EPA Public Access

Author manuscript

Energy Policy. Author manuscript; available in PMC 2020 June 01.

About author manuscripts

Submit a manuscript

Published in final edited form as:

Energy Policy. 2019 June ; 129: 1404–1415. doi:10.1016/j.enpol.2019.02.051.

Simpler is better: Predicting consumer vehicle purchases in the short run

Jacqueline Doremus^{a,*}, Gloria Helfand^b, Changzheng Liu^c, Marie Donahue^d, Ari Kahan^e, Michael Shelby^f

^aEconomics Department, Orfalea College of Business, California Polytechnic State University, San Luis Obispo, 1 Grand Ave, San Luis Obispo, CA 93407, USA

^bOffice of Transportation and Air Quality, U.S. Environmental Protection Agency, 2000 Traverwood Drive, Ann Arbor, MI 48105, USA

^cWalmart Labs, 860 W California Ave, Sunnyvale, CA 94086, USA

^dMinneapolis, MN, USA

^eWashington, DC, USA

^fOffice of Transportation and Air Quality, U.S. Environmental Protection Agency, 1200 Pennsylvania Avenue, N.W. Room 6520J Washington D.C. 20460, USA

Abstract

When agencies such as the US Environmental Protection Agency (EPA) establish future greenhouse gas emissions standards for new vehicles, forecasting future vehicle purchases due to changes in fuel economy and prices provides insight into regulatory impacts. We compare predictions from a nested logit model independently developed for US EPA to a simple model where past market share predicts future market share using data from model years 2008, 2010, and 2016. The simple model outperforms the nested logit model for all goodness-of-prediction measures for both prediction years. Including changes in vehicle price and fuel economy increases bias in forecasted market shares. This bias suggests price increases are correlated with unobserved increases in vehicle quality, changes in preferences, or brand-specific changes in market size but not cost pass-through. For 2010, past shares predict better than a nested logit model despite a major shock, the economic disruption caused by the Great Recession. Observed share changes during this turbulent period may offer upper bounds for policy changes in other contexts: the largest observed change in market share across the two horizons is 6.6% for manufacturers in 2016 and 3.4% for an individual vehicle in 2010.

Keywords

Vehicle demand; Consumer vehicle choice modeling; Validation; Discrete choice modeling

*Corresponding author. jdoremus@calpoly.edu (J. Doremus).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.enpol.2019.02.051>.

1. Introduction

Forecasting how vehicle purchases change in response to changes in fuel economy and prices would provide insights into regulatory impacts of vehicle greenhouse gas and fuel economy standards put forth by the United States Environmental Protection Agency (US EPA) and Department of Transportation (DOT). For high-quality regulatory analysis, agencies seek reliable and replicable forecasts. A “commitment to transparency and parsimony” in policy modeling improves model credibility and clarity over what models can and cannot do (Saltelli and Funtowicz, 2014). The US EPA commissioned a model of consumer vehicle choice from independent researchers. The model they developed, a nested logit model, uses data available to the agencies and is intended to offer “a good compromise between flexibility and simplicity” (Greene and Liu, 2012).

We ask, is the independently developed model of consumer vehicle choice better at forecasting than an even simpler model – one of persistent market shares? We compare prediction accuracy using several goodness of fit measures for two horizons, two years and eight years, and evaluate whether our results are driven by the nested logit model’s sensitivity to parameter values. Our work joins a new research agenda focused on cross-validation and model sensitivity, in contrast to explaining existing variation in the new vehicle fleet and simulating counter-factual scenarios. In transportation policy, recent work has emphasized model sensitivity (Xie and Lin, 2017; Sakti et al., 2017) and the process of model development (Ciuffo and Fontaras, 2017). Machine learning’s emphasis on cross-validation and assessing the quality of model predictions out of sample is spilling over to traditional econometric tools (Athey, 2017).

Our nested logit model fits within a long lineage of models designed to allow consumer substitution across different vehicle types. These models typically describe a static equilibrium; early work lamented that “the treatment of dynamics is not entirely satisfying” (Goldberg, 1995). As an alternative, we offer the simplest dynamic model possible: market shares are persistent and future shares are a function of past shares and an error term. Our simple model is motivated by recent work high-lighting the empirical significance of persistent market shares in oligopoly settings (Sutton, 2007; Bronnenberg et al., 2009).

We use data available to the US EPA, annual vehicle sales data, manufacturer suggested retail prices, and fuel economy, to forecast model market share. We use data from 2008, 2010, and 2016. This time period maximizes the likelihood of a change in consumer vehicle purchases and complements work predicting vehicle sales in a period of economic expansion (Haaf et al., 2014). Economic disruption from the Great Recession reduced consumer spending (Mian et al., 2013), increased unemployment (Mian and Sufi, 2014), and caused particularly severe credit constraints for General Motors and Chrysler (Benmelech et al., 2017). Also during this time, EPA and the Department of Transportation (DOT) announced future GHG and fuel economy standards for light-duty vehicles, though they were not effective until model year 2012. Including 2016 allows us to test a longer horizon that includes the economic recovery from the Great Recession. For both 2010 and 2016, we use changes in vehicle price and fuel economy from 2008 to predict market shares in a nested logit model. Because the model is developed for regulatory analysis, the model focuses on

the variables expected to be affected by standards; we do not incorporate changes in other vehicle characteristics, nor do we observe retail prices.¹ We compare prediction accuracy for the two models using goodness-of- prediction measures including the Kullback-Leibler (KL) divergence, mean squared error and average share error.

We find that, despite major economic and regulatory shocks, the simple model forecast of persistent model market shares outperforms the nested logit for all goodness of prediction measures for both horizons. Our result complements Haaf et al. (2014), who compare predicted sales for 9,000 multinomial logit models and find that a simple forecasting model like ours has the least forecast error in the short run. We differ from Haaf et al. (2014) in several key ways. First, we compare predictions for a model explicitly and independently designed for forecasting by US EPA that uses expert elicitation for nest elasticities and a more complex nesting structure. The model was designed specifically to look at the effects of changing vehicle price and fuel economy, the two variables expected to be affected by regulatory changes.² Second, we offer a theoretical justification for the use of a simple model of persistent market shares by linking them to dynamic models of oligopolies. Third, we test model predictions during the Great Recession, a period of significant economic contraction. In contrast, Haaf et al. (2014) make short run predictions in a period of moderate economic expansion, using data from 2004 to 2006 to predict sales in 2007.³

Together with Haaf et al. (2014), our results suggest a simple model is robust to macroeconomic conditions. While the success of the simple model might not surprise critics of integrated assessment models used in climate change (Pindyck, 2013) or forecasts of oil spot prices (Alquist and Kilian, 2010), it is surprising in the context of the automotive sector, where economists have been concerned about strategic pricing and cost pass-through. Persistent market shares suggest that changes in fleet mix through strategic pricing are an unlikely compliance path to achieve GHG standards.⁴

In our nested logit model, forecasting error arises because increases in manufacturer suggested retail price appear correlated with unobserved increases in a vehicle's relative quality – a result consistent with automakers failing to pass through technology costs, a pattern observed in Europe in a retrospective study (Reynert, 2014). Across the two horizons, the largest observed changes in market share were an increase of 3.4% at the vehicle level and a decrease of 6.6% at the manufacturer level. Given the magnitude of the macroeconomic shocks, these changes serve as upper bounds for structural model policy simulations in other contexts. Our results suggest that model validation may be critical when predicting consumer vehicle purchases. Further, regardless of macroeconomic conditions, it

¹Manufacturers sell vehicles to dealers at an invoice price. The manufacturer's suggested retail price is higher than the average realized retail price due to manufacturer rebates, trade-in value, financing costs, and dealer incentives (e.g. reduced markups). Albuquerque and Bronnenberg (2012) simulate changes in manufacturer and dealer pricing in response to the Great Recession and expect that they would behave similarly: dealer prices would decrease by 13 percent and manufacturer prices would decrease by 11 percent.

²In its regulatory analysis, EPA (2010, p. 47) calculated vehicle costs based on holding other vehicle attributes constant. For instance, changing from a 6-cylinder to a 4-cylinder engine improves fuel economy but reduces power; the cost estimates included the costs of adding a turbocharger to keep power at previous levels.

³Haaf et al. (2014) also predict vehicle sales in 2010, after the Great Recession, but they do so using data from 2004 to 2006, data from a period of economic expansion. In contrast, we use data from the Great Recession to predict sales in 2010.

