

Automakers' Short-Run Responses to Changing Gasoline Prices*

Ashley Langer
University of Michigan

Nathan H. Miller
U.S. Department of Justice

January 2012

Abstract

We provide empirical evidence that automobile manufacturers use cash incentives to offset how gasoline price fluctuations affect the expected fuel expenses of automobile buyers. Regressions based on a database of incentives over 2003-2006 suggests that, on average, manufacturers offset 40% of the change in relative fuel costs between vehicles due to gasoline price fluctuations. The results highlight that (1) carbon taxes and emissions trading programs likely would generate substantial substitution within vehicle classes, and (2) studies that ignore manufacturer discounting likely underestimate consumer demand for fuel economy. The results also have implications for the optimal design of feebate programs.

Keywords: automobile prices, gasoline prices, environmental policy

JEL classification: L1, L9, Q4, Q5

*We thank Severin Borenstein, David Card, Joseph Farrell, Luke Froeb, Richard Gilbert, Ana Maria Herrera, Ryan Kellogg, Patrick Kline, Joshua Linn, Kenneth Train, Clifford Winston, Catherine Wolfram, Florian Zettelmeyer, and seminar participants at the Bureau of Labor Statistics, the College of William and Mary, the George Washington University, the University of California, Berkeley, the University of Michigan, and the University of Wisconsin for valuable comments. Daniel Seigle and Berk Ustun provided research assistance. Langer: Ford School of Public Policy and Department of Economics, University of Michigan, 735 S. State St. #5224, Ann Arbor, MI 48109. Miller: Economic Analysis Group, U.S. Department of Justice, 600 E St. NW, Suite 10000, Washington DC 20530. The views expressed are not purported to reflect those of the U.S. Department of Justice.

1 Introduction

An unusual confluence of events has positioned the transportation sector's reliance on gasoline near the forefront of national policy debate. Retail gasoline prices have exhibited increased volatility over the past decade, including a 17 month period in which prices rose from \$2.21 per gallon to \$4.17 per gallon.¹ The foreign policy and environmental externalities associated with crude oil usage have been put in stark relief due to conflicts in the Middle East and extensive debates regarding climate change policy. And the financial bailout of the American automotive industry has raised questions about the management of the "Big Three" manufacturers and the role of new vehicle production in the broader economy.

The policy interest in automobile demand has been matched by a renewed interest among academic economists in understanding how consumers react to gasoline prices.² This research falls broadly into two groups. The first aims to recover consumer valuations of fuel economy (e.g., Goldberg (1998); Bento et al (2009); Gramlich (2010); Allcott and Wozny (2010); Jacobsen (2010); Beresteanu and Li (2011)). These papers estimate random utility models of demand and focus on the covariance between vehicle market shares and gasoline prices, controlling for suggested retail prices and other vehicle characteristics. The second group seeks to understand how gasoline prices affect equilibrium demand outcomes; these papers generally regress measures of fleet fuel efficiency on gasoline prices and controls (e.g., Li, Timmins and von Haefen (2009); Busse, Knittel and Zettelmeyer (2011); Klier and Linn (2010a)).³ A reasonable synthesis of results is that many market-based interventions, such as moderate carbon and gasoline taxes, are unlikely to produce substantial consumer substitution toward fuel efficient vehicles.

We contribute to this literature by focusing more explicitly on the short run supply-side behavior of automobile manufactures. In particular, we examine the empirical relationship between gasoline prices and the cash incentives offered by manufacturers on a week-to-week basis. The strength of this relationship informs consumer substitution patterns. Intuitively, if the cash incentives available on fuel inefficient vehicles rise with gasoline prices then one can infer that manufacturers are acting to mitigate substitution toward fuel efficient vehicles.

¹This is according to weekly data on all grades, all formulations gasoline prices published by the Energy Information Agency of the US Department of Energy for January 29, 2007 and July 7, 2008. This followed nearly 20 years of steady or declining real gasoline prices.

²The subject also attracted substantial attention from economists following the 1970s oil crises (Blomquist and Haessel (1978); Carlson (1978); Dahl (1979); Greenlees (1980); Wheaton (1982); Kahn (1986)).

³This is a loose characterization. Li, Timmins, and von Haefen (2009) and Busse, Knittel and Zettelmeyer (2011) estimate how gasoline prices affect average vehicles sales in various fuel efficiency quantiles. Klier and Linn (2010a) estimate the how fuel costs affect the sales of individual vehicles.

Furthermore, the relationship between cash incentives and gasoline prices has implications for the proper specification of random utility models that aim to recover consumer valuations of fuel economy directly.

We base our analysis on a theoretical model of Nash-Bertrand competition among manufacturers facing linear demand schedules. We solve the manufacturers' first-order conditions and demonstrate that, in equilibrium, gasoline prices affect an automobile's cash incentives through three main channels: their effect on the vehicle's fuel cost, their effect on the fuel costs of the vehicle's competitors, and their effect on the fuel costs of other vehicles produced by the same manufacturer.⁴ Provided that demand is symmetric, or close to symmetric, the first two channels dominate. It follows that cash incentives should increase with gasoline prices for vehicles that are fuel inefficient relative to their closest competitors, but decrease for relatively fuel efficient vehicles. We manipulate these equilibrium relationships to construct a reduced-form regression equation that we take to data.

In the empirical analysis, we examine of a comprehensive set of manufacturer incentive programs offered by General Motors, Ford, Chrysler, and Toyota over the period 2003-2006. We use these data to construct a measure of the cash incentives available to purchasers of each vehicle, in each week and geographic region. We combine information on vehicle miles-per-gallon (MPG) with information on retail gasoline prices to measure fuel costs. We then regress the cash incentives of each vehicle on the fuel costs of the vehicle, the weighted average fuel costs of the vehicles produced by competitors, and the weighted average fuel costs of other vehicles produced by the same manufacturer. Estimation exploits variation in 230,835 vehicle-week-region observations. The reduced-form coefficients of interest are identifiable even in the presence of vehicle, time, and region fixed effects because gasoline prices affect fuel costs differentially across vehicles (the fuel costs of inefficient vehicles are more responsive to the gasoline prices than the fuel costs of efficient vehicles).

We find that, on average, the cash incentives available for purchasers of a given vehicle increase in the vehicle's fuel costs and decrease in the weighted average fuel costs of vehicles produced by competitors. The net effect is negligible for vehicles that provide similar MPG relative to their close competitors, but can be positive or negative for vehicles that are relatively fuel efficient or inefficient. To quantify these differential effects, we calculate the proportion of changes in the relative cumulative gasoline expenditures across vehicles that are offset by cash incentives.⁵ The results correspond to an average manufacturer offset of

⁴By "fuel cost" we mean the gasoline expense of driving.

⁵Suppose that vehicle A gets 20 miles-per-gallon, vehicle B gets 30 miles-per-gallon, and the gasoline price is \$2.00 per gallon. Then, under plausible assumptions on the discount rate and vehicle usage rates,

40 percent, and we interpret this as a lower bound on the amount that consumers value cumulative gasoline expenditures relative to purchase prices.

The estimated short run relationship between gasoline prices and the cash incentives of manufacturers underscores that market-based policy instruments, such as carbon taxes and emissions trading programs (“cap and trade”), likely would yield substantial abatement from the automobile sector. Further, while our analysis focuses on short run pricing decisions, the results suggest that manufacturers also should adjust research and design investment decisions in response to market-based policy instruments, given the clear relationship between markups and vehicle profitability. We calculate that a one dollar increase in gasoline prices would lead the average markup on vehicles in the highest MPG quartile to increase by \$340 relative to the average markups in the lowest MPG quartile. This channel is well understood to exist but previous efforts to quantify its importance have been scarce (an exception is Klier and Linn (2010b)).⁶

The main results also raise the question of whether the discrete choice literature, which typically has not controlled for supply-side pricing patterns, provides consistent estimates of consumer demand for fuel economy. Intuition suggests that bias exists. For instance, our results show that when gasoline prices rise, manufacturers respond with cash incentives that damp consumer substitution toward fuel efficient vehicles, partially compensating consumers for the differential impact of gasoline prices. If cash incentives are unobserved in the data, the damped consumer shift could be mistaken for consumers being unresponsive to gasoline prices. We derive the bias term formally and show that, for the special case of logit demand, the bias term is obtainable from the covariance between fuel costs and cash incentives. Based on the data, our best estimate is that a downward bias of 13.7 percent is present. We suspect that bias would be exacerbated in the more general nested logit case. Although our data are insufficient to support a point estimate, we provide some evidence that suggests a wide range of possible bias with a possible upper bound (on the downward bias) of 80 percent. It follows that, in equilibrium, one should expect market-based policy instruments to yield more abatement from the automobile sector than some models predict.

We also establish the secondary result that, on average, the cash incentives of a given

the difference in expected cumulative gasoline expenditures between the two vehicles is \$3,762. This gap increases to \$4,703 for gasoline prices of \$2.50 per gallon. If the results indicate that a \$0.50 increase in the gasoline price induces the cash incentives of A to increase by \$375 more than those of B, then we calculate the proportion of relative fuel cost changes that are offset by cash incentives as $\$375/(\$4703 - \$3762) = 40\%$.

