

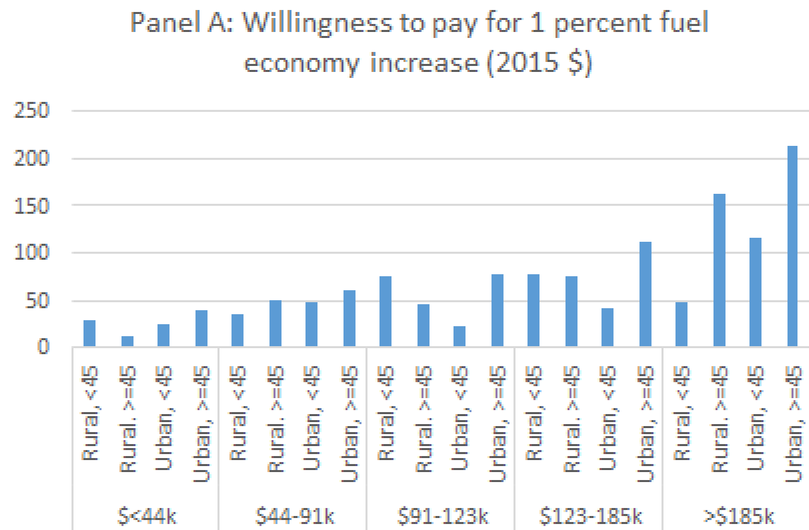

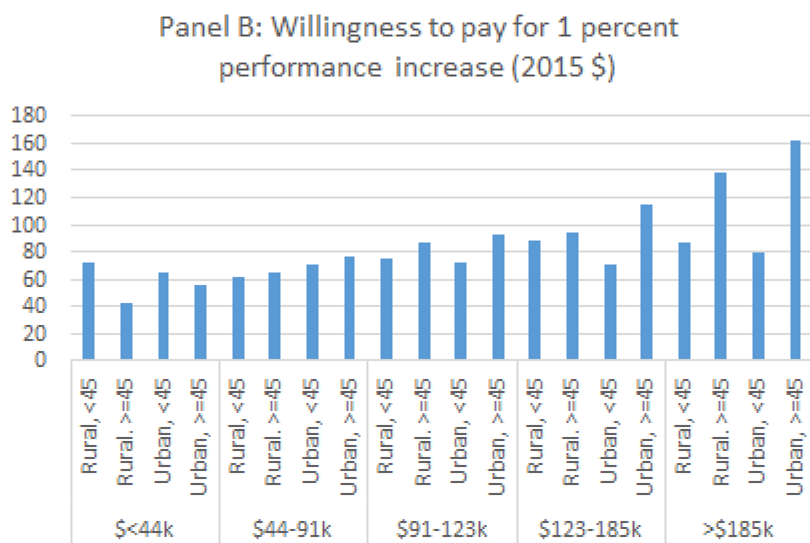

Notes: Panel A reports the WTP for a 1 percent fuel economy increase, and Panel B reports the WTP for a 1 percent performance increase. Each bar shows the estimate for the indicated demographic group. The estimates are computed from the estimated demand model coefficients, and all estimates are weighted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

Figure 9: Comparing Predicted vs. Observed Attributes by Demographic Group and Year

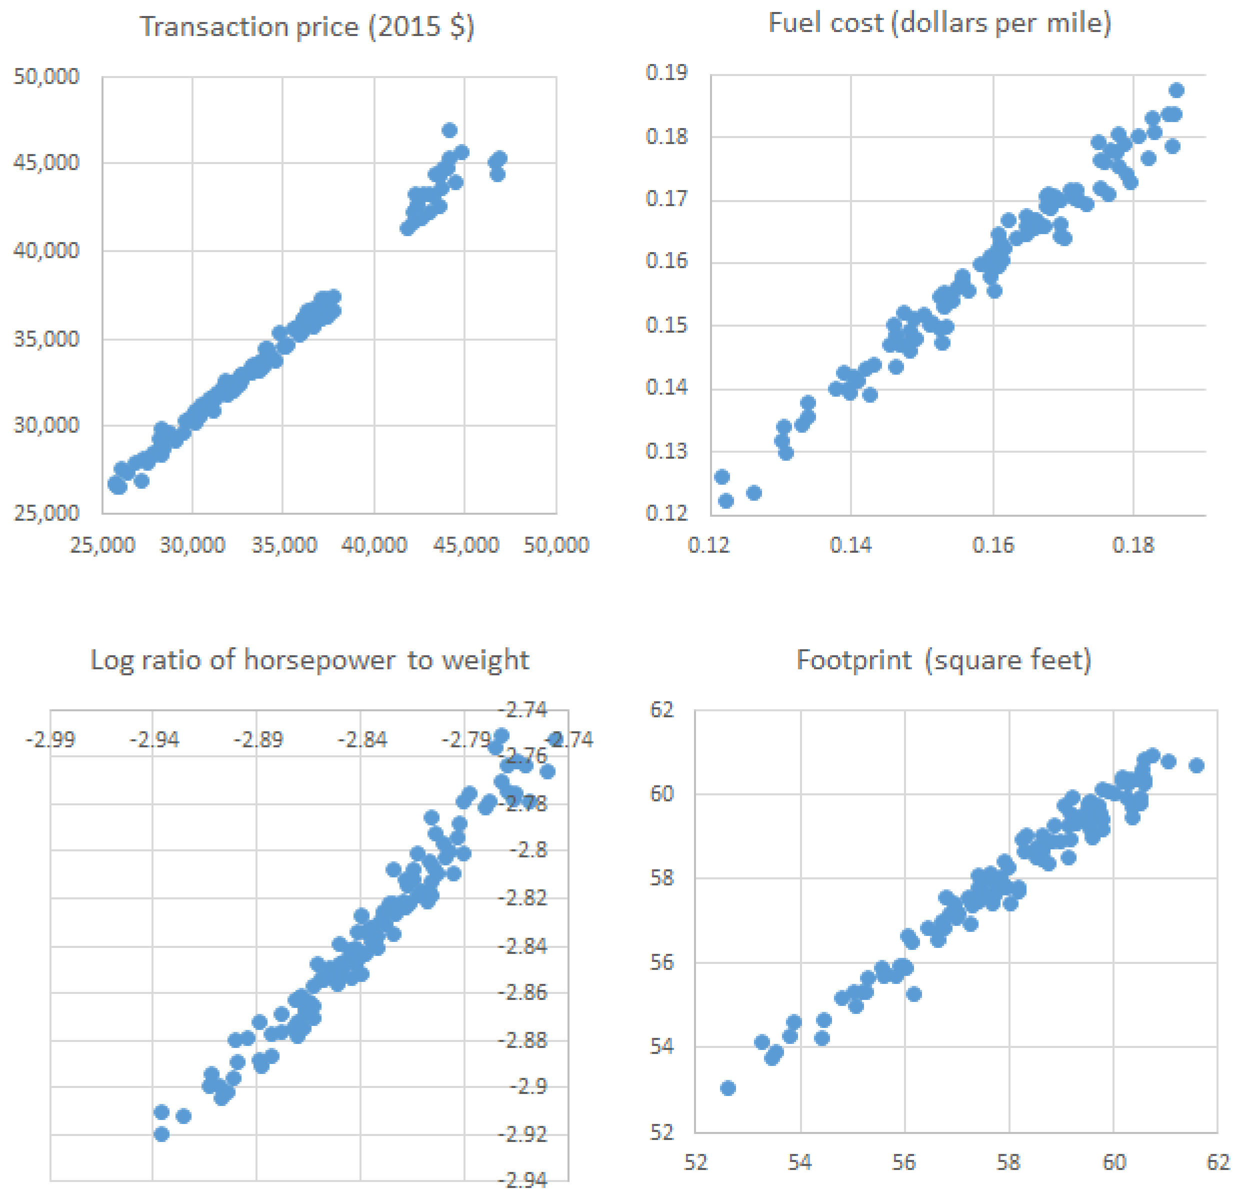

Notes: For each demographic group, we compute the predicted mean attribute indicated in the panel title using the vehicle market shares predicted by the model. The figure plots the predicted mean against the observed sales-weighted mean.

