RESOURCES for the FUTURE

Estimating Preference Heterogeneity in Discrete Choice Models of Product Differentiation

Benjamin Leard

Working Paper 19-01
January 2019;
Revised April 2019

Estimating Preference Heterogeneity in Discrete Choice Models of Product Differentiation

Benjamin Leard*

April 2019

Abstract

Modeling preference heterogeneity in discrete choice models of product differentiation remains computationally challenging. I derive a new method for estimating preference heterogeneity in these models. A key advantage of the method is its simplicity: preference heterogeneity parameters are estimated with a closed-form expression or with a linear regression. I apply the method to estimate parameters of new vehicle demand and to simulate the effects of new vehicle fuel economy standards. The simulation results suggest that a marginal tightening of the standards has a modest impact on total new vehicle sales.

Keywords: discrete choice models, preference heterogeneity, microdata, substitution patterns
JEL codes: C35, C51, C54, L91

1 Introduction

Ever since their development by McFadden (1974), discrete choice models have been applied to analyze many relevant markets and public policies. Recent advances of these models, including by Berry (1994) and Berry et al. (1995), have improved their identification and estimation in two key ways: the estimation of unbiased estimates of average preference parameters and preference heterogeneity. These improvements have allowed researchers and policy makers to answer questions in the industrial organization and policy analysis literature, including the effect of market imperfections (such as market power) on market outcomes (Nevo 2001), how markets are affected by mergers (Thomadsen 2005), how the entry of new products affects producer and consumer welfare (Petrin 2002), and the social welfare effects of regulations (Berry et al. 1999).

*Fellow at Resources for the Future (RFF), leard@rff.org. I thank Maureen Cropper, Joshua Linn, and Christy Zhou for valuable comments on this paper. I am grateful to the Sloan Foundation for supporting the research.

Given their flexible ability to model consumer and producer preferences, these models continue to gain influence in public policy design. Analysis of policies such as federal gasoline taxes (Bento et al. 2009; Grigolon et al. 2018), federal fuel economy standards (Klier and Linn 2012; Jacobsen 2013; Reynaert 2017; Whitefoot et al. 2017; Leard et al. 2019), and subsidies for hybrid and electric vehicle purchases and infrastructure (Beresteanu and Li 2011; Springel 2017; Li 2018; Xing et al. 2019) depend on plausible identification and unbiased estimation of decision maker preferences in these models. Unfortunately, many of these models come with drawbacks. One drawback is that they are, in general, computationally difficult to code and estimate. Many of the models are formed as mixed logit models with product-specific fixed effects, which allow for rich unobserved heterogeneity and are theoretically able to capture any substitution pattern (McFadden and Train 2000). While mixed logit models represent the most flexible form in the class of discrete choice models, their computational complexity renders them unusable for some researchers and policy analysts. This complexity may be a reason why some government agencies do not use them for analyzing policies. For example, the National Highway Travel Safety Administration (NHTSA) and the Environmental Protection Agency (EPA) do not use discrete choice models in their cost-benefit analyses of federal fuel economy and greenhouse gas standards for light-duty vehicles, even though they have publicly stated their interest in using these models.[1] My discussions with representatives from NHTSA confirm that they have attempted to work with these models, but they are so complex as to render them incompatible with their current approaches to analyzing the standards. Furthermore, recent studies have highlighted computational difficulties with traditional mixed logit models following Berry et al. (1995), henceforth referred to as BLP models, resulting in erroneous conclusions. Knittel and Metaxoglou (2014) point out that the generalized method of moments (GMM) objective function for BLP models is not necessarily concave, and find that various routines for optimizing the GMM objective function yield wildly different preference parameter estimates. They find that the resulting difference in model predictions based on the parameter estimates can be significant. This makes it difficult to conclude whether any set of results derived from these models is driven by the underlying data or is biased due to the sensitivity of the computational routine. This finding makes results for policy analysis less believable and therefore less useful.

Other "micro" versions of the BLP models that incorporate household-level data, such as Berry et al. (2004), still do not necessarily have a concave objective function, extending the concerns from

[1] As stated in their regulatory impact analysis of the 2017-2025 standards, "NHTSA also considered developing and using a vehicle choice model to estimate the extent to which sales volumes would shift in response to changes in vehicle prices and fuel economy levels. As discussed Chapter V, the agency is currently sponsoring research directed toward developing such a model. However, that effort has not yet yielded a choice model ready for integration into NHTSA's analysis. If that effort is successful in the future, the agency will consider integrating the model into the CAFE modeling system and using the integrated system for future analysis of potential CAFE standards. If the agency does so, we expect that the vehicle choice model would impact estimated fleet composition not just under new CAFE standards, but also under baseline CAFE standards" (NHTSA 2012).

Knittel and Metaxoglou (2014). Furthermore, versions of these models that use maximum simulated likelihood generally take significant time to estimate due to the required number of computations per household.[2] Depending on the number of household observations used for estimation, these models can take several hours or even days to estimate. This slows down the construction and implementation of these models. Given that government agencies often face extremely tight deadlines for producing cost-benefit analyses, this estimation time issue represents a barrier for adoption.

The computational difficulty encountered when estimating these models arises primarily because of the detailed representation of consumer heterogeneity. Traditional estimation of mixed logit models requires simulating choice probabilities that are based on distributions of preference parameters that vary randomly across consumers. This simulation increases the number of computations and time required for estimation by orders of magnitude relative to the closed-form logit model, and introduces the concerns discussed in Knittel and Metaxoglou (2014). A recently developed method by Fox et al. (2011) simplifies mixed logit estimation by representing consumer heterogeneity as discrete groups, as opposed to the more typical approach of modeling consumer heterogeneity as continuous. Their method does not involve simulation, and models based on their method can be estimated with constrained least squares, which guarantees a global optimum. However, this method requires a control function approach for handling product attribute endogeniety and is not compatible with the more standard product fixed effect approach as in Berry et al. (1995).

In contrast to these techniques, other studies have taken a simpler approach by using a nested logit model derived in Berry (1994) that relates market shares to preference parameters and product attributes in a linear equation.[3] This model is capable of modeling unobserved heterogeneity for mutually exclusive groups of products, such as product classes or types.[4] This model is elegant in that it can be estimated with linear estimation routines, including ordinary least squares and instrumental variables. Therefore, it avoids the drawbacks of mixed logit models while being able to produce plausible substitution patterns and account for unobserved product attributes in the identification of average preference parameters. While the model derived in Berry (1994) is limited in the form of preference heterogeneity that it can accommodate, it often can capture heterogeneity that is relevant for particular policies. Grigolon and Verboven (2014) compare nested logit models and mixed logit models based on their ability to predict changes in prices and market shares due to mergers. They find that the models produce similar results, and they conclude that nested logit models are ideal for accounting for discrete sources of market segmentation not captured by continuous product preference heterogeneity in mixed logit models. Klier and Linn (2012) and Leard et al. (2019) adopt versions of

[2] Recent examples of these models include Train and Winston (2007), Langer (2016), Whitefoot et al. (2017), and Xing et al. (2019).

[3] This is Equation (28) in Berry (1994).

[4] For example, distinct classes for new vehicles can include pickup trucks, SUVs, crossovers, and sedans.

the model derived in Berry (1994) to evaluate fuel economy standards for passenger vehicles. Klier and Linn (2012) estimate unobserved heterogeneity using an instrumental variables approach that uses data on vehicle engine programs and platforms. Leard et al. (2019) estimate observed heterogeneity using correlations between household demographics and vehicle attributes.

In the current paper, I develop a new method that provides a simple approach to identifying and estimating preference heterogeneity that can be used together with conventional methods for obtaining unbiased estimates of average preference parameters. I build on the derivation of the nested logit model developed in Berry (1994) by deriving a initial-stage estimation of the preference heterogeneity. The parameters in the initial stages are identified by incorporating increasingly common microdata. The microdata can be either in the form of second choice data or decision maker characteristics data. The second choice data are based on household survey questions that ask respondents which product they would have bought had the product they purchased been unavailable. The logic of the identification strategy follows Berry et al. (2004) for identifying unobserved preference heterogeneity. Berry et al. (2004) use the correlation in continuous attributes between observed choices and stated second choices to identify the standard deviations of random coefficients for continuous attributes. I use the same correlation approach for identifying unobserved preference heterogeneity for groups of products in a nested logit framework: my approach uses the correlation among the classes of first and second choice products. For example, a buyer of a pickup truck may have a strong unobserved preference for owning a pickup truck. This would be present in the second choice data if many pickup truck buyers stated that they would have purchased a different pickup truck had their purchased truck been unavailable.

Second choice data may not be available in some datasets. In this case, data on decision maker characteristics, such as household demographics, linked to product purchases can be used. These data are becoming more widely available in most marketing datasets. For example, household income of buyers is often recorded in addition to the product chosen by the household. The logic of the identification strategy follows the methods of Berry et al. (2004) in identifying observed preference heterogeneity. Correlation between household demographics and product attribute levels is used to identify the heterogeneous preference parameters. One example of a pattern that may be present in household-level datasets is that households with relatively high income are more likely to purchase expensive products.

In sharp contrast to the approach in Berry et al. (2004) and other mixed logit approaches, the method I derive is simple to estimate and does not involve GMM or the optimization of a likelihood function for estimation. Instead, estimating the preference heterogeneity parameters involves evaluating a closed-form expression that is a function of market share and microdata or estimating a fixed effects linear regression. The remaining "mean utility" parameters are then estimated in a final stage and are consistent with the estimated preference heterogeneity parameters.[5]

This method in this paper is complementary to the approaches presented in Berry et al. (2004) and Fox et al. (2011). The strength of the GMM estimation in Berry et al. (2004) is that it can accommodate virtually any form of preference heterogeneity, both observed and unobserved, and for both discrete and continuous product attributes. Its weaknesses, however, are its computational complexity and its potential for estimation instability. In settings where only certain simpler forms of preference heterogeneity are relevant, the method that I present serves as an alternative that does not share the drawbacks of the Berry et al. (2004) method. The Fox et al. (2011) method also avoids the computational challenges in Berry et al. (2004), but requires specifying a discrete grid of pre-defined preference parameter values and adopting a control function approach for handling product attribute endogeneity. When a researcher has reasons to avoid these modeling requirements, the approach in Berry et al. (2004) or the current paper may be preferable. Together, these distinct approaches provide researchers with a broader toolkit for estimating discrete choice models of product differentiation.

To illustrate the value of the approach for policy evaluation, I use the method to estimate a model of light-duty vehicle demand and simulate the effect of tightening fuel economy standards on new vehicle sales. I find that marginally tightening the standards results in a small reduction in new vehicle sales. The policy relevance of the effect of fuel economy standards on new vehicle sales dates back to the enactment of federal fuel economy standards in the late 1970s and economic analyses of the standards shortly thereafter. Gruenspecht (1982) finds that fuel economy standards for new vehicles, because they only apply to new vehicles and not used vehicles, have unintended effects of lowering new vehicle sales and slowing used vehicle scrappage. This scrappage effect can undo the intended effects of the standards, since used vehicles tend to have lower fuel economy and greater oil consumption than new vehicles. The magnitude of this effect depends on several key factors, including how new vehicle buyers substitute between new and used vehicles and the relationship between used vehicle prices and scrappage (Bento et al. 2018; Jacobsen and van Benthem 2015). I simulate the substitution between new and used vehicles as a result of tightening fuel economy standards by applying the method to estimate unobserved preference heterogeneity for new vehicles. The policy relevance of this issue has only grown since the introduction of the standards, especially during the last last ten years, during which the standards have been revised several times. In 2008, the Obama administration passed legislation to double the stringency of the standards by 2025, and in 2018, the Trump administration proposed legislation to roll back these standards to remain flat at

[5]The parameters are consistent because the heterogeneity parameters are estimated conditional on the values of the mean utility parameters, and the dependent variable for the mean utility estimation equation is defined as a function of the heterogeneity parameters. This reasoning follows prior two-stage BLP estimation approaches, including Berry et al. (2004) and Train and Winston (2007).

2020 levels. A recent analysis of the rollback proposal completed by the EPA and NHTSA finds that the rollback will shrink the entire vehicle fleet and miles traveled, preventing a significant number of vehicle fatalities (EPA 2018). Unfortunately, this analysis has been shown to have severe modeling flaws and limitations. Bento et al. (2018) find that a key flaw in the analysis is that the agencies use an incomplete reduced-form statistical model to simulate the effects of the rollback on new and used vehicle fleet size. This model produces predictions that are inconsistent with basic economic theory, likely because the model's parameters are not estimated based on structural assumptions for consumer or producer decision-making. Bento et al. (2018) suggest following an ideal protocol for analyzing changes in fuel economy standards. The protocol involves using a vehicle choice model that is underpinned by basic economic principles. The method presented in the current paper is a starting point for developing such a model.

The method I build in this paper can be applied to many other settings, where different forms of heterogeneity matter more than others. I focus the application of the method on estimating substitution patterns between inside alternatives, i.e., new vehicles, and the outside option, i.e., used vehicles, for two reasons. First, the identification of this substitution has been neglected in prior studies, which have been focused on substitution patterns among inside alternatives. Second, as I show empirically, substitution to the outside option can have substantial policy implications. The outside option generally measures market size. Policy outcomes dependent on market size therefore depend on the degree of substitution to the outside option.

The remaining sections of the paper are organized in the following manner. In Section 2, I derive the estimation method building on the model presented in Berry (1994). In this section, I present alternative forms of the method that use different types of microdata. I then apply the method to the setting of the US light-duty vehicles market and simulate the effect of tightening fuel economy standards on new vehicle sales in Section 3. In Section 4, I discuss possible extensions and alternative applications of the method, and I make concluding remarks in Section 5.

2 Model Development

In this section, I present three different variations of the method for identifying preference heterogeneity. Each variation differs in the types of heterogeneity that are identified and estimated. All of them share the common methodology of being estimated in multiple simple stages. The first variation adopts a nested logit form based on Berry (1994). The second adopts a logit form that incorporates observed heterogeneity, and the third combines the first and second variations to incorporate both observed and unobserved heterogeneity.

2.1 Model Setup

I derive the first variation of the estimation method that requires second choice data by first deriving the nested logit model beginning with utility maximization. I follow the presentation of the nested logit model in Berry (1994). Assume there are J alternatives indexed as $j = 0, 1, ..., J$, where $\mathbb{J}$ denotes the choice set, and where $j = 0$ denotes the outside good. Alternatives are grouped into $G+1$ groups, indexed by $g = 0, 1, 2, ..., G$. The outside option $j = 0$ is assumed to be the only alternative in group 0. Decision maker i obtains utility u_{ij} when choosing alternative j in group g, where utility is

$$u_{ij} = \delta_j + \xi_{ig} + (1 - \sigma)\epsilon_{ij}. \tag{1}$$

The term δ_j in Equation (1) represents average utility for alternative j, and can be decomposed into two parts: $\delta_j = x_j\beta + \varepsilon_j$. The vector x_j represents values for the attributes of alternative j, and the vector β denotes marginal utilities for the attributes. The term ε_j is an idiosyncratic error term. The term ξ_{ig} is an unobserved component of utility that is common to all alternatives in group g. The last term in Equation (1) is an idiosyncratic error term that is scaled by $(1-\sigma)$. The term σ is often referred to as a nesting parameter and is a measure of within-group correlation with bounds $0 \leq \sigma < 1$. As σ approaches one, the error component of Equation (1) approaches zero and within-group correlation of utility approaches one. Larger values of σ therefore increase the degree of substitution between alternatives that share a group and reduce the degree of substitution between alternatives that do not share a group. Decision makers are assumed to select a single alternative that maximizes their utility.

