Gasoline Prices, Fuel Economy, and the Energy Paradox

Hunt Allcott and Nathan Wozny*

July 21, 2011

Abstract

It is often asserted that consumers undervalue future gasoline costs relative to purchase prices when they choose between automobiles, or equivalently that they have high "implied discount rates" for these future energy costs. We test this by examining how time series variation in gasoline prices affects the relative prices of vehicles with different fuel economy ratings, using a detailed dataset based on 86 million transactions at auto dealerships and wholesale auctions between 1999 and 2008. Our preferred empirical estimate is that vehicle prices move as if consumers are indifferent between one dollar in purchase price and 72 cents in discounted future gas costs, although we document how plausible alternative assumptions can affect this result. We introduce a new approach to behavioral welfare analysis in a discrete choice setting which generates two stark results. First, even if consumers undervalue gasoline prices by a relatively small amount, this distorts vehicle markets more than the failure to impose a Pigouvian carbon price. Second, even if consumers undervalue gas prices by a relatively large amount, CAFE standards increase average fuel economy by much more than can be justified by undervaluation.

JEL Codes: D03, L62, Q41.

We thank Alberto Abadie, Josh Angrist, Orley Ashenfelter, David Autor, Lucas Davis, Henry Farber, Kelly Gallagher, Penny Goldberg, Michael Greenstone, Jerry Hausman, Mark Jacobsen, James Kahn, Ryan Kellogg, Lutz Kilian, David Laibson, David Lee, Henry Lee, Robin Lee, Lars Lefgren, Erin Mansur, Erich Muehlegger, Sendhil Mullainathan, Richard Newell, Ariel Pakes, Nancy Rose, Jesse Rothstein, Stephen Ryan, Jim Sallee, Jorge Silva-Risso, Rob Stavins, Chris Timmins, Roger von Haefen, Glen Weyl, Matt White, and seminar participants at Analysis Group, Carnegie Mellon, Colorado School of Mines, Harvard, Kent State University, Mathematica Policy Research, MIT, NBER Summer Institute, Northwestern, Pontificia Universidad Catolica de Chile, Princeton, Stanford University, Stanford Institute for Theoretical Economics, Texas A&M, University of California Energy Institute, University of Delaware, University of Wisconsin, University of Toronto, and Washington University for their feedback. Funding for this project was provided by Princeton's Industrial Relations Section, Harvard's Energy Technology Innovation Policy group, the Harvard University Center for the Environment, and the Sloan Foundation. We thank JD Power for providing vehicle transaction price data. Thanks to Lonnie Miller at RL Polk for helpful conversations. We thank Manheim Consulting for providing auction data and Gary George, Nirmeet Kacheria, Pete Sauber, and Tom Webb at Manheim for sharing their expertise. We thank Carmen Collyns and Jordon Ricks for research assistance on this and related projects. We especially appreciate meticulous research assistance by Sounman Hong. Code for replicating the estimation is available from Hunt Allcott's website.

*Allcott: New York University and NBER. hunt.allcott@nyu.edu. Wozny: Mathematica Policy Research. NWozny@mathematica-mpr.com.

1 Introduction

In many domains, it appears that consumers choosing between products may be less attentive to ancillary costs than to purchase prices. Consumers on eBay, for example, are less elastic to shipping and handling charges than to the listed purchase price (Hossain and Morgan 2006). Mutual fund investors appear to be less responsive to ongoing management fees than to upfront payments (Barber, Odean, and Zheng 2005). Senior citizens are two to five times more sensitive to a Medicare Part D plan's premium than to its out-of-pocket costs (Abaluck and Gruber 2011). Shoppers are less elastic to sales taxes than to purchase prices (Chetty, Looney, and Kroft 2009).

Similarly, it is often asserted that gasoline costs are not fully salient to automobile consumers when they choose between automobiles with different fuel economy ratings (e.g. Greene *et al.* 2005). If this is true, consumers buy vehicles with lower fuel economy and higher resulting fuel costs than they would in their private optima. In 2007, the median-income American household spent \$2400 on gasoline, and consumers spent \$286 billion in total (U.S. BLS 2007). Misoptimization over such a large expenditure class could cause substantial welfare losses. The purported undervaluation of future gasoline costs would also help explain what Jaffe and Stavins (1994) call the "Energy Paradox": consumers and firms are puzzlingly slow to make seemingly high-return investments in energy efficiency.

Externalities from energy use related to national security and climate change would exacerbate these potential private losses. Policymakers have long debated whether it is preferable to address these externalities through gasoline taxes or Corporate Average Fuel Economy (CAFE) standards, which mandate an increase in the average fuel economy of new vehicles. In the absence of other market failures or misoptimization by consumers, economic analyses typically conclude that CAFE standards are strikingly inefficient relative to gas taxes.[1] If consumers undervalue fuel costs when they choose between vehicles, however, their response to gasoline taxes is not optimal. In this case, CAFE standards might be preferred to gas taxes, as they effectively force consumers to buy the energy efficient vehicles that they would want if they were optimizing. This "paternalistic" argument for energy efficiency policies has long been employed by both academic economists (Hausman

[1] Jacobsen (2010), for example, shows that the CAFE standard has a welfare cost of \$222 per metric ton of carbon dioxide abated, compared to \$92 per metric ton for an increase in the gasoline tax that reduces gasoline consumption by the same amount.

1979, Parry, Walls, and Harrington 2007)[2] as well as the U.S. government. In fact, in the Regulatory Impact Analysis that justifies a recent increase in the CAFE standard, the vast majority of benefits flow through gasoline cost savings to consumers who undervalue the benefits of fuel economy (NHTSA 2010).[3] Put simply, while political feasibility plays an important practical role, paternalism is a leading economic justification for one of the most important and costly public policies affecting the U.S. automotive and energy industries.

One problem with the paternalistic justification for fuel economy standards is the dearth of solid evidence on whether automobile buyers actually are or are not misoptimizing. This paper aims to fill that gap by exploiting a painful but convenient natural experiment, the dramatic fluctuations in gasoline prices over the past ten years. We test the null hypothesis that consumers are willing to pay exactly one dollar more to purchase a vehicle with one dollar less in total forecasted future fuel costs, discounted to present value at their intertemporal opportunity cost of funds. For expositional purposes, we will say that rejecting this hypothesis is evidence that consumers "undervalue or overvalue gasoline costs."[4]

Our question is related to a long literature, dating at least to the energy crises of the 1970s, that estimates consumers' "implied discount rates" for energy efficiency investments and compares them to benchmark consumer discount rates. The typical empirical approach has been to exploit variation in the prices and energy efficiency ratings of a cross-section of energy-using durable goods.

[2] In the Bell Journal of Economics, Hausman (1979) estimates that consumers implicitly use a discount rate of 15 to 25 percent per year when they trade off purchase prices and future energy costs of new air conditioners. He concludes (page 51), "Yet this finding of a high individual discount rate does not surprise most economists. At least since Pigou, many economists have commented on a "defective telescopic faculty." A simple fact emerges that in making decisions which involve discounting over time, individuals behave in a manner which implies a much higher discount rate than can be explained in terms of the opportunity cost of funds available in credit markets. Since this individual discount rate substantially exceeds the social discount rate used in benefit-cost calculations, the divergence might be narrowed by policies which lead to purchases of more energy-efficient equipment."

In the Journal of Economic Literature, Parry, Walls, and Harrington (2007) write, "Higher fuel economy standards significantly increase efficiency only if carbon and oil dependence externalities greatly exceed the mainstream estimates . . . or if consumers perceive only about a third of the actual fuel economy benefits . . . Unfortunately, there is little in the way of solid empirical (as opposed to anecdotal) evidence on this hotly contested issue . . . "

[3] NHTSA (2010, page 2) writes, "Although the economy-wide or "social" benefits from requiring higher fuel economy represent an important share of the total economic benefits from raising CAFE standards, NHTSA estimates that benefits *to vehicle buyers themselves* [original emphasis] will significantly exceed the costs of complying with the stricter fuel economy standards this rule establishes . . . However, this raises the question of why current purchasing patterns do not result in higher average fuel economy, and why stricter fuel efficiency standards should be necessary to achieve that goal. To address this issue, the analysis examines possible explanations for this apparent paradox, including discrepancies between the consumers' perceptions of the value of fuel savings and those calculated by the agency . . . "

[4] Other research has used different words related to what we label as "undervaluation," such as myopia, inattention, biased beliefs, shrouding, salience, and high "implied discount rates."

For example, our null hypothesis in a cross-sectional discrete choice framework would be that, after conditioning on other observed product characteristics, a one dollar increase in a product's purchase price is associated with the same decrease in market share as a one dollar increase in total discounted energy costs. Analogous cross-sectional approaches are used in a seminal paper by Hausman (1979) on air conditioners, as well as analyses of other energy-using durables such as houses (Dubin 1992), water and space heating (Dubin and McFadden 1984), and autos (e.g. Dreyfus and Viscusi (1995), Espey and Nair (2004), and Goldberg (1998)).

For the cross-sectional estimator to be unbiased, the functional form for how other observed product characteristics enter utility must be correctly specified, and any unobserved characteristics must be uncorrelated with energy efficiency. Especially with automobiles, these assumptions appear problematic. Fuel economy is mechanically correlated with weight and horsepower, and it has often proven difficult to separately identify preferences for these different characteristics.[5] Furthermore, fuel economy is highly *negatively* correlated with price in the cross section, suggesting that larger vehicles have more observed and unobserved amenities.

Notice, however, that a vehicle's total future fuel costs vary as a function of both fuel economy and forecasted gasoline prices. As an alternative to the cross-sectional approach, we build an empirical test around the simple intuition that the changes in gas prices over the past decade should affect the relative prices of high- vs. low-fuel economy vehicles. Indeed, media reports and academic analyses have documented that as gasoline prices rise, the relative prices of low-fuel economy vehicles drop (Busse, Knittel, and Zettelmeyer 2009, Langer and Miller 2009). The above null hypothesis, however, does more than predict that gasoline prices should affect vehicle demand: it predicts exactly how much demand should be affected. Intuitively, if relative vehicle prices are not sufficiently responsive to changes in forecasted gasoline costs, this suggests that consumers undervalue gasoline costs when they purchase vehicles.

More precisely, our analysis begins with microdata on 86 million transactions at both auto dealerships and wholesale auctions between 1999 and 2008. For each month of this study period, these data are collapsed to the average price for each new and used vehicle in consumers' choice

[5] At least since Atkinson and Halvorsen (1984), it has been pointed out that the high correlation between weight and fuel economy makes it difficult to separately identify demand for fuel economy. In fact, cross sectional estimation of automobile demand in characteristic space sometimes gives the "wrong" sign on fuel economy.

sets. Each vehicle has a different present discounted value (PDV) of fuel costs, depending on its fuel economy rating, survival probabilities, gasoline price forecasts, discount rates, and annual vehicle-miles traveled. We condition on model-by-age fixed effects, which sweep out all observed and unobserved vehicle characteristics, and test whether relative prices move one-for-one with changes in the PDV of fuel costs.

In our base specification, we estimate that American auto consumers value 72 percent of gasoline costs. That is, we document that the average auto buyer's real intertemporal opportunity cost of funds is six percent, and we then estimate that at this discount rate, vehicle prices move as if consumers are indifferent between one dollar in purchase price and 72 cents in discounted future gas costs. The "implied discount rate," the discount rate for future gas costs that rationalizes market behavior, is 16 percent. The result that consumers undervalue gasoline costs by at least some amount is robust to a number of alternative assumptions around consumers' gasoline price forecasts, unobserved changes in consumer preferences and vehicle characteristics, and alternative empirical strategies that allow market shares and vehicle-miles traveled to vary endogenously with gas prices. However, we emphasize from the outset that it will also be quite possible to construct a plausible combination of alternative assumptions under which we fail to reject the null hypothesis.

To help interpret the empirical results, we build on Bernheim and Rangel (2009) to develop a new and highly-tractable application of behavioral welfare analysis in a discrete choice setting. This is used to analyze a counterfactual "behavioral feebate" policy that conceptually resembles the "internality tax" from O'Donoghue and Rabin (2006): it imposes sales taxes that decrease in fuel economy such that misoptimizing consumers purchase their privately-optimal vehicle. Under our base specification demand parameters and some stylized modeling assumptions, the present discounted value of welfare gains from such a policy is \$5 per potential vehicle consumer per year. Across approximately 240 million potential vehicle consumers in the United States, this sums to \$1.3 billion per year that the policy is in place. Again, we emphasize that these policy simulations should be interpreted cautiously due to the statistical and modeling uncertainty: the extremes of our point estimates under different assumptions are that consumers value 50 or 90 percent of gasoline costs, and in these cases, the gains from the optimal behavioral feebate are \$5.4 billion or \$192 million, respectively. If, contrary to the bulk of our empirical results, the average consumer in fact correctly values gasoline costs, such a policy is not desirable at all.

Before continuing, it is worth doing a simple calibration to demonstrate why it is difficult to overstate the importance of this question. Substantial volumes of academic research and policy discussions have centered on the welfare losses from transport sector carbon emissions and the costs and benefits of different policy responses. Using the U.S. Government's (2010) estimated marginal damage of carbon emissions, gasoline consumption imposes an externality of \$0.18 per gallon, or five to ten percent the current gasoline price. Thus, if the carbon externality is not internalized, consumers account for only 90 to 95 percent of the total social cost of gasoline when they choose between vehicles with different fuel economy ratings. By comparison, our preferred empirical estimates suggest that consumers value only 72 percent of the cost of gasoline when they choose between vehicles. Thus, while undervaluation and uninternalized carbon externalities distort vehicle purchases in the same direction, inducing consumers to buy vehicles that use more gas than in the first best, undervaluation could generate distortions *several times larger* than the distortions from climate change externalities.

The paper progresses as follows. In Section 1.1, we discuss related literature. Section 2 models consumers' utility functions. Section 3 presents the data at our disposal, devoting particular attention to the construction of each vehicle's total discounted gasoline costs. Section 4 details our estimation strategy. Section 5 presents empirical results and a long series of robustness checks. Section 6 formalizes our approach to behavioral welfare analysis, and Section 7 concludes.

1.1 Related Literature

We are not the only researchers to examine how the prices and quantities of new and used vehicles respond in equilibrium to changes in gas prices. While there is a large body of very good related research, our analysis is most comparable to four projects.[6]

Kahn (1986) tests whether relative prices of used vehicles fully adjust to changes in the relative discounted present value of relative gas prices induced by the gasoline price shocks of the 1970s

[6] Other work that examines how vehicle prices adjust in response to gasoline prices include Sawhill (2008), Langer and Miller (2009), and Austin (2008). Verboven (1999) estimates the discount rates implied by differences between the prices of gasoline and diesel vehicles in Europe. Ohta and Griliches (1986) examine whether the 1970s gasoline price shocks affected consumers' valuations of vehicle characteristics.

and 1980s. A working paper by Kilian and Sims (2006) builds on Kahn's approach with updated data. The key difference between our analysis and these two papers is that we use transaction prices from auctions and dealerships instead of data from the National Auto Dealers' Association Used Car Guide or other printed used car price guides such as the Kelley Blue Book. Early in this project, we recognized that using transaction data was crucial, because used car price guides reflect the opinion of a small team of analysts who may or may not fully adjust their price assessments for each vehicle to reflect current market conditions. We decided that using these data could cause us to falsely conclude that vehicle market prices do not fully adjust to changes in gasoline prices.

Like us, Sallee, West, and Fan (2009) exploit transaction data from used vehicle auctions to test whether vehicle prices move one-for-one with the present discounted value of future gasoline costs. The authors' estimation is interesting and complementary to ours in that it substitutes a different identifying assumption in place of what we will label as Assumptions 1 and 2. At the time of this draft, however, there is no working paper available.

A working paper by Busse, Knittel, and Zettelmeyer (2011), hereafter "BKZ," estimates how changes in gasoline prices affect equilibrium prices and quantities of new and used vehicles. BKZ's project is complementary to ours in that their analysis is oriented around reduced form estimates of how these equilibrium effects vary by quartile of the fuel economy distribution. They use these estimates to emphasize how the different structures of new and used vehicle markets have different implications for policy design. BKZ's draft also includes a stylized calculation of the implied discount rates at which the prices of each MPG quartile fully adjust to changes in gasoline costs. Their implied discount rate calculation requires them to make the same types of assumptions that we must make, including parameterizations of substitution patterns between vehicles and assumptions about the elements of a vehicle's future gasoline costs, including survival probabilities, vehicle-miles traveled, and future gasoline prices.

While our analysis is motivated differently and is somewhat more formal, BKZ's stylized calculation is fundamentally comparable to our paper, and they use very similar data. Surprisingly, however, BKZ come to the opposite conclusion: their empirical results show that consumers overvalue gasoline costs, with implied discount rates for used cars that average *negative* 0.3 percent.[7]

[7] This is the average of the nine implied discount rates for used cars in BKZ (2011), Table 7.

In Section 5.3.8, we document that an important reason why BKZ appear to have found this result is that they have a severe endogeneity problem resulting from how their fixed effects and control variables fail to control for differential depreciation of low-MPG vs. high-MPG vehicles. We present two different identification strategies which are robust to this issue, either of which can be used by BKZ or other researchers with comparable data.

2 Model

Our utility function takes the standard discrete choice model from industrial organization (Berry 1994) and adds some trivial dynamics, which help clarify the identifying assumptions to be made later in the analysis. Consumers derive utility from owning a vehicle and from consuming a numeraire good. In each period t, each consumer chooses one vehicle from a set of new and used vehicles, including the vehicle that the consumer already owns. We define a "vehicle" as a model-by-age combination, where j indexes models and a indexes age. Consumers also can choose an outside option, denoted $j = 0$, which is to own no vehicle. The utility of the outside option is normalized to zero.