Figure 10: Comparison of Predicted and Observed 2015 Market Shares by Demographic Group, Brand, and Class: No-Change vs. Demand Model

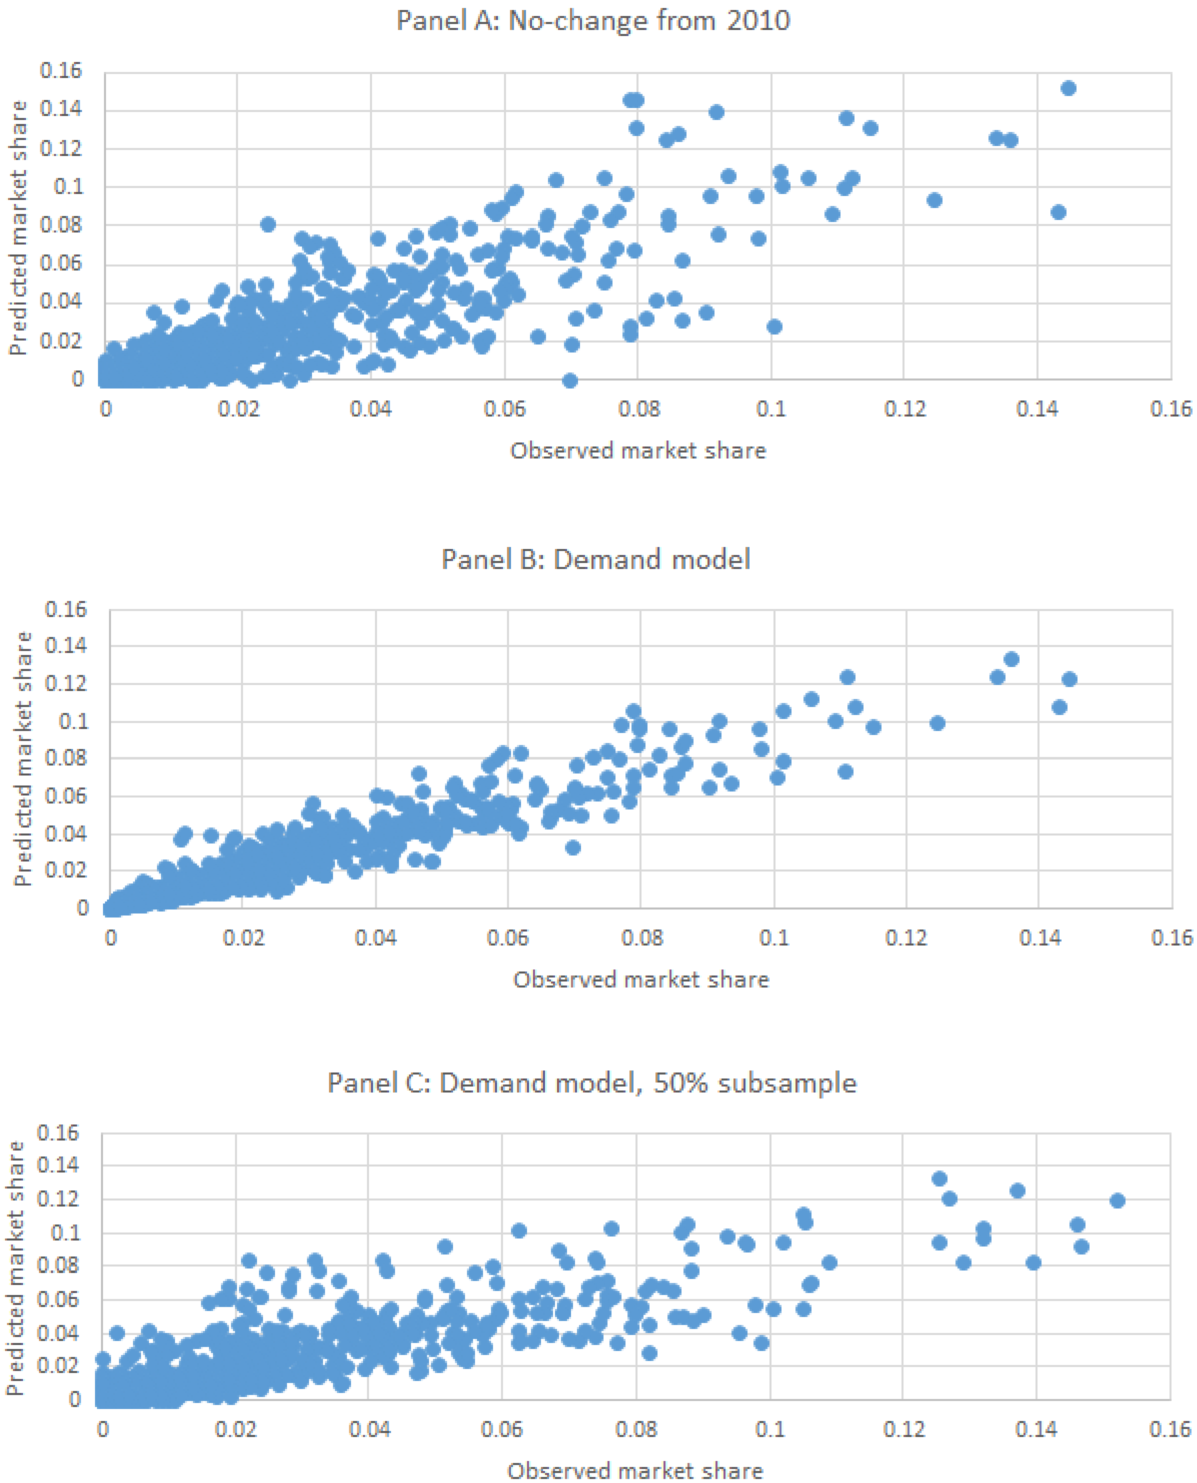

Notes: Vehicles are aggregated by brand, class, and year. The figure plots the predicted against observed market shares by brand, class, and year. In Panel A, the prediction is equal to the observed market share in 2010. In Panel B, the prediction is made using the demand model. In Panel C, the prediction is made using the demand model estimated on a random 50 percent subsample of vehicle by market observations.

Figure 11: Welfare Effects of 1 Percent Fuel Economy Increase

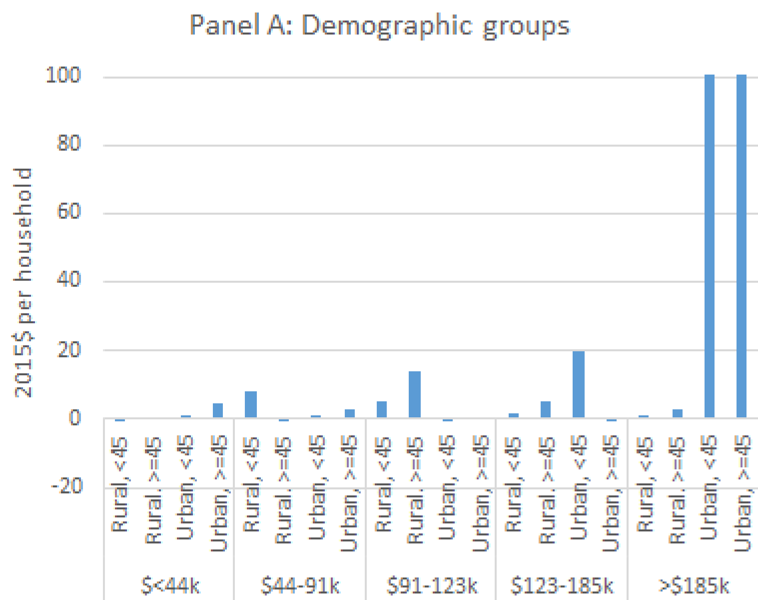

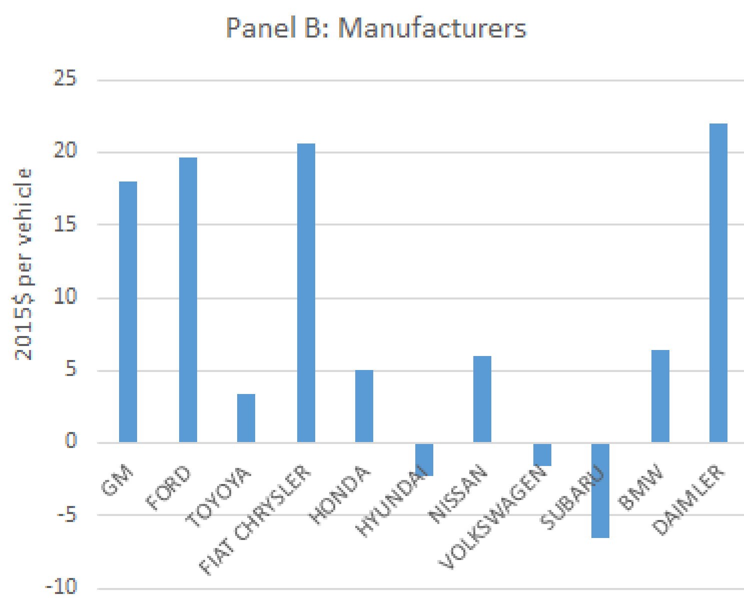

Notes: The figure shows the welfare effects of simulating a 1 percent fuel economy increase for all vehicles sold in 2012. Panel A reports average consumer welfare change per household in 2015$. Panel B reports average change in profits per vehicle for each manufacturer. Panel A cuts off the vertical axis at $100 per household so that the welfare changes for lower income groups are legible. Panel B omits the ”other” category, which accounts for less than 1 percent of total sales.