Based on this setup, Berry (1994) shows that the predicted market share for alternative j equals

$$s_j = \frac{e^{\delta_j/(1-\sigma)}}{D_{g(j)}^{\sigma} \sum_g D_g^{(1-\sigma)}}, \tag{2}$$

where

$$D_{g(j)} = \sum_{k \in \mathbb{J}_{g(j)}} e^{\delta_k/(1-\sigma)} \tag{3}$$

and

$$\sum_g D_g^{(1-\sigma)} = \sum_g \left[\sum_{k \in \mathbb{J}_g} e^{\delta_k/(1-\sigma)} \right]^{(1-\sigma)}. \tag{4}$$

The outside option is assumed to have a normalized mean utility equal to zero, $\delta_0 = 0$, which simplifies its predicted market share:

$$s_0 = \frac{1}{\sum_g D_g^{(1-\sigma)}}. \tag{5}$$

Berry (1994) derives a linear equation relating observed market shares, mean utilities δ_j, and the nesting parameter σ:

$$\ln(s_j) - \ln(s_0) = \delta_j + \sigma \ln(s_{j|g}), \tag{6}$$

where the second term on the right-hand side of Equation (6) within the natural log operator, $s_{j|g}$, is the within-group share:

$$s_{j|g} = \frac{e^{\delta_j/(1-\sigma)}}{D_{g(j)}} = \frac{e^{\delta_j/(1-\sigma)}}{\sum_{k \in \mathbb{J}_{g(j)}} e^{\delta_k/(1-\sigma)}}. \tag{7}$$

Equation (6) can be estimated with linear instrumental variables methods. Since the within-group share $s_{j|g}$ is endogenous, an instrumental variable is required to identify σ. I next show that this parameter can be identified without an instrumental variable for $s_{j|g}$. The method I propose uses second choice data. This alternative method is valuable for three reasons. First, finding a valid instrumental variable for $s_{j|g}$ can be challenging. Second, a valid instrumental variable for $s_{j|g}$ may not produce precise estimates for σ. Third, second choice data provide a source of information that directly measures correlation of unobserved utility among alternatives, which is exactly what σ measures. Therefore, unlike using most candidate instrumental variables, using second choice data serves as an appropriate source of identification.

2.2 Estimating Unobserved Heterogeneity

I now develop a method for estimating unobserved heterogeneity within groups of productsthat uses second choice data. Second choice data provide information on the stated frequency of alternative choices when another alternative is unavailable. Survey questions typically take the following form to elicit the second choice: "If the alternative you chose did not exist, what alternative would you have chosen?" These choices can be aggregated to frequencies and combined with market share data to compute market shares conditional on the removal of an alternative. I convert second choice frequencies to market shares with alternatives removed from the choice set, since these market shares are easily computed with a nested logit model, and this formulation retains all useful information for identifying consumer heterogeneity parameters.[6] Suppose we have data on market sales, denoted by q, shares, and the percentage of consumers choosing alternative j stating they would have chosen alternative k had alternative j not been available, denoted by $s_{k,j}$. Then sales of alternative k with alternative j removed is equal to sales of k (with alternative j present) plus the product of the sales of j and the percentage of alternative j consumers stating they would have chosen k had alternative j not been available:

$$q_{k|j\notin\mathbb{J}} = q_k + s_{k,j}q_j. \tag{8}$$

Dividing both sides by total market size converts the sales terms to market shares:

$$s_{k|j\notin\mathbb{J}} = s_k + s_{k,j}s_j. \tag{9}$$

Dividing both sides of the sales equation by total sales within alternative k's group (assuming that alternative j and k share the same group) yields a within-group sales share conditional on the removal of alternative j from the choice set:

$$s_{k|g,j\notin\mathbb{J}} = s_{k|g} + s_{k,j}s_{j|g}. \tag{10}$$

Equations (9) and (10) can be used to match share predictions from the nested logit model with observed market share data and second choice frequencies. The left-hand side of Equation (9) is the market share prediction from the nested logit model with alternative j removed from the choice set:

$$s_{k|j\notin\mathbb{J}} = \frac{e^{\delta_k/(1-\sigma)}}{D^{\sigma}_{g(k|j\notin\mathbb{J})}\sum_g D^{(1-\sigma)}_{g|j\notin\mathbb{J}}}. \tag{11}$$

The right-hand side of Equation (9) is a combination of micro and macrodata, including aggregate market shares for alternatives j and k as well as the second choice frequency. This matching can be interpreted as a moment condition. More of these conditions can be formed than just the one for alternative k. For example, there are J different versions of Equation (9), where alternative k is replaced with another alternative besides alternative j.

[6] An alternative approach is to directly model the probability of choosing an alternative conditional on the choice of a different alternative. This is typically done in mixed logit models by modeling an ordered logit or exploded logit for the sequence of choice probabilities, where the identification of the unobserved heterogeneity parameters is from the correlation of the attributes of the first and second choices (Berry et al. 2004; Train and Winston 2007).

Equation (11) is a function of the nesting parameter σ and the mean utilities $\delta_1, \delta_2, ..., \delta_J$. The mean utilities can be substituted out by using Equation (6). Solving the alternative k version of Equation (6) for the mean utility yields

$$\delta_k = \ln(s_k) - \ln(s_0) - \sigma \ln(s_{k|g}). \tag{12}$$

The right-hand side of Equation (12) is a function of market share data (s_0, s_k, and $s_{k|g}$) and the nesting parameter. Substituting Equation (12) into Equation (11) yields an expression that is a function of the data and the nesting parameter σ only:

$$s_{k|j\notin\mathbb{J}} = \frac{e^{[\ln(s_k)-\ln(s_0)-\sigma\ln(s_{k|g,})]/(1-\sigma)}}{\widetilde{D}^{\sigma}_{g(k|j\notin\mathbb{J})}\sum_g \widetilde{D}^{(1-\sigma)}_{g|j\notin\mathbb{J}}}, \tag{13}$$

where

$$\widetilde{D}_{g(k|j\notin\mathbb{J})} = \left[\sum_{m\in\mathbb{J}_{g(k|j\notin\mathbb{J})}} e^{[\ln(s_m)-\ln(s_0)-\sigma\ln(s_{m|g})]/(1-\sigma)}\right] \tag{14}$$

and

$$\sum_g \widetilde{D}^{(1-\sigma)}_{g|j\notin\mathbb{J}} = \sum_g \left[\sum_{m\in\mathbb{J}_{g|j\notin\mathbb{J}}} e^{[\ln(s_m)-\ln(s_0)-\sigma\ln(s_{m|g})]/(1-\sigma)}\right]^{(1-\sigma)}. \tag{15}$$

Equation (13) can be solved for a closed-form solution of σ, denoted as $\widehat{\sigma}$:

$$\widehat{\sigma} = \frac{1}{J(J-1)}\sum_{g=1}^{G} s_{j|g} \sum_{j\in\mathbb{J}_g}\sum_{k\in\mathbb{J}_g, k\neq j}\left[1 - \frac{\ln(s_0 + s_{0,j}s_j) - \ln(s_0)}{\ln(s_k + s_{k,j}s_j) - \ln(s_{k,})}\right]. \tag{16}$$

See the appendix for a derivation of this expression. This expression is a weighted average over all j, k pairs of alternatives that share the same group.[7] The weights are equal to within-group market shares, but can be assigned differently by the researcher to account for sampling design. The term within the brackets relates the change in market shares of the outside option and another alternative k sharing the group of the removed alternative j. The data $s_{k,j}$ measures the degree to which decision makers substitute to another alternative k when alternative j in the same group is removed from the choice set. If within-group utility is highly correlated, the removal of an alternative j should lead to a

[7]Because of the way $s_{k|g,j\notin\mathbb{J}}$ is computed, the pairs of alternatives should be limited to those that share the same group.

disproportionally large increase in the share of alternative k that is in the same group as alternative j, which would be reflected by a relatively large value for $s_{k,j}$. This increases the estimate for $\widehat{\sigma}$, since Equation (16) is monotonically increasing in the value of $s_{k,j}$.

To better understand how the data identify the nesting parameter, consider the simplified setting where there is no unobserved heterogeneity. Here the predicted market shares become logit. In this setting, the removal of alternative j from the choice set should not impact the market share of alternative k relative to the market share for the outside option due to the independence of irrelevant alternatives (IIA) property of the logit model (Train 2009). Therefore, the log difference of the outside market shares before and after the removal of any alternative j – which is the numerator of the second term in the brackets – should be identical to the log difference of the alternative k market shares for all k before and after the removal of any alternative j – which is the denominator of the second term in the brackets. This equality implies that the second term in the brackets of Equation (16) is equal to 1 for all j, k pairs, and therefore $\widehat{\sigma} = 0$.

A few caveats for Equation (16) are relevant. First, $\widehat{\sigma}$=1 if for all j, $s_{0,j} = 0$, regardless of the values for $s_{k,j}$. This implies that researchers should be careful to specify the outside option so that $s_{0,j} > 0$ for all or at least a large majority of j. Otherwise, $\widehat{\sigma}$ will be significantly biased toward one. Second, the estimate $\widehat{\sigma}$ is undefined for any $s_{k,j} = 0$.[8] In most contexts, microdata will have some alternatives that have no second choices of other alternatives. The researcher can avoid this issue in two ways. First, one can compute Equation (16) based on a subsample of alternatives where $s_{k,j} > 0$ for all alternatives k, j in the subsample. This strategy should introduce very little (if any) bias in the estimation of $\widehat{\sigma}$ if a small number of alternative pairs have $s_{k,j} = 0$. In cases where many $s_{k,j} = 0$, an alternative strategy is required. One alternative is to impute $s_{j,k}$ with the following function:

$$\widetilde{s}_{k,j} = s_{k|g,j\notin\mathbb{J}} \sum_{k\in\mathbb{J}_g,k\neq j} s_{k,j} = \frac{s_k}{s_g - s_j} s_{g,j}. \tag{17}$$

This function assigns imputed values for $s_{k,j}$ based on the within-group substitution of alternative j, denoted by $s_{g,j}$, and defined as the frequency of decision makers selecting another alternative in group g when alternative j is removed from the group (and choice set). This frequency represents an aggregated measure of within-group correlation of utility, and therefore serves as an intuitive approximation of the alternative specific frequency $s_{k,j}$. Imputed values are scaled by $s_{k|g,j\notin\mathbb{J}} = \frac{s_k}{s_g - s_j}$ so that alternatives with large shares are assigned a relatively large second choice market share. This imputation is useful for several reasons. First, the imputation dramatically reduces the data requirement of observing J within-group second choice shares, as opposed to all $J \times J_g - 1$ shares for

[8] This is because the denominator equals zero in this case.

each alternative pair. Second, the variation in the imputed shares maintains all relevant information for identifying within-group correlation of utility. Third, the imputation is consistent with predictions of the nested logit model, in that the scaling factor implies that within-group substitution is proportional to within-group shares. Fourth, even if second choice data are available for all $s_{k,j}$, if these are computed based on microdata, they likely include substantial sample variance, since most micro datasets are a small fraction of the entire population of decision makers.[9] Alternatively, group-level second choice shares $s_{g,j}$, although still containing sample variance, have less variance than alternative-level second choice shares. This motivates using the imputation strategy for any situation where sample variance may be large.

Substituting $s_{k,j} = \widetilde{s}_{k,j}$ into Equation (16) and simplifying yields

$$\widehat{\sigma} = \frac{1}{J(J-1)} \sum_{g=1}^{G} s_{j|g} \sum_{j \in \mathbb{J}_g} (J_{g(j)} - 1) \left[1 - \frac{\ln\left(1 + \frac{s_{0,j} s_j}{s_0}\right)}{\ln\left(1 + \frac{s_{g,j} s_j}{s_g - s_j}\right)} \right]. \tag{18}$$

This method outlined above yields $\hat{\sigma}$. With this estimate, an equation can be formed to estimate the mean utility preference parameters β by substituting $\hat{\sigma}$ for σ into Equation (6) and re-arranging:

$$\ln(s_j) - \ln(s_0) - \hat{\sigma} \ln(s_{j|g}) = x_j \beta + \varepsilon_j. \tag{19}$$

In summary, the method requires two steps:

1. Estimate the nesting parameter σ with Equation (16) if $s_{k,j} > 0$ for all alternatives

k, j that share the same group. Otherwise, estimate the nesting parameter with Equation

(18).

2. Estimate mean utility parameters based on Equation (19).

This method can be easily extended to estimate more detailed nested logit models that include more than one nesting parameter. For example, in the context of new vehicle demand, the nests can be defined as new or used vehicles, as well as separate nests for each new vehicle class. For example, the new vehicle classes can be defined by the decision to purchase a new car or a new light truck. The data requirement for incorporating multiple nesting parameters is to have second choice data defining the correlation for each nest. For example, the nesting parameter defined by the decision to purchase a new car or new light truck is identified by the frequencies of new car buyers having a different new car as their second choice and new light truck buyers having a different new light truck as their second

[9] As an example, the National Household Travel Survey (NHTS) surveys a little over 100,000 households in each wave, which is about 0.1% of the U.S. population.

choice. In other words, the correlation among the purchased vehicle class and the second choice vehicle class identifies this nesting parameter. The moment conditions are formed by first deriving a linear equation relating market shares, mean utilities, and the nesting parameters. For a nested logit model with a unique nesting parameter for each nest, this equation is

$$\ln(s_j) - \ln(s_0) = \delta_j + \sum_g I_{jg}\sigma_g \ln(s_{j|g}). \tag{20}$$

The appendix includes a derivation of this equation. The term I_{jg} is a dummy variable equal to one if $j \in \mathbb{J}_g$ and zero otherwise. Similar to the equation for the case of a single nesting parameter, this can easily be solved for mean utilities and substituted into predicted market share equations to derive a closed-form expression for the nesting parameters. Using second choice data to define the "moment conditions" yields

$$s_{k|j\notin\mathbb{J}} = \frac{e^{[\ln(s_k)-\ln(s_0)-\sigma_{g(k)}\ln(s_k)]/(1-\sigma_{g(k)})}}{\widetilde{D}^{\sigma_{g(k)}}_{g(k|j\notin\mathbb{J})}\sum_g \widetilde{D}^{(1-\sigma_g)}_{g|j\notin\mathbb{J}}}, \tag{21}$$

where

$$\widetilde{D}_{g(k|j\notin\mathbb{J})} = \sum_{m\in\mathbb{J}_{g(k|j\notin\mathbb{J})}} e^{[\ln(s_m)-\ln(s_0)-\sigma_m\ln(s_m)]/(1-\sigma_m)} \tag{22}$$

and

$$\sum_g \widetilde{D}^{(1-\sigma_g)}_{g|j\notin\mathbb{J}} = \sum_g \left[\sum_{m\in\mathbb{J}_{g|j\notin\mathbb{J}}} e^{[\ln(s_m)-\ln(s_0)-\sigma_m\ln(s_m)]/(1-\sigma_m)}\right]^{(1-\sigma_g)}. \tag{23}$$

The resulting closed-form expression for each $\sigma_{g(k)}$ is similar to the expression for $\widehat{\sigma}$ derived above:

$$\widehat{\sigma_{g(k)}} = \frac{1}{J_g - 1} s_{j|g} \sum_{j\in\mathbb{J}_g}\sum_{k\in\mathbb{J}_g, k\neq j}\left[1 - \frac{\ln(s_0 + s_{0,j}s_j) - \ln(s_0)}{\ln(s_k + s_{k,j}s_j) - \ln(s_{k,})}\right]. \tag{24}$$

Similar to the computation of $\widehat{\sigma}$,the pairs of alternatives j, k used to estimate each $\sigma_{g(k)}$ should be limited to those that share the same group. Unlike the computation of $\widehat{\sigma}$, only data for the pairs of alternatives that are in group g are used to compute each $\sigma_{g(k)}$.

2.3 Numerical Example for Estimating Unobserved Heterogeneity

I provide a simple numerical example of the first two steps to build intuition for the method. Suppose there are $J = 3$ inside alternatives and a single outside option, $j = 0$, for a total of four alternatives. The three inside alternatives all share the same group, $g = 1$, and the outside option is in its own group, $g = 0$. I identify the nesting parameter σ using the frequency with which the second choice is an alternative in group $g = 1$ when another alternative in group $g = 1$ is removed from the choice set.

