Gasoline Prices and the Demand for New Vehicles: Evidence from Monthly Sales Data

Joshua Linn
Department of Economics
University of Illinois at Chicago
jlinn@uic.edu

Thomas Klier
Federal Reserve Bank of Chicago
tklier@frbchi.org

March 2007

Preliminary and Incomplete – Comments Welcome

Abstract

There has been considerable recent interest in understanding the relative merits of alternative policies aimed at reducing gasoline consumption. A number of recent studies have estimated the effect of a gasoline tax on the average fuel efficiency (i.e., miles per gallon) of new vehicles, which rely on a number of strong functional form and exogeneity assumptions. We implement a simple approach to estimate the effect of a gasoline tax on the fuel efficiency of new vehicle purchases, using high frequency sales data. When the expected price of gasoline rises, future driving costs rise by less for fuel efficient vehicles, increasing their demand relative to other new vehicles. We use within model-year variation in vehicle sales and expected driving costs to estimate the effect of the price of gasoline on the demand for fuel efficient vehicles. We find that a one dollar increase in the price of gasoline causes the average fuel efficiency of new vehicles to increase by about 0.5 miles per gallon. This estimate is broadly consistent with previous research, which uses cross-sectional consumer survey data.

I. Introduction

Reducing gasoline consumption has returned as a major policy issue, amid concerns about global warming and the desire to reduce oil imports for national security reasons.[1] There are two broad policy alternatives to reduce consumption: raise the federal gasoline tax or raise the corporate average fuel economy (CAFE) standard, which sets a minimum threshold for the average fuel efficiency of new vehicles. Historically, the policy debate has focused on the CAFE standard, though economists have pointed to several advantages of a gasoline tax. A tax would directly affect drivers of both new and used cars, and there is no "rebound" effect, in which a higher CAFE leads to more driving, because the cost of driving is lower. Recently, however, there has been renewed interest in a gasoline tax among policy makers, and a coinciding increase in research.

Raising the gasoline tax may have both short-run and long-run effects. In the short run, consumers continue to drive the same cars as before the tax increase, but they drive fewer miles as a result of the increase in the cost of driving. In the long run, the increase in the price of gasoline raises the demand for fuel efficient vehicles, both in the new and used car markets. This raises the average fuel efficiency of vehicles in use, further reducing consumption of gasoline.

Estimating both responses is complicated by the fact that a consumer's vehicle choice and gasoline consumption are jointly determined. A large body of research (e.g., Goldberg, 1998 and West, 2004) follows Dubin and McFadden (1984), by first estimating the vehicle choice decision and then estimating the effect of gasoline prices on fuel consumption, conditional on vehicle choice. The approach accounts for the fact that

[1] See, for example, hearings at the House Energy and Commerce committee on March 14, 2007.

characteristics affecting a household's vehicle choice decision may also affect the gasoline consumption decision. For example, a parent that drives children large distances to school may also prefer a large automobile. Typically, this approach uses household level data, such as the Consumer Expenditure Survey, to estimate the short and long run effects. These studies rely on a number of assumptions: geographic variation in the price of gasoline is exogenous to other variables that affect the purchase decision, and that substitution patterns between different types of vehicles (e.g., compact versus subcompact) have a certain structure.[2]

We use a much simpler empirical strategy to estimate the long run effect of a gasoline tax, which circumvents these assumptions. We construct a panel of monthly sales of new vehicles by model (e.g., the Ford Taurus) and monthly gasoline prices from 1980-2006. We use a vehicle choice model similar to Berry, Levinsohn and Pakes (1995), BLP. By using monthly data, we are able to estimate a simple linear model. We control for the average sales of each model-year, focusing on changes in expected driving costs that arise because of within-year variation in the price of gasoline. This approach allows us to control for unobserved variables that affect sales, but which do not vary over the model year, such as engine characteristics, size or retail price. Thus, we account for the potentially endogenous relationship between vehicle sales, price and attributes, which is one of the main motivations for the BLP approach.[3] We identify the effect of the price of gasoline on sales by using cross-model variation in fuel efficiency (i.e., miles per gallon),

[2] Several of the more recent papers, e.g., Bento et al. (2006) have investigated the distributional effects of a gasoline tax, arguing that because driving patterns differ across income groups and race, the policy would differentially affect these groups, and in some cases the effect would be regressive.

[3] The transaction price may vary over the model-year, as shown by Corrado, Dunn and Otoo (2004) and Copeland, Dunn and Hall (2005). We intend to address this issue in future work by using transaction price data.

and time series variation in the price of gasoline. That is, the effect on driving costs of a given change in the expected price of gasoline is smaller for fuel efficient vehicles than for "gas guzzlers". This variation allows us to estimate the effect of a gasoline price on the average fuel efficiency of new vehicles.

We find that new vehicle sales respond to the expected price of gasoline by a significant amount. Our estimate implies that a one dollar increase in the price of gasoline, due to an increase in the Federal gasoline tax, for example, would increase the average fuel efficiency of new vehicles purchased by about 0.5 miles per gallon. Another way to interpret the magnitude is to consider two compact car models: the 2006 Pontiac G6 (21 miles per gallon) and the 2006 Volkswagen Jetta (38 miles per gallon).[4] A one dollar increase in the price of gasoline would reduce sales of the Pontiac by about 15 percent more than the Jetta, corresponding to an elasticity of market share with respect to expected driving costs of about -2.

This empirical strategy has several advantages. First, by controlling for unobserved attributes that are constant during the model-year, we do not rely as heavily on functional form assumptions about consumer demand and the distributions of unobserved variables as previous research. In particular, our estimation approach does not impose restrictions on substitution patterns across vehicle classes (e.g., compact) or within classes. Second, by relying on geographic variation in gasoline prices, other studies have assumed that this variation is uncorrelated with unobserved variation in consumer preferences. By comparison, we exploit monthly variation in gasoline prices, which we believe is more likely to be exogenous to consumer preferences. Crude oil prices drive time series

[4] The Pontiac has a much larger engine, and differs in other dimensions from the Jetta. The empirical strategy controls for such differences, as discussed below.

variation, but geographic variation may be due to a large set of observed and unobserved variables (Chouinard and Perloff, 2007). Third, we estimate the relationship between sales and gasoline prices over a nearly 30 year period, in which there was considerable variation in gasoline prices. By examining sub-periods of the data, we can investigate whether consumers' response to gasoline prices depends on the level of gasoline prices. In practice, we find that the response of vehicle sales has been fairly stable over time.

We would like to point out that our results most directly relate to the question of how a gasoline tax would affect average fuel efficiency of new vehicles. To analyze the long-run effect of raising the CAFE standard we would have to incorporate a model of vehicle supply, such as Goldberg (1998).

II. Conceptual Framework: the Effect of the Price of Gasoline on New Vehicle Purchases

This section describes a model of consumer demand for new vehicles, to illustrate the relationship between gasoline prices, fuel efficiency and equilibrium sales. We use a discrete choice model, similar to BLP.

