Economics Letters 115 (2012) 44–48

Contents lists available at SciVerse ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias[☆]

Antonio M. Bento[a,*], Shanjun Li[a,1], Kevin Roth[b]

[a] Cornell University, Charles H. Dyson School of Applied Economics and Management, 424 Warren Hall, Ithaca, NY 14853, USA
[b] Department of Economics, 404 Uris Hall, Ithaca, NY 14853, USA

ARTICLE INFO

Article history:
Received 21 November 2010
Received in revised form
13 September 2011
Accepted 23 September 2011
Available online 19 October 2011

JEL classification:
Q48
L91

Keywords:
Energy paradox
Fuel economy
Consumer heterogeneity

ABSTRACT

From the previous literature, it can be found that consumers tend to undervalue discounted future energy costs in their purchase decisions for energy-using durables. We show that this finding could, in part, result from ignoring consumer heterogeneity in empirical analyses as opposed to true undervaluation.

1. Introduction

Although economic theory suggests that rational consumers should be willing to pay $1.00 more for a vehicle that saves them $1.00 in discounted future fuel costs, a growing body of literature finds a marginal willingness-to-pay (MWTP) for reduced discounted future fuel costs ranging from $0.35 to $0.79 (Helfand and Wolverton, 2010; Greene, 2010). This perceived undervaluation of future fuel costs is an example of an *energy paradox* in the automobile market. The energy paradox is a general concept used to explain the unexpectedly slow diffusion of apparently cost-effective, energy-efficient technologies that involve similar trade-offs between up-front capital costs and future operating costs (Jaffe and Stavins, 1994). Such a paradox may exist in automobiles and other energy-using durables (Hausman, 1979).[2]

With increasing concerns related to climate change, energy security, and local pollution, many have used this potential market failure to justify policies that promote efficiency-improving technology. Policies that encourage even a small correction in this paradox have the potential to result in sizable decreases in energy use and its related externalities. The magnitude and sources of this paradox have broad implications for any technology that uses energy. In the automobile sector, some of this interest has focused on the debate between the gasoline tax and corporate average fuel economy (CAFE) standards, as well as the design of future CAFE standards. Precise estimates of the MWTP for reduced discounted future fuel costs are central to this debate (Parry et al., 2010). If consumers correctly value future fuel costs, gasoline taxes are found to be less costly than CAFE standards in achieving the targeted fuel reductions (Fischer et al., 2007; Jacobsen, 2010). However, the opposite is true if consumer undervaluation is sufficiently large. Although we focus on the automobile market, this note has implications for the valuation of energy efficiency in a very broad category of purchases.

Our concern with prior literature is that it has often examined the energy paradox ignoring the underlying consumer heterogeneity in MWTP for future reductions in fuel costs. If consumers are heterogeneous in their MWTP, they will be sorted into vehicles based on vehicle fuel efficiency: those with high MWTP for reduced fuel costs will be sorted into fuel-efficient vehicles and those with low MWTP will be sorted into fuel-inefficient ones. We show in

☆ Authors thank Ashley Langer, Lung-fei Lee, Joshua Linn, Ian Parry, Kenneth Small and James Sallee for very helpful comments.

* Corresponding author. Tel.: +1 607 255 0626.
E-mail addresses: amb396@cornell.edu (A.M. Bento), sl2448@cornell.edu (S. Li), kdr42@cornell.edu (K. Roth).

1 Tel.: +1 607 255 1832.

2 The typical magnitude of the energy paradox in appliances requires discount rates of 20% although estimates vary widely (see Hausman, 1979, Dubin and McFadden, 1984, Ruderman et al., 1987).

0165-1765/$ – see front matter © 2012 Published by Elsevier B.V.
doi:10.1016/j.econlet.2011.09.034

this paper that ignoring consumer heterogeneity in the MWTP for future fuel cost in a (multinomial) logit specification could result in heteroskedasticity and bias the estimate of the MWTP toward zero, suggesting spurious undervaluation. The purpose of this note is not to argue whether there is undervaluation of fuel economy or not. Rather our point is that an empirical analysis that ignores consumer heterogeneity may overstate the magnitude of undervaluation. Similar concerns of bias due to sorting were raised in a recent study of the value of a statistical life using labor market data (Deleire et al., 2009).