⁴Greene et al. (2005) find that fuel economy increases are mostly due to adoption of fuel-saving technologies rather than shifts in sales.

appears that a simple model of fixed market shares may be suitable for predicting the future fleet in the short to medium run.

2. Predicting vehicle market share

Rosen (1974) imagined product differentiation as a two-stage game, where firms choose product characteristics and product mix in the first stage and price in the second stage. Recently, scholars have modeled automaker decision-making in the medium-run, allowing automakers to choose price and vehicle characteristics simultaneously, blurring the first and second stages (Knittel, 2011; Klier and Linn, 2012; Whitefoot and Skerlos, 2012; Whitefoot et al., 2017). We focus on the short to medium run, which we define as a period when the fleet of models is relatively fixed. In this horizon, automakers can make minor design changes to existing models but do not have enough time to introduce new models or do significant redesigns in response to the shock. We use the simplest model possible for endogenous product design and compare it with a classic static price-setting model designed to examine sales impacts of changes in fuel economy and vehicle price.

2.1. A simple dynamic model of vehicle market shares

Our simple model assumes the market is in equilibrium and that automakers play a dynamic game. If the market environment is stationary without entry or exit or shocks, optimal market share will be the same in every period. Given shocks, static models overestimate price elasticity and underestimate firm markups compared to dynamic models (Chen et al., 2008). We justify our simple model of persistent market shares by pointing to empirical evidence on the persistence of firm and brand market share in the short run (Sutton, 2007) and across space (Bronnenberg et al., 2009).

Market share persistence may be explained, in part, by brand loyalty, particularly when a brand offers a single model within a class of vehicles (e.g. Ford Focus, Ford's only Compact car in 2008 and 2010). Brand loyalty creates barriers to entry and gives first movers a durable advantage (Bronnenberg et al., 2012). For autos, Anderson et al. (2015) show evidence that brand loyalty is transmitted from parents to children. Mannering et al. (1991) document brand loyalty in new cars and Mannering et al. (2002) find the same pattern in auto leases. Brand loyalty creates a tension between the current period and a stream of future discounted profits: firms may be willing to trade lower markups today for a stream of higher markups in the future, making them less likely to pass through technology costs or pursue a sales-mixing compliance strategy.

Market share persistence may also be explained by endogenous product design if automakers comply with fuel economy standards by changing minor vehicle characteristics in a segment. Because people were satisfied with their previous vehicles, keeping the vehicle apparently constant might keep them in the same market.⁵ If changes in unobserved desirable product attributes targeted to buyers in that market are accompanied by an increase

⁵Work in this area tends to use a mixed logit framework to capture consumer heterogeneity. Their main challenge is to credibly represent unknown consumer and manufacturer tradeoffs between fuel economy technology and other vehicle characteristics (acceleration, weight, footprint, etc.).

in price, this could explain persistent market shares as well as weaker prediction quality in a demand model. This strategy would substantially reduce compliance costs from fuel economy regulation. Evidence suggests that automakers are changing attributes, though it is less clear whether they are unobserved. Knittel (2011) argues that, given annual technological advancement, increasing fuel economy may be more a matter of not continuing to increase weight and acceleration. Whitefoot and Skerlos (2012) model automakers choosing footprint, acceleration, and other fuel-economy related technology as they set prices; they estimate that automakers will comply with US EPA fuel economy standards in part by increasing vehicle footprint. Klier and Linn (2012) model firms choosing power and weight (but not engine design), as well as prices, in their “medium-run” horizon. As compared to a price-only response, they find that endogenous product design decreases the cost of compliance. Reynaert (2014) estimates a mixed logit model for automakers facing Europe’s greenhouse gas emissions standards and concludes that the dominant response was to use technology to decrease emissions, as opposed to strategic price-setting to influence sales-mix; Greene et al. (2005) found a similar result using a calibrated nested logit model. Train and Winston (2007) consider the long run and conclude that changes in vehicle characteristics such as size, power, operating cost, transmission type, reliability, and body type, explain decreases in domestic market share in the US auto market.

Finally, persistent market shares may be driven by behavioral responses by either consumers or dealers. Dealers order their stock of vehicles at the beginning of the model year and may base their order on vehicles sold the previous year, intensifying any underlying persistence in vehicle market shares. Dealers may then use pricing and financing incentives to sell their inventory, which is similar the previous year, even when consumer preferences change. For their part, consumers may be most interested in vehicle characteristics that are persistent across model years, meaning that decreases in price or increases in fuel economy may be less salient than fixed characteristics.

2.2. A static price-setting model of discrete consumer choice

The forecasting model developed for US EPA models vehicles as differentiated products sold by multiproduct oligopolists in a one-shot Bertrand competition. Similar to other vehicle choice papers, the model assumes vehicle characteristics (other than fuel economy and price) are fixed in the short run and focuses on the second stage, price-setting. Heterogeneity in consumer preferences is represented in a nested logit (NL), similar to Goldberg (1995). The main alternatives to NL are mixed logit (e.g. Berry et al., 1995) and a linear system of equations (e.g. Austin and Dinan, 2005).⁶ Each approach estimates a demand system, which connects prices from policy change scenarios to predicted changes in vehicle sales. For NL, the modeler embodies heterogeneity in consumer preferences through the nesting structure instead of estimated random coefficients, as in the mixed logit case. Goldberg used a NL when estimating the impact of fuel economy standards on vehicle

⁶An alternative to discrete choice models is to use a system of linear equations, own-and cross-price supply and demand elasticities by vehicle, to predict sales in response to an increase in price. For example, Austin and Dinan (2005) use proprietary demand elasticities from General Motors (used in Kleit, 2004) and inferred supply elasticities from an equilibrium assumption, observed dealer markups, and an assumption on how these relate to manufacturer markups. They use this system of equations to estimate how vehicle sales mix changes in response to expected changes in vehicle net price as a result of fuel economy standards.

choice (1995, 1998). More recently, Greene et al. (2005) and Harrison et al. (2008) used NL to evaluate the 2011–2015 CAFE standards. NL has also been used to better understand demand for alternative fuel vehicles (Brownstone et al., 1996) and used by other regulatory agencies (Bunch et al., 2011).

US EPA contracted with researchers to develop a model of consumer vehicle choice; the NL framework was justified by Greene and Liu (2012) as “readily calibrated with only a small amount of information ... [and it] allows for substantial flexibility in representing substitutions among vehicle types.”⁷ Vehicle sales are predicted to change in response to changes in net vehicle price, which is calculated as the increase in vehicle cost associated with technologies to reduce GHGs, less a discounted share of future fuel savings associated with those technologies.⁸ Demand elasticities for each vehicle nest are not estimated from an original data set, but rather are based on reviewing estimates in the literature (Greene and Liu, 2012, Table 4). This approach allows for synthesis of the results from multiple analyses, and professional judgment about whether the values are appropriate. The parsimonious model design avoids adding additional uncertainty from projecting changes in other vehicle characteristics and consumers’ marginal willingness to pay for changes in vehicle characteristics. Estimates of consumers’ marginal willingness to pay for vehicle attributes vary greatly and may be sensitive to model formulation and estimation (Greene et al., 2018).

The model is designed for static, same-year analysis of the effects on vehicle sales of adding fuel-saving technologies and their costs; that is, it is intended to compare vehicle sales with and without fuel-saving technologies and additional costs for a single fleet of vehicles. In principle, then, changes in the economy, demographics, or the fleet over time should not affect the ability of the model to predict, because it is predicting against a same-period static counter-factual. Over time, changing conditions might lead to changes in consumer responsiveness.

Though a forecasting model that includes macroeconomic and demographic information, as well as other changes in vehicle characteristics, might appear to be more suitable for projecting the impacts of standards in the future, such a model’s effectiveness depends on the quality of the forecasts for those additional factors, and how they interact with other model parameters.

3. Data and institutional context

The nested logit model was designed to use data assembled by EPA and DOT for their analysis of GHG and fuel economy standards for MYs 2017–25 (U.S. EPA and DOT, 2012). We test goodness-of-prediction using annual vehicle sales and price data for model years

⁷Mixed logit offers a more flexible representation of consumer heterogeneity. Since Berry et al. (1995), it has been used in vehicle choice modeling (e.g. Petrin, 2002; Jacobsen, 2013) but model predictions may be sensitive to start values and the optimization routine (Knittel and Metaxaloglou, 2014). When prices are correlated with unobserved vehicle characteristics, demand models require instruments to consistently estimate preference and cost parameters.