The market for new automobiles contains J varieties, with each variety indexed by j = 1, …, J. Individual i derives utility U_{ij} by purchasing product j according to:

$$U_{ij} = X_j\beta + \alpha(p_j + c_j) + \phi_j + \varepsilon_{ij} \tag{1}$$

The vector X_j consists of attributes of the vehicle that can be observed by the individual and econometrician, such as engine size. The variable p_j is the purchase price of the vehicles; a higher price decreases utility, so α is negative. The variable c_j is the expected cost of operating the vehicle over its lifetime, which also has a negative effect on utility.[5] The last two variables are the mean utility from unobserved vehicle attributes, and an individual- and model-specific error term.

We now discuss these variables in more detail. The attributes in the vector X_j affect the utility of all consumers equally; the coefficient vector β is constant across consumers. This assumption is stronger than the random coefficients specification of BLP, but it is not necessary for the main results, and merely simplifies the notation.

The cost of operating the vehicle, c_j, includes fuel costs and maintenance costs:

$$c_j = F_j + MC_j = \sum_{t=0}^{T} \frac{1}{(1+r)^t} \frac{P_t^{gas}}{MPG_j} M_t + MC_{jt} \tag{2}$$

Total operating costs equal the sum of discounted expected fuel costs and maintenance costs. The lifetime of the car is T years. Fuel costs in year t equal the number of miles driven, M_t, multiplied by the cost of driving the car one mile, P_t^{gas} / MPG_j, where P_t^{gas} is the expected price of gasoline in year t and MPG_j is the fuel efficiency, in miles per gallon (MPG). Total fuel costs, F_j, equal discounted expected fuel costs, using the discount rate r. Maintenance costs in year t equal MC_{jt}, and total maintenance costs, MC_j, equal the sum of discounted expected maintenance costs.

---

[5] Equation (1) imposes the assumption that an increase in purchase price has the same effect on utility as an increase in operating costs, which is made for simplicity.

There are two main approaches to estimate the effect of the price of gasoline on new vehicle demand based on equation (1). The first would be to estimate a nested logit model, derived from an equation similar to (1). This method requires consumer data, and assumptions about substitution patterns across models and the exogeneity of the price of gasoline (or the existence of suitable instruments). Alternatively, the BLP approach would use aggregate data and rely on a valid set of instruments for prices, using a complicated estimation algorithm.

Instead, we assume that the unobserved and observed attributes, as well as the retail price, are constant within a model-year. The assumption is consistent with the typical production process of new vehicles. For most vehicle lines, production begins in July or August, after a brief, one to two week, shutdown period. During that period, the manufacturer may change the characteristics of the vehicle, such as engine size. [6] In practice, changes across model-years range from very minor to a complete overhaul. Under this assumption, a model-year intercept absorbs all unobserved variables in equation (1), as well as the retail price and observed variables that do not vary over the course of the model-year. Note that we assume that deviations from the average retail price are uncorrelated with the other variables in equation (1), though we address this assumption in the empirical work.[7]

This assumption allows us to characterize the aggregate demand for each vehicle in a straightforward manner. Each consumer purchases one of the j vehicles, or the outside

[6] The first month of a model-year may vary by model, which we address in the empirical work.
[7] In future work we hope to incorporate transaction prices in the analysis, addressing the extent to which dealerships respond to demand shocks.

good, which we take to be a used car.[8] We make the standard extreme value assumption for the error term, allowing us to express the market share for good j, s_j, as:

$$s_j = \frac{\exp(\alpha F_j + \widetilde{\phi}_j)}{1 + \sum_{k \in J} \exp(\alpha F_k + \widetilde{\phi}_k)} \tag{3}$$

Following convention, we have normalized the utility of the outside good to zero. The model-specific constant term has absorbed the observable characteristics, X_j, the model price, p_j, and unobserved attributes that do not vary over time. Equation (3) is the standard aggregate logit equation with fixed effects. The market share of vehicle type j decreases when driving costs increase. Note that because of the model fixed effects, we have not restricted cross-price elasticities for the price of output – i.e., the model does not impose independence of irrelevant alternatives. The elasticity of the market share to the price of gasoline, however, is inversely proportional to the fuel efficiency of the model, an assumption we relax below.

III. Estimating the Effect of a Gasoline Tax on Average MPG

If there were only one observation per model-year, it would not be possible to estimate equation (3); the model-year intercepts would be collinear with expected driving costs. In the empirical work we use monthly sales and gasoline price data, giving us 12 observations per model-year.

[8] We plan to directly investigate the effect of the price of gasoline on used vehicle purchases in future work.

We take the natural log of equation (3) and add time and model-year intercepts, to obtain an expression for the log market share of model j in month m and calendar-year y:

$$\ln s_{jmy} = \alpha F_{jmy} + \tau_{my} + \widetilde{\phi}_{jy'} + \nu_{jmy} \qquad (4)$$

The month-year interactions, τ_{my}, control for aggregate shocks to the new vehicle market. The model-year intercepts control for vehicle characteristics that do not change within the model-year, y', such as engine characteristics. Note that the model-year has the subscript y', rather than y. The subscript y' is defined differently for each model, according to the month production begins. For example, the 2005 model-year for the Honda Civic spans October 2004 to September 2005, while the 2005 model-year for the Toyota Prius spans July 2004 to June 2005.

We assume that the price of gasoline follows a random walk.[9] The expected price at time $t' > t$ is equal to the price at time t. As a result, the expected cost of driving the vehicle is proportional to the current price of gasoline, divided by the miles per gallon of the vehicle. We assume that expected maintenance costs do not change within the model-year. We use equation (2) to replace F_{jmy} in equation (4), which yields our estimating equation:

$$\ln s_{jmy} = \alpha \frac{P_{my}^{gas}}{MPG_{jy'}} + \tau_{my} + \widetilde{\phi}_{jy'} + \nu_{jmy} \qquad (5)$$

[9] The estimates are similar if we assume other statistical processes for the price of gasoline, as reported below.

Equation (5) is the baseline estimating equation. The dependent variable is the log sales share of model j in month m and year t. The first independent variable is the expected cost of driving the car one mile at the time of purchase of the vehicle; P_{my}^{gas} is the price of gasoline by month and year, and $MPG_{jy'}$ is the fuel efficiency of model j in model-year y'. We refer to the ratio as dollars-per-mile. The next section describes the variable construction in detail.

The coefficient of interest is α, which is the effect on sales of an increase in the cost of driving one mile. We identify this parameter using time-series variation in the price of gasoline and cross-sectional variation in average fuel efficiency. The driving cost of a fuel efficient vehicle increases less than that of a "gas guzzler" when the expected price of gasoline increases. A consumer deciding between two otherwise identical vehicles is more likely to purchase the fuel efficient vehicle when the expected price of gasoline is high. That is, within-year variation of the price of gasoline affects expected driving costs for all vehicles, but by different amounts. We estimate the effect of driving costs on sales, relative to the average sales of each model.