In Section 2, we analytically illustrate the potential for bias from ignoring consumer preference heterogeneity in future fuel cost in a simplified multinomial logit framework. In Section 3, we provide further evidence with simulations in a richer model of vehicle demand. In doing the simulations, we first generate data from an equilibrium model of the automobile market and then try to recover the average MTWP for fuel cost using a logit model and a random coefficient logit model.[3] Our analysis shows that, when undervaluation of fuel costs is not present in the data-generating mechanism, the logit model could erroneously suggest significant undervaluation, whereas the random coefficients logit model recovers the true average MWTP.

2. Bias analysis from ignoring preference heterogeneity

In the context of vehicle demand, we assume that each consumer chooses to buy a new vehicle, from among J models or products, or not to make any purchase (labeled *choosing the outside good*) in a given period. For ease of exposition, we assume that the utility of consumer i from vehicle choice j only depends on a single dimension of vehicle characteristics, fuel cost (fc). We relax this assumption in the simulations below. The utility of consumer i from vehicle j is

$$u_{ij} = \beta_i fc_j + \varepsilon_{ij}, \tag{1}$$

where the heterogeneous preference β_i has a mean $\bar{\beta}$ and variance σ_β. ε_{ij} has an i.i.d. type I extreme value distribution (conditional on fc_j) with a variance of $\sigma_\varepsilon = \pi^2/6$. The utility function can be rewritten as:

$$u_{ij} = \bar{\beta} fc_j + \tilde{\beta}_i fc_j + \varepsilon_{ij} = \bar{\beta} fc_j + e_{ij}, \tag{2}$$

where the variance of the composite error e_{ij}, $var(e_{ij} \mid fc_j) = \sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2$. Because e_{ij} exhibits heteroskedasticity, the underlying i.i.d. assumption of e_{ij} would be violated if one is to estimate $\bar{\beta}$ using the multinomial logit. To analyze bias of the parameters estimated from the multinomial logit model due to heteroskedasticity, we scale the utility and rewrite (2) as

$$\frac{u_{ij}\sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} = \frac{\bar{\beta} fc_j \sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} + \frac{e_{ij}\sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}}. \tag{3}$$

$$\begin{aligned}
\tilde{u}_{ij} &= \frac{\bar{\beta} fc_j \sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} + \tilde{e}_{ij} \\
&= \bar{\beta} fc_j + \left[\frac{\bar{\beta} fc_j \sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} - \bar{\beta} fc_j \right] + \tilde{e}_{ij} \\
&= \bar{\beta} fc_j + \bar{\beta} fc_j \left[\frac{\sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} - 1 \right] + \tilde{e}_{ij}.
\end{aligned} \tag{4}$$

[3] Given that consumers have multiple vehicle models from which to choose, the empirical methods are multinomial logit models, but we suppress the word multinomial to save space throughout our paper.

The normalization makes $\tilde{e}_{ij}$ homoscedastic with a variance of $\pi^2/6$. Eq. (4) has re-casted the issue of heteroskedasticity to that of omitted variable in a discrete choice model: estimating $\bar{\beta}$ using the multinomial logit model ignoring heteroskedasticity based on Eq. (2) leads to the same problem as estimating $\bar{\beta}$ based on the last line in Eq. (4) while ignoring $z_j = \bar{\beta} fc_j [\frac{\sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}} - 1]$. The omitted variable z_j is positive as long as σ_β^2 is not equal to zero (assuming $\bar{\beta}$ to be negative). It is positively correlated with fc_j. Following Lee (1982) and Yatchew and Griliches (1985) which analyze the omitted variable bias in discrete choice models, the estimate of $\bar{\beta}$ from the multinomial logit model would be biased upward (toward zero). Moreover, a larger σ_β^2 implies a smaller $\frac{\sigma_\varepsilon}{\sqrt{\sigma_\beta^2 fc_j^2 + \sigma_\varepsilon^2}}$ (closer to zero), and a stronger correlation between z_j and fc_j, therefore, the bias would be larger as well.