The method requires data for market shares and second choice frequencies. The example data appear in Table 1. Panel (a) has data for a case of high correlation of utility for alternatives in the same group. The outside option is assumed to have a market share of 0.5 to facilitate a comparison of the shares and frequencies to predicted outcomes with a logit model without preference heterogeneity. I assign second choice frequencies when an inside alternative is removed to reflect strong within-group substitution. When alternative $j = 1$ is removed, 95 percent of consumers would choose a different inside alternative, and only 5 percent would choose the outside option. The second choice frequencies for the inside alternatives are proportional to their respective market shares, reflecting the within-group IIA property (Train 2009). The implied market shares with an inside alternative removed reflect highly correlated within-group utility. The removal of alternative $j = 1$ increases the market share of the outside option by one percentage point. In contrast, the market shares of the remaining inside alternatives increase by 6 to 15 percentage points, despite having a lower market share. For these data, the nesting parameter $\hat{\sigma}$ is estimated to be 0.90, indicating high within-group substitution.

Panel (b) in Table 1 has data for a case of low correlation of utility for alternatives in the same group. The market shares are the same as in the Panel (a) case. Second choice frequencies for the outside option are much higher in that case, reflecting similar substitution among the inside and outside alternatives. The outside option market share with an inside alternative removed is much higher in this case, and the proportional increase in market share is similar for all of the alternatives, resembling more of a logit-type model. For these data, the nesting parameter $\hat{\sigma}$ is estimated to be 0.27, indicating lower within-group substitution.

2.4 Estimating Observed Heterogeneity

In certain empirical settings, second choice data may not be available, preventing the application of the method described in the previous subsections. An alternative approach is to estimate observed heterogeneity that is based on reported characteristics of decision makers, such as consumer demographics. In this subsection, I derive a method for identifying and estimating observed preference heterogeneity. I adopt the same notation used in Section 2.1. Suppose decision makers are assigned to demographic groups denoted by d based on their observed characteristics, such as their age. Decision maker i who belongs to demographic group d obtains utility u_{ij} when choosing alternative j in group g, where utility is

$$u_{ij} = \delta_j + \beta_{dg} + \epsilon_{ij}. \tag{25}$$

The term β_{dg} represents demographic-specific utility for alternatives in group g. Therefore, decision makers in demographic d obtain utility $\delta_j + \beta_{dg}$ when choosing alternative j. The term δ_j maintains the same interpretation from Section 2.1 as being the average utility for alternative j. Decision makers are assumed to select a single alternative that maximizes their utility. Assuming that the idiosyncratic error component ϵ_{ij} is independently and identically distributed type 1 extreme value, in the appendix I show that this form of utility yields a simple linear equation relating market shares and parameters of the decision maker utility function:

$$\ln(s_{dj}) - \ln(s_{d0}) = \delta_j + \beta_{dg}. \tag{26}$$

This equation relates market shares for alternative j by demographic group to the average utility and demographic group-specific utility for alternative j. The demographic-specific utility can be decomposed into an average utility term and a demographic-specific error term, $\beta_{dg} = \bar{\beta}_{dg} + \mu_{dj}$, so that we can form an estimation equation for Equation (26):

$$\ln(s_{dj}) - \ln(s_{d0}) = \delta_j + \bar{\beta}_{dg} + \mu_{dj}. \tag{27}$$

Estimation of this equation requires data on market shares by demographic group. Aggregate market shares by demographic group may not always be available. Data that are typically more commonly available include aggregate market shares for each alternative, s_j, and shares of alternative group market shares by demographic group, s_{dg}. For example, one can aggregate Consumer Expenditure Survey (CEX) microdata to compute market shares of three vehicle groups–new cars, new light trucks, and used vehicles–for various demographic groups, such as those defined by income quintile or by urban or rural residence. But these data do not contain market shares of specific vehicles, such as a new Toyota Prius Plug-In. Market shares for each alternative by demographic group can be imputed without introducing noise in the estimation of the preference parameters in Equation (27). The procedure is as follows: for each alternative j, scale the market share s_j for each demographic group d so that the implied s_{dg} matches the data. To fix this idea, suppose there are three alternatives j=0,1,2, with market shares $s_0 = 0.5$, $s_1 = 0.3$, and $s_2 = 0.2$, where $j = 1, 2$ belong to an alternative group $g = 1$, and $j = 0$ belongs to its own outside good group $g = 0$. There are two demographic groups d=1,2, and we observe demographic by alternative group market shares for demographic group $d = 1$ as 0.4 for the outside good group and 0.6 for the inside good group. This demographic group has a larger market share for the inside good group relative to the entire population. To impute this demographic group's alternative-specific market shares, the inside good market shares for this demographic group are scaled up so that their sum is equal to 0.6. The scaling proportion is equal to the ratio of the inside group market share for the demographic group $d = 1$, which is 0.6, and the inside group aggregate market share, which is $s_1 + s_2 = 0.5$. Multiplying s_1 and s_2 by the scaling factor of $0.6/0.5 = 6/5$ yields imputed alternative market shares of $\widetilde{s}_{11} = 0.36$ and $\widetilde{s}_{12} = 0.24$ and an implied imputed market share of the outside good as $\widetilde{s}_{10} = 0.4$. This procedure is repeated for each demographic group.

The imputed values are used as data to construct the dependent variables in Equation (27). The parameters in this equation are estimated in two stages. In the first stage, the preference heterogeneity parameters $\bar{\beta}_{dg}$ are estimated, and alternative fixed effects δ_j are included. This stage is estimated with a fixed effects regression. The preference heterogeneity parameters enter as product interaction terms, where demographic groups are interacted with alternative groups. The preference heterogeneity parameters are identified from differences in market shares for each demographic group by alternative group pairing, controlling for aggregate mean utility common to each demographic group.

The first stage yields estimates for the fixed effects, denoted by $\hat{\delta}_j$. These values enter in the second stage as the dependent variable. Mean utility preference parameters are estimated in the second stage with the following equation:

$$\hat{\delta}_j = x_j \beta + \varepsilon_j. \tag{28}$$

The following series of steps summarizes the estimation method for obtaining observed preference heterogeneity:[10]

1. Form groups of decision makers and alternatives based on observed characteristics.

[10] This method is a simplified version of the estimation strategy adopted in Leard et al. (2019). Their method uses a similar two-stage estimation strategy to obtain observed heterogeneity and unbiased mean utility parameters of a vehicle demand model. Their method allows for observed heterogeneity for continuous attributes (such as price), which contrasts with the approach outlined above, which only permits observed heterogeneity for discrete, nonoverlapping groups of alternatives. Their method, however, has a significant data requirement for observing alternative market shares for multiple demographic groups, which may be unavailable in certain contexts. This generally requires a massive set of microdata for contexts with a large number of alternatives in the choice set, which is the case for vehicle demand. The data requirement for the method described here is much less demanding, only requiring information on market shares of aggregate alternative groups for different demographic groups. These data are generally common in many contexts and are available in several public datasets, such as the CEX. So, while the method described here is less flexible in modeling certain substitution patterns that are obtained from continuous heterogeneity, it is likely to be much more widely accessible.

2. If only s_j and s_{dg} are available, impute s_{dj} using the procedure above. Otherwise, skip to step 3.

3. Estimate first-stage fixed effects regression in Equation (27).

4. Estimate the second stage with Equation (28), using the estimated $\hat{\delta}'_j s$ from the first stage as the dependent variable.

2.5 A Combined Method for Estimating Observed and Unobserved Heterogeneity

In this section, I formulate a method for estimating observed and unobserved heterogeneity that combines the approaches from Sections 2.1 and 2.4. This is the most data-intensive method, requiring second choice data and decision maker characteristics linked to alternative choices. But it provides more flexibility in modeling heterogeneous preferences. I adopt the same notation used in Sections 2.1 and 2.4. Suppose decision makers are assigned to demographic groups denoted by d based on their observed characteristics. Decision maker i who belongs to demographic d obtains utility u_{ij} when choosing alternative j in group g, where utility is

$$u_{ij} = \delta_j + \beta_{dg} + \xi_{ig} + (1-\sigma)\epsilon_{ij}. \tag{29}$$

The interpretations of the utility function parameters are similar to those stated in prior sections. Assuming that the error term ϵ_{ij} is i.i.d. type 1 extreme value, in the appendix I derive a linear equation relating market shares, observed heterogeneity preference parameters (β_{dg}), and the unobserved heterogeneity preference parameter (σ):

$$\ln(s_{dj}) - \ln(s_{d0}) = \delta_j + \beta_{dg} + \sigma \ln(s_{dj|g}). \tag{30}$$

This equation combines the elements appearing in the estimation Equations (6) and (26). The estimation of the preference parameters proceeds in three stages. In the first stage, the unobserved heterogeneity parameter σ is estimated using moment conditions based on second choice data. These moment conditions are constructed in a similar fashion to the methodology described in Section 2.2. The first step to constructing the moment conditions is to form an expression of market shares with an alternative removed from the choice set. This can be done by demographic group:

$$s_{dk|j\notin\mathbb{J}} = s_{dk} + s_{dk,j}s_{dj}. \tag{31}$$

Note that this expression requires market shares by demographic group and alternative. If these data are unavailable, an imputation strategy outlined in Section 2.4 can be used to obtain imputed market shares.

Equation (32) can be used to match share predictions from the nested logit model with observed market share data and second choice frequencies. The left-hand side of Equation (32) is the market share prediction from the nested logit model with alternative j removed from the choice set

$$s_{dk|j\notin\mathbb{J}} = \frac{e^{(\delta_k+\beta_{dg(k)})/(1-\sigma)}}{D^{\sigma}_{dg(k|j\notin\mathbb{J})}\sum_g D^{(1-\sigma)}_{dg|j\notin\mathbb{J}}}, \tag{32}$$

where

$$D_{dg(k|j\notin\mathbb{J})} = \sum_{m\in\mathbb{J}_{g(k|j\notin\mathbb{J})}} e^{(\delta_k+\beta_{dg(k)})/(1-\sigma)} \tag{33}$$

and

$$\sum_g D^{(1-\sigma)}_{dg|j\notin\mathbb{J}} = \sum_g \left[\sum_{m\in\mathbb{J}_{g|j\notin\mathbb{J}}} e^{(\delta_k+\beta_{dg(k)})/(1-\sigma)}\right]^{(1-\sigma)}. \tag{34}$$

Next, I solve Equation (30) for the parameters that do not represent unobserved heterogeneity, $\delta_k + \beta_{dg}$:

$$\delta_k + \beta_{dg} = \ln(s_{dk}) - \ln(s_{d0}) - \sigma\ln(s_{dk|g}). \tag{35}$$

These terms are then substituted into Equation (32), leaving an expression for the conditional market share with alternative j removed that is a function of data and the unobserved heterogeneity parameter only:

$$s_{dk|j\notin\mathbb{J}} = \frac{e^{[\ln(s_{dk})-\ln(s_{d0})-\sigma\ln(s_{dk})]/(1-\sigma)}}{\widetilde{D}^{\sigma}_{dg(k|j\notin\mathbb{J})}\sum_g \widetilde{D}^{(1-\sigma)}_{dg|j\notin\mathbb{J}}}, \tag{36}$$

where

$$\widetilde{D}_{dg(k|j\notin\mathbb{J})} = \sum_{m\in\mathbb{J}_{g(k|j\notin\mathbb{J})}} e^{[\ln(s_{dm})-\ln(s_{d0})-\sigma\ln(s_{dm})]/(1-\sigma)} \tag{37}$$

and

$$\sum_g \widetilde{D}_{dg|j\notin\mathbb{J}}^{(1-\sigma)} = \sum_g \left[\sum_{m\in\mathbb{J}_{g|j\notin\mathbb{J}}} e^{[\ln(s_{dm})-\ln(s_{d0})-\sigma\ln(s_{dm|g})]/(1-\sigma)} \right]^{(1-\sigma)} . \tag{38}$$

Moment conditions are formed by equating the expressions in Equations (31) and (36). The following closed-form expression for the nesting parameter $\widehat{\sigma}$ can be derived using a similar approach in Section 2.2:

$$\widehat{\sigma} = \frac{1}{J(J-1)} \sum_d \sum_{g=1}^{G} s_{dj|g} \sum_{j\in\mathbb{J}_g} (J_{g(j)}-1) \left[1 - \frac{\ln\left(1+\frac{s_{d0,j}s_{dj}}{s_0}\right)}{\ln\left(1+\frac{s_{dg,j}s_{dj}}{s_{dg}-s_{dj}}\right)} \right] . \tag{39}$$

The parameter $\widehat{\sigma}$ can be computed as an average of the associated micro and macro share data for each combination of d, j, and g, where j is in group g. This computation yields $\hat{\sigma}$. In the second stage, the observed heterogeneity parameters are estimated. They are estimated with the following equation:

$$\ln(s_{dj}) - \ln(s_{d0}) - \hat{\sigma}\ln(s_{dj|g}) = \delta_j + \beta_{dg}. \tag{40}$$

Substituting the decomposition of β_{dg}, $\beta_{dg} = \bar{\beta}_{dg} + \mu_{dj}$, into Equation (40) yields the following estimation equation:

$$\ln(s_{dj}) - \ln(s_{d0}) - \hat{\sigma}\ln(s_{dj|g}) = \delta_j + \bar{\beta}_{dg} + \mu_{dj}. \tag{41}$$

Equation (41) is estimated with a fixed effects regression, with alternative fixed effects δ_j, and with demographic group by alternative group interactions $\bar{\beta}_{dg}$. This yields estimates for the alternative fixed effects, $\hat{\delta}_i$. The third stage is estimated using these alternative fixed effects as the dependent variable in an instrumental variables design of Equation (28). In summary, the method for estimating the model with observed and unobserved heterogeneity includes the following steps:

1. Form groups of decision makers and alternatives based on observed characteristics.

2. If only s_j and s_{dg} are available, impute s_{dj} using the procedure above. Otherwise, skip to step 3.

3. Compute $\widehat{\sigma}$ based on Equation (39) as an average over each combination of d, j, and g, where j is in group g.

4. Estimate second stage fixed effects regression in Equation (41).

5. Estimate the third stage using the estimated $\hat{\delta}_j's$ from the second stage as the dependent variable using Equation (28).

2.6 Further Extensions

In this section, I discuss a series of extensions of the method.

2.6.1 Group-Specific Unobserved Heterogeneity

An extension of the method is that the unobserved heterogeneity parameter, σ, can be computed separately for each alternative group and demographic group, so that each σ_{dg} is estimated. This requires forming separate moment conditions for each group based on second choice data specific to each alternative group and demographic group. Identification requires disaggregated second choice data by alternative group and demographic group. For example, the second choices made by households in urban areas that purchase an SUV may be observed to be different than the second choices made by households in rural areas that purchase an SUV.