We briefly discuss the interpretation of α. An increase in the price of gasoline changes the demand for fuel efficient vehicles, relative to other vehicles. Such a demand shock may have two effects. First, dealers may adjust prices in response to the change in relative demand.[10] Second, firms may vary production of the vehicles they produce within the model-year (Bresnahan and Ramey, 1993 and Copeland and Hall, 2005). Both responses would represent a movement along the supply curve, though previous research

[10] We next discuss changes in transaction prices that are correlated with, but not caused by, changes in gasoline prices. As noted earlier, we hope to incorporate transaction prices in the empirical analysis in a later version of the paper.

suggests the effect should be small. Consequently, we interpret α as the change in equilibrium demand, resulting from the horizontal shift of the demand curve and an upward movement along the supply curve for each vehicle type.[11]

We now discuss the main identification assumptions in equation (5). The model-year intercepts and monthly variation are central to our empirical strategy. The model-year intercepts account for the potentially endogenous relationship between the average transaction price and vehicle characteristics (Nevo, 2000). For a given model-year, we estimate the effect of an increase in expected driving costs on vehicle sales, relative to the average sales of that model-year. We measure the change in driving costs relative to the average, so we are exploiting monthly variation in gasoline prices and variation across models in fuel efficiency, which are quite substantial (see section IV).

In equation (5) the elasticity of the market share with respect to the price of gasoline is inversely proportional to the MPG of the model, and does not depend on the characteristics of other models. Below, we report a number of alternative models, which relax this functional form assumption, yielding similar results.

The main assumption in this framework is that the price of gasoline is uncorrelated with variables that change over the model-year, of which there are several types. First, preferences for fuel-efficient vehicles may change over time. Sales of sport utility vehicles increased dramatically in the 1990s, which may partially be due to relatively low gasoline prices, but may also be due to changes in consumer preferences for vehicle size. In equation (5) we assume that within model-year changes in preferences are uncorrelated

[11] From the model in section II, the coefficient α is proportional to miles driven, which is defined as the number of miles driven per year, conditional on vehicle choice. The price of gasoline may affect this variable; the introduction referred to this effect as the short run response to a gasoline tax. Note that consumers account for this effect in making their purchasing decision.

with the price of gasoline. This assumption appears to be reasonable, given that we obtain similar estimates of α across vehicle types and over time.

The second type of time-varying variable is transaction prices. The model-year intercepts absorb average transaction prices, but as Copeland *et al*. (2005) and Corrado *et al*. (2006) have documented, transaction prices vary within the model-year. Prices decline dramatically over the course of the model-year, and sales follow a "hump-shaped" pattern, peaking in the early summer. These patterns raise the possibility that changes in dollars-per-mile are correlated with transaction prices, biasing the estimate of α.

The baseline specification partially addresses this concern by including month-year interactions. These interactions control for the average sales and transaction price profiles for all models. Copeland *et al*. (2005) and Corrado *et al*. (2006), however, found that sales and price profiles vary across vehicle types (e.g., compact or midsize). We report an additional specification to address this concern, allowing for separate price/sales profiles by vehicle type. We separate models into five categories, based on average fuel efficiency (e.g., less than 18 MPG). We re-estimate equation (5), including interactions of a set of category indicator variables with a set of indicator variables for the number of months since the beginning of the model-year. The coefficient α is identified by within-category changes in dollars-per-mile and sales. We obtain similar results to the baseline, suggesting this is not a likely source of bias.[12]

Finally, equation (5) may omit income effects, either correlated with, or caused by, the price of gasoline. For example, an increase in the price of gasoline might precede a recession, reducing expected income. We have assumed that income effects do not vary by vehicle type, but in practice, demand for smaller, more fuel efficient vehicles, may

[12] In future work we will use transaction price data to explore this issue further.

increase. In that case, we would also expect the demand for used cars to increase. We intend to investigate this issue in future work, analyzing the effect of gasoline prices on the demand for used cars.

IV. Data

We construct the real price of gasoline using data from the Department of Energy and Bureau of Economic Analysis websites. We use the monthly national average retail price of unleaded gasoline from 1980-2006, and the monthly consumer price index (CPI) over the same period. The real price of gasoline, P_{my}^{gas}, is the price of gasoline divided by the CPI, where we normalize the CPI to one for the last month of the sample (May, 2006). Vehicle sales are from Ward's AutoInfoBank. We use monthly sales data by individual model from 1980 to 2006. We supplement the sales data with vehicle characteristics data. This information is available in print in the annual Ward's Automotive Yearbooks (1980-2003). Vehicle characteristics include wheelbase, curb weight, engine size and mileage rating.[13] We match the annual vehicle characteristics with the monthly sales data.[14]

Figure 1a shows the price of gasoline and the sales-weighted average miles per gallon from 1980-2006. Both variables vary considerably over time. The price of gasoline

[13] The mileage ratings reported in the Ward's Automotive Yearbooks are reported for each of the vehicle specifications listed. The mileage rating is shown for the "default", or standard, transmission (either manual or automatic, depending on the model/manufacturer). Because of the greater detail in vehicle specifications reported, the mileage ratings don't exactly match the published EPA ratings, from fueleconomy.gov. We are in the process of substituting the EPA ratings for each model.

[14] The match is not straightforward because the two data sets are reported at different levels of aggregation. Vehicle characteristics data are reported at the "trim level" to recognize differences in the MRSP; for example, the data distinguish the 2- and 4-door versions of the Honda Accord sedan. We aggregate the data to match the model-based sales data, and calculate four statistical moments for the distribution of the vehicle characteristics by car line (minimum, maximum, mean and median). We use the mean value to estimate equation (5), but obtain similar results using other measures.

declined steadily in the early 1980s, after which the price fluctuated around an average of about $1.60, and then rose steadily beginning in 2002. Average miles per gallon increased during the 1980s as the CAFE standard increased, then declined steadily until the late 1990s, remaining constant thereafter. Most of the decrease was due to the rise in sales of sport utility vehicles, which were subject to the lower CAFE standard for light trucks; the average fuel efficiency of new cars has been stable at around 27 MPG, though there is some variation around the mean (not shown).

Figure 1b shows the log price of gasoline and average fuel efficiency, removing annual means from both series. The figure shows that there is a positive correlation between the two variables. Within-year increases in the price of gasoline are associated with increases in the average fuel efficiency of new vehicles. The model-level estimates of equation (5), discussed in the next section, reflect this pattern.

The estimation strategy requires substantial variation of fuel efficiency across vehicles. Figure 2 plots frequency distributions of the number of models available in 1980, 1990 and 2000, by MPG bin. In all three years there is a wide variety of automobiles available. The distributions are fairly stable, with a peak around 22 MPG, and a small number of models available with more than 30 or less than 18 MPG. There appears to be a slight increase over time in the number of models in the high 20s, and around 22 MPG.