Denote vehicle ja's purchase price at time t as p_{jat}. Consumers expect to own the vehicle for a holding period of h years, after which time it will be resold at price $p_{ja,t+h}$. Over that holding period, consumer i has budget constraint w_i. Individuals discount cash flows in future periods by discount factor $\delta = \frac{1}{1+r}$.

G_{jat} is the present discounted value (PDV) of future gasoline costs over the vehicle's remaining life. G_{jat} can be divided into two parts: $G_{jat} = G^o_{jat} + \delta^h G_{ja,t+h}$. The variable G^o_{jat} is the PDV of gasoline costs during consumer i's holding period, and $\delta^h G_{ja,t+h}$ is the PDV over the remainder of the vehicle's life after it is resold, discounted to time t.

The parameter η is the marginal utility of money. The parameter γ is an "attention weight" on fuel costs: if consumers value purchase prices and discounted fuel costs equally, then $\gamma = 1$. If consumers undervalue or overvalue fuel costs, then $\gamma < 1$ or $\gamma > 1$, respectively. We assume that both η and γ are constant.

The variable $\widetilde{\psi}^o_{jat}$ captures the average utility across consumers from owning and using vehicle

ja over the holding period; we will call this the "usage utility." Individual i's unobserved deviation from average utility is ϵ_{ijat}.

Using these variables, consumer i's indirect utility from from purchasing vehicle ja in year t is:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G^o_{jat} + \delta^h p_{ja,t+h}) + \widetilde{\psi}^o_{jat} + \epsilon_{ijat} \tag{1}$$

We define a variable $\widetilde{\psi}_{jat}$, which captures average usage utility $\widetilde{\psi}^o_{jat}$ plus the present discounted utility value of resale price plus remaining gasoline costs:

$$\widetilde{\psi}_{jat} = \widetilde{\psi}^o_{jat} + \eta\delta^h \left(p_{ja,t+h} + \gamma G_{ja,t+h}\right) \tag{2}$$

We further define a variable $\widetilde{\psi}_{ja}$, which captures the average value of usage utility for vehicle ja across all time periods, and an idiosyncratic period-specific deviation $\widetilde{\xi}_{jat} \equiv \widetilde{\psi}_{jat} - \widetilde{\psi}_{ja}$. Utility can now be re-written as:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G_{jat}) + \widetilde{\psi}_{ja} + \widetilde{\xi}_{jat} + \epsilon_{ijat} \tag{3}$$

This utility function is quite natural, and it closely resembles many other discrete choice specifications. In words, utility depends on purchase price, the PDV of gasoline costs over the vehicle's entire remaining life, usage utility, and an unobserved individual-specific shock.

In our estimation later in the paper, we will assume that $\widetilde{\psi}_{ja}$ is constant within a vehicle ja over time up to an idiosyncratic deviation which is uncorrelated with G. Why is this intuitively sensible? Consider the two elements of $\widetilde{\psi}$ defined by Equation (2). The first element, the utility $\widetilde{\psi}^o$ from using the vehicle over a given holding period, is assumed constant for a given model of a given age, independent of the time t when the holding period starts. For example, buying a three-year-old Ford Taurus in 1999 gives the same expected usage utility as buying a three-year-old Ford Taurus in 2000, up to the idiosyncratic error term.

The second element of $\widetilde{\psi}$ contains the total user cost for the next owner: the sum of the vehicle's resale price and remaining fuel cost, weighted by γ. In essence, we assume a weak version of *stationarity*, a common assumption in dynamic analyses of durable goods markets (e.g. Rust (1985), Stolyarov (2002)): consumers believe that resale prices of the same good are constant over time, after adjusting for changes in other user costs and allowing for an idiosyncratic error. Continuing the above example, we assume that consumers believe that if gasoline prices are constant, the resale price of an eight-year-old Ford Taurus in 2004 is the same as the resale value of an eight-year-old Ford Taurus in 2005. If gasoline prices change, the vehicle's resale price will change by the difference in discounted remaining fuel costs $G_{ja,t+h}$, multiplied by the market's attention weight γ. This also illustrates intuitively why a consumer's utility from choosing a particular vehicle depends not just on the gasoline costs that he himself will pay, but on the entire remaining fuel cost incurred by all future owners: the consumer expects that if gasoline prices increase, this will affect both his own gasoline expenditures G^o and also the resale value.

We model substitution patterns using a nested logit framework, which allows consumers' idiosyncratic preferences to be correlated across vehicles within the same predetermined set of vehicles, or "nest": $corr(\epsilon_{ijat}, \epsilon_{ij'a't})$ is nonnegative when ja and $j'a'$ are in the same nest and zero otherwise. A scalar $\sigma \in [0, 1]$ parameterizes these within-nest correlations, with larger σ indicating more highly correlated taste shocks within nests.[8] The nests are specified ex ante, comprising vehicles that the analyst believes are closer substitutes. In our nested logit specifications, we use vehicle size classes as nests because consumers are likely to be more willing to substitute to vehicles of similar sizes, and size is closely connected with fuel economy.[9]

Under the usual extreme value distributional assumptions for ϵ_{ijat}, the nested logit choice

[8] Specifically, the cumulative distribution function for ϵ_{ijat} for all ja for individual i at time t is:

$$F(\cdot) = \exp\left[-\sum_{n\in\mathcal{N}}\left(\sum_{ja\in\mathcal{B}_n} e^{-\epsilon_{ijat}/(1-\sigma)}\right)^{1-\sigma}\right]$$

In this equation, $\mathcal{N}$ is the set of all nests of vehicles, and $\mathcal{B}_n$ is the set of vehicles in nest n. As σ approaches one, the within-nest correlation of utilities approaches one. If $\sigma = 0$, the representative consumer logit model is recovered. This distribution can be extended to accommodate multiple nests or separate σ parameters for each nest.

[9] Another common way of parameterizing unobserved taste shocks is through a random coefficients model, which is more flexible in that it can allow preferences for continuous attributes such as horsepower and weight to vary across the population. The nested logit is equivalent to allowing a random coefficient on the nest indicator. We adopt the nested logit approach because our choice set is unusually large and because it makes our estimation procedure very transparent: we will be able to estimate a simple log-linear relationship between market-level prices and shares.

probabilities can be aggregated over the population to give a market-level relationship between prices and shares. Often, this is rearranged to give the term $\ln s_{jat} - \ln s_{0t}$ on the left:

$$\ln s_{jat} - \ln s_{0t} = -\eta p_{jat} - \eta\gamma G_{jat} + \sigma \ln(s_{jat}/s_{n,jat}) + \widetilde{\psi}_{ja} + \widetilde{\xi}_{jat} \quad (4)$$

It is more unusual but entirely equivalent to write this identity with prices on the left. Rearranging slightly, we have:

$$p_{jat} = -\gamma G_{jat} - \frac{1-\sigma}{\eta} \ln s_{jat} - \frac{\sigma}{\eta} \ln s_{n,jat} + \frac{1}{\eta} \ln s_{0t} + \psi_{ja} + \xi_{jat} \quad (5)$$

This equation includes new variables $\psi_{ja} \equiv \frac{\widetilde{\psi}_{ja}}{\eta}$ and $\xi_{jat} \equiv \frac{\widetilde{\xi}_{jat}}{\eta}$, which intuitively represent the dollar value of the utility represented by $\widetilde{\psi}_{ja}$ and $\widetilde{\xi}_{jat}$. Section 4 will begin with this equation in presenting our identification and estimation strategy.

3 Data

In this section, we detail our dataset, which includes average prices, quantities, and characteristics of all passenger vehicle models registered in the U.S., in monthly cross sections from January 1999 to December 2008.

Fuel economy data were obtained from the U.S. Environmental Protection Agency (EPA), which measures fuel use over a standardized laboratory drive cycle and then adjusts the results to account for the typical consumer's actual in-use fuel economy. We use the EPA's 2008 adjustment for all model years, and we assume that a vehicle's fuel economy degrades over its life at an average of 0.07 MPG per year, due to the findings of a large survey analyzed by Greene *et al.* (2007). Vehicle classes - pickups, sport utility vehicles, minivans, vans, two-seaters, and five classes of cars based on interior volume - are also taken from the EPA's fuel economy dataset. All other vehicle characteristics, including horsepower, curb weight, wheelbase, and Manufacturer's Suggested Retail Price (MSRP), are from Ward's Automotive Yearbook. All dollar figures in this paper are real July

2005 dollars, deflated using the Bureau of Labor Statistics Consumer Price Index series for "All Urban Consumers, All Items Less Energy."

Used vehicle prices are based on microdata obtained from Manheim, a firm which intermediates approximately half of auto auction transactions in the United States. The principal sellers are dealerships, rental car companies, and auto manufacturers re-selling off-lease vehicles. Buyers are typically dealerships which then retail the used vehicles. We observe each of Manheim's approximately 45 million transactions between 1999 and 2008. While only about one in four used vehicles traded passes through an auction (Manheim 2009), the auction market is the largest source of transaction price data. Furthermore, the Kelley Blue Book and other price guides, which are the starting point for price negotiations in many of the non-auction transactions, are largely based on auction prices. We collapse the data to the level of mean price for each vehicle in each month. Adjusting this price for the vehicle's condition, odometer reading, auction location, and method of sale does not affect our results.

New vehicle prices are from the JD Power and Associates (JDPA) "Power Information Network," which collects detailed microdata on approximately 31 percent of US retail auto transactions through a network of more than 9,500 dealers. For each vehicle ja, we observe monthly mean prices adjusted for customer cash rebates and the difference between the negotiated trade-in price and the trade-in vehicle's actual resale value, if any. In specification checks, we use used vehicle prices from JDPA instead of from Manheim. We use the Manheim data in our base specifications because these data include more than twice as many observations, while while there are fewer than 1000 observations in JDPA that are not in Manheim.

We observe national-level registered quantities of each vehicle model in each year from 1999 through 2008 in the National Vehicle Population Profile, a dataset obtained from market research firm R.L. Polk. The quantities represent all vehicles registered to private individuals and to fleets such as taxi and rental car companies and corporate and government motor pools. The registration data are annual snapshots taken on July 1, meaning that some vehicles of the latest model year have not been registered. Since very few vehicles are scrapped in their first year of life, we set the quantity for new vehicles equal to their quantity at age one. These quantity data are matched to new and used vehicle price data using the industry's serial numbers, called Vehicle Identification

Numbers.

We define the market to be all "substitutable" gasoline-fueled light duty vehicles with EPA fuel economy ratings that are less than 25 years old. This includes cars, pickups, SUVs, minivans, and other light trucks, but not motorcycles, cutaway motor homes, limousines, camper vans, chassis cab and tilt cab pickups, hearses, and other unusual vehicles where we expect the substitution elasticity to be very small. In our base specification, we also exclude cargo and passenger vans as well as ultra-luxury and ultra-high performance exotic vehicles[10] due to their low substitutability with the rest of the market. It turns out that including these vehicles does not change the results.

We define a "model" j to capture all possible variation in fuel economy ratings and much of the observable variation in prices. This is more disaggregated than a "nameplate," which we use to refer to a colloquial name such as "Ford Taurus" or "Honda Civic." We define a "model" at the level of make, nameplate, trim level, body type, engine displacement in liters, and the number of cylinders. As a result, the average make and nameplate combination in our dataset includes seven different "models." For example, there are 11 different configurations of model year 2004 cars called the "Honda Civic" that appear in our dataset as separate "models," including coupe and sedan versions of the DX, EX, and LX, the Si Hatchback, the Civic Hybrid, and several others.

Vehicles with the same model name are typically offered for several consecutive model years, although some are offered for many more. As extreme examples, the Ford F-150 and Honda Civic have each been offered in every model year since 1973. Of course, the 1973 and 2008 versions of these models are very different. Every several years, auto manufacturers redesign their models and define a new "generation" of a vehicle. For example, Honda introduced new generations of the Civic for the 1980, 1984, 1988, 1992, 1996, 2001, 2006, and 2012 model years. A "generation" is a well-defined concept, and the generation definitions for each vehicle can even be found on Wikipedia. Within each generation, models do not change significantly. We redefine each new generation of a car or truck as a separate "model" j in our dataset.

Table 1 presents descriptive statistics. The full dataset includes 1,396,254 observations. Of these, 854,248 are used in the base specification. Table 2 describes how the sample is trimmed from

[10] We define exotic vehicles to be the Acura NSX, Audi R8 and TT, Chrysler Prowler and TC, Cadilliac Allante and XLR Roadster, Chevrolet Corvette, Dodge Viper and Stealth, Ford GT, Plymouth Prowler, and all vehicles made by Alfa Romeo, Bentley, Ferrari, Jaguar, Lamborghini, Maserati, Maybach, Porsche, Rolls-Royce, and TVR.

the full dataset to the base estimation sample.

3.1 Discounted Gasoline Costs

We now describe the formulation of the different components of expected discounted gasoline costs G_{jat}, including vehicle-miles traveled (VMT), survival probability, discount rates, and gasoline price forecasts. Our parameter of interest γ will be the coefficient on this variable, and decisions we make here will mechanically affect the parameter estimates. For example, using a lower discount rate than consumers actually face would inflate G_{jat}, thereby biasing $\widehat{\gamma}$ toward zero. Alternatively, understating a vehicle's expected lifetime or usage would deflate G_{jat}, biasing $\widehat{\gamma}$ away from zero.

We assume that G_{jat} is homogeneous for all consumers that choose vehicle ja at time t. In reality, there is substantial variation in vehicle-miles traveled across consumers that own the same vehicle. Furthermore, differences in the proportion of city versus highway driving generate differences in realized fuel economy, and gasoline price expectations may vary across consumers within and between geographic areas. As gasoline prices change, the change in relative prices between two vehicles is determined by a set of marginal consumers.[11] Because there can be many marginal consumers with different draws of ϵ_{ijat} for any pair of vehicles, and because a vehicle's price moves relative to many other vehicles, the set of marginal consumers can be actually quite large. We therefore construct a G_{jat} to reflect the average consumer who owns vehicle ja at time t. We do not believe that this systematically biases our results.

The variable G_{jat} is the present discounted value of lifetime gasoline costs over future years s:

$$G_{jat} = \sum_{s=t+1}^{t+(L-1-a)} \delta^{s-t} \cdot g_s \cdot m_{jas} \cdot f_{jas}^{-1} \cdot \phi_{jas} \tag{6}$$

L denotes the maximum possible lifetime of a vehicle, which we take to be 25 years. The variable g_s is the gasoline price forecast for year s, m_{ja} is expected vehicle-miles traveled, f_{jas}

[11] Note that changes in gasoline prices should cause a re-sorting of vehicles across consumers with different VMT. For example, if gas prices increase, consumers with relatively high VMT are more likely to switch to a vehicle with higher MPG, while consumers with relatively low VMT would switch to a vehicle with lower MPG. Intuitively, this re-sorting among inframarginal consumers does not matter: it is the VMT of the marginal consumers that determines equilibrium prices.

is fuel economy in miles per gallon, ϕ_{jas} is the probability that the vehicle survives to year s conditional on surviving to its current age, and δ is the annual discount factor.

3.1.1 Vehicle-Miles Traveled and Survival Probability

To estimate vehicle-miles traveled (VMT), we use publicly-available data from the 2001 National Household Travel Survey (NHTS). This is a nationally-representative survey of approximately 25,000 households that report, among many other variables, the age, fuel economy, and vehicle class for each of their vehicles. As part of the survey, about 25,000 vehicles in the national sample had their odometers read twice, with several months between readings. These two readings were then used to estimate annualized VMT. We regress annualized VMT on class dummies and vehicle age and use these estimates to fit m_{jas} for all vehicles in our sample. The standard errors in this regression are Huber-White robust; these will later be used to calculate the Murphy-Topel standard errors on $\widehat{\gamma}$.

Our base specification assumes that this fitted m_{jas} does not depend on the gasoline price forecast at time s. Either of two separate arguments can justify this. First, VMT demand is relatively inelastic. Second, the Envelope Theorem can be used to show that since indirect utility u_{ijat} is a function of an optimized value of VMT, changes in VMT caused by marginal changes in gasoline prices have only second-order effects on u_{ijat} and vehicle prices. In the Appendix, however, we derive an alternative approach that allows m_{jas} to be a function of g_s. As we will see, the results are very similar to the base specification.

We use an analogous approach to fit vehicle survival probabilities based on the NVPP registered quantity data. We use a grouped data probit model, where the outcome variable is the number of vehicles of a model and model year registered next year divided by the number of vehicles registered today. We estimate coefficients on age dummies, model year, and fuel economy, with robust standard errors clustered by vehicle, and then predict survival probabilities ϕ_{jas} for all vehicles in our sample.

3.1.2 Discount Rates

The discount rate $r = \delta^{-1} - 1$ should reflect how consumers discount the marginal dollar that could be used to buy a vehicle in the present or buy gasoline in the future. Put differently, r is the rate at which the purchase price is amortized over future years, or the rate at which future gasoline costs are discounted to the present. For vehicle buyers whose marginal dollar comes from a loan or lease, the opportunity cost of paying more to purchase a vehicle in the present is the Annual Percentage Rate (APR). For consumers whose marginal dollar comes from savings, the opportunity cost is the return that could be realized on savings.

Table 3 presents our estimate of the average discount rate across vehicle transactions during the study period. To construct this table, we examined the set of vehicles in the 2001, 2004, and 2007 Surveys of Consumer Finances (SCF) that respondents had bought in the year of the survey. Of these vehicles, 37 percent were new and 63 percent were used. We then calculated the share of transactions paid for by financing, lease, or cash. Of new vehicle transactions, for example, 54 percent were financed, 19 percent were leased, and 27 percent were paid for in cash.