⁸Greene (2010) found highly varied estimates in the literature of consumer willingness to pay (WTP) for additional fuel economy in the vehicle purchase decision, with a number of studies showing WTP less than the expected value of future fuel savings, and some others showing overvaluation. The model allows a user to choose the number of years of expected fuel savings that vehicle buyers are believed to consider in their purchase decisions, as well as the future fuel prices and discount rate they might use for those calculations.

(MY) 2008, 2010 (EPA and DOT, 2012), and 2016 (U.S. EPA, 2018). Price is the manufacturer's suggested retail price (MSRP).⁹ Each model year contains over a thousand unique vehicles at the trim level. For example, there are 20 different versions of the Chevrolet Silverado in the 2008 data, each unique based on engine, footprint, fuel economy, and other attributes. We aggregate trims using the sales-weighted average vehicle price and fuel economy for each vehicle.¹⁰ Total sales fall by about 20% during the recession. We focus on vehicle market share, not the level of vehicles sold.

To predict market share in 2010, each MY 2008 vehicle needed to be matched with its MY 2010 counterpart. This matching is not straightforward.¹¹ Vehicles enter and exit the market between any two model years; for example, Saab dropped out of the market entirely during this time and General Motors dropped 21 models between 2008 and 2010. After aggregation, we match 524 vehicles that capture 94% of the MY 2008 vehicles sold, and 91% of MY 2010 vehicles sold. Dropped model market share is about 13% of total sales in each year.

To predict sales in 2016, we aggregated data to the manufacturer-class level. Match quality was similar to that in 2010. Table 1 compares the disaggregated fleet to the aggregated fleet. Prices and fuel economy do not differ greatly. In total, 108 vehicles remained unmatched because they were manufactured in one year but not in the other; these are dropped in the 2010 analyses.¹² Changes in available models between two years are fairly common in the auto industry, and are among the challenges associated with predicting changes in vehicle purchases.

Annual vehicle-level variation in the percentage change in price and fuel economy for 2010 and 2016 is shown in Fig. 1. The distribution for price is nearly symmetric about zero, whereas the distribution for fuel economy is skewed toward increases in fuel economy. Together, these figures fail to suggest that changes in price are determined by increased costs from fuel economy technology. On average, fuel economy increased by 3% in 2010, in keeping with the trend for fuel economy in the mid-2000s in the U.S. market (US EPA, 2016).¹³

Our model evaluation compares sales in 2010 to sales in 2008, a window of time punctuated by macroeconomic shocks and trends. Housing prices fell, decreasing household wealth; in

⁹MSRP differs from the transacted price – it fails to reflect cash incentives to customers or dealers, which are larger for less fuel efficient vehicles and vary with gas prices (Langer and Miller, 2013). Proprietary datasets with transaction prices for a sample of transactions exist (e.g. Autodata Solutions, the dataset used in Langer and Miller, 2013 or the proprietary data used in Busse et al., 2013). Other researchers use MSRP (e.g., Austin and Dinan, 2005; Bento et al., 2009), in part for its accessibility. In part for accessibility and in part because it is prospective (realized retail prices are only found in retrospect), the EPA model uses MSRP. 2016 prices are deflated to 2008 dollars.

¹⁰This is similar in magnitude to Reynaert (2014), who has about 400 modelengine variants per market and captures 80% of the total sample, dropping vehicles with small market shares.

¹¹In some cases, vehicles change classification across the years, e.g. from Subcompact to Subcompact Prestige. In these cases, we use the 2008 EPA classifications/nests for predictions and observed 2010 vehicle sales outcomes.

¹²As a robustness check, we aggregate trims to each manufacturer-class for 2010 as well. In this specification, for example, we predict future market shares for Nissan compact cars to be the same as current market shares for the suite of compact cars offered by Nissan. This approach allows us to include models that were unmatched across years. However, we are still obligated to drop manufacturers not observed in both periods (e.g. Saab). Our result – that the simple model predicts better than the nested logit – is robust to this aggregation strategy. Results available upon request.

¹³Note that the relationship between improvement in fuel economy and reduced fuel consumption is non-linear. See Larrick and Soll (2008).

response, consumer spending dropped (Mian et al., 2013). Reduced consumer demand caused an increase in unemployment (Mian and Sufi, 2014). Meanwhile, auto credit markets contracted, with the most severe decreases in low income markets (Amronin and McGranahan, 2015) and for leasing companies related to General Motors and Chrysler (Benmelech et al., 2017). Gas prices were very high in early to mid-2008, exceeding \$4 per gallon. Consumers responded by purchasing more fuel efficient vehicles in 2008 (Busse et al., 2013). Lower gas prices in 2010 reduced returns to fuel economy. In 2009 the US government lent Chrysler and GM more than 20 billion dollars to enable them to make payments to workers and creditors. The bailout may have reduced enthusiasm for GM and Chrysler vehicles from American consumers in 2010. Few economists were optimistic that either company would withstand the Great Recession (Goolsbee and Krueger, 2015). Changes between 2010 and 2016 were less dramatic, characterized instead by a steady economic recovery.

Two vehicle policy changes occurred between 2008 and 2010: the federal Cash for Clunkers program and the announcement of new light duty vehicle standards by EPA and DOT. The Cash for Clunkers program increased vehicle purchases in 2009 in part by pulling forward about 300,000 vehicles that would have been purchased in 2010 (Mian and Sufi, 2012; Li et al., 2013).¹⁴ In 2009, the US EPA and DOT announced light-duty vehicle GHG and fuel economy standards for model years 2012–16, finalized in early 2010.¹⁵ Though the standards would not go into effect until model year 2012, it is possible that automakers started to change vehicle characteristics in preparation for compliance with the standard, accelerating the deployment of technology related to fuel economy.¹⁶

4. Methods

Like some others modeling consumer vehicle choice (e.g. Sen et al., 2017; Haaf et al., 2014; Xie and Lin, 2017), we focus on market shares, rather than vehicle sales. In our case, market shares mitigate changes in total sales from the Great Recession. We examine whether the simple model or the static model performed better in predicting changes in the composition of the fleet, as measured by market share, in the face of the Great Recession's significant market shock.

Consider a discrete choice representative consumer framework where utility is derived from unobserved vehicle characteristics A_{jk} and vehicle cost G_j . To simplify notation, we describe the model for two levels of nests. For vehicle j in class k , vehicle cost is a combination of price C_{jk} and a share of the present discounted value of savings F_{jk} from fuel economy.¹⁷

¹⁴ To be eligible for the Cash for Clunkers program, a new car needed to have a price below \$45,000 and get at least 22 MPG. For SUVs, medium-duty passenger vehicles, pickup trucks, minivans and cargo vans, the MPG requirement was 18 MPG. For vans with a wheelbase that exceeded 124 inches or large pickup trucks with a wheelbase that exceeded 115 inches, the requirement was at least 15 MPG. Very large vans and trucks (Category 3) had no MPG minimum but were required to be made no later than MY2001.

¹⁵ These standards were in part a response to the Energy Independence and Security Act of 2007, which required a fleet average fuel economy of at least 35 miles per gallon (mpg) by 2020. The new standards use a footprint-based system, which relates a manufacturer's emissions and fuel economy obligations to the vehicle footprint, defined as the area between the wheelbase and track and measured in square feet.

¹⁶ Reynaert (2014) found that automakers updated vehicle technology in anticipation of EU GHG standards, reaching compliance before the deadline.

$$U_{jk} = A_{jk} + B_k G_j + \varepsilon_{jk} = A_{jk} + B_k (C_{jk} - F_{jk}) + \varepsilon_{jk} \quad (1)$$

If the error term is distributed extreme value, the choice probability for alternative j is $P_{jk} = P_{j|k}P_k$, which is equal to its market share.¹⁸ The conditional probability of choosing alternative j given that an alternative in class k is selected is $P_{j|k}$, defined as

$$P_{j|k} = \frac{\exp(A_{jk} + B_k G_j)}{\sum_{j' \in k} \exp(A_{jk} + B_k G_j)} \quad (2)$$

The marginal probability of choosing an alternative in class k is

$$P_k = \frac{\exp\left(A_k + \frac{B_{root}}{B_k} \ln\left(\sum_{j \in k} \exp(A_{jk} + B_k G_j)\right)\right)}{\sum_k \exp\left(A_k + \frac{B_{root}}{B_k} \ln\left(\sum_{j \in k} \exp(A_{jk} + B_k G_j)\right)\right)} \quad (3)$$

where B_{root} is the generalized cost coefficient for vehicle classes.