Table 1 shows additional summary statistics. The first panel shows the mean and standard deviation of the monthly observations of the price of gasoline and the MPG of available models each decade. The MPG distribution was pretty stable, reflecting the patterns shown in Figure 2. The real price of gasoline was lower in the 1990s and less volatile than in other years. The last two rows in Panel A show the main independent variable, dollars-per-mile, and the dependent variable, model sales. Average sales per model were lower in the 1990s and 2000s than in the 1980s, reflecting the fact that there are more models available in the latter periods. Panel B shows the standard deviations of the dollars-per-mile and sales variables, after removing model year and month-year means (i.e., the interactions in equation (5)). This panel shows that the interactions absorb much of the variation, but there is some variation in all three decades for the two variables.

V. Results

A. Main Specification

Table 2 reports the baseline estimate of α in equation (5). The dependent variable is the log share of sales by model, month and year. The main independent variable is dollars-per-mile, defined as the real price of unleaded gasoline divided by the fuel efficiency of the model, in MPG. The regression includes model-year and month-year interactions, and is estimated by Ordinary Least Squares. The first row of Table 2 reports the estimate of α, the coefficient on dollars-per-mile. Standard deviations are in parentheses, clustered by model.

The estimate of α in column 1 is -8.03 with standard error 2.70, which is significant at the one percent level. To interpret this magnitude, consider the Pontiac G6 (21 MPG) and the Volkswagen Jetta (38 MPG). The estimate of α implies that a one dollar increase in the price of gasoline would reduce sales of the Pontiac by about 15 percent more than the

Jetta. Assuming that both vehicles would last 10 years and would be driven 12,000 miles per year, the estimate implies an elasticity of market share with respect to operating costs of about -2.

We use this result to estimate how much fuel efficiency of new vehicles would change if an increase in the Federal gasoline tax causes the price of gasoline to increase by one dollar.[15] We use the set of models in the market at the end of the sample, May 2006, to calculate the weighted average miles per gallon, where the weights are the predicted market shares under a one dollar price increase. We subtract from this value the sales-weighted average miles per gallon, using market shares in the absence of the policy as weights. We calculate the standard error using the delta method, and report at the bottom of Table 2 that a one dollar price increase would raise average fuel efficiency by 0.48 MPG.[16]

The estimation equation relies on functional form assumptions about the underlying utility function, as well as the distributions of unobserved vehicle, consumer and cost characteristics. We report several additional specifications to check the robustness of these assumptions.

We first use the log of the variable dollars-per-mile in place of the level. The coefficient α is the elasticity of demand with respect to driving costs. Because of the sources of variation of the price of gasoline and fuel efficiency, we cannot estimate this specification with month-year and model-year interactions as in column 1; instead, we include model fixed effects. For comparison, we first show the result of estimating

[15] The actual tax increase that causes a one dollar increase in the retail price of gasoline may be larger than one dollar.
[16] We obtain similar, though not identical, results if we calculate the effect of the tax increase at different points in the sample.

equation (5) with model fixed effects. The estimate of α is similar to the baseline, -8.65 with standard error 2.85. The estimate of α using log dollars-per-mile, in column 3, is -0.42 with standard error 0.30. The estimate is not statistically significant, but the implied effect of a one dollar increase in the price of gasoline is quite similar to column 2.

In column 4 we replace dollars-per-mile with the interaction of the log price of gasoline with log fuel efficiency. The coefficient on this variable is expected to be positive, as an increase in the price of gasoline should raise the demand for vehicles with greater fuel efficiency. The estimated coefficient on the interaction is 0.70 with standard error 0.25, which is significant at the one percent level. The effect of a one dollar price increase is again similar to the baseline, in column 1.

Next, we estimate a similar model to the almost ideal demand system (Deaton and Mullbauer, 1984). We estimate a separate regression for each model, where the dependent variable is the share of vehicle sales by month and year, and the independent variables are the log price of gasoline and a set of year dummies. As in equation (5), we identify the effect of the price of gasoline on vehicle demand using within model-year changes in the price of gasoline and vehicle sales. This specification does not rely on the same functional form equations as equation (5); rather, it can be derived from a second order approximation to an arbitrary expenditure function. Figure 3a shows a histogram of the coefficients, and Figure 3b plots the coefficient against the fuel efficiency for each model. The average coefficient is small and positive, and there is considerable variation across models. The estimated coefficient is positively correlated with average miles per gallon, shown in the bottom figure, meaning that when the price of gasoline increases, the sales of fuel efficient vehicles increase relative to other vehicles. Using these coefficient estimates, a one dollar price increase would raise average miles per gallon by 0.36, which is similar to the baseline.

Finally, we further relax the functional form assumptions in equation (5) by estimating an aggregate equation, using the average miles-per-gallon of new vehicles and the monthly price of gasoline. We estimate the following equation:

$$\ln MPG_{my} = \delta \ln P_{my}^{gas} + \mu_m + \tau_t + \upsilon_{my} \tag{6}$$

The dependent variable is the log of the sales-weighted average MPG of new vehicles sold in month m and year y. The first independent variable is the log price of gasoline in month m and year y. The regression includes month and year dummies, as in Figure 1b. The coefficient δ is the elasticity of average miles per gallon with respect to the price of gasoline. This coefficient is easier to interpret than α in equation (5), but the estimate of δ would be biased if model entry were correlated with the price of gasoline. As column 5 shows, the estimate of δ is -0.03 with standard error 0.01, which is significant at the one percent level. The estimate implies that a one dollar increase in the price of gasoline in 2006 would raise the average fuel efficiency of new vehicles by 0.28 miles per gallon, which is similar to the model-level results.

B. Effect of Gasoline Price by Vehicle Category

In this section we investigate whether the response to a gasoline price increase varies by vehicle type. We first separate vehicles based on the nationality of the company: Honda/Toyota, United States, and other. Over the sample period Honda and Toyota automobiles tend to be smaller and more fuel efficient. Customers of these vehicles may respond by different amounts to a gasoline price change than other customers. Column 1 of Table 3 includes the interactions of the dollars-per-mile variable with two dummy variables, the first equal to one if the vehicle is produced by a US company, and the second equal to one if the company is Honda or Toyota. This specification allows the response to a gasoline price increase to be different for these categories, relative to all other vehicles. We find that purchases of vehicles produced by US companies respond more than other types of vehicles, by a statistically and economically significant amount; the effect of a price increase is about twice as large for US vehicles.

In column 2 we separate cars from light trucks. Column 3 distinguishes vehicles according to the number of cylinders in the engine; 6 cylinders is the most common category in the data set. Column 4 separates vehicles into three categories based on MPG and column 5 separates vehicles into three categories according to retail price (the results are similar if we use narrower category definitions than in columns 4 and 5, not reported).