For new and used vehicles financed with loans, the mean real interest rates reported in the SCF were 3.9 and 6.9 percent, respectively.[12] Interestingly, older versus newer used vehicles do not have different average interest rates. The mean real Annual Percentage Rate for leases was 3.0 percent. For vehicles purchased with cash, we assume that the opportunity cost of funds reflects market returns, and the average real return on the S&P 500 from 1945 to 2008 was 5.8 percent. As shown in Table 3, averaging these three discount rates weighted by their share of transactions gives 5.5 percent for all vehicles and 6.2 percent for used vehicles only. We round this to six percent for our base specification.

This calculation serves simply to generate a sensible guess at consumers' average discount rate, and it is far from inarguable. For example, the real average used vehicle loan APR reported by dealerships through the JD Power Information Network is 8.9 percent, two percentage points higher than the average value reported in the SCF. On the other hand, including years before 1945 or

[12] An alternative source of data on nominal new car loan interest rates is the Federal Reserve Board's G.19 consumer credit survey, available from http://www.federalreserve.gov/releases/g19/hist/cc_hist_tc.txt. The mean real APRs for new car loans over our study period were 3.1 and 5.5 percent from auto finance companies and commercial banks, respectively. These numbers are approximately consistent with the 3.9 percent average from the SCF.

after 2008 would give lower stock market returns. Furthermore, if we modeled consumers with declining marginal utility of consumption, they would want to risk-adjust returns for covariance with the market. Risk adjustment using the Capital Asset Pricing Model would give an interest rate close to the real risk free rate, about 1.6 percent, because changes in gasoline prices have very low correlation with market returns. In the results section, we will show the sensitivity of $\widehat{\gamma}$ to alternative discount rates.

3.1.3 Gasoline Price Forecast

We construct our measure of consumers' gasoline price forecasts g_s using oil futures prices at each time t over the study period. We use U.S. City Average Motor Gasoline Retail Prices for all types of gasoline, as well as Light Sweet Crude Oil prices, from the U.S. Energy Information Administration's (2011) Monthly Energy Review. These data show that oil spot prices predict 94 percent of the monthly variance in gasoline prices, which implies that oil prices are reasonable proxies for gasoline prices. Light Sweet Crude Oil futures prices are from the New York Mercantile Exchange (NYMEX) and the Intercontinental Exchange (ICE).

To calculate the gasoline price forecasts, the oil futures prices are transformed into current dollars using inflation expectations implied by Treasury Inflation Protected Securities, then deflated into real July 2005 dollars, then transformed to dollars per gallon of gasoline using the average historical relationship between oil and gasoline prices. Specifically, this average historical relationship is predicted with a simple linear regression of levels of the Motor Gasoline Retail Prices on the levels of Light Sweet Crude Oil prices. Table 4 presents the annual average real retail gasoline prices and expected prices implied by futures for 1998 to 2008.

Although oil futures contracts are only traded with high liquidity for settlement dates closer than two to three years in the future, Table 4 indicates that there are some trades observed for settlement dates as far as ten years in the future.[13] The futures market does not believe that gasoline prices are a martingale: as illustrated in Figure 1, as gas prices rose between 2003 and 2008 above their

[13] To model expectations for periods beyond the last settlement date observed at each time t, our base specification uses a simple model of mean-reverting expectations, where deviations from a \$1.50/gallon mean decay exponentially using a mean reversion parameter calibrated using all futures data since 1991. The equation fits the data very well: it explains 85% of the variation in the observed futures prices over our 1999-2008 study period.

1990's average of approximately \$1.50 per gallon, the futures market expected prices to eventually return closer to that previous level. In Section 5, we will document the sensitivity of $\widehat{\gamma}$ to different formulations of gasoline price forecasts and will show that our assumption fits the data better than assuming that consumers believe that real gasoline prices are a martingale.

4 Estimation

In this section, we first describe the intuition behind our fixed effects identification. We then describe our base empirical specification, which is remarkably simple and transparent but requires the identifying assumption that market shares are uncorrelated with gas costs. Finally, we present an alternative specification that allows for endogenous market shares.

4.1 Base Specification

We use vehicle fixed effects and look "within" the same vehicle over time as gasoline prices change. The benefit of a panel is that the fixed effects soak up unobserved vehicle characteristics that may be correlated with MPG: mathematically, our estimator is still unbiased even if $E[G\psi] \neq 0$.

We move from Equation (5) to our base specification estimating equation in two steps. First, we define an econometric error term that contains ξ and the market share terms: $\varepsilon \equiv \left(-\frac{1-\sigma}{\eta}\ln s_{jat} - \frac{\sigma}{\eta}\ln s_{nt}\right) + \xi_{jat}$. Second, we add time dummies τ_t, which absorb the outside option share and any shift in the overall market price level. Our base specification is:

$$p_{jat} = -\gamma G_{jat} + \tau_t + \psi_{ja} + \varepsilon_{jat} \tag{7}$$

Equation (7) is quite intuitive: it tests whether the relative vehicle prices move one-for-one with changes in the relative PDV of gasoline costs. If vehicle prices do not respond sufficiently to gasoline costs, we conclude that the market undervalues gasoline costs.

To see our identifying variation visually, consider Figure 2, which shows the average G for an example month by age and fuel economy rating. Of course, newer and lower-MPG vehicles tend to have higher G. Equation (7) is identified by multiplying this cross-sectional variation by the time-series variation in gasoline prices illustrated in Figure 1. As gasoline prices fluctuate over time, the bars in Figure 2 extend and contract proportionally. A given change in gas prices has a larger effect on the level of G for newer and lower-MPG vehicles, and we test for whether relative prices move correspondingly.

Equation (7) consistently estimates γ if $E[G\varepsilon|\tau,\psi] = 0$. This actually contains two different economically-meaningful identifying assumptions. First, we assume that gasoline costs are uncorrelated with market shares:

Assumption 1. $E\left[G \cdot \left(-\frac{1-\sigma}{\eta}\ln s_{ja} - \frac{\sigma}{\eta}\ln s_n\right)\right] = 0.$

We relax this assumption in an alternative specification to be described momentarily. The assumption is more likely to hold for used vehicles than for new vehicles, because while scrappage may respond slightly to changes in gasoline prices (Li, Timmins, and von Haefen 2009), by definition, no additional used vehicles can be produced. For this reason, we exclude new vehicles from the base specification.

Assumption 2 is that changes in discounted gasoline costs are not correlated with changes in preferences and unobserved characteristics over the different model years within a generation of a vehicle:

Assumption 2. $E[G\xi] = 0.$

Later in the paper, we discuss and test several potential violations of Assumption 2 and document that they do not appear to substantially affect the results.

All regressions use the number of transactions as analytic weights to reflect the fact that p_{jat} is an average representing a number of separate observed transactions.

4.1.1 Measurement Error and the Grouping Estimator

Measurement error in G would bias $\widehat{\gamma}$ towards 0, making it appear as if consumers undervalue gasoline costs. In our base specification, we address this with a standard approach called a grouping estimator. The grouping estimator places observations into mutually exclusive and exhaustive groups and instruments for all right-hand-side variables with their group-level averages. Intuitively, by aggregating over observations to the group level, the instruments average out observation-level measurement error. The grouping estimator was introduced by Wald (1940), and Ashenfelter (1984), Angrist (1991), and others have used it to estimate labor supply and other models.

How are the groups defined? Grouping observations into fewer larger groups allows us to average over measurement error that is more severe or correlated across observations. Fewer groups gives larger standard errors, however, and there is a lower bound to the number of groups: there must be at least as many groups as instrumented right-hand-side variables for the regression to be identified. In addition, the groups must be defined so that there is variation between groups in the instrumented variables in order for the regression to be identified. Our base specification groups observations into two fuel economy quantiles and time, giving two groups for each t, or a total of 222 groups. We will show how grouping at less disaggregated levels affects the parameter estimates.

4.2 Endogenous Shares: Nested Logit

Assumption 1 for our base specification was that market shares are uncorrelated with gasoline prices. Especially for new vehicles, this assumption may not hold. In this section, we describe an endogenous quantity estimating equation, which simply adds time dummies and fixed effects to the market-level price-quantity relationship from Equation (5). Put differently, this specification adds the market share terms back into the base specification, Equation (7):

$$p_{jat} = -\gamma G_{jat} - \frac{1-\sigma}{\eta} \ln s_{jat} - \frac{\sigma}{\eta} \ln s_{n,jat} + \tau_t + \mu_{at} + \psi_{ja} + \xi_{jat} \tag{8}$$

Even conditional on our fixed effects, Equation (8) suffers from the usual simultaneity bias: $E[\xi s] \neq 0$. In words, the model year-specific unobservable characteristic ξ_{jat} could still be correlated with market shares if, for example, a feature that is specific to particular model year affects both price and market share. To address this, we need an instrument that generates variation in market shares that is uncorrelated with unobserved quality.

Our instrument exploits the fact that new vehicle market shares respond to gasoline prices. In particular, in years when gasoline prices are high, more high fuel economy vehicles are sold. Figure 3 illustrates the variation that identifies our instrument. As gasoline prices rose between 2004 and 2007, sales of vehicles with fuel economy greater than or equal to 20 MPG increased from 6.5 to 7.7 million, while sales of vehicles rated less than 20 MPG dropped from 8.1 to 6.6 million. This difference in quantities then persists over time. For example, there are more two-year-old high-MPG vehicles on the road in 2008 than in 2006. This difference in quantities is not due to changes in quality ξ; instead, it is due to the different gasoline price forecasts at the time when the vehicles were new.

Our instrument for the market shares of used vehicles is thus $G_{0,jat}$, the predicted lifetime gasoline costs of model j at the time when it was new. More precisely, because a given model year is typically manufactured between September of the year before the model year through August of the model year, we use the average gasoline price forecast for those twelve months to construct $G_{0,jat}$. This instrument acts conditional on the model year dummy variables μ in Equation (8), meaning that vehicles that have high values of $G_{0,jat}$ relative to other vehicles produced in the same year are expected to have lower sales.

In this alternative specification, we therefore allow the market share of vehicle ja to be correlated both with gas costs G and unobservable characteristics ξ. Assumption 1 is now less strong, and it has two parts:

Assumption 1A. $E[G_0\xi] = 0$

Assumption 1B. $E[s_n\xi] = 0$.

Assumption 1B is that nest-level market share is uncorrelated with the vehicle's unobserved characteristic. Mathematically, it is true that higher values of ξ_{jat} cause more sales of vehicle ja, and because vehicle ja is itself a member of nest n, increases in s mechanically imply increases in s_n. This mechanical source of bias is very small, however, as there are a large number of model-by-age combinations in each nest.

Why do we not label this as our "base specification"? The reason is that we cannot implement both the nested logit IV and the grouping estimator IV in the same estimation. In principle, it is indeed possible to estimate Equation (8) using a grouping estimator, while instrumenting for $\ln s$ with the group-level average of G_0. But because the grouping estimator involves a large number of instruments, of which only G_0 has effectively any correlation with $\ln s$, this very naturally generates a weak instruments problem in fitting $\ln s$. As we present results from both the grouping estimator and the nested logit, we will see that measurement error affects the results much more than endogenous quantities, and we therefore label the grouping estimator as our "base specification."

As Figure 3 suggests, much of the variation in the instrument affects model years 2004 and later. Thus, IV parameter estimates are "local" to fixed effect groups with at least some observations of post-2004 model year vehicles. Including additional observations where an instrument has no variation does not change the Local Average Treatment Effect but does reduce the power of the instrument. To maintain sufficient power by the standards of Stock and Yogo (2005), we therefore follow the logic of Lewis (2011) and restrict our nested logit IV sample to include only observations of new and used vehicles beginning in 2004.

5 Results

5.1 Graphical

We begin with graphs that illustrate our data, identification, and results. Figure 4 illustrates a crucial feature of the data: low-MPG vehicles are much more expensive than high-MPG vehicles. The figure shows the results of a regression of Manufacturer's Suggested Retail Price on a set of MPG dummies, for all models in the Ward's data, ranging from model year 1984 to 2008. There are not many vehicles rated lower than 14 MPG, and and many of these are high-priced sports cars. Similarly, there are not many vehicles more than 38 MPG, and these are often higher-cost hybrid vehicles. Between 15 and 30 MPG, which includes nearly all of our dataset, MSRPs decline from $30,000 to $11,000. This figure both corroborates concerns about the traditional cross-sectional approaches to estimating γ and, as we shall see, has important implications for our panel approach.

Figure 5 illustrates average transaction prices for new vehicles, used vehicles in the JDPA and Manheim datasets, and the Manheim used prices demeaned within each ja group. Notice that prices have systematic seasonal patterns, and in particular that used vehicle prices decrease within each year. Because this seasonal depreciation is approximately a percentage of price and lower-MPG vehicles have higher prices, these patterns affect the prices of low-MPG vehicles more than high-MPG vehicles in levels. Since gas price expectations rise on average during the study period, failing to account for within-year trends that affect low- vs. high-MPG vehicles differently could cause us to mis-attribute seasonality to adjustments driven by gas cost changes. In our base specification, we therefore use fixed effects at the level of model by age (in months), not model by age (in years). For example, one ja fixed effect group is the five observations of the four-year-old sixth-generation Honda Civic DX in the month of April for each of the years 2000 through 2004. Notice that this is why the demeaned used vehicle prices in Figure 5 have no seasonality.

Trends in seasonally-adjusted average prices could be manifestations of underlying market shifts that affect low-MPG vs. high-MPG vehicles differently. If correlated with changes in gasoline prices, such shifts could bias our estimator. The double black line in Figure 5 shows that de-meaned within-ja prices declined somewhat from 1999-2003. The decline from 1999-2002 appears to be driven by a decrease in the absolute price level for low-MPG vehicles. This may have been precipitated by an increase in the number of new low-MPG vehicles sold during the middle and late 1990s, which appears to have caused their prices to drop as they became available for resale as used vehicles. Manheim's analysts attribute the nadir in 2003 to the 2001 economic slowdown, during which both low- and high-MPG vehicles were offered at attractive lease terms and an unusually large share were leased instead of sold. When this larger volume of vehicles came off of lease two years later, this depressed resale prices.

Average seasonally-adjusted prices remained relatively steady from 2004-2007 before dropping sharply as the 2008 recession took hold. It is possible that this recession differentially affected low- vs. high-MPG vehicles, for example by differentially affecting people of different income levels, who tend to buy different types of vehicles. We therefore eliminate from our base specification all data beginning with April 2008, when the recession began to affect vehicle market prices. Careful readers will recall that we made this decision beginning with our first working paper version of this project (Allcott and Wozny 2009), and we still agree that this is sensible. Similarly, we would certainly not want to include data from 2009, when the Cash for Clunkers stimulus program significantly changed used vehicle markets, differentially affecting prices of low-MPG vs. high-MPG vehicles. However, we will clearly document in Section 5.3.2 how the choice of time period affects the estimated γ.

Figure 6 moves from describing average price levels to illustrating our identification. On the vertical axis is the difference between the mean transaction price for vehicles with below-median MPG versus above-median MPG. On the horizontal axis is the difference in mean G. This is raw data, unadulterated by fixed effects, time controls, or other manipulations. Remarkably, even in this raw data, it is starkly visible that relative prices and relative gas costs are negatively correlated. The slope of a best fit line would be -1 if $\gamma = 1$ and if fixed effects are uncorrelated with G. In fact, the slopes are -0.86, -0.79, and -1.17 for 1999-2003, 2004-March 2008, and April-December 2008, respectively.

Figure 7 takes the final step between Figure 6 and our base specification. The solid blue line is the difference between average de-meaned G for below- minus above-median MPG vehicles over the months of the study period. The dotted blue line is the same as the solid blue line, except that it is based on the assumption that retail gas price forecasts are consistent with a martingale, not with oil futures prices; we return to this issue later in this section. The double black line is the difference between average de-meaned transaction price for above- minus below-median MPG vehicles. If $\gamma = 1$, the blue and black lines should move in parallel. From this figure, it is again clear that relative prices are responsive to relative gasoline costs.

5.2 Base Specification

Figure 7 exactly represents our base specification, grouping at the level of month by above- vs. below-median MPG, except that the figure also includes data from April-December 2008, which might be confounded by the recession. As shown in "Row 0" of Table 5, the base specification $\widehat{\gamma}$ is 0.72.

5.2.1 Standard Errors

Throughout Tables 5, 6, and 7, we report Huber-White robust standard errors, clustered at the level of model j by age (in years). This standard error is unbiased in the presence of serial correlation over time in the price of, for example, a three-year-old Honda Civic DX sedan. The base specification standard error on $\widehat{\gamma}$ is 0.048. Clustering at the level of "nameplate" by age (in years) gives a standard error is 0.058. This standard error is unbiased in the presence of serial correlation over time in the price of, for example, all three-year-old models of Honda Civic. Clustering at the level of model j gives a standard error of 0.090. Analogously, this is unbiased in the presence of serial correlation in the price of all Honda Civic DX sedans of any age.

These standard errors do not account for the fact that G is a generated regressor estimated from first-step regressions that predict each observation's vehicle-miles traveled and survival probability using ancillary data. Murphy and Topel (1985) show that the true covariance matrix is additively separable in the usual covariance matrix and an additional matrix that accounts for the uncertainty in the first-step parameter estimates. We estimate this additional variance by bootstrapping draws from the estimated distribution of first-step parameters and estimating $\widehat{\gamma}$ for each draw. For the base specification, the standard deviation of this set of $\widehat{\gamma}$ estimates is 0.023. Adding this variance to the square of our estimated IV standard error of 0.048 gives a Murphy-Topel standard error of 0.053. While the Murphy-Topel standard errors are of course larger, the difference is quite small, and their use does not affect the conclusion that $\gamma < 1$ in any of our specifications.