2.6.2 Multi-Level Nesting

The method can be extended to include multiple levels of groups, such as in the three-level nested logit model (Train 2009). Suppose alternatives are assigned to a group g and subgroup h associated with group g. The estimation equation for the three-level nested logit model is

$$\ln(s_j) - \ln(s_0) = \delta_j + \sigma_{hg} \ln(s_{j|h}) + \sigma_g \ln(s_{h|g}). \tag{42}$$

This equation is derived in the appendix. Equation (42) has two nesting parameters, σ_{hg} and σ_g, where σ_{hg} represents within subgroup correlation and σ_g represents within-group correlation. Each of the nesting parameters is multiplied by the natural logarithm of a group share. The share $s_{j|h}$ denotes the share of alternative j within its subgroup h, and the share $s_{h|g}$ denotes the combined share of all alternatives in subgroup h within its group g. Each of these shares are observed in aggregate market data. The nesting parameters can be estimated following the method in Section 2.2. This is done by first solving Equation (42) for the mean utility of alternative j:

$$\delta_j = \ln(s_j) - \ln(s_0) - \sigma_{hg} \ln(s_{j|h}) - \sigma_g \ln(s_{h|g}). \tag{43}$$

This Equation is then substituted into the closed-form predicted market share for the three-level nested logit model with an alternative removed from the choice set, following the approach in Section 2.2. Moment conditions equivalent to Equation (9) are formed by equating observed market shares with an alternative removed with predicted market shares with an alternative removed. Closed-form solutions for the share parameters σ_{hg} and σ_g can be easily (but somewhat tediously) derived. In the appendix I show that the following expression relates market shares and the nesting parameters:

$$\ln(s_{k|j\notin\mathbb{J}}) - \ln(s_{0|j\notin\mathbb{J}}) - [\ln(s_k) - \ln(s_0)] = \sigma_{hg}[\ln(s_{k|h,j\notin\mathbb{J}}) - \ln(s_{k|h})] + \sigma_g[\ln(s_{h|g,j\notin\mathbb{J}}) - \ln(s_{h|g})]. \quad (44)$$

In contrast to the estimation equations associated with prior models, Equation (44) includes multiple nesting parameters. Therefore, additional steps must be taken to identify the nesting parameters from using different combinations of alternatives j, k. Two unique sources of variation can be used to identify σ_{hg} and σ_g. The parameter σ_{hg} measures the degree that utility is correlated among alternatives in subgroup h of group g. This parameter is identified with changes in market shares of alternatives that share the same subgroup. One version of Equation (44) is formed by computing $s_{k|j\notin\mathbb{J}}$, $s_{k|h,j\notin\mathbb{J}}$, and $s_{h|g,j\notin\mathbb{J}}$ for pairs of alternatives j, k that are both in subgroup h of group g. These shares are computed as Equation (9) and the following two equations:

$$s_{k|h,j\notin\mathbb{J}} = \frac{q_{k|j\notin\mathbb{J}}}{q_{h|j\notin\mathbb{J}}} = \frac{q_k + s_{k,j}q_j}{q_h + s_{h,j}q_j}, \quad (45)$$

$$s_{h|g,j\notin\mathbb{J}} = \frac{q_{h|j\notin\mathbb{J}}}{q_{g|j\notin\mathbb{J}}} = \frac{q_h + s_{h,j}q_j}{q_g + s_{g,j}q_j}. \quad (46)$$

The terms q_k, q_h,and q_g denote sales of alternative k, all alternatives in subgroup h, and all alternatives in subgroup g, respectively, and $s_{h,j}$ and $s_{g,j}$ denote second choice frequencies of first choice alternative j for alternatives in subgroup h and group g, respectively. The parameter σ_g measures the degree that utility is correlated among alternatives in subgroup h of group g, conditional on within subgroup correlation. Therefore, this parameter can be identified from changes in market shares of alternatives that share the same group but not the same subgroup. A second version of Equation (44) is formed by computing $s_{k|j\notin\mathbb{J}}$, $s_{k|h,j\notin\mathbb{J}}$, and $s_{h|g,j\notin\mathbb{J}}$ for pairs of alternatives j, k that are both in group g but that are in different subgroups. These shares are computed with Equations (45) and (46).

Using the shares of any two pairs of alternatives where one pair shares the same subgroup and another shares the same group but are in different subgroups as in inputs for the two versions of Equation (44) yields a system of two linear equations and two unknowns, which can be easily solved for the two nesting parameters. All such pairs that satisfy the grouping conditions can be included in the calculation of the nesting parameters, as in Equation (16) or Equation (18). This computation yields estimates for the nesting parameters, denoted as $\hat{\sigma}_{hg}$ and $\hat{\sigma}_g$. The following equation can be then be used to compute mean utilities: $\hat{\delta}_j = \ln s_j - \ln(s_0) - \hat{\sigma}_{hg} \ln(s_{j|h}) - \hat{\sigma}_g \ln(s_{h|g})$. Mean utility parameters can then be estimated in a second stage.

2.6.3 Continuous Attributes Heterogeneity

One of the caveats that I mention above is that the empirical model does not incorporate heterogeneous preferences for continuous product attributes. The method described in Section 2.5 is able to incorporate these preferences if market shares and second choice data are available by demographic group. Continuous attributes heterogeneity parameters are estimated in the observed heterogeneity estimation stage with the $\bar{\beta}$ terms replaced by demographic group by continuous attribute interactions. Denoting the value of the continuous attribute a by z_{ja}, the estimation of observed preference heterogeneity in the combined estimation method from Section 2.5 is replaced by

$$\ln(s_{dj}) - \ln(s_{d0}) - \hat{\sigma} \ln(s_{dj|g}) = \delta_j + \sum_a \bar{\beta}_{da} z_{ja} + \mu_{dj}. \tag{47}$$

Simpler models of continuous attributes heterogeneity without unobserved heterogeneity can be estimated by substituting the $\sum_a \bar{\beta}_{da} z_{ja}$ term for the $\bar{\beta}_{dg}$ in Equation (27).

2.6.4 Identifying and Estimating Unobserved Heterogeneity With Repeated Choice Data

For certain settings, although second choice data may not be available, researchers may observe repeated choices made by decision makers. These repeated choice data have often been used to identify unobserved heterogeneity in discrete choice models.[11] Repeated choice data can be converted into second choice data by assuming an ordering among the repeated choices. For example, if a household buys alternatives 1 and 2, the researcher can assume randomly that the household ranks alternative 1 over alternative 2. Then if alternative 1 were not available, the household would choose alternative 2. The random assignment among the chosen alternatives should not impact the estimation of the preference parameters much unless there is little correlation among the alternatives. This sensitivity can be checked by repeating the estimation many times for different random rankings.

[11] Examples of studies using repeated choices for identification include Bento et al. (2009) and Brownstone and Train (1998).

2.6.5 Identifying and Estimating Observed Heterogeneity at the Market Level

Information on household demographics may not be available for some contexts. The researcher, however, is still able to estimate observed heterogeneity if data on product sales are available in more than one market. For example, product sales may be available by geographic region, such as at the state level. In this case, observed heterogeneity can be estimated by interacting market dummy variables with product groups or attributes in the observed heterogeneity estimation stage. In other words, each market is modeled as a distinct group d as described in Section 2.4. This produces observed heterogeneity at the market level, so that decision maker preferences can vary across markets. When markets are defined as distinct geographic regions, this form of heterogeneity can play an important role for assessing regional policies, or a combination of regional and national policies.[12] If a researcher has a dataset with many markets, they can reduce the number of parameters to estimate by aggregating the definition of the group d (to, for example, the regional level).

2.6.6 Calibration

In some cases, modelers may have a tight deadline for completing an analysis of a policy. Given their time constraint, they may want to estimate only certain parameters and calibrate others based on estimates from the literature. The method in the current paper can be adopted to accommodate calibration. For example, suppose a modeler wants to build a discrete choice model that has heterogeneity and mean utility parameters. If the modeler has access to microdata, they can estimate preference heterogeneity parameters using the methodology from this paper. They can then calibrate mean utility parameters based on estimates from the literature. The mean utility parameter for price and non-price product attributes can be calibrated so that the implied own-attribute elasticity of demands or the implied willingness to pay for each non-price attribute match estimates from the literature.[13]

3 Empirical Application: The Effect of Fuel Economy Standards on New Vehicle Demand

In this section, I apply the method for estimating unobserved heterogeneity described in Section 2.2 by estimating consumer demand for new vehicles. I choose to estimate this version of the method because it best illustrates the ability of the method to accurately address a relevant application of the estimation. I estimate a demand model that accounts for heterogeneity along the new-versus-

[12]An example of a combination of policies in the transportation sector includes the Zero Emissions Vehicle (ZEV) mandate administered by of subset of states in the U.S. and the federal fuel economy (CAFE) standards administered by the federal government.

[13]This calibration process is relatively simple because own-attribute elasticities and willingness to pay values have closed-form solutions for the nested logit model. See the appendix for a derivation of the own-price elasticity of demand.

used vehicle choice dimension. This dimension of vehicle choice has received little attention in prior literature, even though it is often relevant for policy analysis. In particular, this dimension influences the prediction of policy outcomes that depend on aggregate market share impacts, such as the effect of a gasoline tax on new and used vehicle ownership.[14] Prior vehicle demand models often omit this choice margin completely due to computational and data constraints (Train and Winston 2007; Whitefoot et al. 2017; Xing et al. 2019). Other demand models, such those presented in Berry et al. (1995), Berry et al. (2004), and Klier and Linn (2012), include an outside option–either a composite used vehicle or the broad choice to not buy a new vehicle–along with new vehicles in the choice set. But in these models, the substitution between new vehicles and an outside option is identified by differences in new-vehicle attributes, which is likely an inaccurate and misleading source of identifying variation for this choice margin. In contrast, the empirical strategy here uses appropriate identifying variation in the form of used vehicle second choice frequencies of new vehicle buyers.

I use the estimated demand model to quantify the effect of fuel economy and greenhouse gas standards on new and used light-duty vehicle sales. The effect of the standards on vehicle sales has long been of interest to policy makers and analysts, yet little research has addressed this policy question, with a recent notable exception being Linn and Dou (2018). The method in the current paper is ideal for quantifying sales impacts, since they are determined by how new vehicle buyers substitute to used vehicles in response to changes in new vehicle characteristics. This substitution pattern is reflected by the willingness of new vehicle buyers to pay for vehicle attributes, and their propensity to prefer new vehicles over used vehicles. In this section, I estimate these two features of new vehicle buyer preferences.

Fuel economy and greenhouse standards in the United States currently require vehicle manufacturers to achieve a sales-weighted average fuel economy and an equivalent level of greenhouse gas emissions among vehicles sold. In 2008, the Obama administration passed legislation to double the average fuel economy requirement by 2025 relative to 2010 levels. The current Trump administration has since proposed to roll back these standards beginning with the 2020 model year.[15] The federal agencies regulating fuel economy and greenhouse gas emissions for light-duty vehicles, the Environmental Protection Agency (EPA) and the National Highway Traffic Safety Administration (NHTSA), have since released a detailed preliminary impact analysis (PRIA) for the proposed rollback (EPA 2018). The PRIA summarizes a detailed calculation of costs and benefits of the rollback, which

[14] This effect is relevant for understanding the impact of gasoline taxes on total gasoline consumption and greenhouse gas emissions (Bento et al. 2009).

[15] The legislation is titled the Safer Affordable Fuel-Efficient (SAFE) Vehicles Rule for Model Year 2021–2026 Passenger Cars and Light Trucks. See https://www.npr.org/2018/08/02/631986713/white-house-proposal-rolls-back-fuel-economy-standards-no-exception-for-californ and https://www.washingtonpost.com/national/health-science/2018/08/01/90c818ac-9125-11e8-8322-b5482bf5e0f5_story.html?noredirect=on&utm_term=.100ead61f250 for news coverage of the rollback.

finds that the rollback will lead to net social benefits. A recent review of this modeling finds substantial flaws with the assessment, suggesting that the sign and magnitudes of the costs and benefits have been grossly misestimated (Bento et al. 2018). The review finds that the key reason for this result is the agencies' flawed modeling of the effect of the standards on new and used vehicle purchases. The agencies use a reduced-form model of vehicle sales to estimate the effect of the rollback on the composition of new and used vehicles on the road. Bento et al. (2018) indicate large flaws with this model and suggest the agencies take a more structural approach for modeling sales impacts. In particular, they recommend developing a vehicle demand model that has parameters estimated with sales and vehicle characteristics data. The structural approach taken in the current paper is one such example of the model they recommend.

Before continuing to the estimation, it is important to recognize other applications of the methods that incorporate observed heterogeneity in Sections 2.4 and 2.5. The incorporation of observed heterogeneity is not only able to help expand the ability of the model to reflect certain substitution patterns, it can provide modelers the ability to perform distributional analysis according to the assigned demographic groups. For example, if household income is observed, demographic groups can be assigned based on this variable, as in Leard et al. (2019). Policy impacts can then be disaggregated by demographic group.

3.1 Data

I use data on new and used vehicle sales, characteristics, and second choice microdata to estimate a vehicle demand model for the 2015 market year, which corresponds to sales from October 2014 to September 2015. New vehicle sales data are from IHS Automotive. These data are highly disaggregated counts of vehicle registrations by quarter. Each observation is defined by buyer type (household vs. fleet), quarter, model year, make, model, trim/series, fuel type, drive type, body style, and engine size (e.g., four cylinder vs. six cylinder). I drop observations for fleet vehicles since the microdata are only for household buyers.[16] I aggregate the sales data to the market year level, combining observations that share the same variable names but have different quarters or model years.[17] Therefore, each observation represents sales of a vehicle by make, model, trim/series, fuel type, drive type, body style, and engine size during the 2015 market year.

I merge with the sales data vehicle characteristics data from Wards Automotive. These data include information on horsepower, weight, and vehicle dimensions. Based on the vehicle dimensions information, I calculate each vehicle's footprint as the product of the vehicle's wheelbase and its track

[16] Fleet vehicles represent 15 to 20 percent of new vehicle sales, and fleet buyers tend to exclusively purchase new vehicles (Leard et al. 2017). Therefore, there is likely little to no substitution between new and used fleet vehicle demand.

[17] Aggregating over model years avoids issues related to sales and pricing impacts due to inventory effects.

width. These data are merged based on all of the unique vehicle identifiers listed above. I merge fuel economy information from the Environmental Protection Agency's fuel economy database, and I merge annual average gasoline, diesel, and electricity prices from the Energy Information Administration, which are all denominated in 2015$.

I merge transaction prices from household survey data obtained from MaritzCX. This survey includes about 210,000 raw observations for the 2015 market year. These data are self-reported transaction prices for vehicles purchased or leased during the 2015 market year. About one-third of the observations have missing transaction price information, leaving around 140,000 usable prices.[18] I compute average transaction prices by all of the unique vehicle identifiers listed above, which are merged to the sales and characteristics data using the same identifiers.

A key feature of the MaritzCX survey data is that it asks respondents about vehicles that the respondents would have bought had their newly acquired vehicle not existed. This represents the second choice data that can be used to form moment conditions for estimating preference heterogeneity. The exact question is "If the model you acquired did NOT exist, what vehicle would have purchased/leased?" The survey asks respondents for the model year of the second choice vehicle, as well as discrete options for the age of the second choice vehicle: new, used, or pre-owned. I code used and certified pre-owned responses as used vehicles. The data include many additional details about the second choice responses, including make, model, fuel type, engine size, and body style, among other characteristics. About two-thirds of the survey observations have valid responses for these questions.[19] I aggregate the second choice decision for new versus used to the vehicle level. This variable represents the expected likelihood that a new vehicle buyer would buy a used vehicle had their obtained new vehicle been unavailable. After merging, I clean the data, leaving 762 vehicle observations for estimation. See the appendix for a detailed description of the data-cleaning steps taken.

I merge data on used car and light truck sales from the CEX corresponding to the 2015 market year. The CEX surveys about 7,000 households each quarter, and includes questions about household purchases and leases of new and used vehicles. I compute a market share for used vehicles based on the proportion of total vehicle purchases and leases that are used. For the 2015 market year, this

[18]These data are similar to the transaction price data used in Leard et al. (2019). See Leard et al. (2019) for more details on the MaritzCX data.