⁶Whereas other work develops the correlation between vehicle transaction prices and gasoline prices (Busse, Knittel and Zettelmeyer (2011)), to our knowledge our results are the first that explicitly identify manufacturers, rather than automobile dealerships, as the source of vehicle price adjustments.

vehicle correlate more with the fuel costs of similar vehicles than with the fuel costs of dissimilar vehicles – consistent with consumers viewing vehicles of the same type or segment as closer substitutes than vehicles of different type or segment. Thus, the result suggests carbon taxes and emissions trading programs likely would generate more consumer substitution within vehicles types than across vehicles types, e.g., leading consumers to replace relatively fuel inefficient SUVs with fuel efficient SUVs rather than leading consumers to replace SUVs with cars. The result also has implications for the optimal design of revenue-neutral “feebate” programs, under which tax credits are given to purchasers of vehicles with fuel efficiency above some benchmark level while extra fees are charged to purchasers of other vehicles.⁷ In particular, it suggests the possibility that the use of multiple benchmarks (e.g., different benchmarks for cars, SUVs and light duty trucks) could improve effectiveness by generating higher-powered incentives for intra-type substitution.⁸ Finally, the result implies that the immediate impact of electric vehicle subsidies on vehicle choice may be limited because such subsidies are most likely to entice only consumers with strong pre-existing preferences for compact or sub-compact cars.^{9,10}

Finally, our work is largely complementary to Busse, Knittel and Zettelmeyer (2011), which examines a ten percent sample of automobile purchases over 1999-2008 and estimates the mean effect of gasoline prices on the transaction prices of vehicles in each MPG quartile. They find that a one dollar increase in the gasoline price lowers average transaction prices in the lowest MPG quartile by \$250 and raises average transaction prices in the highest MPG quartile by \$104. Our results are similar when comparably aggregated: we find that a one

⁷Denmark, France, the Netherlands and Norway currently use single benchmark feebate programs to encourage purchases of more fuel efficient vehicles. Feebate legislation (AB493) was defeated in the California legislature in January 2008. The optimal design of feebate programs previously has been considered in the academic literature (e.g., see Peters et al (2008)). The use of multiple benchmarks potentially could allay concerns about the fairness of single benchmark feebate programs, related to some consumers being forced to pay additional fees for larger vehicles that they view as necessary for family or work obligations.

⁸Suppose that the maximal fee and rebate are predetermined due to political considerations, and compare a multiple benchmark scheme that interpolates between the maximal fee and the maximal rebate separately for SUVs and cars against a single benchmark scheme that applies the maximal fee to the least efficient SUV, applies the maximal rebate to the most efficient car, and interpolates in between by MPG without consideration of vehicle type. The multiple benchmark scheme likely would encourage greater substitution from inefficient SUVs to efficient SUVs and greater substitution from inefficient cars to efficient cars. While the multiple benchmark scheme could also generate substitution from inefficient cars to efficient SUVs, thereby undermining effectiveness, our results suggest the magnitude of such inter-type substitution likely would be limited.

⁹Electric vehicles have generally not been released as SUVs due to technological barriers.

¹⁰We thank an anonymous referee for pointing out that the secondary result also suggests the within-type and within-segment elasticity estimates of the discrete choice literature may contain more bias than estimates of across-type or across-segment elasticities.

dollar increase in the gasoline price raises average incentives in the lowest MPG quartile by \$248 and lowers cash incentives in highest MPG quartile by \$92. This provides useful corroboration. More generally, the main focus of Busse, Knittel and Zettelmeyer (2011) is on providing a comprehensive analysis of how gasoline prices affect the sales and prices of new and used automobiles of different MPG quartiles. By contrast, we focus exclusively on manufacturer pricing and more fully leverage theory to inform the regression specification.

The paper proceeds as follows. In Section 2, we discuss the data used in the analysis, with a particular focus on the cash incentives, gasoline prices, and vehicle characteristics. We develop the theoretical framework of Bertrand-Nash competition in Section 3. Then, in Section 4, we derive the regression equation, provide a means for interpreting results, and discuss issues related to identification. We present our baseline results together with various sensitivity analyses in Section 5, develop the implications for the existing discrete choice literature in Section 6, and conclude in Section 7.

2 Data

We examine the proprietary data of Autodata Solutions, a marketing research company that maintains a comprehensive list of manufacturer incentive programs. We focus on the national and regional cash incentives offered by General Motors, Ford, Chrysler, and Toyota over the period 2003-2006.¹¹ There are 141,842 incentive-vehicle pairs in the data, each of which provides cash to consumers (“consumer cash”) or dealerships (“dealer cash”) at the time of purchase.¹² Panel A of Table 1 provides summary statistics for these incentives. The mean incentive provides \$1,389 in cash and is offered for 61 days. Just more than half the incentives apply to a single vehicle.

The theoretical framework we introduce provides a reduced-form expression for equilibrium incentive levels, given inter-temporal realizations of supply and demand conditions. Accordingly, we use the data to approximate the cash incentive available to consumers for each vehicle in the data, in each region and week. More than one incentive frequently is

¹¹The German manufacturer Daimler owned Chrysler over this period. We exclude Mercedes-Benz from this analysis since it is traditionally associated with Daimler rather than Chrysler. We consider an incentive to be regional if it is available across an entire Energy Information Agency region. The five EIA regions are East Coast, Gulf Coast, Midwest, Mountain West, and West Coast. See www.eia.doe.org for details.

¹²We focus on cash incentives that are available to the general public. To that end, we exclude incentives that are targeted to specific consumer groups (e.g., the “DaimlerChrysler Farm Bureau Member Certificate”). Employee discounts are excluded, though in 2005 there was a period during which some manufacturers made employee discounts available to non-employees. The inclusion of employee discounts in 2005 does not materially affect the results.

Table 1: Summary Statistics

Panel A: Distribution of Cash-Back Incentives						
Variable	Mean	10%	25%	50%	75%	90%
Cash Amount	1,389	500	500	1,000	2,000	3,000
Duration	61	11	20	40	82	104
# Vehicles	6.5	1	1	2	5	20

Panel B: Distribution of Maximum and Mean Incentive						
Variable	Mean	10%	25%	50%	75%	90%
Maximum Incentive	1,536	0	500	1,000	2,500	3,500
Mean Incentive	917	0	500	750	1,167	1,750

Panel A is based on 141,842 incentive-vehicle pairs over 2003-2006. Cash Amount is in dollars, Duration is in days, and # Vehicles represents the number of vehicles to which the incentive can be applied. Panel B is based on 230,835 vehicle-region-week observations over 2003-2006. Maximum Incentive and Mean Incentive are the maximum and mean cash incentive available for a given vehicle, region, and week, respectively.

available for given vehicle-region-week combinations. This occurs most often when manufacturers pair a broadly applicable incentive (e.g., an incentive for midsize cars) with more specifically targeted incentives. Since consumers likely select among the available incentives, we construct our baseline measure with the maximum incentive. For robustness, we also examine the mean incentive. Panel B of Table 1 provides information on the empirical distributions of the two measures. The maximum incentive has a mean of \$1,536 while the mean incentive has a mean of \$917. Notably, at least one incentive is available in 82.41% of the vehicle-region-week observations.

The second key ingredient to the empirical analysis is the gasoline price. We obtain regional gasoline prices over 2003-2006 from a weekly survey of pump prices conducted by the Energy Information Agency (EIA).¹³ Figure 1 plots gasoline prices over the sample period. That a run-up in gasoline prices occurred over the sample period is well known. The EIA data indicate that national gasoline prices (per gallon) increased from an average of \$1.75 in 2003 to an average of \$2.57 in 2006. The seasonality of the data are also noticeable; prices are higher during summer months and lower during the winter months. We purge the gasoline prices of this seasonality prior to their use in analysis; since manufacturers adjust

¹³The survey methodology is detailed online at the EIA webpage. Pump prices are net of all taxes.




Figure 1: Weekly Pump Prices of Gasoline over 2003-2006.

their prices cyclically over vehicle model-years (e.g., Copeland, Dunn, and Hall (2005)), seasonality in gasoline prices is potentially confounding.¹⁴ The data reveal an upward spike in gasoline prices around September 2005. This is due to the effects of Hurricane Katrina, which temporarily eliminated more than 25 percent of US crude oil production and 10-15 percent of the US refinery capacity (EIA 2006).

Finally, we use certain vehicle characteristics in the analysis. These characteristics are also obtained from Autodata Solutions. To be clear, by “vehicle,” we mean a particular model in a particular model-year. The 2003 Ford Taurus is one vehicle in the data, and we consider it as distinct from the 2004 Ford Taurus. Overall, there are 546 vehicles in the data – including 294 cars, 191 SUVs, and 61 trucks. We observe the manufacturer-suggested retail price (MSRP), miles-per-gallon, horsepower, wheel base,¹⁵ and passenger capacity. We construct a measure of fuel costs by dividing the relevant gasoline price by miles-per-gallon

¹⁴We employ the X-12-ARIMA algorithm, which is also used by the Bureau of Labor Statistics to deseasonalize inputs to the consumer price index. We use data on gasoline prices over 1993-2008 to estimate the seasonal factors, adjusting the regional time-series independently. We specify a multiplicative decomposition, which allows the effect of seasonality to increase with the magnitude of the trend-cycle. The results are robust to log-additive and additive decompositions. For more details on the X-12-ARIMA, see Makridakis, Wheelwright and Hyndman (1998) and Miller and Williams (2004). We refer the reader to the working paper version of this paper (Langer and Miller (2008)) for plots of the deseasonalized gasoline prices. As we discuss below, the regression specification includes a fixed effect for each week in the data, which removes the influence of any seasonality that remains after we apply the X-12-ARIMA.

¹⁵Wheel base is the distance from the center of the front wheel to the center of the rear wheel, and is a standard measure of vehicle size. We report wheel base in inches.

