Refining Equilibria in Counterfactual Simulations: An Application to Corporate Average Fuel Economy Standards

Hunt Allcott and Erich Muehlegger Harvard University

July 16, 2008

Abstract

We simulate how automotive manufacturing firms will respond to the more stringent Corporate Average Fuel Economy standard passed in 2007. We use Bayesian methods to estimate a random coefficients discrete choice model of new automobile demand, introducing a new procedure to exploit both individual-level and market-level data. We then simulate the game in which automakers set prices and characteristics in Nash equilibrium, maximizing profits subject to demand and the new regulation. To select between multiple possible equilibria, we employ a myopic partial best response algorithm, which we argue to be a reasonable stylized representation of new product development in the auto industry.

We thank Gary Chamberlain, Kelly Gallagher, Guido Imbens, Henry Lee, Sendhil Mullainathan, Ariel Pakes, and Robert Stavins for helpful suggestions. Funding for this project was provided by a grant from the U.S. Environmental Protection Agency though the Energy Technology Innovation Policy group at the Harvard Kennedy School of Government. Carmen Collyns, Sounman Hong, Ian Perry, and Jordon Ricks provided excellent research assistance.

1 Motivation

Perhaps the two principal occupations of the empiricist in Industrial Organization are demand system estimation and simulation of counterfactuals. Especially when simulating the behavior of multiproduct firms in differentiated products oligopolies, there is rarely a unique Nash equilibrium, and there is surprisingly little discussion of appropriate equilibrium selection devices. Crucial to such analyses are realistic substitution patterns from a demand model that includes unobserved consumer heterogeneity, but there is substantial recent concern about identification and robustness of estimation results using classical methods. This paper explores promising solutions to these two central problems in the context of one of today's most important policy issues, the strengthening of the Corporate Average Fuel Economy (CAFE) standards.

Ariel Pakes (2008) frames multiple equilibria in counterfactual analysis as one of the most "troubling" problems in the analysis of imperfectly competitive markets. This results from the analyst's fundamental inability to observe an industry well enough to impose a stronger selection mechanism. Our solution is thus to examine our agents' behavior - automakers developing new vehicles - and utilize a tractable equilibrium selection mechanism that captures that process in a stylized but realistic way. Our preferred mechanism is myopic partial best response, in which automakers "propose" new vehicles that best respond to other firms' proposals from the previous iteration; the eventual steady state actions are the equilibrium actions in our game. We illustrate the importance of the multiple equilibria problem by comparing our preferred equilibrium to Nash equilibria selected through other mechanisms.

Authors such as Chib and Greenberg (1998), Train (2003), and Athey and Imbens (2005) have argued that Bayesian estimation methods for random coefficients models have substantial advantages over classical maximization-based routines such as Generalized Method of Moments and Maximum Likelihood. Crucial for identification of this unobserved heterogeneity in any setting is variation in the choice set, which often can come from the addition of market-level cross-sections of aggregate quantities sold. This paper presents an algorithm, based on the standard Gibbs sampler for random coefficients multinomial logit, that allows Bayesian estimation with both micro and macro data. This technique views macro data simply as extensive microdata without individual covariates and nests data augmentation in the Gibbs sampler to impute the distribution of those covariates.

Our application is the response by automotive manufacturing firms to Corporate Average Fuel Economy standards, one of today's most important policy issues due to climate change, energy security concerns, and rising energy prices. In December 2007, President Bush signed the Energy Independence and Security Act, which set a target fuel economy of 35 miles per gallon for the fleet of vehicles sold in 2020. We study the most recent proposed implementation, which would require automakers to improve average fuel economy to 31.6 miles per gallon by 2015. Our applied research question is, "What will be the effects of the new CAFE standards on new vehicle characteristics and firm profits in 2015?" Our answer to this question requires us to address the above two methodological issues, which we believe are more broadly generalizable to applied problems in Industrial Organization.

Because of its economic and environmental importance and the unusual availability of consumer, product, and sales data, the automotive industry has been extensively studied. Papers such as Bresnahan (1987), Berry, Levinsohn, and Pakes (1995, "BLP"), Goldberg (1998), Petrin (2002), Berry, Levinsohn, and Pakes (2004, "MicroBLP"), and numerous others have made important methodological contributions and analyzed questions related to industry structure, entry and exit, or government regulation. An additional set of analyses are specifically focused on fuel economy policies, including Davis et al (1995), National Research Council (2002), Feng, Fullerton, and Gan (2005), Austin and Dinan (2005), Greene, *et al*, (2005), and the US Energy Information Administration's National Energy Modeling System (2007). Perhaps most notable are papers by Jacobsen (2008) and Bento, Goulder, Jacobsen, and von Haefen (2006), which analyze CAFE standards and gasoline taxes while simultaneously incorporating a number of methodological advances. In particular, they use a discrete-continuous model for vehicle choice and gasoline demand, estimate the shadow cost of the CAFE standard for constrained automakers, and simulate new and used car markets in equilibrium.

One of the most consistent and compelling stylized facts from this body of work is that the response to fuel economy regulation will depend heavily on the introduction of new models. Greene, et al, and Davis, et al, both find that approximately 90 percent of the change in average new vehicle fuel economy results from changes in vehicle attributes as opposed to changes in relative sales of different vehicle classes. Pakes, Berry, and Levinsohn (1993) find that the average fuel economy of vehicles within each of their six classes increased by five to 21 percent in response to the gasoline price shock of the early 1970s. In an extension to his simulations, Jacobsen shows that the welfare costs of the CAFE standard per gallon of gasoline saved are 64% lower if the model allows automakers to adjust the fuel economy of their vehicles along with the prices. It seems that any realistic analysis of fuel economy regulation must allow vehicle characteristics to be modified along a production function.

This desire for realism is at odds with what the present economic theory allows. Although Caplin and Nalebuff (1991) prove that there is a unique equilibrium in prices in differentiated products oligopolies with single-product firms, there may be multiple Nash equilibria in a game such as ours where multiproduct firms set prices and characteristics. In studies of the auto industry and elsewhere, this problem has rarely been addressed directly. Berry, Levinsohn, and Pakes (2004) do not allow competing firms to best respond in prices or attributes in their simulations of the exit of Oldsmobile and the introduction of the luxury Sport Utility Vehicle. In simulating responses to fuel economy regulation, Austin and Dinan (2005) choose an "equilibrium" that maximizes the sum of all automakers' profits. The US National Energy Modeling System models only two aggregated firms and does not require them to best respond in equilibrium. Jacobsen (2008), while adopting a counterfactual equilibrium computation approach that appears to be in principle similar to ours, highlights other methodological issues instead of this one.

Pakes (2008) and Lee and Pakes (2008) propose two types of solutions to the multiple equilibrium problem in applied counterfactual simulations. The first is to compute all possible equilibria, but in many applications the size of the action space may be too large or the set of equilibria too diverse to draw any qualitative conclusion. The second is to select an equilibrium using a learning mechanism such as those discussed in Fudenberg and Levine (1998) and the works cited therein; ideally, this selection mechanism would reflect agents' actual behavior.

We simulate a game in which the six largest automakers maximize profits in Nash equilibrium, subject to consumer demand and the CAFE regulation, by changing prices and attributes along an exogenous production function. An examination of the industry's product development process suggests that firms propose the important attributes of vehicles - powertrain technology, horsepower, size - several years before they are sold. These "proposals" are public - presented at auto shows and industry conferences - and iterative. The public and iterative nature of this process, combined with the computational tractability of myopia, motivates our use of a myopic partial best response algorithm to select the equilibrium set of attributes and prices.

Simulating equilibrium sales and markups requires an estimate of the new vehicle demand system and substitution patterns. Since BLP (1995), discrete choice models in characteristics space have typically generated more realistic substitution patterns by modeling unobserved heterogeneity in consumer preferences through random coefficients. More recent work, such as Knittel and Metaxoglou (2008), has shown that, because of flat regions of the objective function and local extrema, classical optimization-based methods for estimating random coefficients models are sensitive to different starting values and optimization methods. As Athey and Imbens (2005) also point out, classical asympotic confidence intervals have poor properties in finite samples with irregular likelihood functions.

Bayesian methods, whose posterior parameter estimates have fundamentally the same interpretation as classical parameters (Geweke 1989), are increasingly popular because they do not require maximization and are thus robust to local extrema. When applied to random coefficients discrete choice models, they are also often more computationally efficient, especially as the number of parameters rises. Since random coefficients are identified from variation in the choice set, however, it is often helpful to exploit more than a single cross section of microdata. This additional data often comes only in the form of repeated cross-sections of market-level quantities sold. Imbens and Lancaster (1994) and Hellerstein and Imbens (1999) show the usefulness of such additional data using classical estimation strategies, in settings other than discrete choice. Unfortunately, there is no Bayesian equivalent of MicroBLP (2004) that provides an algorithm for computing the posterior parameter distribution of a random coefficients discrete choice model using both micro and macro data.

Our new procedure follows very naturally from two Bayesian Markov Chain MonteCarlo (MCMC) techniques, the Gibbs sampler for multinomial logit random coefficients models (Allenby and Lenk 1994) and data augmentation (Tanner and Wong 1987). We pool microdata with observations lacking covariates which represent the choices of consumers in the aggregate data. In each iteration of the Gibbs sampler, we impute a distribution of the missing covariates conditional on a vehicle purchased using the marginal population covariate distribution and draws from the current distribution of utility parameters.

We find that, even absent new regulation, high gasoline prices create an incentive for automakers to improve vehicle fuel economy. In our base case, which omits the proposed CAFE standards, fleet fuel economy rises approximately 38 percent over the 2007 model-year, to approximately 31 miles per gallon. With the CAFE standards, fleet fuel economy rises to 34.7 miles per gallon. Although the regulations do bind for American automakers, we estimate that overall fleet fuel economy exceeds the 2015 target of 31.6 mpg.