[19]The survey also includes a third and fourth choice option, with the same vehicle characteristics questions. Third and fourth choice data are less frequently provided than the second choice information, but could be used for identification of preference heterogeneity. For example, Train and Winston (2007) use up to four stated second choices by survey respondents to estimate preference heterogeneity among new vehicle buyers.

proportion is 0.681, or a little over two-thirds of the entire light-duty market. This market is consistent with recent reports on sales of new and used vehicles.[20]

Summary statistics for the data appear in Table 2. Average transaction prices are around $40,000. This is substantially higher than the median transaction price in the sample (about $32,000) due to the logarithmic shape of the new vehicle price distribution. The second choice data suggest a strong within-group preference for new vehicles. About 92 percent of new vehicle buyers state that they would have acquired a different new vehicle had their acquired new vehicle not been available. Only 8 percent of these buyers stated they would buy a used vehicle as their second choice. Comparing these proportions to the used vehicle market share confirms that new vehicle buyers have a strong preference for new vehicles. A benchmark comparison is with a logit model, which does not account for shared within-group utility. A logit model would predict that the proportion of new vehicle buyers substituting to a used vehicle would be approximately equal to the market share for used vehicles for a small market share of the removed alternative.[21] The fact that the substitution is much lower suggests a high correlation in utility among new vehicles. The minimum and maximum values for these variables suggest some heterogeneity among vehicles. Curiously, buyers of new 2015 Mini Coopers tend to favor used vehicles as their second choice. This vehicle is the only vehicle observation with a second choice new frequency below 50 percent. The correlation coefficient for transaction price and the second choice new variable is 0.38, suggesting that buyers of inexpensive vehicles are more likely to substitute to a used vehicle. This is consistent with lower income households having a higher price elasticity of demand and opting to buy either inexpensive new vehicles or used vehicles.[22]

3.2 Estimation Results

I specify utility to be a linear function of cost per mile, performance measured as the ratio of horsepower to weight, size measured by footprint, and the natural log of transaction price. I include a control variable for the average model year of each vehicle, and I include fixed effects for fuel type, body style (e.g., pickup truck), and drive type (e.g., all-wheel drive). I estimate utility parameters

[20] Used car and light truck sales in the US are typically around 40 million per year. For example, see https://www.edmunds.com/about/press/used-vehicle-sales-hit-record-high-in-2017-according-to-latest-edmunds-used-car-report.html. New car and light truck sales in the US were about 17 million in 2017: https://www.automobilemag.com/news/u-s-auto-sales-totaled-17-25-million-calendar-2017/.

[21] The proportion is typically slightly larger than the market share for used vehicles. To see this, we know that $s_{0|j\notin\mathbb{J}}^{logit} = s_0^{logit} + s_{0,j}^{logit} s_j^{logit}$, where the superscript denotes the shares are based on the logit model. Solving this equation for the substitution proportion $s_{0,j}^{logit}$ yields $s_{0,j}^{logit} = \frac{s_{0|j\notin\mathbb{J}}^{logit} - s_0^{logit}}{s_j^{logit}}$. A few steps of algebra shows that $s_{0|j\notin\mathbb{J}}^{logit} - s_0^{logit} = s_{0|j\notin\mathbb{J}}^{logit} s_j^{logit}$. Substituting this into the expression for $s_{0,j}^{logit}$ yields $s_{0,j}^{logit} = s_{0|j\notin\mathbb{J}}^{logit}$. The share $s_{0|j\notin\mathbb{J}}^{logit}$ satisfies $s_{0|j\notin\mathbb{J}}^{logit} > s_0^{logit}$ and $s_{0|j\notin\mathbb{J}}^{logit} \approx s_0^{logit}$ for small values of s_j^{logit}.

[22] See Leard et al. (2019) for estimates of observed household demand heterogeneity that are consistent with this pattern.

with a series of logit and nested logit models. The demand estimation results appear in Table 3. Columns (1) and (2) include logit model specifications that do not include a first stage estimation of the nesting parameter. The results appearing in column (1) are estimated with ordinary least squares, and column (2) shows results for an instrumental variables (IV) specification. The coefficient estimates for the vehicle characteristics have expected signs. Vehicle buyers prefer lower prices, lower fuel costs, higher performance, and larger vehicles. For all IV specifications, I construct instruments following Train and Winston (2007), using the sum of continuous characteristics of other vehicles sold by the same manufacturer and the sum of the continuous characteristics of other vehicles sold by other manufacturers in the same body style category, as well as the squares of these sums. For consistency with my simulation exercises, I deviate from Train and Winston (2007) by relaxing the assumption that cost per mile is exogenous, so that only performance and footprint are used as instruments. Therefore, for the IV specifications, both the log of price and cost per mile are instrumented. The logit results appearing in column (2) show that treating price and cost per mile as endogenous increases the price sensitivity, which is consistent with results from prior literature showing that unobserved vehicle characteristics tend to bias the price coefficient toward zero.

Columns (3) and (4) report estimation results for nested logit specifications. The first-stage estimation of the nesting parameter shows a strong within-group correlation, with $\hat{\sigma} = 0.955$. This value is consistent with the high within-group share reported in Table 2. The price coefficient in instrumented nested logit specification appearing in column (4) is about twice as large in magnitude relative to the OLS estimate appearing in column (3). The implied own-price elasticity of demand for the IV specification is -3.58, which is within the range of price elasticity estimates from prior literature (Berry et al. 1995; 2004; Train and Winston 2007).[23] It is also similar to a recent estimate from Leard et al. (2019) that uses a similar level of vehicle aggregation and several years of data. T

To infer household demand for vehicle attributes,I calculate implied willingness to pay (WTP) for a 1 percent change in vehicle attributes and report these figures in Table 3.[24] Households are willing to pay \$60 for a one percent reduction in cost per mile. This estimate is similar to the WTP for fuel cost reductions in Leard et al. (2019). Assuming that the associated lifetime fuel cost savings are \$249 based on calculations from Leard et al. (2017), the implied fuel cost valuation ratio is 0.24. This valuation ratio is defined by a noisy estimate of the cost per mile coefficient. For this reason, in the simulations I vary WTP for fuel cost savings over a range of values from recent literature. Households are willing to pay \$101 for a 1 percent increase in vehicle performance. This is similar in magnitude to the WTP for performance reported in Leard et al. (2017). Households are willing to

[23]See the appendix for a derivation of this elasticity.

[24]Willingness to pay for a unit change in an attribute is the ratio of the marginal utility of the level of an attribute to the marginal utility of price. Obtaining willingness to pay for a 1 percent change requires normalizing the unit change calculation by 1 percent of the level of the attribute.

pay \$546 for a 1 percent increase in footprint, which is consistent with WTP for particular subgroups of the population reported in Leard et al. (2019). The WTP estimates for performance and footprint are near the median of the distribution of estimates reported in Greene et al. (2018).

In Table 3, I report the implied total new vehicle market price elasticity of demand.[25] This elasticity is approximately equal to the percentage change in aggregate new vehicle sales due to a one percent change in all new vehicle prices. The market price elasticity of demand defines the change in new vehicle sales due to a policy that causes changes in new vehicle prices. Therefore, this elasticity can be used to estimate the sales impacts of tightening or relaxing fuel economy standards. For the IV specification, this elasticity is equal to -0.11, suggesting an inelastic market demand response. This response is smaller than the central value from Berry et al. (2004), equal to -1 based on private information from General Motors. However, the central market elasticity assumed in Berry et al. (2004) implies an extraordinarily large (in absolute value) own-price semi-elasticity of demand equal to -10.56. Berry et al. (2004) also calibrate their model with a market elasticity of -0.4, which yields an own-price semi-elasticity equal to -3.94. This implied own-price elasticity is in line with prior estimates and the estimate from this paper, suggesting that -0.4 is a more appropriate market elasticity. The estimate from this paper of -0.11, although smaller in magnitude, is consistent with this inelastic market price elasticity.[26]

The total new vehicle cost per mile elasticity of demand is equal to -0.02, which is about a fifth of the total new vehicle market price elasticity of demand. Therefore, fuel economy standards that lower fuel costs by over a factor of five or more than the associated increase in purchase prices should increase total new vehicle sales. For example, the regulatory impact analysis of the Obama $2017-2025$ fuel economy standards predicted an increase in purchase prices of \$1,800, or roughly 6 percent of the sticker price for model year 2025 vehicles, with an associated reduction in fuel costs of \$5,700 to \$7,400, or roughly 38 to 49 percent of lifetime fuel costs (EPA 2012). The estimated demand parameters applied to the predictions would imply a small increase in new vehicle sales due to tightening standards.[27]

A final note about the estimation is that it is computationally fast. The first stage for the nested logit models, which is computed with a closed-form expression, take less than one second to estimate.[28] Of course, the computational time necessary to estimate the first stage will be longer if a larger dataset

[25] See the appendix for a derivation of this elasticity.

[26] This elasticity is smaller than but similar in magnitude to the demand response found in EPA (2018), which uses a reduced-form time series model to estimate the effect of changes in new vehicle prices on new vehicle sales. The implied market price elasticity from that model falls in the range of -0.2 to -0.3, suggesting an inelastic demand response.

[27] Equivalently, the recent proposal to roll back $2021-2025$ standards should be expected to reduce new vehicle sales, which is contrary to findings from the preliminary regulatory impact analysis of the rollback (EPA 2018). Of course, these results critically depend on the forecasted changes in purchase prices and fuel costs.

[28] For each model, the final second stage takes under one second.

is used or if more than one nesting parameter is estimated. But the time here is orders of magnitude faster than most standard BLP or micro-BLP estimation routines, which can take hours or even days to estimate. For modelers and policy analysts who want to build a discrete choice model by running multiple specifications or apply many different specifications for the purposes of policy simulation, this short estimation time is likely to prove useful.

3.3 Simulation of Tightening Fuel Economy Standards

I use the estimated demand model to simulate the sales impacts of tightening fuel economy standards. I take a stylized approach to quantify the effect of the standards on sales. I assume a simple supply-side response by manufacturers, which pass the costs of the standards on to new vehicle buyers in the form of higher vehicle prices. I further assume that the standards affect vehicle prices uniformally. I consider the case of a 1 percent increase in the stringency of the standards relative to 2015 fuel economy levels. This is modeled by increasing each vehicle's fuel economy by 1 percent. Following estimates implied by Leard et al. (2019) that are based on engineering relationships between fuel-savings from technology adoption and manufacturing costs, for the benchmark simulation I assume that this increase in stringency is accompanied by an increase of vehicle prices by 0.25 percent for cars and 0.18 percent for light trucks.[29] I reference this setting as using engineering technology costs. I abstract from a non-uniform increase in prices due to pricing competition to focus on the impact of varying demand modeling assumptions on sales.[30]

I consider four simulation scenarios. In the first scenario, changes in vehicle prices are defined by the engineering technology cost relationships defined above and where vehicle buyers do not value changes in cost per mile. This scenario is consistent with assumptions made in the new vehicle sales simulation model adopted by federal agencies in quantifying the effects of the recently proposed rollback of fuel economy standards (EPA 2018). Under this scenario, tighter standards increase new vehicle prices and lower costs per mile, but new vehicle demand only responds to the increase in new vehicle prices. This is equivalent to assuming that the cost per mile coefficient is equal to zero, or assuming that the change in cost per mile is equal to zero. For the second and third scenarios, I assume that changes in the present value of fuel costs equal changes in vehicle prices as a result of a marginal tightening of the standards. This scenario requires calibrating the relationship between cost per mile and the present value of lifetime fuel costs. Discounted fuel costs equal the product of cost per mile and the present discounted miles driven. For the latter, I assume that cars and light trucks are driven 195,264

[29]Leard et al. (2019) estimate an elasticity of vehicle manufacturer marginal costs to fuel economy of about 0.25 for cars and 0.18 for light trucks. Assuming that changes in marginal costs are fully passed on to new vehicle buyers in the form of higher prices yields the assumptions made in the current paper.

[30]Recent examples of modeling efforts to incorporate pricing effects include Jacobsen (2013), Reynaert (2017), and Leard et al. (2019).

and 225,865 miles, respectively, following EPA (2012). Therefore, a one percent change in cost per mile increases new vehicle purchase prices by the product of cost per mile and 1,952.64 for cars and 2,258.65 for trucks.[31] These scenarios represent a setting where all technologies where the associated fuel cost savings exceed installation costs have already been adopted. This setting implies there is no market failure on the supply side of the market for fuel economy. In the second scenario, I continue to assume that consumers do not value fuel cost savings. In the third scenario, I assume that changes in fuel costs equal changes in vehicle prices and consumers value fuel cost savings according to the demand model estimates from the nested logit IV specification (column (4) in Table 3). In the fourth scenario, I assume that changes in fuel costs equal changes in vehicle prices and consumer value 75 percent of fuel cost savings, which is approximately three times the implied valuation from the third scenario.[32]These three scenarios present a wide range of alternative assumptions regarding technology costs and consumer demand.

The simulated effects of a marginal tightening of fuel economy standards appear in Figure 1. The vertical axis measures the percentage change in new vehicle sales. In the benchmark scenario where changes in prices are defined by engineering estimates and where vehicle buyers do not value fuel cost savings, new vehicle sales fall by about 0.023 percent as a result of a 1 percent tightening. This magnitude is smaller than the change in new vehicle sales estimated in Linn and Dou (2018): they use a reduced-form approach relating new vehicle sales to fuel economy stringency over time and estimate that a 1 percent tightening reduces new vehicle sales by 0.1 percent. However, in the second scenario where changes in prices equal changes in fuel costs and where consumers do not value fuel cost savings, new vehicle sales fall by 0.082 percent, which is close to the estimate in Linn and Dou (2018). For the third scenario where consumers value fuel cost savings according to the IV nested logit demand estimates, new vehicle sales fall less, by 0.061 percent, as reduction in demand from higher prices is tempered by higher demand due to lower fuel costs. If consumers value 75 percent of fuel cost savings, a 1-percent tightening of fuel economy standards reduces new vehicle sales by 0.021 percent. Across all of the scenarios, the change in new vehicle sales is quite modest: the elasticity of new vehicle sales to fuel economy stringency is highly inelastic. This is due to the limited substitution between new and used vehicles, given the inelastic demand of the new vehicles estimate reported in Table 3.

[31] This could be overestimate of the cost increase since this calculation includes an implicit assumption that the real discount rate is equal to zero. Therefore, this scenario can be interpreted as an upper bound for the change (in magnitude) in new vehicle sales as a result of a marginal tightening of the standards. An alternative approach is to compute present discounted miles driven using an annual miles schedule estimated from household data in National Highway Travel Survey and household survey data on auto loans, as in Leard et al. (2017).

[32] This valuation ratio is consistent with benchmark estimates from Allcott and Wozny (2014).

3.4 Comparison with Alternative Models

To address how the method for identifying and estimating preference heterogeneity is relevant for assessing policies like fuel economy standards, I compare the simulation results to outcomes derived from alternative models. I consider two alternative models. The first has a smaller nesting parameter σ equal to one-half of the estimated parameter in Table 3. This alternative represents a nested logit model with parameters estimated based on macrodata alone using an instrumental variables strategy as suggested in Berry (1994). It can also represent a mixed logit model that has a random parameter for the outside option (used vehicles in the current context) that is estimated without second choice microdata. The second alternative model is an IV logit version of the benchmark model, which I define as the set of model parameters from column (4) in Table 3, where vehicle buyers do not value changes in cost per mile (represented by the left bar in Figure 1). For both of these alternative models, I adjust the price coefficient so that the implied own-price elasticity of demand equals -3.58 to be consistent with the benchmark model.

The simulated effects of tightening standards on new vehicle sales for these alternative models appear in Figure 2. The benchmark model results appear as the left-most bar for comparison. The change in new vehicle sales is substantially larger for both of the alternative models. For the logit model, the simulated change in new vehicle sales is about 20 times as large as the benchmark. Although the predicted change is not as extremely different for the model with a smaller assumed value for the nesting parameter σ, it is about an order of magnitude larger. These differences are much larger than the differences in sales impacts implied by adjusting assumptions about how tightening the standards affects vehicle prices or how vehicle buyers value fuel costs. These stark differences highlight the importance of using microdata for identifying the preference heterogeneity coefficients.