(this ratio is the gasoline expense of a single mile of travel).¹⁶

Table 2 provides summary statistics for these vehicle characteristics, both for the full sample and separately for cars, SUVs, and trucks. The unit of observation in each case is at the vehicle-region-week level. Fuel Cost is the ratio of the gasoline price to miles-per-gallon; its mean of 0.10 indicates that gasoline expenses are roughly 10 cents per mile on average. The means of MSRP, miles-per-gallon, horsepower, wheel base, and passenger capacity are \$29,118, 22.90, 218.39, 111.75, and 5.08, respectively. The subsample statistics are consistent with the generalization that cars are smaller, more fuel efficient, and less powerful than SUVs and trucks. As we discuss below, the regression specification includes vehicle fixed effects to account for vehicle heterogeneity (both observed and unobserved) but the vehicle characteristics nonetheless play an important role.

Table 2: Means of Variable Characteristics

Variables	All Vehicles	Cars	SUVs	Trucks
Fuel Cost	0.10	0.09	0.11	0.11
MSRP	29,118	28,543	32,131	22,331
Miles-Per-Gallon	22.90	25.91	19.42	19.83
Horsepower	218.39	205.43	237.58	218.13
Wheel Base	111.75	107.79	114.29	122.12
Passenger Capacity	5.08	4.84	5.88	3.74

Means are based on vehicle-region-week observations over the period 2003-2006. There are 230,835 observations on 546 vehicles in the full sample. Subsample means are based on 121,860 car observations, 82,600 SUV observations, and 26,375 truck observation, representing 294 cars, 191 SUVs, and 61 trucks, respectively. Fuel Cost is the gasoline price divided by miles-per-gallon. Fuel Cost is in dollars per mile, MSRP is in dollars, and Wheel Base is in inches.

¹⁶When more than one set of attributes exist for a vehicle (e.g., due to option packages or multiple trim levels), we use the attributes corresponding to the lowest MSRP, although the results are fundamentally unchanged when we use the attributes of the vehicle with the MSRP closest to the mean MSRP for the model nameplate in a model year. We impute the period over which each vehicle is available to consumers as beginning with the start date of production, as listed in Ward’s Automotive Yearbook, and ending after the last incentive program for that vehicle expires. When the start date of production is unavailable, we set the start date at August 1 of the previous year. As an example, we would set the start date of the 2006 Civic Hybrid to be August 1, 2005. We impose a maximum period length of 24 months. In robustness checks, we used an 18 month maximum; the different period lengths do not affect the results.

3 Theoretical Framework

We derive our regression equation from a model of Bertrand-Nash competition between automobile manufacturers that face a linear demand schedule. We take as given that there are F automobile manufacturers and J vehicles. Each manufacturer produces some subset \mathbb{J}_f of the vehicles and prices to maximize short run variable profits:

$$\pi_{ft} = \sum_{j \in \mathbb{J}_f} (p_{jt} - c_{jt}) * q(p_{jt}, \mathbf{p}_{-j,t}), \quad (1)$$

where for each vehicle j and period t , the terms p_{jt} , c_{jt} , and $q(p_{jt}, \mathbf{p}_{-j,t})$ are the manufacturer price, the marginal cost, and the quantity sold respectively. We assume constant returns to scale for simplicity, and abstract from the manufacturers' selections of vehicle attributes and fleet composition, which is more important to long run analyses.

We assume that consumer demand depends linearly on manufacturer prices, expected lifetime fuel costs, and certain exogenous demand shifters that include vehicle attributes, maintenance costs, and other factors:

$$q(p_{jt}, \mathbf{p}_{-j,t}) = \sum_{k=1}^J \alpha_{jk} (p_{kt} + x_{kt}) + \mu_{jt}, \quad (2)$$

where the α_{jk} is a demand parameter, x_{kt} captures fuel costs, and μ_{jt} captures the net effect of the demand shifters. We impose the normality conditions that demand is downward sloping ($\alpha_{jj} \leq 0$), vehicles are substitutes ($\alpha_{jk} \geq 0$ for $k \neq j$), and a price increase common to all vehicles lowers demand ($|\alpha_{jj}| \geq \sum_{k \neq j} \alpha_{jk}$ for all j).

The equilibrium manufacturer prices in each period can be characterized by J first-order conditions. We solve these first-order equations to obtain equilibrium manufacturer prices as functions of the exogenous factors.¹⁷ The resulting manufacturer “price rule” is a

¹⁷The solution technique is simple. Turning to vector notation, one can rearrange the first-order conditions such that $Ap = b$, where A is a $J \times J$ matrix of demand parameters, p is a $J \times 1$ vector of manufacturer prices, and b is a $J \times 1$ vector of “solutions” that incorporate the fuel costs, marginal costs, and demand shifters. Provided that the matrix A is nonsingular, Cramer’s Rule applies and there exists a unique Nash equilibrium in which the equilibrium manufacturer prices are linear functions of all the fuel costs, marginal costs, and demand shifters.

linear function of the fuel costs, marginal costs, and demand shifters:

$$p_{jt}^* = \phi_{jt}^1 x_{jt} + \sum_{k \notin \mathbb{J}_f} \phi_{jkt}^2 x_{kt} + \sum_{l \in \mathbb{J}_f, l \neq j} \phi_{jlt}^3 x_{lt} \\ + \phi_{jt}^4 c_{jt} + \phi_{jt}^5 \mu_{jt} + \sum_{k \notin \mathbb{J}_f} (\phi_{jkt}^6 c_{kt} + \phi_{jkt}^7 \mu_{kt}) + \sum_{l \in \mathbb{J}_f, l \neq j} (\phi_{jlt}^8 c_{lt} + \phi_{jlt}^9 \mu_{lt}). \quad (3)$$

The reduced-form coefficients $\phi^1, \phi^2, \dots, \phi^9$ are nonlinear functions of the structural demand parameters. The price rule makes it clear that the equilibrium price of a vehicle depends on its characteristics (i.e, its fuel cost, marginal cost, and demand shifter), the characteristics of vehicles produced by competitors, and the characteristics of other vehicles produced by the same manufacturer. For the time being, we collapse the second line of the price rule into a vehicle-period-specific factor, which we denote γ_{jt} .

Estimation based on equation 3 is infeasible because the J^2 fuel cost coefficients per period cannot be identified with J observations per period. However, the price rule can be manipulated to obtain an expression in weighted averages:

$$p_{jt}^* = \phi_{jt}^1 x_{jt} + \phi_{jt}^2 \sum_{k \notin \mathbb{J}_f} \omega_{jkt}^2 x_{kt} + \phi_{jt}^3 \sum_{l \in \mathbb{J}_f, l \neq j} \omega_{jlt}^3 x_{lt} + \gamma_{jt}. \quad (4)$$

In this reformulation, the equilibrium price of a vehicle depends on the vehicle's fuel cost, the weighted average fuel cost of vehicles produced by competitors, the weighted average fuel cost of vehicles produced by the same manufacturer, and the vehicle-time-specific factor.¹⁸ This reduces dramatically the number of coefficients to be estimated.

The weights in equation 4 are functions of the structural demand parameters or, equivalently, the own-price and cross-price elasticities. Reduced-form analysis can proceed even when these structural parameters are unknown and cannot be estimated reliably, provided that reasonable approximations to the weights can be made. (Of course, if the structural parameters were known then reduced-form analysis would be more difficult to motivate.) Analytical solutions for the weights are obtainable through the theory – though the algebraic burden increases nonlinearly in the number of vehicles. With three vehicles, the weights that vehicles 2 and 3 receive in the determination of vehicle 1's equilibrium price are given

¹⁸The weights are $\omega_{jkt}^i = \phi_{jkt}^i / \phi_{jt}^i$, for $i = 2, 3$, and the coefficient ϕ_{jt}^i is the sum of the ϕ_{jkt}^i coefficients ($\phi_{jt}^i = \sum \phi_{jkt}^i$). Thus, the weights sum to unity for each vehicle-period combination: $\sum_{k \notin \mathbb{J}_f} \omega_{jkt}^2 = \sum_{l \in \mathbb{J}_f, l \neq j} \omega_{jlt}^3 = 1$.

by:

$$\omega_{12}^2 = \frac{A_{12}}{A_{12} + A_{13}} \quad \text{and} \quad \omega_{13}^2 = 1 - \omega_{12}^2, \quad \text{where} \quad A_{12} = \frac{\alpha_{12}}{\alpha_{11}} - \frac{1}{2} \frac{\alpha_{13} \alpha_{32}}{\alpha_{11} \alpha_{33}}. \quad (5)$$

Here, the demand parameters $(\alpha_{11}, \alpha_{22}, \alpha_{12}, \dots)$ are as specified in equation 2.¹⁹ It follows that, in determining the equilibrium price of a vehicle, the fuel costs of more readily substitutable vehicles receive greater weight. To see this, note that the ratio α_{jk}/α_{jj} is a diversion ratio and can be interpreted as the proportion of consumers purchasing vehicle j that considers vehicle k as the next best option.²⁰ As shown, the weight that vehicle k receives in the determination of price j increases in the diversion ratio between the two vehicles and decreases in the diversion ratios between vehicle j and other vehicles.