Overall, the results suggest that the demand for small, fuel efficient and less expensive vehicles is more sensitive to the price of gasoline. Given that these vehicles' characteristics are highly correlated with one another, it is difficult to distinguish which features are most important.

C. Additional Specifications

Table 4 reports several additional robustness checks. As noted earlier, in the baseline specification we assume that the price of gasoline is uncorrelated with other time-varying attributes of the model, such as sales price. Previous work, e.g., Copeland *et al*. (2005)

has found that transaction prices vary differentially within the model-year for different vehicle types. We address this possibility by separating models into five categories, based on fuel efficiency. We then construct a variable for the number of months since the beginning of the model-year. For example, a model that begins selling in September has a value of 1 in October, and 2 in November. By comparison, a model that begins selling in August has a value of 2 in October. We interact a set of dummy variables for the number of months with a set of category indicators. This specification allows models in each fuel efficiency category to have different transaction price and sales profiles. The estimate of α using this specification is reported in column 1 of Table 4. The result is quite similar to the baseline.

The sales data includes sales of consecutive model-years. For example, the sales of the Honda Civic in September 2005 may include sales of the 2005 Civic and the 2006 Civic. This data limitation creates two potential difficulties. First, sales of each model-year may follow the hump-shaped pattern documented by Corrado *et al*. (2003) and Hall and Copeland (2005), but the combined sales of the two model-years may follow a different pattern. However, as long as models in the same fuel efficiency category (as defined in column 1) follow similar patterns, and the beginning of the model-year does not change systematically over time, the specification in column 1 should address this concern.

Second, the fuel efficiency of some models changes over time. Therefore, the sales for a given model in a particular year and month may include vehicles with different fuel efficiencies. This would create measurement error in the dollars-per-mile variable, biasing the estimate of α. Given that such changes in fuel efficiency are generally quite small, the bias is not likely to be large. We can assess the possible magnitude of the bias by using the average fuel efficiency for the current and previous year to compute the dollars-per-mile variable. The results in column 2 are nearly identical to the baseline.

Columns 3 and 4 use alternative measures of the price of gasoline.[17] The baseline specification is unbiased if consumers assume the price of gasoline follows a random walk; in that case, the expected price is equal to the future price. This assumption is consistent with the statistical process the price followed over the sample period, but it is possible that the average consumer uses an alternative model to forecast the price. For comparison, in column 3 we assume that consumers use an autoregressive process, with two lags, to forecast prices. In column 4 we assume that consumers use the average price over the previous three months to forecast the future price. The estimate in column 3 is significant at the one percent level, but the magnitude reflects the greater variance of the current price, relative to the forecasted price. The estimate of α is also somewhat smaller using the three month average.

Finally, we consider whether the response of demand to the price of gasoline varies over time. We construct a dummy variable, equal to one beginning in January, 1990. We add to equation (5) the interaction of this variable with the dollars-per-mile variable. The coefficient on the interaction allows the demand response to differ in the second time period. The estimate of this interaction, reported in column (5), is small and not statistically significant. We obtain similar results using other time periods.

---

[17] In future work, we plan to further explore the sensitivity to alternative measures of the price of gasoline. For example, we hope to include crude oil futures prices and to consider nonlinear models. For example, the demand effect may be larger when the price of gasoline crosses some threshold.

VI. Conclusion

We estimate the effect of the price of gasoline on the average fuel efficiency of new vehicles. The empirical strategy combines time series variation of the price of gasoline with cross sectional variation of fuel efficiency, exploiting the fact that the effect of a given price change on future driving costs is inversely proportional to fuel efficiency. We find that the price of gasoline has a significant effect on the demand for fuel efficient vehicles. If an increase in the Federal gasoline tax causes the price of gasoline to rise by one dollar, average fuel efficiency of new vehicles would rise by about 0.5 miles per gallon.

These results are similar to other studies, e.g., Bento *et al*. (2006), which jointly estimate a demand and production model. Those papers estimate the effect of a gasoline tax on equilibrium sales, allowing both producers and consumers to respond, but holding the menu of available vehicle models fixed. The results of this paper are directly comparable, because we measure the effect of the price of gasoline on sales, allowing both producers and consumers to respond to the price, but holding fixed the set of vehicles on the market. An open question is whether an increase in the price of gasoline would stimulate the entry of new vehicles with different fuel efficiency characteristics.

In future work we will incorporate several sources of additional data. First, we will use transaction price data to estimate the effect of gasoline prices on sales prices. Second, we will incorporate data on used car prices and sales to investigate the effect of the price of gasoline on the new/used vehicle margin. Finally, we hope to include monthly sales data from the 1970s, when the price of gasoline varied dramatically.

References

1. Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen and Roger H. von Haefen (2006), "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes."
2. Berry, Steven, James Levinsohn and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*: Vol. 63, n4, 841-890.
3. Chouinard, Hayley H. and Jeffrey M. Perloff (2007), "Gasoline Price Differences: Taxes, Pollution Regulations, Mergers, Market Power, and Market Conditions," *The B.E. Journal of Economic Analysis & Policy*: Vol. 7 : Iss. 1 (Contributions), Article 8
4. Copeland, Adam, Wendy Dunn and George Hall (2005), "Prices, Production and Inventories Over the Automotive Model Year."
5. Copeland, Adam and George Hall (2005), "The Response of Prices, Sales and Output to Temporary Changes in Demand."
6. Corrado, Carol, Wendy Dunn and Maria Otoo (2006), "Incentives and Prices for Motor Vehicles: What Has Been Happening In Recent Years?", *Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System*, 2006-09.
7. Deaton, Angus and John Muellbauer (1980), "An Almost Ideal Demand System," *The American Economic Review*: Vol. 70, n3, 312-326.
8. Dubin, J. A. and D. L. McFadden (1984), "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption," Econometrica: Vol. 52, n2, 345-362.
9. Goldberg, Pinelopi Koujianou (1998), "The Effects of the Corporate Average Fuel Efficiency Standards in the US," *The Journal of Industrial Economics*: Vol. 46, n1, 1-33.
10. Nevo, Aviv (2000), "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics and Management Strategy*: Vol. 9, n4, 513-548.
11. Ward's Automotive Yearbook, 1980-2003, Ward's Communications
12. Ward's AutoInfoBank, Ward's Automotive Group
13. West, Sarah E. (2004) "Distributional Effects of Alternative Vehicle Pollution Control Policies," *Journal of Public Economics*: Vol. 88, 735-757.

