



E2e Working Paper 002

Gasoline Prices, Fuel Economy, and the Energy Paradox

Hunt Allcott and Nathan Wozny

November 2012

Published in

the Review of Economics and Statistics, 2014, 96(10), 779 - 795

This paper is part of the E2e Project Working Paper Series.

E2e is a joint initiative of the Energy Institute at Haas at the University of California, Berkeley, the Center for Energy and Environmental Policy Research (CEEPR) at the Massachusetts Institute of Technology, and the Energy Policy Institute at Chicago, University of Chicago. E2e is supported by a generous grant from The Alfred P. Sloan Foundation.

The views expressed in E2e working papers are those of the authors and do not necessarily reflect the views of the E2e Project. Working papers are circulated for discussion and comment purposes. They have not been peer reviewed.



**Massachusetts
Institute of
Technology**

Gasoline Prices, Fuel Economy, and the Energy Paradox
Hunt Allcott and Nathan Wozny
NBER Working Paper No. 18583
November 2012
JEL No. D03,D12,L62,Q41

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 show how this can be tested by measuring whether relative prices of vehicles with different fuel economy ratings fully adjust to time series variation in gasoline price forecasts. We then test the model using a detailed dataset based on 86 million transactions at auto dealerships and wholesale auctions between 1999 and 2008. Over our base sample, vehicle prices move as if consumers are indifferent between one dollar in discounted future gas costs and only 76 cents in vehicle purchase price. We document how endogenous market shares and utilization, measurement error, and different gasoline price forecasts can affect the results, and we show how to address these issues empirically. We also provide unique empirical evidence of sticky information: vehicle markets respond to changes in gasoline prices with up to a six month delay.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Nathan Wozny
2354 Fairchild Dr
Department of Economics and Geosciences
Suite 6K144
USAFA, CO 80840
nathan.wozny@usafa.edu

1 Introduction

In some situations, it appears that consumers choosing between products are 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 add to 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 highly inefficient relative to gas taxes.¹ If consumers undervalue fuel costs when they choose between vehicles, however, CAFE standards can increase welfare, 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

¹Jacobsen (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.

economists (e.g. Hausman 1979)² as well as the U.S. government.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. However, one problem with the paternalistic justification for fuel economy standards is that there is not much solid evidence on whether automobile buyers are actually misoptimizing (Parry, Walls, and Harrington 2007).

The relevant null hypothesis is that consumers are willing to pay 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 say that rejecting this hypothesis is evidence that consumers "undervalue or overvalue gasoline costs." This hypothesis 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.⁴ The typical empirical approach in this literature has been to exploit variation in the prices and energy efficiency ratings of a cross-section of energy-using durable goods. 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.

²Hausman (1979) finds 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 then argues that "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."

³For example, the Regulatory Impact Analysis that justifies the 2011-2015 increase in the CAFE standard argues that about \$15 billion per year in net benefits will flow to consumers who undervalue the benefits of fuel economy (NHTSA 2010).

⁴The implied discount rate is simply a shorthand way of capturing how people appear to trade off upfront costs against an analyst's estimate of future cash flows. Many underlying factors other than consumers' actual discount rates might affect the implied discount rate, including time horizons, beliefs, and inattention. In our setting, as in many others, these factors are empirically indistinguishable. We therefore prefer the "undervaluation or overvaluation" language, as it is more explicitly agnostic. Regardless of the choice of language, the analyst must eventually take a stand on consumers' actual discount rates in order to determine whether people's decisions conform to the theoretical prediction.

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.

This paper formalizes one potentially-promising alternative approach, which exploits the significant fluctuations in gasoline price forecasts over the past fifteen years. The approach exploits the fact that a vehicle's future fuel costs vary as a function of both fuel economy and forecasted gasoline prices, so changes in gas price forecasts should affect the relative value 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 2012, Langer and Miller 2012). The above null hypothesis, however, does more than predict that gasoline price forecasts 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. Conversely, if vehicle prices respond to gas price forecasts more than theory predicts, this suggests that consumers overvalue gasoline costs.

We begin from the primitives of a discrete choice utility function and show the empirical assumptions required to test the null hypothesis. We then implement this test using data from 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 sets.⁶ Each vehicle has a different present discounted value

⁵At 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 a negative sign on fuel economy, which would imply that consumers dislike fuel economy.

⁶By "vehicle," we mean a model-by-age combination, where "model" is a highly disaggregated definition capturing essentially all variation in make, nameplate, trim, body type, fuel economy, engine displacement, number of cylinders, and design generation.