4 The Empirical Model

4.1 Regression equation

The theoretical framework developed above motivates the regression equation that we take to the data:

$$INC_{jtr} = \beta_1 \frac{gp_{tr}}{mpg_j} + \beta_2 \sum_{k \notin \mathbb{J}_j} \tilde{\omega}_{jkt}^2 \frac{gp_{tr}}{mpg_k} + \beta_3 \sum_{l \in \mathbb{J}_j, l \neq j} \tilde{\omega}_{jlt}^3 \frac{gp_{tr}}{mpg_l} + \gamma_{jtr}^*, \quad (6)$$

in which the composite error term γ_{jtr}^* is specified as follows:

$$\gamma_{jtr}^* = \mathbf{z}'_{jt} \boldsymbol{\theta} + \kappa_j + \delta_t + \eta_r + \epsilon_{jt}. \quad (7)$$

The dependent variable, INC_{jtr} , is the maximum cash incentive available for vehicle j in week t and region r . The main independent variables are own fuel costs (i.e., the ratio of gasoline price to miles-per-gallon), the weighted average fuel costs of vehicles produced by competitors, and the weighted average fuel costs of vehicles produced by the same manufacturer. The empirical weights, $\tilde{\omega}_{jkt}^2$ and $\tilde{\omega}_{jkt}^3$, play a crucial role in the construction of the latter two variables, and we discuss the weights in detail shortly. The composite error term, which accounts for demand and cost shifters, includes a third-order polynomial in the number of weeks the vehicle has been on the market and analogous third-order polynomials

¹⁹We derive this result in the working paper.

²⁰Diversion ratios are used frequently in antitrust analysis to measure product substitutability because they can be more easily discerned from data than own-price and cross-price elasticities.

for vehicles produced by competitors and other vehicles produced the same manufacturer.²¹ The composite error term also includes vehicle, week, and region fixed effects.

Although the regression equation is tightly linked with equation 4 from the theoretical framework, some differences exist. For instance, the dependent variable is based on cash incentives rather than vehicle prices. This switches the signs of the coefficients but does not have broader implications as the vehicle fixed effects absorb the constant portion of vehicle prices. Also, we use the ratio of gasoline price to MPG (gasoline expenditure per mile) as an empirical proxy of expected cumulative fuel costs, a nearly ubiquitous practice in the empirical literature (e.g., Goldberg (1998); Bento et al (2009); Jacobsen (2010); Gramlich (2010); Li, Timmins and von Haefen (2009); Sallee, West and Fan (2009); Beresteanu and Li (2011)). The empirical proxy should be accurate if automobile consumers treat the current gasoline price as a forecast of future prices. There is some evidence that this is the case: Anderson, Kellogg and Sallee (2011) examine survey data on individuals' gasoline price forecasts over 1993-2008 and determine that the average individual's forecast is statistically indistinguishable from a "no change" forecast.^{22,23}

We estimate the regression equation with ordinary least squares and cluster the standard errors at the vehicle level to account for autocorrelation and other potential correlations in the residuals.²⁴ The theory suggests that a vehicles' incentives should increase with its fuel costs and decrease with the fuel costs of vehicles produced by competitors. The fuel costs of other vehicles produced by the same manufacturer have no effect if demand is symmetric;

²¹Copeland, Dunn and Hall (2005) document that vehicle prices fall approximately nine percent over the course of the model-year.

²²Current prices do not always reflect expectations. One motivating example is Hurricane Katrina, which temporarily stymied crude oil production and refinement, and created a wedge between oil prices and (distant) futures prices. Our regression results are robust to the exclusion of observations from August, September, and October 2005 from the data sample.

²³The academic literature has sought to determine whether retail gasoline prices and crude oil prices actually follow a random walk, without clear resolution (e.g., see Davis and Hamilton (2004); Geman (2007); Hamilton (2009); Kilian (2009)). What is clear is that price changes are difficult to predict and that "no change" forecasts perform well relative to forecasts based on futures prices and forecasts based on simple econometric models (e.g., Alquist and Kilian (2010)). In our data, an OLS regression of fuel costs on lagged fuel costs, eight lags of fuel cost changes, and the control variables shown in equation 7 yields a coefficient on lagged fuel costs of 0.5719 (standard error of 0.0014). Given the critical values reported in Hamilton (1994, chapter 17), this would seem to reject the hypothesis that fuel costs follow a random walk in our data. We are wary of interpreting this too result too strongly, however, since our data cover a short sample period relative to the data examined elsewhere.

²⁴We have experimented with Tobit regressions that account for the fact that cash incentives are censored at zero (i.e., incentives are never negative). Maximum likelihood routines have weaker small sample properties, however, and the bevy of vehicle, week, and region fixed effects in the specification leads to multi-collinearity problems in estimation. We are skeptical that censoring is problematic because we observe positive incentives in more than 80 percent of the observations.

otherwise the implications of these fuel costs are theoretically ambiguous.²⁵ Formally, the theory provides the following three hypotheses: $\beta_1 \geq 0$, $\beta_2 \leq 0$, and $\beta_1 \geq |\beta_2|$.

4.2 Quantifying the impact of gasoline prices

Of particular interest is the proportion of fuel cost changes that are offset by cash incentives. Given regression results, the difference in responsiveness between two vehicles, j and i , can be calculated as follows:

$$\begin{aligned} \frac{\partial(\widehat{INC}_j - \widehat{INC}_i)}{\partial gp} &= \widehat{\beta}_1 \left(\frac{1}{mpg_j} - \frac{1}{mpg_i} \right) + \widehat{\beta}_2 \left(\sum_{k \notin \mathbb{J}_j} \widetilde{\omega}_{jk}^2 \frac{1}{mpg_k} - \sum_{k \notin \mathbb{J}_i} \widetilde{\omega}_{ik}^2 \frac{1}{mpg_k} \right) \\ &\quad + \widehat{\beta}_3 \left(\sum_{l \in \mathbb{J}_j, l \neq j} \widetilde{\omega}_{jl}^3 \frac{1}{mpg_l} - \sum_{l \in \mathbb{J}_i, l \neq i} \widetilde{\omega}_{il}^3 \frac{1}{mpg_l} \right), \end{aligned} \quad (8)$$

where we have suppressed the week and region subscripts for simplicity. By focusing on differences, we isolate the fuel cost channels through which gasoline prices affect cash incentives. Gasoline prices fluctuations could also affect cash incentives due to changes in real consumer income, production costs, or the desirability of used vehicles. These other effects are controlled for but not estimated directly in our regression model, and they cancel when the incentive derivatives are expressed in differences.²⁶

We calibrate these differences against the differential impacts that gasoline prices have on the cumulative fuel costs that consumers expect to incur over their vehicles' lifetimes:

$$OFFSET_{ji} \equiv \frac{\partial(\widehat{INC}_j - \widehat{INC}_i)}{\partial gp} \bigg/ \frac{\partial(\widetilde{x}_j - \widetilde{x}_i)}{\partial gp}, \quad (9)$$

where \widetilde{x}_j is a measure of cumulative fuel costs that we define as follows:

$$\widetilde{x}_{jt} = \sum_{y=1}^Y \left[\left(\frac{1}{1+r} \right)^{y-1} * MPY * \frac{gp_t}{mpg_j} \right]. \quad (10)$$

where Y is vehicle lifespan, r is the consumer discount rate, and MPY is the miles per year that vehicles are driven. Following Greene (2010) and statistics calculated by the National

²⁵We derive these relationships for the case of $J = 3$ in the working paper under mild regularity conditions.

²⁶These effects are present if $\partial\gamma_t/\partial gp_t \neq 0$. Gicheva, Hastings, and Villas-Boas (2010) tests for income effects using scanner data on grocery purchases in California over 2000-2005, and finds that a 100% increase in gasoline prices leads to a 5-11% decrease in the net price paid per grocery item.

Traffic Safety Administration (NHTSA 2006), we assume a vehicle lifespan of 14 years, that cars are driven 12,061 miles per year, and that SUVs and Trucks are driven 13,436 miles per year. We also assume a consumer discount rate of seven percent. Since the metric of interest, $OFFSET_{ji}$, depends on these assumptions, we conduct sensitivity analysis using discount rates of five and ten percent and vehicle lifetimes of ten and 18 years.

The ratio derived in equation 9 can be interpreted as the proportion of relative fuel cost changes that manufacturers offset with cash incentives. A value of one would indicate that manufacturers fully compensate consumers for changes in the relative fuel costs of vehicles j and i , while a value of zero would indicate that manufacturers are not responsive to the relative fuel costs of the two vehicles. To build intuition, consider two hypothetical cars produced by different manufacturers. Car A gets 20 miles-per-gallon and car B gets 30 miles-per-gallon. With a gasoline price of \$2.00 per gallon, the difference in expected cumulative gasoline expenditures is \$3,762.²⁷ This gap increases to \$4,703 for gasoline prices of \$2.50 per gallon. Thus, if the regression results indicate that a \$0.50 increase in gasoline prices induces the cash incentives of A to increase by \$375 more than those of B, we would calculate the proportion of relative fuel cost changes that are offset by cash incentives (the “offset percentage”) as $\$375/(\$4703 - \$3762) = 40\%$.

4.3 Empirical weights

We approximate the weights using data on vehicle attributes. Our assumption is that the degree of substitutability between vehicles decreases in the Euclidean distance between their attributes. Or, stated more simply, that consumers tend to substitute among vehicles that have similar characteristics. In industrial organization, the linking of product characteristics to consumer substitution dates to Lancaster (1966), and seminal contributions use vehicle characteristics to estimate demand elasticities in the automobile industry (e.g., Berry, Levinsohn and Pakes (1995, 2004), Petrin (2002)). The critical distinction is that we make assumptions regarding the relative importance of the vehicle characteristics, whereas more structural approaches estimate the relative importance based on the data.

In our application, we treat each of the available vehicle characteristics – MSRP, miles per gallon, horsepower, passenger capacity, and wheelbase – equally in the construction of the empirical weights.²⁸ Formally, we take M vehicle attributes, which we denote

²⁷Expected cumulative gasoline expenditures are \$11,286 and \$7,524 for the two vehicles, respectively.

