Gas Prices, Fuel Efficiency, and Endogenous Product Choice in the U.S. Automobile Industry

Jacob Gramlich*
Georgetown University

Abstract

I develop and estimate a model of the U.S. automobile industry in which firms choose the fuel efficiency of their new vehicles. I use the model to analyze the 2008 gas price increase, and to computer the "gas-tax equivalent" of the 35 miles-per-gallon (*mpg*) CAFE proposals. Firms face a technological frontier between providing fuel efficiency and other quality, and the gas price shifts incentives to locate along this frontier. Where firms locate along this frontier is of environmental and policy significance. Demand is nested logit, supply is differentiated products oligopoly. The model is estimated using data from the US automobile market from 1971-2007. The model matches 2008 sales decline well, and suggests that an after-tax gasoline price of $4.55 would achieve fuel efficiency of 35 *mpg*, the stated goal of new CAFE proposals. Contributions to previous work include modeling product choice, relaxing restrictive identifying assumptions, and obtaining more realistic estimates of fuel efficiency preference. The model can be used to predict market equilibrium at any after-tax gasoline price.

KEYWORDS: Automobiles, endogenous product choice, environmental policy

J.E.L. CLASSIFICATION: D21, D12, H23, L1, Q38

---

*Georgetown University, McDonough School of Business, jpg72@georgetown.edu
I thank my advisors Steven Berry and Philip Haile for guidance. I also thank Justine Hastings, Patrick Bayer, Philip Keenan of General Motors, Martin Zimmerman, William Brainard, Costas Arkolakis, Larry Samuelson, Alvin Klevorick, Fiona Scott-Morton, Yuichi Kitamura, Donald Andrews, Robert Shiller, Joshua Lustig, Christopher Conlon, Myrto Kalouptsidi, and participants in the Yale microeconomics prospectus workshop for helpful comments. Steven Berry, Jim Levinsohn, and Ariel Pakes generously shared data for this research. The Cowles Foundation generously purchased data for this research.

1 Introduction

The industry for new automobiles is large and environmentally significant. 2007 revenues exceeded \$400 billion, or 3% of U.S. GDP. Passenger vehicles consume 20% of our nation's energy[1] and emit 20% of our nation's CO_2.[2] Automotive fuel efficiency[3] is one of the most direct measures of a vehicle's environmental impact, and has played a prominent role in our domestic energy policy for over 30 years.[4]

Recently fuel efficiency has received increased attention due to changes in both the economic and policy environment. Gasoline prices rose steadily for a decade before spiking dramatically in the summer of 2008. Also in 2008, the federal government passed the first significant increase in the CAFE Standards since they were instituted over 30 years ago.[5]

This paper investigates how changes in the economic and policy environment affect firms' decisions about product characteristics. In particular in the auto industry, changes in gasoline prices and taxes are likely to affect the fuel efficiency firms place in their new vehicles. This paper seeks to address the questions: how, and by how much, will changes in the economic and policy environment affect firms' choices of fuel efficiency in their new vehicles? Understanding how firms choose fuel efficiency is crucial to informing our energy policy.

I develop a model of the U.S. automobile industry in which firms choose the fuel efficiency of their new vehicles. Previous models held fuel efficiency fixed, and thus were unable to make predictions as to how this choice variable would change. Endogenizing this choice enables predictions of fuel efficiency in response to policy and market changes. The model gives quantitative predictions for market fuel efficiency at any gas price and tax combination.

The basic workings of the model are that consumers care more about fuel efficiency when gas prices are high, less when they are low, and firms face a technological frontier between providing fuel efficiency and other quality. The level of the gas price, therefore, working through consumer

[1] NEED (2008).

[2] EPA (2008).

[3] Fuel efficiency is the rate of gasoline consumption. In this analysis it is miles (traveled) per gallon (of gasoline consumed) - *mpg*.

[4] Fuel efficiency is the subject of the CAFE standards (Corporate Average Fuel Efficiency), which were instituted in 1975. These standards stipulate minimum sales-weighted fuel efficiencies for automobiles sold in America, and carry fines for non-complying firms.

[5] This increase was, at least in part, a response to rising gasoline prices.

demand, shifts firms' optimal locations along this frontier.[6]

Automobile manufacturers can increase fuel efficiency in their vehicle fleet in three ways. First, because they are multi-product firms, they may introduce more fuel efficient automobiles and/or discontinue less efficient models. Second, they may increase fuel efficiency of an existing vehicle, holding other quality constant. Third, they may increase fuel efficiency by trading off quality of the vehicle.[7] This paper focuses on the third mechanism. In the "medium run" (3-5 years) the first two mechanisms are limited because they change the automobile's "class" or "segment."[8] These are larger changes and take longer to implement. In this analysis I hold vehicle segments constant and focus on the third margin, the tradeoff between quality and efficiency within vehicle segment. Changes in the gas price induce changes in firms' optimal locations along this efficiency v. quality frontier.

Empirical models of industry, including seminal work on automobiles (Berry, Levinsohn, and Pakes, 1995), have focused on controlling for the correlation of price with unobservable quantities.[9] However, a limitation of these models has been the treatment of product characteristics. The limitation has been twofold. First, characteristics have been assumed to be fixed exogenously. This is a limitation when changes in these characteristics (such as fuel efficiency) are of direct interest. Second, in contrast to prices, the models have not allowed characteristics to be correlated with the unobserved cost and quality shocks. In fact, the identifying assumption of these models has been strict orthogonality between characteristics and the unobserved error terms. This restriction is implausible for the same reason it is with prices: prices and characteristics are both choice variables of the firm, and are chosen in response to an observed economic environment which includes econometrically unobserved quantities. So while the orthogonality assumption is useful in estimation, it is implausible in many industries including automobiles.

My model addresses both of these limitations in the context of automobiles. I develop a model in which firms choose product characteristics and this allows me to analyze changes in fuel efficiency itself. Second, I relax the restrictive identifying assumption commonly used to estimate demand. I instead construct estimation moments based on the timing of events in the auto industry. These moments are based on the assumption that ex-post regret in product planning is not systematic, or predictable, by the firms in any way. The reason there is ex-post regret in product planning is that

[6] In the wake of the 1979 gas price increase, GM President Elliott M. Estes noted that consumer demand placed pressure on firms to offer more fuel efficient choices. He stated that consumers were reasserting themselves as the auto companies' "No.1 taskmaster".

[7] Most notably, fuel efficiency is increased by decreasing engine power or vehicle weight. Engine power provides performance, and vehicle weight can provide luxury and/or safety.

[8] Segments are vehicle designations such as Sport Utility Vehicle (SUV) or Small Car.

[9] Price is correlated with "unobserved" portions of cost and quality. "Unobserved" means unobserved to the econometrician, but still observable to market participants. Allowing for this correlation has improved model fit and predictions.

firms must commit to product planning in a stochastic environment. They choose characteristics before the year of sale, a this means they do not know the gas price that will prevail during the year of sale. This has the potential to cause ex-post regret in characteristic choice.[10] The timing of the model implies that this regret is not foreseeable, systematic, or predictable. The regret is uncorrelated with anything known at the time the decision was made(Hansen and Singleton, 1982). If the regret were systematic, that would imply sub-optimal decision making (inconsistent with the model) rather than decision-making in a stochastic environment (consistent with the model). This implication of the timing becomes an estimation moment and allows me to relax the implausible assumption that unobserved qualities are orthogonal to observed characteristics. Given that I allow correlation between product characteristics and unobserved qualities, I report these correlations in the estimation results and show them to be quantitatively significant.

A third contribution of the model relates to the representation of consumer preferences for fuel efficiency. Previous models, including seminal work,[11] have had some difficulty finding parameter estimates that show consumers care about fuel efficiency. Parameters on preference for fuel efficiency have been biased towards zero. The reason for this is that in automobiles, fuel efficiency (*mpg*) is negatively correlated with other characteristics that provide utility. Some of these characteristics are observed and easily controlled for, such as horsepower and weight. Others, however, are not. Engine characteristics such as timing of transmissions and gear ratios can also affect the fuel efficiency v. performance tradeoff but are harder to capture in data. I correct for this aspect of unobserved product quality by controlling for both the economic and quality effects of fuel efficiency in consumer preferences. Unlike previous work, my resulting parameter estimates indicate that consumers care strongly about fuel efficiency.[12] There are quality tradeoffs in providing fuel efficiency, even beyond the characteristics most commonly observed and controlled for in automotive data. I propose a method for controlling for these unobservables.

The rest of the paper is structured as follows. Section 2 provides a brief discussion of two strands of literature to which this analysis is related. In Section 3 I describe the model. This includes the market participants, their objective functions, and the timing of the game. In Section 4 I discuss the data and industry. Section 5 discusses estimation - both the estimation moments and the estimation methods used. Section 6 presents estimation results and Section 7 discusses various robustness tests. Section 8 discusses the model's predictions for counterfactual scenarios of market equilibrium at various after-tax gasoline prices. I analyze summer 2008 gasoline prices, and the after-tax price of gasoline needed to achieve 35 *mpg*, the goal of recent CAFE proposals. (The level is $4.55). Section 9 concludes.

[10]Note this is not ex-ante regret. Ex-ante regret would mean firms are making sub-optimal decisions, but ex-post regret does not. Ex-post regret simply means firms must optimize in an uncertain environment. The uncertainty for auto manufacturers in this context is the gas price.

[11]Berry, Levinsohn, and Pakes (1995).

[12]I discuss the magnitude and distribution of these preferences further in the estimation results (Section 6).

2 Relation to Literature

This paper is related to literatures on both the automotive industry and endogenous product selection. The automotive literature has recently focused on endogeneity of prices [Berry (1994), Berry, Levinsohn, and Pakes (1995)], as well as various policy questions [Goldberg (1995), Gruenspecht (1982), Berry, Levinsohn, and Pakes (1999), Kleit (1990) and numerous others]. This paper investigates a similar question to Pakes, Berry, and Levinsohn (1993) (the effects of a gas price increase on the automobile market) but adds a model of product choices in order to analyze fuel efficiency change.[13]

This work is also related to a growing literature on endogenous product selection. This literature reflects the importance of changes to product characteristics, not just changes to prices. The industries and applications within which this question have been studied are various, and the modeling obstacles and adaptations are as varied. Some examples are early work in this area by Mazzeo (2002) which studied binary and trinary entry decisions in the motel market. Lustig (2008) uses cross sectional variation in market structure to investigate health insurance quality. Sweeting (2007) uses a dynamic framework to analyze the radio industry, and Crawford and Shum (2007) use the theory of monopoly screening to study choices in the cable television industry.

The aim of this paper is to bring the two literatures together. There are important product choices in the automobile industry that affect energy consumption. My goal is to advance our understanding of how these choice are made, and how they might respond to economic and policy changes.

3 Model

3.1 Demand

Consumer demand is based on the unit of the household. Each year, U.S. households take as given both the gas price and the product offerings of firms. They choose a new automobile (or no new automobile - the outside good) to maximize a conditional indirect utility function:

[13] Modeling choices of non-price attributes is essentially new to the automobile literature. Goldberg (1998) models the proportion of domestic (vs. foreign) production in response to CAFE standards.

$$u_{ijt} = u(p_j, econ_j, qual_j, X_{jt}, \xi_j, \tilde{\epsilon}_{ijt}) \tag{1}$$

The subscripts are i for individual, j for automobile model, and t for time (year).[14] I use the terms "model," "product," and "vehicle" (subscript j) interchangeably. Vehicle "type" (v), vehicle "segment" (s), and vehicle "sub-segment" (ss) refer to specific levels of auto classification that come in my data.[15] These are displayed in Figure 2 at the end of the text. Throughout the discussion, scalar quantities (such as p_j) will be plain text, while vector quantities (such as $\mathbf{p}$) will be in bold.

The utility specification in Equation (1) indicates that within an automotive sub-segment, consumers have preferences over price (p), fuel economy ($econ$), and "other quality" $qual$. Other quality includes things such as power, weight, acceleration, electronics, sportiness, interior room, etc. X contains terms controlling for sub-segment, vehicle origin, and macroeconomic variables such as GDP growth. The macroeconomic variables are intended to capture the utility of the "outside good" which is purchasing no new automobile. ξ is unobservable quality not captured in the other utility covariates. $\tilde{\epsilon}$ is a nested logit error term.

Fuel efficiency (mpg) affects both the fuel economy ($econ$) and "other quality" ($qual$) of a vehicle through the technological tradeoff. The measure I use for $econ$ is dollars-per-mile (dpm), which is equal to the gas price over the fuel efficiency ($\frac{p_{gas}}{mpg}$).[16] This is a key component of the economic cost to operating a vehicle.[17] This form implies the natural result that the marginal benefit of fuel efficiency, $\frac{\partial u}{\partial mpg}$, is increasing in the price of gas.[18] The measure I use for $qual$ is $ln(mpg)$. This may seem counterintuitive to use the very quantity that *trades off with* quality to proxy for quality itself. But it is a very strong empirical regularity that, conditional on both dpm and sub-segment, higher mpg is associated with lower "other quality" of a vehicle. This tradeoff can be pictured with a technology frontier for a given manufacturer within a given sub-segment, Figure 1 in the text. I will discuss first the concept of the frontier, then the suitability of $ln(mpg)$ as a proxy for the quality axis.

[14] Any j subscript could itself be subscripted by t because model attributes change from year to year, but I suppress this notation for ease of exposition. Where t subscripts do appear, they are to emphasize that some variables are common across all j in a given year (such as macroeconomic variables in X_{jt}) or to emphasize the frequency of some errors or decisions (such as $\tilde{\epsilon}_{ijt}$).

[15] t would be the natural subscript for vehicle "type," but it already subscripts time.

[16] Note the distinction between the technological parameter fuel **efficiency** (mpg) and the consumer preference parameter fuel **economy** (dpm).

[17] I do not separately model consumers' vehicle utilization decisions.

[18] I.e., $\frac{\partial^2 u}{\partial mpg \partial p_{gas}} > 0$.

Figure 1: Technology Frontier Within an Automobile Sub-Segment:

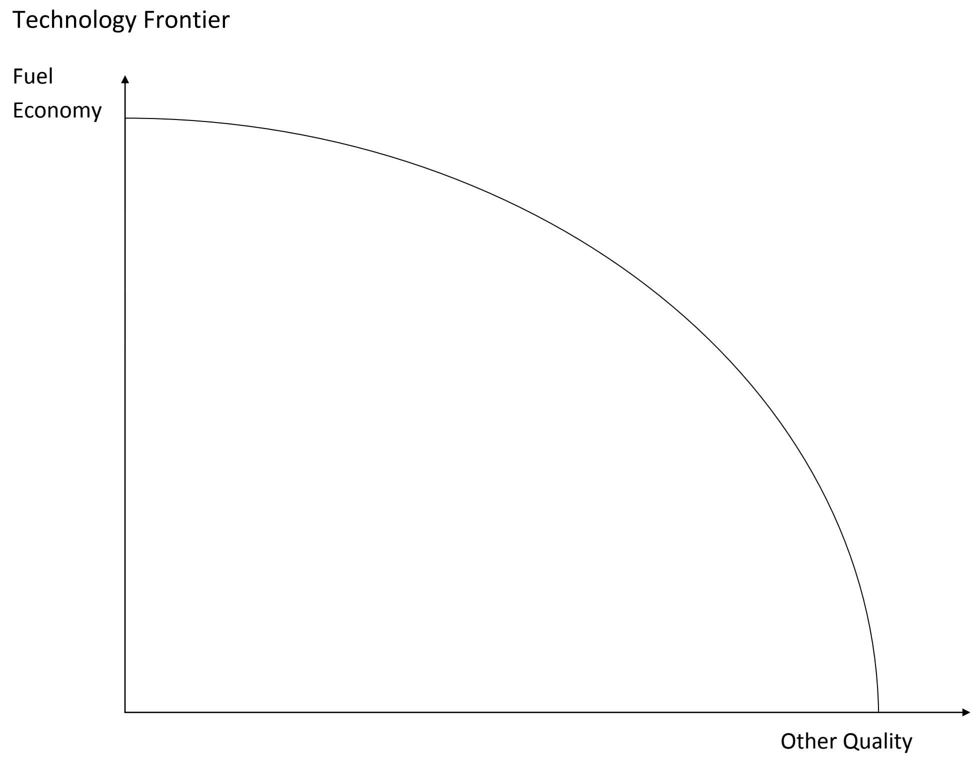

This frontier reflects the tradeoffs (within a particular sub-segment) between providing efficiency and other attributes such as power (performance) and weight (luxury/safety). This tradeoff has been documented in previous work (Kleit, 1990). It is also evident in Figure 3 at the end of the text, which shows the relationship between $ln(mpg)$, $ln(hp)$[19], and $ln(weight)$ in my data. The tradeoff between $ln(mpg)$ and (these two) attributes of quality holds true at any level of vehicle aggregation, though for purposes of my model it is important that it holds within a sub-segment (the lower right panel of the Figure). Table 4 shows this relationship in the form of a regression. $ln(mpg)$ is regressed on $ln(hp)$, $ln(weight)$, and a time trend to capture advancing technology. The regression shows a strong negative relationship between fuel efficiency and both power and weight. The regression numbers in the Table are pooled over all vehicle types, but at finer levels of aggregation the results change little. Despite the decreasing numbers of observations, the coefficients remain large and significant, and the R-squared generally remains above .5. R-squared remains high despite decreasing numbers of observations because the technological relationship captured is increasingly precise.[20] In my model

[19] hp is horsepower, a measure of engine power.

[20] The regression in Table 4 is of equilibrium choices of these three magnitudes by firms. Therefore, by itself the regression is not conclusive regarding the underlying production function. However, given the competitive structure in the automobile industry the relationship does suggest a tradeoff. Within an automotive sub-segment, firms cannot

I assume this tradeoff between fuel efficiency and other quality holds exactly within a sub-segment.

There are improvements that can be made to vehicle power and/or weight which sacrifice less fuel efficiency (or vice-versa). However, they are expensive and amount to changing the sub-segment of the vehicle. Luxury cars will have more horsepower than a mid-sized car with similar fuel efficiency, but it is much more expensive to produce the luxury car and it therefore does not compete in the same sales segment. Another way to think of the technology tradeoff is that isocost curves exhibit this shape, and competition within a sub-segment occurs along a single isocost curve. A firm can, to some extent, shift out the frontier, but only by incurring higher costs and switching segments of the market.