3.5 Caveats

The empirical results presented here come with several caveats. The model is a highly styled version of the light-duty vehicle market, so that the results should be interpreted more qualitatively than quantitatively. Many of the assumptions made are likely to affect the overall magnitudes of the sales impacts. In particular, a marginal tightening of the federal fuel economy standard is unlikely to uniformly raise vehicle prices. The imperfectly competitive nature of the new vehicle market likely makes the standards have heterogeneous effects across vehicles. Vehicle buyers are also assumed to have homogeneous preferences for new vehicle characteristics. Prior literature has shown that vehicle buyers have quite heterogeneous preferences for vehicle characteristics (Berry et al. 1995; 2004; Train and Winston 2007; Leard et al. 2019; Xing et al. 2019). This form of consumer heterogeneity may affect the relationship between fuel economy standards and new vehicle sales. It is also likely to play a key role in evaluating the incidence impacts of the standards (Jacobsen 2013; Leard et al. 2019). I also do not model tradeoffs between new vehicle characteristics, including the tradeoffs between vehicle performance, fuel economy, and weight (Knittel 2011; Klier and Linn 2012; Leard et al. 2019). Leard et al. (2017) show that the tradeoff between vehicle performance and fuel economy has significant implications for assessing the welfare and sales impacts of fuel economy standards. This is because manufacturers tend to forego performance increases to meet the tightened standards and new vehicle buyers have a relatively high valuation of performance. Therefore, tightening standards can lower new vehicle buyer welfare and sales if the value of the sacrificed performance is sufficiently large. A more detailed simulation model should incorporate this tradeoff to achieve a more accurate assessment of the sales changes due to tightening standards.

4 Applications in Other Settings

The methodology here can be applied to other empirical settings beyond the vehicles market. One application is the estimation of a mode travel choice model, in which households choose a travel mode–such as taking the bus–for their daily commute to work. Empirically estimated mode choice models have been developed in the transportation engineering literature to address impacts of various policy interventions, such as the effect of subsidizing public transit. But they have not been widely developed in the economics literature, partly due to a lack of quality market-level data on mode choices.[33] The 2017 wave of the National Household Travel Survey (NHTS) may be used to estimate a national mode choice model. This version of the survey includes questions that are useful for identifying observed and unobserved heterogeneity parameters in a mode choice model. The survey has mode choices linked with household demographics. These data can be used to identify and estimate observed heterogeneity, as described in Section 2.4. The survey asks, "If you were unable to use your household vehicle(s), which of the following options would be available to you to get you from place to place?" The options include walking, biking, taking a bus, taking a train, and taking a rideshare. Although this question is not a clean-cut request for the traveler's preferred second choice, the question's responses do contain similar information about household substitution patterns among the different modes. The responses from this question along with stated mode choices can be aggregated and used for estimating unobserved heterogeneity in a mode choice model following the method described in Section 2.2.

Many empirical settings have aggregate sales data and microdata aggregated to different levels. For example, the IHS new vehicle sales data that I use are highly disaggregated, including trim and engine size configurations for each model. This contrasts with many public sets of microdata of vehicle choice, such as the CEX, which only has vehicle identifiers at the make-by-class level. This makes combining the data for estimation challenging, although recent research has derived methods for incorporating

[33] One example of a study in the economics literature on mode choice is Parry (2009), which includes a calibration exercise for a model of the choice between commuting by car, bus, or light rail.

datasets that are aggregated differently (Brownstone and Li 2018).[34] The method that I develop is easily capable of handling different levels of aggregation. The researcher can define demographic and vehicle groups by the levels of aggregation of the microdata (which tends to be more aggregated). These generally are sufficiently disaggregated for modeling an appropriate amount of heterogeneity. For example, the CEX data on new and used vehicle purchases used in this paper also differentiates between new and used cars and trucks, which I could add to the empirical model to reflect substitution between these classes.

The methods presented in the current paper are especially useful for settings with extremely large choice sets. One example of such a setting is the explicit modeling of the decision to buy a new or used vehicle, in which all new and used vehicles are represented as unique alternatives. A recent example of this type of model is Bento et al. (2009), which estimates observed and unobserved heterogeneity in a new and used vehicle demand model with household level data from the 2001 wave of the NHTS. A key benefit of this approach is that the model is capable of predicting compositional changes in the used vehicle market in response to various policies, such as a gasoline tax or a tightening of fuel economy standards (Jacobsen 2013). These used vehicle market changes have been shown to be relevant for cost-benefit analysis of federal fuel economy standards (EPA 2018; Jacobsen 2013; Jacobsen and van Benthem 2015; Bento et al. 2018). Bento et al. (2009) adopt a Bayesian estimation approach and aggregate their vehicle choice set to avoid computational constraints, creating a choice set of 270 alternatives. This aggregation likely masks relevant substitution patterns, and it may bias implied elasticities that are relevant for policy analysis. An alternative to their approach is to adopt the simplified estimation method from this paper, exploiting the household-level data that the NHTS has to offer for identifying observed and unobserved heterogeneity. These data include household demographics linked with vehicle ownership, and they include all of the vehicles owned by each household. The household demographics linked with vehicle ownership data can be used to identify and estimate observed heterogeneity based on the method from this paper. The vehicle portfolio can be used for identifying unobserved heterogeneity by assuming that vehicle ownership is a separate, repeated choice, following the logic described in Bento et al. (2009). As I explain in Section 2.6, the repeated choice data can be converted to second choice data, which then can be used to form moment conditions for estimating unobserved heterogeneity based on the method from this paper.

[34] A traditional method for accounting for different levels of aggregation is to aggregate the more disaggregated dataset to the level of the most aggregated dataset (Bento et al. 2009; Klier and Linn 2012). But doing so often masks relevant variation that can be used for identifying model parameters and may even bias parameter estimates (Brownstone and Li 2018).

5 Conclusion

Using discrete choice models for differentiated products to address questions about market and policy outcomes remains both theoretically and computationally challenging. In this paper, I derive a simple approach for the identification and estimation of observed and unobserved heterogeneity parameters that helps create plausible substitution patterns. The method requires an additional source of identification in the form of microdata, but is estimated with basic estimation routines, rendering it easily accessible and computationally fast. The accessibility of the method should lower the entry barrier for it to be adopted by other analysts and policy makers. Furthermore, this method can be combined with recent innovations for estimating unbiased mean utility parameters in a final estimation stage, such as using optimal instruments (Reynaert and Verboven 2014; Reynaert 2017; Grigolon et al. 2018).

I illustrate the method by estimating a vehicle demand model that incorporates vehicle buyer heterogeneity along the new versus used dimension. I use second choice data to identify the heterogeneity parameter, finding a strong correlation in utility among new vehicles. I then evaluate the implications of this heterogeneity by simulating the sales impacts of federal fuel economy standards. I find that the model predicts a small sales impact from a marginal tightening of the fuel economy standard. This is in contrast to the simulations I perform with alternative models that have limited or no heterogeneity along the new versus used dimension, which predict a sales impact that is an order of magnitude larger.

The method that I have developed can be applied to estimate parameters of choice models that can be used to perform cost-benefit analysis calculations for major social policies, such as fuel economy and greenhouse gas standards for light-duty vehicles. This application can address the weaknesses highlighted in Bento et al. (2018) in the most recent analyses of federal fuel economy standards made by federal agencies by providing an economic modeling framework that predicts plausible vehicle substitution patterns.

References

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

Bento, A., K. Gillingham, M. R. Jacobsen, C. Knittel, B. Leard, J. Linn, V. McConnell, D. Rapson, J. Sallee, A. van Benthem, and K. Whitefoot (2018). Flawed analyses of U.S. auto fuel economy standards. *Science 362*(6419), 1119–1121.

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.

Bento, A., K. Roth, and S. Zhou (2018). Vehicle lifetime trends and scrappage behavior in the used car market. *Energy Journal 39*(1).

Beresteanu, A. and S. Li (2011). Gasoline prices, government support, and the demand for hybrid vehicles in the united states. *International Economic Review 52*(1), 161–182.

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 (1999). Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review 89*(3), 400–430.

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.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics 25*(2), 242–262.

Brenkers, R. and F. Verboven (2006). Liberalizing a distribution system: The european car market. *Journal of the European Economic Association 4*(1), 216–251.

Brownstone, D. and P. Li (2018). A model for broad choice data. *Journal of Choice Modeling 27*, 19–36.

Brownstone, D. and K. Train (1998). Forecasting new product penetration with flexible substitution patterns. *Journal of Econometrics 89*(1-2), 109–129.

EPA (2012). EPA and NHTSA set standards to reduce greenhouse gases and improve fuel economy for model years 2017 - 2025 cars and light trucks. Technical report, Environmental Protection Agency Office of Transportation and Air Quality.

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*. https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/ld-cafe-co2-nhtsa-2127-al76-epa-pria-180823.pdf.

Fox, J. T., K. I. Kim, S. P. Ryan, and P. Bajari (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics 2*, 381–418.

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.

Grigolon, L., M. Reynaert, and F. Verboven (2018). Consumer valuation of fuel costs and tax policy: Evidence from the European car market. *American Economic Journal: Economic Policy 10*(3), 193–225.

Grigolon, L. and F. Verboven (2014). Nested logit or random coefficients logit? A comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics 96*(5), 916–935.

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

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

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

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.

Knittel, C. (2011). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *American Economic Review 101*(7), 3368–3399.

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

Langer, A. (2016). (Dis)incentives for demographic price discrimination in the new vehicle market. *Working Paper*.

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 K. Springel (2019). Pass-through and welfare effects of regulations that affect product attributes. *Resources for the Future Working Paper 19-06*.

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., V. McConnell, and C. Zhou (2017). The effect of fuel price changes on fleet demand for new vehicle fuel economy. *Resources for the Future working paper 17 - 25*.

Li, J. (2018). Compatibility and investment in the U.S. electric vehicle market. *Working Paper*.

Linn, J. and X. Dou (2018). How do US passenger vehicle fuel economy standards affect purchases of new and used vehicles? *Resources for the Future Report*.

McFadden, D. (1974). *Conditional logit analysis of qualitative choice behavior.* Frontiers in Econometrics, edited by P. Zarembka. New York: Academic Press.

McFadden, D. and K. Train (2000). Mixed MNL models for discrete response. *Journal of Applied Econometrics 15*(5), 447–470.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereals industry. *Econometrica* (69), 307–342.

NHTSA (2012). Final regulatory impact analysis for corporate average fuel economy for MY 2017-MY 2025 passenger cars and light trucks. Technical report, National Highway Traffic Safety Administration.

$https://one.nhtsa.gov/staticfiles/rulemaking/pdf/cafe/FRIA_2017-2025.pdf$.

Parry, I. (2009). Should urban transit subsidies be reduced? *American Economic Review 99*(3), 700–724.

Petrin, A. (2002). Quantifying the benefit of new products: The case of the 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*.

Reynaert, M. and F. Verboven (2014). Improving the performance of random coefficients demand models: The role of optimal instruments. *Journal of Econometrics 179*(1), 83–98.

Springel, K. (2017). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *Working Paper*.

Thomadsen, R. (2005). The effect of ownership structure on prices in geographically differentiated industries. *RAND Journal of Economics 36*(4), 908–929.

Train, K. (2009). *Discrete Choice Methods with Simulation* (2nd ed.). Cambridge, UK: Cambridge University Press.

Train, K. E. and C. Winston (2007). Vehicle choice behavior and the declining market share of U.S. automakers. *International Economic Review 48*, 1469–1496.

Whitefoot, K. S., M. L. Fowlie, and S. J. Skerlos (2017). Compliance by design: Influence of acceleration tradeoffs on CO2 emissions and costs of fuel economy and greenhouse gas regulations. *Environmental Science and Technology 51*(18), 10307–10315.

Xing, J., B. Leard, and S. Li (2019). What does an electric vehicle replace? *Working Paper*.

Figures

Figure 1: Predicted Percentage Changes in New Vehicle Sales Due to a 1-Percent Increase in New Vehicle Fuel Economy

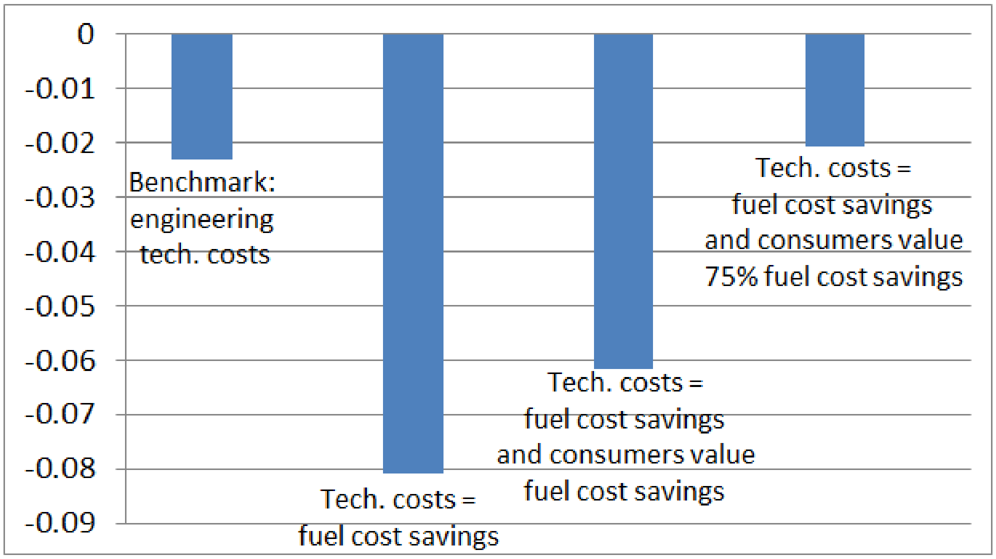

Notes: The figure reports the simulated change in new vehicle sales as a result of a 1-percent tightening of new vehicle fuel economy standards. The change in new vehicle sales is measured as a percentage change relative to 2015 new vehicle sales. The leftmost bar represents a setting where the change in technology costs are defined by engineering costs as described in the text, and where vehicle buyers do not value changes in cost per mile. In this setting, a tightening of fuel economy standards increases new vehicle purchase prices, which reduces new vehicle sales. The second bar represents a setting where changes in technology costs equal changes in associated present value lifetime fuel costs and where vehicle buyers do not value changes in cost per mile. In this setting, tightening of fuel economy standards increases new vehicle purchase prices and reduces fuel costs per mile of new vehicles. The third bar represents a setting where changes in technology costs equal changes in associated present value lifetime fuel costs and where vehicle buyers value changes in cost per mile according to parameter estimates from the nested logit IV demand model estimation. The rightmost bar represents a setting where changes in technology costs equal changes in associated present value lifetime fuel costs and where vehicle buyers value 75% of changes in fuel cost savings.