Consistent with previous work, the vast majority of the change in fuel economy comes through vehicle fuel economy improvements, rather than consumer vehicle selection. In addition, we find that allowing firms to adjust the entire set of vehicle attributes introduces important variation in how firms comply with CAFE regulations. Even absent CAFE, increased gasoline prices incentivize Honda and Toyota to substantially increase fuel economy. While Honda and Toyota introduce technology to improve vehicle fuel efficiency, approximately half of the increase in fuel economy comes from decreasing vehicle power. While the American automakers also incorporate more fuel efficient technology, absent the proposed CAFE standards the automakers increase power as well as fuel economy, resulting in more modest fuel economy improvements. The addition of the proposed CAFE standards lead firms furthest from compliance to reallocate fuel efficiency gains from power to fuel economy. Sacrificing power for fuel economy brings the firms initially furthest from compliance either into or very close to compliance - Ford slightly misses its 2015 CAFE target (by 0.3 mpg). Interestingly, while trading off power for fuel economy does affect the attractiveness of the vehicle to consumers, it does not entail substantial additional investment in fuel efficiency technology, beyond that automakers choose to install absent CAFE regulations. This suggests that studies holding vehicle attributes constant may overestimate the cost of CAFE compliance, as well as the degree to which firms investment in new engine technology.

Subsequent to this motivation, our paper has three central sections. In Section 2, we estimate the auto market demand system using both micro and macro data. In Section 3, we detail our model of the new vehicle development game, including the production function estimate and the myopic partial best response algorithm. Section 4 presents our estimates of how CAFE standards will affect the characteristics sold and automakers' profits.

2 Demand

We eventually wish to simulate firms maximizing profits, which are the sum of the products of quantities sold and markups for each products. The demand system must therefore give reasonable estimates both of quantities sold and substitution patterns, which automakers use to set Bertrand markups. In this section, we thus must estimate the US new auto market demand system from micro and aggregate data, using the Gibbs sampler. In the next four subsections, we present the data, set up the indirect utility function and identifying assumptions, describe the estimation routine, and discuss results.

Our indirect utility function and the resulting likelihoods are standard, as are our identification assumptions. We depart from the traditional framework as we include aggregate data in our individual-level likelihood functions. Our approach can be thought of as a "glass is half empty" approach to macro data.

Typically, one approaches such data as an extra opportunity and modifies the estimation strategy by adding additional moment conditions. In our framework, we consider macrodata in the same Gibbs sampler with microdata, but we have a missing data problem: although we observe the choices of every individual in the population, we only observe covariates for those individuals in the microdata. We thus add a data augmentation routine within the Gibbs sampler to impute missing covariates for each observation from the macro data.

In many settings, there are limited cross sections of micro data but more frequent cross sections of market-level price and quantity data featuring additional variation in the choice set. In these cases, the econometrician wishes to identify utility function parameters on interactions between individual and product attributes primarily using microdata, and then use additional variation in aggregate data to help identify random coefficients.

The combined use of both micro and macro data is particularly helpful in our application. We will need to identify the mean and variance of heterogeneous demand parameters for fuel economy, weight, and horsepower, the attributes that firms will modify in our simulated game. Figures 1 and 2 show these three characteristics in three-dimensional space. Figure 1 illustrates that weight and the inverse of fuel economy are highly collinear, and Figure 2 shows that weight and horsepower have a similar, but less severe, problem. To aid the econometric model, utility will thus be a function of weight (in tons), horsepower per ton of weight, and dollars of gasoline cost per mile (the inverse of fuel economy times the gasoline price) per ton of weight. The addition of any new choice sets helps identify all the paramters, especially because there is little variation in the latter variable in any particular year.

A natural way of identifying demand for fuel economy is to exploit time series variation in gasoline prices[1]. A complementary analysis by Allcott (2008) uses both microdata and annual cross sections from 1976 to 2007 of prices, quantities, and attributes of all new and used vehicles on the road to estimate the auto market's "implicit discount rate" for future fuel expenditures, in the spirit of Hausman (1979). By modeling the new and used car markets in equilibrium over several different gasoline price shocks, that complementary paper generates an even more credible estimate of the utility function parameter on fuel economy. This will later be incorporated into the present analysis; for the moment we model only the new vehicle market and two years of additional macro data.

2.1 Data

There are especially rich sources of data on the automotive industry. In 1983, 1990, 1995, 2001, and 2008, the US Energy Information Administration administered the National Household Travel Survey (NHTS) - or its predecessor, the National Personal Transportation Survey - a nationally-representative survey of approximately 25,000 households. Market level data dating to the 1960s,

[1] Although there is cross-sectional variation in gas prices across cities and states in the U.S., this is likely to be correlated with unobserved factors.

including prices, quantities, and product attributes, is available from multiple sources.

Unfortunately, the microdata cross sections do not offer substantial variation in the choice sets, especially in terms of fuel cost in dollars per mile. The 1983 NHTS has fewer data points and is considered to be anomalous, the 2008 data are still being gathered, and there is little variation in gas prices between 1990, 1995, and 2001. Using market-level data for years with gasoline price variation but no microdata is thus potentially helpful for our identification. In this analysis, we use only the 2001 and 2005 quantities sold for new cars and trucks, from market research company R.L. Polk; even this one additional year generates useful additional variation in the choice set.

Our microdata is taken from the 2001 NHTS, the most recent public source of individual-level data on household characteristics and vehicle ownership. We observe 26 thousand households in the nationally representative core sample. Each household reported total annual income and urban/suburban/rural status, the age, gender, and education of each household member, and the make, model, model year, Vehicle-Miles Traveled, and primary driver for each vehicle owned. In total, the households in the sample own 53 thousand vehicles. We model a choice occasion for each household member 16 years or older[2]. Since we wish to model one year of a new vehicle market, we consider all individuals who had purchased a new vehicle within 12 months of being surveyed; all other individuals are considered to have chosen the outside option of no new vehicle purchase. Since the NHTS was administered over a period of months, there were a total of 741 separate "new" models observed, of model years 2000, 2001, and 2002. Table I shows individual-level descriptive statistics for the NHTS microdata as well as the 2001 Current Population Survey (CPS), which along with the 2005 CPS will be used in our econometric procedure.

For each vehicle, we take fuel economy, horsepower, and weight data from Ward's Automotive Yearbook. Importantly, retail prices average 15 percent below list prices, and that discount varies by manufacturer. We use data from CNW Marketing, an auto industry research firm, to transform list prices into average purchase prices. Table 2 shows descriptive statistics for vehicles observed in the microdata and in the market-level data. Using this data, we can now continue to a description of our model.

2.2 The Model and Identification

Our indirect utility function and the derivation of the individual's likelihood is entirely standard. We model consumer i's indirect utility from vehicle j as:

$$V_{ij} = \beta_i X_{ij} + \eta P_j + \xi_j + \varepsilon_{ij} \tag{1}$$

[2]Because a small number of households have more vehicles than choice occasions, we may instead follow Jacobsen (2008) by modeling households with a number of choice occasions equal to the number of members over 16, plus one.

Where:
β_i = Random utility parameters, $\beta_i \sim N(b, W)$
X_{ij} = Interactions of individual and vehicle characteristics
η = Price coefficient
P_j = Log of vehicle price
ξ_j = Vehicle-specific unobservable characteristic
ε_{ij} = Logit error, distributed iid Extreme Value

As in typical, our model includes a vehicle-level unobservable ξ_j that enters utility linearly and is constant across consumers[3]. The standard concern in demand estimation is that these unobservable features are correlated with price, biasing $\widehat{\eta}$. In a GMM setting, BLP (1995) introduced instruments for vehicle prices based on the characteristics of all vehicles in the market. Their 2004 analysis (MicroBLP) compared those instruments to the assumption that the aggregate price elasticity in the new car market was -1, as suggested by analysts at General Motors. The MicroBLP instrumental variables estimates had high variance, with point estimates that differed from the industry analysts' belief by several hundred percent. In this analysis, we thus follow microBLP and again calibrate η to achieve an aggregate elasticity of -1.

Nearly all other discrete choice demand models assume that non-price product and individual characteristics X_{ij} are exogenous, i.e. that $corr(\xi, X) = 0$. This is not particularly problematic in other applications where out-of-sample market shares are predicted using the same product characteristics. This is more awkward in this analysis given our simulations, in which automakers will endogenously modify characteristics along with prices. Nonetheless, we maintain the (strong) assumptions that indirect utility is linear in X and that X is exogenous.

Our likelihood function is also standard; we include a brief exposition for clarity. Integrating up over the logit error gives the probability the consumer i purchases vehicle j. This is the standard logit purchase probability, where k indexes vehicles from 1 to J.

$$P_{ij} = \frac{e^{V_{ij}}}{\sum_{k=1}^{J} e^{V_{ik}}} \tag{2}$$

The likelihood for consumer i that has purchased vehicle y_i is thus:

$$L_i(y_i \mid \beta_i, X_{ij}) = \prod_k P_{ik}^{Y_{ik}} = P_{iy_i} \tag{3}$$

[3] As Athey and Imbens (2005) point out, the assumption that the product-level unobservable enters all consumers' utility functions equally is stronger than needed, as microdata can identify product-level unobservables that vary as a function of individual observables. This issue, however, is not central to our analysis, so we make the standard, stronger assumption.

Where:
$Y_{ij} = \{1,\ y_i = j;\ 0, \text{ otherwise}\}$

Since the distribution of the β_i are determined by hyperparameters b and W, the sample likelihood can be written as:

$$\mathcal{L}(y \mid b, W, X) = \prod_{i=1}^{N} L_i(y_i \mid b, W, X_i) \tag{4}$$

2.3 Estimation

A classical strategy is to find the parameters b and W that maximize the above likelihood function. In the Bayesian approach, we update diffuse prior beliefs, denoted g, about parameters based on observed data, giving posterior distributions of b and W conditional on the data[4]:

$$G(b, W \mid X, y) \propto \prod_{i=1}^{N} L_i(y_i \mid b, W, X_i) \cdot g(b, W) \tag{5}$$

The Gibbs sampler is a Markov Chain Monte Carlo (MCMC) method that allows the econometrician to draw from posterior distribution G. This procedure breaks parameters b, W, and β into disjoint sets and draws iteratively from each set. After and initial "burn-in" period, the simulation reaches an ergodic state, and draws of each parameter set will be in proportion to the parameters' marginal distribution[5]. Since the "glass is half empty" and we only observe covariates X_i for the part of the population observed in the microdata, we must impute the additional covariates. Fortunately, the MCMC framework is natural for missing data problems. Using data augmentation, missing data are viewed as additional parameters, and another step to the Gibbs sampler is added in which their distributions are computed conditional on the most recent draws of the other parameters.