Figure 1a: Average MPG and Real Price of Gasoline, 1980-2006

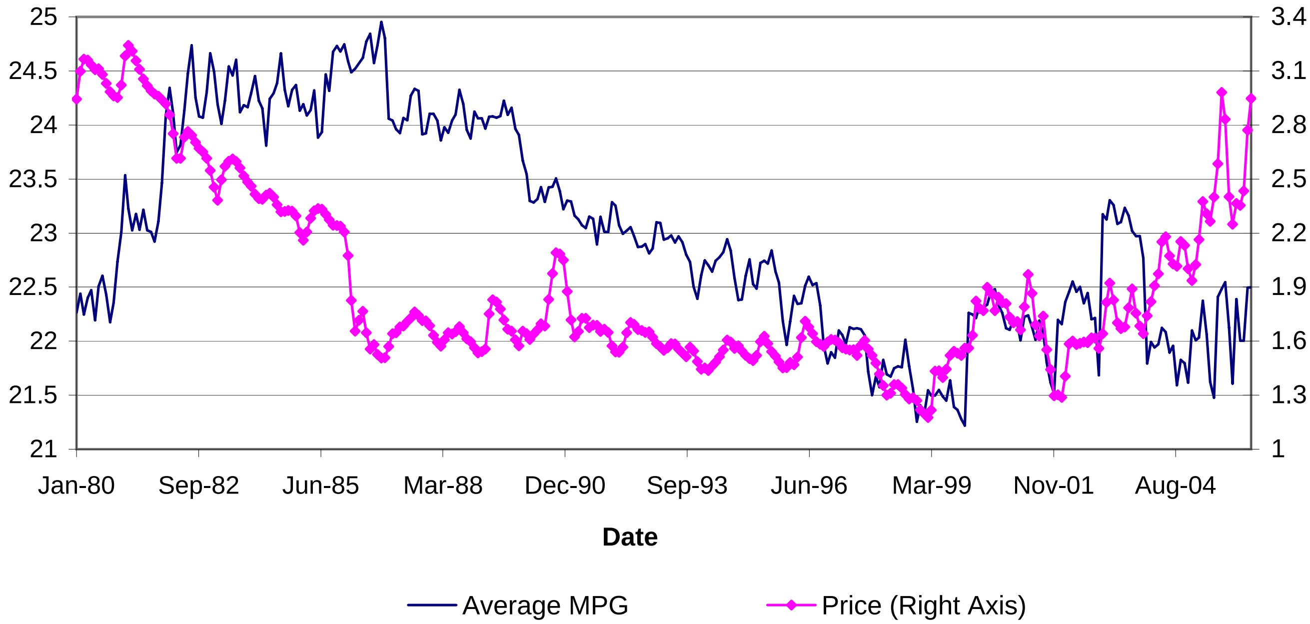

Figure 1b: Demeaned Log Average MPG and Log Gasoline Price

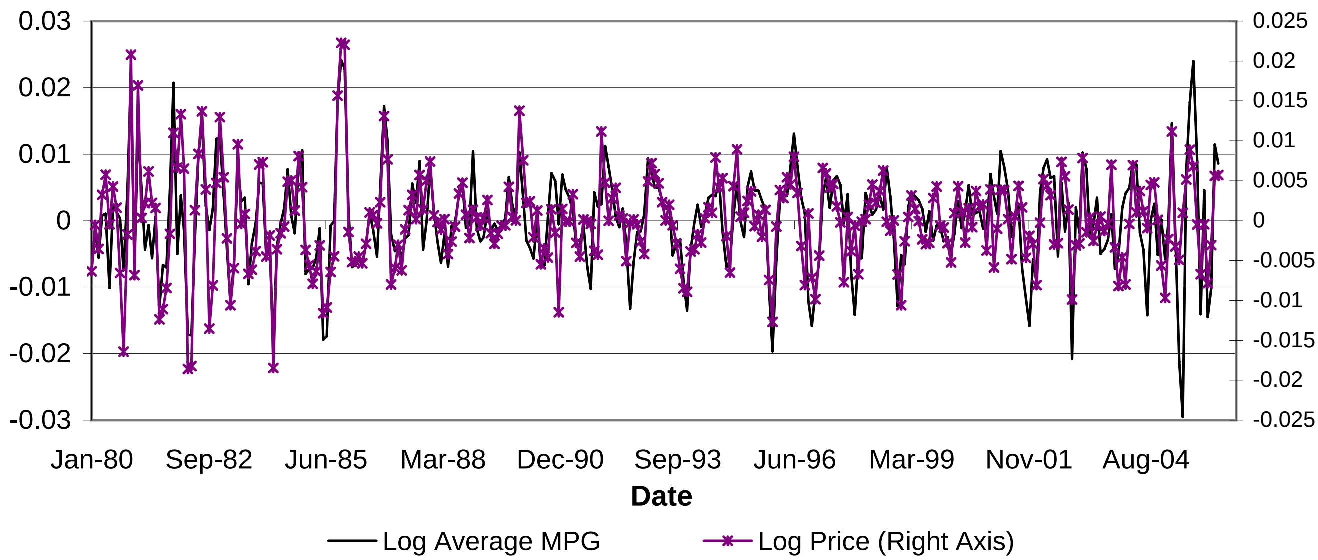

Notes: Average miles per gallon is the sales-weighted average miles per gallon of vehicles sold in the Wards database, by year and month. Figure 1a plots average miles per gallon from 1980-2006, and Figure 1b plots the monthly log of average miles per gallon, after subtracting the annual mean of the variable. The real price of gasoline is the price of unleaded gasoline divided by the consumer price index, using the national average gasoline price from the Department of Energy and the consumer price index (CPI) from the Bureau of Labor Statistics. The CPI is normalized to one for May, 2006. Figure 1a plots the real price of gasoline, in 2006 dollars, and Figure 1b plots the log real monthly price, after subtracting the annual mean (see text for details).