²⁸We also include 13 indicator variables for the segment of the vehicle. The car segments are subcompact, compact, intermediate, luxury, sport, luxury high, and luxury sport. The SUV segments are compact, intermediate, large, and luxury. The truck segments are small pick-up and large pick-up.

z_{jm} for $m = 1, \dots, M$, and standardize each to have a variance of one. We then sum the squared differences between each attribute to calculate the effective “distance” in attribute space. We form initial weights as follows:

$$\omega_{jk}^* = \frac{1}{\sum_{m=1}^M (z_{jm} - z_{km})^2}.$$

To finish, we set the initial weights to zero for vehicles of different types (i.e., cars, SUVs, and trucks) and normalize such that the weights sum to one for each vehicle-period. We perform this weighting procedure separately for vehicles produced by the same manufacturer and vehicles produced by competitors to obtain $\tilde{\omega}_{jkt}^2$ and $\tilde{\omega}_{jkt}^3$, respectively. Thus, the weighting scheme is based on the inverse Euclidean distance between vehicle attributes among vehicles of the same type.²⁹

In Table 3, we provide a matrix of competitor weights for eight selected 2005 model-year vehicles – four large pickup trucks and four small pickup trucks. The elements in each row are the weights used to predict the cash incentives for the vehicle listed at the left of the row.³⁰ The weights are for the week of January 3, 2005. As shown, vehicles of the same segment typically have weights that are roughly an order of magnitude larger than vehicles of different segments. To model the incentives on the Silverado, a large pick-up truck, we place weights of 0.0938, 0.1110, and 0.0545 on the F-150, the Ram, and the Tundra (all large pickups) and weights of 0.0033, 0.0203, 0.0009 on the Ranger, the Dakota, and the Tacoma (all small pickups). There is substantial variation in the weights that vehicles within the same segment receive. The Colorado and the Tacoma appear as particularly close competitors due to similarity in their attributes: the GM Colorado has 24.3 MPG, 175 horsepower, 111” wheelbase, passenger, and an MSRP of \$15,095, while the Toyota Tacoma has 24.3 MPG, 164 horsepower, 109” wheelbase, and an MSRP of \$13,415. Neither is as close to the Dakota, another small pickup truck, because the Dakota has 19.3 MPG, 210 horsepower, 131” wheelbase, and an MSRP of \$19,885.³¹

²⁹Although the initial weights are constant across time for any vehicle pair, the final weights vary due to changes in the set of vehicles available on the market.

³⁰Three properties of the matrix are readily apparent: First, the matrix has a block diagonal structure because vehicles produced by the same manufacturer receive a competitor weight of zero. Second, the matrix is asymmetric because the weighting scheme does not impose symmetry. Finally, the weights do not sum to unity across rows because the vehicles compete with four other 2005 model-year trucks, as well as with vehicles from the 2004 model-year. The omitted 2005 model-year trucks include the GM Canyon, the GM Sierra, the GM Avalanche, and the Ford F-150 Supercrew.

³¹All of the pickup trucks shown have a passenger capacity of three.

Table 3: Matrix of Competitor Weights for Selected Model-Year 2005 Pickup Trucks

Vehicle	Segment	Silverado	Colorado	F-150	Ranger	Ram 1500	Dakota	Tundra	Tacoma
GM Silverado	Large	0	0	0.0938	0.0033	0.1110	0.0203	0.0545	0.0009
GM Colorado	Small	0	0	0.0026	0.1675	0.0009	0.0767	0.0544	0.2407
Ford F-150	Large	0.0370	0.0022	0	0	0.0431	0.0216	0.1455	0.0008
Ford Ranger	Small	0.0013	0.1169	0	0	0.0008	0.0588	0.0035	0.0527
Chrysler Ram	Large	0.1642	0.0024	0.1302	0.0031	0	0	0.0872	0.0009
Chrysler Dakota	Small	0.0014	0.0093	0.0031	0.0102	0	0	0.0051	0.0029
Toyota Tundra	Large	0.0163	0.0021	0.2626	0.0027	0.0176	0.0217	0	0
Toyota Tacoma	Small	0.0014	0.6412	0.0024	0.2010	0.0009	0.0636	0	0

The appropriateness of treating each vehicle characteristic as an equal driver of consumer behavior is not clear *a priori* and, furthermore, weights based on observed characteristics likely understate the competitive influence of vehicles with popular unobserved characteristics. We construct a number of alternative weighting schemes to assess the sensitivity of the regression results. First, we construct weights that exclude each of the vehicle characteristics in turn. Second, we examine equal weights across all vehicles of the same segment (i.e., compact car or luxury SUV), equal across vehicles of the same type (i.e., cars, SUVs, trucks) and equal weights across all vehicles. We discuss the results of these robustness checks in Section 5.2.

4.4 Identification

We estimate the average responsiveness of vehicles' cash incentives to their fuel costs, the fuel costs of vehicles produced by competitors, and the fuel costs of other vehicles produced by the same manufacturer.³² The fuel cost coefficients are identifiable even in the presence of time, vehicle, and region fixed effects because changes in the gasoline price over time (and across regions) affect the fuel costs of vehicles differentially. That is, identification rests on the observation that the fuel costs of fuel efficient vehicles are less responsive to changes in the gasoline price than the fuel costs of fuel inefficient vehicles.

It follows that the empirical weights are central to identification – the weights determine how the fuel cost regressors incorporate heterogeneity in fuel efficiency. To build intuition, suppose that there are three vehicles produced by a different manufacturers. Vehicles A and B are identical compact cars while vehicle C is a luxury car. If, in the determination of A's cash incentive, the fuel costs of B receive a weight of one and the fuel costs of C receive a weight of zero, then the fuel cost of A is collinear with the average fuel costs of A's competitors and the fuel cost coefficients are not be separably identifiable. However, if the fuel costs of B receives slightly less weight than one, with C receiving the remaining weight, then the fuel cost of A differs from the average fuel costs of A's competitors and the fuel cost coefficients are separably identifiable.³³ We have established that the optimal weighting scheme weights B and C according to their competitive significance.

³²Heterogeneity likely exists in these effects across vehicles and time periods. This is evident, for example, in the vehicle-time-specific coefficients of equation 4, which are combinations of the underlying structural demand parameters. We use subsample regressions to capture some of this heterogeneity.

³³This identification strategy requires rich variation in the data – otherwise the high degree of collinearity between fuel costs and competitor fuel costs could hinder estimation. As we develop in Section 5, the data we examine do indeed provide sufficient variation, as evidenced by the precision of the point estimates and the robustness of the results across different specifications and samples.

Ordinary least squares regression based on equation 6 generates unbiased estimates provided that the regressors are uncorrelated with the vehicle-period-region specific residual (which captures deviations in demand and production costs). This condition is reasonable given the set of fixed effects included in the specification. Consider the potential feedback between automobile demand and gasoline prices. The strength of demand likely has a small effect on the global price of oil but the time fixed effects account for the overall effect, so only changes in the distribution of demand (e.g., greater demand for efficient vehicles) could create bias. Fuel costs are the most obvious source of such relative demand changes, by incorporating fuel costs as regressors, they are removed from the residual. Analogously, manufacturers likely adjust vehicle characteristics with the gasoline price, but the inclusion of vehicle fixed effects restricts identification to changes in the gasoline price that occur within the model-year; and characteristics are fixed within the model-year.

5 Empirical Results

5.1 Main regression results

We present the main regression results in column 1 of Table 4. The table also shows results when week fixed effects or vehicle fixed effects are excluded (columns 2 and 3), when regional variation in cash incentives and gasolines prices is discarded (column 4), and when the dependent variable is constructed as the mean cash incentive rather than the maximum cash incentive (column 5). In each case, we run OLS and cluster the standard errors at the vehicle level to account for heteroscedasticity, autocorrelation, and any other correlations among the residuals of each vehicle.

We discuss the main results first. The own fuel cost coefficient of \$44,535 is positive, as predicted by theory, and statistically significant. This captures the intuition that manufacturers partially compensate consumers for higher gasoline expenditures. Considered in isolation, this coefficient would indicate that a \$1.00 increase in the gasoline price would increase cash incentives by \$4,454 for a vehicle with fuel costs of \$0.10 per mile. But the coefficient should not be considered in isolation. As shown, the competitor fuel cost coefficient of $-\$43,318$ is negative, also as predicted by theory, and precisely estimated. This indicates that increases in competitors' fuel costs motivate manufacturers to reduce cash incentives. The net effect of these two channels depends on the fuel efficiency of a vehicle relative to its rivals; for a vehicle with fuel costs of \$0.10 per mile and average competitor fuel costs of \$0.10 per mile, the coefficients imply that the cash incentive would increase only \$122 due to

Table 4: Cash Incentives and Fuel Costs

Variables	Maximum Incentive + Full Sample			National	Mean
	(1)	(2)	(3)	Sample	Incentive
	(4)	(5)			
<i>Selected Coefficients and Standard Errors</i>					
Own fuel cost	44,535*** (12,475)	37,924*** (12,745)	10,810** (5,132)	50,841*** (14,034)	27,238*** (8,638)
Competitor fuel cost	-43,318*** (14,100)	-41,422*** (13,162)	-12,196** (6,199)	-41,630** (18,494)	-26,731*** (9,734)
Same-firm fuel cost	-516 (2,867)	-255 (3,046)	14,042*** (4,511)	-1,171 (2,943)	-998 (2,245)
<i>Specification of Fixed Effects</i>					
Week	yes	no	yes	yes	yes
Vehicle	yes	yes	no	yes	yes
Region	yes	yes	yes	no	yes
<i>Mean Offset</i>					
	40%	34%	13%	45%	24%
R^2	0.6200	0.6064	0.1202	0.6496	0.5624

Results from OLS regressions. The dependent variable in columns 1-4 is the size of the maximum cash incentive and the dependent variable in column 5 is the size of the mean incentive. There are 230,835 observations at the vehicle-week-region in columns 1-3 and column 5, and 46,167 observations at the vehicle-week level in column 4. All regressions include third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

a \$1.00 increase in the gasoline price.³⁴ Although this net effect is positive, as predicted by theory, it is not statistically significant. Finally, the same-firm fuel cost coefficient is small and not statistically significant, consistent with roughly symmetric demand.