Figure 12: Fuel Economy Willingness to Pay and Own-Price Elasticity of Demand Versus Pass-Through Rate

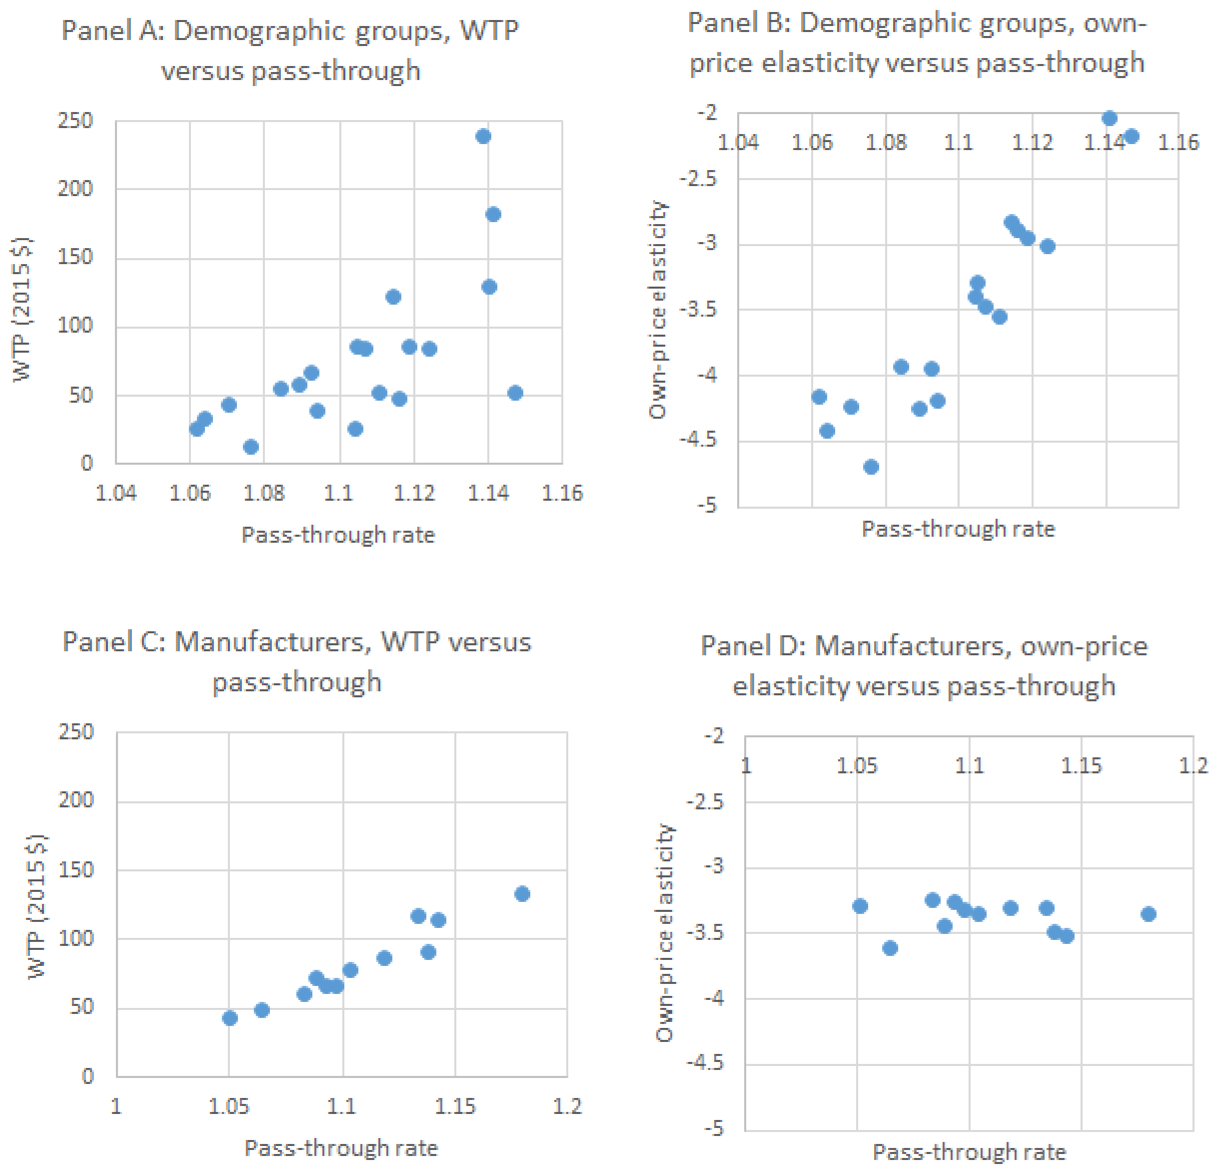

Notes: The figure plots the results of simulating a 1 percent fuel economy increase for all vehicles sold in 2012. In Panels A and B each point is a demographic group, and in Panels C and D each point is a manufacturer. Panels A and C plot WTP against pass-through rate, and Panels B and D plot own-price elasticity against pass-through rate. Own-price elasticity of demand and willingness to pay are the same as reported in Figure 7 for demographic groups. For manufacturers, own-price elasticity and willingness to pay are the means across consumers purchasing the vehicles, weighted by predicted sales. The pass-through rate is the ratio of the vehicle price change to the marginal cost change.

Figure 13: Own-Price Elasticities, Willingness to Pay, and Welfare

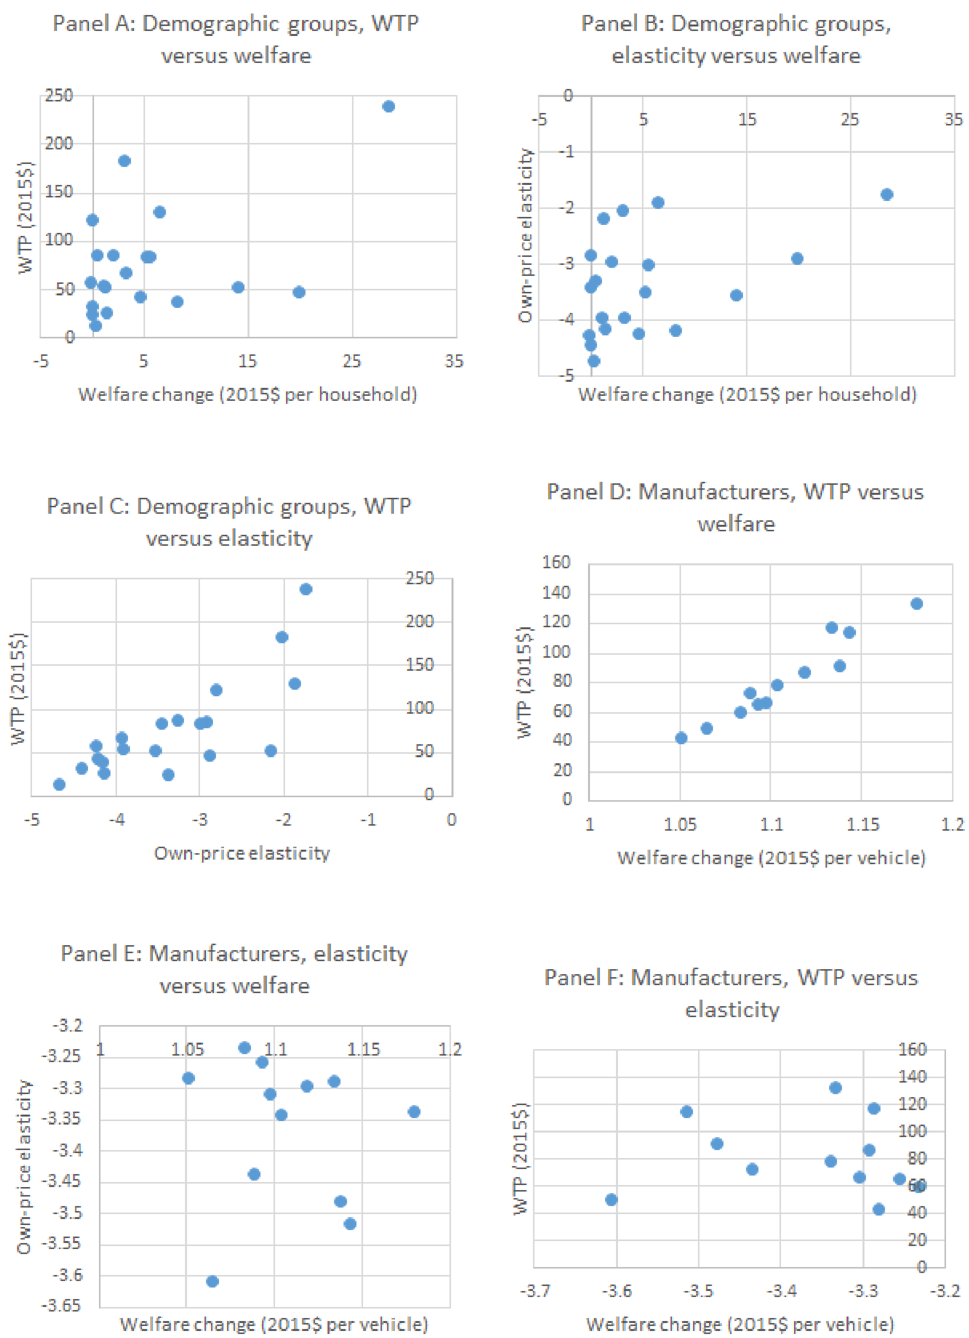

Notes: The figure plots the results of simulating a 1 percent fuel economy increase for all vehicles sold in 2012. Each panel plots the outcomes and preference estimates indicated in the title. All numbers are computed as in Figure 12.