Independent of the question of whether technology tradeoffs exist, there is the question of whether $ln(mpg)$ is an appropriate proxy for "other quality." It seems counterintuitive because these would appear to be opposite axes of the same technology tradeoff. But there is reason to use $ln(mpg)$ to capture quality. The reason is that, *conditional* on *dpm*, $ln(mpg)$ is the best empirical proxy for the quality tradeoffs. Any demand specification I estimate, no matter what observables are controlled for, puts a negative and significant sign on $ln(mpg)$ (*provided dpm* is controlled for). This means there are attributes of quality that are a) not measured in the dataset, and b) negatively correlated with $ln(mpg)$. Trying to construct a metric for other quality based on all non-*mpg* observables misses some of these.[21] One implication of this for future work is that specifications of automobile demand ought to include both *dpm* and $ln(mpg)$, not simply one nor the other. (At least until the datasets are more rich).

As a final note on the subject, there is both conceptual and empirical reason to choose $ln(mpg)$, rather than simply *mpg* or any other functional form, to represent "other quality." Conceptually the quality tradeoffs required to move from 10 *mpg* to 20 *mpg* are larger than those required to move from 40 *mpg* to 50 *mpg*. The decreases in quality are decreasing in constant *mpg* improvements. Another way to think of this that increasing quality is more costly (in terms of efficiency sacrificed) as quality levels become higher. The second, empirical, reason to favor the logarithm is that it insures the relationship between quality and efficiency always stays negative, whereas higher order expansions do not.[22]

push past the frontier without incurring extra costs and essentially changing the sub-segment within which the vehicle competes. If firms could attain more of both without a tradeoff, and/or without incurring costs, then the regression relationship would not hold.

[21]These unmeasured qualities likely include engine characteristics that can tweak *mpg*, conditional on power and weight. These may include timing of transmissions, gear ratios, turbo charging, cylinder shutdown, variable valve timing, etc.

[22]The second order expansion bends back on itself around 55-60 *mpg*, and the third order expansion does reverse this because there is not enough data in that *mpg* range to calibrate it. So for high levels of *mpg*, increasing *mpg* improves quality and this leads to non-sensible predictions in the counterfactual.

Given the above definitions, and with an additional linearity and separability assumption, consumer utility in Equation (1) becomes:

$$u_{ijt} = \alpha p_j + \beta_d dpm_j + \beta_m ln(mpg_j) + \xi_j + BX_{jt} + \tilde{\epsilon}_{ijt} \tag{2}$$

I expect $\alpha, \beta_d, \beta_m < 0$. The nesting structure designations in Figure 2 come from Wards Automotive. The structure is a 3-level nest, where each level is an increasingly specific level of automobile classification. The top level, vehicle type (v), are Cars, Utility Vehicles, and Trucks/Vans. The next level, segment (s), consists of designations such as Small Car and Sport Utility Vehicle (SUV). The third and most specific level, sub-segment (ss) designates categories such as Lower Small Car, or Luxury SUV. Finally, the model (j) designate names such as Honda Accord or Ford Explorer.

The nested logit error term $\tilde{\epsilon}$ contains parameters to be estimated ($\sigma_{ss}, \sigma_s, \sigma_v$) which describe the similarity of choices within a nest. To be consistent with utility maximization, these nesting parameters must be between 0 and 1 (Cardell, 1997). Appendix A contains details on the calculation of nested logit shares and σ's. The nested logit does restrict cross-price elasticities in a specific way.[23] This means I do not measure substitution patterns as flexibly as if I were to forego the nesting structure and include more utility covariates. As I will discuss in the supply side, however, the nested logit structure is instrumental in enabling my model of endogenous characteristic determination, and past empirical work in the auto industry has used the nested logit with some success (Goldberg, 1995). So I find the tradeoff a reasonable price to pay.

Note that similar to other applied work, ξ_j from Equation (2) will be used in estimation moments. However, as I will also discuss in the supply section, I will use less restrictive assumptions on its distribution.

3.2 Supply

The timing of the supply side model is depicted in Figure 4 at the end of the text. One year before the year of sale,[24] firms observe gas prices, cost shocks, and demand shocks. In response they commit to product characteristics for each vehicle, still one year in advance, to allow for a production lag. Then, in the interim, gas prices change. Finally, in the year of sale, firms choose

[23]Choice probabilities within a nest, for the choices down one level, are logit based on inclusive value.

[24]As a robustness check I also estimate the model assuming 3 and 5 year lags in Section 7.

a price for each vehicle and consumers make purchase decisions. The information sets (H and I) as well as the ex-post regret (R) will play a role in the estimation moments as discussed in Section 5. Note that this timing allow firms to observe their cost and demand shocks when choosing their product characteristic, mpg, which is plausible but not allowed in previous models of the automobile industry.

Firms take actions to maximize expected profit. Firms' two choice variables (mpg and p) form a subgame-perfect Nash equilibrium. Each firm f's problem is:

$$\max_{\mathbf{mpg_f}} E_{p_{gas}} \left[\max_{\mathbf{p_f}} \Pi_f \right] \tag{3}$$

Firm f's profits are given by:

$$\Pi_f = \sum_{j \in \Im_f} M_t s_j(\mathbf{mpg}, \mathbf{p}; \theta) \left[p_j - mc_j(mpg_j; \theta) - \lambda_{ft}\{bind_{ft}\}(CAFE_t - \overline{mpg_{ft}}) \right] \tag{4}$$

Subscript f is for firm, and recall that j is a model. $\Im_f$ is the set of cars produced by firm f, as auto manufacturers are multiproduct firms. M_t is the size of the market (the number of U.S. households), s_j is the market share of model j, and is a function of **mpg** and **p**, the entire industry vectors. θ is a vector of demand and cost parameters to be estimated. $CAFE_t$ are the stipulated CAFE standards, and $\overline{mpg_{ft}}$ is the firm's average fuel efficiency for determining CAFE compliance.[25] $\{bind_{ft}\}$ indicates whether the standards are binding for the firm, and λ_{ft} is a firm and time-specific shadow cost of violating the standards. Recall that the t subscripts have been suppressed in a number of places for ease of exposition.

Because of the Nash equilibrium assumption, derivatives of the profit function (with respect to the choice variables) lead to first-order conditions of optimality. These hold with respect to the both choices, p_j and mpg_j, for each model j:

[25] CAFE standards calculate firm fuel efficiency as a harmonic average rather than a raw average. $\overline{mpg_f} = \frac{q_f}{\sum_{j \in \Im_f} \frac{q_j}{mpg_j}}$. Harmonic averages are used to average rates, as they indicate the rate of fuel consumption when all automobiles are driven the same distance. When averaging any (unequal) numbers, harmonic averages are mechanically lower than raw averages, since low rates pull down more than high rates push up.

$$s_j + \sum_{r \in \Im_f} \left[(p_r - mc_r) \frac{\partial s_r}{\partial p_j} + \lambda\{bind\} \frac{\partial \overline{mpg_f}}{\partial p_j} \right] = 0 \tag{5}$$

$$E_{p_{gas}} \left[\sum_{r \in \Im_f} \left[s_r \left(\frac{\partial p_r}{\partial mpg_j} - \frac{\partial mc_r}{\partial mpg_j} \right) + (p_r - mc_r) \frac{\partial s_r}{\partial mpg_j} + \lambda\{bind\} \frac{\partial \overline{mpg_f}}{\partial mpg_j} \right] \right] = 0 \tag{6}$$

Note that a single choice of *mpg* along the sub-segment frontier, or isocost curve, determines both the fuel efficiency ($dpm_j = \frac{p_{gas,t}}{mpg_j}$) and other quality ($ln(mpg_j)$) for consumer demand. A challenge in modeling product choice in any industry is finding exogenous variation to shift incentives to offer characteristics. Characteristics-based models of auto demand have tended to include approximately six preference characteristics of automobiles. Because choices of these attributes are interrelated (for example increasing fuel efficiency by decreasing engine power), a model that allows firms to choose any of these characteristics must allow firms to choose them all.[26] Having said this, however, it is difficult to find instruments to shift firms incentives' to provide all six characteristics in the auto industry. Some industries have the benefit of having many local, geographic markets in the cross section that provide variation in market structure (e.g. Lustig (2008)) to shift incentives. However, automobile supply is one, national market each year, so any shifter must have national time-series variation.

My solution to this problem is to use the nested logit structure, the technology tradeoff, and an assumption of medium run time horizons. First, the nested logit structure captures consumer substitution patterns without having to use a full set of characteristic controls. Consumers in my model are free to have preferences over other/all auto characteristics, but I do assume these preferences influence the sub-segment choice itself, not the choice within sub-segment. The choice within sub-segment is based only on two non-price characteristics (*econ* and *qual*).[27] Having used the nests to reduce the endogenous choices to two, I use the technology tradeoff to parameterize these two choices into one choice. Third, I restrict firms from changing the sub-segment of their vehicles. This essentially amounts to assuming a medium run horizon of perhaps 5 years. Changing the sub-segment of a car essentially means rebranding and reintroducing, and this is likely to be difficult to do in under 4-5 years.[28] So these three modeling features - the nesting structure, the technology tradeoff, and the medium run focus - allow me to pare down the endogenous choice variables to a single choice *mpg* (which determines the two utility characteristics), and there is a clear stochastic and exogenous

[26] Allowing firms to choose only some characteristics, while treating others as exogenous, would bias measured incentives.

[27] Which is more general than saying that competition throughout all vehicles is only based on these two characteristics.

[28] On the other hand, modifying an engine to gain/lose fuel efficiency is something that can, and does, happen more frequently.

shifter for this characteristic choice is the gas price.

There are costs of this approach. Though the nested logit captures some relevant consumer behavior, there is likely loss of richness in the demand system by not retaining more preference characteristics. Also, because I limit my attention to the medium run, there are interesting longer-run decisions and dynamics which I will not capture in the model. I find think these costs are worth the price to pay for being able to model fuel efficiency choice. Nested logits work better when the more nesting groups are relevant to consumers' decisions, and there is some indication in the trade press and previous empirical literate[29] that this is at least somewhat true with autos. In terms of the medium-run restriction of the analysis, this is at least an improvement on previous work which was unable to analyze fuel efficiency choice on any time horizon. I trust that future work will develop methods to look at issues of new model introduction.

Firms face constant marginal costs of producing automobile j:

$$ln(mc_j) = mc(mpg_j, X_{jt}, \omega_j, t) \tag{7}$$

The specific functional form I use is:

$$ln(mc_j) = \gamma_1 + \gamma_2 ln(mpg) + \Gamma X_{jt} + \omega_j \tag{8}$$

Because firms locate along an isocost curve, the standard interpretations of γ_2 does not necessarily apply. Normally in empirical work a term like γ_2 is expected to be greater than zero because it is attached to a good. But here, $ln(mpg)$ is a parameterization of a tradeoff between 2 goods, rather than a good itself. Within a sub-segment, cars with relatively more fuel efficiency may or may not be more expensive to produce. In fact, the better the model approximation of reality, the more likely that $\gamma_2 = 0$. The model does not impose this restriction, but it is an implication check of the model.[30] X is a full set of sub-segment dummies, country of origin, and a time trend. ω_j is the unobserved product cost (unobserved but inferred from the data).

Note that ω_j, like ξ_j in demand, are used in estimation moments. However ω_j, also like ξ_j, is used in a less restrictive way than in previous work.

[29]Goldberg (1995) and Goldberg (1998).

[30]If my proxies for economy and quality closely match consumer valuations, and competition within the sub-segment is nearly on a single-isocost curve, then γ_2 should estimate out to 0.

3.3 Modeling CAFE Standards

The federal government's fines for violating CAFE standards in a given year are \$55 per car sold per 1 *mpg* below the standard that the firm average falls. It is commonly thought, however, that there may be additional reputation and political costs to American manufacturers of violating CAFE standards. The Big Three[31] have influence over the CAFE standards through their relationship with the federal government. One thought is that these manufacturers may risk this influence by violating the standards, or might lose reputation with American consumers by violating our country's own domestic energy policy. I leave the fine (λ in Equations (5) and (6)) as a shadow cost to be estimated. The estimation results can thus speak to whether the implicit costs are higher than the fine alone.

For simplicity, I assume CAFE binds on all American manufacturers in all years after 1977, and on no other manufacturers in any years. This is not exactly true, but is a reasonable approximation of both a somewhat complicated system of carry-overs and credits used in calculating CAFE fines, and historical evidence on sales-weighted fuel efficiency. For Asian manufacturers, CAFE has not bound. They have never paid fines, and have never been particularly close to the standards. Some European firms have paid \$55 per car-*mpg* fines. (In fact, European firms are the only firms to have ever paid CAFE fines.) However, for these European firms there is likely less political pressure to abide by U.S. domestic energy policy than there is for U.S. firms. These fines to European manufacturers have tended to be on small fleets of luxury and sports cars where the pecuniary penalties are a small portion of the list price. I simply add them as an extra tax to these models, rather than estimate a shadow cost of compliance. American firms, in contrast to both Asians and Europeans, are likely to have been genuinely constrained over the years. They face unique pressure in being the domestics, and they have hovered close to the regulation for many, many years. No American manufacturer has ever actually paid a fine, but this is only due to the system of carry-overs and credits from year to year. The standard looks to have been binding or near binding for the whole time period since 1977 when the standards began. The model is general enough to handle a rich set of year-specific, firm-specific shadow costs, but I estimate only one shadow cost. Interpreting yearly costs would be muddy because of the system of carryovers, and each shadow cost is an extra non-linear search parameter for the GMM search routine. So for computational reasons fewer terms is preferable. I simply assume one shadow cost on all 3 American manufacturers, starting in 1977.

[31] General Motors, Ford, and Chrysler are the three American manufacturers.

4 Data & Industry

The model is estimated using data on all new automobiles sold in the U.S., and macroeconomic data, from 1971-2007. I have data on all sales and characteristics of passenger vehicles during that time period with two exceptions. First, I'm missing data from 1991-1995. The collection of the earlier data stopped in 1990, and the electronic data-keeping by Wards Automotive did not begin until 1996. Second, I'm missing truck data for the early years (1971-1990), which were not collected with the original data set.[32]

The early year data, 1971-1990, are the same data used in Berry, Levinsohn, and Pakes (1995), and were generously shared with me by the authors. These data are from Automotive News Market Data Book, and contain information on sales and base model characteristics of all passenger cars sold in the U.S. I supplement these with segment information that I collected from Wards Automotive Yearbooks. The later data, 1996-2007, are from Wards Automotive. They contain sales, characteristics, and segment data on all passenger vehicles sold in the U.S. during those years. Table 5 in the back of the text shows popular cars in the various sub-segments for the later years.

Prices are list prices, fuel efficiency is city fuel efficiency as measured by the EPA (U.S. Environmental Protection Agency). The unit of observation is the vehicle model, the level at which sales data are reported. I use the characteristics of the base model for each model-year.[33]

I also collect macroeconomic data from various government agencies: the Energy Information Administration (gas prices), Bureau of Labor Statistics (CPI, number of households, unemployment, income distribution), and the Bureau of Economic Analysis (GDP and GDP growth).

Table 6 contains selected summary statistics from all data sources. Sample average *mpg* is 20.7. (At a finer level of detail this would be higher for Cars and lower for Trucks and Utility Vehicles.) There is substantial variation in the gas price, both in levels and in changes. Note that year-to-date through August 2008 (outside the sample summarized in the Table) the gas price jumped from $2.80 to $3.43, which is above the highest real gas price of the sample. This jump also nearly equalled the highest change of the sample, which is $0.68 in 1980. The summer of 2008 is analyzed in the counterfactual Section (8). There are 4,820 model-years over the course of 32 years, for an average of 151 models offered each year. There are 3 vehicle types, 9 vehicle segments, and 28 vehicle sub-segments.

[32] I intend the collect the missing years of data to supplement the analysis.

[33] Models have even finer distinctions of various trim levels. These offer slightly different configurations of engines and accessories. I ignore this in my analysis, as I do not have sales data at this level. The base model is usually the most fuel efficient trim of each model.

Figure 5 shows some trends in the industry over time. The number of U.S. households (M, the market size) is the bottom line in the Figure. It has grown at a relatively constant rate throughout the sample. Aggregate sales (middle line) has grown at more or less the same rate, albeit with substantially more variation around the trend. The number of vehicles (top line) has grown at a faster rate than the market size. The number of model offerings (and even segments) have proliferated throughout the sample period.[34]

Figure 6 shows the industry's response to the gas price and regulation. The gas price is the solid line which moves by itself. The two dotted lines are the CAFE standards for passenger cars and light trucks (the standards apply separately for these groups). The solid lines which track them are the respective sales-weighted fuel efficiencies for the vehicle types. There have been 3 major gas price increases in the sample: 1973 (the OPEC embargo), 1979 (the Iranian Revolution), and 2007-2008 (supply interruptions and growing global demand). There have also been periods of historically low fuel prices in the late 1990's and early 2000's. This amounts to large variation in the gas price over the sample. CAFE standards were enacted in 1975 in response to the 1973 energy crisis. They imposed a standard which rose gradually until the early 1980's, and has remained largely unrevised until 2008. Recently the federal government has stipulated higher standards by 2020, and a stricter treatment of SUVs and trucks in calculating averages.[35] Note that average fuel efficiency has tended to outpace regulation during hight gas prices (1979 and into the 2000's), but the regulations appear to have mattered during low gas prices (in the mid 1980's and 1990's). This is especially true for American manufacturers, although that is not directly observable from this aggregated chart.

Figure 7 shows yearly characteristics for one long-lived base model in the sample, the Toyota Celica. There is some pattern of inverse movements between *mpg* and both *hp* and *weight*, as well as positive movements between *mpg* and the gas price. However, these patterns do not necessarily move in lockstep. Do note also that product characteristics change quite frequently, so it is reasonable to think that firms are able to change these characteristics yearly by making adjustments to engines, materials, amenities, etc.