Figure 2: Alternative Model Predictions of Percentage Changes in New Vehicle Sales

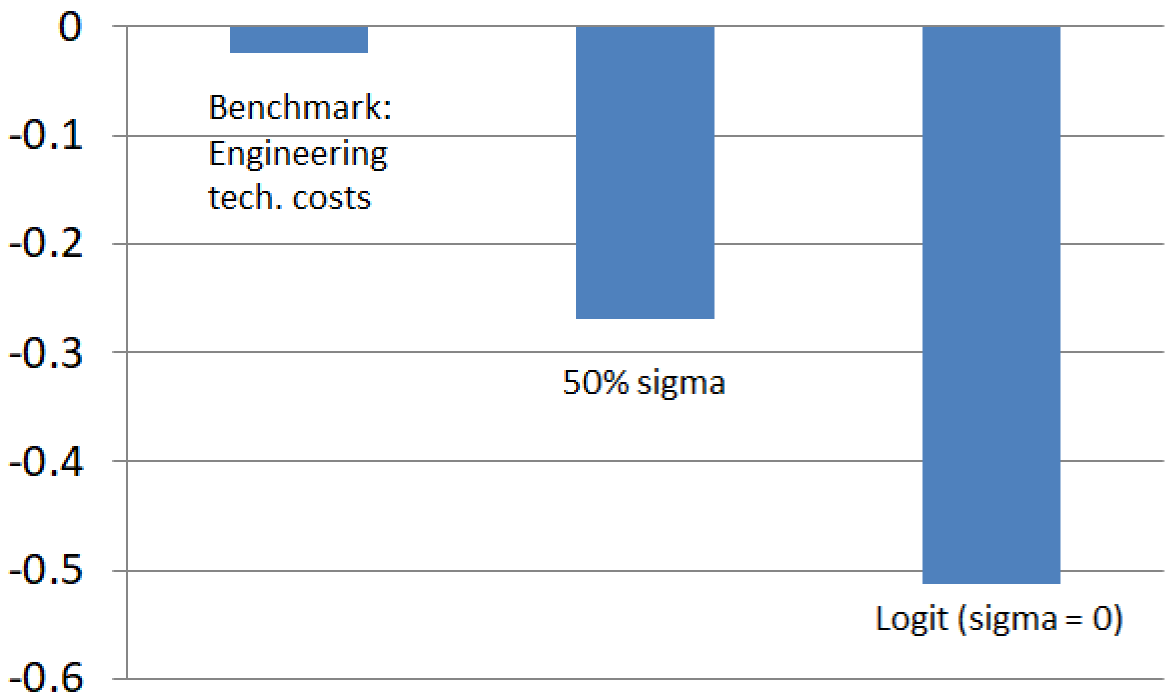

Notes: The figure reports the simulated change in new vehicle sales as a result of a 1-percent tightening of new vehicle fuel economy standards for alternative assumptions on the degree of vehicle buyer heterogeneity. The change in new vehicle sales is measured as a percentage change relative to 2015 new vehicle sales. Each scenario represents a setting where vehicle buyers do not value changes in cost per mile. The leftmost bar indicates the sales impact in the benchmark setting, which is equivalent to the simulation results from the leftmost bar in Figure 1. This sales impact is based on the heterogeneity estimates from the nested logit IV demand estimation. The middle bar represents the sales impacts predicted by a model with less consumer heterogeneity, as measured by recalibrating the demand model with a value for σ that is equal to 50 percent of the estimated value reported in Table 3. The rightmost bar represents the sales impacts predicted by a model with no consumer heterogeneity, as measured by recalibrating the demand model with a value of $\sigma = 0$. In both recalibrations, the own-price elasticity of demand is recalibrated to match the implied own-price elasticity from the nested logit IV demand estimation. In each of the simulations, changes in new vehicle prices are defined by the engineering technology cost relationships as defined in the text.

Tables

Table 1: Data for Numerical Example Estimation of Unobserved Heterogeneity

Panel (a): High Correlation

Alternative	Market share	Second choice freq. with j removed			Market share with j removed		
		$j=1$	$j=2$	$j=3$	$j=1$	$j=2$	$j=3$
$j=0$	0.50	0.05	0.24	0.10	0.51	0.52	0.52
$j=1$	0.22	–	0.38	0.66	–	0.25	0.35
$j=2$	0.08	0.27	–	0.24	0.14	–	0.13
$j=3$	0.20	0.68	0.38	–	0.35	0.23	–
All inside alternatives	0.50	0.95	0.76	0.90	0.49	0.48	0.48
Estimated $\hat{\sigma} = 0.90$							

Panel (b): Low Correlation

Alternative	Market share	Second choice freq. with j removed			Market share with j removed		
		$j=1$	$j=2$	$j=3$	$j=1$	$j=2$	$j=3$
$j=0$	0.50	0.41	0.50	0.625	0.59	0.54	0.625
$j=1$	0.22	–	0.26	0.275	–	0.24	0.275
$j=2$	0.08	0.17	–	0.10	0.12	–	0.10
$j=3$	0.20	0.42	0.24	–	0.29	0.22	–
All inside alternatives	0.50	0.59	0.50	0.375	0.41	0.46	0.375
Estimated $\hat{\sigma} = 0.27$							

Notes: The table reports example data and estimated unobserved preference heterogeneity parameters. In each panel, the outside option is assumed to have a market share of 0.5. Panel (a) includes example data for a setting where utility for alternatives in the same group is highly correlated. The example data in this setting are calibrated to reflect a strong within-group substitution when an inside alternative is removed from the choice set. Panel (b) includes example data for a setting where utility for alternatives in the same group is weakly correlated. The example data in this setting are calibrated to reflect a weak within-group substitution when an inside alternative is removed from the choice set. The unobserved heterogeneity parameter σ is estimated with the method outlined in Section 2.2. In Panel (b), the case with alternative $j = 3$ removed from the choice set has data generated from a logit model.

Table 2: Summary Statistics for 2015 Vehicle Sample

Variable	Mean	Std. Dev.	Min	Max
Sales	15,961	30,110	107	299,101
Transaction price	40,790	17,620	14,673	98,749
Cost per mile	0.114	0.026	0.029	0.197
Horsepower/weight	0.065	0.019	0.012	0.184
Footprint	8.151	1.016	4.513	13.152
Second choice new	0.920	0.078	0.448	1
Second choice used	0.080	0.078	0	0.551
All-wheel drive	0.278	0.448	0	1
Sedan	0.324	0.468	0	1
SUV	0.382	0.486	0	1
Hybrid	0.035	0.185	0	1
Plug-in hybrid or electric	0.016	0.125	0	1
Used vehicle market share	0.681	0	0.681	0.681

Notes: The table reports summary statistics of characteristics and sales for new vehicles sold during the 2015 market year. The total number of vehicle observations is 762. Vehicle transaction prices are from the MaritzCX microdata. Non-price attributes are from Wards Automotive. Cost per mile is defined as the average annual fuel price divided by fuel economy. For gasoline vehicles, this is the average annual gasoline price (from the Energy Information Administration) divided by the vehicle's fuel economy. For electric vehicles, this is the average annual electricity price (from the Energy Information Administration) divided by the vehicle's electricity use per mile. For plug-in hybrid vehicles, a weighted average approach following Leard et al. (2017) is used to construct cost per mile. Vehicle prices and costs per mile are denominated in 2015$. Second choice new and second choice used are variables constructed from MaritzCX microdata. These variables represent the frequency of second choice vehicles being either new or used, respectively. The used vehicle market share is computed based on used and new car purchase and lease data from the Consumer Expenditure Survey.

Table 3: Demand Estimation Results

Variables	(1) Logit OLS	(2) Logit IV	(3) NLogit OLS	(4) NLogit IV
First stage				
σ			0.955	0.955
Second stage				
Ln(Price)	-1.774	-3.588	-0.0795	-0.161
	(0.261)	(0.765)	(0.0117)	(0.0343)
Cost per mile	-13.4	-5.803	-0.6	-0.260
	(5.153)	(18.37)	(0.231)	(0.823)
Horsepower/weight	5.634	17.94	0.252	0.804
	(6.119)	(11.13)	(0.274)	(0.499)
Footprint	0.491	0.711	0.022	0.0318
	(0.118)	(0.283)	(0.00529)	(0.0127)
Constant	-1,936	-2,923	-87.46	-131.7
	(741.6)	(985.4)	(33.23)	(44.15)
Observations	762	762	762	762
R-squared	0.271	0.202	0.271	0.202
Own-price elasticity of demand			-1.77	-3.58
Own-cost per mile elasticity of demand			-1.45	-0.63
Total new vehicle market price elasticity of demand			-0.05	-0.11
Total new vehicle market cost per mile elasticity of demand			-0.04	-0.02
WTP for a 1% reduction in cost per mile			282	60
WTP for a 1% increase in horsepower/weight			64	101
WTP for a 1% increase in footprint			764	546

Notes: Standard errors are reported in parentheses and are clustered by vehicle model, e.g., Toyota Prius. Vehicle prices and cost per mile are denominated in 2015$. The instruments used for specifications in columns (2) and (4) include the sales-weighted sum of horsepower/weight and footprint for all other vehicles sold by the same firm and for all other vehicles sold by other firms sharing the same vehicle body style (e.g., SUV), as well as the squares of these sums. The own-price elasticity of demand is calculated according to the formula in Appendix A.8. It is calculated as the average across all vehicle models and is weighted by vehicle sales. The own-cost per mile elasticity of demand is calculated using a similar formula. The total new vehicle market price elasticity of demand is calculated according to the formula in Appendix A.9. The total new vehicle market cost per mile elasticity of demand is calculated using a similar formula. The willingness to pay (WTP) calculations are based on the ratio of the estimated marginal utility for a vehicle attribute to the marginal utility of vehicle price. All WTP values are reported in 2015$.

Appendix

A.1 Derivation of Closed-Form Expression for Unobserved Heterogeneity Parameter

Denote the outside share with alternative j removed as

$$s_{0|j\notin\mathbb{J}} = \frac{1}{\sum_g \widetilde{D}_{g|j\notin\mathbb{J}}^{(1-\sigma)}}. \tag{A.1}$$

Taking the difference of the natural log of $s_{k|j\notin\mathbb{J}}$ and the natural log of $s_{0|j\notin\mathbb{J}}$ yields

$$\ln(s_{k|j\notin\mathbb{J}}) - \ln(s_{0|j\notin\mathbb{J}}) = \frac{\ln(s_k) - \ln(s_0) - \sigma\ln(s_{k|g})}{1-\sigma} - \sigma\ln(\widetilde{D}_{g(k|j\notin\mathbb{J})}). \tag{A.2}$$

Adding and subtracting $\sigma \ln e^{[\ln(s_k) - \ln(s_0) - \sigma\ln(s_{k|g})]/(1-\sigma)}$ gives

$$\ln(s_{k|j\notin\mathbb{J}}) - \ln(s_{0|j\notin\mathbb{J}}) = \frac{\ln(s_k) - \ln(s_0) - \sigma\ln(s_{k|g})}{1-\sigma} - \sigma\frac{\ln(s_k) - \ln(s_0) - \sigma\ln(s_{k|g})}{1-\sigma} + \sigma\ln(s_{k|g,j\notin\mathbb{J}}), \tag{A.3}$$

where $s_{k|g,j\notin\mathbb{J}}$ is the alternative k within-group share with alternative j removed from the choice set, defined as $s_{k|g,j\notin\mathbb{J}} = \frac{e^{[\ln(s_k)-\ln(s_0)-\sigma\ln(s_{k|g})]/(1-\sigma)}}{\widetilde{D}_{g(k|j\notin\mathbb{J})}}$. Equation (A.3) simplifies to

$$\ln(s_{k|j\notin\mathbb{J}}) - \ln(s_{0|j\notin\mathbb{J}}) = \ln(s_k) - \ln(s_0) - \sigma\ln(s_{k|g}) + \sigma\ln(s_{k|g,j\notin\mathbb{J}}). \tag{A.4}$$

Appendix Equation (A.4) can be solved for σ:

$$\sigma = \frac{\ln(s_{k|j\notin\mathbb{J}}) - \ln(s_{0|j\notin\mathbb{J}}) - [\ln(s_k) - \ln(s_0)]}{\ln s_{k|g,j\notin\mathbb{J}} - \ln(s_{k|g})}. \tag{A.5}$$

Substituting Equations (9) and (10) into Appendix Equation (A.5) yields

$$\sigma = \frac{\ln(s_k + s_{k,j}s_j) - \ln(s_0 + s_{0,j}s_j) - [\ln(s_k) - \ln(s_0)]}{\ln(s_{k|g} + s_{k,j}s_{j|g}) - \ln(s_{k|g})}. \tag{A.6}$$

We know that $\ln(s_k + s_{k,j}s_j) - \ln(s_k) = \ln(s_{k|g} + s_{k,j}s_{j|g}) - \ln(s_{k|g})$ since $\frac{s_j}{s_k} = \frac{s_{j|g}}{s_{k|g}}$. Therefore, Appendix Equation (A.6) simplifies to

$$\sigma = 1 - \frac{\ln(s_0 + s_{0,j}s_j) - \ln(s_0)}{\ln(s_{k|g} + s_{k,j}s_{j|g}) - \ln(s_{k|g})}. \tag{A.7}$$

We can further simplify the expression by substituting the denominator for $\ln(s_k+s_{k,j}s_j)-\ln(s_k)$:

$$\sigma = 1 - \frac{\ln(s_0 + s_{0,j}s_j) - \ln(s_0)}{\ln(s_k + s_{k,j}s_j) - \ln(s_k)}. \tag{A.8}$$

Taking a weighted average over all j, k pairs, where alternatives j and k share the same group and the weights are equal to market shares, yields an estimate for σ:

$$\hat{\sigma} = \frac{1}{J(J-1)} \sum_{g=1}^{G} s_j \sum_{j \in J_g} \sum_{k \in J_g, k \neq j} \left[1 - \frac{\ln(s_0 + s_{0,j}s_j) - \ln(s_0)}{\ln(s_k + s_{k,j}s_j) - \ln(s_k)} \right]. \tag{A.9}$$

A.2 Derivation of Multinesting Parameter Share Equation

Alternatives are grouped into $G+1$ groups, indexed by $g = 0, 1, 2, ..., G$. The outside option $j = 0$ is assumed to be the only alternative in group $g = 0$. Consumer i obtains utility u_{ij} when choosing alternative j in group $g(j)$, where utility is

$$u_{ij} = \delta_j + \xi_{ig} + (1 - \sigma_{g(j)})\epsilon_{ij}. \tag{A.10}$$

Assuming that ϵ_{ij} is identically and independently distributed extreme value, the predicted market share for alternative j is[35]

$$s_j = \frac{e^{\delta_j/(1-\sigma_{g(j)})}}{\left(\sum_{k \in \mathbb{J}_{g(j)}} e^{\delta_k/(1-\sigma_{g(k)})}\right)^{\sigma_{g(j)}} \sum_g \left(\sum_{k \in \mathbb{J}_g} e^{\delta_k/(1-\sigma_g)}\right)^{1-\sigma_g}}. \tag{A.11}$$

Given the nested logit specification, the predicted within-group share for alternative j is a logit formula:

$$s_{j|g} = \frac{e^{\delta_j/(1-\sigma_{g(j)})}}{\sum_{k \in \mathbb{J}_{g(j)}} e^{\delta_k/(1-\sigma_{g(k)})}}. \tag{A.12}$$

[35] This can be derived following more conventional approaches to defining the nested logit model, such as the definition in Train (2009).

The outside option has utility normalized to zero, so that its predicted market share is

$$s_0 = \frac{1}{\sum_g \left(\sum_{k \in \mathbb{J}_g} e^{\delta_k/(1-\sigma_g)}\right)^{1-\sigma_g}}. \tag{A.13}$$

Taking the difference between the natural logarithm of the predicted market share for alternative j and the natural logarithm of the predicted market share for the outside option yields

$$\ln(s_j) - \ln(s_0) = \delta_j/(1-\sigma_{g(j)}) - \sigma_{g(j)} \ln\left(\sum_{k \in \mathbb{J}_{g(j)}} e^{\delta_k/(1-\sigma_{g(k)})}\right). \tag{A.14}$$

Adding and subtracting $\sigma_{g(j)} \ln[e^{\delta_j/(1-\sigma_{g(j)})}]$ to the right-hand side of Appendix Equation (A.14) and substituting the definition of the predicted within-group share yields

$$\ln(s_j) - \ln(s_0) = \delta_j/(1-\sigma_{g(j)}) - \sigma_{g(j)}\delta_j/(1-\sigma_{g(j)}) + \sigma_{g(j)} \ln(s_{j|g}). \tag{A.15}$$

Combining like terms and making cancellations yields

$$\ln(s_j) - \ln(s_0) = \delta_j + \sigma_{g(j)} \ln(s_{j|g}). \tag{A.16}$$

Converting the term $\sigma_{g(j)} \ln(s_{j|g})$ into a summation with dummy variables yields Equation (20).