4.1. Simple model forecast: persistent market shares

A simple econometric model of vehicle market shares inspired by Sutton (2007) predicts that oligopolists' vehicle market shares are persistent across time. Formally, this corresponds to the coefficient B_k equal to zero in equation (1). Vehicle shares are still a function of the inferred vehicle constant term A_{jk} , though the constant is not calculated. In the simple model, forecasted shares are equal to past shares:

$$E[P_{j,t+1}] = P_{j,t} \quad (4)$$

As in the nested logit model, our unit of analysis is annual market share for a model made by a manufacturer for 2010 and manufacturer-nest for 2016. Data are at the annual level to reflect the types of data available to US EPA. Assuming persistence at a fine level like the vehicle is in keeping with recent evidence that consumers shift between vehicles with the same footprint in response to gas shocks (Leard et al., 2017).¹⁹

¹⁷ Great debate exists about the role of fuel savings in consumer vehicle purchases (Greene, 2010; Greene et al., 2018). Some researchers find that vehicle buyers consider the expected full lifetime of fuel savings in deciding on how much fuel economy to buy (e.g., Busse et al., 2013), while others argue that buyers put much less emphasis on fuel economy (Gallagher and Muehlegger, 2011).

¹⁸ Formally, the probability is $P(j \cap k) = P(j|k)P(k)$.

¹⁹ As noted above, as a robustness check, we aggregated the data to the manufacturer-class level. Results available upon request.

4.2. Nested logit model

Greene and Liu (2012) developed a vehicle choice model for EPA intended to estimate changes in total sales and fleet mix in response to GHG and fuel economy standards. Given a vehicle fleet, the model allows regulators to compare vehicle sales and fuel economy with and without fuel economy standards.²⁰ The nested logit model (NL) forecasts vehicle market share using observed shares, fuel economy, and price in 2008 and changes in fuel economy and price in 2010. The NL does not estimate price elasticities in calculating the cost coefficients B_k . Instead, the cost coefficients are approximated using elasticities from the literature:

$$B_k = \frac{\eta_k}{\bar{p}_k(1 - \bar{S}_k)} \quad (5)$$

where, for vehicles in class k , η_k is the own-price elasticity, $\bar{p}_k$ is average price, and $\bar{S}_k$ is average conditional market share for vehicles in class k . This expression is derived from the definition of elasticities and logit model equations (Train, 2009). Tables 2 and 3 provides the elasticities used in the analysis, which, as discussed above, are the synthesis of estimates from multiple sources (Greene and Liu, 2012). The nested logit restricts demand elasticities for the nests: responsiveness to price must be highest at the individual-vehicle level, and decrease at each higher nest.

Given nest cost coefficients, each vehicle's constant term is calibrated to fit baseline sales data. For any two vehicles within a class, the ratio of their probabilities is equal to the ratio of their shares.

$$\frac{P_{i|k}}{P_{j|k}} = \frac{S_i}{S_j} = \frac{e^{A_{ik}}}{e^{A_{jk}}}, \quad \forall i, j \in k \quad (6)$$

$$A_{ik} - A_{jk} = \ln S_i - \ln S_j, \quad \forall i, j \in k \quad (7)$$

One vehicle's constant is normalized to zero, giving $A_{ik} = \ln S_{ik} - \ln S_{1k}$, $\forall i \in k$, where $\ln S_{1k}$ is the share of the normalized constant. Observed shares and cost coefficients from prices and fuel economy in 2008 pin down vehicle and class-level constants. The constant for not buying a vehicle is assumed to be zero.

²⁰ EPA regulates GHG emissions from vehicles; the Department of Transportation regulates vehicle fuel economy. Because the primary way to reduce GHG emissions is to improve fuel economy, the agencies have largely harmonized their regulations (U.S. EPA and DOT, 2010, 2012). The model uses fuel economy rather than GHG emissions, because fuel economy is more salient an attribute to vehicle buyers and directly affects expenditures.

As mentioned above, Greene (2010) found highly varied estimates in the literature of consumer willingness to pay (WTP) for additional fuel economy in the vehicle purchase decision. Fuel Savings for vehicle j in model year t relative to its baseline configuration are

$$F_j(t) = \sum_{\tau=t+1}^{t+L} \frac{1}{(1+r)^{\tau-1}} FP(\tau) M(\tau-1) \left[\frac{1}{\xi MPG_i^0(t)} - \frac{1}{\xi MPG_i^1(t)} \right] \quad (8)$$

where r is the consumer discount rate, L is the payback period in years (that is, the number of years of future fuel savings a buyer considers in the purchase decision relative to up-front costs), $FP(\tau)$ is the price of fuel in year, ξ is the OnRoad discount factor that discounts fuel economy in order to reflect real-world driving conditions,²¹ and $M(\tau-1)$ are the annual miles traveled for a vehicle of age $\tau-1$.

The default payback period is assumed to be five years, consistent with “the length of a typical new light-duty vehicle loan” (U.S. EPA and DOT, 2010, p. 25517). Future fuel savings use a default discount rate of 3%.²² Future fuel prices are reported in 2008 dollars. Year 1 fuel prices are that year’s realized fuel prices and come from US Energy Information Administration’s (EIA) Annual Energy Outlook (2012). Projected fuel prices come from EIA’s Annual Energy Outlook Petroleum Product Prices (2008, 2010).²³ In Fig. A-I in the appendix, we see that in 2008 analysts expected both gasoline and fuel prices to decrease over a twenty-year period. In 2010, however, analysts expected gasoline prices to increase over the next twenty years. Annual miles traveled differ for cars and trucks and decrease over time; see the second panel of Fig. A-I. These values are used by EPA in estimating the effects of MY 2012–16 vehicle fuel economy/GHG standards on vehicles (US EPA and DOT, 2010).²⁴

4.2.1. Nesting framework—The five-nest structure, shown in Fig. 2, was independently developed by Greene and Liu for US EPA (2012). The structure is similar to Bunch et al. (2011) and NERA (2009). The first layer constitutes the buy/don’t buy decision. Next it distinguishes between passenger vehicles, cargo vehicles, and ultra-prestige vehicles.²⁵ The model then separates passenger vehicles into Two Seaters, Prestige Cars, Standard Cars, Prestige SUVs, Standard SUVs, and Minivans, and cargo vehicles into Pickup Trucks and Vans. The next level – the last nest – continues the division into classes. At this nest, the model uses a simple logit model. Within this nest are the individual vehicles for 2010 and

²¹ “OnRoad Discount’ is used in fuel cost calculation to discounts EPA fuel economy (MPG) test value, which is displayed in fuel economy window stickers and used in the [model], to better reflect fuel economy under real-world driving conditions” (Greene and Liu, 2012, p. 39).

²² In the technical appendix, we investigate the sensitivity of model predictions to variation in the elasticities, payback period, and discount rate.

²³ Motor gasoline prices are calculated based on “... sales weighted-average price for all grades. Includes Federal, State, and local taxes” and “Diesel fuel for on-road use. Includes Federal and State taxes while excluding county and local taxes.”

²⁴ Since the payback period is five years, only initial differences in realized and projected gas prices and miles traveled change the vehicle’s net price.

²⁵ In the model, sport-utility vehicles and minivans are included as passenger vehicles, although many of these vehicles are considered light-duty trucks for regulatory purposes. Consumers commonly consider these to be passenger vehicles; it is more likely, for instance, that people consider an SUV to be a substitute for a large or midsize car than for a pickup truck. Because the model is meant to reflect consumer decision processes, it was considered appropriate to nest SUVs and minivans as passenger vehicles rather than cargo vehicles.

the weighted averages for manufacturers in 2016. Within nests, logit exhibits “independence of irrelevant attributes:” the ratio of probabilities (or market shares, in this model) of two options does not vary if a third option is added to the mix. As a result, an increase in the market share of one alternative within a nest draws proportionately from all other alternatives (Train, 2009). Across nests, independence of irrelevant attributes does not hold.