5.3 Robustness Checks

5.3.1 Level of Grouping

Correcting for measurement error appears to be extremely important. Rows 1 through 6 of Table 5 present the results of increasingly disaggregated grouping estimators, while Row 7 is the ungrouped OLS estimator. As we reduce the aggregation, the estimated values of γ drop steadily to the OLS estimate of 0.53. This suggests that measurement error significantly biases the OLS estimates, and even some of the relatively disaggregated grouping estimators. This pattern can arise either if the variance of the measurement error is quite large or if it is correlated across observations.

Specifications 11-14 follow the base specification in grouping all ages together and grouping MPG in two quantiles, but they additionally group time periods with increasing aggregation, from two to six months per group. As a concrete example, Specification 14 has 38 independent groups: 19 six-month periods by two MPG bins per period. The estimated values of γ remain almost unchanged from the base specification. The fact that increasing the level of aggregation beyond that in the base specification is consistent with the idea that the base level of aggregation has addressed as much measurement error as possible. This is the reason why we label this level of aggregation as our "base specification": it has the most disaggregated grouping that appears to fully address measurement error.

Figures 8 through 11 are partial regression plots of the second stage of the grouping estimators in Table 5 Rows 7, 5, 0 and 12, respectively. Each point reflects the group average of p_{jat} and G_{jat} after partialing out fixed effects and time dummies. For ease of viewing, the figures include the group average of p_{jat}, while the regression itself of course has variation in p_{jat} within a group; this mathematically does not affect the estimated coefficients or the best fit line. The best fit line is the estimated slope of the relationship between p and G conditional on the fixed effects and time dummies, which corresponds to $-\widehat{\gamma}$ for the respective specification in Table 5. Notice that the more aggregated specifications have less dispersion in group-level average G, showing intuitively how increased aggregation eliminates both measurement error and potentially-useful variation in G.

Although the scales of the graphs change, the horizontal and vertical axis are magnified proportionally. This makes it easier to see that as we increase the level of aggregation over the first three figures, the slope of the best fit line increases in absolute value from the OLS estimate of -0.53 to -0.63 to the base specification slope of -0.72. As shown in Figure 11, however, grouping over four-month periods does not further steepen the slope.

5.3.2 Alternative Time Periods

As illustrated earlier, vehicle markets changed over the study period. How sensitive is $\widehat{\gamma}$ to the choice of time periods? Specifications 21 and 22 of Table 5 repeat the base specification for different samples of the data: 2004 through March 2008 and the entire available data from 1999 through the end of 2008. In both cases, $\widehat{\gamma}$ is closer to one than the 1999-March 2008 base specification sample, although both are still statistically less than one. In Specification 21, the fact that excluding the early period does not significantly change the result mitigates our earlier concern that market-level trends during 1999-2003 could have biased our estimator.

There are two explanations for the fact that including the latter part of 2008 moves the coefficient closer to one. The first is that, per our concerns discussed earlier, the recession causes market forces that bias this estimated γ. The second potential explanation is that γ truly did increase over this period as gas prices fluctuated dramatically. This would be consistent with other work that models "extraordinary" events about which information diffuses instantly (Reis 2006) or that cause consumers to update beliefs between coarse categories, for example from gas costs being "inconsequential" to gas costs being "high" (Mullainathan 2002).

In summary, although one should be aware of potential concerns from market-level shocks that differentially affect low- vs. high-MPG vehicles, this issue does not affect the qualitative conclusion that $\gamma < 1$. However, including "questionable" data from the beginning of the recession in April through December 2008 does make the qualitative conclusion less certain. Of course, including this "questionable" time period does not affect the conclusion that consumers undervalued gas costs before April 2008.

5.3.3 Alternative Discount Rates

While our calculations show that a six percent discount rate appears to to be most sensible, reasonable people may disagree. Specifications 31-33 in Table 5 show the sensitivity of $\widehat{\gamma}$ to the assumed discount rate, ranging from 0.65 at three percent to 0.97 at 15 percent. At discount rates of larger than about 13 percent, we fail to reject $\gamma = 1$ with 90 percent confidence using the Murphy-Topel standard errors. Our "implied discount rate," the discount rate that "rationalizes" the data by giving $\gamma = 1$, is 16 percent.

5.3.4 Endogenous VMT

In the Appendix, we re-derive our estimating equation while allowing vehicle-miles traveled to be an endogenous function of gasoline prices. In summary, this alternative specification takes an assumed elasticity of VMT to gasoline prices, modifies G_{jat} to allow VMT to be based on the forecasted future gasoline prices at time t, and introduces an additional term that captures the fact that usage utility also varies as a function of the number of miles driven per year. Table 5 Rows 41-44 present the set of $\widehat{\gamma}$ estimates from this procedure under different assumptions. The point estimate of γ hardly changes from the base specification with an elasticity of 0.2, which is larger than or equal to recent empirical estimates by Hughes, Knittel, and Sperling (2007), Small and Van Dender (2007), and Gillingham (2010). Even with an elasticity of 0.5, $\widehat{\gamma}$ does not change significantly, and these results are not sensitive to the assumption of linear vs. constant elasticity VMT demand. Intuitively, the reason for this is that the Envelope Theorem approximately holds: when gasoline prices increase and consumers choose to drive less, the resulting decrease in gas costs is roughly equal to the decrease in the monetary value of usage utility.

5.3.5 Changes in Characteristics

Assumption 2 would be violated if a model's characteristics change in ways that are correlated with G. For example, since gasoline prices rise on average during the study period, if characteristics improve more over model years for high-MPG vs. low-MPG models, we would mis-attribute these vehicles' increased desirability to changes in gasoline prices. As discussed in the Data section, redefining a new generation of a model as a new model j addresses the bulk of these concerns. Even within a generation, however, there can be some small variation in observable characteristics; see, for example, Knittel (2011, figure 1).

An additional suggestive test is to to add controls for observable characteristics to the estimation. Row 51 is the estimate of Equation (7) while defining separate groups by integers of vehicle age and by 20 MPG quantiles. This increased disaggregation relative to the base specification is necessary in order to retain sufficient variation to estimate coefficients on observable characteristics. Rows 52 through 55 add progressively more controls for observable characteristics: horsepower, curb weight, wheelbase, Anti-Lock Brakes, Stability Control, and Traction Control. The point estimates move slightly away from unity, although they are not statistically different. Although whether controlling for observables affects a parameter estimate does not directly tell us whether controlling for unobservables would affect the estimate, it is plausible that changes in observable and unobservable characteristics are correlated, by logic similar to that of Altonji, Elder, and Taber (2005). This is therefore suggestive evidence that changes in unobservable characteristics do not bias our estimates.

5.3.6 Changes in Preferences

In addition to representing unobserved vehicle characteristics, variation in the utility function parameter ξ_{jat} over time can also represent changes in consumer preferences for a vehicle with the same characteristics. Since gasoline prices rise on average during the study period, differential time trends in preferences for low- vs. high-MPG vehicles would violate Assumption 2 and bias $\widehat{\gamma}$.

One potential concern is that consumers became increasingly "green," or environmentally-oriented, over the study period, resulting in increased preference for high fuel economy vehicles independent of the financial savings. This would bias $\widehat{\gamma}$ upward, as it would increase the prices of high-MPG vehicles over the study period as gasoline prices rose. To test this, we exclude hybrids and the top 60 most "green" vehicles from a recent ranking (Yahoo 2009). Rows 61 and 62 of Table 5 show that this does not affect the results.

An opposite concern is that preferences for particular classes of large vehicles, for example SUVs or pickups, strengthened over the study period. Analogously, this would bias $\widehat{\gamma}$ downward, as it would increase the prices of low-MPG vehicles over time, attenuating the decrease in relative prices that the model would expect as gasoline prices rose. To test this, Specifications 63 through 65 exclude SUVs, minivans, and all cars, respectively, from the estimation. Interestingly, eliminating SUVs from the estimation actually decreases $\widehat{\gamma}$. The estimated γ is larger when considering only trucks in line 64, although this is not statistically different, and the coefficient is identical when minivans are excluded.

5.3.7 Retail Price Data

The wholesale market in the Manheim auctions is of course directly connected to the retail market. Auto dealerships, many of which are surveyed by JD Power, are the primary auction buyers, and many of the vehicles observed in the Manheim transaction data are transacted shortly thereafter in the JDPA dealership data. If changes in G_{jat} are passed through from wholesale to retail prices in levels, the use of wholesale vs. retail data should not affect the estimated γ. As an additional specification check, however, we estimate our base specification with used vehicle retail transaction prices from the JD Power data.

To do this, we first find the sample of jat observations that are common to both datasets. As shown in Row 71, $\widehat{\gamma}$ estimated from Manheim data in the common sample is 0.76, slightly larger than the point estimate for the base specification sample. Using the JDPA data, $\widehat{\gamma} = 0.53$. The $\widehat{\gamma}$ estimated with JDPA data does not change much when restricting to the common sample because there are so few jat observations in JDPA that are not observed in Manheim.

Could this be because the retail market did not experience the same time trends as the wholesale market? Figure 5 suggests that mean price trends in JDPA and Manheim data were similar for 1999-2003, 2004-March 2008, and April-December 2008. To confirm this, we include Specifications 74-77, which are analogous to Specifications 21 and 22 in their analysis of different time periods. Specifications 74 and 75 show that the coefficients estimated with the Manheim wholesale data in the common sample are closer to one than the base specification, as we had seen with the full Manheim data in Specifications 21 and 22. Specifications 76 and 77 show that the coefficients also move closer to one in the JDPA retail data, although excluding 1999-2003 makes somewhat more of a difference in JDPA than in Manheim. In summary, the JD Power $\widehat{\gamma}$ is farther from unity, which strengthens the case that the market undervalues gas costs, and this result is insensitive to the time period analyzed.

5.3.8 Fixed Effects

Why do we use model-by-age fixed effects instead of model-by-model year fixed effects? The reason is that the latter approach suffers from a mechanical and particularly severe endogeneity problem. To see this, examine Figure 12, which illustrates average depreciation patterns for model year 1999 and 2002 vehicles with above-median and below-median MPG. Depreciation for vehicles of a given age scales roughly as a percentage of price: more expensive vehicles depreciate more in levels. Notice in Figure 12 that because low-MPG vehicles cost more than high-MPG vehicles, they depreciate faster in levels.

Consider now what happens to price, the left hand side variable in our regressions, after conditioning on model-by-model year fixed effects. Prices residual of model-by-model year fixed effects drop over time for all vehicles, and they drop faster for low-MPG vehicles. Meanwhile, although there was substantial fluctuation, gasoline prices tended to rise over the study period between 1999 and 2008. As a result, the estimator with model-by-model year fixed effects falsely attributes conditional decreases in the relative prices of low-MPG vehicles to the increase in gasoline prices over the study period, biasing $\widehat{\gamma}$ away from zero.

Roughly how severe might this problem be? Consider the years 2000 and 2001, a pair of years in which average gasoline prices were close to the same. Across all used vehicle auction transactions in 2001, the average prices of below- and above-median MPG vehicles were \$12,800 and \$8400, respectively. The average levels of depreciation between 2000 and 2001 for these two MPG groups were \$3380 and \$1950, respectively. This means that even over a year when gasoline prices did not increase, low-MPG vehicles lost about \$1400 more value than high-MPG vehicles. Using model-by-model year fixed effects, this substantial differential depreciation would be falsely attributed to the increase in gas prices over the years of the study period.

There are two ways to address the differential depreciation patterns of low-MPG vs. high-MPG vehicles. The first is to use model-by-age fixed effects. The second is to use model-by-model year fixed effects and include dummy variables for age interacted with MPG group. We chose the former approach as our base specification because the latter requires us to disaggregate the grouping estimator groups by age to identify the controls. As we saw earlier, the problem with this increased disaggregation is that it attenuates our estimated γ. For example, Rows 0 and 1 in Table 5 show that grouping into age groups four years wide appears to attenuate $\widehat{\gamma}$ from 0.72 to 0.66.

Empirically, the bias from failing to control for differential depreciation appears to be large. In Row 81, we implement our base specification with three changes: we use model-by-model year fixed effects, group at the level of time by four year age groups by above- and below-median MPG, and add dummy variable for each age group. This specification, like all others in this paper, includes a full set of time dummies τ_t. The coefficient is 1.61, suggesting substantial overvaluation. In Row 82, we replace the age dummy variables with age group-by-MPG group dummies. The estimated γ drops to 0.59. The estimation in Busse, Knittel, and Zettelmeyer (2011) is analogous to our Row 81: they similarly have model-by-model year fixed effects with age dummies and time controls, but they do not include the age group-by-MPG group dummies. This appears to bias their estimates toward finding overvaluation, which is part of what explains their finding of a slightly *negative* implied discount rate.[14]

Row 1 in Table 5 has the specification with model-by-age fixed effects that is equivalent to Row 82 in the sense of having the same level of grouping. Importantly, notice that the $\widehat{\gamma}$ in Row 82, which is 0.59, is within the 90 percent confidence interval of the $\widehat{\gamma}$ in Row 1, which is 0.66. This means that when grouped at the same level of aggregation in the grouping estimator, our two approaches give comparable results.

In summary, we have shown two approaches to consistently estimate γ in the presence of depreciation patterns that vary systematically by fuel economy. We label the model-by-age fixed effects as our "base specification," however, because this allows us to omit age-by-MPG group dummies and group at a more aggregate level, which eliminates bias from measurement error.[15]

[14] We calculate that the γ implied by BKZ's negative 0.3 percent estimated discount rate is 1.15, not the 1.61 we estimate in Row 81. Several other differences between Row 81 and BKZ's specifications are also important. First, BKZ assume that consumers forecast that gasoline prices are a martingale, while we assume that consumers forecast that gasoline prices will follow oil futures. In the next section, we provide some evidence to substantiate our assumption and show that the martingale assumption moves γ downward substantially, by 0.22. Second, BKZ use JDPA data, which as we documented in Section 5.3.7 can give different estimates compared to the Manheim data. Third, BKZ's study period ends June 30th, 2008, while we include only through the end of March. Fourth, BKZ use data from the transaction level, not the jat level, and add a series of transaction-level controls. This could generate other differences and also makes it impossible for us to exactly replicate BKZ.

[15] In Section 5.3.5, we provided evidence that changes in vehicle characteristics between generations that could be correlated with gasoline prices do not appear to bias our estimated γ. In Section 5.5, we will show that potential changes in market shares correlated with gasoline prices also do not substantially bias our estimates. These empirical results help to support the validity of our Assumptions 1 and 2.

Some readers have suggested, however, that model-by-model year fixed effects (presumably with the necessary age-by-MPG group controls) would fully address these two concerns. Superficially, this seems appealing: if a change in the quality or market share of, say, a model year 1999 Honda Civic increases its price by $100, can't a model-by-model year fixed effect exactly soak up that $100 difference?

Unfortunately, the fact that vehicles depreciate means that the model-by-model year fixed effect does not fully address the issue. Consider again the example model year 1999 Honda Civic. As it depreciates over the years of the study period, that incremental $100 similarly depreciates. Thus, a model-by-model year fixed effect undercorrects for the unobserved change early in the study period and overcorrects late in the period. Thus, if quality or market shares differentially change for low-MPG vs. high-MPG models as gasoline prices rise over the study period, gas costs G are still correlated with this error conditional on model-by-model year fixed effects.

In practice, the performance of the model-by-model year fixed effect depends on how much the hypothetical $100 depreciates over the period that the car is observed. If it does not depreciate at all, then the model-by-model year fixed effect entirely solves the potential endogeneity problem. If it depreciates fully, on the other hand, then the model-by-model year fixed effect is less helpful. Figure 12 gives empirical examples of depreciation observed in our dataset. The solid lines are the extreme cases in our data: model year 1999 cars and trucks observed for the entire ten years. They lose 82 to 85 percent of their value over the study period. On average in our data, a model of a particular model year is observed for 6.7 years, and it loses 73 percent of its initial value over that period.

The takeaway here is that regardless of the fixed effects that we use, we still must maintain our Assumptions 1 and 2 and test their robustness to the same classes of concerns.

5.4 Martingale versus Futures

In our base specification, we formulate G under the assumption that auto consumers' gas price forecasts are consistent with oil futures markets. Table 6 illustrates why this is an important assumption. All specifications in the table use the grouping estimator at the same level of aggregation as the base specification. Column 1 reproduces the base specification, while Column 2 replaces G with G_m, the gas cost variable constructed under the assumption that gas prices are a martingale. While the base specification gives $\widehat{\gamma} = 0.72$, the martingale specification in Column (2) gives $\widehat{\gamma} = 0.50$, much further from unity. The reason for this is that as current gas prices rose over the study period, the oil futures market expected prices to eventually revert downwards. Relative to the martingale assumption, using futures thus attenuates the variation in G as gas prices rise, requiring smaller changes in prices to rationalize this variation and giving a larger $\widehat{\gamma}$.

Which specification is more reasonable? Should we model that the futures market more closely reflects consumers' beliefs, perhaps because the media often covers oil markets and experts' predictions of future prices? Or should we instead assume that consumers only see the current price at the gas station and believe prices are a martingale? One approach to answering this question is to elicit consumers' beliefs through a survey, such as the Vehicle Ownership and Alternatives Survey (Allcott 2010, 2011) or the Michigan Survey of Consumers (Anderson, Kellogg, Sallee, and Curtin 2011). A revealed preference approach, however, is to test whether vehicle prices tend to move more closely with G or G_m. More specifically, we can test which assumption about gasoline price forecasts explains more of the variation in vehicle prices.

As we saw in Figure 2, although oil futures prices move with retail gasoline prices, there is substantial month-to-month variation in current retail prices that the oil futures market believes is transitory. This means that there is variation in G_m conditional on G. Figure 7 illustrates our empirical test and foreshadows two results. First, vehicle prices appear to track more closely with futures prices, as they do not move fully with the larger swings in current gas prices. Second, however, the market does appear to respond somewhat to current gas prices in addition to futures prices, especially in the early part of the study period.