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

Under the assumptions that consumers' gas price forecasts match those of oil futures markets and that consumers' real intertemporal cost of funds is six percent, auto consumers appear to be indifferent between one dollar in discounted future gas costs and 76 cents in purchase price. The corresponding "implied discount rate," the discount rate for future gas costs that rationalizes market behavior, is just under 15 percent. We show that the results are robust to a number of potential confounding factors, such as endogenous changes in utilization patterns and market shares in response to gas prices, changes in underlying preferences for "green vehicles," and changes in characteristics over model years.

However, some empirical issues are extremely important. We use the grouping estimator (Wald 1940, Angrist 1991) to address potential measurement error, and we show that failing to do so would cause significant attenuation bias. We show that the result that consumers undervalue gas costs is largely driven by older vehicles: prices for vehicles aged 11-15 years appear to be highly insensitive to gasoline prices, while prices for relatively-new used vehicles move much more closely to the theoretical prediction. In addition, undervaluation appears to be more severe when current gas prices instead of oil futures are used as a proxy for consumers' forecasts of future gas prices.

The data also show clear evidence of sticky information. In models such as Mankiw and Reis (2002), Sims (2003, 2010), and Woodford (2003), information about prices, inflation, and other economic variables takes time to diffuse through the population of consumers and firms, causing delays in responses to news. In this context, sticky information would cause vehicle prices to adjust with a delay to changes in gas price forecasts. Indeed, when the model is estimated using monthly differences instead of the demeaning estimator, vehicle markets appear to be highly insensitive to the current month's change in gas price forecasts. Four to six month lags of gas price changes are conditionally associated with the current month's change in vehicle prices, suggesting that sticky information keeps vehicle markets from updating immediately. These results add a unique microdata-based analysis to the limited set of empirical tests of the sticky information model,

which includes Coibion (2010), Khan and Zhu (2006), Klenow and Willis (2007), and Mankiw and Reis (2003). Aside from being of interest in macroeconomics and finance, sticky information also has important implications for how the undervaluation hypothesis can be tested: identifying off of high-frequency gas price variation can cause the econometrician to falsely conclude that consumers undervalue fuel economy as a product attribute when they are instead inattentive to changes in gas prices.

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 estimated marginal damage of carbon emissions (Greenstone, Kopits, and Wolverton 2011), gasoline consumption imposes an externality of \$0.18 per gallon, or about five percent of the current gasoline price. Thus, if the carbon externality is not internalized, consumers account for about 95 percent of the total social cost of gasoline when they choose between vehicles with different fuel economy ratings. By comparison, an implied discount rate of 15 percent, which is well within the range of Hausman (1979) and other estimates in the literature, would suggest that consumers value only 75 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 our data, 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 details the evidence of sticky information, and Section 7 concludes.

1.1 Related Literature

Within the broader literatures on consumer choice of energy-using durables⁷ and the effects of gasoline prices on vehicle markets,⁸ our analysis is most closely related to four papers. Kahn (1986) tests whether relative prices of used vehicles fully adjust to changes in the relative discounted present value of relative gas prices caused by the gasoline price shocks of the 1970s 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 used car price guides such as the Kelley Blue Book or the National Auto Dealers’ Association Used Car Guide. It is crucial to use transaction data 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. Data from used car price guides could cause us to falsely conclude that vehicle market prices do not fully adjust to changes in gasoline prices.

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. Their identification strategy is different and complementary to ours: they exploit the fact that autos with different odometer readings have different remaining lifetimes, and an increase in gas prices has a larger effect on the remaining PDV of gas costs for autos with lower current odometer readings. One benefit of their approach is that it allows a more extensive set of fixed effects than we use. By construction, however, they are testing a different empirical hypothesis: they ask whether consumers correctly value differences in odometer readings within vehicle models when gas prices change, while we ask whether consumers correctly value differences in fuel economy across models when gas prices change. Conceptually, our hypothesis is closer to the policy-relevant question of

⁷This literature includes Hausman (1979), Dubin and McFadden (1984), Davis and Kilian (2011a), Beresteanu and Li (2011), and others.

⁸Other work that examines how vehicle prices adjust in response to gasoline prices include Sawhill (2008), Langer and Miller (2012), 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. Edelstein and Kilian (2009) document how gas prices affect consumer expenditures, including demand for domestic vs. foreign automobiles. Li, Timmins, and von Haefen (2009) and Klier and Linn (2010) test how the price of gasoline affects sales of new vehicles with different fuel economy ratings.

whether consumers undervalue energy efficiency.