My utility specification (Equation (2)) effectively assumes that consumers expectations of the future gas price are summarized by the current price of gas. The reason I assume this is that gas prices are difficult to predict. Oil futures are a relatively recent invention, they do not exist for time periods longer than 12 months, and are usually quite close to the spot price. The spot price incorporates a large amount of the available information about future oil prices. Consumers may still, of course, react to oil prices in other ways besides considering current levels, but Busse, Knittel, and Zettelmeyer (2008) show that sales-weighted fuel efficiency purchases track quite closely with the

[34]The entire SUV and CUV segments did not exist before the mid 1990's, but in 2007 comprised 30% of sales.

[35]All firms' vehicle averages, both for passenger cars and for light trucks, must exceed 35 *mpg* by 2020 (and proposals being considered could accelerate this to 2016). As well, vehicles over 8,500 lbs. of gross vehicle weight are no longer exempt from the standards. The new exemption is 10,000 lbs.

current gas price, not outpacing it nor lagging it. So this is the specification I use in the estimation results, and I discuss other specifications in the robustness Section (7).

5 Estimation

The estimation routine is Generalized Method of Moments, and the moments are constructed from the timing of the model.

5.1 Estimation Moments I

To help understand the estimation moments, it is useful to refer again to Figure 4, the model timing. The two information sets defined in the Figure are snapshots of information at various points in the year leading up to production. These information sets, when they are visible, are visible to all market participants. Loosely speaking, these sets are:

$$H_{t-1} = \{\text{Information known about time } t \text{ at the ``beginning'' of time } t-1\} \tag{9}$$

$$I_{t-1} = \{\text{Information known about time } t \text{ at the ``end'' of time } t-1\} \tag{10}$$

Take the example year t=2008. Let H_{t-1} be the knowledge at the beginning of 2007, before 2008 characteristics are chosen, about what the auto industry will look like in 2008. This includes the set of all 2008 models, their manufacturers, their sub-segments, and the 2007 gas price. H_{t-1} does *not* include the cost and demand shocks, *mpg* choices, or gas prices and macroeconomic variables that will prevail in the 2008 marketplace.

Then, still in 2007, the cost and demand shocks, ω and ξ, that will affect the 2008 market are revealed to all manufacturers on all models. They are common knowledge. In response to these, and still in 2007, *mpg* is chosen for all 2008 models. Now take another snapshot and redefine this information set as I_{t-1}. It contains H_{t-1}, but is also updated to include shocks and *mpg*. It still, however, does not include the final pieces of the 2008 market - gas prices and macroeconomic variables. So, more specifically:

$$H_{t-1} = \{1, p_{gas,t-1}, \text{ownership/sub-segment of year } t \text{ models}\} \tag{11}$$
$$I_{t-1} = \{1, p_{gas,t-1}, \text{ownership/sub-segment of year } t \text{ models}, \omega_t, \xi_t, mpg_t\} \tag{12}$$

ω_t, ξ_t, and mpg_t are vectors for the entire industry, observed by all participants. The inclusion of "1" in these information sets simply indicates that anything (shocks or ex-post regret) orthogonal to these information sets should be mean zero.

Two estimation moments that are implications of this timing (rather than new assumptions), are that demand and cost shocks are uncorrelated with H_{t-1}:

$$E[\xi_{jt} * H_{t-1}] = 0 \tag{13}$$
$$E[\omega_{jt} * H_{t-1}] = 0 \tag{14}$$

These moments hold model by model - no vehicles' demand or cost errors should be correlated with the exogenous observables. Errors are always orthogonal to observables in empirical work - the logic of this assumption is standard. However the formulation here is less restrictive because I exclude the endogenous choice variable *mpg* from the instrument set. This means that in my model firms are able to observe market shocks when choosing their *mpg*, and this allows for a free correlation between errors and *mpg*. The standard demand identifying assumption would be to add mpg_t into what I call H_{t-1}, which would mean firms are unaware of the cost and demands shocks when they commit to characteristic decisions. This scenario is implausible in the auto industry because automobile models survive multiple years and there is learning. For example, the Toyota Camry sells quite well relative to other cars with comparable characteristics. This indicates a high ξ for the Camry, and all auto firms are likely to be aware of this when they choose their own product characteristics.[36]

[36] In addition to observing the shocks, firms might actually be choosing these shocks, or at least a portion of them. In the automotive industry, firms likely choose a portion of these shocks. I do not allow for this in my model, however, because doing so would require additional assumptions in order to split a single unobservable quantity into two unobservable portions (chosen and unchosen), and/or additional exogenous variation as an instrument. This is an interesting pursuit for future work.

5.2 Estimation Moments II

The second set of estimation moments are, like the first set, implications of the timing of the model rather than new assumptions. They are the first-order conditions on optimal firm choice of p and mpg in (see Equations (5) and (6)). Because the first-order condition for mpg has an expectation, the empirical counterpart will not hold exactly equal to zero. I have two choices as to how to handle the wedge. One is that I can simulate the empirical distribution of gas price change to approximate the integral in the expectation, and make this integral hold equal zero. Because this is computationally burdensome, and requires me to commit to my belief as to firms' beliefs concerning gas prices, I instead use a second method outlined in Hansen and Singleton (1982).

To understand the Hansen and Singleton moments, define:

$$R_{jt} = \frac{\partial \Pi_{ft}}{\partial mpg_{jt}} \tag{15}$$

$$R_{jt} = \sum_{r \in \Im_f} \left[(p_r - mc_r)\frac{\partial s_r}{\partial m_j} + s_j\left(\frac{\partial p_r}{\partial m_j} - \frac{\partial mc_r}{\partial m_j} + \lambda\{bind\}\frac{\partial \overline{mpg_f}}{\partial m_j}\right)\right] \tag{16}$$

R_{jt} is the derivative of the profit function with respect to mpg_{jt}. It is the ex-post regret in mpg choice. In Figure 4, this appears at the extreme right of the timetable. It is the object which is set to zero *in expectation* in the firms' first-order condition. But because the gas price changes after mpg_{jt} is chosen, R_{jt} will not generally equal zero in the year of sale. It represents the ex-post regret (thus the choice of the letter R) in the mpg_{jt} choice. Another way of saying this is that if firms could readjust mpg_{jt} during the year of sale, they would. This is *not* the same as saying that firms fail to optimize. Firms do optimize, they simply must do so at a time when not all the information they would wish to see (the gas price) has been revealed.

With this definition, Hansen and Singleton (1982) note that an implication of stochastic optimization is:

$$E[R_{jt} * I_{t-1}] = 0 \tag{17}$$

This says there are no patterns of ex-post regret (R_{jt}) which are systematic, or in any way predictable at time $t-1$. If there were such patterns, then this would, indeed, imply that firms are not optimizing correctly. That would be inconsistent with the model.

The benefit of this estimation approach is twofold. First, it is far less computationally burdensome than simulating the empirical distribution of gas price changes. Second, it allows me to remain agnostic about firms' beliefs as to future gas prices.

As I have mentioned, I do experiment in the robustness section with placing I_{t-3} and I_{t-5} into Equation (17) to reflect the possibility that firms have to choose product characteristics more than 1 year before the year of sale. Note that all that is different in these two information sets is that the older gas prices ($p_{gas,t-3}$ and $p_{gas,t-5}$, as opposed to $p_{gas,t-1}$). I discuss the results of this robustness test further in Section 7.

5.3 Estimation Mechanics

Three moments, (13), (14) and (17) are stacked for estimation. The fourth equation, the pricing first-order condition (in Equation (5)), holds exactly for any parameter values because of the continuous functions $\xi(\theta)$ and $\omega(\theta)$ (Berry, 1994). So that equation does not enter estimation as a moment, though it is used within each moment evaluation to back out the error terms. Estimation is by 2-stage GMM (Generalized Method of Moments). Parameter values are chosen to minimize the following objective function:

$$min_\theta \begin{pmatrix} \xi_t(\theta) * H_{t-1} \\ \omega_t(\theta) * H_{t-1} \\ R_{jt}(\theta) * I_{t-1}(\theta) \end{pmatrix} W_n \begin{pmatrix} \xi_t(\theta) * H_{t-1} \\ \omega_t(\theta) * H_{t-1} \\ R_{jt}(\theta) * I_{t-1}(\theta) \end{pmatrix}' \tag{18}$$

There are only 4 parameters that require non-linear search - the three nesting parameters (σ's in the utility function), and the shadow cost on CAFE compliance (λ). Conditional on these 4 parameters, the rest of the parameters in θ can be calculated by Two Stage Least Squares to minimize the value of the moments. I choose starting values for the 3 nesting parameters based on Two Stage Least Squares estimates of demand with just the first two moments. I choose starting values for the CAFE standard based on the pecuniary fine \$55, though the search routine is not sensitive to any of these starting values.[37]

In the first stage of GMM, W_n is a block diagonal matrix containing the inverse of the normalized instrument matrices H_{t-1}, H_{t-1}, and $I_{t-1}(\theta)$. The only non-standard part of the routine is that weight matrix is updated for each parameter guess, because I_{t-1} is a function of θ by virtue of containing the elements $\xi(\theta)$ and $\omega(\theta)$. In the second stage of GMM, W_n is the standard block diagonal matrix containing the inverses of the variance-covariance matrices of the moments themselves.

[37] The objective function has a clear global minimum, which the estimation finds.

I do ignore two equilibrium effects in the calculation of R_{jt} (Equation (16)) because they are likely to be small in magnitude and yet add considerable computational complexity. The first is the $\frac{\partial \mathbf{p}}{\partial mpg_j}$ term - the effect that changing one's fuel efficiency will have on the equilibrium prices (of all vehicles) played in the subsequent game. I assume this term to be zero. The second omission is that in calculating $\frac{\partial \overline{mpg_f}}{\partial m_j}$, I ignore $\frac{\partial \mathbf{q}}{\partial mpg_j}$ - the effect that changing one's fuel efficiency will have on equilibrium quantities in the downstream game.[38] In sum, I ignore effects of mpg choices on the equilibrium of the downstream pricing and purchasing game. This means that I am only approximating the subgame perfect Nash equilibrium, but this approximation is a reasonable price to pay for tractability. These terms are likely to be small because of the number of automobiles in the market is large.

6 Estimation Results

The results of the GMM estimation are reported in Table 7. They exhibit sensible and statistically significant parameter estimates. The nesting parameters (σ's) are between one and zero. α, β_d, and β_m, are significant and negative as expected.[39] There are three findings that are of particular interest: the shadow cost of the CAFE standards, the "cost" of providing fuel efficiency, and the results for fuel efficiency preference by segment.

First, the parameter estimate on the CAFE standard indicates that the shadow cost to American manufacturers of violating the CAFE standards is \$347 per vehicle per year. This is significantly larger than the non-compliance fine of \$55 per vehicle per *mpg*. A common hypothesis is that these shadow costs are due to political and reputation pressure, either from the federal government and/or domestic consumers. Whatever the cause, the parameter estimates are consistent with the notion that American manufacturers incur lost profits due to the regulations. This also suggests that during this time period domestic manufacturers would have set lower fuel efficiency in the absence of the CAFE standards. American manufacturers did not actually pay any fines during this time period - the shadow cost estimate is rather an indication that they incurred significant resources trying not to.

The second result of interest is the coefficient on $ln(mpg)$ in cost. It is not insignificant. This result was not imposed by the model, but rather is an implicit test of model specification that turns out favorably. Unlike most empirical cost functions, when firms provide more *mpg* they also provide less quality. They slide along a production frontier, or isocost curve, for the given sub-segment, rather

[38]This term also includes another instance of the first ignored term, $\frac{\partial \mathbf{p}}{\partial mpg_j}$.

[39]Recall that higher dpm means less fuel economy, higher $ln(mpg)$ means less "other quality," and p always enters utility negatively.

than incurring more cost to provide more *mpg* ceteris paribus. The sub-segment dummies do the lion's share of the work on the cost side, not surprising given that competition within sub-segments is along an isocost curve.

The third result of interest is the set of findings concerning preference for fuel economy. Preference for fuel economy (*dpm*) is allowed to vary by segment. Note that all segments have a preference for fuel economy (all 9 coefficients are negative), but some moreso others. Conditional on sub-segment of purchase, Utility Vehicle purchasers are actually the most sensitive to fuel economy. Luxury Car and Van purchasers are the least sensitive to fuel economy. Trucks, and the remaining four car segments are in between.

It may seem counterintuitive that purchasers of Utility Vehicles (Crossover Utility Vehicles, CUVs, and Sport Utility Vehicles, SUVs) are the most conscious of fuel efficiency, because these vehicles (especially SUVs) are rather fuel inefficient. However, it is actually a sensible result, given that it is conditional on sub-segment purchased. A driver saves more money by increasing *mpg* on a gas-guzzler than on an efficient vehicle. Adding 2 *mpg* to the mean SUV (16.2 *mpg* in 2007) saves more gas money than adding 2 *mpg* to the mean middle sized car (28.2 *mpg* in 2007), conditional on equal mileage driven by each car.[40] Despite, and perhaps even because of, UVs being inefficient compared to other segments, competition *within* the segment is quite sensitive to fuel efficiency. Another way to interpret the result is that drivers have chosen this segment for reasons other than efficiency (they may need the space, or prefer the safety or amenities of an SUV), but once in the segment the returns to fuel efficiency are high.[41] In the summer of 2008 when gas prices were quite high, the market saw large declines in SUV purchases relative to other segments. This corroborates, rather than refutes, the notion that UV purchases are efficiency-minded.

I put a time trend on CUV and SUV utility. This is not standard - generally the literature assumes preferences are stable over time. One of the virtues of a hedonic demand system is that it can tease out a stable set of preferences despite large product turnover. However, there is a historical coincidence that makes a time trend on these segments necessary, in my judgement. That coincidence is that the two UV segments were introduced in the mid 1990's, and grew steadily and rapidly in sales (0 to 30% market share within 10 years) at the same time that the gas price was also steadily and monotonically rising. With no time trend in the model, the estimation attributes this correlation to consumers of these vehicles having a strong *dis*like for fuel economy. I think it is more likely the case that utility for these vehicles was growing because consumers were gaining experience and familiarity with the vehicles. So I have chosen to add a time trend to utility to capture experience and familiarity. This restores the natural result that consumers prefer lower gas expenditure, not

[40] This is the reason many countries express fuel efficiency as Volume per Distance rather than Distance per Volume: fuel costs are linear in Vol/Dist, but concave in Dist/Vol. Some have argued that we ought to adapt the same system in the United States(Larrick and Soll, 2008).

[41] $\frac{\partial^2 u}{\partial mpg \partial mpg} < 0$.

higher.

To interpret the magnitudes of these parameter estimates, I have calculated willingness-to-pay (WTP) for fuel efficiency increases in Table 8. WTP depends on a number of factors: segment, starting *mpg*, ending *mpg*, and the price of gas.[42] Concerning segments, I display all 9. Concerning starting and ending *mpg*, I consider a 20% improvement from the segment-average *mpg*. Because the segment averages are different, this means that both the starting *mpg*, and the absolute *mpg* improvement, are different across segments. This is still the fairest comparison, because the alternative is to compare a single *mpg* jump (such as 20 to 24 *mpg*) across all segments. But any single jump means very different things across segments. Lastly, concerning gas prices, I choose three gas prices at which to display WTP. These are $1.50, $3.43, and $4.55. $1.50 is a relatively low gas price, $3.43 was the price from the summer of 2008, and $4.55 is the gas price we'd need to achieve 35 *mpg* as an industry. ($3.43 and $4.50 are the two prices that I feature in the section on policy implications (Section 8)).

As anticipated by the unequal segment preferences for *dpm*, there are unequal segment WTP for *mpg*. Within a segment, higher gas price always means higher willingness to pay. At $3.43 gas, SUV consumers are willing to pay $7,059 for the 20%-from-average *mpg* improvement. CUV consumers $5,246. Luxury car purchasers were almost indifferent to fuel efficiency at average levels, and most segments' consumers were willing to pay between $1,000 - $3,000 for the 20% increase.

Note that for low gas prices WTP for higher fuel efficiency can be negative. This is because of the quality tradeoffs associated with higher efficiency. At a high enough gas price, all WTPs are positive, and at a low enough gas price all WTPs are negative. Table 9 shows, for comparison, the willingness to pay for efficiency improvement *without* quality loss. This Table does not reflect how the model works, nor how I believe the real industry works, but I include it for comparison. Without quality adjustment, all WTPs are both a) higher than their quality-adjusted counterparts, and b) necessarily higher than 0.

GDP variables are statistically significant, especially GDP growth per capita. "autonews" is a dummy indicating the variable is from the early years of the data when Truck/Van data are not collected. Therefore vehicles (cars) from this time period are deemed to have a higher intercept. Costs are increasing over time, and are higher for Asian and European vehicles.

The model implies a mean Lerner index ((p-mc)/p) of 32.5%, and a minimum of 3.9%.

[42] WTP is higher during times of high gas price because the marginal benefit of fuel efficiency, $\frac{\partial u}{\partial mpg_j}$, is increasing in the price of gas. $\frac{\partial^2 u}{\partial mpg \partial p_{gas}} > 0$.

Table 1: Estimated Correlation of Shocks and Characteristics

	mpg	p	ω
ξ	-0.28	0.45	0.25
ω	-0.26	0.61	

Allowing for correlation between errors and characteristics proves to have quantitative significance. The correlations are reported in Table 1, in the text. Previous literature had allowed for the correlation with price (second column), but not the correlation with characteristics (first column). The new correlations are not as large as the existing correlations, but this is unsurprising. All quality is priced to the consumer, but not all quality correlates with all other dimensions of quality. Cars with high demand and/or cost errors tend to be of the less efficient - more quality variety.

7 Robustness

There are a number of robustness checks of the model that I run. First, one might assume that firms set product characteristics before year $t-1$. In other words, there may be a longer lead time in production than 1 year. I have tested this alternative assumption by assuming product characteristics are set 3 and 5 years from the year of sale. This changes the estimation results surprisingly little. Though the change is large conceptually, mathematically all it means is that a single element of the instrument set (p_{gas}) is changed. Furthermore, the 3 and 5 year lagged gas price is (serially) correlated with the 1 year lag, so the instrument set is largely unchanged. This is a strength of the model - that it is fairly robust to when, exactly, firms must lock in their production decisions.