Figure 14: Own-Price Elasticity and Consumer Welfare Changes for Model with Zero WTP for Fuel Economy

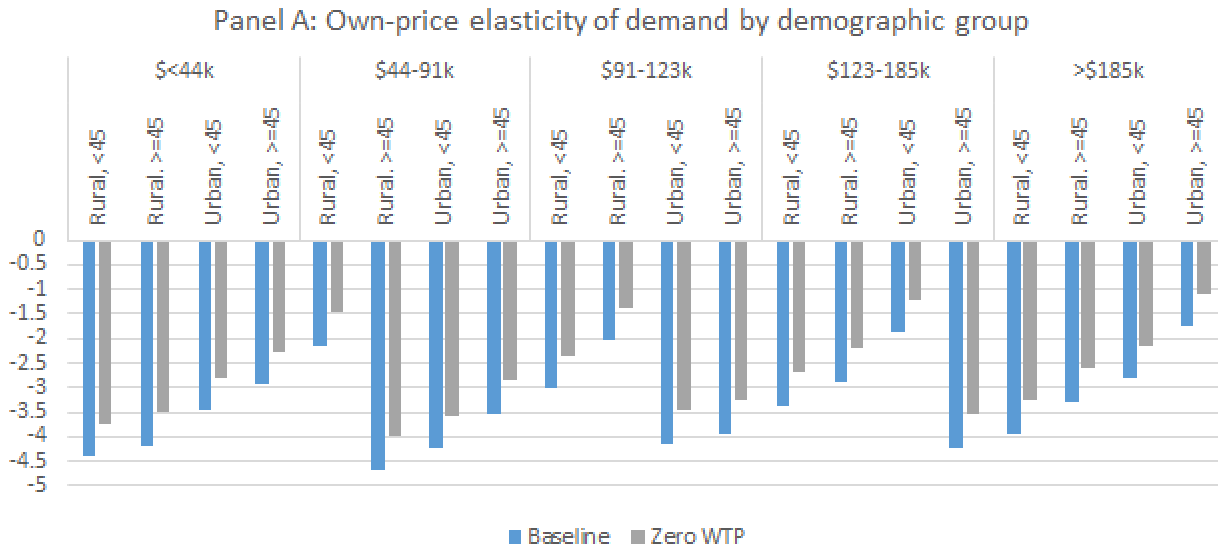

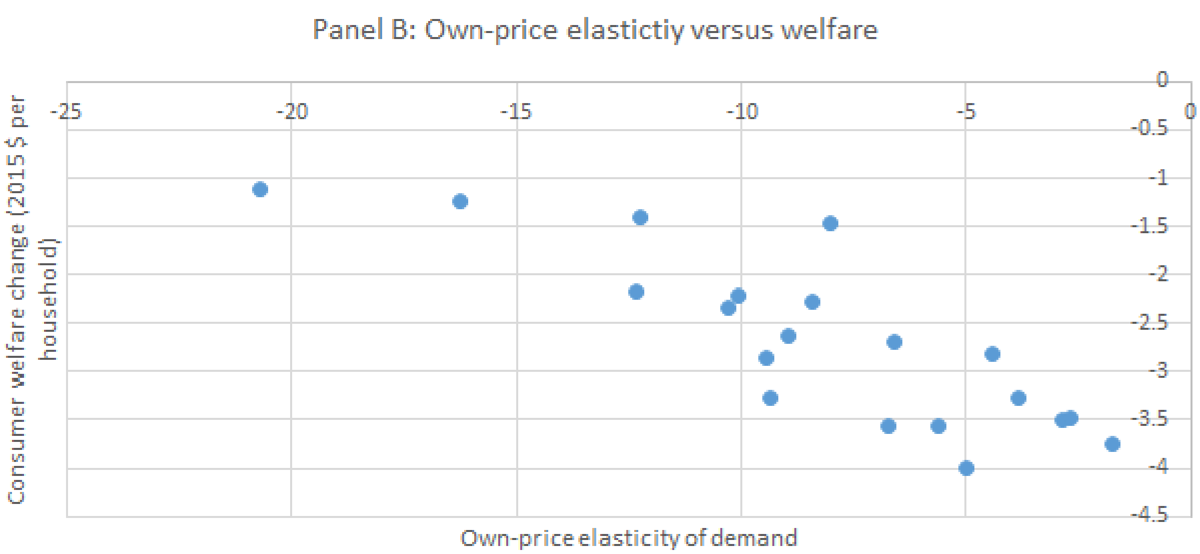

Notes: The figure reports results from a version of the model that imposes the restriction that consumer WTP for fuel economy equals zero. We re-estimate the demand model imposing this restriction, and plot the estimated own-price elasticity of demand by demographic group in Panel A. Using the new demand estimates we re-compute marginal costs and simulate the welfare effects of a 1 percent fuel economy increase. Panel B plots the consumer welfare change against the own-price elasticity of demand for each demographic group.

Figure 15: Welfare Implications of Fuel Economy/Performance Trade-Off

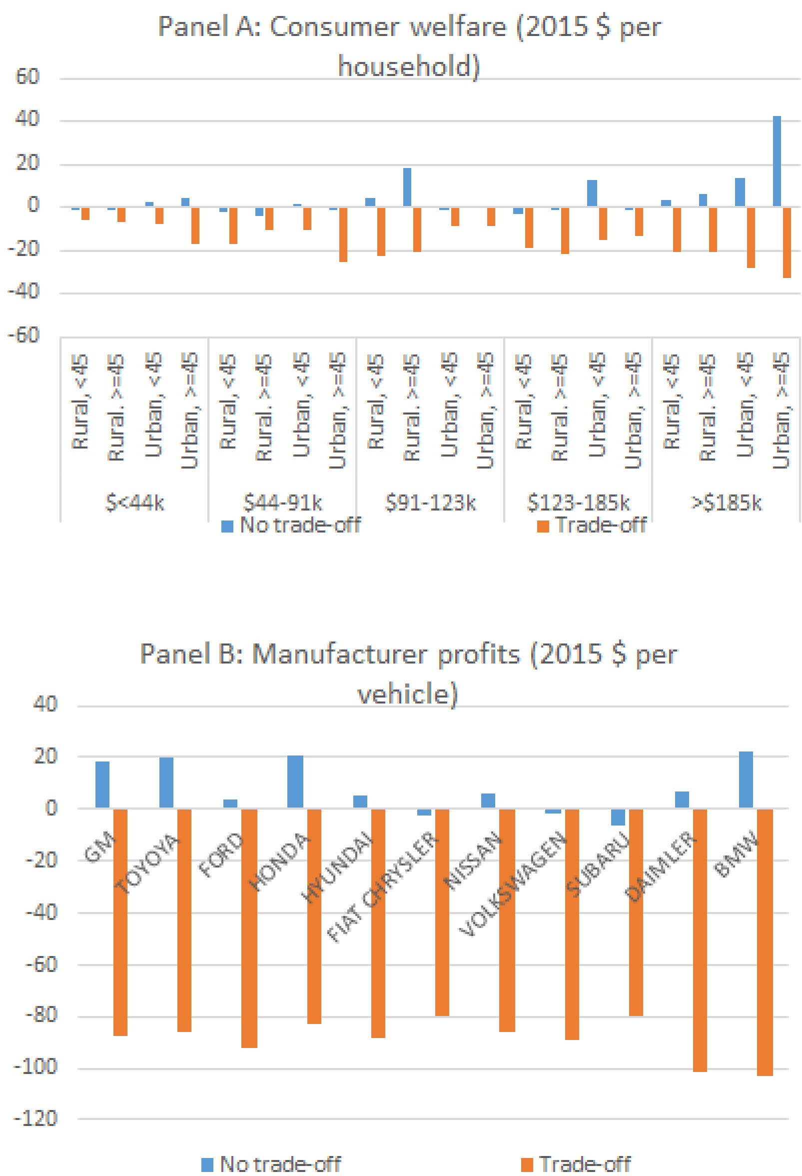

Notes: The figure compares the welfare effects by demographic group and manufacturer for the policy scenario in Figure 12 (no trade-off) with a policy scenario that includes a trade-off between fuel economy and performance.

Tables

Table 1: Attribute Comparison for Highest Selling 2015 Engine Twin Vehicles

Name	Engine displ. (liters)	Price (2015 $)	Miles per gallon	Horsepower / weight (pounds)	2015 Sales	Sales share (%)
Toyota Camry FWD	2.5	26,171	28	0.045	293,249	2.18
Toyota Camry FWD	3.5	31,648	25	0.077	14,919	0.11
Nissan Altima FWD	2.5	26,041	28	0.057	240,552	1.79
Nissan Altima FWD	3.5	31,015	26	0.079	8,064	0.06
Ford F150 XL Gasoline	2.7	36,486	22	0.079	46,401	0.34
Ford F150 XL Gasoline	3.5	36,227	19.5	0.085	197,240	1.47
Subaru Forester Touring	2.5	28,232	27	0.050	154,107	1.15
Subaru Forester Touring	2	32,974	25	0.069	9,564	0.07
Chevrolet Silverado 1500	4.3	34,280	18.3	0.059	58,277	0.43
Chevrolet Silverado 1500	5.3	39,877	18.6	0.073	84,493	0.63

Notes: The table reports vehicle characteristics and sales of the top five highest-selling vehicle twins during the 2015 market year. Twins are vehicle pairs that share the same model year, make, model, trim name, fuel type, body style, and drive type but have different engine configurations. The sales share is the vehicle's sales multiplied by 100 and divided by total new vehicle sales in 2015.