The following two issues would add complications to the above analysis. First, the omitted variable introduces possible misspecification into the multinomial logit model. Second, a utility function with more vehicle characteristics as defined in Eq. (6) in the next section adds confounding factors to the bias in $\bar{\beta}$ when other vehicle characteristics are correlated with the fuel cost variable. Nevertheless, the bias analyzed above is likely to be the dominant issue in addressing our research question. Because the bias does not have a closed-form solution due to the nonlinear nature of the model even in the simple specification analyzed above, next we use Monte Carlo simulations that incorporate those two issues to provide further support.

3. Simulations

3.1. The equilibrium model of automobile market

The equilibrium model for data generation is composed of a demand side and a supply side. In the demand side, the utility of consumer i from vehicle j is defined as

$$u_{ij} = \alpha_i p_j + \beta_i fc_{ij} + \gamma_i x_j + \varepsilon_{ij} \tag{5}$$

where α_i, β_i, and γ_i are individual-specific taste parameters. We define $\theta = \{\alpha_i, \beta_i, \gamma_i\}.p_j$ is price of model j. fc_{ij} is the present value of the total expected discounted fuel cost of the vehicle; it is defined by

$$fc_{ij} = \sum_{t=0}^{T_j} \delta_i^{t*} AVMT_{itj}^{*} gp_{it}^{e} / MPG_j \tag{6}$$

where T_j is the expected lifetime of vehicle model j, δ_i is an individual-specific discount factor, $AVMT_{itj}$ is annual vehicle miles of travel in year t (which usually decreases with the vehicle's age), and gp_{it}^e is the expected gasoline price at year t of consumer i. Heterogeneity can arise from any of the elements used to calculate the lifetime fuel cost of the vehicle. x_j is a vector of other vehicle attributes, and ε_{ij} is assumed to have a type I extreme value distribution. We normalize the utility from the outside good, u_{i0}, to zero. The probability of household i choosing vehicle j is given by

$$P_{ij} = \frac{\exp(\bar{u}_{ij})}{1 + \sum_h \exp(\bar{u}_{ih})} \tag{7}$$

where $\bar{u}_{ij} = u_{ij} - \varepsilon_{ij}$. Given individual choice probabilities, the aggregate demand can be obtained through summation.

The supply side is composed of several firms, each producing multiple vehicle models. They engage in Bertrand competition in that each firm chooses prices to maximize its total profit in a given year, taking the products available as fixed. Following the literature, we assume that the marginal cost of each product is constant. The total profit of firm f is

$$\pi^f = \sum_{f \in F} [(p_j - mc_j) q_j(p, \theta)] \quad (8)$$

where F is the set of all products produced by firm f, mc_j is the marginal cost, and q_j is the aggregate demand. p is the price vector and it is obtained through the first-order conditions in equilibrium

$$p = mc + \Delta^{-1} q(p, \theta) \quad (9)$$

where the element of Δ, Δ_{jr} is zero if j and r are produced by different firms. Otherwise, it is equal to $-\partial q_r / \partial p_j$. Given the underlying preference parameters and marginal cost, this equation can be used to compute equilibrium prices and sales.

3.2. Data generation

Through our data generation approach, we aim to mimic the US auto market. Vehicle information comes from the 2001 Ward's Automotive Yearbook; vehicle characteristics include miles per gallon (MPG), horsepower, weight, and manufacturer. We construct marginal cost, a function of MPG, horsepower, and weight, for each model based on estimates from Berry et al. (1996).[4] We randomly choose a set of vehicle models (25 in the baseline simulation) and assume that these models are available in each year from 2001 to 2006, the time span for our analysis.