To see this formally, refer again to Table 6. The reported "Partial R^2" is the R^2 of the fit of price on gas cost, after both variables are conditioned on ja fixed effects and time dummies and collapsed to the average within each time by MPG quantile group. By construction, the partial R^2 reported in Column 1 is the same as the R^2 of the base specification partial regression plot in Figure 10. The first key result is that the partial R^2 is higher for Column 1 than Column 2: oil futures forecasts explain more of the variation in relative vehicle prices than current retail prices. Aside from being of interest *per se* in understanding how consumers forecast energy prices, this also helps to validate the assumption made to construct G for our base specification.

Although this partial R^2 "best fit test" implies that it may be more accurate to base G on oil futures instead of current retail gas prices, it is also possible that the market responds to variation in current gas prices in addition to futures. Column 3 tests this by regressing prices on both G and $[G_m - G]$. The results show that vehicle markets do respond significantly to current gasoline price changes that are not reflected in futures markets.

5.5 Nested Logit IV

The base specification and robustness checks in Table 5 are unbiased under Assumption 1, which was that the market share terms were uncorrelated with gas cost and could therefore be excluded. In this section, we detail the results of the alternative specification that allows for endogenous quantities. The results suggest that Assumption 1 turns out to be reasonable in this application.

Column 3 presents the results of estimating Equation (8) using the sample of new and used vehicles from 2004 through March 2008. The lower half of the table details the first stage estimates. As should be expected, the instrument G_0 is negatively correlated with market share, conditional on the vehicle fixed effects and the model year dummies. The coefficient estimate implies that a 1000 dollar increase in lifetime gas costs relative to other vehicles of the same model year reduces a vehicle's market share by 0.0794 percent. For example, the average gas price forecast for September 2004 through August 2005, the period when most 2005 model year vehicles were manufactured, was \$2.03 per gallon, while the gasoline price forecast for model year 2006 vehicles was \$2.40. Between those two years, the value of G_0 for a Ford F-150 pickup went from \$16,057 to \$19,362, while it went from \$6865 to \$8204 for a Honda Civic sedan. As a result, the Civic market share is predicted to decrease by 16 percent relative to the F-150.

The first stage Cragg and Donald (1993) statistic is 23.0, meaning that the estimates do not suffer from a weak instruments problem (Stock and Yogo 2005). Including earlier years, however, reduces the Cragg-Donald statistic to a point that suggests that weak instruments could bias the results. This is not surprising, since there is less gasoline price variation, and thus less variation in the instrument, for observations before 2004.

Turn now to the to the estimated IV coefficients on $\ln s$ and $\ln s_n$ in Column 3 of the upper half of Table 7. The $\ln s$ coefficient implies that if the quantity available of a given new or used vehicle increased by ten percent, its equilibrium price would decrease by \$242.70. Conversely, this also implies that if the price of a given vehicle were increased by \$1000, its market share would decrease by 39 percent. Because we define a vehicle at a finely disaggregated submodel level, there are many close substitutes for a given "vehicle," so a high degree of price sensitivity is expected. The $\ln s_n$ coefficient in Column 3 is -3380, meaning that a vehicle's price drops by \$338 if the quantity of substitute vehicles in the same nest increases by ten percent. Recall that the σ parameter in the nested logit model ranges from zero to one, reflecting the degree of correlation in consumers' tastes for vehicles in the same nest. Using the algebraic definitions of the coefficients in Equation (8), it is easy to calculate that the coefficients on $\ln s_{jat}$ and $\ln s_{n,jat}$ give a $\widehat{\sigma}$ of 0.58.

As discussed in the estimation section, these specifications do not use the grouping estimator, as this produces a weak instruments problem. As a result, the estimated $\widehat{\gamma}$ is attenuated toward zero due to measurement error. Therefore, although we have discussed the magnitudes of the estimated coefficients in Column 3, the more important part of Table 7 is actually how the estimated values of γ compare across columns.

In particular, now consider Column 1. This is the OLS estimate of Equation (8), except without $\ln s$ or $\ln s_n$ on the right-hand-side, for the same sample: new and used vehicles for 2004 through March 2008. Comparing this to Column 3 shows that endogenizing market shares changes $\widehat{\gamma}$ from 0.40 to 0.36. These point estimates are not statistically different, which indicates that relaxing the exogenous shares assumption, while theoretically important, does not make a substantive difference in the results.

Why is this the case? Recall that Equation (5) shows that there are two ways that shares affect prices of vehicle ja, through $\ln s$ and through $\ln s_n$. The "own share effect" is that more availability of vehicle ja reduces its equilibrium price. The "substitute share effect" is that more availability of close substitutes in the same nest reduces the equilibrium price. Stated precisely, Assumption 1 is that the sum of these two effects is uncorrelated with G, and thus does not affect the estimated γ.

The difference between Columns 1 and 2 illustrates the own share effect: when we include $\ln s$ in the estimation but exclude $\ln s_n$, the point estimate of γ moves toward one. This is to be expected: when gas prices increase, the market share of new low-MPG vehicles decreases. Because demand is downward-sloping, the marginal consumer of a low-MPG vehicle now has higher willingness to pay for the vehicle. Thus, while its price decreases relative to a high-MPG vehicle, it does not decrease as much as it would have if the same consumer were on the margin before and after the change. The exogenous shares model that does not allow for this effect therefore needs a larger change in relative prices to rationalize a given change in relative gas costs, and the resulting estimate of γ is biased slightly toward zero.

The difference between Columns 2 and 3 illustrates the second effect: when $\ln s_n$ is included, the estimated γ moves away from unity. The explanation for this appears to be the increase in the market shares of low-MPG vehicles during the mid-to-late 1990s. As the shares of these vehicles in resale markets increased over the study period, their prices dropped. The estimator that excludes $\ln s_n$ falsely attributes this decrease in prices to the increase in gasoline prices over the study period, moving $\widehat{\gamma}$ toward one. Allowing for the substitute share effect by including $\ln s_n$ moves $\widehat{\gamma}$ away from one. When the two share effects are combined, on net it happens that endogenizing market shares does not significantly change the estimated γ.

The nested logit model, including our specific choice of nests, imposes a particular structure on substitution patterns which might or might not be a good approximation of reality. Recall, however, that the nested logit model is mathematically equivalent to allowing random coefficients on nest indicator variables. This means that different nest structures can approximate any number of relatively flexible substitution patterns. If the estimated γ is robust across these different nest structures, this suggests that our *a priori* assumption about substitution patterns is not driving our results.

We have experimented with many different nest structures, and $\widehat{\gamma}$ proves to be quite insensitive. Columns 4 and 5 of Table 7 present two example alternative nest structures; $\widehat{\gamma}$ and the $\ln s$ coefficient are both essentially unchanged. We also constructed eleven other nest structures, using various combinations of bins of horsepower, weight, and fuel economy, as well as luxury and class indicators. In only one of these eleven other specifications did the estimated γ fall outside the 90 percent confidence interval of the $\widehat{\gamma}$ under exogenous shares in Column 1. While we do not include all of these specifications to conserve space, we are happy to provide the results upon request, or to experiment with other nest structures.

Column 6 shows the importance of instrumenting for market shares: if we estimate the nested logit model in Equation (8) in OLS without the G_0 instrument, the coefficient on $\ln s$ is positive. This is the usual apparently upward-sloping demand symptomatic of simultaneity bias, the correlation of equilibrium market share with the unobserved demand shifter ξ.

To recapitulate, it would ideally be possible to both allow endogenous market shares and address measurement error in the same specification. As we have seen, however, endogenizing market shares does not significantly change the estimated γ, while measurement error appears to have a large effect. It is for this reason that we are comfortable labeling the grouping estimator as our "base specification" and maintaining Assumption 1. However, for future empirical studies that include used vehicle markets during periods when there is variation in gas price forecasts, this instrument could be an appealing alternative to the Berry, Levinsohn, and Pakes (1995) instruments, which depend on assumptions about the nature of the supply-side price setting game.

5.6 Magnitude of Mispricing

One way of interpreting the magnitude of an estimated γ is to calculate the dollar value of the misadjustment in relative prices that occurs when gas prices change. Consider a hypothetical set of used vehicles that have different fuel economy ratings but would all be driven 12,000 miles per year for the remaining seven years of their lifetimes. The blue "Predicted Price Change" line in Figure 13 illustrates the change in the relative prices of these vehicles in response to a permanent \$1 increase in gas prices, with $\widehat{\gamma} = 0.72$ from the base specification. For example, relative to a 25 MPG vehicle such as a Honda Accord with a 2.4 liter engine, a 21 MPG vehicle decreases in price by \$485. If $\gamma = 1$, however, the relative price should decrease by \$670, or \$185 more than it does. As another example, the price of a 30 MPG vehicle such as a Honda Fit should increase by \$2,230 relative to a 15 MPG vehicle such as a Ford F-150 pickup. The base specification results suggest that the relative price adjusts by 72 percent of that amount, or \$1610. The dollar values of the apparent mispricing are certainly not trivial.

6 Welfare Analysis

For this section only, we set aside the empirical discussion and assume that it is indeed true that consumers are misoptimizing by at least some amount when they purchase vehicles. When consumers misoptimize, their willingness to pay no longer measures their true welfare, so traditional approaches to welfare analysis no longer apply (Bernheim and Rangel 2009). This section formalizes a new approach to behavioral welfare analysis in a discrete choice setting and simulates the gains from a corrective policy.

6.1 Theory

Following the language of Kahneman (1994), we distinguish between *decision utility*, the utility function that consumers act as if they are maximizing at the time of choice, and *experienced utility*, the actual utility that consumers are expected to realize as a result of the choice. For rational consumers, choices maximize experienced utility, so decision utility and experienced utility are equivalent. The original utility function in Equation (3) was the decision utility function.

Decision utility is what can be estimated using observed choices. Specifying an experienced utility function that differs from decision utility requires the economist to take a stand on how the consumer is misoptimizing. We assume that an optimizing consumer sets $\gamma = 1$: experienced utility u^e_{ijat} depends only on the usage utility and the total consumption of the numeraire good, and a dollar of gasoline cost reduces numeraire good consumption by the same amount as a dollar in purchase price[16]:

[16] This can be framed more formally as an application of Bernheim and Rangel (2009). In their language, vehicle purchase is a "Generalized Choice Situation" in which consumer i chooses between a set of vehicles with total discounted user costs $p_{jat} + G_{jat}$ and usage utilities $\widetilde{\psi}_{jat} + \epsilon_{ijat}$. Whether the user cost flows through p or G is an "ancillary condition," meaning that while it may affect choices by agents who misoptimize, it is not material to welfare. We estimate elasticity to total discounted user cost η from only the "welfare-relevant domain," which we assume to be the consumers' response to variation in purchase prices. Conversely, we assume that variation in total discounted user cost resulting from variation in G is "suspect," meaning that it should not be used to infer utility functions. This set of assumptions implies that the true marginal utility of money is $\widehat{\eta}$ and that $\gamma = 1$ in the experienced utility function.

$$u^e_{ijat} = \eta(w_i - p_{jat} - 1 \cdot G_{jat}) + \widetilde{\psi}_{jat} + \epsilon_{ijat} \tag{9}$$

We adopt a term from Herrnstein, Loewenstein, Prelec, and Vaughan (1993) and describe the difference between decision utility and experienced utility as an "internality," denoted u^b_{ijat}. In our application, u^b captures the utility value of the portion of future gasoline costs that the consumer did not appropriately value in the discrete choice. This can be thought of as consumption of the numeraire good that the consumer anticipated having at the time of the discrete choice, but does not actually have because of additional expenditures on gasoline:

$$u^b_{ijat} = u^d_{ijat} - u^e_{ijat} = \eta(1-\gamma)G_{jat} \tag{10}$$

Aggregating over the populaton, experienced consumer surplus is decision utility net of the internality:

$$CS^e = CS^d - CS^b \tag{11}$$

We have now written experienced utility as the sum of two terms that can each be easily aggregated over consumers to give experienced consumer surplus. Summing over the choices made by consumers of market size M and transforming from utility to dollar terms by dividing by η, we have the average internality CS^b:

$$CS^b = \frac{1}{\eta} \cdot \frac{1}{M} \sum_{i=1}^{M} u^b_i = (1-\gamma)G_{jat} \cdot s_{jat} \tag{12}$$

Define the variable $\delta_{jat} = -\eta p_{jat} - \gamma\eta G_{jat} + \widetilde{\psi}_{jat}$ as the average decision utility for product ja at time t. We can integrate up over the logit error to find "decision consumer surplus" using the standard formula from Small and Rosen (1981), modified for the nested logit. Up to a constant, this is:

$$CS^d = \frac{1}{\eta} \ln \left[\sum_{n \in \mathcal{N}} \left[\sum_{ja \in \mathcal{B}_n} \exp\left(\frac{\delta_{jat}}{1-\sigma} \right) \right]^{1-\sigma} \right] \quad (13)$$

In this equation, n indexes nests, $\mathcal{N}$ refers to the set of all nests, and $\mathcal{B}_n$ refers to the set of vehicles in nest n.

The appeal of this approach is the resulting simplicity: it allows the use of the Small and Rosen (1981) analytical formula for average consumer surplus instead of requiring the analyst to simulate out the unobserved taste shocks ϵ_{ijat}. This approach can be used as long as the internality is additively separable from decision utility. It is general to any discrete choice setting and could be easily extended to random coefficients models.

6.2 Simulation

What are the welfare costs of misoptimization? Put differently, what are the gains in experienced consumer surplus from a policy that moves consumers to their private optima? Consider an "internality tax," in the spirit of O'Donoghue and Rabin (2006) and Allcott, Mullainathan, and Taubinsky (2011), that raises the relative prices of low-fuel economy vehicles. Our counterfactual tax has three properties, which give it the form of a fee-and-rebate policy that transportation analysts often call a "feebate." First, the policy changes vehicle prices by amounts F_j^* such that consumers choose the same vehicle as they do in their private optima. Second, because we do not wish make assumptions required to model substitution in and out of the new vehicle market, our counterfactual policy includes a rebate R on all new vehicle sales that adjusts the price level so as to hold total new vehicle sales constant. The feebate amount is:

$$F_j^* = (1-\gamma)\, G_j - R \quad (14)$$

The third property of our counterfactual feebate is that it is revenue neutral: any deficit or surplus revenues are recycled to all consumers (including those who choose the outside option) as a lump sum tax or refund.

To see how this policy induces consumers to choose the same vehicle as their private optima, observe that when F_j^* is substituted into the decision utility function, price p_j and gasoline cost G_j have the same attention weight, just as they do in the experienced utility function:

$$\begin{aligned} u_{ij}^d &= \eta(w_i - p_j - \gamma G_j - F_j^*) + \widetilde{\psi}_j + \epsilon_{ij} \\ &= \eta(w_i - p_j - \gamma G_j - ((1-\gamma)\, G_j - R)) + \widetilde{\psi}_j + \epsilon_{ij} \\ &= \eta(w_i - p_j - G_j + R)) + \widetilde{\psi}_j + \epsilon_{ij} \end{aligned} \tag{15}$$

Of course, if such a feebate policy were introduced, it would affect many aspects of vehicle markets. For example, the prices of used low fuel economy vehicles would increase as consumers substitute away from new low fuel economy vehicles, and auto manufacturing firms would likely offer a wider variety of high MPG vehicles and invest more in R&D to improve fuel economy. Simulating these effects is well beyond the scope of this paper. Our simulation uses model year 2007 new vehicles as the choice set and aggregates all other choices into an outside option. The simulation takes demand parameters $\widehat{\eta}$ and $\widehat{\sigma}$ from the nested logit specification, backs out the set of $\widehat{\psi}_j$ that rationalize observed market shares, adds the feebate F_j^* while holding vehicle prices and characteristics constant, resimulates market shares using Equation (4), and calculates the change in market outcomes and consumer surplus. The simulation institutes the feebate for model year 2007 sales only and reports effects over the lifetimes of new vehicles sold in that one year. Note that when we report effects in dollars per "consumer," this denominator reflects all consumers potentially in the market, both those that actually purchase a new vehicle and those who choose the outside option: the total number of consumers is the 240 million Americans over age 16.

6.3 Results

Table 8 presents the simulation results for $\gamma = 0.9$, $\gamma = 0.75$, and $\gamma = 0.5$. Focus first on the results for $\gamma = 0.75$, which maps most closely to our base specification results. The "pivot" of the feebate is 19.1 MPG: all vehicles with lower fuel economy have $F_j^* > 0$, with the tax increasing as fuel economy worsens. All vehicles with fuel economy higher than the pivot have a net subsidy $F_j^* < 0$, due to the rebate included to keep new vehicle sales constant. Total sales of vehicles with fuel economy better than the pivot increase by about 15 percent, while sales of vehicles worse than the pivot MPG decrease by 21 percent.

Instituting the policy for one year reduces decision consumer surplus CS^d by a present discounted value of about \$5.50 per consumer, as the change in relative prices moves consumers away from their perceived optimum. The average internality CS^b, however, is \$10.90 lower. Experienced consumer surplus therefore increases by just over \$5 per consumer. This is \$77 for each person who buys a new vehicle, and \$1.3 billion in total. For comparison, the total value of the new vehicle market in 2007 was \$350 billion. Intuitively, these welfare gains accrue as consumers spend less money on gasoline and more on some combination of higher fuel economy vehicles and the numeraire good. Since the optimal feebate achieves the first best in this model, these welfare gains are also equal to the welfare losses from misoptimization.

At 2007 gasoline prices, which were about \$2.40 per gallon, instituting the policy for one year reduces gasoline costs by a PDV of \$44 per consumer over the lifetimes of the vehicles purchased. Notice that these gasoline cost savings are eight times larger than the welfare gains. This is because although high-MPG vehicles save money, consumers like larger, lower-MPG vehicles more, and there are substantial losses in usage utility ψ as the counterfactual policy moves consumers into higher-MPG vehicles.