Our procedure thus combines these two different MCMC methods to draw from the posterior distribution of b and W. We begin by generating "observations" from the aggregate data; each vehicle sold is represented by s observations, each with frequency weight w_j such that $s \cdot w_j$ equals the quantity sold of that vehicle. These observations, without covariates, are "stacked" on the complete observations from the microdata. We then iterate through the following four steps:

[4] As this is a new combination of two well-known procedures, our exposition follows Train (2003) and Tanner (1996). These are excellent sources of additional detail.

[5] Casella and George (1992) provide an intuitive explanation of the process and why it works.

1. Draw from $f(b \mid \beta, W)$. With a diffuse normal prior, $b \sim N(\overline{\beta} \mid W/N)$.
2. Draw from $f(W \mid \beta, b)$. With an Inverse Wishart (IW) prior on W that is as diffuse as possible while still integrating to one, $W \sim IW(K_b + N, \frac{K_b I + NS}{K_b + N})$[6].
3. Impute missing covariates X^a for macro data, as detailed below.
4. For each individual, draw from $f(\beta_i \mid b, W, X_{ij}, y_i)$. This is done via a Metropolis-Hastings algorithm: a trial value of β_i, labelled β_i, is drawn using $\widetilde{\beta}_i = \beta_i + \rho C r$. Here, C is the Choleski factor of W, r is a vector of draws from the standard normal distribution, and ρ is a scalar calibrated to maximize the information content of the draws (see Rossi, Allenby, and McCulloch 2005). The posterior is $G(\beta_i \mid b, W, X_{ij}, y_i) \propto L_i(y_i \mid \beta_i, X_{ij}) \cdot \phi(\beta_i \mid b, W)$, where ϕ denotes the normal probability density function. The new draw is accepted if $\frac{L_i(y_i|\widetilde{\beta}_i, X_{ij}) \cdot \phi(\widetilde{\beta}_i|b,W)}{L_i(y_i|\beta_i, X_{ij}) \cdot \phi(\beta_i|b,W)}$ is larger than a random draw from $U(0,1)$.

The third step is the only departure from a very standard Gibbs sampler for random coefficients models, and it is this that allows us to incorporate both macro and micro data. This step requires draws from the distribution of each "macro" observation's missing covariates, $f(X_i^a \mid y_i^a, b, W)$, conditional on the most recent draws of hyperparameters b and W and vehicle choice y_i^a. We implement this by using our approximation of the marginal population distribution of X contained in the 2001 Current Population Survey (CPS). By projecting b and W onto the covariates for each individual in the CPS, we predict each CPS individual's purchase probabilities $\widehat{P}_{ij}^{CPS} \mid b, W, X_i^{CPS}$. For each observation in the macro data that has purchased vehicle j, we then generate our draw X_i^a from the distribution of individual characteristics by sampling an individual from X^{CPS}. The choice of the CPS individual is random, with probability proportional to the chance that they purchased vehicle j, $\widehat{P}_{ij}^{CPS} \mid b, W, X_i^{CPS}$.

More precisely, the subparts of step 3 are:

1. Using the most recent iteration of b and W, draw β_i^{CPS} for each observation in the CPS.
2. Compute each individual's purchase probabilities $P_{ij}^{CPS} \mid \beta_i^{CPS}$
3. Compute predicted sales of each vehicle, $S_j = \sum_{i=1}^{N_{CPS}} P_{ij}^{CPS}$.
4. For each vehicle, compute the probability that it was purchased by each individual, $s_{ij} = P_{ij}^{CPS}/S_j$.

[6]In the present analysis, we use a diagonal W, whose parameters are drawn through a slightly different procedure.

5. For each observation in the macro data that chose vehicle j, draw covariates from an observation i in the CPS with probability s_{ij}.

Although one of the motivations for the Bayesian estimation strategy is computational efficiency, implementing this estimation algorithm with a relatively large choice set required several simplifications for tractibility. Since aggregating choices into larger classes reduces the exploitable variation in the choice set, we focused on reducing the number of individuals over which choice probabilities must be computed. First, since over 90 percent of the observations in both the microdata and the macro data do not purchase a new vehicle, a subset of the individuals who choose the outside option are randomly sampled for inclusion; their frequency weights are adjusted appropriately. Second, since the CPS includes on the order of 100 thousand individuals (depending on the year), we also subsampled from the CPS to generate our population distribution. Finally, the data augmentation routine typically involves a number of imputations of each X_i^a in each iteration, with subsequent likelihoods computed from the mixture of distributions conditional on each imputation. Because we have $s = 20$ observations in the macro data of "individuals" purchasing each vehicle and those 20 imputations will form a distribution of $X^a \mid y^a$, we impute each X_i^a only once.

2.4 Results

2.4.1 Simulated Data

We first present results from applying the procedure to simulated data. In this simulation, utility is a function of five product characteristics and the interaction of each product characteristic with one of five individual characteristics. All individual and product characteristics are uncorrelated random draws from $U(0,1)$; note that this eases identification relative to the typical application in which characteristics are correlated. The dataset was designed so that the microdata alone could theoretically identify the parameters, but the additional macro data would be especially useful. The number of products is 24 in the microdata, while the macro data includes 48. The sample size is 300 in the microdata, and 1500 simulated observations from the "CPS" are used in the aggregate data.

Table 3 presents the results of the two runs: the mean and standard deviation of the draws from the posterior distribution. The leftmost column is the true utility parameter used in the data generating process. Even with only 300 observations, the microdata estimation gives reasonable estimates of the b parameters. The random coefficients W, however, are extremely difficult to identify with such a limited choice set and sample size. The addition of aggregate data aids in pinning down W, and the more precise estimate of W then helps in estimating b.

Figure 3 shows the draws of the first five b parameters in the Gibbs sampler with microdata only, which have "true" values $\{-2,-1,0,1,2\}$. The starting values were set to $\overrightarrow{0}$, and after about 20,000 iterations the simulation reaches ergodicity. The micro-macro simulations burn in starting from the microdata-only parameter estimates; Figure 4 shows this procedure, in which the inclusion of macro data has substantially reduced the variance of the draws[7]. Figures 5 and 6, which have very different y-axis scales, illustrate that these additional data aid substantially in pinning down the variance W of two of the random coefficients; the other 8 look similar.

Having illustrated the procedure's effectiveness, we now continue to our application to new vehicle demand.

2.4.2 Actual Data

To illustrate the data, we first run a simple fixed coefficient representative consumer logit estimated with the standard OLS equation:

$$\log(Q_j/Q_0) = \beta X_j + \eta P_j + \xi_j \tag{6}$$

Table 4 reports the $\widehat{\beta}$ estimates in both the micro and combined years of macro data. Within each regression, the relative parameter magnitudes are meaningful, so the results are quite similar, and more precisely are not statistically different.

Table 5 shows the results of the Gibbs sampler with microdata; our micro-macro results are still being finalized. We fix $\eta = -2.3$ for all consumers, which sets the market price elasticity to -1. All coefficients have reasonable signs: indirect utility decreases in dollars per mile, increases in HP (especially for non-urban individuals), and increases in weight (especially for individuals in larger households). Interestingly, the first W parameter indicates that wealthier individuals have stronger preferences for high fuel economy cars; this result (as well as the signs and rough magnitudes of the other coefficients) holds consistently across alternative specifications.

Having modeled the new vehicle demand system, we can now continue on to our simulations of automaker responses to fuel economy regulation.

[7] There is a remaining issue with the "tuning" of ρ in the procedure, which determines the "stickiness" of the parameter draws. We follow a rule of thumb discussed in Rossi, McCulloch, and Allenby (2005), which is to set ρ such that the Metropolis-Hastings algorithm accepts 23% of the β_i draws. As they note, however, this is not actually the proper way to set ρ, and we suspect that the addition of macro data may affect this.

3 Counterfactual Simulation and Equilibrium Selection

The objective of our simulations is to estimate how automakers will adjust endogenous vehicle attributes (technology, fuel economy, power, weight and price) in response to regulation. We simulate vehicle and fleet attributes for four different policy scenarios - the Corporate Average Fuel Economy standards proposed by the National Highway Traffic Safety Administration ("NHTSA", 2008), an incremental carbon-based gas tax, a feebate based on vehicle emission and a base case, in which no new regulation is enacted. We focus specifically on the 2015 model-year - approximately one product cycle from current vehicles, as well as the final year in which the NHTSA proposal defines fuel economy standards.

We first present a game in which automakers endogenously select vehicle attributes for which many possible equilibria exist. In the second section, we discuss an equilibrium refinement, myopic partial best response, which reasonably approximates the iterative process by which automakers decide on future vehicle attributes and which we use to iteratively select one of the many possible equilibria. We then present the details of the four policy scenarios we consider, followed by our policy results and results examining the effect of alternative the equilibrium selection criteria.

3.1 Model

We focus on the static Nash equilibria of a game of complete information, in which six automakers, $f \in \{Chrysler, Ford, GM, Honda, Toyota, Other\}$, choose four attributes (fuel economy (mpg), power (hp), weight, and price) for five classes of model-year 2015 vehicles, $c \in \{Compact, Mid/FullSize, SUV, Truck, Minivan)\}$. For computational tractability, we aggregate all automakers not explicitly modeled into the category, "Other." Used vehicles and non-purchasing are aggregated into a generic outside good with utility, U_0. For computational tractability, we make the strong assumption that used vehicle markets do not affect the equilibrium in new vehicle markets, which we acknowledge will affect our results. We denote the triplet of physical attributes for a particular vehicle and automaker as x_{fc} and denote the price as p_{fc}.[8]

Letting x_f denote the vector of attributes for firm f, each firm observes demand and cost and chooses attributes and prices to maximize profits given by:

[8] We initially consider the case in which automakers choose price and attributes at the same time - we are currently working on an extension in which automakers first fix the fuel economy, power and weight of each vehicle, and then set Bertrand prices for all vehicles. In this case, automakers announce only the triplet of physical attributes which maximize profits conditional on the physical attributes opponents' announced in the previous period. Prices are then set optimally for all automakers, conditional on all vehicles' physical attributes.

$$\Pi_f(x_f, p_f) = \sum_c [p_{fc} - MC(x_{fc})] \int_\eta N(\eta) s_{fc}(p_f, x_f, p_{-f}, x_{-f}, \hat{\beta}, \eta) d\eta, \quad (7)$$

where s_{fc} is the market share of automaker f's vehicle class c, η is a set of characteristics of an individual drawn from the CPS, $N(\eta)$ is the person-weight used to scale the CPS record up a nationally representative sample, and $\hat{\beta}$ are the estimated structural coefficients from our demand model.