Busse, Knittel, and Zettelmeyer (2012) estimate how changes in gasoline prices affect equilibrium prices and quantities of new and used vehicles in different quartiles of the fuel economy distribution. They then plug these estimates into an Excel spreadsheet which outputs the implied discount rates at which the average prices of vehicles in each MPG quartile fully adjust to changes in gasoline costs. Depending on their assumptions, they find implied discount rates ranging from negative 6.2 percent to positive 20.9 percent. When using assumptions that correspond most closely to ours, they find an implied discount rate for used vehicles of 13 percent.⁹ Our analysis differs by being somewhat more formal: for example, we specify the undervaluation hypothesis from the primitives of a utility function, which helps in understanding the economic meaning of the identifying assumptions. We also uncover additional empirical findings, such as the evidence of sticky information in Section 6.

2 Model

In this section, we specify a static discrete choice model. In Appendix A, we show that under the assumption of stationarity, this simple model can be derived from a dynamic model in which consumers buy and re-sell their vehicles. Consumers derive utility from owning a vehicle and from consuming a numeraire good. We define a "vehicle" as a model-by-age combination, where j indexes models and a indexes age in years. Consumers also can choose an outside option, denoted $j = 0$, which is to own no vehicle. The utility of this outside option is normalized to zero. Consumer i has budget constraint w_i .

The purchase price of vehicle ja at time t as p_{jat} . G_{jat} is the present discounted value (PDV) of future gasoline costs over the vehicle's remaining life. The average consumer's utility from owning and using vehicle ja over its remaining life is $\tilde{\psi}_{jat}$; we call this the "average usage utility." Individual i 's usage utility also has an unobserved deviation from average usage utility, denoted ϵ_{ijat} . The variable $\tilde{\psi}_{ja}$ captures the mean value of average usage utility for vehicle ja across all time periods, and $\tilde{\xi}_{jat} \equiv \tilde{\psi}_{jat} - \tilde{\psi}_{ja}$ is the period-specific deviation.

⁹This 13 percent figure is from Table 7 of Busse, Knittel, and Zettelmeyer (2012). It is the mean of the implicit discount rates for the three different quartile pairs for used vehicles when using vehicle-miles traveled and survival probabilities from NHTSA, which uses the same underlying data as we do.

Consumer i 's indirect utility from vehicle ja at time t is:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G_{jat}) + \tilde{\psi}_{ja} + \tilde{\xi}_{jat} + \epsilon_{ijat} \quad (1)$$

The marginal utility of money is η . The "valuation weight" on fuel costs is γ : 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. This framework is analogous to the attention weight models in Chetty, Looney, and Kroft (2009), DellaVigna (2009, page 349), and other analyses. We assume that both η and γ are constant.

In specifying the basic model, we assume that ϵ takes the extreme value distribution, giving the traditional representative-consumer logit model. Integrating over the distribution of ϵ gives a market-level relationship between prices and shares, which are denoted by s . Typically, this identity is arranged to give the term $\ln s_{jat} - \ln s_{0t}$ on the left:

$$\ln s_{jat} - \ln s_{0t} = -\eta p_{jat} - \eta \gamma G_{jat} + \tilde{\psi}_{ja} + \tilde{\xi}_{jat} \quad (2)$$

Of course, this identity can also be re-arranged with prices on the left:

$$p_{jat} = -\gamma G_{jat} - \frac{1}{\eta} (\ln s_{jat} - \ln s_{0t}) + \psi_{ja} + \xi_{jat} \quad (3)$$

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

3 Data

Our dataset includes average prices, quantities, and characteristics of all used passenger vehicle models registered in the U.S., in monthly cross sections from January 1999 to December 2008. Table 1 presents descriptive statistics.

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. For all model years, we use the EPA's most recent adjustment, which was calculated in 2008. Per Greene *et al.* (2007), we also assume that fuel economy degrades over each vehicle's life at 0.07 MPG per year. Vehicle class designations - 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."

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. To get p_{jat} , we simply collapse the data to the level of mean price for each vehicle in each month, after excluding vehicles rated as low-quality or scrap quality.