For ease of exposition, we make several demand-side assumptions. For preference parameters, we assume that all consumers have the same preference on all characteristics except fuel costs. In calculating fuel costs, we assume that the discount factor δ, annual vehicle miles of travel $AVMT$, and expected gasoline price gp^e are all constant across consumers for any given vehicle. We assume a 10% yearly discount rate. Vehicle lifetime and age-specific annual miles of travel for passenger cars and light trucks are from Lu (2006). We further assume that expected gasoline prices during a vehicle's lifetime are equal to current annual gasoline price (i.e., gasoline price follows a random walk (Anderson et al., 2010)). Annual gasoline prices during 2001–2006 are from the Energy Information Administration. These simplifying assumptions, innocuous for our conclusion, imply that consumer heterogeneity is manifested only through the consumer-specific taste parameter on fuel cost, β_i. In the baseline simulation, we assume that β_i has a uniform distribution; the range of the distribution affects the degree of consumer heterogeneity. We choose two levels of dispersion for the taste parameter $[-4, 0]$, and $[-3, -1]$. Anderson et al. (2010) using survey data examine the dispersion of predicted gasoline price, defined as the standard deviation of the predictions divided by the mean. They find this dispersion ranges from 30% to 60% in recent years, which roughly correspond to the dispersion of the two uniform distributions, noting that our distribution assumptions are different. Heterogeneity on discount rates, VMT and vehicle lifetime will further increase the dispersion on the parameter, β_i.

We generate data in two steps. First, we generate equilibrium prices for each model, assuming the whole market with 50,000 consumers in each year. Second, based on equilibrium prices, we generate vehicle choices for 20,000 consumers in each year. The choices of these consumers as well as equilibrium prices are taken as data for the estimation.

[4] We also add a random error term to the marginal cost of each attribute and to the marginal cost of each product based on the standard errors estimated by Berry et al. (1996). All costs are converted to 2001 dollars.

3.3. Estimation

The goal of the estimation is to recover the underlying preference parameters and to obtain consumers' MWTP for reduced fuel costs. For ease of exposition, we assume that the econometrician observes all vehicle characteristics relevant to consumers.[5] We employ two methods: a logit model and a random coefficients logit model. The logit model is estimated using the standard maximum likelihood method. As discussed in Train (2003), the appeal of the random coefficients model comes from its ability to incorporate unobserved consumer heterogeneity, which in our context avoids sorting bias. This model is estimated using the simulated maximum likelihood method. To conduct numerical integration in the simulated method, we employ Halton sequences, which are more efficient than direct Monte Carlo sampling.

3.4. Results

We find three main results from the Monte Carlo analysis.

Result 1. *In the presence of heterogeneity, the logit model suggests undervaluation of the MWTP for reduced future fuel costs, even when undervaluation is not present in the data.*

Support. Panel A in Table 1 shows that consumers undervalue fuel costs by 29%. The parameter estimates on vehicle price and fuel cost implies that consumers are only willing to pay \$0.71 for a \$1.00 reduction in discounted future fuel costs. The bias comes from individuals sorting into vehicles based on their MWTP: those very averse to fuel costs (e.g., with very negative MWTP) purchase vehicles with low fuel costs. The correlation between fuel cost and the average MWTP among consumers who purchase corresponding vehicles is depicted on the left panel of Fig. 1 (the correlation coefficient is 0.83). We believe that at least part of the undervaluation found in prior literature could be attributable to this type of sorting bias.

Result 2. *The random coefficients logit model correctly identifies the MWTP.*

Support. Table 1, Panel A shows that, by explicitly modeling consumer heterogeneity, the random coefficients logit model is able to recover the underlying parameters on vehicle price and fuel cost. The implied MWTP is -1, indicating that consumers are willing to pay \$1.00 for a \$1.00 reduction in discounted future fuel costs, consistent with our model assumption.

Result 3. *The greater the heterogeneity, the larger the bias from the logit model.*

Support. The underlying data-generating process in Panel A of Table 1 implies twice the heterogeneity of Panel B. As a consequence, the undervaluation for the logit model in Panel A, 29%, is larger than the 10% undervaluation in Panel B.

Table 2 presents Monte Carlo results for alternative specifications. Panel A suggests that increased market power magnifies the bias from the logit model, with the undervaluation going to 37% from 29% in the baseline model in Table 1. Increasing the number of vehicle draws (Table 2, Panel B) slightly decreases the undervaluation from 29% to 27%. The three findings discussed above still hold when the distribution of MWTP takes a log-normal distribution (Table 2, Panel C).

[5] In real applications, it is important to control for unobserved product attributes. Most recent literature on the energy paradox has explicitly dealt with this issue.

Table 1
Monte Carlo results.