Vehicle fuel economy, power and weight affect consumer demand as well as the marginal cost of production. Because data on production costs is limited in this industry, we estimate production functions using an engineering approach instead of the typical econometric approaches (e.g. Olley and Pakes 1996; Levinsohn and Petrin 2003). We decompose marginal cost additively into two parts: the marginal cost for automaker f to produce vehicle class c, MC_{fc}^{base}, derived from the first-order conditions of the demand system, and the incremental cost of the necessary fuel efficiency[9] improvement to achieve fuel economy, horsepower and weight x_{fc}[10]:

$$MC(x_{fc}) = MC_{fc}^{base} + e^{\phi_c \Delta FuelEff(x_{fc})} \quad (8)$$

We estimate class-specific incremental cost curves for fuel efficiency based on engineering estimates from National Research Council (2002). The NRC study estimates the cost, fuel efficiency improvement and availability of different vehicle technologies for ten different vehicle classes. From the engineering data, we construct cost curves for each of the ten vehicle class and estimate a log-log relationship between fuel efficiency and incremental marginal cost for each of our five vehicle classes. Table 7 present the class-specific coefficients (ϕ_c) - generally, the larger the vehicle, the greater the number of technological improvements exist and the lower the cost necessary to achieve a given fuel efficiency improvement.

Letting FE_{fc}, HP_{fc} and W_{fc} denote the fuel economy, horsepower and weight of vehicle fc, and letting FE_{fc}^0, HP_{fc}^0 and W_{fc}^0 denote the physical attributes of automaker f's 2007 model-year vehicle in class c, we assume the necessary percentage increase in fuel efficiency is given by the following relationship[11]:

[9] Fuel efficiency is defined as amount of energy generated per unit of fuel. Fuel economy is the distance traveled per unit of fuel. Engine, transmission and design improvements increase fuel efficiency - increased fuel efficiency allows automoakers to improve a vehicle's fuel economy, power and/or weight.

[10] Note that new technology in this context does not shift out the production possibility frontier - by investing in new technology, the automaker can build a vehicle with better attributes at a higher cost.

[11] We assume that $\alpha_{FE} = 1$, $\alpha_W = 1$ - conditional on weight and power, a increase in fuel efficiency translates into an equivalent increase in fuel economy. Similarly, holding fuel

$$\Delta Eff_{fc} = \left(\frac{FE_{fc}}{FE^0_{fc}}\right)^{\alpha_{FE}} \left(\frac{HP_{fc}}{HP^0_{fc}}\right)^{\alpha_{HP}} \left(\frac{W_{fc}}{W^0_{fc}}\right)^{\alpha_{W}}. \tag{9}$$

The static Nash equilibria to the game are all sets of vehicle attributes $\{x_f, p_f \forall f\}$ such that no firm has the incentive to unilaterally devatiate.

3.2 Equilibrium Selection and Simulation

Many possible sets of attributes satisfy the equilibrium condition to the game specified above. In order to simulate counterfactuals, two approaches have been taken in the literature. The first approach follows Lee and Pakes (2008) which solves for all possible equilibria to the ATM entry game presented in Ishii(2006) and then restricts the set of equilibria to those robust to cost shocks. In the case of automaker attribute selection, the size of the action space is sufficiently large to prevent a systematic search for all possible equilibria to the game. The second approach, which we follow, is to choose a equilibrium refinement reasonably representing the automaker decision-making process and use the refinement to select one of the many possible equilibria to the game.

We consider an equilibrium refinement, reinforcement learning, from Fudenberg and Levine (1988).[12] Reinforcement learning defines a sequence of myopic "announcements" converging to a unique equilibrium. At the beginning of the game, each firm f chooses the action which is a myopic best response to starting actions of all other players, x_0. In each subsequent round t, firm f chooses the action $x_{f,t}$ which is a myopic best response to some function of the actions played by opponents in all previous rounds, $\{x_{-f,0}, x_{-f,1} x_{-f,t-1}\}$. The optimal choice of $x_{f,t}$ satisfies

$$x_{i,t} = argmax \Pi_i(x_{i,t}, f(x_{-i,t-1})) \tag{10}$$

Depending on the starting point x_0 and the function by which a firm aggregates opponents' actions in previous rounds, the reinforcement learning equilibrium, x^*, is the set of actions to which the sequence of actions converges - that is $x^* = lim_{t\to\infty}\{x_{i,t}, x_{-i,t}\}$. We consider two specific types of reinforcement learning: (1) myopic partial best response, in which each firm only considers the opponents' last actions, and plays the myopic best response to $x_{-i,t-1}$, and (2)

economy and power constant, energy generated per unit of fuel is proportional to the weight of the vehicle moved. Alternatively, we could estimate the relationship between fuel efficiency and the three attributes of interest from empirical data, if we could reliably observe fuel efficiency.

[12] In our context, we are interested not in the learning dynamics, but in the game's steady state: the steady state of successive proposed actions is a Nash equilibrium. Rather than focusing on the dynamic by which the new equilibrium is reached, we use the refinement to identify a specific equilibrium from the set of all possible equilibria.

fictitious play, in which each firm averages each opponent's sequence of actions, and plays the myopic best response to the the set of average actions.[13]

Based on research and discussion with auto industry analysts, we believe that myopic partial best response reasonably approximates the automaker decision process. The sequence of actions we consider are a sequence of non-binding announcements made by the automaker prior to setting the final attributes of each vehicle. Although non-binding, the announcements provide information about the characteristics and technology of future vehicles. In our simulations, we select the limit of the announcements as the equilibrium.

Automakers begin to plan general characteristics of their entire line of vehicles ten to fifteen years in advance as part of a "cycle plan." In a cycle plan, firms begin to plan future models' powertrain technology (such as plug-in hybrids or hydrogen vehicles) and other general attributes. Analogous to the actions chosen in each round of the reinforcement learning equilibrium, automakers make repeated non-binding announcements at auto shows and conferences which reveal information about the characteristics and technology of future vehicles. In addition, trade press closely tracks and anticipates vehicle attributes and technology. Binding vehicle attribute decisions begin approximately three years prior to a particular model year. At this point, the automaker sets the attributes and technology of a vehicle, including "market segment and competitive positioning, expected sales volume and price, and key vehicle attributes including size, performance, drivetrain and other major technology options."[14]. Two years in advance, the automaker engineers model-specific components and develops the manufacturing capacity to produce the vehicle. In the final year, the automaker sets marketing and prices. The degree to which myopic partial best response reasonably represents automotive new product development hinges on four assumptions. First, automaker announcements must be observable to other automakers, revealing the potential attributes of their future product lines. Second, firms must respond strategically and play a best response to the announcements of competitors. Third, firms must base their decisions on the most recent "announcements" of competitors. Finally, firms must behave myopically. We consider each of these in order.

Automakers provide substantial information about future product attributes years in advance of production. For example, GM and Ford currently provide information about the attributes of the Volt and Edge HySeries, a plug in hybrid and fuel cell vehicle respectively. In addition, trade press closely tracks future vehicle design - Motor Trend currently reports on attributes of vehicles three model years into the future. In addition, automakers often announce new technology prior to incorporating it in consumer vehicles. For example, Ford has reported that it "intends to improve the efficiency of its internal combustion engines through the addition of direct-injection and turbocharging technologies, resulting in a ten to twenty percent improvement in fuel economy and better

[13] Although in each case, convergence selects a particular equilibrium, the two refinements could select different equilibria. We examine this possibility in an extension of our simulation.

[14] Center for Automotive Research (2007)

performance." [15] Thus, while styling and appearance of new models may difficult to observe in advance of production, major changes in the underlying characteristics that we model in a product line - fuel economy, power, weight - can reasonably be observed in advance. Moreover, the vast majority of automaker announcement relate to active projects - there is little evidence that automakers make non-binding announcement about products they do not intend to pursue. Not only do many announcement incorporate some investment costs, but automakers use the announcements to publically signal their future product plans to financial markets.

Strategic, profit maximizing response should be uncontroversial, and the industry literature supports this view.

> A product's attributes and sub-attributes are typically evaluated against a defined set of competing products. Automakers closely track what the competition is doing to have a good idea of likely competing product profiles several years down the road to determine where their product's profile needs to be in order to meet or exceed its sales volume goals. [16]

Moreover, the sequential nature of the announcements suggests that automakers pay most attention to recent attribute announcements rather than initial announcements made early in the cycle-plan.[17] The final assumption which must hold is that firms play myopically. While there is relatively little evidence that automakers make non-binding announcements which they do not plan to execute in order to influence competitor product development, we cannot rule out the possibility that automakers act strategically in certain niche vehicle markets. Strategic play dramatically complicates counterfactual simulation - thus, for our purposes, we implicitly assume firms do not anticipate future rounds of announcements and act strategically to influence the future announcements of other automakers.[18]

The pattern of successive announcement through the trade press, conferences, auto shows and company press maps reasonably closely to the myopic partial best response algorithm. Automakers successive disclosures about product attributes well in advance of finalized decisions about the attributes of a particular model-year are roughly analogous to the iterative process by which we select the equilibrium attributes. Through repeated disclosures, firms adjust the attributes of their vehicle lines so as to best compete with the anticipated

[15] Automotive Design and Production (2008)

[16] Center for Automotive Research (2007)

[17] As an robustness check, we simulate results using several fictitious play selection criteria, each of which discounts previous announcements to a different degree. While the actual point of convergence depends to a small degree on the selection criterion, our policy implications are unchanged.