Table 2: Estimated Valuation Ratios by Demographic Group

Panel A: Rural

Income	Age < 45	Age >= 45
< 44k	0.10	0.06
44k - 91k	0.10	0.17
91k - 123k	0.18	0.14
123k - 185k	0.19	0.21
> 185k	0.10	0.44

Panel B: Urban

Income	Age < 45	Age >= 45
< 44k	0.09	0.23
44k - 91k	0.17	0.27
91k - 123k	0.08	0.31
123k - 185k	0.13	0.41
> 185k	0.35	0.76

Notes: The table reports the valuation ratio for each demographic group using the WTP estimates from Figure 8. See appendix for details on the calculations.

Table 3: Estimated Effects of Fuel Economy on Marginal Costs

Dependent variable is log of marginal costs						
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Cars						
Log of Fuel Economy	0.247	0.249	0.255	0.260	0.239	0.220
	(0.097)	(0.095)	(0.099)	(0.112)	(0.100)	(0.098)
R-squared	0.985	0.985	0.986	0.988	0.986	0.986
Observations	2,508	2,508	2,508	2,508	2,508	2,508
Panel B: Light trucks						
Log of Fuel Economy	0.178	0.169	0.161	0.111	0.184	0.181
	(0.075)	(0.077)	(0.076)	(0.093)	(0.077)	(0.076)
R-squared	0.973	0.973	0.974	0.977	0.973	0.974
Observations	2,901	2,901	2,901	2,901	2,901	2,901
Year fixed effects interacted with:	Body type fixed effects	Body type and drive type fixed effects	Body type and fuel type fixed effects	Body type and make fixed effects	Body type and engine size fixed effects	Body type and cylinders fixed effects

Notes: Each column and panel shows the results of a separate regression. Observations are by vehicle and market. The dependent variable is the log marginal costs estimated from Equation (26). The table shows the coefficient on log fuel economy with the standard error in parentheses, robust to heteroskedasticity. All regressions include the log of weight, the log of horsepower, vehicle fixed effects, and the independent variables described in the bottom row of the table.

Table 4: Changes in Aggregate Consumer Welfare and Manufacturer Profits by Counterfactual Scenario

Welfare (million 2015$)	Primary counterfactual, 1% fuel economy increase	Zero cost-per-mile coef., 1% fuel economy increase	Proportional fuel economy increase	Horsepower / fuel economy trade-off, 1% fuel economy increase
Consumer surplus	172	-614	206	-1,336
Manufacturer profits	133	-475	163	-1,072
Total	305	-1,089	369	-2,408

Notes: The table reports the changes in aggregate consumer welfare and manufacturer profits for each scenario discussed in Section 6. All scenarios include a uniform one percent fuel economy increase, except the scenario with a proportional fuel economy increase. In that scenario, each vehicle's fuel economy increase is proportional to its fuel economy requirement, and the fuel economy increases are calibrated to achieve the same fuel consumption reduction as the primary counterfactual. The total welfare change is the sum of consumer welfare and profits. See text for further explanations of the counterfactuals.

Appendix

Extension of the analytical model: Internalities

In Section 2, we assume that consumers' expected utility from purchasing the product equals the realized utility. If there is an internality, consumers' expectations include systematic errors, causing the realized and expected utilities to differ from one another. For example, consumers may underestimate the utility they will receive from a vehicle that offers higher fuel economy than another vehicle.

Here, we briefly consider the implications of this possibility for the conclusions in Section 2. Because the equilibrium price and quantity depend on consumers' expected utility rather than realized utility, an internality does not affect the results for pass-through and profits. However, the change in consumer surplus would have an additional term under the integral that equals the internality; that is, the difference between the realized and expected utility change caused by the attribute change. Consequently, the conclusion that WTP positively affects the consumer surplus change would only fail to hold if WTP and the internality are sufficiently strongly negatively correlated with one another. That is, as long as the correlation is positive or weak, the main conclusion about consumer surplus is unaffected by an internality.

Data

This section provides further details on the data construction. For each unique vehicle we compute the average transaction price in MaritzCX and assign this as the vehicle's purchase price. For the remaining 10 percent of vehicles that do not have reported transaction prices, we estimate the vehicle's purchase price based on the observed transaction price for the most closely related vehicle in the data with an observed transaction price. For example, for a vehicle with a missing transaction price, suppose that some other observations of vehicles with the same make, model, trim, model year, and purchase year have reported transaction prices. For the matching vehicles with non-missing data, we compute the average difference between MSRP and transaction price (in dollars). For the vehicle with a missing transaction price, the imputed price is the sum of its MSRP and the calculated difference between transaction price and MSRP. We repeat this procedure by sequentially aggregating the vehicle definition until we obtain imputed prices for all vehicles in the MaritzCX data.

As noted in the main text, we use the CEX data to weight household observations in the MaritzCX data. The CEX samples about 7,000 households in the United States each quarter, and importantly for our purposes the survey data include information about household demographics as well as whether the household purchased a new or used vehicle in the current quarter or the quarter prior to being surveyed. We use the CEX survey weights to compute the total numbers of new and used vehicles obtained by demographic group and year.

We construct weights for the MaritzCX household observations in three steps. First, we construct a weight variable so that the total new purchases by year and demographic group matches total new purchases by year and demographic group in the CEX. Second, we adjust the household weights so that the vehicle's share of sales in total sales by year is equal to the corresponding share according to the IHS data. Third, we adjust the household weights so that total new vehicles obtained by year in the MaritzCX data match total vehicles obtained by year in the IHS data. After constructing these weights, we compute the total new vehicles obtained by year, vehicle, and demographic group.

Note that by taking this approach, we assume implicitly that variation in survey response rates across demographic groups is orthogonal to variation in response rates across vehicles. Reversing the order has little effect on the estimated parameters of the consumer demand model.

Model and estimation

Derivation of market share equation

We derive the market share Equation (16). We begin by observing that the household's utility from purchasing a vehicle is the same as the utility another household in the same demographic group g would experience from choosing the same vehicle, so that $v_{ijt} \equiv v_{gjt}$. Consequently, the probability the household chooses a particular vehicle is the same for all households belonging to the same demographic group g, and we can aggregate the choice probabilities to market shares by demographic group, vehicle, and market:

$$s_{gjt} = \frac{e^{v_{gjt}}}{\sum\limits_{k} e^{v_{gkt}}}. \tag{A.1}$$

Taking the natural log of both sides of Equation (A.1) yields

$$\ln(s_{gjt}) = v_{gjt} - \ln(\sum_{k} v_{gkt}). \tag{A.2}$$

Normalizing the utility of the outside good to zero for each demographic group implies that $\ln s_{g0t} = -\ln(\sum_k v_{gkt})$. Substituting this log share into Equation (A.2), re-arranging, and substituting the definition of utility from Equation (14) yields the market share equation in the main text, Equation (16).

Definition of the outside option

In this section, we review how the previous literature handles the outside option and argue that in our setting conditioning on buying a vehicle (new or used) is appropriate. Berry et al. (1995; 2004) and Petrin (2002) define the outside good as the decision to not buy a new vehicle. Consequently, their choice models apply to all households in the United States during their sample periods. The models can be used to simulate the effects of a policy or vehicle entry and exit on total new vehicle sales and consumer welfare.

In contrast to that approach, many recent vehicle demand models exclude an outside good because utility from this option does not represent a structural parameter (Train and Winston 2007). Utility from the outside good is often normalized to zero (Berry et al. 1995; 2004), and therefore does not change in response to policy changes. As a result, including the outside good in the demand model estimation and post-estimation simulation exercises may provide inconsistent inferences about welfare. Furthermore, the definition of utility from the outside good is broad and not well-defined. Households deciding not to buy a new vehicle in a given time period do so for many different reasons. For example, some may not drive at all, some may be financially constrained, some may be satisfied with their current vehicle portfolio, and some may decide to buy a used vehicle.

We reduce these concerns by narrowing the definition of the outside good to purchasing a used vehicle. Conditioning on buying a (new or used) vehicle sharpens the interpretation of the outside good utility: Households that choose the outside good receive the utility from the new vehicle that is common to their demographic group. While this definition does not provide a structural interpretation for the outside good utility, it does make our inferences about welfare more consistent than the standard method of defining the outside good. Because our demand model is estimated using several years of data, our estimated substitution patterns between new and used vehicles are identified by changing attributes of new and used vehicles. And because the average attributes of used vehicles do not change much relative to the market for new vehicles, changes in new vehicle attributes identify the substitution patterns between new and used vehicles. For example, during our sample period, tightening new vehicle fuel economy standards caused new vehicle fuel economy to increase rapidly. These fuel economy changes help identify substitution patterns between new and used vehicles for each demographic group.

Calculation of valuation ratios

The valuation ratios reported in Table 2 are the ratio of the WTP for a 1 percent fuel economy increase to the present discounted value of the resulting fuel savings. The fuel savings equal the difference in the present discounted value of fuel costs with and without the fuel economy increase.

We begin by calculating the present discounted value of the fuel costs without the fuel economy increase, which is given by $PDV_{gjt} = \sum_{\tau=t}^{t+T} \frac{\pi_{j\tau} V_{g\tau} f_{\tau}}{m_{jt}(1+r_g)^{\tau}}$. T is the maximum lifetime of the vehicle, $\pi_{j\tau}$ is the probability that the vehicle is not retired before year τ (which is sometimes referred to as the survival probability rate), $V_{g\tau}$ is the number of miles the vehicle is driven in year τ, f_{τ} is the real fuel price in year τ, m_{jt} is the vehicle's fuel economy, and r_g is the real discount rate.