A second robustness check is whether the addition of the sub-segment level of the nest is too fine a partition for automobiles. When nesting groups are too finely partitioned, too much of the utility can end up being generated by the nest shares. This can lead to minimal substitution across nests. I re-estimate the model omitting sub-segments, but the results also change little. The nesting parameters are still relatively high. Substitution patterns in the counterfactuals are plausible with the 3-levels of nest, so I prefer this specification as it adds more richness to the demand specification.

A third robustness check is to incorporate consumers who are forward (or backward) looking with respect to the price of gas. This is done with interactions of fuel efficiency preference and lagged gas prices and/or movements. Upswings could, for example, make people more aware of gas prices and thus more sensitive to fuel efficiency. On the other hand, upswings could make people less sensitive to fuel efficiency if they expect prices to mean revert. It turns out that in this sample the former is somewhat true, though the effect is not large and steals some statistical precision from the *dpm* term itself. Therefore I prefer the specification with only the spot price of gas (as the numerator of *dpm*). Other work has documented that consumers' preferences for fuel efficiency follows the spot price of gasoline rather closely (Busse, Knittel, and Zettelmeyer, 2008).

I have also experimented with other controls (in X) on both demand and cost. I left out most which were insignificant. These included other functional forms of *mpg*, interaction terms with *mpg* and gas prices, and macroeconomic variables attempting to further control for the utility of the outside good (such as the Federal Funds Rate, unemployment, and election year dummies). The final specification leaves in the terms with explanatory power.

8 Counterfactuals: Summer 2008, and 35 *mpg*

8.1 Introduction

Having developed a model in which firms are free to choose the fuel efficiency of their new vehicles, I now use the model to learn about how the market would respond to various gasoline prices and taxes. Higher gasoline prices lead to a higher consumer sensitivity to fuel efficiency, and this in turn yields incentives for firms to raise fuel efficiency in their vehicles even though it requires making quality tradeoffs. As we consider ways to curtail energy expenditure and pollution, one option is gasoline taxes. Gasoline taxes raise the price of gasoline that consumers pay. Consumers would therefore want fuel efficient cars, and firms would have incentive to offer them.

There are other policy instruments that raise fuel efficiency - most notably the CAFE standards. The CAFE standards stipulate minimum sales-weighted fuel efficiency, firm by firm, and levy fines for non-compliance. The United States has had the CAFE standards in place since 1975, and there is evidence these have affected the fuel efficiency of new vehicle purchases. However, there have also been criticisms of the CAFE standards both in the academic literature and in the press.[43] Some criticisms apply equally to any policy that would increase fuel efficiency - for example that they encourage smaller cars which are less safe. Other criticisms of CAFE apply specifically vis-a-vis gasoline taxes. These criticisms have a common theme - that because the Standards do not target the specific behavior that causes the negative environmental externality (burning a gallon of gasoline), they have perverse effects and do not achieve their intended ends. One example is that both trucks and heavy vehicles receive more lenient Standards. This gives perverse incentive to produce more of these inefficient vehicles. Another example is that while the Standards may affect

[43] Gruenspecht (1982), Crandall and Graham (1989), and the series of citations herein by Kleit are some of the many academic papers which discuss the CAFE standards.

Table 2: International Gasoline Taxes

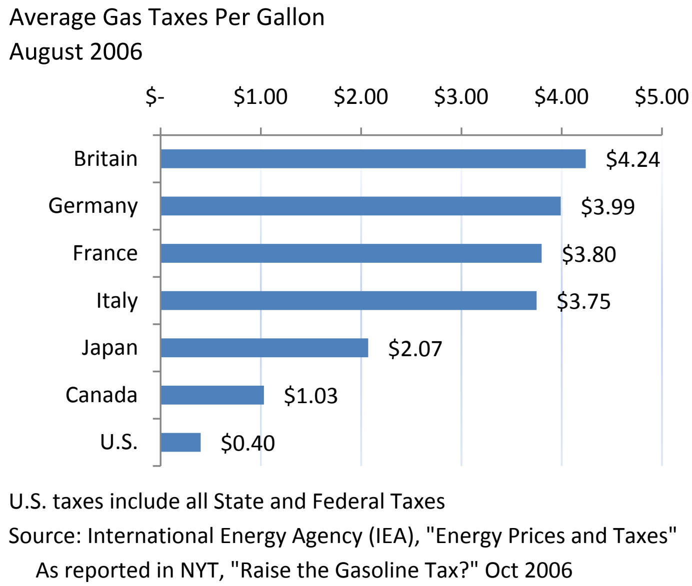

new vehicle purchases, they do not necessarily curb driving (the gas-burning) behavior. In fact, the effect on driving could be perverse: shifting people into more fuel efficient cars, while gas prices are still relatively low, could induce more driving. Taxing gasoline, on the other hand, would attach the policy instrument directly to the externality causing behavior of driving. This would solve both of the above problems. Some argue, therefore, that taxes are a more effective (less costly) way to achieve any given level of fuel conservation.

There is international precedent for using gasoline taxes, either in addition to or instead of regulation such as CAFE. Consumers in many industrialized nations face considerably higher gasoline taxes than the United States. (See Table 2 set in the text.) Many of these nations also have considerably higher fuel efficiency than the United States. The European Union sees over 40 sales-weighted *mpg*, compared to the United States which is under 25. This correlation does not establish causation, and there are many factors which affect nations' fuel efficiency.[44] However, the link between gasoline taxes and fuel efficiency in noteworthy, and gives context to our own consideration of domestic gasoline taxes. Higher gasoline taxes in European countries mean these consumers routinely pay $6-$8 per gallon of gasoline, and even more during global oil price spikes such as in the summer of 2008.

This section will consider two policy scenarios in this country. First, what would happen to domestic

[44]Some of these nations even have regulations similar to the CAFE standards (ICCT, 2007).

Table 3: Actual and Predicted Sales Declines, Summer 2008

Aggregation Level	Actual	Model Prediction
Cars	-3%	-9%
Utility Vehicles	-19%	-23%
Trucks	-16%	-13%
Total	**-13%**	**-14%**

fuel efficiency if the gas price were \$3.43? Second, what would the gas price need to be in order to achieve 35 *mpg* industry wide? The first scenario is of interest because this was the summer 2008 gasoline price when consumers because to react dramatically. During summer months the price was even higher than \$3.43, but \$3.43 was the annualized average price through August so I can compare YTD sales with model predictions. The second question of 35 *mpg* is of interest because recent policy proposals have stated a goal of raising the CAFE standards to 35 *mpg* by 2016. To preview the results of both scenarios, the \$3.43 gasoline (a 23% price increase over 2007) would increase average fuel efficiency offered to 26.9 *mpg* (a 31% increase) and sales-weighted fuel efficiency to 26.4 *mpg* (a 17% increase). The level of gas price (after tax) that would be require to boost market-wide efficiency to 35 *mpg* would be \$4.55 per gallon.

8.2 What Happens at \$3.43 Gasoline?

When the price of gasoline rose quickly to \$3.43, as it did in the summer of 2008, the first response is from consumers. Firms do not have a chance to adjust their product lines in the middle of the year, but consumers can adjust their purchasing behavior overnight. In the summer of 2008, the did so. All car sales were down, but even moreso in the SUV and truck segments as compared to cars. Table 3 in the text reports actual and model-predicted sales declines for August year-to-date 2008 automobile sales. The predictions at the finer aggregations levels are somewhat less precise, but the general shape/ordering of these predictions matches well, as does the overall level.

Table 3 is a condensed version of a table at the back of the text, Table 10. That Table tells the complete story of the counterfactual. The counterfactual can be thought of in three stages. First, 2007 is used as a baseline before the gas price increase. Then, 2008 is the year of the gas price increase - consumers have time to respond but firms do not. Then, in 2009, firms have the chance to update their characteristics. The 2009 column represents the new equilibrium of *mpg* offerings and purchases under \$3.43 gasoline.

A few notes are in order about this new equilibrium. First, because gas prices fell back down towards the end of 2008, it is not realistic to expect these "2009" fuel efficiency improvements to have materialized in 2009. These are simply what the model predicts would have happened if high gas prices had remained long enough to be reflected in new model years.

Second, despite calling this the "2009" equilibrium, I prefer to remain agnostic as to exactly how/how quickly would firms would arrive at this equilibrium, and whether or not it would be within a year. I have been able, in the model and in estimation, to remain rather agnostic about production lags. This is a strength in the model, though it means less specificity in the timing of the counterfactual. This tradeoff is reasonable, because I believe the level of the new equilibrium to be more interesting than the timing of the adjustment path to arrive there. Gas prices and product planning are admittedly much more continuous processes than the once-a-year snapshot of data that I observe, so trying to overextending the model's ability to predict timing would add little value. Firms may adjust more quickly or slowly based on when (in the model planning year) information about gas prices is revealed, whether a policy was pre-announced, and depending on production lead times. But the model indicates that at $3.43 gas, under the assumption that firms' historical record of ex-post regret was not foreseeable, there is an equilibrium of fuel efficiency offerings and purchases as described in column 2009.

The third and final note is that I do not model scenarios for gas price innovations from $3.43 gasoline. Gas prices are volatile, so it is true they will not simply move to $3.43 and stay there. (They did not in the summer of 2008). But the *mpg*-equilibrium is a function of the gas price, and the goal of the counterfactual is to understand how the levels are related.

Turning to the predictions in Table 10 for "2009," the new equilibrium, they indicate that at $3.43 gasoline, there is a new Nash equilibrium (offerings by firms and choices by consumers) which has average offerings of 26.9 *mpg* (31% higher than 2007) and average purchases of 26.4 *mpg* (17% higher than 2007). This is substantial improvement given that $3.43 gasoline was only 23% higher than 2007. Fuel efficiency changes are different for different segments in accordance with different willingness-to-pay for *dpm* (see Section 6 and Table 8). CUV's (both offered and purchased) have *mpg*'s in the low 40's, SUV's in the high 30's. These levels are high but plausible. They are due to the high consumer sensitivity to *dpm* in these segments. Trucks, Vans, Middle Cars and Large Cars show more modest *mpg* improvements. Specialty, Luxury, and Small Cars actually reduce fuel efficiency. These consumers have the lowest willing-ness-to-pay for *dpm*, there are more fuel efficient substitutes now available, and there is more room under the CAFE standards (the standards cease to bind for all manufacturers in this equilibrium). To put the overall efficiency in context, it would be historically high fuel efficiency for this country, but still considerably lower than the European Union's fuel efficiency of over 40 *mpg*. $3.43 is, after all, still a far lower price than what prevails in most of the EU.

8.3 Notes on Computing New Equilibria

Before turning to the other counterfactual scenario, this subsection contains comments on computing the new equilibrium described above. I use two different algorithms to converge to the new equilibrium. One is to iterate on best responses, the other is to slide there slowly. Conceptually, iteration on best responses (by full jumps) may be the "quicker" way to arrive at a destination fixed point. But with concern over multiple equilibria (fixed points), it may be prudent to take the slow and cautious approach. For my purposes it does not matter - both algorithms converge to the same equilibrium. This would appear to indicate that the nesting and ownership information are sufficient to avoid both "different-shaped" equilibria, and "same-shaped" equilibria but with interchanged identities. In other words, it suggests there may only be one fixed point. (Though it certainly does not prove it).

The first algorithm, iteration on best responses, is as follows. Conditional on old optimal choices, introducing a change in the economic environment (a gas price/tax) gives every vehicle a new optimal *mpg*. Calculate each of these new optima, vehicle by vehicle and holding all others fixed, and then move all vehicles at once to their new optima. Now that all the optima have changed presumably so have the best responses. Recalculate new best responses (again, one by one, holding others fixed), and again update all vehicles at once. Repeat this process until a fixed point attains.

There are two potential short-comings of this algorithm. First, and most importantly, it can involve large jumps of optima from iteration to iteration. This could potentially change the "shape" of the new equilibrium, and/or the identities of the agents in the case that they leapfrogged each other. If there are multiple equilibria, a desirable property of a new equilibrium would be that it find the "same" equilibria as was played before. A second potential short-coming is that finding a new optimal choice is computationally intensive compared to, for example, simply evaluating the profit derivative at the current choice. Finding the new optimal choice involves evaluating the profit function (or derivative) as many times as one wants to bisect the choice space (or update a Newton Method). The more desired precision for the new optima, the more evaluations. Evaluating the profit derivative at only the current choice is considerably faster.

The second algorithm I use attempts to address both of these potential short-comings of the iteration method. First, to establish the next "jump," I only evaluate the profit derivative once (at the current choice) rather than multiple times (to find the actual best response). The profit function is smooth and globally concave in the choice variable (second derivative always negative), so the first derivative is a good indicator of distance-to-root. (See Figure 8 for a picture of profit derivatives at two different levels of gas price for the Toyota Avalon.) Second, I only move in small steps towards new optima. The *relative* size of the move is in proportion is profit derivative, and the *absolute* size of moves is scaled to be a small distance. There is a tradeoff here - the smaller the jumps the slower the convergence but the less likely to jump to another equilibrium.

As I mentioned, both algorithms converge to the same equilibrium. So the multiple-equilibria consideration is moot in relation to selecting a convergence algorithm. This does not prove the equilibrium is unique, but it is a comforting finding nonetheless. In other contexts of industries the convergence algorithm may matter. The iteration by best responses did exhibit some large jumps in optima from iteration to iteration. The second, small-stepping equilibrium, can provide visual depiction of non-leapfrogging convergence. Figure 9 shows a sequence of scatter plots of *mpg* (Y-axis) and profit derivative (X-axis). The points slide smoothly and regularly towards the X-origin, which is the new Nash equilibrium.[45]

Another note on the equilibrium is that I do not allow firms to adjust prices of their vehicles in the new equilibrium. I did this for computational reasons. Finding one new equilibrium in *mpg* takes a few thousand iterations, and 20-40 minutes. Conceptually speaking, finding a new equilibrium would in prices would be similar to finding the 20-40 minute equilibrium *on top of each* of the iterations of *mpg*. By not allowing firms to adjust prices I am only approximating the Nash Equilibrium, but I have two reasons to believe this approximation may be a good one. First, in Pakes, Berry, and Levinsohn (1993) the authors found that new equilibrium prices were scarcely affected by a gas price change of similar size (21%) in the mid-1970's. So the gas price change itself is likely not to affect optimal prices overly much. The characteristics changes are also unlikely to affect prices much, because in my model these are repositioning along the same isocost curve. So optimal choices of *mpg* do not affect cost, and are therefore unlikely to affect optimal significantly. I would be more concerned if I were adjusting a characteristic that did affect cost. The approximation of the Nash Equilibrium is a reasonable price to pay for significant computational reduction.

8.4 Achieving 35 *mpg*

In 2007 Congress passed legislation to increase the CAFE standards to 35 *mpg* by 2020.[46] As of the time of writing, the Obama Administration has proposed accelerating the timetable to 2016. Whether by 2020 or 2016, my model allows me to put this CAFE improvement in terms of "gasoline price equivalent." In other words, I can answer the question: what gasoline price/tax combination would induce the market to achieve 35 *mpg* without the help of CAFE? The answer is $4.55. This is higher than the historical highs in the late 1970's and summer of 2008, though not by much. It is also still relatively low compared to European gasoline prices.[47]

[45] Here $\frac{\partial \pi}{\partial mpg} = 0$ for each model.

[46] This was the first substantial change to the standards in over 30 years since they were first instituted.

[47] Though, again, 35 *mpg* is also relatively low compared to the E.U.

A similar table of predictions to the \$3.43 scenario is provided in the back of the text in Table 11. I have left the column headings the same - 2007, 2008, and 2009 - though, again, these need not be thought of as literal years. They can be thought of respectively as the baseline, with consumer adjustment but not firms', and once firms adjust. Again, I remain agnostic as to exactly how/how quickly we might arrive at this equilibrium. The Table indicates the level equilibrium efficiency associated with \$4.55 gasoline. The 35 *mpg* statistic is in the lower right - it is the sales-weighted average fuel-efficiency of all vehicle purchases. The composition effects are qualitatively similar to the previous scenario of \$3.43, but the efficiencies are higher everywhere. Utility Vehicles are still pushed highest, and perhaps to levels that begin to strain plausibility. On the other hand, they are not far out of the range of existing *mpg*'s in these classes, and as green technology continues to develop they become even more attainable.

8.5 A Note on Gas Prices and Taxes

As I have mentioned, the \$3.43 or \$4.55 levels of gasoline could be achieved either by market forces (if the price of crude oil increased) or analogously through an increase in gasoline taxes. In the model, gasoline prices and taxes enter the same way. There is a wealth of information concerning domestic gasoline prices and taxes on the website of the Energy Information Administration (EIA). One rough way to summarize a large amount of information is that the U.S. pump price for a gallon of unleaded gasoline is roughly \$1.00 + 1/42 times the price of a barrel of crude oil. The \$1.00 is somewhat stable - \$0.40 for taxes (\$0.18 Federal, \$0.22 average state), \$0.30 refining, \$0.30 distribution and marketing. A barrel of crude (the more volatile component of price) contains 42 gallons, thus the 1/42 factor.[48] So, when a barrel of crude is \$42, we pay in the neighborhood of \$2 at the pump. When it is \$126, we pay in the neighborhood of \$4 at the pump. There is large regional and temporal variation in this relationship, but this heuristic fixes magnitudes. Raising the gasoline tax could add either to the volatile portion of the price (if it were a percentage tax) or to the more stable portion (if it were a fixed tax per gallon).

Gas prices will always exhibit some volatility because of their link to crude oil. Per gallon gasoline taxes, though, are part of the relatively stable base of prices exhibiting less volatility. Increases in per gallon gas taxes, therefore, seem more likely to induce fuel efficiency gains than spikes in the crude oil price because they are more likely to persist.

Recall also that firms in the counterfactual scenarios are constrained from introducing new automobiles. It may be, therefore, that the fuel efficiency gains I report are lower bounds, especially over greater lengths of time. The more time elapsed from a gas tax increase, presumably the more

[48]Only half of a barrel of crude will turn into gasoline for motor vehicles, but no crude is wasted. The rest turns into jet fuel, kerosene, plastics, etc.

resources would be shifted towards developing and introducing relatively fuel efficient vehicles.

9 Conclusion

The auto industry is large, represents a significant portion of the nation's energy needs, and contributes significantly to carbon emissions. Understanding the impacts that various policies and markets (particularly the gasoline market) have on automotive fuel efficiency is crucial to informing energy policy.