Vehicles are classified into nests by type, size, and relative price. Vehicles with prices higher than or equal to unweighted average price in the corresponding EPA class are classified as Prestige vehicles.²⁶ Ultra-Prestige vehicles have prices that exceed \$75,000. Non-prestige SUVs are divided into small, midsize and large SUVs by vehicle footprint.²⁷

4.2.2. Nested logit forecasts—The NL uses data from 2008 to estimate the cost coefficient B_k , determined from its own-price elasticity η_k from the literature and average price $\bar{p}_k$ and average conditional market share $\bar{s}_k$ within class k according to expression (5). Given cost coefficients B_k , the model calibrates vehicle-level constants A_{jk} to fit shares in 2008. To forecast shares, the model uses these cost coefficients and vehicle-level constants as well as the 2010 vehicle cost G_j in expressions (2) and (3) to calculate the choice probability for alternative j , $P_{jk} = P_{jk} P_k$. 2010 vehicle cost G_j is determined by the vehicle’s 2010 price and fuel economy.

4.3. Measuring goodness-of-prediction

We use several measures to compare goodness-of-prediction across the simple and nested multinomial logit models. The first is Kullback-Leibler (KL) divergence, which measures the gap between an observed distribution and a predicted distribution (Kullback and Leibler, 1951). Maximum likelihood estimation attempts to minimize the KL divergence, defined as

$$KL(P|Q) \equiv \sum_i P_i \ln \left(\frac{P_i}{Q_i} \right) \quad (9)$$

where P_i is the observed market share and Q_i is the predicted market share. Next, we use two related goodness-of-prediction measures, mean squared error and average share error:

$$\text{mean squared error} \equiv \frac{1}{N} \sum_i \{P_i - Q_i\}^2 \quad (10)$$

$$\text{average share error} \equiv \frac{1}{N} \sum_i |P_i - Q_i| \quad (11)$$

²⁶ For example, the Prestige Two-Seater class is the set of relatively expensive vehicle configurations in EPA’s class of Two-Seaters, those with prices higher than or equal to the unweighted average price of EPA Two-Seaters.

²⁷ Small SUVs have a footprint less than 43 square feet, Midsize SUVs have a footprint between 43 square feet and 46 square feet, and Large SUVs have a footprint greater than or equal to 46 square feet.

Like Haaf et al. (2014), we plot a cumulative distribution function (CDF) of absolute error, $|P_i - Q_i|$, to compare model performance at different error thresholds. The CDF of error shows the percentage of model predictions that have error less than a given threshold. To measure the NL's accuracy in predicting growth or decline in market share, we create a binary variable equal to one if the direction of the NL prediction matches the observed change in vehicle market share and assess the fraction of predictions correctly reported.

5. Results

Table 4 reports our goodness-of-prediction measures across the two models. For all measures except Direction, a low value indicates less model error; for both years, for all error measures, the simple model outperforms the Nested Logit (NL). Kullback-Leibler Divergence measures the difference between a predicted distribution and the true distribution. It must be positive and would be zero if the model perfectly predicted the observed distribution. Table 4 shows that the simple model has a lower Kullback-Leibler divergence as compared to the NL. Mean squared error shows a similar pattern: the mean squared error for the simple model is nearly half that of the NL model. Mean squared error uses a quadratic loss function, which penalizes outliers. Average share error for the simple model is about two-thirds that of the NL. The NL correctly predicted the direction of market share change about two-thirds of the time. It appears that the simple model's error distribution is less biased and its average error is lower than the NL.²⁸

Next, we turn to the graph of the cumulative distribution of errors for both models for 2010 in Fig. 3. The cumulative error function reports the share of predictions that fall within a certain error tolerance. For example, 85–90% of the models' predictions are within an error tolerance of 0.002. A higher value for a given error tolerance implies that the model has a greater share of total predictions with an error less than the tolerance. The simple model has a higher share of predictions within the tolerance for all error tolerances. For very low tolerances, such as those below 0.0005, the difference between the two models is small.

The better forecasting performance of our simple model is in line with other papers that test out-of-sample predictions of vehicle choice, particularly those that compare estimates to a simple model of static market shares. Haaf et al. (2016) use data from MY 2004–6 vehicles to estimate a number of different econometric models and test their predictions against MY 2007 and 2010 vehicle sales. They find that a simple model like ours – that is, one that assumes constant market shares – performs well compared to other models for one year forecasts of MY 2007. Haaf et al. (2014) also find that a simple model outperforms multinomial, mixed, and nested logit models, even when varying the functional form of the utility function and the set of attributes used for the prediction. Like Haaf et al. (2016), we find that a smaller cost coefficient – zero, in our case – improves our predictions. Including changes in price biases estimates in the wrong direction. Price increases correlated with increases in unobserved vehicle quality, changes in preferences, or brand-specific changes in market size may explain this finding, due to their expected opposite effects on the price coefficient.

²⁸ These results hold even when we aggregate vehicles to the manufacturer-class level.

For the short run, these results suggest that persistence of market shares is a strong predictor of future market shares. For a horizon of at least 4 years, Haaf et al. (2014) find that including vehicle attributes improves prediction accuracy over the simple model. They cite entry by new vehicles as the reason for the simple model's poor prediction accuracy over a longer horizon. A model including vehicle attributes, however, requires forecasts of those attributes in the future, which increases uncertainty. Uncertainty grows non-linearly when these forecasted parameters interact, potentially making the model less useful for policy (Saltelli and Funtowicz, 2014). We find that the simple model continues to outperform the nested logit in 2016.

5.1. Relationship between changes in price and market share

To get a better sense of the relationship between changes in price and market share, Fig. 4 shows patterns for the six largest manufacturers in 2010. The x-axis is the percentage change in price (MSRP) in 2010, relative to 2008. The y-axis is the percentage change in market share in 2010, relative to 2008.²⁹ To make the graphs easier to read, the percentage changes are top-coded at 200% for the share increase and 40% for the price increase. The NL model has a cost coefficient term B_k that is negative and approximated using elasticities reported in the literature. If there is a negative relationship between price and share, the scatter plots should be downward sloping from left to right. Honda comes closest to exhibiting this relationship; the other automakers fail to show such a pattern. For example, for Toyota the relationship between price and market share appears to be positive. This would be consistent with prices accompanying improved vehicle quality or an increase in demand.³⁰

For the remaining automakers, the relationship is unclear. Chrysler and General Motors shares appear shifted downward, consistent with a brand-specific shock such as limited non-bank financing for these two companies (Benmelech et al., 2017). Ford may have benefited from these restrictions, with several vehicles experiencing large increases in market share and price in the upper right quadrant. Toyota, Honda, and Nissan had a small set of vehicles that experienced increases in market share. For Toyota, the Camry, Prius, and Corolla saw increased sales in 2010 relative to 2008. Marginal buyers who abstained from purchasing during the Great Recession may have been those that preferred manufacturers' less popular vehicles.

Announced standards may have led to product design changes to increase fuel economy that failed to pass technology costs through to price. Fig. 1 fails to suggest a strong relationship between price and fuel economy changes. Reynert (2014) shows evidence that automakers in Europe responded to announcements of standards enough to reach compliance early, and they complied via technology, not sales-mixing. For the U.S., trends in fuel economy technology penetration support this prediction: Gasoline Direct Injection was used in less

²⁹ For example, if a vehicle had a market share of 0.5% in 2008 and 0.4% in 2010, the relative change would be 20%.

³⁰ We use the 2008 MSRP and the change in MSRP instead of transacted price, which includes cash incentives for dealers and customers. Using MSRP will bias the nested logit predictions if automakers chose MSRP and cash incentives jointly, instead of using cash incentives for unanticipated responses to fluctuations in demand. For example, if automakers increased MSRP but added cash incentives such that the transacted price falls, this could increase sales (either directly, from a lower price, or indirectly, through reference dependence). We know that automakers use cash incentives defensively, to address regional demand shocks or gas price shocks (Langer and Miller, 2013), but we are not aware of evidence that they are jointly determined. If they are, this may be another reason to use a simpler model when predicting vehicle market share.

than 3% of vehicles in MY 2008 and is projected to reach half of the MY 2016 models (US EPA, 2016). The number of car models with a minimum of 30 MPG was 28 in 2008, 40 in 2010, and expected to be nearly 70 in 2016 (EPA, 2016). The pattern is similar for SUVs; the number of SUV models with at least 25 MPG was 10 in 2008, 14 in 2010, and nearly 50 in 2016. The diffusion of fuel-saving technology to rivals is good for reducing greenhouse gas emissions; Sutton (2007) describes a similar pattern in Japanese oligopolies. However, diffusion also reduces returns to investment in innovation.