[18] As a partial test of the importance of strategic play, we consider a model in which firms make sequential announcements rather than simultaneous announcements. Although firms still behave myopically, we find that changing the sequence of play does affects the selected equilibrium attributes, suggesting that incorporating strategic play may be a valuable extension for future work.

attributes of other automakers. Following final selection of attributes for all vehicles, firms then choose retail prices, along with customer and dealer incentives.

In the context of our simulation model, we select our equilibrium by assuming that automakers make successive, simultaneous "announcements" about the attributes of future vehicles. To form each set of announcements, automakers determine the myopic best response to the attributes announced by opponents in the previous round, analogous to the actions chosen in each round of the reinforcement learning equilibrium refinement. To estimate each firm's announcement in each round t, we numerically simulate demand using the estimates from our demand system ($\hat{\beta}$ and drawing individuals from the Current Population Survey.[19] We discretize the set of vehicle attributes and prices into single percentage point increments relative to the base attributes of the model-year 2007 vehicles. In each round, we search the attribute space for the set or attributes which maximizes automaker i's payoff conditional on the configurations chosen by other automakers in the previous iteration, the cost of innovation and simulated demand. We iterate the process of announcements until the model converges to the myopic partial best response equilibrium.

3.3 Policy Scenarios

We initially consider four counterfactual simulations of future policies: (1) the recently proposed CAFE standards for passenger cars and light trucks, (2) a carbon tax on gasoline, (3) a feebate based on carbon emissions and (4) a base case, absent new CAFE standards or a carbon-based fuel tax or feebate. In each case, we use our equilibrium refinement to identify a particular set of equilibrium vehicle attributes for the 2015 model-year, approximately one product cycle away from the current production year.

The first scenario replicates the recent CAFE proposal made by NHTSA (2008).[20]The proposed standard replace the existing CAFE standards, defining more stringent requirements based on vehicle footprint. Specifically, the proposal specifies footprint-based fuel economy standards, where the footprint of the vehicle is the product of the wheelbase and the track (distance between the front and back axles). The footprint targets for 2011 to 2015 are designed to raise passenger car fuel economy from 31.2 mpg to 35.7 mpg and light truck fuel economy from 25.0 mpg to 28.6 mpg. The NHTSA estimates that fleet fuel

[19] We implicitly assume that the automakers' expectation of future demographics is accurately captured by current demographics. In addition, we hold our sample of individuals constant (including all randomly drawn coefficients on vehicle attributes) throughout all iterations to ensure convergence.

[20] Although the ultimate goal of the proposed CAFE standards is to define a path by which the fuel economy of the fleet of vehicles sold in 2020 will exceed the 35 miles per gallon objective set forth in the Energy Independence and Security Act ("EISA") of 2007, the rulemaking only proposed defined fuel economy standard through 2015. Thus, we focus on estimating the characteristics of fleet of 2015 model-year vehicles, although our methodology could easily be extended to estimate hypothetical standards for 2020.

economy will rise from 27.8 mpg in 2011 to 31.6 mpg in 2015. Ultimately, the 2015 targets are designed to ease industry transition to the fleet-wide 35 mpg standard in the Energy Independence and Security Act of 2007. For each automaker, the relevant standard for an automaker is the harmonic average of the footprint-based standards, weighted by the proportion of an automaker's fleet with each footprint. Automakers failing to meet the standards are subject to current CAFE-violations fines, adjusted for inflation - specifically, an automaker is fined $55 per vehicle for each miles per gallon its fleet of vehicles falls below the fleetwide CAFE standard.[21]

To estimate the footprint-based standards for each of our five relevant vehicle classes, we estimate the footprint of each vehicle class based on 2007 model-year vehicle dimensions in each of our five vehicle classes. We then calculate the relevant standard for each class of vehicle based on the proposed footprint standard for 2015. Our estimates for the 2015 standards for our five vehicle classes (Compact, Mid/Full, SUV, Truck and Van) are approximately 39, 32.3, 31.3, 26.6, and 28.3 miles per gallon respectively. As discussed in the previous sections, automakers choose the degree to which they want to use more fuel efficient technology in their vehicles, but are allowed to allocate fuel efficiency gains towards fuel economy, power or vehicle weight. Due to concern in the proposal that footprint-based standards will incentivize safety-reducing downweighting by automakers, we constraint the degree to which automakers can reduce weight to increase fuel economy.[22]

We simulate our second and third counterfactuals to compare the effect of the 2015 CAFE standard to other policy mechanisms which primarily incentivize consumer purchase of high fuel economy vehicles. Specifically, we consider a carbon tax on gasoline and a feebate based on vehicle carbon emissions generated during the first 60 thousand miles. In each case, we set the price of carbon to $43 per ton, taken from Tol (2005)[23]. Translating the tax on a per ton basis to a tax on a per gallon of fuel, we estimate that a $43 per ton tax is equivalent to approximately a $0.42 tax per gallon of gasoline. [24] As an alternative, we calculate a feebate, based on the carbon emissions associated with the first fifty thousand miles of vehicle ownership. We assume that the feebate is assessed as an additional marginal cost of production on the automaker. The feebate varies with the inverse of fuel economy - we assess feebates of $1050, $700, and $525 on vehicles with fuel economies of 20 mpg, 30 mpg and 40 mpg respectively.

Finally, we simulate a counterfactual absent the proposed CAFE standards, as well as the carbon tax on gasoline or the feebates based on vehicle carbon

[21] The proposal allows automakers to trade CAFE credits, but specifies a minimum fuel economy for domestically produced passenger vehicles each automaker purchsing credits must meet. In our simulations, we currently abstract away from interautomaker trade.

[22] We limit the reduction to 5 percent of the base weight for the relevant vehicle an automaker produces for a particular class. In practice, the majority of vehicles are unconstrained by the floor on weight.

[23] Although many studies estimate the long-run marginal damages of carbon emissions, the estimates in Tol(2005) are cited by the Working Group II in their contribution to the Fourth Assessment Report of the IPCC

[24] Implicitly, we assume that consumers bear the full burden of a carbon tax on fuel.

emissions. The "base case" allows us to benchmark the policies relative to changing demand resulting from changing gasoline prices and vehicle preferences.

4 Results

We present the results of the simulation in tables 8 and 9, which aggregate vehicle characteristics in each simulation by class and automaker respectively.[25] For each of the four simulations table 8 presents equilibrium attributes identified by the reinforcement learning selection criterion. In addition, we also include the mean fuel economy, horsepower, weight and price for the 2007 model-year of each vehicle class.

In each of the four counterfactuals, including the "base case" in which regulations are absent, the average fuel economy for the entire fleet of vehicles exceeds the fleet-wide fuel economy goal of 31.6 mpg in the NHSTA proposed rulemaking - mean fleet-wide fuel economy in the base case, CAFE simulation, Feebate simulation and carbon gas tax simulation are 31.8, 34.7, 33.7 and 33.5 mpg respectively. Interesting, although the fleet-wide fuel economy exceeds the rulemaking objective of 31.6 mpg, we find variation in vehicle class compliance in all four simulations.

In the base case, automakers incorporate new technology to improve the fuel economy of compact and mid/full size passenger cars - both categories, on average, substantially overcomply with the footprint standards. Relative to the 2007 model year, we estimate that automakers improve the fuel economy of the fleet of small and midsize vehicles by approximately 60 to 65 percent - to achieve the fuel economy gains, automakers improve fuel efficiency (through new technology) approximately 40 percent and reduce power by approximately 20 percent. Automakers overcomply with the standards (even in the base case) to target the subset of the consumers who place substantial value on fuel economy.

Automakers also invest in new technology to improve fuel economy of sport-utility vehicles (SUVs), vans and trucks. In the base case, fuel efficiency for SUVs, vans and trucks rise 68 percent, 48 percent and 51 percent respectively. Unlike Compact and Midsize sedans in which automakers focus the fuel efficiency gains on improving fuel economy, automakers increase fuel economy, power and weight of SUVs, vans and trucks. We estimate that automakers increase fuel economy in SUVs, trucks and vans by 25 percent, five percent and 14 percent respectively, increase power by 16, 23 and 14 percent respectively and weight by approximately 15 percent in all three cases. As a result in the base case, mean SUV, truck and van fuel economy (23.6, 18.7 and 22.4 mpg, respectively) falls substantially below the estimated footprint-based standards of 31.3, 26.6 and 28.3 mpg.

[25] When using myopic partial best response as our equilibrium selection mechanism, convergence is reached after 5-6 sets of announcements. Using fictitious play where firms respond to the vector of average opponents'actions in all previous periods, actions converge after 8-10 announcements.

When we simulate the effect of the 2015 CAFE standards, total fleet fuel economy improves. Compact and Midsize vehicles continue to substantially surpass footprint-based thresholds. For SUVs and vans, automakers facing non-compliance reallocate fuel efficiency technology towards fuel economy rather than power. With the 2015 CAFE standards, average fuel economy for SUVs and Vans improves 48 and 38 percent respectively - while the vehicle class fuel economy averages still fail to meet the footprint-based thresholds (which automakers meet by trading CAFE credits from Compact and Midsize vehicles), reallocation of fuel efficiency technology improves fuel economy 4.5 mpg for SUVs and vans relative to the base case. Fuel economy for trucks is largely unchanged in our CAFE simulation relative to the base case. Although fuel economy for 2015 model-year truck fleet is slightly higher than the 2007 model-year truck fleet, it still falls substantially short of the footprint-based CAFE standard.

Table 9 presents similar summary statistics at the automaker-level. Variation in compliance among automakers also provides insights as to how automakers change vehicle attributes in response to the CAFE regulations. In the base case, only Toyota and Honda exceed their automaker standards under the proposed CAFE rulemaking. As in the 2007 model-year, Toyota and Honda tend to sell vehicles with higher fuel economy and lower power and weight than American automakers. In the 2015 model-year, Toyota and Honda further differentiate their vehicles along this dimension, tending to lower power and weight, while increasing fuel economy across all vehicle classes. In the base case, Ford, GM and "Other" automakers make more modest improvements to fuel economy (26, 37, and 29 percent, respectively), in addition to increasing vehicle power. Ford and the "Other" automakers fall particularly short of the proposed CAFE standards in the base case. We estimate Ford's fleet-wide fuel economy to be 25.3 mpg, relative to a benchmark of 31.9 mpg and estimate the "Other" automakers' fleet fuel economy to be 28.6 mpg relative to a benchmark of 34.5 mpg.