While one could imagine enriching the model in various ways, even this stylized analysis generates two stark insights. First, although climate change has been a primary motivator of academic research and policy action in this domain, misoptimization could be causing much larger distortions. To see this, notice that carbon externalities and undervaluation of fuel costs distort the new vehicle market in the same way, by causing consumers to favor low-MPG vehicles more than the social optimum by an amount proportional to gasoline costs G. Thus, the magnitudes of the two optimal corrective taxes are directly comparable, and they are sufficient statistics for the welfare gains. We can therefore use the carbon content per gallon of gasoline to calibrate what the marginal damage from carbon emissions would need to be in order for the optimal carbon tax to be equal to the optimal internality tax.

As shown in Table 8, if $\gamma = 0.75$, the optimal taxes, and thus the distortions without the taxes and the welfare gains from implementing them, are equal only if marginal damages are $70 per metric ton of carbon dioxide. For comparison, the United States Government's (2010) social cost of carbon is approximately $20 per ton. Even if $\gamma = 0.9$, the distortion from misoptimization is greater than the distortion from the unpriced carbon externality unless the marginal damages are $28 per ton. The same insight can also be drawn by comparing the welfare gains of the optimal feebate that flow through the change in experienced consumer surplus CS^e to the gains that flow through a reduced carbon externality under the assumed $20 per ton marginal damages. When $\gamma = 0.75$, for example, Table 8 shows that the feebate increases private welfare by about $5 per consumer while reducing climate change damages by only $2.70 per consumer.

The second fundamental insight from this model is that although misoptimization could in principle justify some corrective policy, CAFE standards appear to be much more aggressive than can be justified by misoptimization alone. To see this, notice that our feebate policy has the same effects on quantities demanded as a CAFE regulation in a stylized model where the choice set is fixed and automakers comply by raising the prices of low-MPG vehicles and cross-subsidizing high-MPG vehicles. Under those assumptions, the optimal CAFE standard to correct misoptimization increases MPG above baseline by the same amount as the optimal feebate. From Table 8, the optimal increase in fuel economy above the baseline market equilibrium is 1.0 MPG if $\gamma = 0.75$ and 2.1 MPG if $\gamma = 0.5$. By comparison, the revised CAFE standard in the 2007 Energy Independence and Security Act increases average fuel economy by more than seven MPG above the existing CAFE standard, which itself had raised average fuel economy above some baseline equilibrium. Certainly, other market failures might justify CAFE standards, but this argument is tempered by our first result above, which was that misoptimization appears to be a relatively important distortion.

7 Conclusion

At least since the energy crises of the 1970s, economists and policymakers have been interested in how consumers trade off future costs of energy using durable goods with their purchase prices. This paper builds on this literature but reframes the question, the identification strategy, and the policy implications. We adopt a different and more plausible approach to estimating demand for energy efficiency, which exploits the interaction of time series variation in gasoline prices with cross-sectional variation in fuel economy ratings of different vehicles. We introduce a new instrument for used vehicle market shares, which we expect to be a useful alternative to the Berry, Levinsohn, and Pakes (1995) instruments, which rely on particular assumptions about how firms set prices. We also develop a new and very tractable approach to behavioral welfare analysis in discrete choice models.

In our base specification, we estimate that consumers are indifferent between one dollar in vehicle purchase prices and only 72 cents in the present discounted value of future gasoline costs. The result that consumers undervalue gasoline costs appears robust to many factors, including additional potential measurement error, the potential endogeneity of market shares and vehicle-miles traveled, and potential changes over time in consumer preferences or vehicle characteristics. Some alternative specifications, such as the use of retail price data from JD Power or a lower discount rate, strengthen the qualitative conclusion that $\gamma < 1$. However, other alternative specifications move $\widehat{\gamma}$ closer to one, including the use of plausible higher discount rates and modifying the time period analyzed. We caution that although no individual assumption we examined moves $\widehat{\gamma}$ to unity, it is not hard to imagine combinations of alternative assumptions that could generate this result.

While the undervaluation result is therefore not unambiguous, two relatively strong qualitative conclusions can be drawn by combining our empirical study and stylized welfare simulations. First, although the primary regulatory justification for Corporate Average Fuel Economy standards is that consumers are misoptimizing (NHTSA 2010), the regulation appears to be much more stringent than can be justified by even our smallest plausible $\widehat{\gamma}$. Second, however, even if γ is much closer to one than our base specification suggests, misoptimization distorts vehicle markets more than the failure to internalize climate change externalities. Given the amount of policy attention and academic interest surrounding climate change, this remarkable result suggests that economists should be devoting significant additional effort to understanding how consumers value energy costs and understanding the policy consequences of undervaluation.

References

[1] Abaluck, Jason, and Jonathan Gruber (2011). "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review*, Vol. 101, No. 4 (June), pages 1180-1210.

[2] Adams, William, Liran Einav, and Jonathan Levin (2009). "Liquidity Constraints and Imperfect Information in Subprime Lending." *American Economic Review*, Vol. 99, No. 1, pages 49–84.

[3] Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2010). "Externalizing the Internality." Working Paper, MIT Department of Economics (December).

[4] Allcott, Hunt (2010). "Beliefs and Consumer Choice." Working Paper, MIT Department of Economics (November).

[5] Allcott, Hunt (2011). "Consumers' Perceptions and Misperceptions of Energy Costs." *American Economic Review*, Vol. 101, No. 2 (May).

[6] Allcott, Hunt, and Nathan Wozny (2009). "Gasoline Prices, Fuel Economy, and the Energy Paradox." Working Paper, Princeton University (November).

[7] Altonji, Joseph, Todd Elder and Christopher Taber (2005). "An Evaluation Of Instrumental Variables Strategies For Estimating The Effects Of Catholic Schooling." *Journal of Human Resources*, Vol. 40, No. 4 (Fall), pages 791-821.

[8] Anderson, Soren, Ryan Kellogg, James Sallee, and Richard Curtin (2011). "Forecasting Gasoline Prices Using Consumer Surveys." *American Economic Review*, Vol. 101, No. 2 (May).

[9] Angrist, Joshua (1991). "Grouped-Data Estimation and Testing in Simple Labor-Supply Models." *Journal of Econometrics*, Vol. 47, No. 2-3 (February), pages 243-266.

[10] Ashenfelter, Orley (1984). "Macroeconomic Analyses and Microeconomic Analyses of Labor Supply." *Carnegie-Rochester Conference Series on Public Policy*, Vol. 21, No. 1 (January), pages 117-156.

[11] Atkinson, Scott, and Robert Halvorsen (1984). "A New Hedonic Technique for Estimating Attribute Demand: An Application to the Demand for Automobile Fuel Efficiency." *Review of Economics and Statistics*, Vol. 66, No. 3 (August), pages 417-426.

[12] Austin, David (2008). "Effects of Gasoline Prices on Driving Behavior and Vehicle Markets." U.S. Congressional Budget Office Working Paper (January).

[13] Barber, Brad, Terrance Odean, and Lu Zheng (2005). "Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows." *Journal of Business*, Vol. 78, No. 6 (November), pages 2095-2120.

[14] Bernheim, Douglas, and Antonio Rangel (2009). "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics*, Vol. 124, No. 1 (February), pages 51-104.

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

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

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

[18] Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer (2009, 2011). "Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing." Working Paper, Northwestern University.

[19] Chetty, Raj, Adam Looney, and Kory Kroft (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, Vol. 99, No. 4 (September), pages 1145-1177.

[20] Cragg, John, and Stephen Donald (1993). "Testing Identifiability and Specification in Instrumental Variable Models," *Econometric Theory*, Vol. 9, No. 2 (April), pages 222-240.

[21] Dreyfus, Mark, and Kip Viscusi (1995). "Rates of Time Preference and Consumer Valuations of Automobile Safety and Fuel Efficiency." *Journal of Law and Economics*, Vol. 38, No. 1, pages 79-98.

[22] Dubin, Jeffrey (1992). "Market Barriers to Conservation: Are Implicit Discount Rates Too High?" Social Science Working Paper 802, California Institute of Technology (July).

[23] Dubin, Jeffrey, and Daniel McFadden (1984). "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica*, Vol. 52, No. 2 (March), pages 345-362.

[24] Ellison, Glenn (2005). "A Model of Add-on Pricing." *Quarterly Journal of Economics*, Vol. 120, No. 2 (May), pages 585-638.

[25] Espey, Molly, and Santosh Nair (2005). "Automobile Fuel Economy: What Is It Worth?" *Contemporary Economic Policy*, Vol. 23, No. 3 (July), pages 317-323.

[26] Finkelstein, Amy (2009). "E-ZTAX: Tax Salience and Tax Rates." *Quarterly Journal of Economics*, Vol. 124, No. 3 (August), pages 969-1010.

[27] Fischer, Carolyn, Winston Harrington, and Ian Parry (2007). "Do Market Failures Justify Tightening Corporate Average Fuel Economy (CAFE) Standards?" *The Energy Journal*, Vol. 28, No. 4, pages 1-30.

[28] Gabaix, Xavier, and David Laibson (2006). "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, Vol. 121, No. 2, pages 505-540.

[29] Gillingham, Kenneth (2010). "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock in California." Working Paper, Stanford University (September).

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

[31] Greene, David, Janet Hopson, R. Goeltz and Jia Li, (2007). "Analysis of In-Use Fuel Economy Shortfall Based on Voluntarily Reported Mile-per-Gallon Estimates," *Transportation Research Record*, No. 1983, pages 99-105.

[32] Greene, David, Philip Patterson, Margaret Singh, and Jia Li (2005). "Feebates, Rebates, and Gas-Guzzler Taxes: A Study of Incentives for Increased Fuel Economy." *Energy Policy*, Vol. 33, No. 6 (April), pages 757-775.

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

[34] Hossain, Tanjim, and John Morgan (2006). "...Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay." *Advances in Economic Analysis and Policy*, Vol. 6.

[35] Hughes, Jonathan, Christopher Knittel, and Daniel Sperling (2008). "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *The Energy Journal*, Vol. 29, No. 1, pages 113-134.

[36] Jacobsen, Mark (2010). "Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity." Working Paper, University of California at San Diego (September).

[37] Jaffe, Adam, and Robert Stavins (1994). "The Energy Paradox and the Diffusion of Conservation Technology." *Resource and Energy Economics*, Vol. 16, pages 91-122.

[38] Kahn, James (1986). "Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach." *Quarterly Journal of Economics*, Vol. 101, No. 2 (May), pages 323-340.

[39] Kilian, Lutz and Eric Sims (2006). "The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data," Working Paper, University of Michigan (April).

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

[41] Knittel, Christopher (2011). "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review*, forthcoming.

[42] Langer, Ashley, and Nathan Miller (2009). "Automakers' Short-Run Responses to Changing Gasoline Prices and the Implications for Energy Policy." Working Paper, University of California at Berkeley (September).

[43] Lewis, Randall (2011). "Smallest Superset of Identification." Working Paper, Yahoo Research (February).

[44] Li, Shanjun, Christopher Timmins, and Roger von Haefen (2009). "Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy*, Vol. 1, No. 2 (August), pages 113-137.

[45] Manheim (2009). "Used Car Market Report." 14th Edition. Atlanta, GA: Manheim Consulting.

[46] Mullainathan, Sendhil (2002). "Thinking Through Categories." Working Paper, MIT (April).

[47] Murphy, Kevin M., and Robert Topel (1985). "Estimation and Inference in Two-Step Econometric Models." *Journal of Business and Economic Statistics*, Vol. 3, No. 4 (October), pages 370-379.

[48] National Highway Traffic Safety Administration (NHTSA) (2010). "Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2012-MY 2016 Passenger Cars and Light Trucks." Office of Regulatory Analysis and Evaluation, National Center for Statistics and Analysis (March).

[49] O'Donoghue, Ted, and Matthew Rabin (2006). "Optimal Sin Taxes." *Journal of Public Economics*, Vol. 90, No. 10-11 (November), pages 1825-1849.

[50] Ohta, Makoto, and Zvi Griliches (1986). "Automobile Prices and Quality: Did the Gasoline Price Increases Change Consumer Tastes in the U.S.?" *Journal of Business and Economic Statistics*, Vol. 4, No. 2 (April), pages 187-198.

[51] Parry, Ian, Winston Harrington, and Margaret Walls (2007). "Automobile Externalities and Policies." *Journal of Economic Literature*, Vol. 45, pages 374-400.

[52] Reis, Ricardo (2006). "Inattentive Consumers." *Journal of Monetary Economics*, Vol. 53, No. 8 (November), pages 1761-1800.

[53] Rust, John (1985). "Stationary Equilibrium in a Market for Durable Assets." *Econometrica*, Vol. 53, No. 4 (July), pages 783-805.

[54] Sallee, James, Sarah West, and Wei Fan (2011). "Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations." Work in Progress, Macalester College (January).

[55] Sawhill, James (2008). "Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market." Working Paper, University of California at Berkeley (April).

[56] Small, Kenneth, and Kurt Van Dender (2007). "Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect." *The Energy Journal*, Vol. 28, No. 1, pages 25-51.

[57] Small, Kenneth, and Harvey Rosen (1981). "Applied Welfare Economics with Discrete Choice Models." Econometrica, Vol. 49, No. 1 (January), pages 105-130.

[58] Staiger, Douglas, and James Stock (1997). "Instrumental Variables Regression with Weak Instruments." *Econometrica*, Vol. 65, No. 3 (May), pages 557-586.

[59] Stock, James, and Motohiro Yogo (2005). "Testing for Weak Instruments in Linear IV Regression." In Stock, James, and Donald Andrews, Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg. Cambridge, England: Cambridge University Press.

[60] Stolyarov, Dmitriy (2002). "Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?" Journal of Political Economy, Vol. 110, No. 6 (December), pages 1390-1413.

[61] U.S. Bureau of Labor Statistics (BLS) (2007). "Consumer Expenditure Survey." ftp://ftp.bls.gov/pub/special.requests/ce/standard/2007/quintile.txt

[62] U.S. Department of Transportation (2001). "National Household Travel Survey." http://nhts.ornl.gov/download.shtml#2001.

[63] U.S. Energy Information Administration (2011). "Monthly Energy Review." http://www.eia.doe.gov/totalenergy/data/monthly/.

[64] U.S. Environmental Protection Agency (2008). "Light-Duty Automotive Technology and Fuel Economy Trends: 1975 Through 2008. Appendix A: Database Details and Calculation Methods." http://www.epa.gov/otaq/cert/mpg/fetrends/420r08015appx-a.pdf.

[65] U.S. Government, Interagency Working Group on Social Cost of Carbon (2010). "Technical Support Document: - Social Cost of Carbon for Regulatory Impact Analysis - Under Executive Order 12866 -." http://www.epa.gov/oms/climate/regulations/scc-tsd.pdf

[66] Verboven, Frank (1999). "The Markets for Gasoline and Diesel Cars in Europe." Centre for Economic Policy Research Discussion Paper 2069 (February).

[67] Wald, Abraham (1940). "The Fitting of Straight Lines if Both Variables are Subject to Error." *Annals of Mathematical Statistics*, Vol. 11, pages 284-300.

[68] Yahoo Autos (2009). "Top 100 Cars by Green Rating." http://autos.yahoo.com/green_center-top100/, accessed August 2009.

Tables

Table 1: Summary Statistics

	Complete Dataset	2007, New Vehicles	Base Estimation Sample
Year	2003.4	2007.0	2003.1
	(2.9)	(0.0)	(2.6)
Model Year	2000.0	2007.0	1998.8
	(4.4)	(0.0)	(4.0)
Age	3.4	0.0	4.3
	(3.5)	(0.0)	(3.3)
Price	13,485	25,826	10,077
	(10,491)	(10,781)	(6,925)
Quantity	63,958	57,073	65,670
	(65,002)	(62,932)	(65,693)
MPG	19.2	20.7	18.9
	(4.0)	(5.4)	(3.8)
G	9,462	13,878	8,506
	(4,200)	(3,771)	(3,529)
Horsepower	187	219	180
	(56)	(68)	(51)
Weight	3,621	3,878	3,573
	(1,079)	(923)	(981)
Wheelbase	111.5	114.3	111.2
	(13)	(15)	(13)
Percent Car	56	48	56
Observations	1,396,254	10,453	854,248

Notes: Weighted by transaction quantities. Standard deviations in parenthesis. The complete dataset includes monthly observations of all light duty vehicles between January 1999 and December 2008. Column 2 includes 2007 model year vehicles observed in 2007.

Table 2: Sample Determination

Sample	N
Full Dataset	1,396,254
Exclude Vans and Exotics	1,248,324
Exclude New Vehicles	1,132,891
Exclude April-December 2008	1,039,798
Exclude if Zero Registered Quantity	1,039,263
Exclude if Only One Observation per Fixed Effect Group	854,248

Table 3: Discount Rates

Payment Method	**Share of Vehicles**	**Discount Rate**
New		
Financed	54%	3.9%
Leased	19%	3.0%
Cash	27%	5.8%
Total	100%	
Weighted Average		4.2%
Used		
Financed	37%	6.9%
Leased	0%	N/A
Cash	63%	5.8%
Total	100%	
Weighted Average		6.2%
Weighted Average	37% New, 63% Used	5.5%

Notes: Share of Vehicles and Finance and Lease Discount Rates are real averages from the 2001, 2004, and 2007 Survey of Consumer Finances. Cash Discount Rate is the real average return to the S&P 500 between 1945 and 2008.