5.2. Changes in nest and automaker market share

Fig. 5 reports observed absolute changes in market share between 2008 and 2010 and 2008 and 2016 by nest and manufacturer in order to offer upper bounds for structural model policy simulations. We focus on the changes in 2010 here; patterns are mostly similar for 2016. In 2010, sales grew slightly more concentrated: the Herfindahl-Hirschman Index grew from 1,271 in 2008 to 1,308 in 2010. We report results for automakers and the lowest level nests that had at least 4% market share.³¹ We find that the Great Recession increased the share of Standard Cars at the expense of prestige cars, prestige SUVs, and pickups, among others. The Standard Car class market share increased by almost 5 percentage points, mainly from increased sales for four vehicles: the Ford Fusion, and the Toyota Camry, Corolla, and Prius. Corolla sales tripled in 2010, despite flat prices and fuel economy and lower priced rivals.

In contrast, GM and Chrysler/Fiat lost 5% and 1.5% market share. These decreases in market share may be driven by less access to non-bank lenders associated with GM and Chrysler, who were constrained during the Great Recession (Benmelech et al., 2017). They could also be driven by brand-specific shocks due to the auto bailout (Goolsbee and Krueger, 2015). The net effect was a loss of market leadership by GM, a result that echoes the importance of idiosyncratic, industry-specific shocks as a driver of changes in market leadership in Japanese oligopolies (Sutton, 2007). Toyota and Ford market share increase by about 4% across the two years.

5.3. Nested logit model sensitivity

One possible reason for the superior performance of the simple model is if the nested logit model is sensitive to parameter values and we use the wrong parameter values. This is a potential problem with more complex models in policymaking – they create more opportunities for uncertainty and disagreement about parameter values, potentially rendering them less credible in the eyes of the public (Saltelli and Funtowicz, 2014). To investigate model sensitivity, we took the base fleet in 2008 and simulated a 20 percent increase in fuel economy to all vehicles and no increase in price. We then compared predicted market shares with our preferred parameter values and a 20 percent increase in fuel economy to: variation in the model's initial fleet; the expert elicited elasticities; the discount rate; and the payback period. Full results from the sensitivity analysis are reported in the appendix. We found that the model's predictions were fairly insensitive to changes in parameter values and transformations of the initial fleet, making it unlikely that the simple model's success is

³¹ The full set of classes and manufacturers are used in the fleet-level results.

driven by parameter values. It could, however, be driven by the nesting structure, the interaction of parameter values, or data quality (i.e. manufacturer's suggested retail price instead of realized retail prices).

6. Discussion

A priori, the way that auto manufacturers will achieve GHG/fuel economy standards is ambiguous. On the one hand, the automakers can add fuel-saving technologies, which are likely to increase prices if automakers pass through costs.³² Another possibility is that standards may induce strategic pricing, where firms achieve compliance by shifting the mix of vehicles sold through prices—a strategy known as “sales-mixing” (e.g., Goldberg, 1998; Jacobsen, 2013).³³

We can reconsider the choice between adding technology and using price in the context of a dynamic oligopoly model where the equilibrium is persistent market shares. If consumers form habits, sales today affect future sales. When choosing a vehicle, changes in price or fuel economy may be less salient to consumers than persistent, unobserved vehicle characteristics. Instead of passing through technology costs, firms may choose lower markups today in order to avoid losing sales both today and in the future. Likewise, though a sales-mixing strategy might reduce the cost of compliance today, it may bring large future opportunity costs. If markups are higher on less fuel efficient vehicles, and consumers form habits, shifting consumers to more efficient vehicles results in fewer future sales on vehicles with higher markups.

Given tradeoffs between current and future profits, automakers may respond to standards with increased innovation in and adoption of fuel economy technology. In this way, they could improve fuel economy at lower cost, keeping current and future markups high. An innovation response would be in line with the Porter Hypothesis, which states that environmental regulations may “trigger innovation that may partially or more than fully offset the costs of complying with them” (Porter and Van der Linde, 1995, 98, cited in Ambec et al., 2013). An innovation response may be more likely in oligopoly settings, where market leaders may prefer to “coast” and neglect investment in research and development unless their market leader positions becomes threatened (Ericson and Pakes, 1995). For industries in Japan, Sutton (2007) documents persistent market share and, for some industries, shares remain stable because successful innovations are quickly imitated by rivals.

For vehicles, there is evidence of an innovation response. Klier and Linn (2016) show an increased rate of adoption of fuel-saving technology in response to fuel economy standards. Reynaert (2014) found that firms used a technology response to comply with European

³² For instance, improving the fuel economy of a minivan from 18 miles per gallon (mpg) to 22 mpg, as occurred with the Honda Odyssey between 2010 and 2014, is expected to save \$400/year, according to fuelconomy.gov. The 2014 model is listed as about \$2000 - \$4000 more expensive than the earlier version; it has changed in characteristics other than fuel economy during that time.

³³ In the medium to long run, they may change vehicle characteristics such as 1) improved technology (Klier and Linn, 2012) or changes in vehicle size (Whitefoot and Skerlos, 2012; Whitefoot et al., 2017) or 2) changing the vehicles they offer in their fleet (introducing more efficient models or retiring models that fail the standards).

GHG standards. In this case, fuel economy standards would increase innovation, lowering the costs of achieving the standards.

7. Conclusion and policy implications

Within the Great Recession, we assessed the relative forecasting performance of a nested logit model and a simple model of persistent market shares. The nested logit model was developed for predicting the future vehicle fleet given fuel economy standards and technology costs. The simple model was motivated by empirical evidence of persistence in oligopolies. Using vehicle sales in 2008, 2010, and 2016, we compared each model's goodness-of-prediction using three error measures. For each measure, for each horizon, we found that the simple model's forecasts were less biased than those from the nested logit model.

Comparing the simple model to the nested logit framework, we learn that including information on vehicle price and fuel economy increased prediction bias. Sensitivity analyses of the nested logit model, reported in the appendix, make it unlikely that this result is driven by parameter values. Though it is possible the nested logit's poor performance could be driven by model design, e.g. nest structure or combination of parameter values, or the discrepancy between manufacturer's suggested retail price and realized retail prices, better prediction performance from the simple model is consistent with other vehicle modeling comparisons (Haaf et al. 2014, 2016) and in other contexts (e.g. healthcare, Bayati et al., 2018, or oil spot prices, Alquist and Kilian, 2010). Given these results, more research is needed to validate vehicle choice models before they can be reliably used in a regulatory context. Researchers modeling consumer vehicle choice should consider further validation as part of their research. In the meantime, using past market shares may be a good approximation for vehicle markets.