With the introduction of the proposed CAFE standards, Ford and the "Other" automakers optimize vehicle attributes - reducing vehicle power and increasing fuel economy. In the CAFE simulation, "Other" automakers come into compliance, while Ford comes very close to compliance with a fleet-wide fuel economy of 31.6 (relative to a benchmark of 31.9 mpg) GM and Chrysler, on the other hand, fail to come into compliance. We estimate that GM and Chrysler fall approximately 4 and 1.8 mpg short of the standard, respectively. For GM, in particular, part of the non-compliance arises from an increase in consumer purchases of GM SUVs and vans, both of which fall short of their respective footprint-based CAFE standards. As Ford and "Other" automakers increase fuel economy of their SUVs and vans at the expense of power, consumers who value power over fuel economy substitute to GM vehicles.

Examining the feebate and carbon gas tax simulations, we find evidence that the programs create different incentives for automakers to change vehicle attributes. Automakers respond to a carbon tax on gasoline by increasing fuel economy for Compact cars from 46.4 mpg (in the base case) to 51.8 mpg, targeted at consumers who place a high value on fuel economy, or alternatively experience substantially disutility from high gasoline prices. As expected the change in fuel economy for compact cars, comes primarily through Honda and Toyota, who offer the most fuel efficient vehicles. Changing the characteristics of the vehicles [26] With feebates, automakers increase fuel economy of SUVs and vans, which otherwise incur $400-$500 higher fixed charges than Small and MidSize sedans. Interestingly, the incremental feebate of $400-$500 vehicles is large relative to potential CAFE non-compliance penalties of $55 per mpg per vehicle - as a result, the feebate generates similar behavior to that generated by CAFE for automakers with fleets with relatively low fuel economy.

5 Conclusion

This paper studies the 2015 CAFE standard, contributing to a growing policy literature on fuel economy regulation by simulating how firms will adjust vehicle attributes and prices in equilibrium. We have also considered two alternative policies: feebates based on a vehicles' anticipated carbon emissions and a carbon tax levied on gasoline. We find that, even absent new regulation, automakers have a strong existing incentive to improve the fuel economy of Compact and Midsize sedans if consumers expect gasoline prices to stay relatively high. In the base case, the incentive to improve fuel economy is sufficiently great to push Honda and Toyota's fleet fuel economy above their relevant CAFE standards. The addition of the proposed CAFE standards increases the incentive for other automakers not in compliance in the base case to improve fuel economy, rather than power or vehicle weight. In total, we estimate fleet fuel economy to be 34.7 miles per gallon with the CAFE regulations. Although the regulations are binding at the automaker-level, fleet fuel economy exceeds the proposed rule's stated goal of 31.6 mpg.

We emphasize that it is not possible to write a paper that advances all parts of a literature as well-developed as the body of work on the economics of the auto industry. Indeed, in making our two methodological advances, it became necessary to abstract away from issues that other analyses of the auto industry have done better. For example, we ignore the introduction of alternative fueling stations, which the National Energy Modeling System covers in detail, causing us to overstate the utility of new powertrain technologies. We also ignore the equilibrium effects on used car prices and scrappage, causing us to overstate equilibrium demand for the smaller number of new low-fuel economy vehicles being sold under the strengthened CAFE standard. We additionally assume a myopic learning dynamic for computational tractibility. In making these and other simplifications, we are modest about our contribution to literature on the effects of fuel economy regulation.

[26] The calculated CAFE standards for each automaker are based on predicted fleet sales in each simulations. Thus, the fact that CAFE standards do not significantly rise with the addition of a carbon gas tax or feebate implies that consumers do not substantially change their vehicle

This applied problem, however, has motivated two contributions to the econometrics and Industrial Organization literature on demand estimation and counterfactual simulation. In estimating demand, we introduce a new procedure to exploit variation in both consumer-level micro data and market-level quantities and choice sets. While Bayesian strategies are increasingly popular in demand estimation, there was no existing method for combining these two types of data. In simulating counterfactuals, we address the problem of multiple equilibria by using myopic partial best response as a selection mechanism, which we argue to be a reasonably stylized model of the industry's product development process. We consider how the choice of equilibrium refinement affects the seleted equilibrium by also considering the fictitious play learning dynamic. Our paper is one of the first to consider how the choice the equilibrium refinement affects outcomes in counterfactual simulations.

6 Bibliography

Athey, Susan, and Guido Imbens (2005). "Discrete Choice Models with Multiple Unobserved Choice Characteristics." Working Paper, Harvard University.

Automotive Design and Production (2007). “Going Green: The Challenges and the Solutions.” January 1, 2008.

Allenby, Greg, and Peter Lenk (1994). "Modeling Household Purchase Behavior with Logistic Normal Regression." Journal of the American Statistical Association, Vol. 88, pages 669-679.

Austin, David, and Terry Dinan (2005). "Clearing the air: The costs and consequences of higher CAFE standards and increased gasoline taxes." Journal of Environmental Economics and Management, Vol. 50, pages 562-582.

Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen (2006). "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes." Working Paper, Stanford University.

Berry, Steven (1994). "Estimating Discrete Choice Models of Product Differentiation." RAND Journal of Economics, Vol. 23, No. 2, pages 242-262.

Berry, Steven, and Ariel Pakes (2005). "The Pure Characteristics Demand Model." Working Paper, Harvard University.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium." Econometrica, Vol. 63, No. 4 (July), pages 841-890.

Berry, Steven, James Levinsohn, and Ariel Pakes (2004). "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." Journal of Political Economy, Vol. 112, No. 1, pages 68-105.

Bresnahan, Tim (1987). "Competition and collusion in the American Automobile Industry: The 1955 Price War." Journal of Industrial Economics, Vol. 35, No. 4, pages 457-482.

Caplin, Andrew, and Barry Nalebuff (1991). "Aggregation and Imperfect Competition: On the Existence of Equilibrium." Econometrica, Vol. 59, No. 1 (January), pages 25-59.

Casella, George, and Edward George (1992). "Explaining the Gibbs Sampler." The American Statistician, Vol. 46, No. 3 (August), pages 167-174.

Center for Automotive Research (2007). “How Automakers Plan Their Products.” July 2007.

Chib, Siddhartha, and Edward Greenberg (1998). "Analysis of Multivariate Probit Models." Biometrika, Vol. 85, No. 2, pages 347-361.

Davis, W.B., M.D. Levine, Kenneth Train, and K.G. Duleep (1995). "Effects of Feebates of Vehicle Fuel Economy, Carbon Dioxide Emissions, and Consumer Surplus." Working Paper, Office of Policy, US Department of Energy.

Energy Information Administration (2007). "Transportation Sector Module of the National Energy Modeling System: Model Documentation 2007." United States Departemtn of Energy. http://tonto.eia.doe.gov/FTPROOT/modeldoc/m070(2007).pdf

Espey, Molly, and Santosh Nair (2005). "Automotive Fuel Economy: What is it Worth?" Contemporary Economic Policy, Vol. 23, No. 3 (July), pages 317-323.

Feng, Ye, Don Fullerton, and Li Gan. "Vehicle Choices, Miles Driven, and Pollution Policies." National Bureau of Economic Research Working Paper No. 11553.

Fudenberg, Drew, and David Levine. "The Theory of Learning in Games." Cambridge, MA: MIT Press.

Geweke, John (1989). "Bayesian Inference in Econometric Models Using Monte Carlo Integration." Econometrica, Vol. 57, pages 1317-1339.

Goldberg, Pinelopi (1995). "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." Econometrica, Vol. 63 (July), pages 891-951.

Goldberg, Pinelopi (1998). "The Effects of the Corporate Average Fuel Economy Standards in the US." Journal of Industrial Economics, Vol. 46, pages 1-33.

Greene, David, Philip Patterson, Margaret Singh, and Jia Li (2005). "Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy." Energy Policy, Vol. 33., pages 757-775.

Hausman, Jerry (1979). "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." Bell Journal of Economics, Vol. 10, No. 1, pages 33-54.

Hellerstein, Judith, and Guido Imbens (1999). "Imposing Moment Restrictions from Auxiliary Data by Weighting." The Review of Economics and Statistics, Vol. 81, No. 1 (February), pages 1-14.

Hotelling, Harold (1929). "Stability in Competition." The Economic Journal, Vol. 39, No. 153 (March), pages 41-57.

Howarth, Richard B (2004). "Fuel Economy Standards." The Journal of Economic Perspectives, Vol. 18, No. 2 (Spring), pages 272-273.

Imbens, Guido, and Tony Lancaster (1994). "Combining Micro and Macro Data in Microeconometric Models." Review of Economic Studies, Vol. 61, No. 4 (October), pages 655-680.

Ishii, Joy (2006). "Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry." Working Paper, Stanford University.

Jacobsen, Mark (2008). "Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity." Working Paper, University of California at San Diego.

Jaffe, Adam, and Robert Stavins. "The Energy Paradox and the Diffusion of Conservation Technology." Resource and Energy Economics, Vol. 16, pages 91-122.

Lee, Robin, and Ariel Pakes (2008). "Multiple Equilibria, Selection, and Learning in an Applied Setting." Working Paper, Harvard University.

Levinsohn, James, and Amil Petrin (2003). "Estimating Production Functions Using Inputs to Control for Unobservables." The Review of Economic Studies, Vol. 70, pages 317-341.

Klier, Thomas, and Joshua Linn (2008). "The Price of Gasoline and the Demand for Fuel Efficiency: Evidence from Monthly New Vehicles Sales Data." Working Paper, University of Illinois at Chicago.

Knittel, Christopher, and Konstantinos Metaxoglou (2008). "Estimation of Random Coefficient Demand Models: Challenges, Difficulties, and Warnings." Working Paper, University of California at Davis.

National Highway Traffic Safety Administration. "Average Fuel Economy Standards for Passenger Cars and Light Trucks Model Years 2011-2015." Docket No. NHTSA-2008-0089. April 21, 2008.