Table 4: Gas Prices

Year	**Spot**	**Future Year**								
		0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10
1998	1.34	1.44	1.46	1.47	1.47	1.46	1.46	1.46	1.45	-	-
1999	1.43	1.50	1.46	1.45	1.44	1.44	1.43	1.43	1.39	-	-
2000	1.77	1.73	1.61	1.56	1.51	1.49	1.47	1.46	-	-	-
2001	1.69	1.65	1.59	1.55	1.52	1.51	1.49	1.48	1.47	-	-
2002	1.56	1.63	1.58	1.55	1.53	1.51	1.50	1.48	1.50	-	-
2003	1.74	1.71	1.62	1.59	1.58	1.57	1.56	1.55	1.59	-	-
2004	1.99	1.95	1.84	1.78	1.74	1.70	1.68	1.66	1.75	-	-
2005	2.34	2.33	2.28	2.22	2.15	2.10	2.06	2.03	2.02	-	-
2006	2.55	2.55	2.55	2.49	2.42	2.36	2.32	2.27	-	-	-
2007	2.68	2.59	2.55	2.50	2.46	2.42	2.39	2.40	2.37	2.35	2.56
2008	3.00	3.12	3.10	3.08	3.06	3.04	3.02	2.99	2.97	2.95	2.67

Notes: All prices are in dollars per gallon and are inflation adjusted to 2005 dollars. Futures prices are deflated to 2005 dollars using inflation expectations implied by Treasury Inflation-Protected Security yields, then transformed from oil prices to retail gasoline prices using their historical average relationship.

Table 5: Estimation Results

Row	Specification:	$\widehat{\gamma}$	SE
	Grouping		
0	Time by MPG Quantile; 2 MPG Quantiles	0.72	0.048
1	Time by MPG Quantlle by Age Group; 2 MPG Quantiles and 4-year age groups	0.66	0.038
2	Time by Age by MPG Quantile; 2 MPG Quantiles	0.67	0.038
3	Time by Age by MPG Quantile; 5 MPG Quantiles	0.63	0.035
4	Time by Age by MPG Quantile; 10 MPG Quantiles	0.64	0.034
5	Time by Age by MPG Quantile; 20 MPG Quantiles	0.63	0.035
6	Time by Age by Firm by MPG Quantile; 20 MPG Quantiles	0.61	0.031
7	OLS (No grouping)	0.53	0.035
	Grouping Over Time		
11	Time Group by MPG Quantile; 2 month groups and 2 MPG Quantiles	0.73	0.049
12	Time Group by MPG Quantile; 3 month groups and 2 MPG Quantiles	0.74	0.050
13	Time Group by MPG Quantile; 4 month groups and 2 MPG Quantiles	0.73	0.050
14	Time Group by MPG Quantile; 6 month groups and 2 MPG Quantiles	0.73	0.053
	Time Periods		
21	2004-March 2008	0.77	0.046
22	1999-End 2008	0.88	0.038
	r		
31	r=3%	0.65	0.043
32	r=10%	0.83	0.055
33	r=15%	0.97	0.066

Table 5 continues on the next page.

Table 5 (Continued): Estimation Results

Row	Specification:	$\widehat{\gamma}$	SE
	Intensive margin		
41	Constant elasticity = 0.2	0.71	0.064
42	Linear elasticity = 0.2	0.72	0.058
43	Constant elasticity = 0.5	0.68	0.104
44	Linear elasticity = 0.5	0.70	0.078
	Characteristics (Grouping at Time by 20 MPG Quantiles by Age)		
51	Sample with characteristics	0.68	0.037
52	Control for HP	0.69	0.037
53	Control for HP, weight	0.68	0.036
54	Control for HP, weight, wheelbase	0.69	0.036
55	Control for HP, weight, wheelbase, ABS, stability, traction	0.64	0.036
	Preferences		
61	Exclude hybrid vehicles	0.72	0.048
62	Exclude "green" vehicles	0.71	0.050
63	Exclude SUVs	0.61	0.074
64	Exclude all cars	0.83	0.099
65	Exclude minivans	0.72	0.041
	JDPA Prices		
71	JDPA-Manheim Common Sample: Manheim Data	0.76	0.052
72	JDPA-Manheim Common Sample: JDPA Data	0.53	0.044
73	JD Power Used Vehicle Prices	0.53	0.044
74	Common Sample, Manheim, 2004-March 2008	0.79	0.049
75	Common Sample, Manheim, 1999-End 2008	0.92	0.041
76	Common Sample, JDPA 2004-March 2008	0.65	0.044
77	Common Sample, JDPA 1999-End 2008	0.66	0.032
	Fixed Effects (Grouping at Time by 2 MPG Quantiles by Four-Year Age Groups)		
81	Model-by-Model Year; Age Controls	-1.61	0.051
82	Model-by-Model Year; Age-by-MPG Controls	-0.59	0.045

Notes: Unless otherwise stated, all specifications include 1999 through March 2008. For all specifications covering 1999 through March 2008 with Manheim data, N=854,258, and there are 260,689 ja fixed effect groups. For the specifications including 1999 through end 2008 with Manheim data, N=934,860. All specifications include a full set of month-by-year time dummies. Observations weighted by number of observed transactions. Unless otherwise stated, all specifications use the grouping estimator with groups at the level of time by MPG quantile, with two MPG quantiles. Robust standard errors, clustered at the level of model by age (in years), are in parenthesis.

Table 6: Alternative Gas Prices

	(1)	(2)	(3)
Specification:	**Futures Forecast**	**Martingale Forecast**	**Both**
G	-0.72		-0.67
	(0.048)		(0.049)
G_m		-0.50	
		(0.032)	
$[G_m - G]$			-0.226
			(0.058)
Observations	854,248	854,248	854,248
Partial R^2	0.435	0.419	0.453

Notes: Column 1 replicates Row 0 of Table 4, the base specification. All other columns have the same structure. This means that they include 1999-March 2008 data, include month-by-year time dummies, and use the grouping estimator, with groups at the level of time by MPG quantile, with 2 MPG quantiles. Observations weighted by number of observed transactions. Robust IV standard errors, clustered at the level of model by age (in years), are in parenthesis.

Table 7: Nested Logit

	(1)	(2)	(3)	(4)	(5)	(6)
Model:	**Exog s**	**Logit**	**NL**	**NL**	**NL**	**NL: OLS**
G	-0.40	-0.44	-0.36	-0.36	-0.36	-0.34
	(0.028)	(0.036)	(0.035)	(0.035)	(0.035)	(0.028)
$ln(s)$		-2565	-2427	-2414	-2448	78
		(1003)	(934)	(960)	(956)	(39)
$ln(s_n)$			-3380	-3264	-3142	-2612
			(802)	(1194)	(1217)	(514)
$ln(s_n)$, 2nd Nest				-87	948	
				(570)	(757)	
$ln(s_n)$, 3rd Nest					-1130	
					(494)	
First Stage						
$G_0/10^6$		-78.4	-79.4	-77.0	-77.3	
		(20.9)	(16.6)	(20.7)	(20.7)	
$G/10^6$		26.0	34.3	33.1	32.7	
		(13.1)	(12.7)	(12.6)	(12.7)	
$ln(s_n)$			-0.33	-0.59	-0.61	
			(0.22)	(0.30)	(0.30)	
$ln(s_n)$, 2nd Nest				0.20	0.10	
				(0.14)	(0.16)	
$ln(s_n)$, 3rd Nest					0.12	
					(0.09)	
Cragg-Donald Stat		22.2	23.0	21.7	21.8	
Observations	420,905	420,905	420,905	420,905	420,905	420,905
R^2	0.35					0.36
Nest structure			Class	Class/Age	Class/Age/Luxury	Class

Notes: All specifications include model-by-age (in months) fixed effects, month-by-year time dummies, and model year dummies. Observations weighted by number of observed transactions. Robust standard errors, clustered at the level of model by age (in years), are in parenthesis.

Table 8: Welfare Effects

γ	**0.9**	**0.75**	**0.5**
Effects on New Vehicle Market			
Pivot (MPG)	18.7	19.1	19.7
ΔSales Above Pivot (%)	5	15	31
ΔSales Below Pivot (%)	-9	-21	-36
ΔAverage MPG	0.4	1.0	2.1
Welfare Effects, Excluding CO2			
ΔCS^d: Change in Decision Consumer Surplus ($/consumer)	-0.9	-5.5	-21.6
ΔCS^b: Change in Internality ($/consumer)	-1.7	-10.9	-44.2
ΔCS^e: Change in Experienced Consumer Surplus ($/consumer)	0.8	5.3	22.6
Gasoline and CO2 Effects			
ΔGasoline Use (gallons/consumer)	-7.6	-19.8	-40.7
ΔGasoline Costs ($/consumer)	-17	-44	-88
ΔCO2 Emissions (metric tons/consumer)	-0.07	-0.17	-0.36
ΔClimate Externality ($NPV/consumer)	-1.0	-2.7	-5.5
Comparison to Climate Change Welfare Losses			
Marginal Damage for Equivalent Carbon Tax ($/metric ton CO2)	28	70	140

Notes: All numbers are net present values over the lifetimes of new vehicles sold, for a counterfactual policy that affects one model year of sales. For the ΔClimate Externality, marginal damages of CO2 emissions are assumed to be $20 per metric ton (U.S. Government 2010). Change in decision consumer surplus also includes the recycled net revenues from the feebate policy.