Of particular interest is the proportion of relative fuel cost changes that are offset by cash incentives. We calculate this for each vehicle-pair in the data. To create a single summary statistic, we first calculate the weighted average offset between vehicle j and all other vehicles produced by competitors, using the empirical weights to focus more on vehicles with a high degree of substitutability:

$$\overline{OFFSET}_j = \sum_{i=1}^J \tilde{\omega}_{ji}^2 OFFSET_{ji}, \quad (11)$$

where $OFFSET_{ji}$ is defined in equation 9. We then take the mean across vehicles to form the “mean offset” among vehicles produced by competing manufacturers:

$$\overline{\overline{OFFSET}} = \frac{1}{J} \sum_{j=1}^J \overline{OFFSET}_j. \quad (12)$$

This statistic measures the proportion of fuel costs changes that are offset by cash incentives. An offset of one would indicate that manufacturers fully compensate consumers for changes in fuel costs, on average, while an offset of zero would indicate that manufacturers are not responsive to fuel cost changes.³⁵

The main results generate a mean offset of 40 percent. The bootstrapped standard error is 11.4 percent, based on 10,000 simulation draws, and the corresponding 95 percent confidence interval is 21.2% to 58.6%.³⁶ In this offset calculation, we assume a discount rate of seven percent and an expected vehicle lifespan of 14 years. Table 5 provides sensitivity checks for discount rates of five, seven, and ten percent and an expected lifespan of ten, 14, and 18 years. As shown, the mean offset varies from 31 percent to 56 percent.

An alternative metric is the net effect of gasoline prices on cash incentives that accrues

³⁴The net effect for such a vehicle is simply $\hat{\beta}_1 - \hat{\beta}_2 = 0.10 * (\$44,535 - \$43,318) = \122 .

³⁵We calculate the offset percentage using vehicles in the data for the week of December 25, 2006. This offset calculation explicitly takes into account the fact that manufacturers may offset changes in relative operating costs differently for vehicles with greater consumer substitutability than for vehicles that are less close substitutes. In fact, we find that manufacturers offset a greater percentage of relative operating cost for more substitutable vehicles than they do for less substitutable vehicles, as might be expected in the theoretical model.

³⁶The standard error may be a lower bound because vehicle lifespan and usage rates are treated as deterministic.

Table 5: Sensitivity Analysis of Mean Offsets

Vehicle Life	Discount Rate		
	5%	7%	10%
10 years	46%	50%	56%
14 years	36%	40%	47%
18 years	31%	35%	42%

Based on the regression coefficients that appear in column 1 of Table 4.

through the fuel cost variables.³⁷ We calculate the net effect of a one dollar increase in the price of gasoline for each vehicle-week-region observation in the data using the regression coefficients from the baseline specification (column 1 of Table 4). We then aggregate the predictions to construct the mean net effect of each MPG quartile per region-week. We find that a one dollar gasoline price increases the mean incentive of the least efficient quartile by \$248. The mean incentives of the second and third least efficient quartile increase by \$126 and \$13, respectively, and the mean incentive in the most efficient quartile decreases by \$92. This is consistent with the intuition that adverse gasoline price shocks reduce demand for fuel inefficient vehicles and raise demand for fuel efficient vehicles. Comparing across quartiles, the markup on vehicles in the most efficient quartile increase by \$340 relative to the markup on vehicles in the least efficient quartile.

These statistics have the added benefit of being directly comparable to Busse, Knittel, and Zettelmeyer (2011), which examines a ten percent sample of automobile purchases over 1999-2008 and estimates the conditional mean effect of gasoline prices on the transaction prices of vehicles in each MPG quartile. They find that a one dollar increase in the gasoline price lowers average transaction prices by \$250 in the least efficient quartile, \$96 in the second least efficient quartile and \$11 in the third least efficient quartile, but increases average transaction prices by \$104 in the most efficient quartile. These results are similar to our own, both in terms of sign and magnitude, and we interpret them as a useful corroboration.³⁸

We now return to Table 4 and explore the implications of some basic specification and

³⁷As we discuss in Section 4.2, the full net effect is not identifiable given our estimation strategy because the time fixed effects absorb any variation due to income effects, production cost effects, or used vehicles.

³⁸The cash incentives we examine tend to be somewhat sticky, in that there is a tendency for the incentives of given vehicles to be constant over several weeks and then jump, as manufacturers incur menu and advertising costs. The similarity between our results and those of Busse, Knittel, and Zettelmeyer (2011) helps rule out serial correlation as a major source of inconsistency in estimation.

sample choices. As shown in column 2, the fuel cost coefficients are not materially different than the baseline results when week fixed effects are omitted from the specification. By contrast, when vehicle fixed effects are excluded, the own and competitor fuel cost coefficients are much smaller, the coefficient on same-firm fuel costs enters meaningfully, and the mean offset drops to 13 percent. Of course, vehicle characteristics are important determinants of demand and production costs and the exclusion of vehicle fixed effects could lead to bias. The “national sample,” which uses national gasoline prices and nationally-available cash incentives, produces fuel costs coefficients that are similar to the baseline coefficients and a mean offset of 45 percent. Finally, when the dependent variable is constructed as the mean incentive, the fuel cost coefficients are somewhat smaller and the mean offset is 24 percent.³⁹

5.2 Alternative Empirical Weights

The empirical weights that we employ in the baseline results follow from the assumption that the degree of substitutability between vehicles can be approximated by evaluating the similarity of the vehicles’ attributes. We now examine how the results change under alternative weighting schemes, namely equal weights across all vehicles of the same segment (i.e., compact car or luxury SUV), equal across vehicles of the same type (i.e., cars, SUVs, trucks) and equal weights across all vehicles.⁴⁰

Table 6 presents the results. Columns 1-3 show the results obtained from each alternative weighting scheme, in turn, and columns 4-6 show the results obtained from horse-races between the baseline weights and each of the alternative weighting schemes. As shown, when weights are equal among all vehicles of the same segment (column 1) the fuel cost coefficients are similar to those obtained from the baseline weights. The fuel cost coefficients are somewhat smaller when weights are equal among all vehicles of the same type (column 2) but the coefficients remain statistically significant. The mean offset is 43 percent and 24 percent in these two columns, respectively. By contrast, when weights are equal among all vehicles regardless of segment or type, the competitor fuel cost coefficient is close to zero and not statistically significant. The implied mean offset is 13 percent.

³⁹We view the maximum incentive as the more appropriate dependent variable because consumers typically select among the available incentives (when multiple incentives are available). If the maximum incentive is indeed the object of interest then one would expect mean incentives be less responsive to fuel costs.

⁴⁰We also construct a series of weights, following the procedure outlined in Section 4.3, which exclude each of the observed vehicle characteristics in turn. The resulting fuel cost and offset percentage generally are quite similar to those of the baseline results. The exception is wheelbase – when it is excluded from the weights the fuel costs coefficients are smaller and the implied offset percentage falls to 18 percent. Wheelbase is a standard measure of vehicle size, an important determinant of consumer choice.

These patterns are precisely what one should expect, provided that competition between vehicles is indeed localized in attribute space, because the inclusion (or over-weighting) of distant competitors introduces measurement error that biases regression coefficients toward zero. As an example, consider the cash incentives of a Toyota Prius. If competition is localized then potential consumers of the Prius are selecting among relatively fuel efficient vehicles. Thus, Toyota should adjust its Prius incentives with the fuel costs of efficient vehicles (e.g., the Ford Focus) but not the fuel costs of inefficient vehicles (e.g., the Hummer). The inclusion of inefficient vehicles would then create measurement error and the estimated coefficients would be too small in magnitude.⁴¹ By contrast, weighting efficient vehicles more heavily would reduce measurement error and produce more accurate estimates.

To inform whether competition is indeed localized in attribute space, we conduct horse races between the baseline weights and the alternative weighting schemes. The results are shown in columns 4-6 of Table 6. Column 4 includes two sets of competitor and same-firm fuel cost variables, constructed respectively with the baseline weights and equal weights among vehicles of the same segment. As shown, the own fuel cost coefficient is similar to that of the baseline regression (Table 4, column 1). Of more interest are the two competitor fuel costs coefficients. Since each is about half of what is estimated in the baseline regression, the combined effect is similar in magnitude. The two coefficients are jointly statistically significant at the one percent level though neither is significant alone. In columns 5 and 6, the competitor fuel cost variables constructed with the baseline weights strictly dominate the variables constructed with equal weights among vehicles of the same type and equal weights among all firms, respectively. In both cases, the net effect of competitor fuel costs is similar to that of the baseline regression. We interpret these results as evidence that more localized weighting schemes (e.g., the baseline weights and equal weights within segment) have more explanatory power than more global weighting schemes, and that the substitutability of vehicles increases in the similarity of attributes.

⁴¹The econometric intuition is standard: since variation in the Hummer's fuel costs exists but does not correlate strongly to Prius incentives, weighting the Hummer heavily would lead to the inference that Prius incentives are unresponsiveness to competitor fuel costs.