The present discounted value is calculated using assumptions on fuel prices, vehicle miles traveled (VMT), scrappage rates, and discount rates. For fuel prices, we assume that the price in market t is equal to the real fuel price in all subsequent years τ. Using household-level data from the 2017 National Household Travel Survey (NHTS), we estimate unique VMT-by-age schedules separately for cars and light trucks, as well as for each of the 20 demographic groups. Appendix Figures A.1 and A.2 display estimated VMT schedules for a few of the demographic groups.

We estimate unique scrappage rates for cars and light trucks using R.L. Polk vehicle registration data from 2002 through 2014. The data provide vehicle counts by class (car or light truck), age, and year. From these data, we compute annual average scrappage rates as the difference in vehicle counts divided by prior year vehicle counts for each vehicle class. The scrappage rates are identical to those that appear in the appendix of Leard et al. (2017).

We use vehicle loan rates in the MaritzCX data to compute average discount rates by demographic group. The loan rates are presented in Appendix Tables A.1 and A.2. The tables show that the loan rates vary considerably across demographic groups; low income households typically face higher loan rates.

Having computed the present discounted value of fuel costs for each demographic group, vehicle, and market, we compute the change in fuel costs caused by a 1 percent fuel economy increase. We compute the average change in fuel costs for each demographic group using as weights the vehicle market shares predicted by the demand model for all markets in the estimation sample. The valuation ratios are then computed as the ratio of the estimated WTP for fuel cost savings and the average change in fuel costs.

Comparison of our cost estimates with NHTSA estimates

An alternative to using Equation (26) to estimate the relationship between marginal costs and fuel economy is to use NHTSA estimates of the costs of fuel-saving technologies. As part of the 2016 Technical Assessment Report, NHTSA uses its technology cost model to estimate the costs of meeting the standards (EPA and NHTSA 2016). The estimation algorithm begins with a set of vehicles in year 2016, recording the fuel economy, retail price, and set of technologies for each vehicle. The agency uses the model to simulate compliance with future fuel economy standards. In the simulation, over time technologies are added to the vehicles so that manufactures achieve the specified standards. The model keeps track of changes over time in fuel economy and technology costs. The agency assumes that the vehicle's price increases in proportion to the cost increase. The result of the simulation is a panel data set of vehicles over time, including fuel economy and retail price.

We obtained these data and for each vehicle and year $t > 2016$ we compute the log ratio of the vehicle's price in that year to its price in year 2016, as well as the log ratio of fuel economy in that year to fuel economy in 2016. Because the agency assumes that retail prices are a fixed markup over production costs, regressing the log price ratio on the log fuel economy ratio is equivalent to regressing the log production cost ratio on the log fuel economy ratio. The fuel economy coefficient in this regression is typically about 0.15, depending on the additional controls we include and the sample. This estimate is fairly similar to the estimates reported in Table 3 using Equation (26).

There are two important differences between the estimates in Table 3 and the estimates using NHTSA data. First, the NHTSA cost estimates pertain to technologies the agency projected would be adopted after 2016, whereas the estimates in Table 3 reflect the cost of adding technologies during the sample period of 2012 through 2015. In principle, the NHTSA cost estimates could be higher or lower than those estimated using the sample period of 2012 through 2015. On the one hand, in the NHTSA model the manufacturers adopt technologies roughly in order of decreasing cost effectiveness. Consequently, the cost of a given fuel consumption improvement increases over time, which causes the NHTSA data to yield higher cost estimates than our data. On the other hand, the NHTSA model includes technological change over time, which reduces the cost of adopting a particular technology after 2016 compared to the cost prior to 2016. These two effects oppose one another.

The second difference is that in Table 3, the dependent variable is the marginal cost of producing the vehicle, whereas in the NHTSA data the dependent variable is the average production cost. For the reasons provided in the main text, we prefer to use the estimates from Equation (26) in the policy simulations.

Computing vehicle prices in policy counterfactuals

We explain the algorithm for computing the profit-maximizing equilibrium vehicle prices in the policy counterfactuals. The equilibrium prices solve each firm's first order condition for vehicle price (that is, Equation (19)), so that the prices represent the best responses of each manufacturer given prices of all other manufacturers.

The policy counterfactual raises each vehicle's fuel economy as well as its marginal costs according to Equation (26). The derivatives in Equation (19) depend on predicted vehicle market shares, which in turn depend on vehicle prices. Consequently, the first-order condition is an implicit nonlinear function of the vehicle prices, which creates challenges for finding the equilibrium prices. We circumvent this nonlinearity by constructing an initial guess of the set of equilibrium prices that uses market shares computed from the baseline equilibrium prices, as well as the new marginal costs and fuel economy in the policy counterfactual. Given these derivatives, marginal costs, and other parameters, Equation (19) is linear in the vehicle price. We solve this equation for each vehicle's price, which constitutes our initial guess of the equilibrium prices in the policy counterfactual.

Next, we use the new prices to recompute market shares and derivatives, and solve the first order condition for a new set of prices. We iterate the procedure until the change in equilibrium prices between one iteration and the next is less than a specified tolerance. Finally, we check that the first-order conditions are satisfied for all vehicles and that the second order conditions indicate that the equilibrium represents a maximum.