National Research Council (2002). "Effectiveness and Impacts of Corporate Average Fuel Economy (CAFE) Standards." Washington, DC: National Academy Press.

Nevo, Aviv (1998). "A Research Assistant's Guide to Random Coefficients Discrete Choice Models of Demand." National Bureau of Economic Research Technical Working Paper 221.

Oak Ridge National Laboratory (ORNL) (2001). "The National Household Travel Survey." http://nhts.ornl.gov/.

Olley, Steven, and Ariel Pakes (1996). "The Dynamics of Productivity in the Telecommunications Equipment Industry." Econometrica, Vol. 64, No. 6 (November), pages 1263-1297.

Pakes, Ariel (2008). "Theory and Empirical Work on Imperfectly Competitive Markets." Working Paper, Harvard University.

Petrin, Amil (2002). "Quantifying the Benefits of New Products: The Case of the Minivan." Journal of Political Economy, Vol. 110, No. 4, pages 705-729.

Rossi, Peter, Robert McCulloch, and Greg Allenby (1996). "The Value of Purchase History Data in Target Marketing." Marketing Science, Vol. 15, No. 4, pages 321-340.

Rossi, Peter, Greg Allenby, and Robert McCulloch (2005). "Bayesian Statistics and Marketing." Sussex, England: John Wiley & Sons.

Tanner, Martin (1996). "Tools for Statistical Inference: Methods for Exploration of Posterior Distributions and Likelihood Functions." New York: Springer-Verlag.

Tanner, Martin, and Wing Hung Wong (1987). "The Calculation of Posterior Distributions by Data Augmentation." Journal of the American Statistical Association, Vol. 82, pages 528-540.

Tol, Richard (2005), "The Marginal Damage Costs of Carbon Dioxide Emissions: An Assessment of the Uncertainties." Energy Policy, Vol. 33, pages 2064-2074.

Train, Kenneth (2003). "Discrete Choice Methods with Simulation." New York: Cambridge University Press.