Figure 2: MPG Distribution for 1980, 1990 and 2000

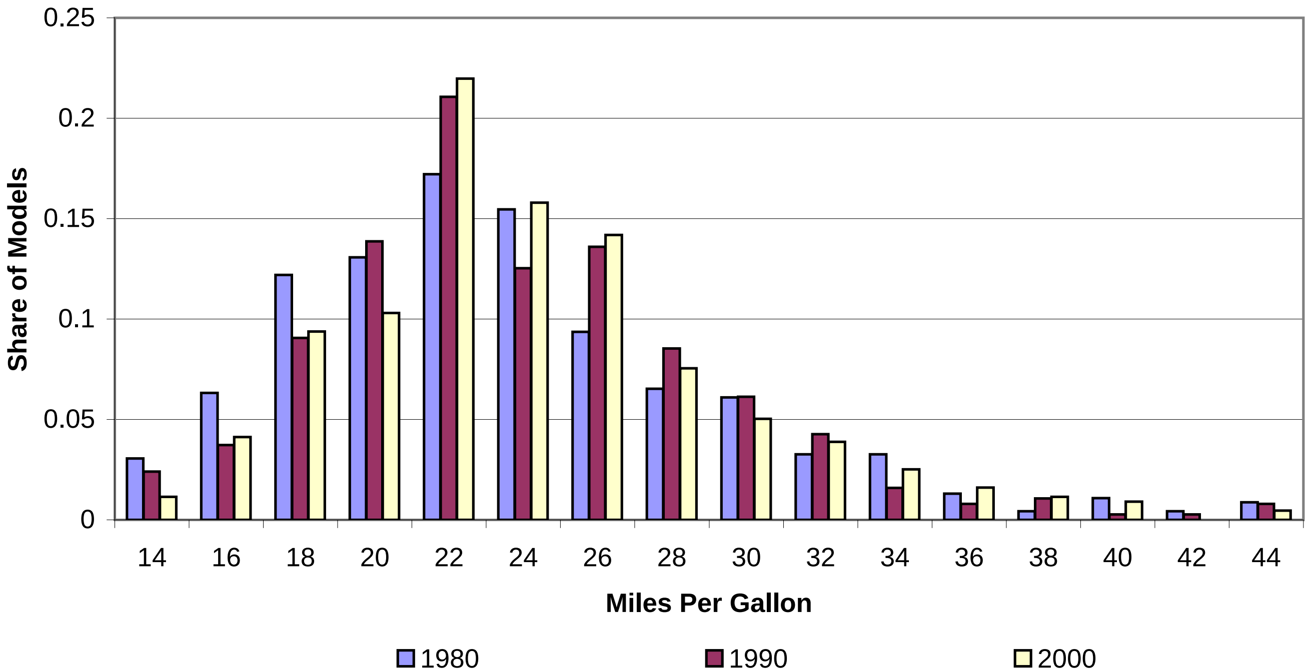

Notes: The figure plots a histogram of the fuel efficiency, in miles per gallon, of vehicles in the Wards database in the indicated years. The vertical axis is the share of models with positive sales for the corresponding year. The horizontal axis labels the maximum fuel efficiency of vehicles in the bin.

Figure 3a: Histogram of Coefficient on Log Gasoline Price

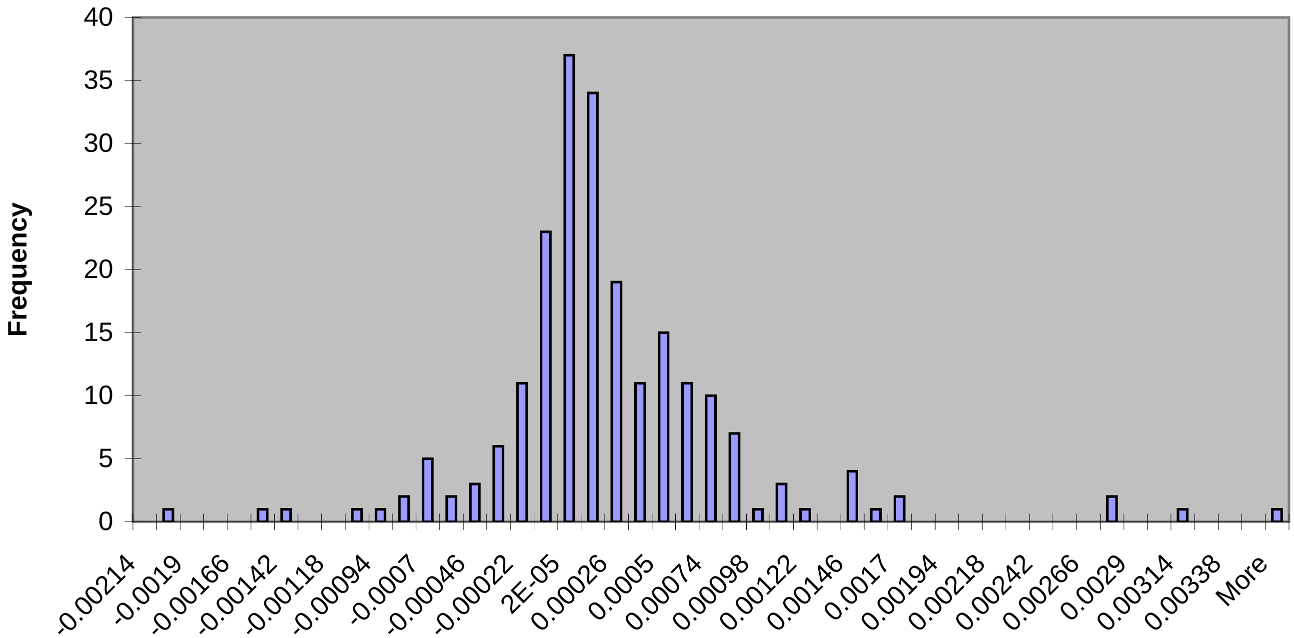

Figure 3b: Coefficient on Log Gasoline Price vs. Average Fuel Efficiency

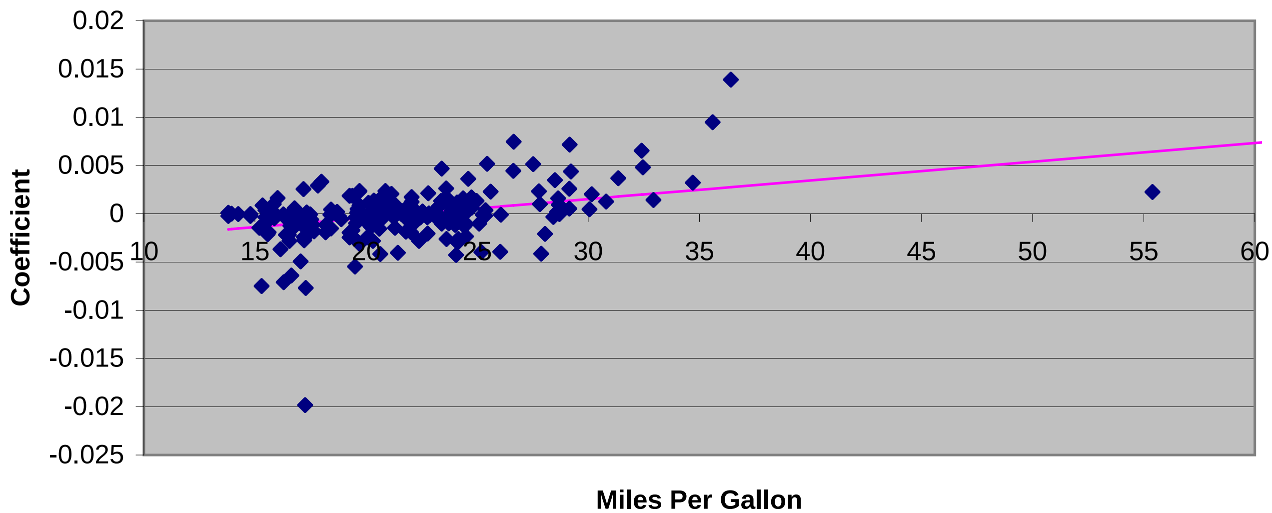

Notes: Figure 3a reports a histogram of the estimated coefficients from separate model level regressions. The dependent variable is the share of sales in total sales by month and year for the particular model. The independent variables are a set of year dummies and the log real price of gasoline. The histogram shows the number of models for which the estimated coefficient falls within the indicated bin. Figure 3b plots the coefficients against the miles per gallon of the model, and plots the fitted values of an OLS regression of the coefficient on miles per gallon.