8 Figures

Figure 1: Gasoline Prices and Expectations Over Time

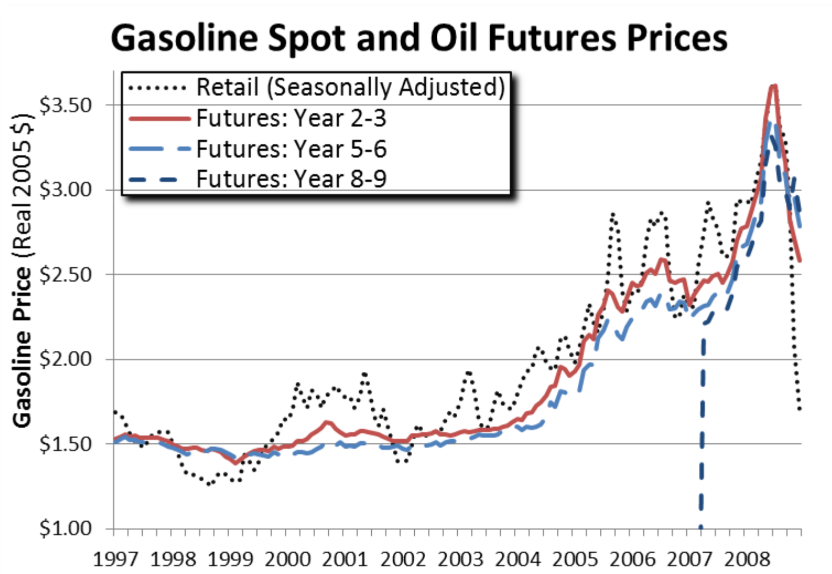

Figure 2: Identifying Variation in Gas Cost

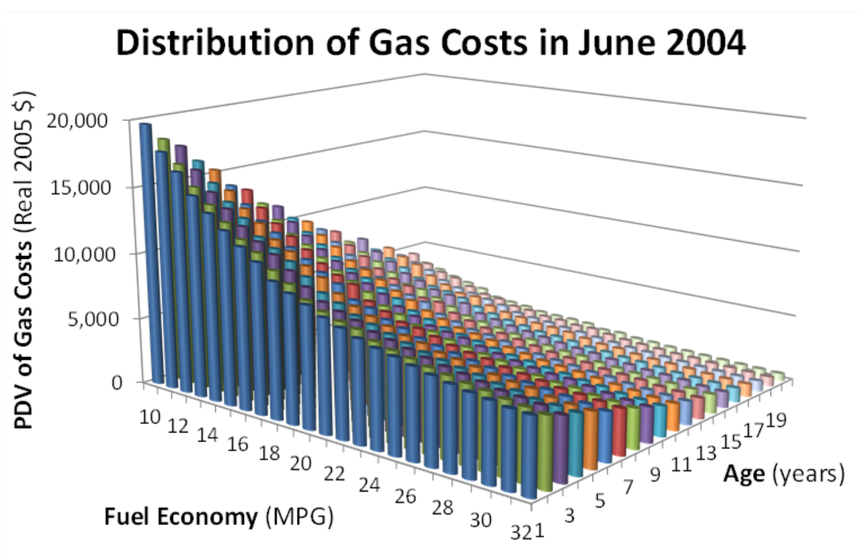

Figure 3: New Vehicle Sales by MPG Rating

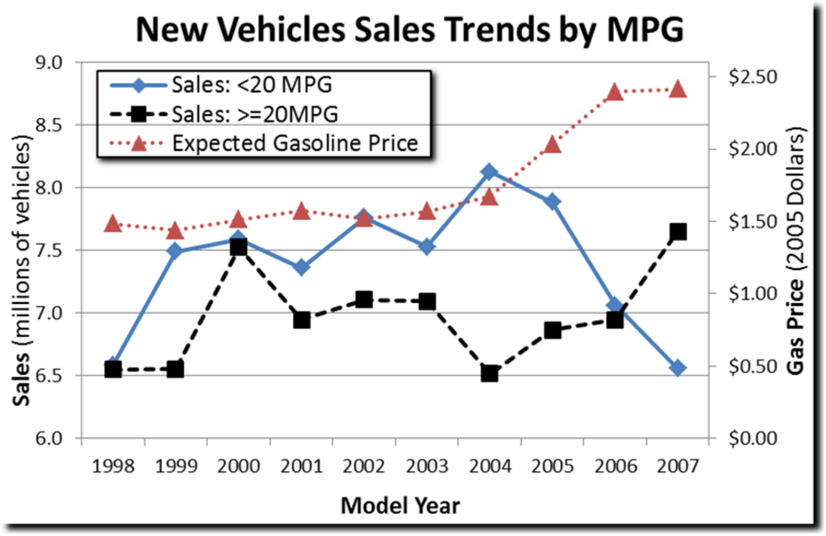

Figure 4: Manufacturer's Suggested Retail Price vs. Fuel Economy

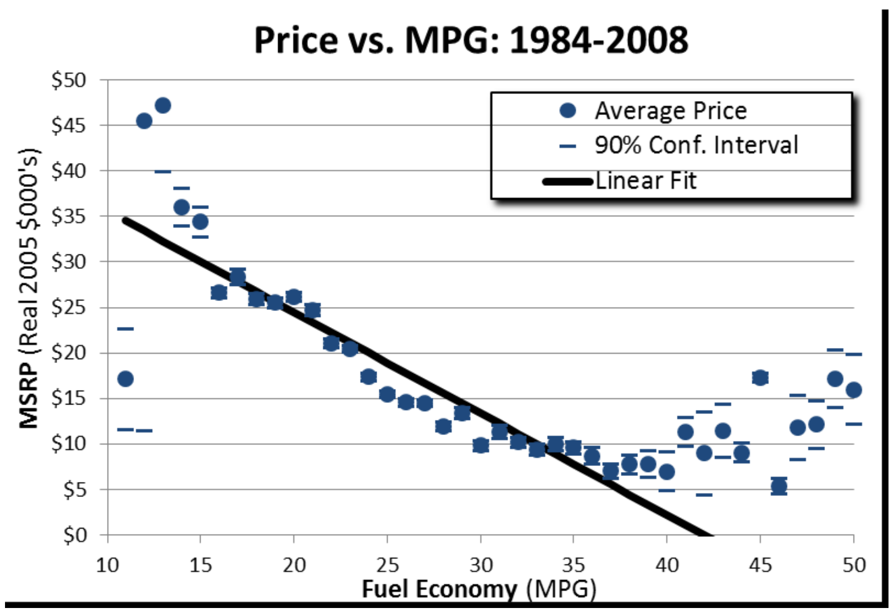

Figure 5: Vehicle Price Trends

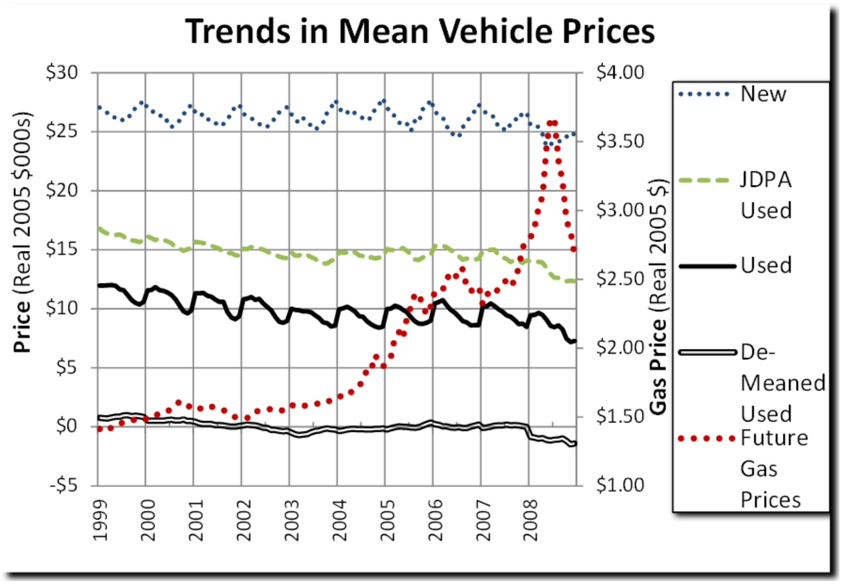

Figure 6: Raw Data Scatterplot

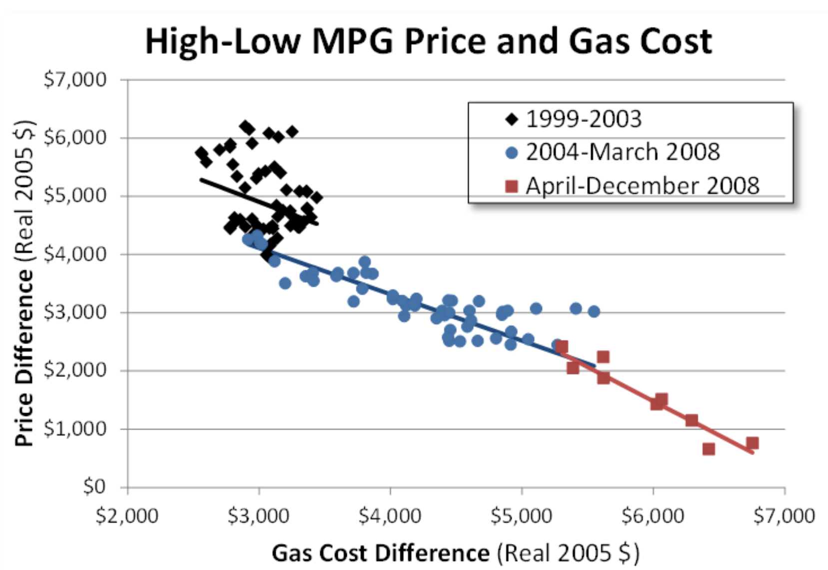

Figure 7: Conditional Variation Over Time

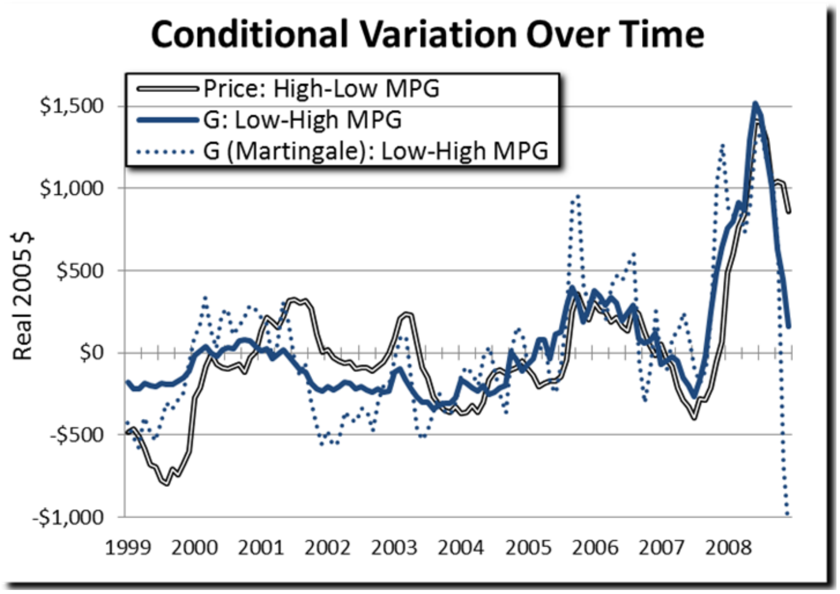

Figure 8: Partial Regression Plot: Ungrouped OLS

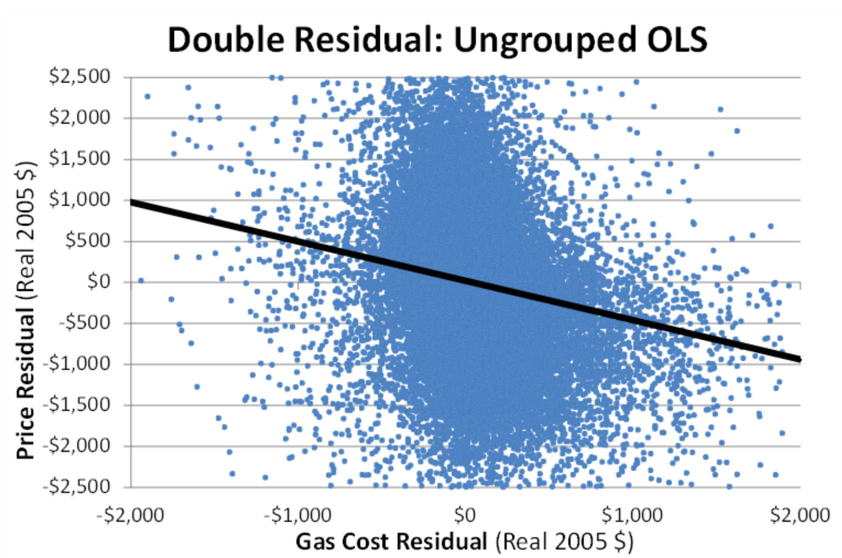

Note: For presentational reasons, this graph includes a randomly-selected 10 percent sample of observations.

Figure 9: Partial Regression Plot: Disaggregated Groups

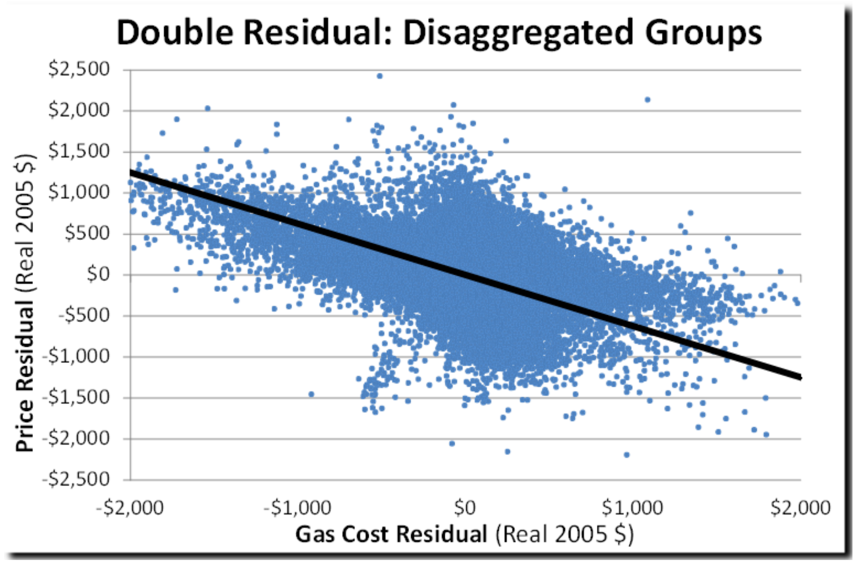

Figure 10: Partial Regression Plot: Base Specification

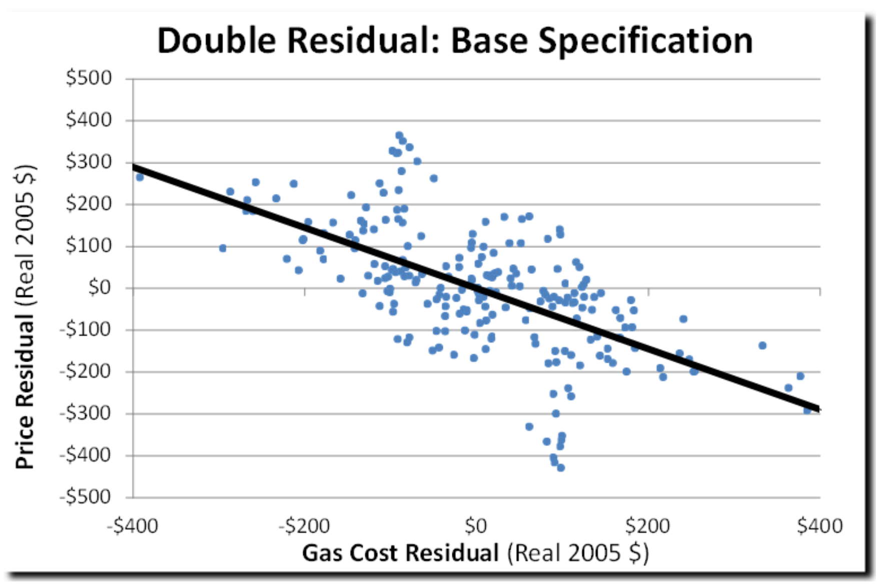

Figure 11: Partial Regression Plot: Aggregated Groups

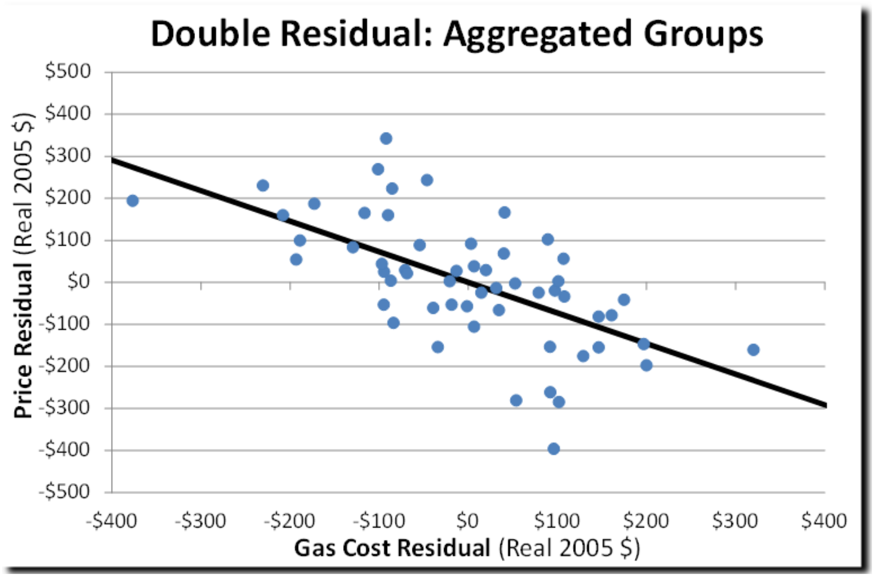

Figure 12: Depreciation Patterns

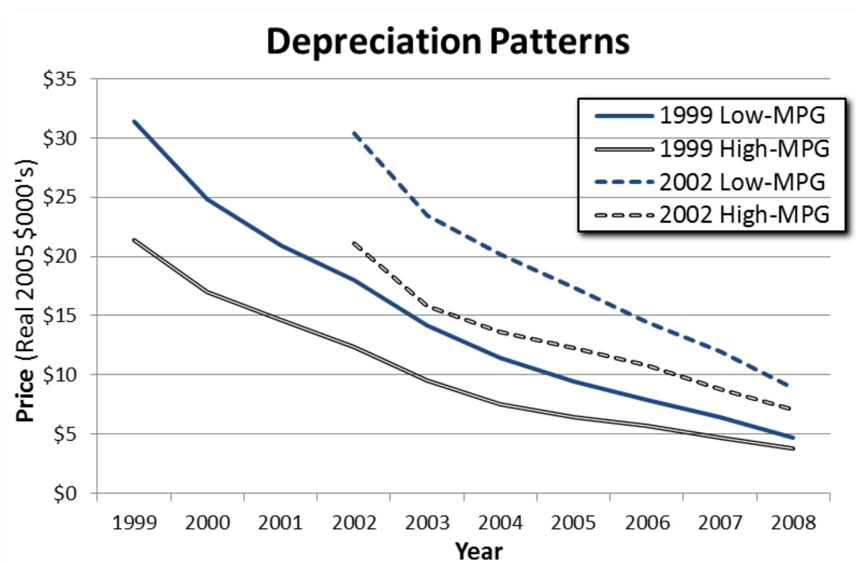

Note: This figure uses JDPA data in order to consistently present retail prices for new and used vehicles.

Figure 13: Predicted Mispricing

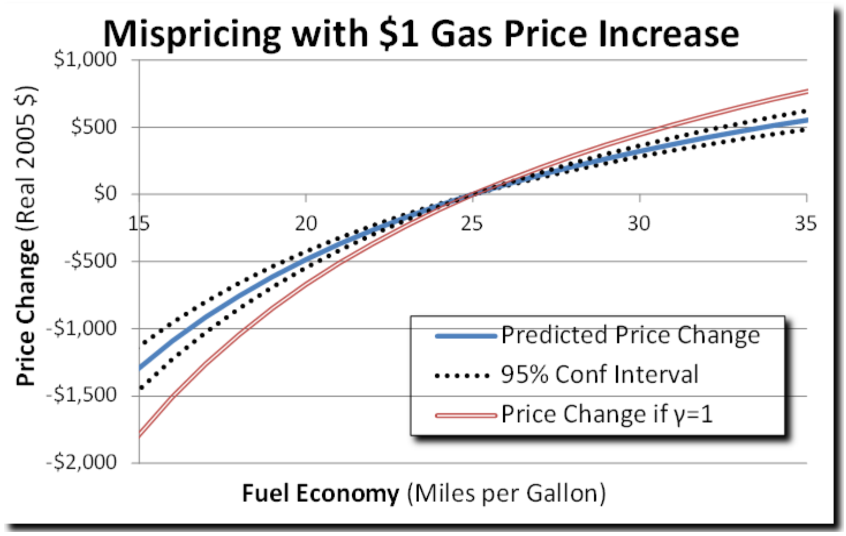

A Appendix: Endogenous Vehicle-Miles Traveled

Our primary specification assumes that Vehicle-Miles Traveled m_{jas} is independent of gasoline prices. In reality, demand for VMT is inelastic, but not fully inelastic, and the changes in g over the study period could be large enough that our Envelope Theorem approximation does not hold. This section presents a simple and intuitive alternative specification that accounts for endogenous VMT.

As illustrated in Appendix Figure 1, the change in VMT resulting from a change in gasoline prices affects consumer i's utility u_{ijat} through two channels. First, changes in VMT affect gasoline expenditures G_{jat}. The NHTS survey that measured VMT was carried out in 2001; at 2001 gas price g_{2001}, the cost to drive vehicle ja one mile is g_{2001}/f_{jas}, and the consumer chooses VMT $m_{ja,2001}(g_{2001})$. At time s with higher gasoline prices g_s, the consumer reduces VMT to $m_{jas}(g_s)$. The correct annual gasoline cost is now the shaded blue rectangle bounded by the g_s/f_{jas} and m_{jas}.

The second effect is that usage utility $\widetilde{\psi}_{jat}$ also decreases when an increase in gas price reduces VMT. For example, the utility from owning a vehicle and driving it 12,000 miles per year is different than the utility of owning a vehicle and driving it 11,500 miles per year. In Appendix Figure 1, the consumer's total willingness to pay for vehicle use is the area under the VMT demand curve. As gasoline prices increase from g_{2001} to g_s and the consumer's utility-maximizing VMT decreases, this total willingness to pay decreases by the solid green area.

To build on this intuition mathematically, explicitly denote $G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t), \cdot)$ and $\widetilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))$ as functions of $\mathbf{m}_{jat}(\mathbf{g}_t)$, where $\mathbf{g}_t$ and $\mathbf{m}_{jat}$ in boldface represent the vectors of forecasted gas prices and utility-maximizing VMTs in each future period of the vehicle ja's life beginning with time t. The utility function in Equation (3) can be written as:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G(\mathbf{m}_{jat}(\mathbf{g}_t), \cdot)) + \widetilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_{2001})) + \int_{\mathbf{g}_{2001}}^{\mathbf{g}_t} \frac{\partial \widetilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))}{\partial \mathbf{g}_t} d\mathbf{g}_t + \epsilon_{ijat} \tag{16}$$

The per-period dollar value of this new term with the integral is the solid green area in Appendix Figure 1. It can be calculated under an assumed functional form of VMT demand. We assume constant elasticity ζ, with the demand function intercept κ_{ja} pinned down by the vehicle's fitted VMT from the 2001 NHTS data and the gas prices from that year.[17] Based on recent empirical estimates,[18] we present specifications with $\zeta = -0.05$ and $\zeta = -0.2$. Denote the dollar value of the usage utility changes over the vehicle lifetime as $I_{jat}(\mathbf{g}_t)$:

[17] This can also be stated mathematically. Under constant elasticity of demand ζ, we have VMT at any gasoline price:

$$\ln(m_{jas}) = \zeta \cdot \ln(g_s) + \kappa_{ja} \tag{17}$$

The constant κ_{ja} is pinned down by $m_{ja,2001}$, the fitted VMT from the 2001 National Household Travel Survey data, and the gas prices at that time:

$$\kappa_{ja} = \ln(m_{ja,2001}) - \zeta \cdot \ln(g_{2001}) \tag{18}$$

[18] We draw on three recent estimates of the elasticity of VMT to gasoline prices, or equivalently the "short-run" elasticity of gasoline demand (holding vehicle capital stock fixed). Hughes, Knittel, and Sperling (2007) find that between 2001 and 2006, this elasticity was between -0.034 and -0.077. Small and Van Dender (2007) find that with covariates at their 1997-2001 levels (the latest years in their study period), the elasticity is -0.022. Using data from California between 2001 and 2008, Gillingham (2010) estimates a short-run elasticity of -0.15 to -0.2.

$$I_{jat}(\mathbf{g}_t) \equiv \frac{1}{\eta}\int_{\mathbf{g}_{2001}}^{\mathbf{g}_t} \frac{\partial\widetilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))}{\partial \mathbf{g}_t} d\mathbf{g}_t \tag{19a}$$

$$= \sum_{s=t+1}^{t+(L-1-a)} -\left\{\left(m_{jas}^{1+1/\zeta} - m_{ja,2001}^{1+1/\zeta}\right) \cdot \frac{\exp\left(\frac{-\kappa_{ja}}{\zeta}\right)}{1+1/\zeta}\right\} \cdot \frac{\phi_{jas}}{f_{jas}} \cdot \delta^{s-t} \tag{19b}$$

Using the same steps as in the derivation of the primary specification, we carry this term through from the utility function to an alternative estimating equation. Because I_{jat} and G_{jat} are highly correlated, we move I_{jat} to the left hand side. The estimating equation is:

$$p_{jat} - I_{jat}(\mathbf{g}_t) = -\gamma G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t)) + \tau_t + \psi_{ja}(\mathbf{m}_{ja}(\mathbf{g}_{2001})) + \xi_{jat} \tag{20}$$

In this equation, $G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))$ is calculated as before in Equation (6), except with VMT as a constant elasticity function of gas price. As before, the term $\psi_{ja}(\mathbf{m}_{ja}(\mathbf{g}_{2001}))$ represents a vehicle fixed effect.

Appendix Figure 1: VMT Demand