7 Graphs

7.1 Figure 1: Weight and Fuel Economy in Attribute Space

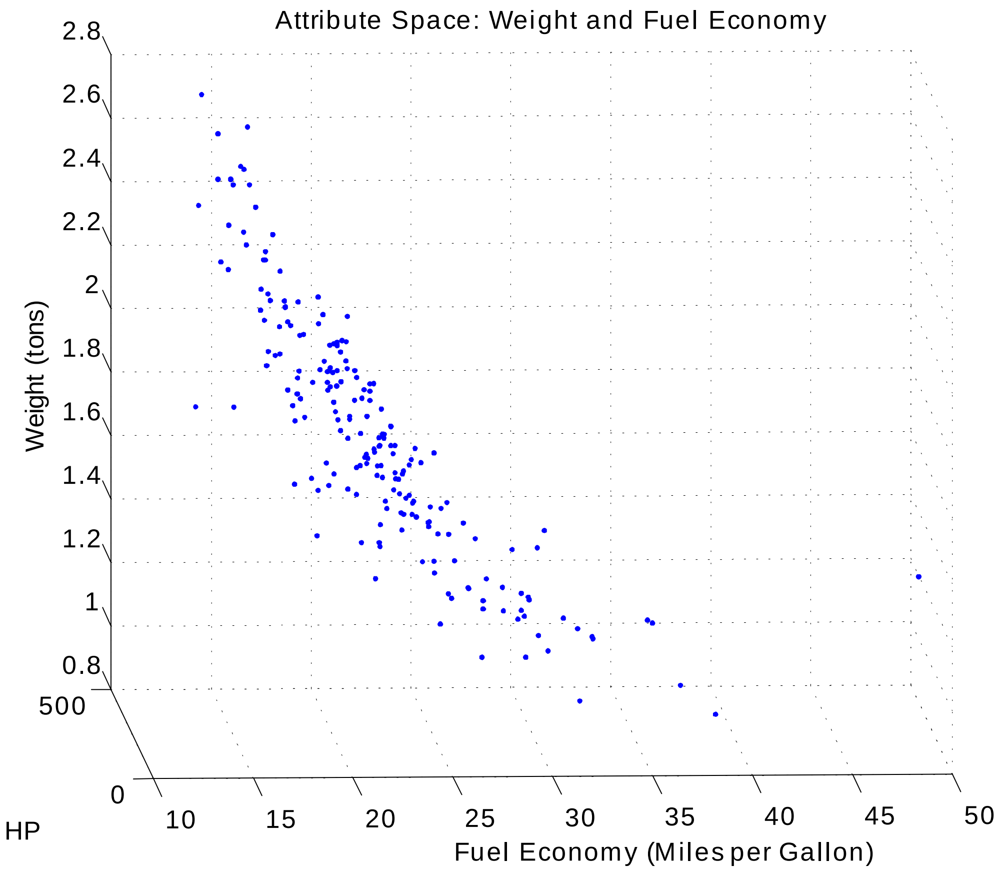

7.2 Figure 2: Weight and HP in Attribute Space

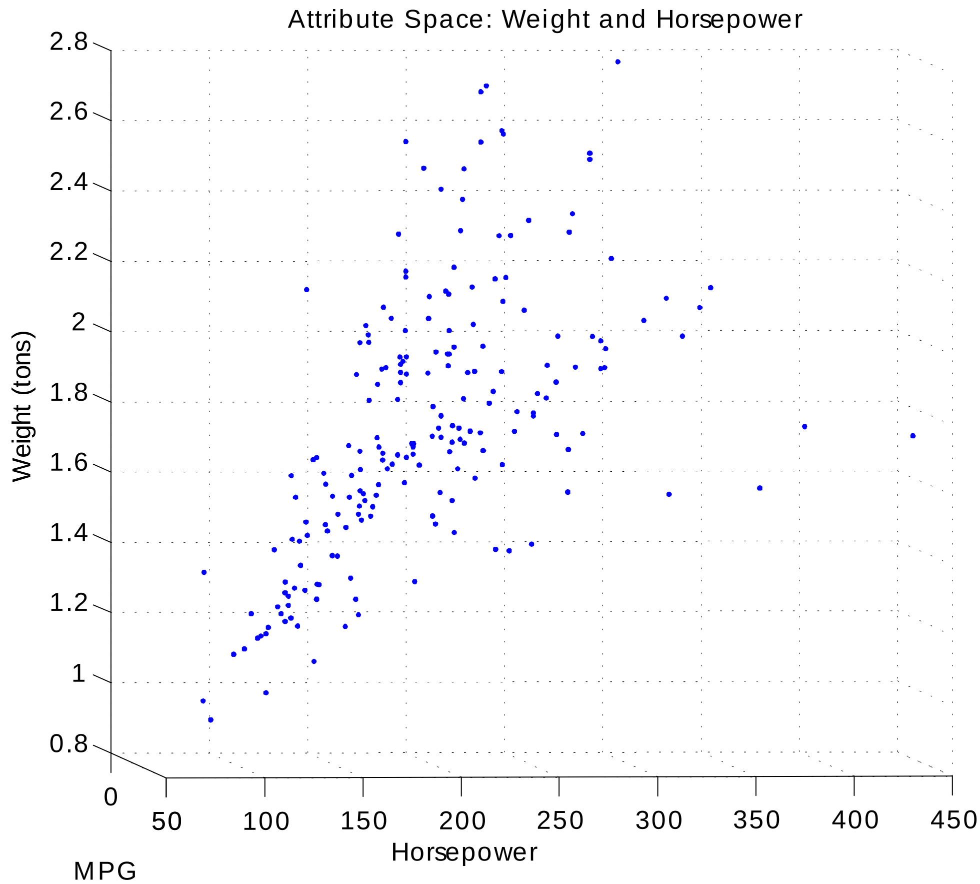

7.3 Figure 3: Gibbs Sampler Mean Coefficients: Microdata Only

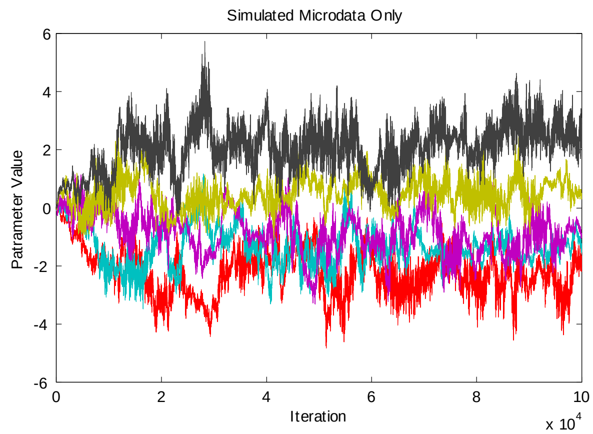

7.4 Figure 4: Gibbs Sampler Mean Coefficients: Micro and Macro Data

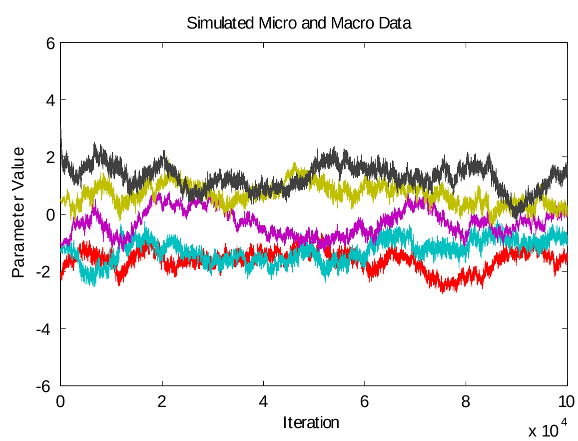

7.5 Figure 5: Gibbs Sampler Variance Parameters: Microdata Only

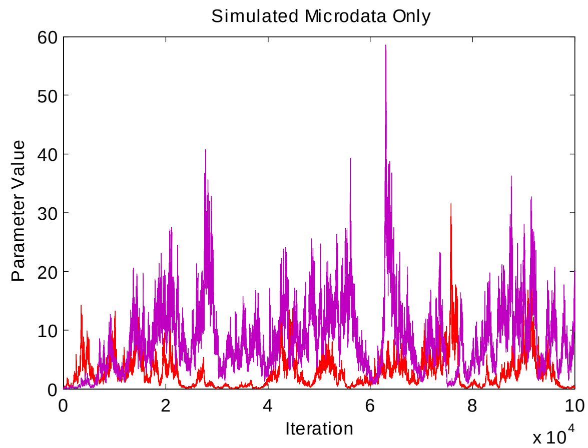

7.6 Figure 6: Gibbs Sampler Variance Parameters: Micro and Macro Data

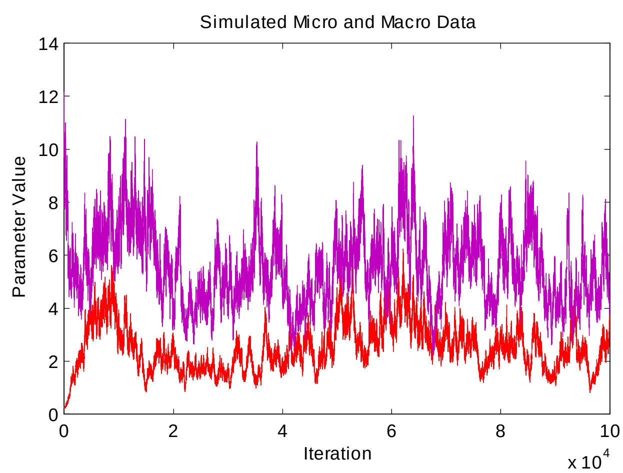

8 Tables

8.1 Table 1: Individual-Level Descriptive Statistics

8.1.1 From NHTS Microdata

	Obs	Mean	SD	Min	Max
HH Income (1000s)	6581	72.9	46.8	2.65	159
1(Urban)	6581	.273	.445	0	1
1(Suburban)	6581	.266	.441	0	1
Age	6581	48.3	17.2	16	88
Household Size	6581	2.77	1.33	1	14
Sample Weight	6581	89,902	89,452	3384	1,042,969

8.1.2 From CPS

	Obs	Mean	SD	Min	Max
HH Income (1000s)	97,912	74.0	69.8	0	875
1(Urban)	97,912	.278	.417	0	1
1(Suburban)	97,912	.444	.461	0	1
Age	97,912	44.2	18.0	16	90
Household Size	97,912	2.96	1.63	1	16
Sample Weight	97,912	2165	1154	67.1	18578

8.2 Table 2: Vehicle-Level Descriptive Statistics

8.2.1 From NHTS Microdata

	Obs	Mean	SD	Min	Max
Quantity (1000's)	741	31.1	55.8	0.43	643
Miles per Gallon	741	22.3	4.91	13.7	48.6
Dollars per Mile	741	.071	.015	.031	.111
Horsepower	741	188	48.9	67	368
Weight (tons)	741	1.75	.363	.958	2.79
Price (1000's)	741	24.7	9.78	8.24	84.4

8.2.2 From Macro Data

	Obs	Mean	SD	Min	Max
Quantity (1000's)	454	68.6	87.2	158	798
Miles per Gallon	454	21.7	5.31	12.0	60.3
Dollars per Mile	454	.098	.028	0.033	.194
Horsepower	454	217	73.9	66	617
Weight (tons)	454	1.85	0.40	0.80	2.94
Price (1000's)	454	30.9	27.8	8.38	368

8.3 Table 3: Micro-Macro Procedure with Simulated Data

	Microdata Only		Micro and Macro	
True Value	Mean	Std Dev	Mean	Std Dev
b Parameters				
-2	-2.4	0.6	-1.7	0.4
-1	-1.4	0.5	-1.1	0.4
0	-1.0	0.6	-0.3	0.4
1	0.6	0.4	0.6	0.3
2	2.2	0.7	1.3	0.5
2	3.5	1.0	1.6	0.8
1	2.7	1.0	1.2	0.7
0	2.5	1.2	0.8	0.8
-1	-1.0	0.8	-0.6	0.7
-2	-0.4	1.2	0.2	0.9
W Parameters				
2	5.1	3.6	3.9	1.1
1	1.8	2.1	6.2	1.4
0	3.2	3.1	2.7	0.8
1	3.3	3.7	4.8	1.2
2	9.2	6.3	5.6	1.3
2	6.9	8.1	2.1	0.7
1	15.1	10.9	2.4	1.0
0	17.8	14.7	2.9	1.1
1	3.6	5.6	2.6	0.9
2	4.3	5.9	2.9	0.9

Note: Results are calculated from every tenth draw in the sequence after a burn-in of 50,000 iterations.

8.4 Table 4: Aggregate Logit

	Aggregate	Micro
	(1)	(2)
log(Dollars per Mile/Weight)	-2.517	-.620
	(0.32)***	(0.506)
log(HP/Weight)	2.073	0.714
	(0.398)***	(0.342)**
log(Weight)	2.909	1.257
	(0.421)***	(0.346)***
log(Price)	-2.584	-1.036
	(0.221)***	(0.218)***
Const.	-19.231	-10.639
	(2.062)***	(2.357)***
Obs.	454	741
R^2	0.413	0.028
F statistic	59.548	6.388

Robust standard errors.

8.5 Table 5: Estimation Results

	Microdata Only	
	Mean	Std Dev
b Parameters		
log(Dollars per Mile)	-0.0584	0.5444
log(HP/Weight)	1.951	0.1714
log(Weight)	1.6979	0.2669
log(Price)	-2.3	0
log(DPM/Weight)*log(Income)	-0.3151	0.1354
log(HP/Weight)*1(Urban)	-0.679	0.3078
log(Weight)*(Household Size)	0.3363	0.0779
W Parameters		
log(Dollars per Mile)	1.0094	0.4098
log(HP/Weight)	0.2993	0.124
log(Weight)	0.3476	0.0867
log(DPM/Weight)*log(Income)	0.1998	0.0548
log(HP/Weight)*1(Urban)	0.7267	0.4254
log(Weight)*(Household Size)	0.1102	0.0382

8.6 Table 6: Fuel Efficiency Investment Cost

Dep Var: Log(Incr. Cost)	Compact	Mid/Fullsize	SUV	Truck/Van
Log(Incr. Efficiency)	1.76	1.723	1.725	1.713
	(0.033)***	(0.031)***	(0.026)***	(0.030)***
Constant	0.802	0.872	0.866	0.876
	(0.105)***	(0.102)***	(0.085)***	(0.101)***
Obs	56	58	85	84
R-squared	0.99	0.98	0.98	0.98

Table 8: Simulated Vehicle Characteristics, by Vehicle Class

	Fuel Economy Standard	Fuel Economy (mpg)	Power (hp)	Weight (tons)	Price (000s)
Starting Point - Model-year 2007					
Compact		27.84	181.60	1.51	20.57
Mid/Full		23.49	207.68	1.73	23.80
SUV		18.94	234.76	2.16	27.33
Truck		17.77	226.29	2.19	21.02
Van		19.62	216.36	2.25	23.54
2015 Base Case					
Compact	39.0	46.4	145.3	1.6	32.0
Mid/Full	32.3	38.0	166.1	2.0	37.2
SUV	31.3	23.6	273.0	2.5	42.6
Truck	26.6	18.7	277.6	2.5	32.7
Van	28.3	22.4	247.4	2.6	36.7
CAFEnew					
Compact	39.0	50.96	145.32	1.46	32.5
Mid/Full	32.3	36.80	165.97	2.04	37.5
SUV	31.3	28.00	224.94	2.53	43.2
Truck	26.6	18.38	279.32	2.56	33.0
Van	28.3	27.02	196.58	2.64	37.0
Feebate					
Compact	39.0	45.3	145.6	1.6	32.4
Mid/Full	32.3	37.6	166.0	2.0	37.0
SUV	31.3	28.2	218.7	2.5	42.6
Truck	26.6	18.3	279.0	2.5	32.8
Van	28.3	29.2	174.8	2.6	36.6
CarbonTax					
Compact	39.0	51.8	145.1	1.5	32.5
Mid/Full	32.3	37.7	165.9	2.0	37.5
SUV	31.3	23.8	276.0	2.5	43.0
Truck	26.6	18.6	275.3	2.6	32.9
Van	28.3	24.7	226.4	2.6	37.0

Table 9: Simulated Vehicle Characteristics, by Automaker

	Fuel Economy Standard	Fuel Economy (mpg)	Power (hp)	Weight (tons)	Price (000s)
Starting Point - Model-year 2007					
Ford		20.1	216.2	2.0	25.49
GM		20.9	228.2	2.0	25.01
Honda		26.8	195.5	1.7	23.03
Other		22.0	218.8	1.9	24.42
Toyota		25.1	198.0	1.8	23.38
2015 Base Case					
Chrysler	30.7	29.3	197.1	2.3	34.35
Ford	31.9	25.3	233.3	2.3	39.51
GM	32.3	28.8	251.7	2.3	39.76
Honda	35.2	46.4	174.5	1.8	35.69
Other	34.5	28.6	233.6	2.1	37.85
Toyota	32.5	40.9	171.6	2.0	36.25
CAFEnew					
Chrysler	30.6	28.8	199.4	2.4	34.37
Ford	31.9	31.6	173.1	2.3	39.49
GM	32.3	28.0	255.8	2.3	40.04
Honda	35.2	47.8	156.6	1.8	35.72
Other	34.5	37.0	181.6	2.0	38.74
Toyota	32.5	40.9	171.2	2.0	37.12
Feebate					
Chrysler	30.6	28.9	199.3	2.4	34.37
Ford	31.9	30.7	173.0	2.3	39.50
GM	32.4	29.7	240.2	2.2	38.79
Honda	35.2	48.0	156.3	1.8	35.66
Other	34.5	32.1	181.2	2.2	38.75
Toyota	32.5	40.7	171.1	2.0	36.24
Carbon Tax					
Chrysler	30.7	29.7	196.8	2.3	34.38
Ford	31.9	25.2	240.3	2.3	39.53
GM	32.3	28.9	251.4	2.3	39.73
Honda	35.2	49.3	156.3	1.8	35.64
Other	34.5	32.8	235.8	2.1	38.88
Toyota	32.5	41.8	169.8	2.0	37.18