Table 6: Regression Results with Alternative Weighting Schemes

Variables	(1)	(2)	(3)	(4)	(5)	(6)
Own fuel cost	39,346*** (10,277)	28,015*** (7,020)	18,217*** (5,504)	48,504*** (12,261)	45,922*** (12,778)	45,377*** (12,358)
Competitor fuel cost baseline weights				-20,560 (19,249)	-36,248* (20,062)	-43,896*** (16,620)
Same-firm fuel cost baseline weights				-642 (2,855)	-172 (2,864)	86 (2,856)
Competitor fuel cost, equal weights in segment	-29,701*** (11,374)			-21,415 (14,689)		
Same-firm fuel cost, equal weights in segment	7,442 (8,122)			-5,644 (8,205)		
Competitor fuel cost, equal weights in type		-20,873** (8,848)			-4,886 (12,273)	
Same-firm fuel cost, equal weights in type		-8,054 (8,540)			-6,271 (8,418)	
Competitor fuel cost, equal weights			-284 (13,756)			13,711 (15,285)
Same-firm fuel cost, equal weights			-17,026 (12,978)			-15,170 (12,749)
<i>Mean Offset</i>	43%	24%	13%	41%	40%	49%
R^2	0.6204	0.6194	0.6189	0.6206	0.6201	0.6203

Results from OLS regressions. The dependent variable is the size of the maximum cash incentive and the sample includes 230,835 vehicle-week-region observations. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

5.3 Additional Regression Results

First, we explore heterogeneity in the responsiveness of cash incentives to the fuel cost variables using sub-sample regressions for cars, SUVs, and trucks.⁴² Table 7 shows the results. For cars, the own fuel cost coefficient is substantially larger than the coefficient obtained from full sample (see column 1 of Table 4), while the competitor and same-firm fuel cost coefficients are similar in magnitude. Together, these coefficients imply a mean offset of 61 percent. For SUVs, the fuel cost coefficients are similar in magnitude to those obtained from the full sample and the mean offset of 30 percent is slightly smaller. Finally, for trucks, the fuel costs coefficients roughly halve in magnitude relative to the full sample, statistical significance is not maintained, and the mean offset is only 18 percent. Thus, the results indicate that the cash incentives of cars appear to be more responsive to fuel costs than those of SUVs, which appear to be more responsive than those of trucks. Our estimation approach does not provide a clean explanation for this pattern, but we speculate that it could be due to differences in the intensity of competition (e.g., the car industry could be more densely populated in characteristic-space) or differences in preferences among consumers of the vehicle types (e.g., car buyers could be more sensitive to fuel expenditures).

Second, we explore the timing implied by the baseline regression specification, which implicitly imbeds the assumptions that consumers use current gasoline prices to forecast future prices and that cash incentives adjust immediately with current gasoline prices. Column 1 of Table 8 provides results from an alternative specification in which cash incentives are regressed on fuel cost variables constructed as averages over the previous four weeks. As shown, the own and competitor fuel cost coefficients are slightly larger than those produced by the baseline specification, and the mean offset rises to 55 percent. In column 2, we pair the “current” fuel cost variables with the “lagged” fuel cost variables. The own fuel cost coefficients are each roughly half the size of the fuel cost coefficient of column 1, so the combined effect is similar, and the same is true for the competitor fuel cost coefficients. The results are suggestive that consumers construct forecasts using recent gasoline prices and that manufacturers respond with some delay to gasoline price fluctuations. The larger offset percentages indicate that our baseline results may be conservative.

⁴²Heterogeneity in responsiveness is suggested by the theoretical model. For instance, consider the vehicle-specific coefficients of equation 4, each of which is a combination of the underlying structural demand parameters. We cannot fully estimate these heterogeneous effects because the $3J$ coefficients per region-week are not identifiable with J observations per region-week, and our baseline regressions estimate the average responsiveness of cash incentives to the fuel cost variables.

Table 7: Regression Results for Vehicle Type Subsamples

Variables	Cars	SUVs	Trucks
Own fuel cost	62,738*** (21,185)	43,036** (19,012)	25,588 (21,399)
Competitor fuel cost	-44,781* (24,323)	-48,471** (21,860)	-19,464 (22,170)
Same-firm fuel cost	-5,840 (6,092)	5,402 (5,547)	-3,154 (2,216)
<i>Mean Offset</i>	61%	30%	18%
R^2	0.5928	0.6495	0.6593
Observations	121,860	82,600	26,375

Results from OLS regressions. The dependent variable is the size of the maximum cash incentive and the units of observation are at the vehicle-week-region level. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Table 8: Regression Results with Lagged Fuel Costs

Variables	(1)	(2)
Fuel cost		22.33** (10.67)
Competitor fuel cost		-18.12 (12.27)
Same-firm fuel cost		0.75 (1.99)
Lagged fuel cost	47.62*** (13.45)	25.70** (11.97)
Lagged competitor fuel cost	-50.05*** (15.10)	-30.34** (12.48)
Lagged same-firm fuel cost	-0.67 (3.17)	-1.49 (3.22)
<i>Mean Offset</i>	54%	55%
R^2	0.6201	0.6214

Results from OLS regressions. The dependent variable is the size of the maximum cash incentive (in thousands). There are 230,835 observations, representing 546 vehicles, at the vehicle-week-region level. Lagged variables are constructed as the mean over the previous four weeks. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

6 Implications for Discrete Choice Estimation

Our results indicate that manufacturers adjust their cash incentives in response to changes in the fuel costs of their vehicles and the fuel costs of vehicles produced by their competitors. This raises the question of whether the discrete choice literature, which typically does not control for these supply-side responses, provides consistent estimates of consumer demand for fuel economy. Intuition suggests that bias exists. For instance, our results show that when gasoline prices rise, manufacturers respond with cash incentives that damp consumer substitution toward fuel efficient vehicles, partially compensating consumers for the differential impact of gasoline prices. If cash incentives are unobserved in the data, the damped consumer shift could be mistaken for consumers being unresponsive to gasoline prices.

In this section, we formalize this logic and approximate the magnitude of bias. The extant literature largely relies on random utility models such as the nested logit model (e.g., Goldberg (1998); Gramlich (2010); Allcott and Wozny (2010)) and the random coefficients logit model (e.g., Bento et al (2009), Jacobsen (2010), Beresteanu and Li (2011)). We focus on the nested logit model, which yields a linear expression for vehicle market shares:

$$\log(s_{jt}/s_{0t}) = \psi_p(p_{jt} - p_{0t}) + \psi_x(x_{jt} - x_{0t}) + \sigma \log(s_{jt/gt}) + \kappa_j + \delta_t + \mu_{jt}, \quad (13)$$

where s_{jt} is the share of vehicle j , and $s_{jt/gt}$ is the share of vehicle j within nest g . The outside good, which is often interpreted as the option to purchase a used vehicle, is included as vehicle $j = 0$. The main regressor of interest, x_{jt} , represents expected cumulative fuel expenditures. The remaining terms are defined as in Sections 3 and 4.

Price can be decomposed into a constant portion (e.g., MSRP) and a time-varying negotiated discount (e.g., cash incentives). Denoting the constant portion of price as M_j and the discount as d_{jt} , the model can be re-written as follows:

$$\log(s_{jt}) = \psi_x x_{jt} + \sigma \log(s_{jt/gt}) + \kappa_j^* + \delta_t^* + \mu_{jt}^*, \quad (14)$$

where $\kappa_j^* = \kappa_j + \psi_p(M_j - M_0)$ is a composite vehicle fixed effect that absorbs the influence of time-invariant prices, $\delta_t^* = \delta_t + \log(s_{0t}) - \psi_p p_{0t} - \psi_x x_{0t}$ is a composite time fixed effect that absorbs the influence of the outside good, and $\mu_{jt}^* = \mu_{jt} - \psi_p(d_{jt} - d_{0t})$ is a composite error term that includes discounts. The vehicle fixed effects can be replaced with MSRP and other vehicle characteristics when variation in the data is more limited.

This formulation makes it apparent that the main regressor of interest, the expected cumulative fuel cost, is correlated with the composite error term due to the supply-side dis-

counting behavior of manufacturers. This produces inconsistency in estimators that require orthogonality between fuel costs and the residual; and inconsistency exists even if vehicle and period fixed effects are included. These fixed effects account for the average price of each vehicle and the average effect of fuel costs on vehicle prices, respectively, but do not account for differential impact of fuel costs on discounts across vehicles. Standard econometric manipulations yield an analytical expression for the bias of OLS estimates:

$$\widehat{\psi}_x \xrightarrow{p} \psi_x \left(1 - \underbrace{\frac{\text{Cov}(x_{jt}, d_{jt} | s_{jt/gt}, \kappa_j^*, \delta_t^*)}{\text{Var}(x_{jt} | s_{jt/gt}, \kappa_j^*, \delta_t^*)}}_{\text{bias term}} \right). \quad (15)$$

In the special case of the standard logit (i.e., $\sigma = 0$), the bias term simplifies to the covariance between fuel costs and discounts, conditional on the fixed effects (but not on shares within nest) and normalized by the variance of fuel costs. This is obtainable as the regression coefficient from an OLS regression of discounts on expected cumulative fuel costs, controlling for vehicle and time fixed effects. Sales information is unneeded.