Appendix Figures

Figure A.1: Estimated Car VMT Schedules for Selected Demographic Groups

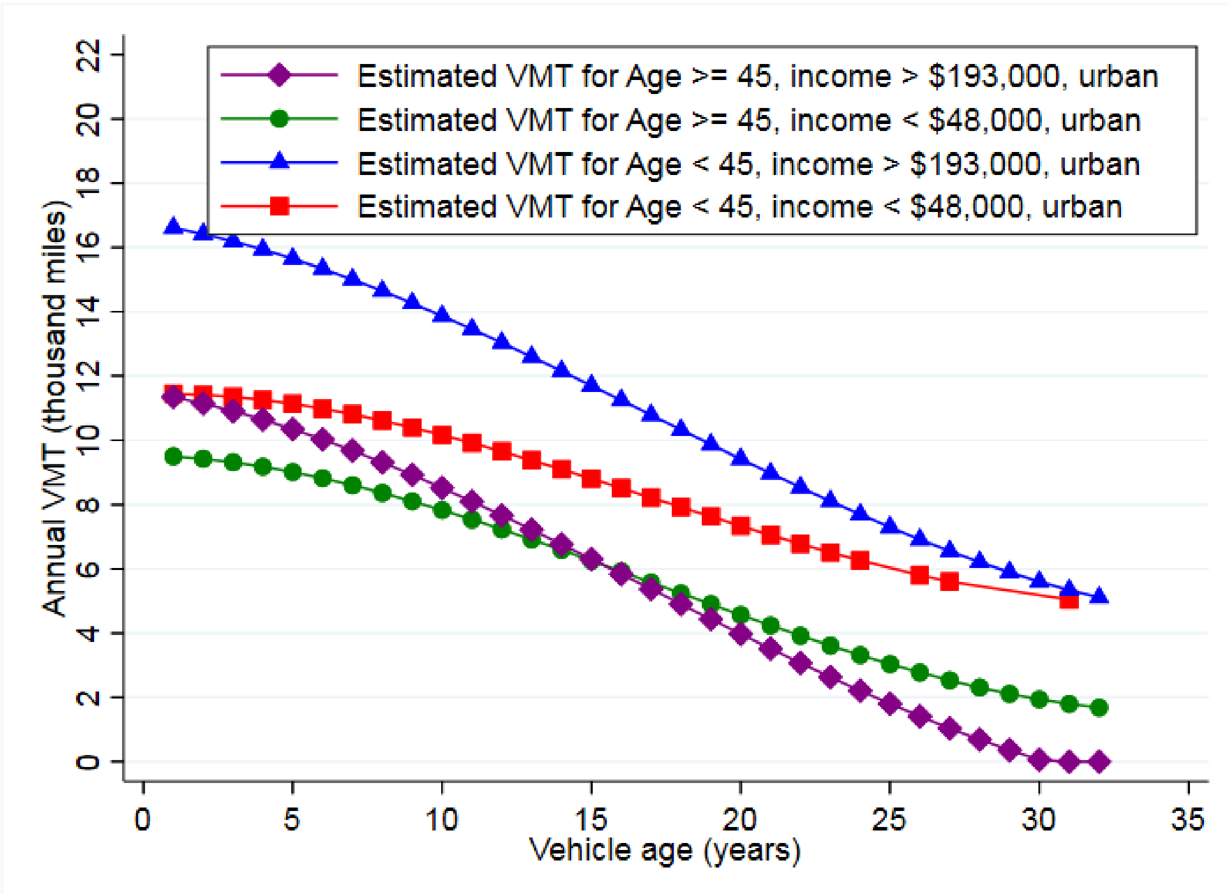

Figure A.2: Estimated Light Truck VMT Schedules for Selected Demographic Groups

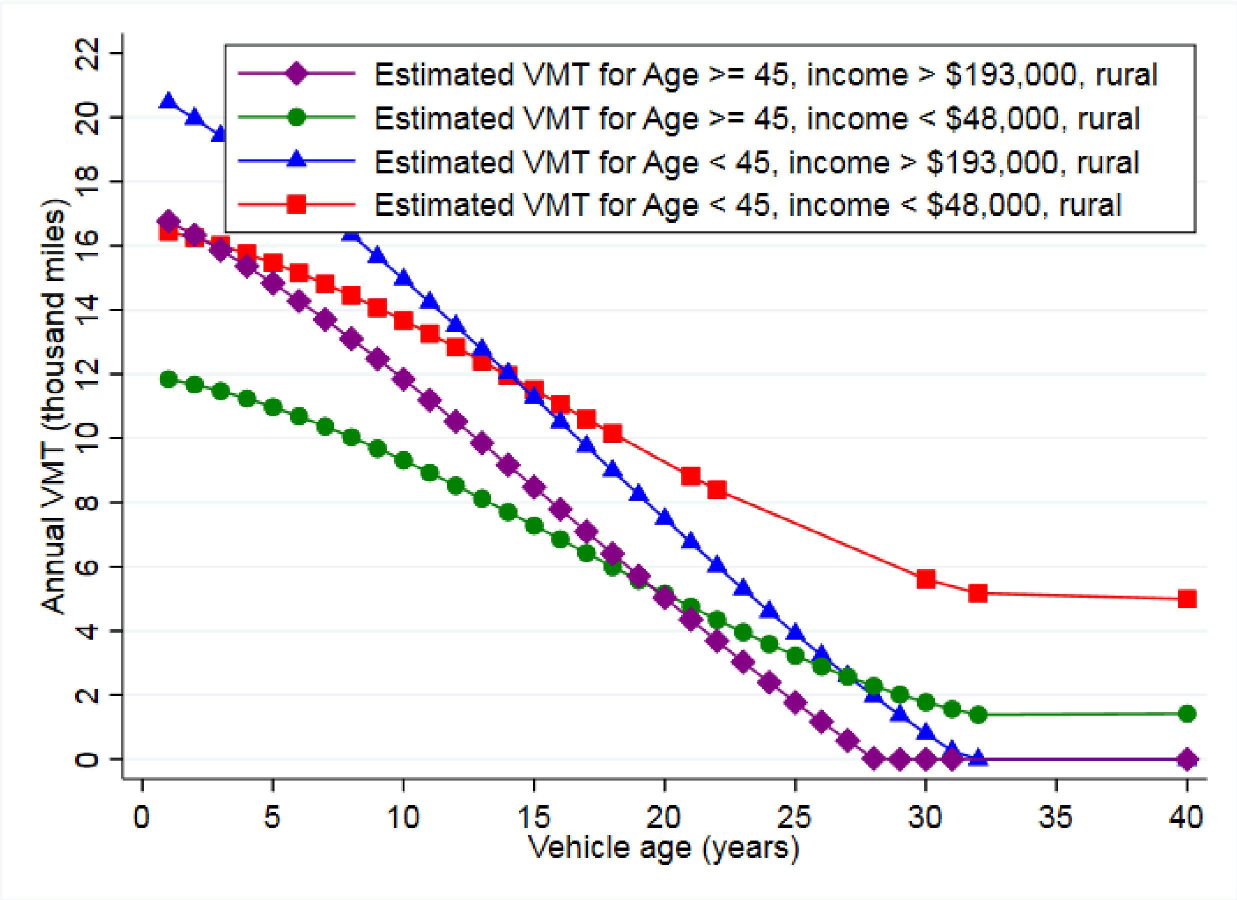

Appendix Tables

Table A.1: Vehicle Attributes by Demographic Group - Rural

Income	Age group	Price (2015 $)	APR	Share light truck	Share hybrid or plug-in	Fuel economy	Horsepower / weight	Foot-print	2016 fuel economy requirement	Share used
< 44k	< 45	26,950	5.33	0.36	0.015	25.9	0.056	54.9	31.2	0.95
		(7,881)				(6.0)	(0.011)	(7.4)	(3.5)	
44k - 91k	< 45	30,677	4.81	0.52	0.018	23.9	0.059	58.0	29.7	0.92
		(9,182)				(6.4)	(0.013)	(9.1)	(3.6)	
91k - 123k	< 45	33,270	4.03	0.63	0.018	22.6	0.059	59.6	28.8	0.90
		(9,544)				(6.5)	(0.012)	(9.4)	(3.6)	
123k - 185k	< 45	36,708	3.54	0.67	0.023	22.1	0.061	60.4	28.4	0.81
		(11,132)				(6.6)	(0.012)	(9.2)	(3.5)	
> 185k	< 45	43,614	3.16	0.72	0.038	21.3	0.063	61.5	27.9	0.77
		(14,693)				(7.4)	(0.014)	(8.7)	(3.3)	
< 44k	>= 45	29,465	4.60	0.52	0.018	24.1	0.058	57.9	29.7	0.85
		(8,133)				(5.9)	(0.012)	(9.0)	(3.6)	
44k - 91k	>= 45	32,091	4.01	0.59	0.028	23.1	0.059	59.3	29.1	0.85
		(9,042)				(6.3)	(0.013)	(9.2)	(3.5)	
91k - 123k	>= 45	34,637	3.58	0.61	0.038	22.9	0.061	59.7	28.8	0.81
		(10,291)				(6.6)	(0.014)	(9.2)	(3.4)	
123k - 185k	>= 45	37,039	3.32	0.62	0.040	22.4	0.062	60.1	28.7	0.80
		(11,684)				(6.8)	(0.014)	(9.1)	(3.4)	
> 185k	>= 45	43,716	2.96	0.64	0.049	21.8	0.065	61.1	28.3	0.73
		(15,104)				(7.4)	(0.015)	(8.9)	(3.3)	

Notes: The table shows the purchases-weighted mean attribute or market share for each demographic group, with standard deviations in parentheses. The sample includes all vehicles purchased between 2010 and 2015.

Table A.2: Vehicle Attributes by Demographic Group - Urban

Income	Age group	Price (2015 $)	APR	Share light truck	Share hybrid or plug-in	Fuel economy	Horsepower / weight	Foot-print	2016 fuel economy requirement	Share used
< 44k	< 45	26,529	5.07	0.28	0.023	26.9	0.055	53.9	31.8	0.92
		(7,567)				(6.0)	(0.011)	(6.7)	(3.3)	
44k - 91k	< 45	30,253	4.45	0.42	0.030	25.3	0.058	55.7	30.6	0.90
		(9,199)				(6.8)	(0.012)	(7.3)	(3.4)	
91k - 123k	< 45	33,071	3.92	0.51	0.040	24.5	0.059	57.0	29.8	0.85
		(10,286)				(8.0)	(0.013)	(7.6)	(3.4)	
123k - 185k	< 45	36,421	3.39	0.58	0.049	23.8	0.060	58.2	29.2	0.77
		(11,626)				(8.8)	(0.013)	(7.9)	(3.4)	
> 185k	< 45	43,263	2.97	0.63	0.057	23.1	0.063	58.8	28.7	0.63
		(15,183)				(9.7)	(0.014)	(7.1)	(3.3)	
< 44k	>= 45	27,965	4.65	0.39	0.028	25.7	0.057	55.4	30.8	0.81
		(7,921)				(6.0)	(0.012)	(7.3)	(3.4)	
44k - 91k	>= 45	31,371	4.13	0.48	0.040	24.8	0.059	56.7	30.0	0.78
		(9,442)				(6.6)	(0.013)	(7.6)	(3.3)	
91k - 123k	>= 45	33,739	3.62	0.50	0.050	24.3	0.060	57.1	29.8	0.79
		(10,749)				(7.1)	(0.013)	(7.4)	(3.2)	
123k - 185k	>= 45	36,539	3.29	0.52	0.061	24.1	0.062	57.4	29.6	0.74
		(12,285)				(7.9)	(0.014)	(7.4)	(3.2)	
> 185k	>= 45	43,699	2.89	0.55	0.079	23.4	0.064	58.5	29.2	0.62
		(15,627)				(9.1)	(0.015)	(7.2)	(3.2)	

Notes: The table shows the purchases-weighted mean attribute or market share for each demographic group, with standard deviations in parentheses. The sample includes all vehicles purchased between 2010 and 2015.