	True	Estimates			
		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Panel A: baseline model					
Constant	1	0.60	0.05	1.05	0.07
Price	−2	−2.02	0.01	−2.00	0.01
Fuel cost	−2	−1.43	0.03	−2.01	0.07
Weight	4	4.49	0.15	3.83	0.17
Horsepower	8	7.68	0.14	8.18	0.15
Sigma[a]	4			4.18	0.26
Log-likelihood		228,335		228,268	
Implied valuation for $1 drop in fuel cost		$0.71		$1.00	
Implied undervaluation		29%			
Panel B: smaller heterogeneity					
Constant	1	0.93	0.05	1.08	0.06
Price	−2	−2.01	0.01	−2.01	0.01
Fuel cost	−2	−1.82	0.03	−2.03	0.06
Weight	4	3.99	0.15	3.80	0.16
Horsepower	8	8.05	0.14	8.21	0.14
Sigma[a]	2			2.31	0.31
Log-likelihood		225,942		225,933	
Implied valuation for $1 drop in fuel cost		$0.90		$1.01	
Implied undervaluation		10%			

[a] Sigma measures the degree of heterogeneity for MWTP for fuel cost. The value of sigma is multiplied by random draws from a uniform distribution [−0.5, 0.5]. In Panel A, the range of the MWTP for fuel cost is [−2, 0], whereas in Panel B it is [−1.5, −0.5].

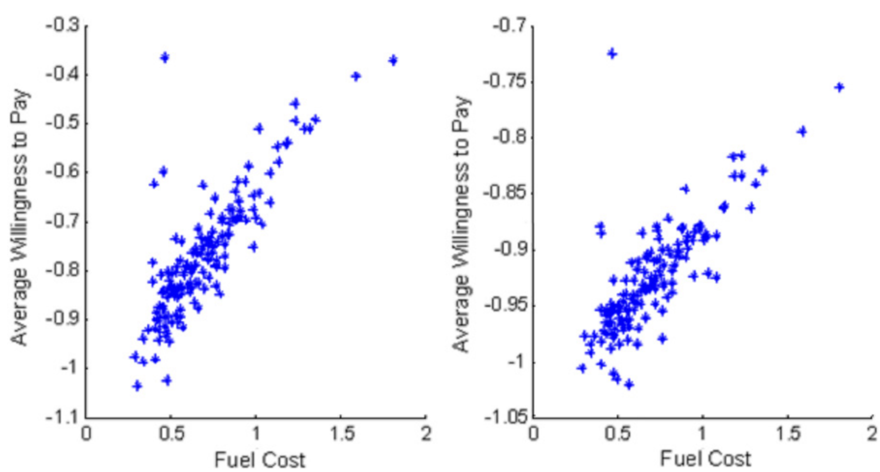

Notes: Figure 1 plots the average MWTP for reduced future fuel cost among consumers who purchase vehicles with a given fuel cost. Fuel cost on the *x*-axis is the lifetime discounted fuel cost divided by 10,000. The left figure corresponds to Panel A in Table 1 where the MWTP has a uniform distribution [–2,0] while the right figure corresponds to Panel B where the MWTP has a uniform distribution [–1.5, –0.5]

Fig. 1. Fuel cost and average marginal willingness to pay among buyers.

4. Conclusion

Our analysis shows that, if not accounted for, unobserved consumer heterogeneity can significantly affect the estimated MWTP for discounted future fuel costs. We believe that this may partly explain consumer undervaluation of future fuel costs and the wide range of estimates found in the literature. Here we have modeled consumer heterogeneity through the valuation of fuel economy. However, this is only one of many potential ways of representing consumer heterogeneity. For example, the heterogeneity could also arise from expected future fuel costs. While ignoring this source of heterogeneity would create a similar bias as the one identified in this paper, the implications for policy (whether or not there are consumer mistakes that constitute market failure) may be different. To properly evaluate the existence, source and magnitude of the energy paradox, further econometric analysis that explicitly models consumer heterogeneity from multiple sources by using random coefficient models in either a discrete choice or hedonic framework (e.g. Berry et al., 1995; Bajari and Benkard, 2005), are needed.

Table 2
Robustness checks.