I develop a model of the automobile industry to analyze these issues. In the model, firms provide more or less fuel efficiency depending on the stochastically changing gas price. The model contributes to both the methodological literature and the policy discussion. There are three modeling contributions. First, unlike previous work I allow firms to choose product characteristics of their new vehicles. This allows me to analyze changes in vehicle characteristics themselves. Second, I relax identifying assumptions commonly used in empirical demand estimation. These identifying assumptions are convenient but implausible in the automobile market, and I relax them by forming moments based on the timing of product selection. Third, I provide parameter estimates that imply that consumers care about fuel efficiency. Previous work has had some difficulty doing so, due to a negative correlation between the economic and quality effects of fuel efficiency. Controlling for both of these effects simultaneously restores the expected preference pattern.

The model makes contributions to policy discussion, as well. The model can be used to predict the level of fuel efficiency that a given gas price/tax combination may achieve. Gasoline taxes are an oft-proposed alternative to CAFE standards, though gas taxes are still relatively low in this country. I analyze two hypothetical gas price scenarios - the high prices from the summer of 2008, and achieving 35 *mpg* with taxes. In the summer of 2008 gasoline prices were \$3.43. If gas prices were to have stayed at this level, the model indicates that sales-weighted fuel efficiency would increase 17% to 26.4 *mpg*. The other question I address is: what level of gas price would we need in order to achieve 35 *mpg* fleet-wide? 35 *mpg* is the goal of recent CAFE proposals that will take effect between 2016 and 2020. The model indicates that a gas price of \$4.55 would take us to 35 *mpg* absent CAFE. This price is high for U.S. standards, but modest compared to other developing nations where gas taxes, gas prices, and fuel efficiency are all higher than in the U.S.

The estimation results, incidentally, indicate that the CAFE standards have been quite costly to American manufacturers. The model estimates a shadow cost imposed on domestic manufacturers, since 1977, of \$347 per car. That is much larger than the \$55 per-*mpg* fine. If such a shadow cost is true it indicates that the domestic firms are strongly averse to violating the standards, and/or that these firms' fuel efficiencies would have been considerably lower without the CAFE statutes.

The model is best thought of as a medium-run analysis. While I do control for improvements in the technology frontier over time, I do not allow firms to introduce new vehicle models. On a longer time horizon, firms would have time to not only adjust characteristics on existing vehicles but also develop new models. This extra degree of freedom would presumably magnify fuel efficiency movements in either direction.

The main message of the paper is that historical data strongly indicate consumers care about fuel efficiency. They do so all the more when gas prices are high. This gives firms an incentive to offer efficient vehicles when consumers pay a lot at the pump. One way to increase price at the pump is through gas taxes. Taxes have the benefit of directly pricing the environmentally costly behavior - burning gasoline. My model suggests that taxes would also have the effect of raising the fuel efficiency of our domestic fleet considerably. Raising fuel efficiency has been a long-standing policy objective and appears likely to remain one into the future.

References

Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242–242.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–841.

——— (1999): "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, 89(3), 400–30.

Busse, M., C. Knittel, and F. Zettelmeyer (2008): "Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing," *Working Paper*.

Cardell, N. (1997): "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," *Econometric Theory*, 13, 185–213.

Crandall, R., and J. Graham (1989): "Effect of Fuel Economy Standards on Automobile Safety, The," *Journal of Law and Economics*, 32, 97.

Crawford, G., and M. Shum (2007): "Monopoly Quality Degradation and Regulation in Cable Television," *The Journal of Law and Economics*, 50(1), 181–219.

EPA (2008): "Environmental Protection Agency: Inventory of US Greenhouse Gas Emissions and Sinks, 1990 - 2006," .

Federal Reserve Board, C. (2003): "Predicting Oil Prices," .

Goldberg, P. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the US Automobile Industry," *Econometrica*, 63, 891–891.

——— (1998): "The Effects of the Corporate Average Fuel Efficiency Standards in the US," *Journal of Industrial Economics*, 46(1), 1–33.

Goldberg, P., and F. Verboven (2001): "The Evolution of Price Dispersion in the European Car Market," *The Review of Economic Studies*, 68(4), 811–848.

Gruenspecht, H. (1982): "Differentiated Regulation: The Case of Auto Emissions Standards," *American Economic Review*, 72(2), 328–331.

Hansen, L., and K. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50(5), 1269–1286.

Hotelling, H. (1929): "Stability in Competition," *Economic Journal*, 39(153), 41–57.

ICCT (2007): "Passenger Vehicle Greenhouse Gas and Fuel Economy Standards: A Global Update," *The International Council on Clean Transportation*.

Kleit, A. (1990): "The effect of annual changes in automobile fuel economy standards," *Journal of Regulatory Economics*, 2(2), 151–172.

——— (2004): "Impacts of Long-Range Increases in the Fuel Economy (CAFE) Standard," *Economic Inquiry*, 42(2), 279–294.

Kleit, A. N. (2002): "CAFE Changes, by the Numbers," *Regulation*, 25(3).

Larrick, P. L., and J. B. Soll (2008): "The MPG Illusion," *Science*, 320(20 June), 1593–1594.

Lustig, J. (2008): "The Welfare Effects of Adverse Selection in Privatized Medicare," *Working Paper*.

Mazzeo, M. (2002): "Product Choice and Oligopoly Market Structure," *RAND Journal of Economics*, 33(2), 221–242.

NEED (2008): "National Energy Education Development Project: Intermediate Energy InfoBook," .

Pakes, A., S. Berry, and J. Levinsohn (1993): "Applications and Limitations of Some Recent Advances in Empirical Industrial Organization: Price Indexes and the Analysis of Environmental Change," *American Economic Review*, 83(2), 240–246.

Seim, K. (2006): "An empirical model of firm entry with endogenous product-type choices," *RAND Journal of Economics*, 37(3), 619–640.

Sweeting, A. (2007): "Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry," *Forthcoming Econometrica*.

Thomadsen, R. (2005): "The Effect of Ownership Structure in Geographically Different industries," *RAND Journal of Economics*, 36(4), 908–929.

Wu, T., and A. McCallum (2005): "Do Oil Futures Prices Help Predict Future Oil Prices?," *Federal Reserve Board San Francisco Economic Letter*, (38).

Appendix A: Nested Logit Shares

This appendix details the nested logit errors, shares, and nesting parameters σ.

Rewrite the utility function in Equation (2) as:

$$u_{ijt} = \delta_j + \tilde{\epsilon}_{ijt} \tag{19}$$

The nested logit error term $\tilde{\epsilon}_{ijt}$ is:

$$\tilde{\epsilon}_{ijt} = \epsilon_{i,v} + (1-\sigma_v)\epsilon_{i,s} + (1-\sigma_s)\epsilon_{i,ss} + (1-\sigma_{ss})\epsilon_{ij} \tag{20}$$

Each individual i has a shock not only on the auto model (ϵ_{ij}), but also a shock common to all vehicles within the same sub-segment ($\epsilon_{i,ss}$), segment ($\epsilon_{i,s}$), and vehicle type ($\epsilon_{i,v}$). These error terms have a distribution such that the cumulative error term within any group follows the type-I extreme value distribution. Therefore, market shares for product j are:

$$s_j = \frac{e^{\frac{\delta_j}{1-\sigma_{ss}}}}{D_{ss}} * \frac{D_{ss}^{\frac{1-\sigma_{ss}}{1-\sigma_s}}}{D_s} * \frac{D_s^{\frac{1-\sigma_s}{1-\sigma_t}}}{D_t} * \frac{D_t^{1-\sigma_t}}{D_a} \tag{21}$$

where D's are inclusive values for nests:

$$D_{ss} = \sum_{j \in ss} e^{\frac{\delta_j}{1-\sigma_{ss}}}$$

$$D_s = \sum_{ss \in s} D_{ss}^{\frac{1-\sigma_{ss}}{1-\sigma_s}}$$

$$D_t = \sum_{s \in t} D_s^{\frac{1-\sigma_s}{1-\sigma_t}}$$

$$D_a = \sum_{t} D_t^{1-\sigma_t} + 1$$

Taking logs and rearranging yields:

$$ln\frac{s_j}{s_o} = \sigma_{ss} ln\frac{s_j}{s_{ss}} + \sigma_s ln\frac{s_{ss}}{s_s} + \sigma_v ln\frac{s_s}{s_v} + \delta_j \tag{22}$$

δ_j customarily includes a demand error, ξ_j, that is used to form estimation moments.

Appendix B: Derivation of First-Order Condition

This appendix details the derivation of the first-order condition of profit with respect to *mpg* choice. This includes the handling of the shadow cost, λ. The derivation of the first-order condition for price p is similar but reduces to a simpler form. *mpg* is abbreviated as m for parsimony:

$$\Pi_f = \sum_{j \in \Im_f} M_t s_j(\mathbf{mpg}, \mathbf{p}; \theta) \left[p_j - mc_j(mpg_j; \theta) - \lambda_{ft}\{bind_{ft}\}(CAFE_t - \overline{mpg_{ft}}) \right]$$

$$\begin{aligned}
0 &= \frac{\partial \Pi_{ft}}{\partial mpg_{jt}} \\
&= \sum_{r \in \Im_f} \left[(p_r - mc_r)\frac{\partial s_r}{\partial m_j} + s_j\left(\frac{\partial p_r}{\partial m_j} - \frac{\partial mc_r}{\partial m_j} + \lambda\{bind\}\frac{\partial \overline{mpg_f}}{\partial m_j}\right) \right] \\
&= \sum_{r \in \Im_f} \left[(p_r - mc_r)\frac{\partial s_r}{\partial m_j} + s_j \lambda\{bind\}\frac{\partial \overline{mpg_f}}{\partial m_j} \right] - s_j \frac{\partial mc_j}{\partial m_j} \\
&= \sum_{r \in \Im_f} \left[(p_r - mc_r)\frac{\partial s_r}{\partial m_j} + s_j \lambda\{bind\} \left[M \sum_{r \in \Im_f} \left[\frac{\partial s_r}{\partial m_j}\left(\frac{\overline{mpg_f}}{q_f} - \frac{\overline{mpg_f}^2}{q_f m_r}\right) \right] + \frac{q_r \overline{mpg_f}^2}{q_f m_j^2} \right] \right] \\
&\qquad - s_j \frac{\partial mc_j}{\partial m_j} \\
&= \sum_{r \in \Im_f} \frac{\partial s_r}{\partial m_j} + \lambda\{bind\} q_f T_1 + \frac{\overline{mpg_f}^2 q_r}{M m_j^2}\lambda\{bind\} - s_j \frac{\partial mc_j}{\partial m_j}
\end{aligned} \tag{23}$$