Table 1:

Summary Statistics

	1980-1989	1990-1999	2000-2006
		Panel A: Sample Statistics	
Miles Per Gallon	23.26	22.99	22.57
	(5.79)	(5.53)	(6.38)
Gasoline Price	2.27	1.56	1.90
	(0.56)	(0.16)	(0.37)
Dollars-Per-Mile	0.10	0.07	0.09
	(0.04)	(0.02)	(0.03)
Log Sales Per Model	7.90	7.48	7.71
	(1.56)	(1.97)	(1.81)
		Panel B: Standard Deviation After Removing Model-Year Means	
Dollars-Per-Mile	0.0013	0.0009	0.0017
Sales Per Model	0.39	0.51	0.43

Notes: Cells report means, with standard deviations in parentheses. The first row of Panel A reports the average miles per gallon of all models sold in the indicated decade. The second row reports the monthly real gasoline price, computed as in Figure 1. The third row reports the ratio of the price of gasoline to miles per gallon, using all monthly observations. The fourth row reports log monthly sales by model. Panel B reports the standard deviation of the indicated variables, after removing model-year means.

Table 2:

Effect of Gasoline Price on New Vehicle Sales, 1980-2006

	Baseline: Equation (5)	Model Fixed Effects	Log Cost Per Mile With Model Fixed Effects	Interact Log Gas Price With Log MPG	Estimate Aggregate Equation, (6)
	(1)	(2)	(3)	(4)	(5)
Dollars-Per-Mile	-8.03 (2.70)	-8.65 (2.85)			
Log Dollars-Per-Mile			-0.42 (0.30)		
Log Gas Price X Log MPG				0.70 (0.25)	
Log Gas Price					0.03 (0.01)
Number of Observations	63,098	63,098	63,098	63,098	317
R^2	0.94	0.62	0.62	0.94	0.96
Dependent Variable	Log Share of Sales	Log Share of Sales	Log Share of Sales	Log Share of Sales	Log Avg MPG
Effect of One Dollar Price Increase on MPG	0.48 (0.25)	0.42 (0.25)	0.45	0.29	0.28 (0.08)

Notes: Standard deviations are in parentheses, clustered by model in columns 1-4 and robust to heteroskedasticity in column 5. Columns 1-4 report the results from estimating equation (5), and column 5 reports equation (6). The dependent variable in columns 1-4 is the log share of sales, by model, month and year. The regressions include month-year interactions, model fixed effects in columns 2 and 3, and model-year interactions in columns 1 and 4. Columns 1 and 2 report the coefficient on dollars-per-mile, defined as in Table 1. Column 3 reports the coefficient on log dollars-per-mile, and column 4 reports the coefficient on the interaction of the log real price of gasoline with the log miles per gallon of the model. In column 5 the dependent variable is the sales-weighted average miles per gallon of all vehicles sold, by month and year. The regression includes month-year interactions, and the column reports the coefficient on the log price of gasoline. The bottom row of the table reports the effect of a one dollar increase in the price of gasoline on the average miles per gallon of new vehicles sold. In column 1-4, the calculation uses the predicted market shares of models sold in May, 2006, with and without the price increase. The standard deviation is in parentheses, calculated using the delta method. The effect of the price increase in column 5 is the change in average miles per gallon if the price increases by one dollar, relative to the price in May, 2006.

Table 3:

Effect of Gasoline Price by Vehicle Type

	(1)	(2)	(3)	(4)	(5)
			Dependent Variable: Log Sales		
Dollars-Per-Mile	-4.76 (3.05)	-6.98 (2.88)	-3.82 (5.05)	-11.00 (3.94)	-14.27 (2.80)
Dollars-Per-Mile X USA	-3.36 (1.31)				
Dollars-Per-Mile X Honda/Toyota	0.54 (1.98)				
Dollars-Per-Mile X Car		1.53 (1.36)			
Dollars-Per-Mile X 6 Cylinder			-3.52 (2.18)		
Dollars-Per-Mile X 8 Cylinder			2.10 (2.51)		
Dollars-Per-Mile X Med Efficiency				-2.22 (2.06)	
Dollars-Per-Mile X High Efficiency				-2.47 (4.09)	
Dollars-Per-Mile X Mid-Price					4.06 (0.53)
Dollars-Per-Mile X Expensive					4.93 (0.57)
Number of Observations	63,098	63,098	63,098	63,098	63,098
R^2	0.94	0.94	0.94	0.94	0.94

Notes: Standard deviations are in parentheses, clustered by model. The table reports the results of estimating equation (5), with the same dependent and independent variables as in column 1 of Table 2. The regressions include interactions of dollars-per-mile with different indicator variables. Column 1 includes variables for whether the model is produced by a US company, and whether it is produced by Honda or Toyota. Column 2 uses a dummy equal to one if Wards labels the model as a car. Column 3 includes dummies equal to one if the engine has six or eight cylinders. Column 4 uses a dummy equal to one if the model attains 18-26 miles per gallon and a dummy if the car attains greater than 26 miles per gallon. Column 5 uses dummy variables for the top third of the distribution of model retail prices and a dummy for the middle third of the distribution.

Table 4:

Additional Specifications

	Include Model-Month by MPG Category Interactions	Use Average MPG Across Models	Use Predicted Gasoline Price From AR(2)	Use Average Gasoline Price for Previous Three Months	Include Post-1990 Interaction
	(1)	(2)	(3)	(4)	(5)
			Dependent Variable: Log Sales		
Dollars-Per-Mile	-8.69 (2.75)	-8.04 (2.78)	-0.29 (0.09)	-4.76 (3.06)	-7.10 (3.12)
Dollars-Per-Mile X Post-1990					-1.27 (2.52)
Number of Observations	63,098	63,098	63,098	62,854	63,098
R^2	0.94	0.94	0.94	0.94	0.94

Notes: Standard deviations are in parentheses, clustered by model. The dependent variables is log sales by model, month and year. The independent variables are the same as in column 1 of Table 2 in columns 1 and 5. Column 2 uses the average miles per gallon for the previous and current years to construct cost per mile. The forecasted gasoline price in a given month and year is the average discounted expected gasoline price for the subsequent 10 years, using an autoregressive specification with two lags to compute the expected price. Column 3 uses the forecasted gasoline price in place of the current price. Column 4 uses the average price over the previous three months in place of the current price. The MPG category is a set of indicator variables, equal to one if the average fuel efficiency falls in the corresponding range: less than 18 miles per gallon (MPG), 18-22 MPG, 23-26 MPG, 27-30 MPG and greater than 30 MPG. Column 1 includes the interactions of these category variables with a set of indicator variables for the number of months since the beginning of the model-year. Column 5 includes the interaction of dollars-per-mile with an indicator equ