During the Great Recession, automakers failed to pass along costs associated with improved fuel economy. Between 2008 and 2010, automakers increased fuel economy by 3% but prices remained about the same. Between 2008 and 2016, fuel economy increased by 20% and prices changed by 5.5%. Changes in vehicle price appear uncorrelated with changes in fuel economy. However, given prediction bias in the nested logit model, price changes appear correlated with changes in unobserved vehicle quality, changes in consumer preferences, or brand-specific changes in market size. We fail to find evidence of automakers anticipating compliance via sales mixing. Instead, automakers may comply with greenhouse gas emissions standards through product design by adding fuel-saving technologies, consistent with the Porter Hypothesis (Porter and Van der Linde, 1995). A technology response is consistent with the European experience (Reynert, 2014), with US regulatory agencies' models (US EPA and DOT, 2010, 2012), and with recent work evaluating the product design response to US and European standards since 2003 (Klier and Linn, 2016). Automaker technology adoption appears to maintain market share, consistent with predictions in oligopoly settings (Sutton, 2006).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Albuquerque P, Bronnenberg BJ, 2012 Measuring the impact of negative demand shocks on car dealer networks. *Market. Sci.* 31 (1), 4–23.
- Alquist Ron, Kilian Lutz, 2010 What do we learn from the price of crude oil futures? *J. Appl. Econom.* 25 (4), 539–573.
- Ambec S, Cohen MA, Elgie S, Lanoie P, 2013 The Porter hypothesis at 20: can environmental regulation enhance innovation and competitiveness? *Rev. Environ. Econ. Pol.* 7 (1), 2–22.
- Amromin G, McGranahan L, 2015 The Great recession and credit trends across income groups. *Am. Econ. Rev.* 105 (5), 147–153.
- Anderson ST, Kellogg R, Langer A, Sallee JM, 2015 The intergenerational transmission of automobile brand preferences. *J. Ind. Econ.* 63 (4), 763–793.
- Athey Susan, 2017 “The Impact of Machine Learning on Economics” *Economics of Artificial Intelligence*. University of Chicago Press.
- Austin D, Dinan T, 2005 Clearing the air: the costs and consequences of higher CAFE standards and increased gasoline taxes. *J. Environ. Econ. Manag.* 50, 562–582.
- Bayati M, Bhaskar S, Montanari A, 2018 Statistical analysis of a low cost method for multiple disease prediction. *Stat. Methods Med. Res.* 27 (8), 2312–2328. [PubMed: 27932665]
- Benmelech E, Meisenzahl R, Ramcharan R, 2017 The real effects of liquidity during the financial crisis: evidence from automobiles. *Q. J. Econ.* 132 (1), 317–365.
- Bento AM, Goulder L, Jacobsen M, von Haefen R, 2009 Distributional and efficiency impacts of increased US gasoline taxes. *Am. Econ. Rev.* 99 (3), 667–699.
- Berry S, Levinsohn J, Pakes A, 1995 Automobile prices in market equilibrium. *Econometrica* 63 (4), 841–940.
- Bronnenberg B, Dhar S, Dubé J, 2009 Brand history, geography, and the persistence of brand shares. *J. Polit. Econ.* 117 (1), 87–115.
- Bronnenberg B, Dube J, Gentzkow M, 2012 The evolution of brand preferences: evidence from consumer migration. *Am. Econ. Rev.* 102 (6), 2472–2508.
- Brownstone D, Bunch D, Golob T, Ren W, 1996 A transactions choice model for forecasting demand for alternative-fuel vehicles. *Res. Transport. Econ.* 4, 87–129.
- Bunch D, Greene D, Lipman T, Shaheen S, 2011 Potential Design, Implementation, and Benefits of a Feebate Program for New Passenger Vehicles in California. State of California Air Resources Board and the California Environmental Protection Agency, Sacramento, California available at: http://76.12.4.249/artman2/uploads/1/Feebate_Program_for_New_Passenger_Vehicles_in_California.pdf.
- Busse M, Knittel C, Zettelmeyer F, 2013 Are consumers myopic? Evidence from new and used car purchases. *Am. Econ. Rev.* 103 (1), 220–256.
- Chen J, Esteban S, Shum M, 2008 Demand and supply estimation biases due to omission of durability. *J. Econom.* 147 (2), 247–257.
- Ciuffo B, Fontaras G, 2017 Models and scientific tools for regulatory purposes: the case of CO₂ emissions from light duty vehicles in Europe. *Energy Policy* 109, 76–81.
- Energy Information Administration (US) (Ed.), 2012 Annual Energy Outlook 2012: With Projections to 2035. Government Printing Office.
- Ericson R, Pakes A, 1995 Markov-perfect industry dynamics: a framework for empirical work. *Rev. Econ. Stud.* 62 (1), 53–82.
- Gallagher KS, Muehlegger E, 2011 Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology. *J. Environ. Econ. Manag.* 61 (1), 1–15.
- Goldberg P, 1995 Product differentiation and oligopoly in international markets: the case of the U.S. Automobile industry. *Econometrica* 63 (4), 891–951.

- Goldberg P, 1998 The effects of the corporate average fuel efficiency standards in the U.S. *J. Ind. Econ.* 46 (1), 1–33.
- Goolsbee A, Krueger A, 2015 A retrospective look at rescuing and restructuring general Motors and chrysler. *J. Econ. Perspect.* 29 (2), 3–24.
- Greene D, 2010 How Consumers Value Fuel Economy: A Literature Review. Office of Transportation and Air Quality, U.S. Environmental Protection Agency EPA-420-R-10–008.
- Greene D, Hossain A, Hofmann J, Helfand G, Beach R, 2018 Consumer willingness to pay for vehicle attributes: what do we Know? *Transport. Res. Pol. Pract.* 118, 258–279.
- Greene D, Liu C, 2012 Consumer Vehicle Choice Model Documentation. U.S. Environmental Protection Agency EPA-420-B-12_052. <http://www.epa.gov/otaq/climate/documents/420b12052.pdf>.
- Greene D, Patterson P, Singh M, Li J, 2005 Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy. *Energy Policy* 33 (6), 757–776.
- Harrison D, Nichols A, Reddy B, 2008 “Evaluation of NHTSAs Benefit-Cost Analysis of 2011–2015 CAFE Standards. National Economic Research Associates, Boston, Massachusetts June 30. Available at: http://www.nera.com/67_5346.htm.
- Haaf C, Michalek J, Morrow W, Liu Y, 2014 Sensitivity of vehicle market share predictions to discrete choice model specification. *J. Mech. Des.* 136 121402–121402–9.
- Haaf C, Morrow W, Azevedo I, Feit E, Michalek J, 2016 Forecasting light-duty vehicle demand using alternative-specific constants for endogeneity correction versus calibration. *Transport. Res. Part B* 84, 182–210.
- Jacobsen M, 2013 Evaluating U.S. Fuel economy standards in a model with producer and household heterogeneity. *Am. Econ. J. Econ. Policy* 5 (2), 148–187.
- Kleit AN., 2004 Impacts of long-range increases in the fuel economy (CAFE) standard. *Econ. Inq.* 42, 279–294.
- Klier T, Linn J, 2012 New-vehicle characteristics and the cost of the corporate average fuel economy standard. *Rand J. Econ.* 43 (1), 186–213.
- Klier T, Linn J, 2016 The effect of vehicle fuel economy standards on technology adoption. *J. Public Econ.* 133, 41–63.
- Knittel CR, 2011 Automobiles on steroids: product attribute trade-offs and technological progress in the automobile sector. *Am. Econ. Rev.* 2012 (101), 3368–3399.
- Knittel CR, Metaxoglou K, 2014 Estimation of random-coefficient demand models: two empiricists’ perspective. *Rev. Econ. Stat.* 96 (1), 34–59.
- Kullback S, Leibler RA, 1951 On information and sufficiency. *Ann. Math. Stat.* 22 (1), 79–86.
- Langer Ashley, Miller Nathan H., 2013 Automakers’ short-run responses to changing gasoline prices. *Rev. Econ. Stat.* 95 (4), 1198–1211.
- Larrick RP, Soll JB, 2008 The Mpg illusion. *Science* 320 (5883), 1593–1594. [PubMed: 18566271]
- Leard Benjamin, Joshua Linn, Virginia McConnell, 9 2017 Fuel prices, new vehicle fuel economy, and implications for attribute-based standards. *J. Associat. Environ. Resour. Econ.* 4 (3), 659–700.
- Li S, Linn J, Spiller E, 2013 Evaluating “Cash-for-Clunkers”: program effects on auto sales and the environment. *J. Environ. Econ. Manag.* 65 (2), 175–193.
- Mannering F, Winston C, Griliches Z, Schmalensee R, 1991 “Brand loyalty and the decline of American automobile firms” *Brookings Papers on Economic Activity. Microeconomics* 67–114.
- Mannering F, Winston C, Starkey W, 2002 An exploratory analysis of automobile leasing by US households. *J. Urban Econ.* 52 (1), 154–176.
- Mian A, Sufi A, 2012 The effects of fiscal stimulus: evidence from the 2009 cash for Clunkers program. *Q. J. Econ.* 127 (3), 1107–1142.
- Mian A, Sufi A, 2014 “What explains the 2007–2009 drop in employment? *Econometrica* 82 (6), 2197–2223.
- Mian A, Rao K, Sufi A, 2013 Household balance sheets, consumption, and the economic slump. *Q. J. Econ.* 128 (4), 1687–1726.