Figure 2: Nesting Structure

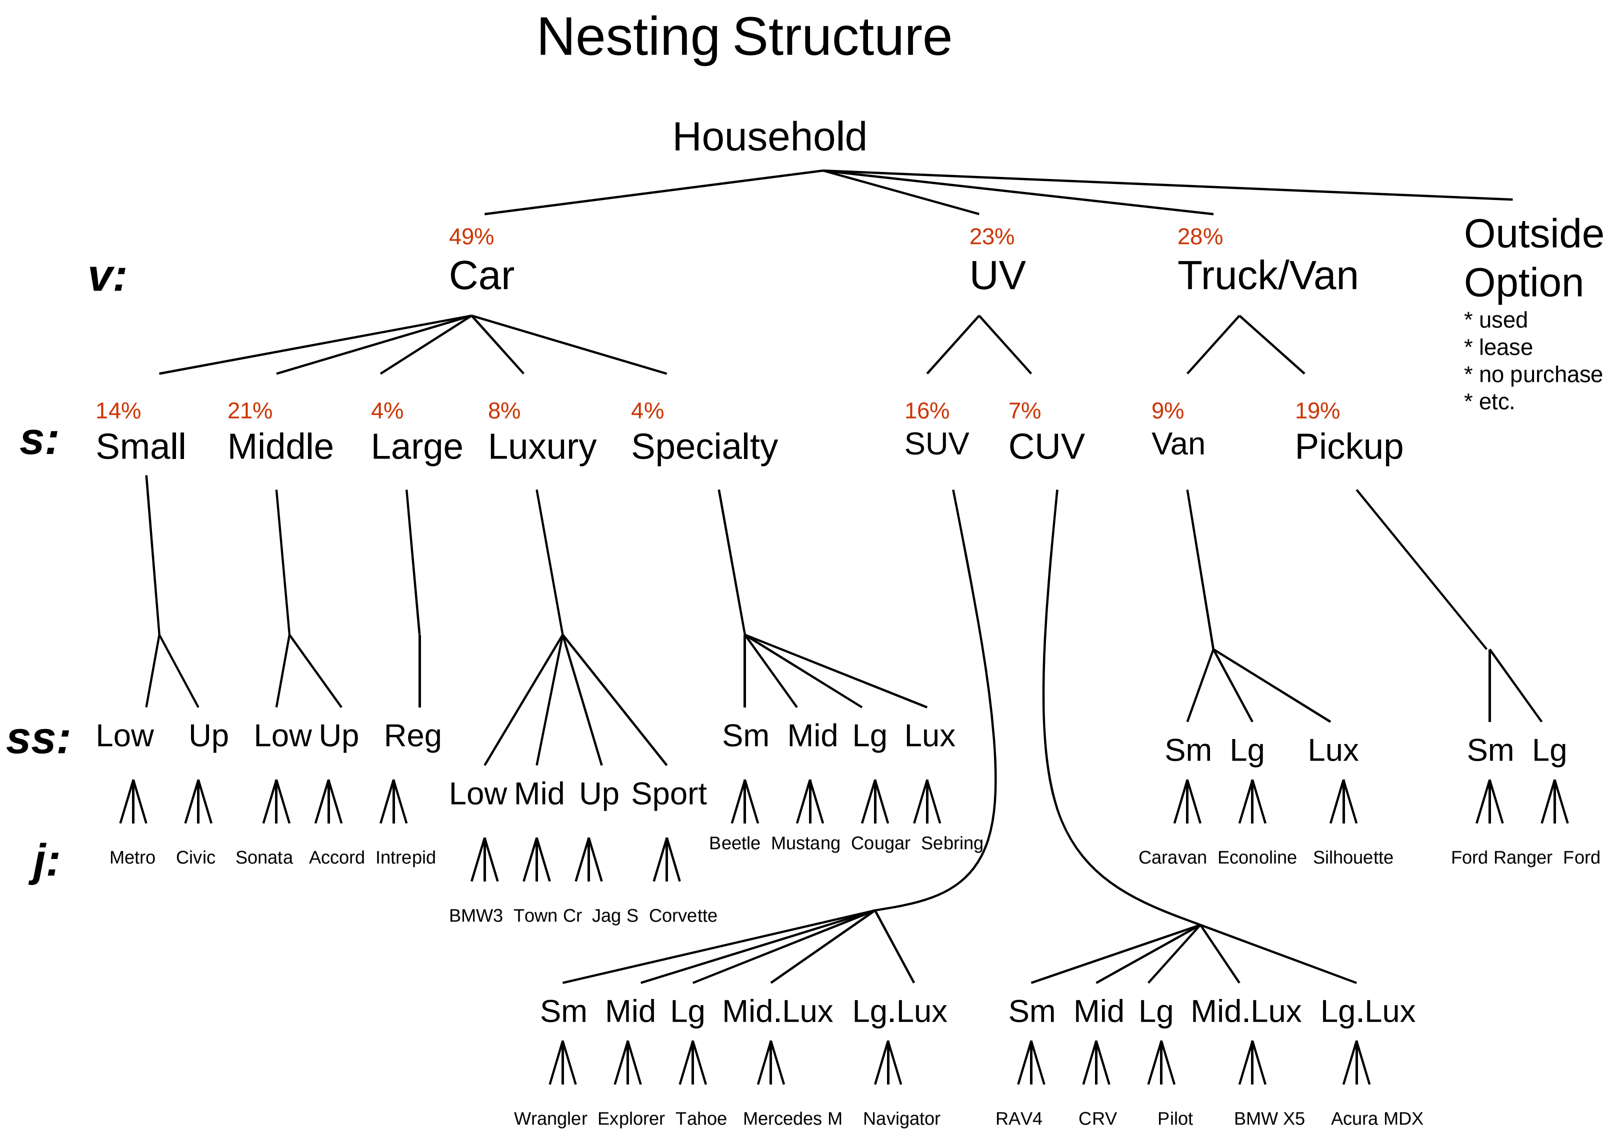

Figure 3: Tradeoff Between MPG, HP, and Weight

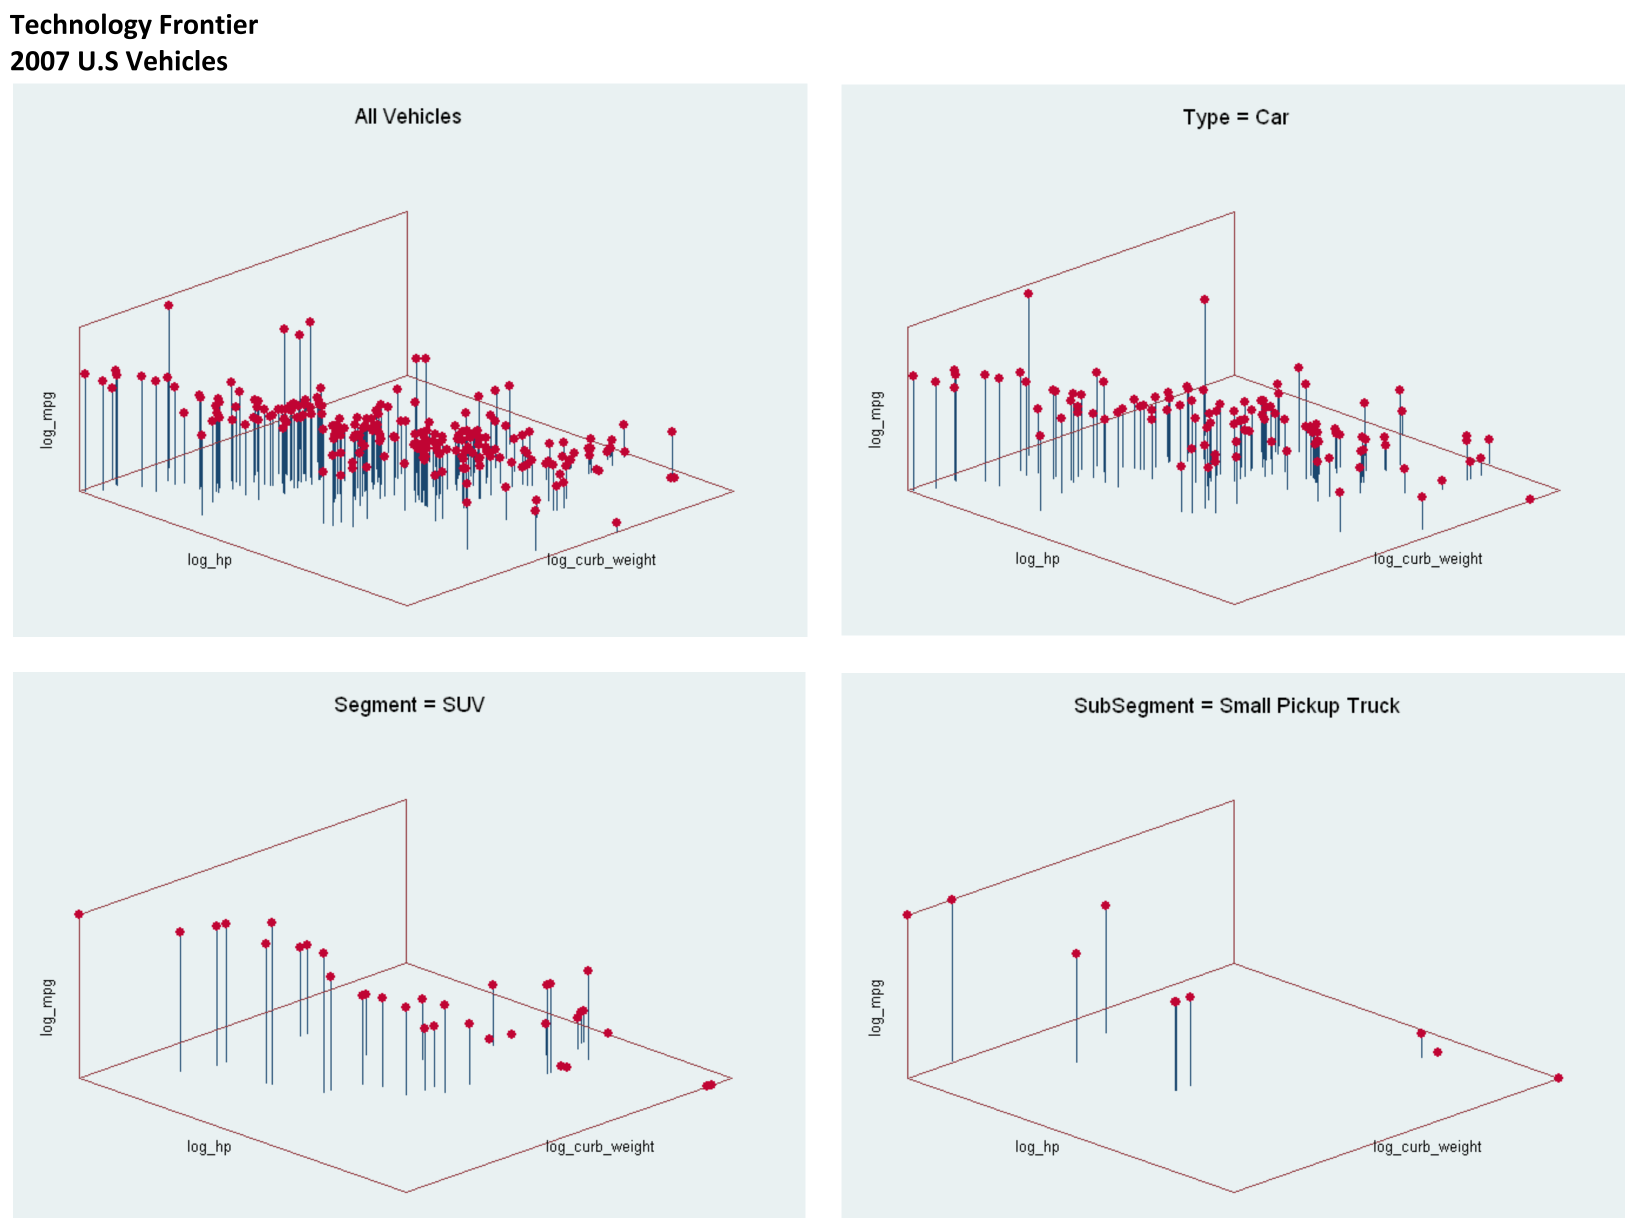

Figure 4: Timing of the Game

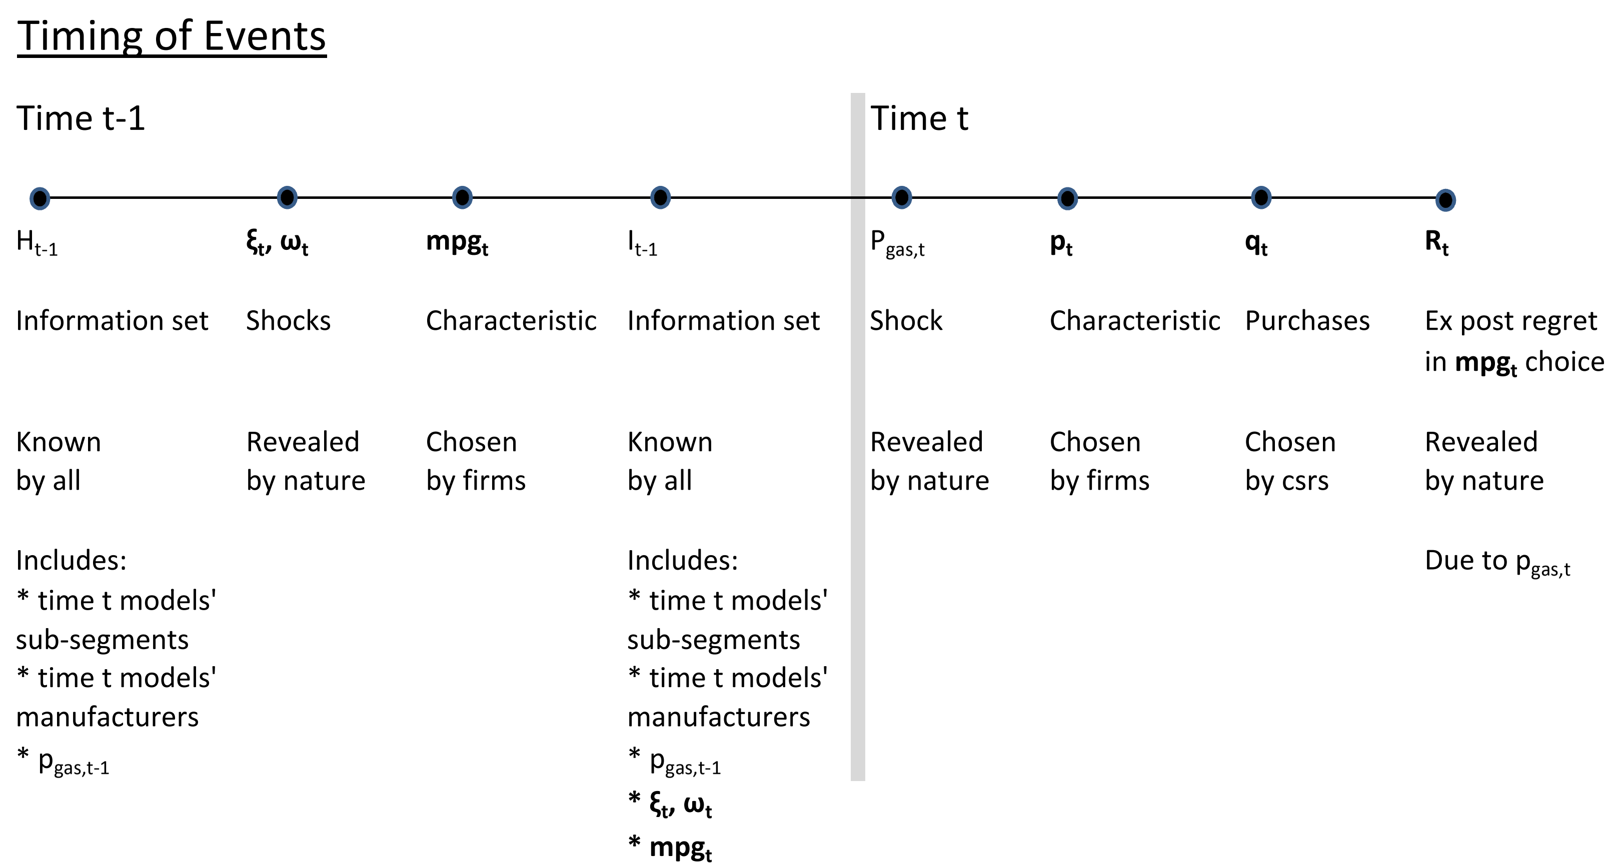

Figure 5: Industry Change over Time

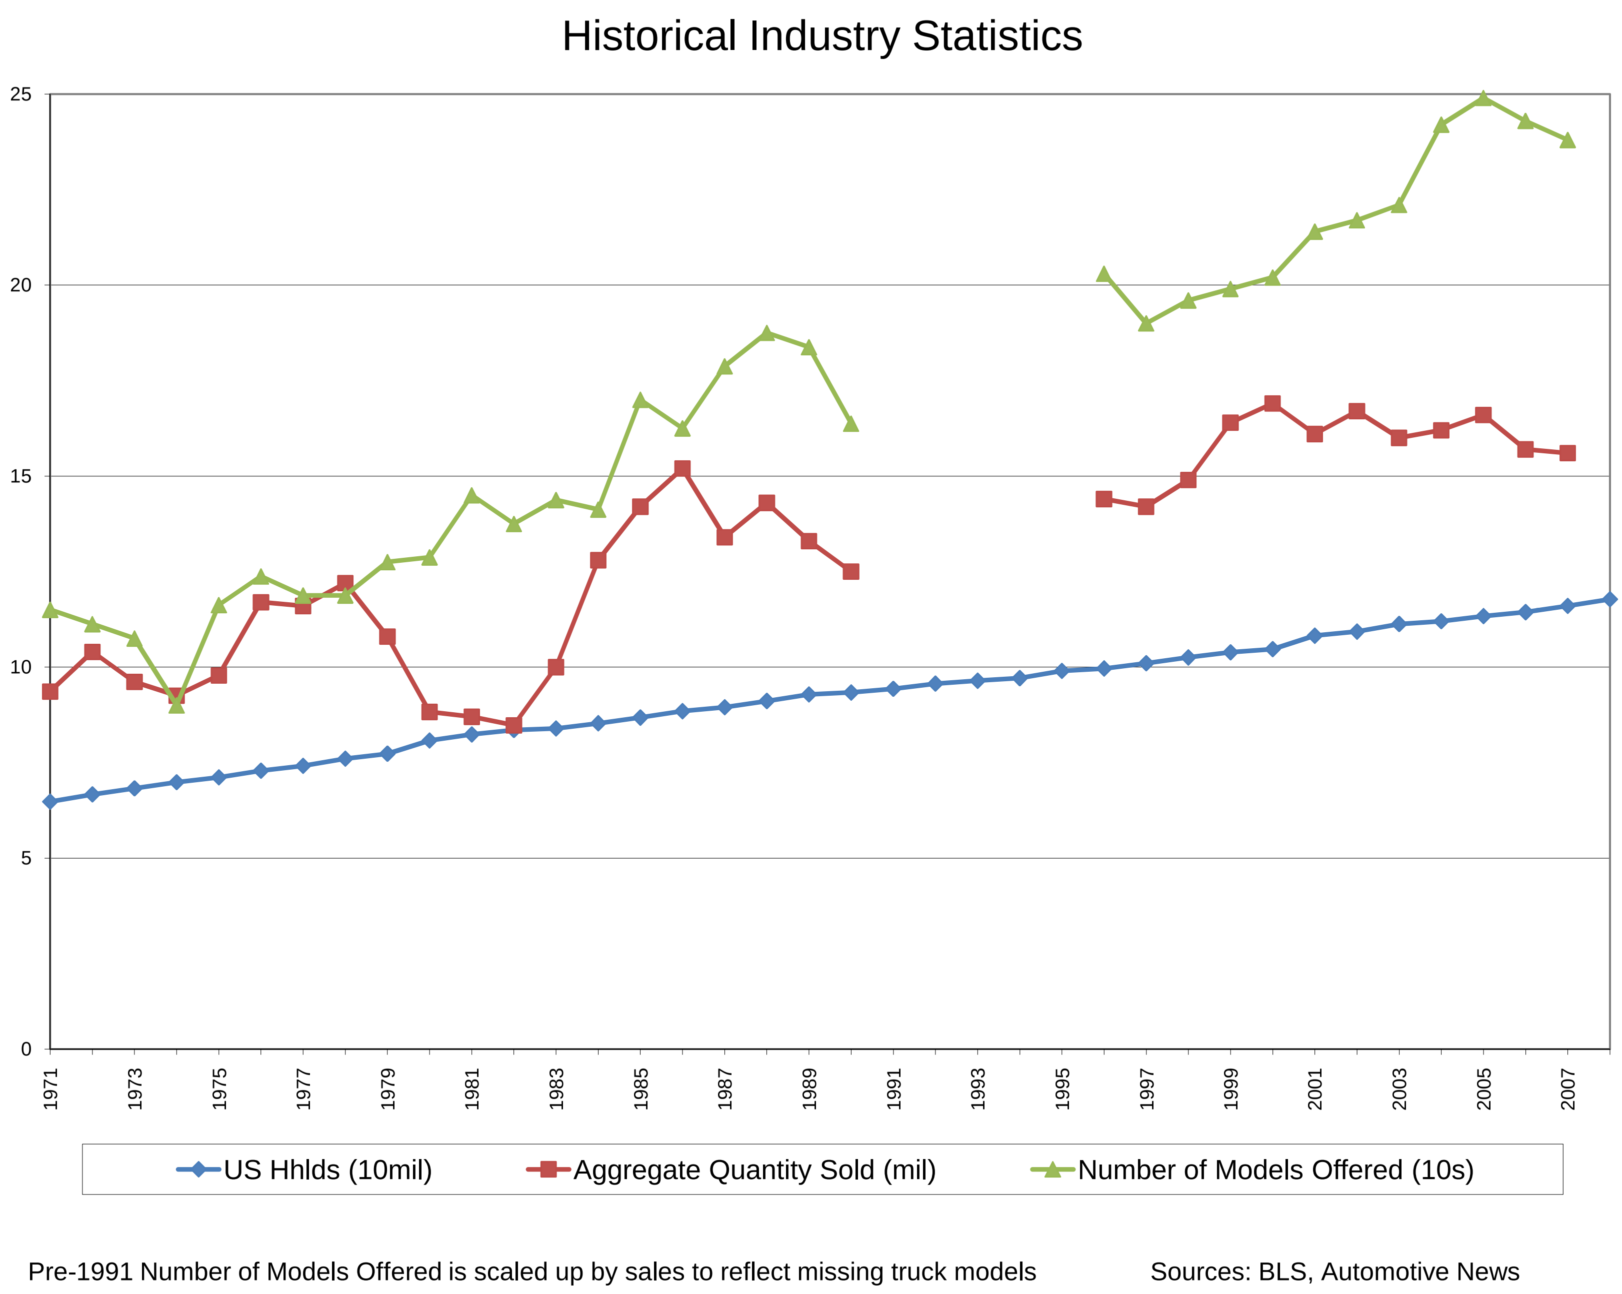

Figure 6: History of Fuel Efficiency and Regulation