Table A.3: Vehicle Attributes by Manufacturer

Firm	Market share	Price (2015 $)	Share light truck	Share hybrid or plug-in	Fuel economy	Horsepower / weight	Footprint	2016 fuel economy requirement
GM	0.180	35,801	0.60	0.009	21.8	0.061	61.6	28.4
		(11,962)			(5.7)	(0.015)	(9.6)	(3.6)
Ford	0.170	32,381	0.59	0.022	22.8	0.061	60.0	29.0
		(8,710)			(5.9)	(0.013)	(9.7)	(3.6)
Toyota	0.141	31,523	0.40	0.182	27.5	0.054	55.6	30.7
		(10,485)			(8.6)	(0.012)	(7.4)	(3.3)
Fiat/Chrysler	0.112	31,007	0.71	0.002	20.6	0.064	59.3	28.5
		(6,709)			(5.0)	(0.013)	(7.9)	(3.3)
Honda	0.107	28,852	0.46	0.011	26.4	0.056	54.4	30.6
		(6,990)			(4.9)	(0.007)	(4.8)	(3.2)
Hyundai	0.085	25,331	0.24	0.018	26.7	0.058	53.1	31.9
		(5,441)			(3.5)	(0.009)	(3.3)	(2.6)
Nissan	0.084	29,248	0.38	0.016	27.3	0.057	54.4	31.0
		(9,671)			(10.7)	(0.013)	(5.3)	(3.0)
Volkswagen	0.036	36,797	0.20	0.004	26.2	0.056	53.5	32.0
		(15,671)			(4.2)	(0.012)	(4.1)	(2.7)
Subaru	0.029	28,729	0.56	0.006	26.0	0.052	52.5	30.9
		(2,915)			(3.1)	(0.006)	(2.5)	(2.3)
BMW	0.023	51,020	0.30	0.002	24.8	0.066	54.8	31.3
		(14,821)			(5.8)	(0.013)	(5.2)	(3.3)
Daimler	0.019	54,633	0.41	0.003	22.3	0.066	57.4	30.0
		(14,411)			(6.2)	(0.015)	(8.4)	(3.8)
Other	0.003	41,252	0.53	0.031	25.3	0.060	54.2	30.5
		(21,354)			(13.3)	(0.014)	(4.3)	(3.1)

Notes: The table shows the purchases-weighted mean attribute for each manufacturer, with standard deviations in parentheses. The sample includes all vehicles purchased between 2010 and 2015.

Table A.4: Demographics Shares by Manufacturer

Manufacturer	< 44k	44k - 91k	91k - 123k	123k - 185k	> 185k	Age >= 45	Urban
GM	0.35	0.29	0.13	0.14	0.09	0.74	0.56
Ford	0.34	0.29	0.13	0.15	0.09	0.69	0.62
Toyota	0.34	0.28	0.13	0.15	0.10	0.65	0.70
Fiat/Chrysler	0.37	0.30	0.13	0.13	0.07	0.62	0.58
Honda	0.35	0.27	0.13	0.15	0.09	0.60	0.71
Hyundai	0.45	0.29	0.12	0.10	0.04	0.65	0.67
Nissan	0.38	0.28	0.12	0.14	0.08	0.59	0.69
Volkswagen	0.21	0.24	0.14	0.19	0.22	0.52	0.77
Subaru	0.27	0.28	0.16	0.18	0.10	0.61	0.65
BMW	0.10	0.19	0.13	0.24	0.34	0.63	0.80
Daimler	0.10	0.17	0.13	0.23	0.36	0.71	0.81
Other	0.28	0.21	0.11	0.17	0.23	0.58	0.74

Notes: The table shows the purchases-weighted market share for each manufacturer, with standard deviations in parentheses. The sample includes all vehicles purchased between 2010 and 2015.

Table A.5: Estimated Price Elasticities and Willingness to Pay by Demographic Group

Panel A: Rural

	Own-price elasticity		WTP for 1 percent fuel economy increase		WTP for 1 percent horsepower increase		WTP for 1 percent footprint increase	
Income	Age < 45	Age >= 45	Age < 45	Age >= 45	Age < 45	Age >= 45	Age < 45	Age >= 45
< 44k	-4.49	-4.78	30.34	12.73	72.63	41.97	240.68	339.57
44k - 91k	-4.22	-4.27	35.20	50.97	61.89	64.58	380.65	462.06
91k - 123k	-3.52	-3.57	76.24	46.71	75.14	87.07	560.17	536.36
123k - 185k	-2.98	-3.04	78.36	75.61	87.65	94.52	673.55	673.01
> 185k	-2.24	-2.10	48.13	162.86	87.31	138.00	902.81	1074.56

Panel B: Urban

	Own-price elasticity		WTP for 1 percent fuel economy increase		WTP for 1 percent horsepower increase		WTP for 1 percent footprint increase	
Income	Age < 45	Age >= 45	Age < 45	Age >= 45	Age < 45	Age >= 45	Age < 45	Age >= 45
< 44k	-4.23	-4.29	24.10	39.69	64.00	55.42	165.60	253.28
44k - 91k	-3.95	-4.00	48.98	60.80	70.48	76.31	251.84	335.17
91k - 123k	-3.44	-3.34	23.17	77.62	71.37	92.60	320.33	409.91
123k - 185k	-2.92	-2.89	42.14	111.15	71.14	114.98	445.34	505.70
> 185k	-1.95	-1.81	116.71	213.17	79.39	161.26	770.68	836.77

Table A.6: Effects of One Percent Fuel Economy Increase on Consumer Welfare

Panel A: Rural

Age group	Income group	Own-price elasticity of demand	WTP for 1 percent fuel economy increase	Pass-through rate	Percentage change in used vehicle purchases	Welfare change (2015$ per person)
< 45	< 44k	-4.41	33.61	1.064	0.007	-0.48
< 45	44k - 91k	-4.18	39.35	1.094	0.008	-0.61
< 45	91k - 123k	-3.47	84.96	1.107	-0.029	2.78
< 45	123k - 185k	-2.94	86.87	1.118	-0.036	4.18
< 45	> 185k	-2.16	53.10	1.147	0.011	-1.81
>= 45	< 44k	-4.69	14.01	1.076	0.059	-3.83
>= 45	44k - 91k	-4.25	59.09	1.089	-0.014	1.10
>= 45	91k - 123k	-3.54	53.47	1.111	0.002	-0.19
>= 45	123k - 185k	-3.00	85.30	1.124	-0.037	4.29
>= 45	> 185k	-2.03	183.65	1.141	-0.103	17.94

Panel B: Urban

Age group	Income group	Own-price elasticity of demand	WTP for 1 percent fuel economy increase	Pass-through rate	Percentage change in used vehicle purchases	Welfare change (2015$ per person)
< 45	< 44k	-4.15	26.85	1.062	0.018	-1.28
< 45	44k - 91k	-3.93	55.37	1.084	-0.005	0.44
< 45	91k - 123k	-3.38	26.02	1.104	0.037	-3.56
< 45	123k - 185k	-2.89	48.23	1.116	0.014	-1.67
< 45	> 185k	-1.88	131.15	1.140	-0.066	12.38
>= 45	< 44k	-4.22	44.12	1.070	0.005	-0.38
>= 45	44k - 91k	-3.93	68.38	1.092	-0.043	3.49
>= 45	91k - 123k	-3.28	87.32	1.105	-0.056	5.70
>= 45	123k - 185k	-2.82	123.56	1.114	-0.114	13.86
>= 45	> 185k	-1.74	239.74	1.138	-0.209	42.73

Notes: The table reports the welfare results from the same simulation as in Figure 11. Own-price elasticity and WTP are the same as reported in Appendix Table 3. Pass-through is the ratio of the vehicle price change to marginal cost change, weighted by predicted sales. Welfare change per person is the same as in Figure 11.

Table A.7: Effects of One Percent Fuel Economy Increase on Manufacturer Profits

Manufacturer	Own-price elasticity of demand	WTP for One percent fuel economy increase	Pass-through rate	Profits change (2015$ per vehicle)	Profits change (million 2015$)
GM	-3.48	92.60	1.137	18	38
Ford	-3.29	87.46	1.118	20	41
Toyota	-3.31	67.80	1.097	3	5
Fiat / Chrysler	-3.34	79.51	1.103	21	30
Honda	-3.23	60.86	1.083	5	7
Hyundai	-3.28	43.88	1.050	-2	-3
Nissan	-3.26	66.70	1.093	6	6
Volkswagen	-3.44	73.28	1.088	-1	-1
Subaru	-3.61	50.46	1.064	-7	-2
BMW	-3.51	115.18	1.142	6	1
Daimler	-3.33	133.81	1.179	22	5
Other	-3.29	118.07	1.133	132	5

Notes: The table reports the welfare results from the same simulation as in Figure 11. Own-price elasticity and WTP are the means across consumers purchasing vehicles sold by the manufacturer, weighted by predicted sales. Pass-through is the ratio of the vehicle price change to marginal cost change, weighted by predicted sales. Profits change per vehicle is the same as in Figure 11.

Table A.8: Comparing Manufacturer Profits with Uniform and Proportional Fuel Economy Increases

Manufacturer	Fuel economy change, uniform (mpg)	Per-vehicle profits change, uniform	Percentage of sales, uniform	Fuel economy change, proportional (mpg)	Per-vehicle profits change, proportional (2015 $)	Percentage of sales, proportional
GM	0.21	18	17.20	0.25	22	17.20
Toyota	0.26	20	17.01	0.19	25	17.02
Ford	0.21	3	13.22	0.25	1	13.22
Honda	0.25	21	11.92	0.17	30	11.93
Hyundai	0.27	5	10.52	0.17	2	10.52
Fiat/Chrysler	0.19	-2	9.73	0.30	-2	9.73
Nissan	0.24	6	7.95	0.24	5	7.95
Volkswagen	0.26	-1	4.56	0.23	-2	4.56
Subaru	0.25	-7	2.68	0.22	-7	2.68
Daimler	0.21	6	1.90	0.28	7	1.90
BMW	0.23	22	1.90	0.26	29	1.90
Other	0.23	132	1.41	0.28	288	1.41

Notes: The table reports the changes in profits per vehicle and the percentages of sales by manufacturer for the uniform and proportional scenarios. See text for description of the scenarios.