A.3 Derivation of Estimation Equation (26)

Given the assumption that ϵ_{ij} is i.i.d. type 1 extreme value, the predicted choice probability for decision maker i in demographic group d choosing alternative j and market share for demographic group d and alternative j is

$$s_{dj} = \frac{e^{\delta_j + \beta_{dg(j)}}}{\sum_k e^{\delta_k + \beta_{dg(k)}}}. \tag{A.17}$$

I assume that decision makers in each demographic group obtain utility equal to zero when selecting the outside option:

$$u_{d0} = \delta_0 + \beta_{d0} = 0. \tag{A.18}$$

Therefore, the outside good market share for demographic group d is

$$s_{d0} = \frac{1}{\sum_k e^{\delta_k + \beta_{dg(k)}}}. \tag{A.19}$$

Taking the difference of the natural logarithms of Appendix Equations (A.17) and (A.19) yields Equation (26).

A.4 Derivation of Estimation Equation (30)

Assuming that the error term in Equation (29) is i.i.d. type 1 extreme value, demographic group d's predicted market share for alternative j is

$$s_{dj} = \frac{e^{(\delta_j+\beta_{dg})/(1-\sigma)}}{\left(\sum_{k\in\mathbb{J}_{g(j)}} e^{(\delta_k+\beta_{dg}/(1-\sigma)}\right)^{\sigma} \sum_g \left(\sum_{k\in\mathbb{J}_g} e^{(\delta_k+\beta_{dg})/(1-\sigma)}\right)^{1-\sigma}}. \tag{A.20}$$

The within-group predicted market share is

$$s_{dj|g} = \frac{e^{(\delta_j+\beta_{dg})/(1-\sigma)}}{\sum_{k\in\mathbb{J}_{g(j)}} e^{(\delta_k+\beta_{dg})/(1-\sigma)}}. \tag{A.21}$$

For every demographic group, the outside option utility is normalized to zero:

$$\delta_0 + \beta_{d0} = 0. \tag{A.22}$$

Therefore, the predicted market share for the outside option is

$$s_{d0} = \frac{1}{\sum_g \left(\sum_{k\in\mathbb{J}_g} e^{(\delta_k+\beta_{dg})/(1-\sigma)}\right)^{1-\sigma}}. \tag{A.23}$$

Taking the difference of the natural logarithm of Equations (A.20) and (A.23) yields

$$\ln(s_{dj}) - \ln(s_{d0}) = (\delta_j + \beta_{dg})/(1-\sigma) - \sigma \ln\left(\sum_{k\in\mathbb{J}_{g(j)}} e^{(\delta_k+\beta_{dg})/(1-\sigma)}\right). \tag{A.24}$$

Adding and subtracting $\sigma \ln\left(e^{(\delta_k+\beta_{dg})/(1-\sigma)}\right)$ to the right-hand side of Appendix Equation (A.24) and substituting the definition of the within-group share from Equation (A.21) and making cancellations yields

$$\ln(s_{dj}) - \ln(s_{d0}) = (\delta_j + \beta_{dg})/(1-\sigma) - \sigma(\delta_j + \beta_{dg})/(1-\sigma) + \sigma\ln(s_{dj|g}). \tag{A.25}$$

This equation simplifies to estimation Equation (30).

A.5 Derivation of Three-Level Nested Logit Estimation Equation (42)

I adopt the presentation of the three-level nested logit model based on Brenkers and Verboven (2006). I maintain the same model notation from prior sections, so that j denotes alternatives and g denotes groups of alternatives. I denote subgroups by h, so that h is a subgroup of group g. The predicted share for this model can be expressed as the product of conditional probabilities:

$$s_j = s_{j|h} s_{h|g} s_g = \frac{e^{\delta_j/(1-\sigma_{hg})}}{e^{I_{hg}/(1-\sigma_{hg})}} \frac{e^{I_{hg}/(1-\sigma_g)}}{e^{I_g/(1-\sigma_g)}} \frac{e^{I_g}}{e^I}, \tag{A.26}$$

where

$$I_{hg} = (1-\sigma_{hg}) \ln \sum_{j=1}^{J_{hg}} e^{\delta_j/(1-\sigma_{hg})}, \tag{A.27}$$

$$I_g = (1-\sigma_g) \ln \sum_{h=1}^{H_g} e^{I_{hg}/(1-\sigma_g)}, \tag{A.28}$$

and

$$I = \ln \sum_{g=1}^{G} e^{I_g}. \tag{A.29}$$

The parameters σ_{hg} and σ_g measure within subgroup and within-group correlation of utility, respectively. The terms J_{hg}, H_g, and G denote the number of alternatives in subgroup h of group g, the number of subgroups in group g, and the number of groups, respectively. The conditional shares $s_{j|h}$ and $s_{h|g}$ represent the within subgroup h share of alternative j and the within-group g share of subgroup h, respectively. The share s_g denotes the group g share. I assume that the outside option is mean utility equal to zero, $\delta_0 = 0$, so that its predicted market share is

$$s_0 = \frac{1}{e^I}. \tag{A.30}$$

Taking the natural logarithm of both sides of Appendix Equation (A.26) and subtracting the natural logarithm of the outside good share yields

$$\ln(s_j) - \ln(s_0) = \frac{\delta_j}{1-\sigma_{hg}} + \frac{I_{hg}}{1-\sigma_g} + I_g - \frac{I_{hg}}{1-\sigma_{hg}} - \frac{I_g}{1-\sigma_g}. \tag{A.31}$$

Combining like terms gives

$$\ln(s_j) - \ln(s_0) = \frac{\delta_j}{1-\sigma_{hg}} + \frac{\sigma_g - \sigma_{hg}}{(1-\sigma_g)(1-\sigma_{hg})} I_{hg} - \frac{\sigma_g}{1-\sigma_g} I_g. \tag{A.32}$$

Adding and subtracting $\frac{\sigma_g}{1-\sigma_g} I_{hg}$ to and from the right-hand side of Appendix Equation (A.32) and substituting the definition of $\ln(s_{h|g})$ yields

$$\ln(s_j) - \ln(s_0) = \frac{\delta_j}{1-\sigma_{hg}} + \frac{\sigma_g - \sigma_{hg}}{(1-\sigma_g)(1-\sigma_{hg})} I_{hg} - \frac{\sigma_g}{1-\sigma_g} I_{hg} + \sigma_g \ln(s_{h|g}). \tag{A.33}$$

Combining like terms and simplifying gives

$$\ln(s_j) - \ln(s_0) = \frac{\delta_j}{1-\sigma_{hg}} + \frac{\sigma_{hg}}{1-\sigma_{hg}} I_{hg} + \sigma_g \ln(s_{h|g}). \tag{A.34}$$

Adding and subtracting $\frac{\sigma_{hg}}{1-\sigma_{hg}} \delta_j$ to and from the right-hand side of Appendix Equation (A.34) and substituting the definition of $\ln(s_{j|h})$ yields

$$\ln(s_j) - \ln(s_0) = \frac{\delta_j}{1-\sigma_{hg}} + \frac{\sigma_{hg}}{1-\sigma_{hg}} \delta_j + \sigma_{hg} \ln(s_{j|h}) + \sigma_g \ln(s_{h|g}). \tag{A.35}$$

This simplifies to

$$\ln(s_j) - \ln(s_0) = \delta_j + \sigma_{hg} \ln(s_{j|h}) + \sigma_g \ln(s_{h|g}). \tag{A.36}$$

A.6 Derivation of Equation (44)

This derivation follows closely the steps in Appendix Section (A.5). For the three-level nested logit model, the market share for alternative k belonging to subgroup h and group g conditional on the removal of alternative j from the choice set is

$$s_{k,j\notin\mathbb{J}} = s_{k|h,j\notin\mathbb{J}} s_{h|g,j\notin\mathbb{J}} s_{g,j\notin\mathbb{J}} = \frac{e^{\delta_k/(1-\sigma_{hg})}}{e^{I_{hg}/(1-\sigma_{hg})}} \frac{e^{I_{hg}/(1-\sigma_g)}}{e^{I_g/(1-\sigma_g)}} \frac{e^{I_g}}{e^{I}}, \tag{A.37}$$

where I_{hg},I_g and I are defined in Appendix Equations (A.27), (A.28), and (A.29), respectively, with the exception that alternative j is removed from the choice set. Taking the natural logarithm of both sides and subtracting the natural logarithm of the outside good share yields

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \frac{\delta_k}{1-\sigma_{hg}} + \frac{I_{hg}}{1-\sigma_g} + I_g - \frac{I_{hg}}{1-\sigma_{hg}} - \frac{I_g}{1-\sigma_g}. \tag{A.38}$$

Combining like terms gives

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \frac{\delta_k}{1-\sigma_{hg}} + \frac{\sigma_g - \sigma_{hg}}{(1-\sigma_g)(1-\sigma_{hg})} I_{hg} - \frac{\sigma_g}{1-\sigma_g} I_g. \tag{A.39}$$

Adding and subtracting $\frac{\sigma_g}{1-\sigma_g} I_{hg}$ to and from the right-hand side of Appendix Equation (A.39) and substituting the definition of $\ln(s_{h|g,j\notin\mathbb{J}})$ yields

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \frac{\delta_k}{1-\sigma_{hg}} + \frac{\sigma_g - \sigma_{hg}}{(1-\sigma_g)(1-\sigma_{hg})} I_{hg} - \frac{\sigma_g}{1-\sigma_g} I_{hg} + \sigma_g \ln(s_{h|g,j\notin\mathbb{J}}). \tag{A.40}$$

Combining like terms and simplifying gives

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \frac{\delta_k}{1-\sigma_{hg}} + \frac{\sigma_{hg}}{1-\sigma_{hg}} I_{hg} + \sigma_g \ln(s_{h|g,j\notin\mathbb{J}}). \tag{A.41}$$

Adding and subtracting $\frac{\sigma_{hg}}{1-\sigma_{hg}}\delta_k$ to and from the right-hand side of Appendix Equation (A.41) and substituting the definition of $\ln(s_{k|h})$ yields

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \frac{\delta_k}{1-\sigma_{hg}} + \frac{\sigma_{hg}}{1-\sigma_{hg}}\delta_k + \sigma_{hg} \ln(s_{k|h,j\notin\mathbb{J}}) + \sigma_g \ln(s_{h|g,j\notin\mathbb{J}}). \tag{A.42}$$

This simplifies to

$$\ln(s_{k,j\notin\mathbb{J}}) - \ln(s_{0,j\notin\mathbb{J}}) = \delta_k + \sigma_{hg}\ln(s_{j|h,j\notin\mathbb{J}}) + \sigma_g \ln(s_{h|g,j\notin\mathbb{J}}). \tag{A.43}$$

Substituting $\delta_k = \ln(s_k) - \ln(s_0) - \sigma_{hg}\ln(s_{k|h}) - \sigma_g \ln(s_{h|g})$ (which is the alternative k re-arranged version of Appendix Equation (A.36)), factoring common terms and re-arranging yields Equation (44).

A.7 Further Details on Data Used to Estimate Vehicle Demand

Given the highly disaggregated definition of a vehicle, the vehicle sample after merging the data sets is 1,413. I take several steps to purge the data of observations that may bias demand coefficient estimates. First, I drop extremely expensive vehicles that have a transaction price exceeding $100,000. This drops 82 observations from the data. I then drop observations that have more than 50 percent of sales that are for a prior or future model year. For example, a vehicle sold during the 2015 market year can include 2014 and 2016 versions. These are usually sold at highly discounted prices to clear out inventory for the current model year version. In some cases, most or all of the sales of a model are from the prior model year, sometimes due to the model being discontinued. To prevent any inventory effects biasing the demand coefficient estimates, I limit the sample to models that have a majority of sales from the same model year, i.e., 2015 model year versions. This drops 520 observations, leaving 811. A small number of remaining observations are dropped due to missing transaction price or second choice data. I limit the sample to vehicles with at least 100 sales, and those that have sufficient observations for constructing instrumental variables. This drops an additional 35 observations. The final observation count is 762.

Table A.1: Data Cleaning

Reason for dropping observations	Observations dropped	Remaining observations
Transaction price exceeding $100,000	82	1,331
Over 50% of sales are for a prior or future model year	520	811
Missing transaction price or second choice data	14	797
Fewer than 100 sales	26	771
Insufficient observations for constructing instrumental variables	9	762

Notes: The final sample size is 762. The initial sample size is 1,431.

A.8 Derivation of Own-Price Elasticity of Demand

In this section of the appendix, I derive a closed-form expression for the own-price elasticity of demand. The average own-price elasticity of demand is equal to

$$\varepsilon_{own-price} = \frac{1}{J}\sum_j \frac{dq_j}{dp_j}\frac{p_j}{q_j}, \tag{A.44}$$

where q_j denotes sales of vehicle j. Sales of vehicle j are $q_j = Ns_j$, where N is the number of new and used vehicle buyers, i.e., the market size. Sales are given by

$$q_j = N\frac{e^{\delta_j/(1-\sigma)}}{\left[\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right]^{\sigma} + \sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}}. \tag{A.45}$$

Differentiating q_j with respect to price p_j yields

$$\begin{aligned}\frac{dq_j}{dp_j} = {} & N\frac{e^{\delta_j/(1-\sigma)}\frac{d\delta_j}{dp_j}/(1-\sigma)}{\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma} + \sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}} \\ & - N\frac{e^{\delta_j/(1-\sigma)}\left[\sigma\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma-1} e^{\delta_j/(1-\sigma)}\frac{d\delta_j}{dp_j}/(1-\sigma) + e^{\delta_j/(1-\sigma)}\frac{d\delta_j}{dp_j}/(1-\sigma)\right]}{\left[\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma} + \sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right]^2}.\end{aligned}$$

This simplifies to

$$\frac{dq_j}{dp_j} = N\left[s_j\frac{d\delta_j}{dp_j}/(1-\sigma) - s_j\frac{d\delta_j}{dp_j}/(1-\sigma)s_j\left(\sigma\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma-1} + 1\right)\right]. \tag{A.46}$$

Factoring common terms yields

$$\frac{dq_j}{dp_j} = Ns_j\frac{d\delta_j}{dp_j}/(1-\sigma)\left[1 - s_j\left(\sigma\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma-1} + 1\right)\right]. \tag{A.47}$$

Substituting this expression into Appendix Equation (A.44) and making cancellations yields

$$\varepsilon_{own-price} = \frac{1}{J}\sum_j \frac{\beta_{\ln(Price)}}{1-\sigma}\left[1 - s_j\left(\sigma\left(\sum_{k\in J_{g(j)}} e^{\delta_k/(1-\sigma)}\right)^{\sigma-1} + 1\right)\right], \tag{A.48}$$

where $\beta_{\ln(Price)}$ is the preference coefficient for the log of purchase price.

A.9 Derivation of Total New Vehicle Market Price Elasticity of Demand

Total new vehicle sales are

$$\sum_{j=1}^{J} q_j = N(1 - s_0), \tag{A.49}$$

where N is the market size and where s_0 is the share of the outside option. The total new vehicle market price elasticity of demand is

$$\varepsilon_{total-price} = \sum_{j=1}^{J} \frac{dq_j}{dp_j}\frac{p_j}{q_j} = \sum_{j=1}^{J} \frac{d(1-s_0)}{dp_j}\frac{p_j}{(1-s_0)}. \tag{A.50}$$

Evaluating $\frac{d(1-s_0)}{dp_j}$ yields

$$\frac{d(1-s_0)}{dp_j} = s_0^2(1-\sigma)\left[\sum_{j=1}^{J} e^{\delta_j/(1-\sigma)}\right]^{-\sigma} e^{\delta_j/(1-\sigma)}\frac{\beta_{\ln(Price)}}{(1-\sigma)p_j}, \tag{A.51}$$

where $\beta_{\ln(Price)}$ is the preference coefficient for the log of purchase price. Substituting Appendix Equation (A.51) into Appendix Equation (A.50) and simplifying yields

$$\varepsilon_{total-price} = \frac{s_0^2}{1-s_0}\beta_{\ln(Price)}\left[\sum_{j=1}^{J} e^{\delta_j/(1-\sigma)}\right]^{1-\sigma}. \tag{A.52}$$