Some alternative specifications employ prices 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 obtained monthly mean prices adjusted for customer cash rebates and, if the transaction included a trade-in, the difference between the trade-in vehicle's actual resale value and the negotiated trade-in price. We use the Manheim used vehicle price data in our base specifications because these data include more than twice as many observations at the jat level, while there are fewer than 1000 observations in JDPA that are not in Manheim.

We observe national-level registered quantities of each vehicle 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. These quantity data are matched to the price data using vehicle-specific serial numbers called Vehicle Identification Numbers.

We define the choice set to include all "substitutable" used gasoline-fueled light duty vehicles that have EPA fuel economy ratings and are less than 25 years old. By "substitutable," we mean 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 vehicles¹⁰ 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 refers to a colloquial name such as "Ford Taurus" or "Honda Civic." We define a "model" at the level of make, nameplate, trim, body type, fuel economy, engine displacement, 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-

¹⁰These excluded exotic vehicles are the Acura NSX, Audi R8 and TT, Chrysler Prowler and TC, Cadillac 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.

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.

3.1 Discounted Gasoline Costs

The variable G_{jat} is the present discounted value of lifetime gasoline costs over future years y , for the average driver of vehicle ja at time t :

$$G_{jat} = \sum_{y=0}^L \delta^y \cdot g_y \cdot m_{jay} \cdot f_{jay}^{-1} \cdot \phi_{jay} \quad (4)$$

L denotes the maximum possible lifetime of a vehicle, which we take to be 25 years. The variable g_y is the gasoline price forecast for year y , m_{jay} is expected vehicle-miles traveled, f_{jay} is fuel economy in miles per gallon, ϕ_{jay} is the probability that the vehicle survives to year y conditional on surviving to its current age a , and δ is the annual discount factor.

Our parameter of interest γ is the coefficient on this variable, and decisions we make here mechanically affect the parameter estimates. For example, using a lower discount rate than consumers actually face would inflate G_{jat} , thereby biasing $\hat{\gamma}$ toward zero. Alternatively, understating a vehicle's expected lifetime or usage would deflate G_{jat} , biasing $\hat{\gamma}$ away from zero.

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 includes the age, fuel economy, and vehicle class for each vehicle owned. 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_{jay} for all vehicles in our sample. When we adjust for uncertainty in the first-step estimates of G_{jat} (Murphy and Topel 1985), we use the Huber-White robust standard errors from this regression.

Our base specification assumes that this fitted m_{jay} does not depend on the gasoline price forecast in year y . Either of two separate arguments can justify this. First, empirical estimates of VMT demand show that it is relatively inelastic (Hughes, Knittel, and Sperling 2007, Small and Van Dender 2007, Gillingham 2010, Bento *et al.* 2009, Kilian and Murphy 2011, Davis and Kilian 2011b). 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 Appendix C, however, we derive an alternative approach that makes m_{jay} a function of g_y and allows for non-marginal changes. 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 vehicle age dummies, model year, and fuel economy, with robust standard errors clustered by vehicle, and then predict survival probabilities ϕ_{jay} for all vehicles in our sample.

3.1.2 Discount Rate

The discount rate $r = \delta^{-1} - 1$ is the intertemporal opportunity cost of money: the rate at which a vehicle's purchase price is amortized over future years or future gasoline costs are discounted to the present. To implement this empirically, we calculate the weighted average δ across used vehicle buyers. For vehicle buyers whose marginal dollar comes from a loan or lease, the opportunity cost of paying more to purchase a vehicle 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 2 details the calculation. Using data from the 2001, 2004, and 2007 Surveys of Consumer Finances (SCF), we estimate that 37 percent of used vehicles are financed, while 63 percent are paid for in cash. We transform interest rates reported in the SCF from nominal to real by deflating by inflation expectations implied by Treasury Inflation-Protected Securities. In the SCF, the average real interest rate for used vehicle loans is 6.9 percent. For vehicles purchased with cash, we assume

that the opportunity cost of funds is equity market returns, and the average real return on the S&P 500 between 1945 and 2008 (inclusive) was 5.8 percent. As shown in Table 2, averaging these two discount rates weighted by their share of transactions gives a 6.2 percent discount rate, which we round to six percent for our base specification.

While we think that this is the most reasonable and straightforward approach, it surely is debatable. 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 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. Because annual changes in gasoline prices have very low correlation with market returns, risk adjustment using the Capital Asset Pricing Model would give an interest rate close to the real risk-free rate: about 1.6 percent. In the results section, we show the sensitivity of $\hat{\gamma}$ to alternative discount rates.