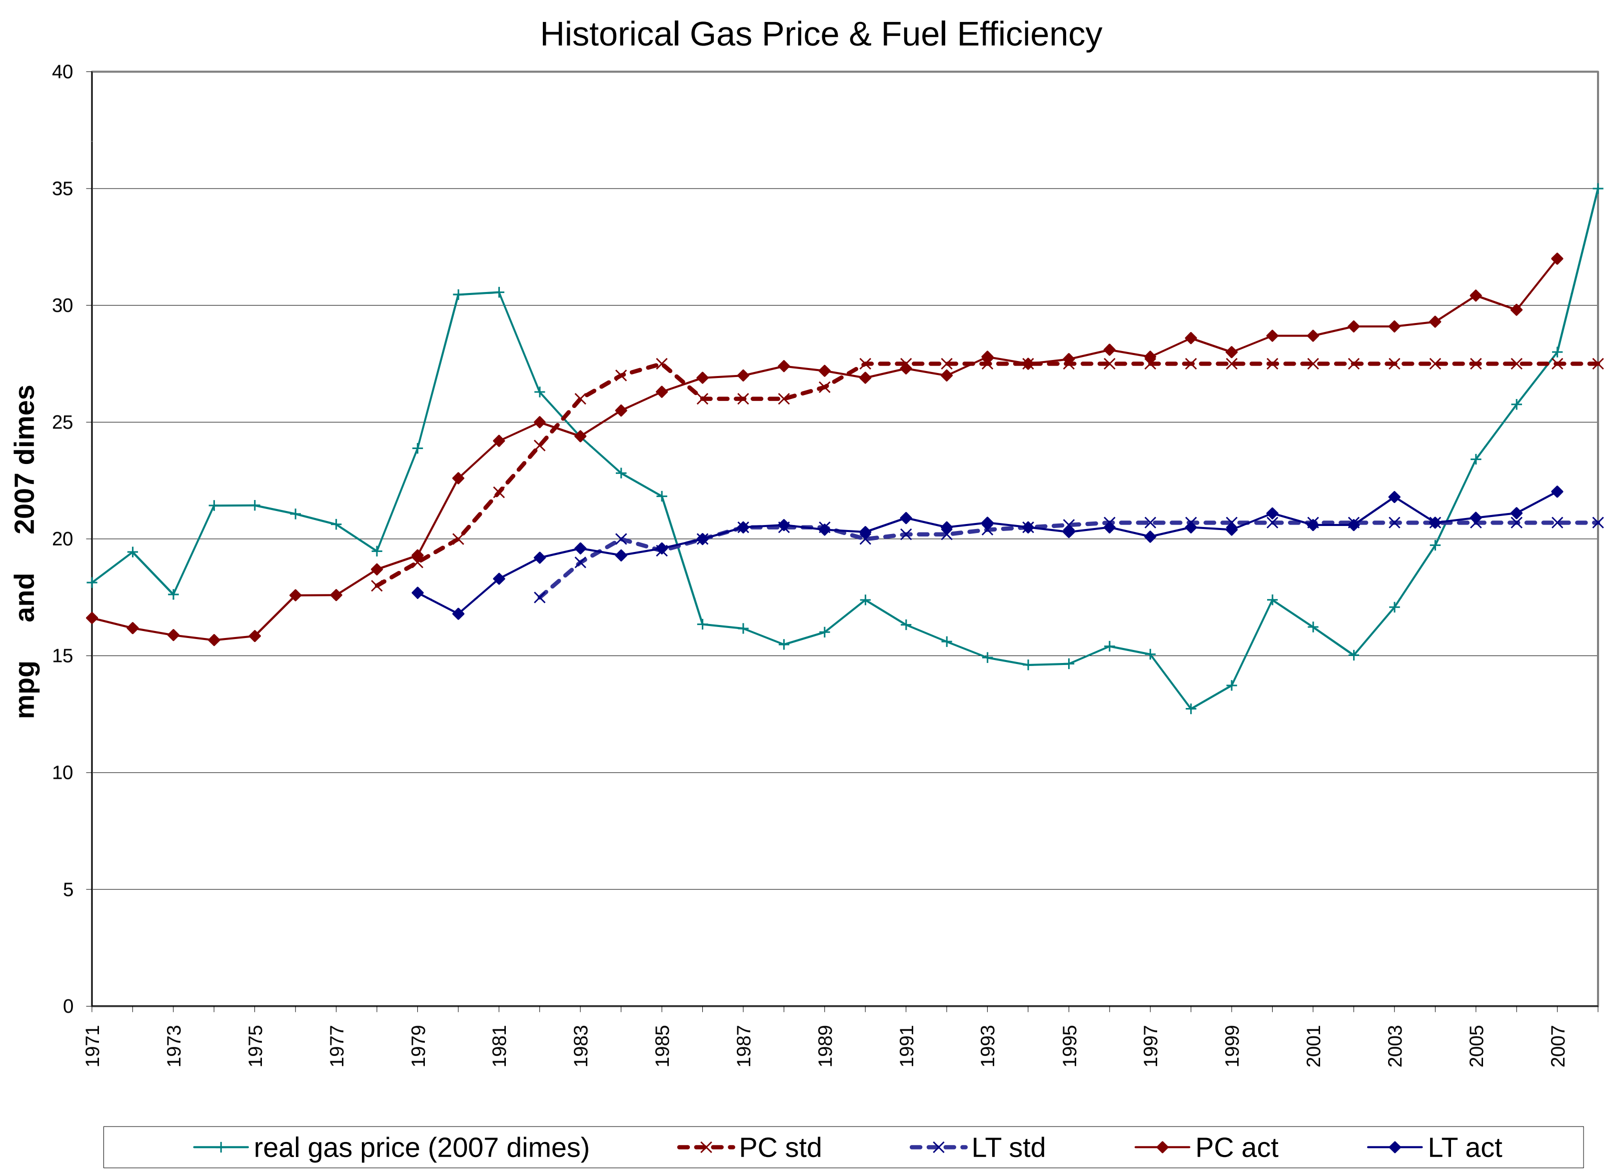

PC = Passenger Cars, LT =Light Trucks

Sources: Wards, Automotive News, EIA, NHSTA

Figure 7: Yearly Characteristics Changes

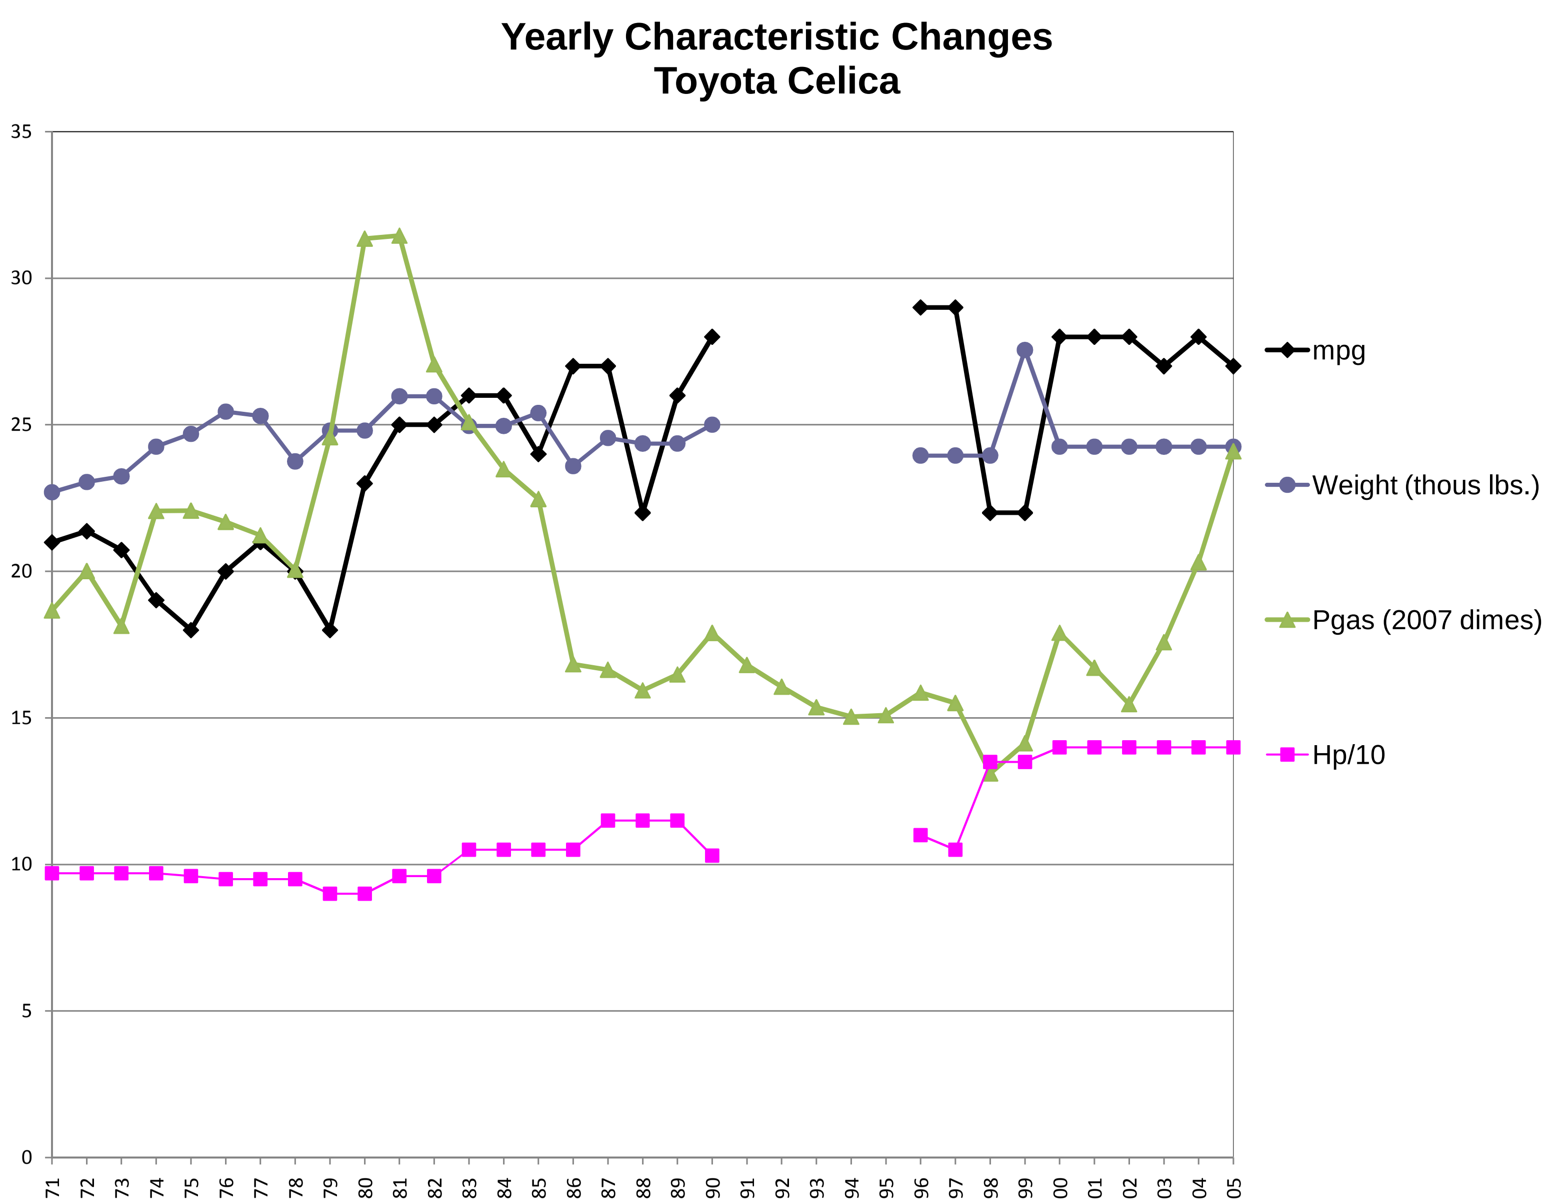

Figure 8: Changing Gas Price, Changing Optima

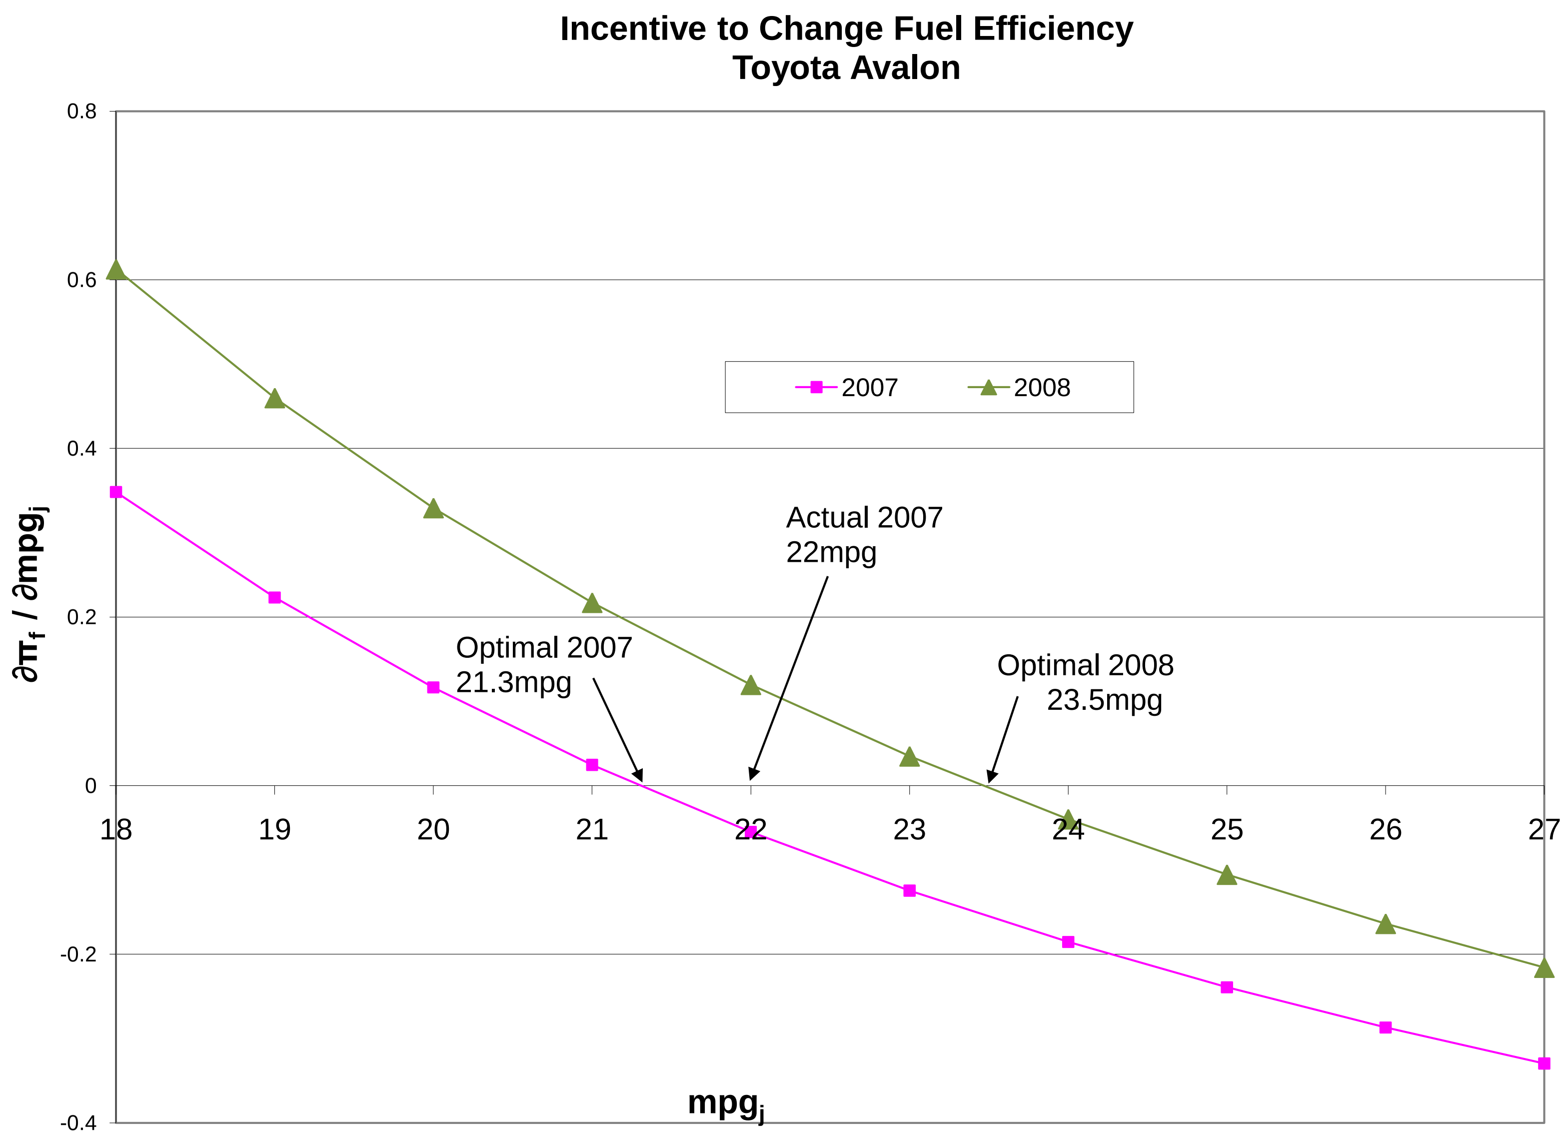

Figure 9: Convergence to New Equilibrium: Y-axis = mpg, X-axis = $\frac{\partial \pi}{\partial mpg}$