We turn to the data for an empirical estimate of the bias term in standard logit models of demand. We regress the maximum incentive for a given vehicle-week observation on the measure of cumulative fuel costs that we develop in Section 4.2, controlling for vehicle and time fixed effects. We use the national sample of Table 4 (column 4) because discrete choice models typically use national data. We estimate with OLS and cluster the standard errors at the vehicle level. The resulting fuel cost coefficient of 0.1372 (standard error of 0.0537) indicates a bias term of 13.7 percent.⁴³

We suspect that bias would be exacerbated in the more general nested logit case, which features more intense localized competition. Here the bias term must be conditioned on the within-nest market shares ($s_{jt/gt}$). This makes empirical estimates infeasible in the absence of data on vehicle sales and, instead, we attempt to construct upper bounds by estimating the bias that would arise in an “extreme” model within which consumers never substitute across vehicles types but exhibit logit behavior within type. To this end, we regress cash incentives on cumulative fuel costs and the fixed effects, separately for cars, SUVs, and trucks. The resulting fuel cost coefficients are 0.7796 (standard error of 0.1714) for cars, 0.2486 (standard error of 0.1034) for SUVs, and 0.1745 (standard error of 0.1399) for trucks. This suggests wide range of possible bias for nested logit models, in which some consumer

⁴³This is still an approximation of the bias in a logit model, since in our regression there is one observation per vehicle-week, while in most discrete choice analyses observations will be weighted by sales.

substitution across nests is incorporated.

This bias is difficult to confront. Instrumental variables methods, such as two stage least squares, are inapplicable because the unobserved manufacturer price responses are literally functions of the observed fuel costs. It follows that any instrument with power is likely invalid.⁴⁴ And relying on regional variation in gasoline prices rather than inter-temporal variation in gasoline prices (e.g., as in Bento et al (2009)) may not suffice because manufacturers often vary their cash incentives at the local and regional level. Thus, we suspect the most promising path for discrete choice estimation involves the acquisition of high quality transaction price data, such as that of Busse, Knittel, and Zettelmeyer (2011).⁴⁵ Alternatively, interpretation can be softened. This is the approach of Klier and Linn (2010a), which estimates an regression along the lines of equation 14 and interprets the regression coefficient as a reduced-form estimate of how fuel cost changes affect vehicle sales.

7 Conclusion

We provide empirical evidence that automobile manufacturers adjust relative vehicle prices in response to changes in the price of retail gasoline. In particular, we show that the vehicle incentives tend to increase in their own fuel costs and decrease in the fuel costs of their competitors. The net effect is such that manufacturers offset through changes in relative incentives 40% of the change in relative fuel costs between any pair of vehicles. These differential price changes should incent firms investment in fuel economy research and design as gasoline prices increase or with the implementation of a gas tax. Additionally, we find that manufacturers' price responses may lead to downward bias of at least 13% in some discrete choice estimates of consumer demand for fuel economy. Both of these effects lead us to believe that gas taxes will be more effective at improving fleet fuel economy than previously suggested. The results do not speak, however, to the optimal magnitude of any policy responses; we leave that important matter to future research.

⁴⁴This statement might be too strong insofar as it assumes perfect knowledge on the part of manufacturers. Variables that affect consumer fuel cost forecasts and are unobserved by manufacturers could be both powerful and valid. Whether such instruments can be found is another matter.

⁴⁵The use of transaction prices in discrete choice models is not without difficulty because estimation requires a price for every vehicle considered not just each vehicle purchased. We refer the reader to Langer (2011) for one approach to dealing with this problem.

References

- [1] **Allcott, Hunt and Nathan Wozny** 2010. “Gasoline Prices, Fuel Economy, and the Energy Paradox” Manuscript: <http://web.mit.edu/allcott/www/papers.html>
- [2] **Alquist, Ron and Lutz Kilian.** 2010. “What Do We Learn from the Price of Crude Oil Futures? *Journal of Applied Econometrics*, 25: 539-573.
- [3] **Anderson, Soren T., Ryan Kellogg, and James Sallee.** 2011. “What Do Consumers Believe about Future Gasoline Prices?” NBER Working Paper 16974.
- [4] **Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen.** 2009. “Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes.” *American Economic Review*, 99(3): 667-669.
- [5] **Berry, Steven, James Levinsohn, Ariel Pakes** 1995. “Automobile Prices in Market Equilibrium” *Econometrica*, 63(4): 841-890
- [6] **Berry, Steven, James Levinsohn, Ariel Pakes** 2004. “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market” *Journal of Political Economy*, 112(1): 68-105
- [7] **Beresteanu, Arie and Shanjun Li.** 2011. “Gasoline Prices, Government Support, and Demand for Hybrid Vehicles in the U.S.” *International Economic Review*, 52(1): 161-182
- [8] **Blomquist, Ake G. and Walter Haessel.** 1978. “Small Cars, Large Cars, and the Price of Gasoline.” *Canadian Journal of Economics*, 11 (3): 470-489.
- [9] **Busse, Megan R., Christopher R. Knittel, and Florian Zettelmeyer.** May 2011. “Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing.” http://web.mit.edu/knittel/www/papers/gaspaper_latest.pdf
- [10] **Carlson, Rodney L.** 1978. “Seemingly Unrelated Regression and the Demand for Automobiles of Different Sizes, 1965-1975: A Disaggregate Approach.” *Journal of Business*, 51 (2): 243-262.
- [11] **Copeland, Adam, Wendy Dunn, and George Hall.** 2005. Prices, Production and Inventories over the Automotive Model Year.” NBER Working Paper 11257.

- [12] **Dahl, Carol A.** 1979. “Consumer Adjustment to a Gasoline Tax” *Review of Economics and Statistics*, 61: 427-432.
- [13] **Davis, Michael C. and James D. Hamilton.** 2004. “Why Are Prices Sticky? The Dynamics of Wholesale Gasoline Prices.” *Journal of Money, Credit and Banking*, 36 (1): 17-37.
- [14] **Energy Information Agency.** 2006. “A Primer on Gasoline Prices.” <http://www.eia.doe.gov/bookshelf/brochures/gasolinepricesprimer/>
- [15] **Geman, Hélyette.** 2007. “Mean Reversion Versus Random Walk in Oil and Natural Gas Prices.” In *Advances in Mathematical Finance*, ed. Michael C. Fu, Robert A. Jarrow, Ju-Yi Yen, and Robert J. Elliott, 219-228. Boston: Birkhauser.
- [16] **Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas.** 2010. “Revisiting the Income Effect: Gasoline Prices and Grocery Purchases.” *American Economic Review, Papers and Proceedings*, 100(2): 480-484.
- [17] **Goldberg, Pinelopi.** 1998. “The Effects of the Corporate Average Fuel Economy Standards in the Automobile Industry.” *Journal of Industrial Economics*, 46: 1-33.
- [18] **Gramlich, Jacob.** 2010. “Gas Prices, Fuel Efficiency, and Endogenous Product Selection in the U.S. Automobile Industry.” http://faculty.msb.edu/jpg72/Autos_GramlichJacob.pdf
- [19] **Greene, David L.** 2010. *How Consumers Value Fuel Economy: A Literature Review*, EPA-420-R-10-008, Washington, D.C.
- [20] **Greenlees, John S.** 1980. “Gasoline Prices and Purchases of New Automobiles.” *Southern Economic Journal*, 47: 167-178.
- [21] **Hamilton, James D.** 1994. Time Series Analysis. Princeton University Press. Princeton, New Jersey.
- [22] **Hamilton, James D.** 2009. “Understanding Crude Oil Prices.” *The Energy Journal*, 30(2): 179-206.
- [23] **Jacobsen, Mark.** 2010. “Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity.” http://econ.ucsd.edu/~m3jacobs/Jacobsen_CAFE.pdf

- [24] **Kahn, James A.** 1986. "Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach." *Quarterly Journal of Economics*, 101 (2) 323-339.
- [25] **Klier, Thomas and Joshua Linn.** 2010a. "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly New Vehicles Sales Data." *American Economic Journal: Economic Policy*, 2(3): 134-153.
- [26] **Klier, Thomas and Joshua Linn.** 2010b. "New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard." Resources for the Future Discussion Paper 10-60.
- [27] **Lancaster, Kelvin J.** 1966. "A New Approach to Consumer Theory." *Journal of Political Economy* 74: 132-157.
- [28] **Langer, Ashley.** 2011. "Demographic Price Discrimination and New Vehicle Sales" Manuscript available on www.ashleylanger.com.
- [29] **Langer, Ashley and Nathan H. Miller.** 2008. "Automobile Prices, Gasoline Prices, and Consumer Demand for Fuel Economy." EAG Working Paper 08-11.
- [30] **Li, Shanjun, Christopher Timmins, and Roger H. von Haefen.** 2009. "How Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy*, 1(2): 113-137.
- [31] **Makridakis, Spyros, Steven C. Wheelwright, and Rob J. Hyndman.** 1998. *Forecasting Methods and Applications*, 3rd ed. New York: Wiley.
- [32] **Miller, Don M. and Dan Williams.** 2004. "Damping Seasonal Factors: Shrinkage Estimators for the X-12-ARIMA Program." *International Journal of Forecasting*, 20 (4): 529-549.
- [33] **Peters, Anja, Michel G. Mueller, Peter de Haan, and Roland W. Scholz.** 2008. "Feebates Promoting Energy-Efficient Cars: Design Options to Address More Consumers and Possible Counteracting Effects." *Energy Policy*, 36 (4): 1355-1365.
- [34] **Petrin, Amil.** 2002. "Quantifying the Benefits of New Products: The Case of the Minivan" *Journal of Political Economy*, 110(4): 705-729
- [35] **Sallee, James M., Sarah E. West, and Wei Fan.** 2009. "Consumer Valuation of Fuel Economy: A Microdata Approach." *Proceedings of the National Tax Association Annual Conference on Taxation*.

- [36] **U.S. Department of Transportation, National Highway Traffic Safety Administration.** 2006. *Vehicle Survivability and Travel Mileage Schedules*, DOT HS 809 952, Washington, D.C.
- [37] **Wheaton, William C.** 1982. “The Long-Run Structure of Transportation and Gasoline Demand.” *Bell Journal of Economics*, 13: 439-454.