	True	Estimates			
		Logit		Random coef. logit	
		Para	S.E.	Para	S.E.
Panel A: monopoly instead of oligopoly					
Constant	1	0.62	0.06	1.10	0.09
Price	−2	−2.02	0.02	−2.01	0.02
Fuel cost	−2	−1.27	0.03	−2.14	0.14
Weight	4	4.38	0.19	3.80	0.21
Horsepower	8	7.86	0.18	8.32	0.19
Sigma	4			4.55	0.42
Log-likelihood		158,480		158,437	
Implied valuation for $1 drop in fuel cost		$0.63		$1.07	
Implied undervaluation		37%			
Panel B: 50 vehicle models instead of 25					
Constant	1	0.62	0.03	1.00	0.05
Price	−2	−2.02	0.01	−2.00	0.01
Fuel cost	−2	−1.48	0.03	−1.99	0.06
Weight	4	4.18	0.12	3.99	0.13
Horsepower	8	7.75	0.11	8.01	0.11
Sigma	4			4.01	0.21
Log-likelihood		326,436		326,351	
Implied valuation for $1 drop in fuel cost		$0.73		$1.01	$0.99
Implied undervaluation		27%			
Panel C: lognormal MWTP instead of uniform distribution					
Constant	1	0.74	0.07	0.94	0.10
Price	−2	−2.01	0.02	−2.00	0.02
Fuel cost		−1.65	0.04	N/A	
Weight	4	4.38	0.21	4.21	0.21
Horsepower	8	7.78	0.19	7.96	0.20
Mean of underlying normal distribution	0.57			0.59	0.05
Sigma of underlying normal distribution	0.50			0.45	0.09
Implied mean of the lognormal distribution	2			−2.01	
Log-likelihood		113,980		113,976	
Implied valuation for $1 drop in fuel cost		$0.82		$1.00	
Implied undervaluation		18%			

References

Anderson, S., Kellogg, R., Sallee, J., 2010. What Do Consumers Believe About Future Gasoline Prices? NBER Working Paper #16974.

Bajari, P., Benkard, L., 2005. Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach. Journal of Political Economy 113, 1239–1276.

Berry, S., Kortum, S., Pakes, A., 1996. Environmental change and hedonic cost functions for automobiles. Proceedings of the National Academy of Sciences 93, 12731–12738.

Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. Econometrica 63, 841–890.

Deleire, T., Khan, S., Timmins, C., 2009. Roy Model Sorting and Non-Random Selection in the Valuation of a Statistical Life. Duke University Working Paper.

Dubin, J.A., McFadden, D.L., 1984. An econometric analysis of residential electric appliance holdings and consumption. Econometrica 52 (2), 345–362.

Fischer, C., Harrington, W., Parry, I., 2007. Do market failures justify tightening corporate average fuel economy (CAFE) standards? The Energy Journal 28 (4), 1–30.

Greene, D., 2010. How Consumers Value Fuel Economy: A Literature Review. Office of Transportation and Air Quality, US EPA. Report EPA-420-R-10-008.

Hausman, J.A., 1979. Individual discount rates and the purchase and utilization of energy-using durables. The Bell Journal of Economics 10, 33–54.

Helfand, G., Wolverton, A., 2010. Evaluating the Consumer Response to Fuel Economy: A Review of the Literature. US EPA.

Jacobsen, M., 2010. Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity. University of California — San Diego Working Paper.

Jaffe, A., Stavins, R., 1994. The energy paradox and the diffusion of conservation technology. Resource and Energy Economics 16 (2), 91–122.

Lee, Lung-Fei, 1982. Specification error in multinomial logit models. Journal of Econometrics 20, 197–209.

Lu, S., 2006. Vehicle Scrappage and Travel Mileage Schedules. National Highway Traffic Safety Administration Technical Report.

Parry, I., Evans, D., Oates, W., 2010. Do Energy Efficiency Standards Increase Social Welfare? Working Paper.

Ruderman, H., Levine, M.D., McMahon, J.E., 1987. The behavior of the market for energy efficiency in residential appliances including heating and cooling equipment. Energy Journal 8 (1), 101–124.

Train, K., 2003. Discrete Choice Methods with Simulation. Cambridge University Press.

Yatchew, A., Griliches, Z., 1985. Specification error in probit models. The Review of Economics and Statistics 67, 134–139.