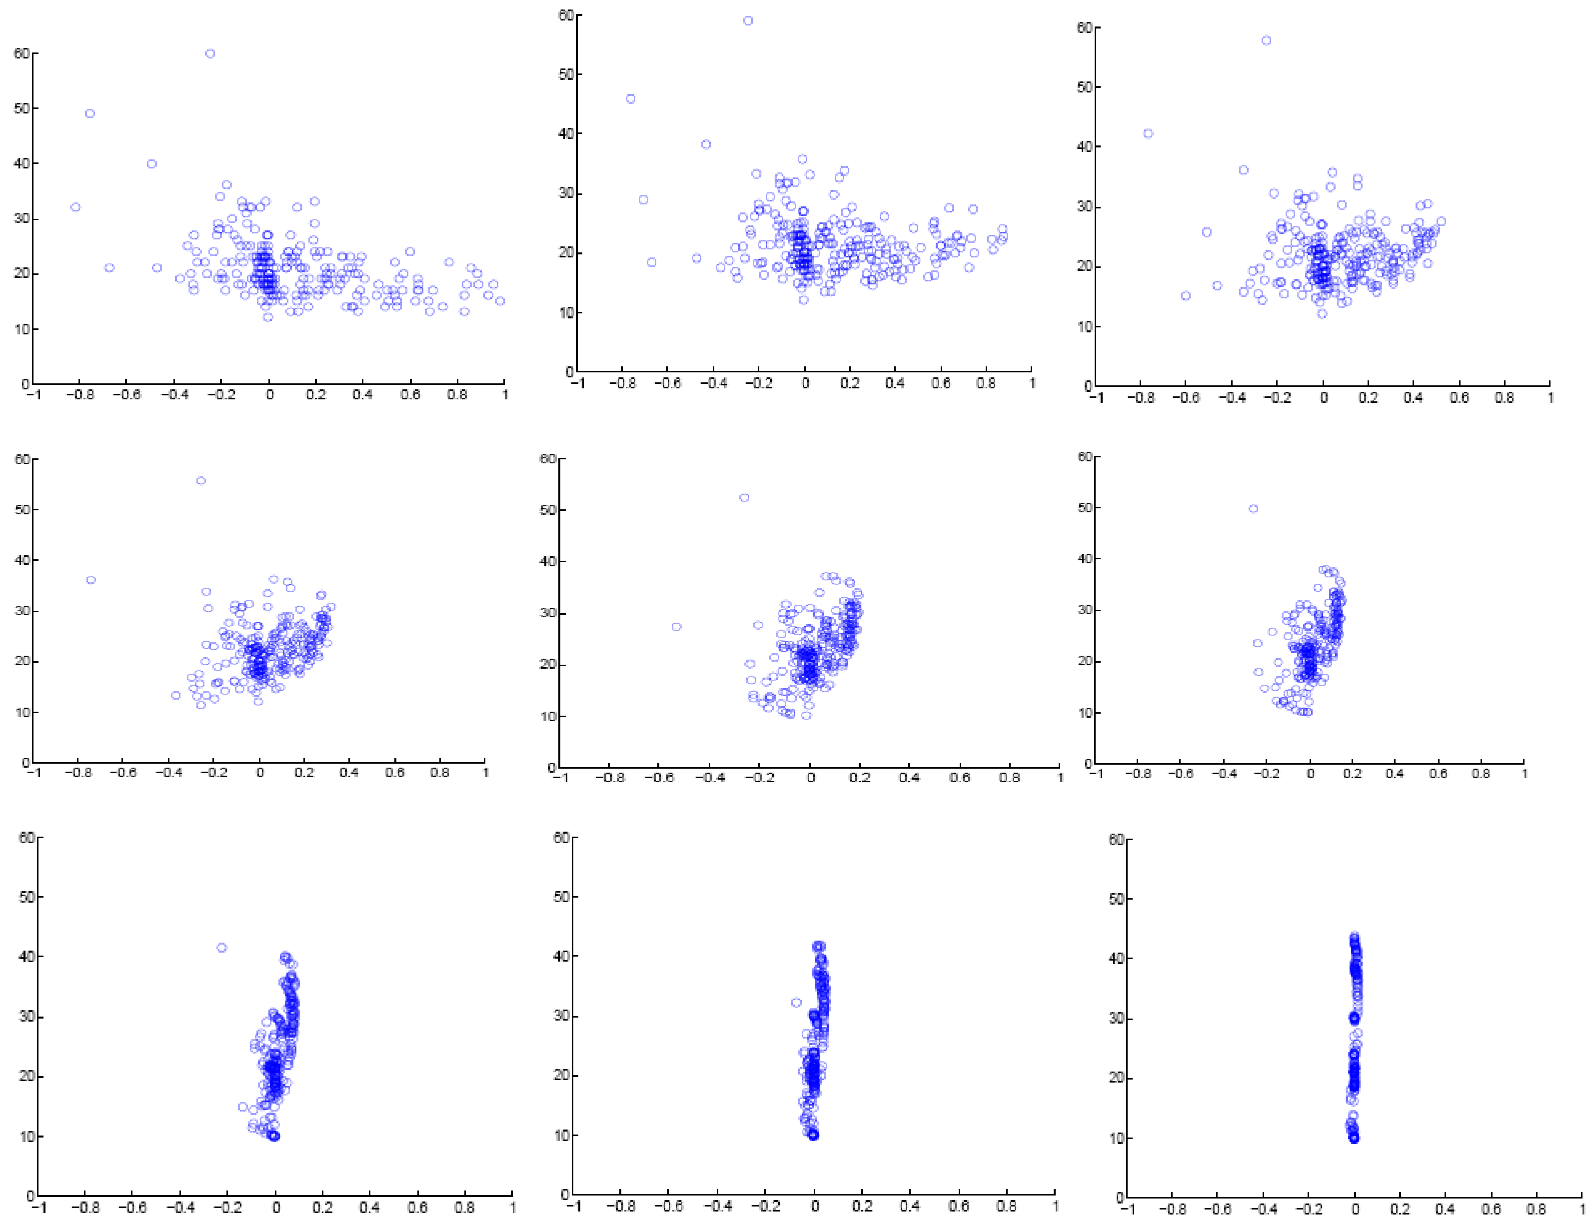

Table 4: Tradeoff Between MPG, HP, and Weight

OLS Rregression

Dep Var = ln(mpg)	Coef	SE	t
ln(hp)	-0.279	0.008	-37
ln(curb_weight)	-0.596	0.012	-52
year	0.013	0.000	66
const	-16.682	0.397	-42
Nobs	4820		
R-sq	0.78		

R-sq between .4 and .85 for
same regression within subsegment

Table 5: Popular Cars by Automotive Sub-Segment

Automotive Types, Segments and Sub-Segments

Type	Segment / Sub-Segment	% of Model-Yrs	% of Sales	Popular Models				
Cars	**Small**	**10.9%**	**13.5%**					
	Lower Small	3.6%	1.1%	GMC Metro	Hyundai Accent	Kia Sephia		
	Upper Small	7.3%	12.3%	Honda Civic	GMC Cavalier	Ford Escort	Toyota Corolla	
	Middle	**11.5%**	**20.5%**					
	Lower Middle	4.7%	5.8%	Honda Civic	GMC Grand Am	GMC Malibu	Hyundai Sonata	Nissan Altima
	Upper Middle	6.8%	14.7%	Toyota Camry	Honda Accord	Ford Taurus	GMC Impala	Nissan Altima
	Large	**3.1%**	**3.8%**					
	Large Regular	3.1%	3.8%	GMC Lesabre	Chrysler Intrepid	Chysler 300 Series	Ford Grand Marquis	Ford Crown Victoria
	Luxury	**23.6%**	**7.6%**					
	Lower Luxury	7.6%	3.5%	BMW 3	Toyota Avalon	Volvo 70	Acura TL	Lexus ES330
	Middle Luxury	6.4%	2.6%	GMC Fleetwood Deville	Ford Town Car	GMC Park Avenue	BMW 5	Mercedes C
	Upper Luxury	3.8%	0.8%	Mercedes E	GMC Seville	BMW 5	Lexus LS430	Jaguar S
	Luxury Sport	5.7%	0.7%	Corvette	BMW Z3	Chrysler Crossfire	Porsche 911	
	Specialty	**10.8%**	**4.1%**					
	Small Specialty	3.3%	0.8%	Beetle II	Toyota Scion	Toyota Celica	BMW Mini Cooper	Mazda MX3
	Middle Specialty	4.5%	2.7%	Ford Mustang	Mitsubishi Eclipse	GMC Camaro	GMC Monte Carlo	Toyota Celica
	Large Specialty	0.3%	0.1%	Ford Thunderbird	Ford Cougar	these are the only two cars in this class & they only exist 96-98		
	Luxury Specialty	2.6%	0.5%	Chrysler Sebring CNVT	Acura CL	Mercedes CLK	GMC Eldorado	BMW 6
Utility Vehicles	**CUV**	**10.6%**	**6.5%**					
	Small CUV	1.4%	1.2%	Chrysler PT Cruiser	GMC HHR	Toyota RAV4	Honda Element	
	Middle CUV	4.3%	3.5%	Ford Escape	Honda CRV	GMC Equinox	Toyota Highlander	
	Large CUV	1.0%	0.6%	Honda Pilot	Chrysler Pacifica	Ford Freestyle		
	Middle Luxury CUV	3.0%	1.0%	Lexus RX330	BMW X5	Volvo XC90	Mercedes M	
	Large Luxury CUV	0.9%	0.3%	Acura MDX	GMC SRX	Mercedes R Class		
	SUV	**13.0%**	**16.0%**					
	Small SUV	1.4%	0.9%	Chrysler Wrangler	Kia Sportage	GMC Tracker	Suzuki Vitara	
	Middle SUV	4.9%	9.3%	Ford Explorer	Jeep Grand Cherokee	GMC Trailblazer	Chrysler Liberty	
	Large SUV	2.3%	4.5%	Ford Expedition	GMC Tahoe	Chrysler Durango		
	Middle Luxury SUV	2.6%	0.7%	Mercedes M	Lexus GX470	GMC Bravada	Infiniti QX4	
	Large Luxury SUV	1.9%	0.6%	Ford Navigator	GMC Escalade	GMC Hummer 2		
Trucks/Vans	**Van**	**9.7%**	**9.4%**					
	Small Van	6.1%	6.6%	Chrysler Caravan	Ford Windstar Pass.	Chrysler Town & Country	Toyota Sienna	
	Large Van	2.4%	2.2%	Ford Econoline	GMC Express Cargo			
	Luxury Van	1.2%	0.6%	Chrysler Town & Country	GMC Silhouette	Mazda MPV Pass. Van	Toyota Previa	
	Pickup	**5.4%**	**18.5%**					
	Small Pickup	3.3%	5.3%	Ford Ranger	GMC S10 Pickup	Toyota Tacoma		
	Large Pickup	2.1%	13.2%	Ford F Series	GMC Silverado	GMC Chevy CK Pickup	Chrysler Ram Pickup	
	Other	**1.4%**	**0.1%**					
	Comm. Chassis	1.4%	0.1%	Chevy P Model	Isuzu NPR			

Table 6: Summary Statistics

Summary Statistics

Variable	Mean	Std. Dev.	Min	Max
Year	1992.5	11.1	1971	2007
Mpg	20.7	5.8	9.13	61
Dpm	$0.10	$0.04	$0.03	$0.24
Price	$30,059	$21,195	$7,038	$115,000
q	75,566	95,043	123	890,790
Price gas	$2.02	$0.48	$1.31	$3.15
Price gas change	$0.05	$0.24	-$0.56	$0.68
GDP Growth	3.1%	1.7%	-1.9%	7.2%
N Models in Year	150.6	55.9	72	249

N Model-Years	4,820
N Years	32
N Types	3
N Segments	9
N Sub-segments	28

Note: Prices are Real 2007 Dollars

Table 7: Estimation Results

GMM Estimation Results

Demand parameters:

Variable		Coeff	S.E.	T-stat
price ($10k)	α	-0.33	0.02	-18.41
ln mpg	β_m	-0.75	0.17	-4.51
dpm Car small	β_d	-4.70	0.66	-7.13
dpm Car middle	β_d	-5.55	0.55	-10.01
dpm Car large	β_d	-4.68	0.65	-7.16
dpm Car luxury	β_d	-2.13	0.43	-4.98
dpm Car specialty	β_d	-4.02	0.44	-9.17
dpm CUV	β_d	-9.07	2.95	-3.08
dpm SUV	β_d	-8.10	1.13	-7.20
dpm Truck	β_d	-4.17	0.70	-5.98
dpm Van	β_d	-3.65	0.70	-5.20
gdp gr	β	2.06	0.36	5.79
gdp per cap	β	0.02	0.01	1.45
cuv x year	β	0.08	0.03	2.50
suv x year	β	0.11	0.01	8.11
autonews datasource	β	0.34	0.04	8.33
asian	β	-0.06	0.01	-4.52
euro	β	0.14	0.03	5.30
partial year of sales	β	-0.57	0.02	-27.06
$\ln(s_{j\|ss})$ †	σ_{ss}	0.78	0.00	388
$\ln(s_{ss\|s})$ †	σ_s	0.92	0.00	1081
$\ln(s_{s\|t})$ †	σ_v	0.83	0.00	360
Subsegment_dums ‡ (range -4.07 to 0.56)	β	1.28	0.23	4.76

Cost parameters:

Variable		Coeff	S.E.	T-stat
ln mpg	γ	-0.03	0.11	-0.32
year	γ	0.03	0.00	14.72
autonews	γ	0.32	0.05	6.14
asian	γ	0.16	0.02	6.78
euro	γ	0.60	0.03	21.96
Subsegment_dums ‡ (range 9.02 to 10.74)	γ	9.86	0.21	47.87
CAFÉ	λ	$ 347	$ 3.40	102

Notes:

† In addition to being statistically different from 0, log shares are also statistically different from 1. They are, respectively, 108, 97, and 76 standard errors from 1.

‡ There are 31 Subsegment Dummies. Rather than report them all, I report means of absolute values of Coeff, S.E., and T-stat.

N observations = 5,244
R2 from 2sls on demand (for starting values) is 0.95

Table 8: Willingness to Pay

Willingness to Pay for 20% MPG increase

By Segment and By Gas Price

Segment	Small	Middle	Large	Luxury	Specialty	CUV	SUV	Truck	Van
Mean MPG *	31.2	28.3	20.2	20.3	22.0	22.8	16.2	17.2	18.1
20% of Mean MPG	6.2	5.7	4.0	4.1	4.4	4.6	3.2	3.4	3.6
p_gas									
$4.55	$ 1,554	$ 3,457	$ 4,030	$ 556	$ 2,380	$ 7,571	$ 9,964	$ 4,440	$ 3,413
$3.43	$ 720	$ 2,154	$ 2,586	$ (33)	$ 1,342	$ 5,256	$ 7,059	$ 2,895	$ 2,121
$1.50	$ (718)	$ (91)	$ 98	$ (1,047)	$ (446)	$ 1,265	$ 2,054	$ 233	$ (105)

* 2007 Sales-Weighted, Segment-Average MPG

Negative entries are due to the quality reductions required to improve MPG (within a sub-segment).

Table 9: Willingness to Pay without Quality Increase

Without Quality Decrease:
Willingness to Pay for 20% MPG increase

By Segment and By Gas Price

Segment	Small	Middle	Large	Luxury	Specialty	CUV	SUV	Truck	Van
Mean MPG *	31.2	28.3	20.2	20.3	22.0	22.8	16.2	17.2	18.1
20% of Mean MPG	6.2	5.7	4.0	4.1	4.4	4.6	3.2	3.4	3.6
p_gas									
$4.55	$ 3,390	$ 5,293	$ 5,866	$ 2,392	$ 4,215	$ 9,407	$ 11,799	$ 6,276	$ 5,249
$3.43	$ 2,556	$ 3,990	$ 4,422	$ 1,803	$ 3,178	$ 7,091	$ 8,895	$ 4,731	$ 3,957
$1.50	$ 1,118	$ 1,745	$ 1,934	$ 789	$ 1,390	$ 3,101	$ 3,890	$ 2,069	$ 1,730

* 2007 Sales-Weighted, Segment-Average MPG

These WTP's:

1) hold quality constant (while improving MPG)
2) therefore do not reflect the economic model
3) are included for comparison to the previous table (which do reflect the economic model)
4) are strictly larger than their counterparts in the previous table
5) are strictly larger than 0

Table 10: Model Predictions for $3.43 gas

Model Predictions for $3.43 gas

Vehicle Classification			Quantity †			MPG Offered †			MPG Purchased †		
Type	Segment	# Models	2007	2008	2009	2007	2008 ‡	2009	2007	2008	2009
Car		140	7,567	6,861	7,266	22.6	22.6	19.0	26.8	27.6	19.3
				-9%	-4%		0%	-16%		3%	-28%
	Small	29	2,312	2,357	2,350	28.3	28.3	21.1	31.2	31.8	21.0
	Middle	27	2,893	2,136	1,714	24.4	24.4	29.8	28.3	30.4	30.1
	Large	14	674	479	320	19.9	19.9	24.1	20.2	20.3	23.8
	Luxury	49	1,192	1,440	2,551	19.6	19.6	10.3	20.3	20.4	10.0
	Specialty	21	496	451	332	21.4	21.4	18.9	22.0	22.5	18.5
UV		93	4,644	3,577	4,716	18.5	18.5	41.4	19.9	21.9	41.0
				-23%	2%		0%	124%		10%	106%
	CUV	54	2,636	2,320	2,058	20.6	20.6	43.2	22.8	24.7	43.4
	SUV	39	2,008	1,257	2,658	15.5	15.5	39.0	16.2	16.7	39.2
Truck/Van		37	3,673	3,181	3,129	18.0	18.0	20.6	17.4	17.8	21.0
				-13%	-15%		0%	14%		2%	20%
	Pickup	21	2,730	2,264	2,290	18.4	18.4	21.7	17.2	17.6	21.6
	Van	16	943	916	839	17.6	17.6	19.1	18.1	18.4	19.2
All Vehicles		**270**	**15,883**	**13,619**	**15,110**	**20.6**	**20.6**	**26.9**	**22.6**	**23.8**	**26.4**
				-14%	**-5%**		**0%**	**31%**		**5%**	**17%**

† All percentages (shaded) are percentage changes from 2007 levels
‡ 2008 Fuel Efficiency Offerings are fixed at 2007 levels for comparability

Quantity = Sales (1000s)
MPG Offered = Raw-average MPG across models
MPG Purchased = Sales-weighted MPG

2007 - actual
2008 - predicted (before firms adjust characteristics - only consumers respond)
2009 - predicted (after firms adjustt characteristics)

For comparison, actual 2008 (Aug-YTD) Quantities were:
-3%, -19%, -16%, and -14%, for, respectively,
Cars, Utility Vehicles, Trucks/Vans, and All Vehicles.

Table 11: Model Predictions for $4.55 gas

Model Predictions for $4.55 gas

Vehicle Classification			Quantity †			MPG Offered †			MPG Purchased †		
Type	Segment	# Models	2007	2008	2009	2007	2008 ‡	2009	2007	2008	2009
Car		140	7,567	5,878	6,020	22.6	22.6	25.0	26.8	28.6	25.6
				-22%	-20%		0%	11%		7%	-5%
	Small	29	2,312	2,229	1,947	28.3	28.3	27.8	31.2	33.2	27.8
	Middle	27	2,893	1,252	1,420	24.4	24.4	39.6	28.3	35.0	39.9
	Large	14	674	233	265	19.9	19.9	31.6	20.2	20.4	31.6
	Luxury	49	1,192	1,811	2,114	19.6	19.6	13.4	20.3	20.6	13.2
	Specialty	21	496	353	275	21.4	21.4	25.0	22.0	23.6	24.6
UV		93	4,644	2,348	3,907	18.5	18.5	55.1	19.9	25.8	54.4
				-49%	-16%		0%	198%		30%	173%
	CUV	54	2,636	1,873	1,705	20.6	20.6	57.4	22.8	28.0	57.5
	SUV	39	2,008	475	2,202	15.5	15.5	51.8	16.2	17.3	51.9
Truck/Van		37	3,673	2,475	2,593	18.0	18.0	27.3	17.4	18.6	27.8
				-33%	-29%		0%	51%		7%	59%
	Pickup	21	2,730	1,622	1,897	18.4	18.4	28.7	17.2	18.5	28.7
	Van	16	943	853	695	17.6	17.6	25.3	18.1	18.7	25.4
All Vehicles		**270**	**15,883**	**10,701**	**12,520**	**20.6**	**20.6**	**35.7**	**22.6**	**25.7**	**35.0**
				-33%	**-21%**		**0%**	**73%**		**14%**	**55%**

† All percentages (shaded) are percentage changes from 2007 levels
‡ 2008 Fuel Efficiency Offerings are fixed at 2007 levels for comparability

Quantity = Sales (1000s)
MPG Offered = Raw-average MPG across models
MPG Purchased = Sales-weighted MPG

2007 - actual
2008 - predicted (before firms adjust characteristics - only consumers respond)
2009 - predicted (after firms adjustt characteristics)