- NERA Economic Consulting, 2009 Evaluation of NHTSA's Benefit-Cost Analysis of 2011–2015 CAFE Standards. <http://www.nera.com/publications/archive/2008/evaluation-of-nhtsas-benefitcost-analysis-of-20112015-cafe-st.html>.
- Petrin A, 2002 Quantifying the benefits of new products: the case of the minivan. *J. Polit. Econ.* 110 (4), 705–729.
- Pindyck RS, 2013 Climate change policy: what do the models tell us? *J. Econ. Lit.* 51 (3), 860–872.
- Porter ME, Van der Linde C, 1995 Toward a new conception of the environment- competitiveness relationship. *J. Econ. Perspect.* 9 (4), 97–118.
- Reynaert M, 2014 Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market. KU Leuven Center for Economic Studies Discussion Paper Series DPS14.31.
- Rosen S, 1974 Hedonic prices and implicit markets: product differentiation in pure competition. *J. Polit. Econ.* 82 (1), 34–55.
- Sakti Apurba, Azevedo Inês ML., Fuchs Erica RH., Michalek Jeremy J., Gallagher Kevin G., Whitacre Jay F., 2017 Consistency and robustness of forecasting for emerging technologies: the case of Li-ion batteries for electric vehicles. *Energy Policy* 106, 415–426.
- Saltelli Andrea, Funtowicz Silvio, 2014 When all models are wrong. *Issues Sci. Technol.* 30 (2), 79–85.
- Sen Burak, Noori Mehdi, Tatari Omer, Oct. 2017 Will corporate average fuel economy (CAFE) standard help? Modeling CAFE's impact on market share of electric vehicles. *Energy Policy* 109, 279–287.
- Sutton John, 2007 Market share dynamics and the “persistence of leadership” debate. *Am. Econ. Rev.* 97 (1), 222–241.
- Train K, 2009 Discrete Choice Methods with Simulation. Cambridge University Press <http://elsa.berkeley.edu/books/choice2.html>.
- Train K, Winston C, 2007 Vehicle choice behavior and the declining market share of U.S. Automakers. *Int. Econ. Rev.* 48, 1469–1496.
- U.S. Energy Information Administration, 2008 Annual Energy Outlook 2008. DOE/EIA-0383. <http://www.eia.gov/oiaf/aeo/pdf/0383%282008%29.pdf>.
- U.S. Environmental Protection Agency, 2010 Final Regulatory Impact Analysis: Rulemaking to Establish Light-Duty Vehicle Greenhouse Gas Emissions and Corporate Average Fuel Economy Standards. US Environmental Protection Agency EPA-420-R-009.
- U.S. Environmental Protection Agency, 2016 Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 through 2016 Executive Summary. US Environmental Protection Agency, Washington, DC EPA-420-R-16-010.
- U.S. EPA, 2018 2016–2035 Production Summary Baseline 10.15.2018 Docket.” Supporting and Related Material for the Safer Affordable Fuel-Efficient (SAFE) Vehicles Proposed Rulemaking. <https://www.regulations.gov/document?D=EPA-HQ-0AR-2018-0283-3264>, Accessed date: 12 April 2018.
- U.S. Environmental Protection Agency and Department of Transportation, 2010 Light- duty vehicle greenhouse gas emissions standards and corporate average fuel economy standards; final rule. *Fed. Regist.* 75 (88), 25324–25728.
- U.S. Environmental Protection Agency and Department of Transportation, 2012 Joint Technical Support Document: Final Rulemaking for 2017–25 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. EPA-420-R-12-901, (Chapter 1). <http://www.epa.gov/otaq/climate/documents/420r12901.pdf>.
- Whitefoot KS, Fowlie ML, Skerlos SJ, 2017 Compliance by design: influence of acceleration trade-offs on CO2 emissions and costs of fuel economy and greenhouse gas regulations. *Environ. Sci. Technol.* 51 (18), 10307–10315. [PubMed: 28825797]
- Whitefoot K, Skerlos S, 2012 Design incentives to increase vehicle size created from the U.S. Footprint-based fuel economy standards. *Energy Policy* 41 402–411.0.
- Xie Fei, Lin Zhenhong, 2017 Market-driven automotive industry compliance with fuel economy and greenhouse gas standards: analysis based on consumer choice. *Energy Policy* 108, 299–311.

U.S. Environmental Protection Agency, 2014 Operational Performance Evaluation of Air Quality Model Simulations. <http://www.epa.gov/AMD/Research/Air/operationalEval.html>, Accessed date: 2 September 2014.

Frequency of Percentage Change

Fig. 1.
1 Annual Variation in price, fuel economy, and market share.

Fig. 2.
2Nested logit structure of consumer choice model.

Fig. 3.
3Cumulative error distribution (2010 results).

Fig. 4.
4Price and market share change (2010 Results).

Fig. 5. 5Percentage Point Change in Share for Nests and Automakers. Note: Automakers restricted to those with at least 4% market share.

Table 1

Summary statistics for observed changes in fleet.

	2008			2010			2016		
	Observed	Aggregated (2010)	Aggregated (2016)	Observed	Aggregated	Observed	Aggregated	Observed	Aggregated
Number unique vehicles	1,302	524	179	1,171	524	1,229	179		
Total sales	13,851,770	12,976,769	12,741,662	11,190,181	10,199,188	16,262,536	15,094,485		
Percentage sales matched	-	93.7%	92.0%	-	91.1%	-	92.8%		
Weighted average price	27,873	27,702	27,868	26,767	26,624	29,407	29,404		
Minimum price	11,783	11,783	13,455	9,970	11,923	11,504	14,872		
Maximum price	1,734,000	1,734,000	385,279	1,700,000	1,700,000	495,323	467,813		
Weighted avg fuel economy	26.2	27.4	25.61	28.4	28.3	31.5	28.3		
Minimum fuel economy	12.0	12.0	13.1	12.0	12.0	15.0	15.7		
Maximum fuel economy	65.8	65.8	65.8	70.8	70.8	218.1	70.1		

Table 2

Vehicle class definition and elasticities in the consumer vehicle choice model.

Model Class	Elasticity	Parent Nest	EPA Fleet
Prestige Two-Seaters	-3.8	Two-Seater	Car
Prestige Subcompact Cars	-3.5 *	Prestige Car	Car
Prestige Compact Cars and Small Station	-3.5	Prestige Car	Car
Prestige Midsize Cars and Station Wagons	-3.6 *	Prestige Car	Car
Prestige Large Cars	-3.5	Prestige Car	Car
Two-Seater	-3.5	Two-Seater	Car
Subcompact Cars	-5	Standard Car	Car
Compact Cars and Small Station Wagons	-5	Standard Car	Car
Midsize Cars and Station Wagons	-5	Standard Car	Car
Large Cars	-5	Standard Car	Car
Prestige SUVs	-3.7	Prestige SUV	Truck
Small SUVs	-4.9	Standard SUV	Truck
Midsize SUVs	-5.1	Standard SUV	Truck
Large SUVs	-5.1	Standard SUV	Truck
Minivans	-4.9	Minivan	Truck
Cargo/Large Passenger Vans	-5.1	Cargo Van	Truck
Small Pickup Trucks	-5.1	Cargo Pickup	Truck
Standard Pickup Trucks	-5.1	Cargo Pickup	Truck
Ultra Prestige Vehicles	-3.9	Ultra Prestige	Truck

* Notes: Parameter values differ for 2016 predictions due to nest elasticity requirement. Prestige Subcompact Cars elasticity was -4.1 and Prestige Midsize Cars was -3.6. These were the closest values to the original set that were feasible.

(1) Prestige and non-prestige classes are defined by vehicle price: the prestige are vehicles whose prices are higher than or equal to unweighted average price in the corresponding EPA class, and vice versa for non-prestige vehicles, e.g., Prestige Two-Seater class is the set of relatively expensive vehicle configurations in EPA class of two seaters with prices higher than or equal to the unweighted average price of EPA two seaters.

(2) Non-prestige SUVs are divided into small, midsize and large SUVs by vehicle's footprint (small: footprint < 43; midsize: 43 ≤ footprint < 46; large: footprint ≥ 46).

(3) Ultra Prestige class is defined as the set of vehicles whose prices are higher than or equal to \$75,000.

Table 3

Default elasticities, nests 1–3.

Level	Name	Elasticity	Parent Category
Choice Among 19 Vehicle Classes within Vehicle Type			
3	Two-Seater	−1.3	Passenger
3	Prestige Car	−2.2	Passenger
3	Standard Car	−3	Passenger
3	Prestige SUV		Passenger
3	Standard SUV	−2.7	Passenger
3	Minivan		Passenger
3	Cargo Van		Cargo
3	Cargo Pickup	−2	Cargo
3	Ultra Prestige		Ultra Prestige
Choice of Vehicle Type within Passenger or Cargo Categories			
2	Passenger	−1.1	Buy
2	Cargo	−0.7	Buy
2	Ultra Prestige		Buy
Choice of Passenger, Cargo or Ultra Prestige Vehicle			
1	Buy	−0.7	Root
1	No Buy		Root

Table 4

Goodness-of-prediction measures.

	2010		2016	
	NML	Simple	NML	Simple
Kullback-Leibler Divergence	0.351	0.212	0.659	0.271
Mean Squared Error	9.820E-06	5.330E-06	1.330E-04	2.510E-05
Average Share Error	0.001127	0.000842	0.004599	0.002967
Right Direction	63%		72%	