3.1.3 Gasoline Price Forecast

We use three measures of vehicle consumers' gasoline price forecasts: seasonally-adjusted retail prices, oil futures prices, and beliefs elicited via surveys. The seasonally-adjusted retail price measure is equivalent to assuming that consumers forecast that gasoline prices are a martingale with seasonal trends. To construct this measure, we take the U.S. City Average Motor Gasoline Retail Prices for all types of gasoline from the U.S. Energy Information Administration's (2011) Monthly Energy Review and convert to real July 2005 dollars.¹¹ We then eliminate seasonal trends by regressing the level of monthly gasoline prices on eleven month dummies, subtracting the fitted values, and re-adjusting so that the mean monthly gasoline price is unchanged.

Our second measure assumes that oil futures markets reflect vehicle consumers' gasoline price forecasts. To construct this measure, we use Light Sweet Crude Oil spot prices from the U.S. Energy Information Administration's (2011) Monthly Energy Review and futures prices from the New York Mercantile Exchange (NYMEX) and the Intercontinental Exchange (ICE). The oil futures prices are transformed into current dollars using inflation expectations implied by Treasury Inflation

¹¹Edelstein and Kilian (2009) and Kilian (2010) also use data on the real price of gasoline.

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.¹² Appendix Table A1 reports annual average retail gas prices and futures-based gas price forecasts. In every year of our sample, there are futures trades for settlement dates up to 6-7 years in the future, and in later years, trades are observed for dates as far as 10 years out.¹³

Our third measure is derived from the Michigan Survey of Consumers (MSC), which elicits beliefs about future gasoline prices, inflation, and other economic variables from a nationally representative sample. Respondents are asked whether they think the price of gasoline will go up, go down, or stay the same during the next five years. Those who think it will change are asked how much the price will increase or decrease. Anderson, Kellogg, and Sallee (2011) analyze these data, and interested readers should refer to their paper for more details. For our analysis, the authors provided us with the mean real gasoline price forecast across all respondents for each month between 1999 and March 2008. Each respondent's forecast is deflated to current dollars as of the date of the survey using his or her own inflation expectation.

Figure 1 plots seasonally-adjusted retail gasoline prices along with the gasoline prices implied by oil futures at three different horizons. The figure illustrates that futures prices differ from retail prices in two important ways. First, the futures market does not believe that gasoline prices are a martingale. Especially when spot prices spiked in mid-2008, the futures market expected prices to eventually decrease at least slightly. As a result, G averages seven percent larger under the martingale forecast than the futures forecast. This suggests that the $\hat{\gamma}$ estimated under the martingale forecast will be smaller than the $\hat{\gamma}$ with the futures-based forecast. Second, while the two series follow the same trends, retail prices are substantially more volatile from month to month. This is not a spurious result of adjusting for seasonality: retail prices are also more volatile when we do not condition on month-of-year effects. As we shall see, this will cause $\hat{\gamma}$ to differ substantially

¹²Specifically, 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. The reason we transform oil futures prices into gasoline prices instead of directly using gasoline futures is that gasoline futures are typically not traded for time horizons longer than three years, while oil futures are traded, albeit with limited liquidity, at time horizons of up to ten years.

¹³To 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.

between the two forecasts when using a monthly difference estimator.

It is important to acknowledge several limitations of oil futures price data in light of recent work in this area. First, Kilian (2010) has shown that in the short run, there can be important deviations between crude oil and gasoline prices. Over our sample period, monthly oil spot prices predict 94 percent of the variance in monthly gasoline prices, not 100 percent. However, what matters for our analysis is the expected future relationship between oil and gasoline: oil futures could be a good measure of long-run gasoline price forecasts even if the short-run relationship does not have an R^2 of one. Second, Alquist and Kilian (2010), Alquist, Kilian, and Vigfusson (2011), and Baumeister and Kilian (2012) have shown that oil futures prices are typically not as good as current oil prices at predicting actual future oil prices at both short-term horizons of 12 months and longer-term horizons of up to seven years. This suggests that factors other than oil price expectations can affect futures markets, which would introduce noise in our estimates. Third, oil futures are only traded with high liquidity for settlement dates less than two to three years in the future. While this is a concern, it is not clear if this would bias the estimates in any particular direction.

3.2 Vehicle Price Trends

Three related facts about vehicle prices have important implications for our empirical strategy. First, as Figure 2 illustrates, low-MPG vehicles are 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 indicator variables, for all models of model years 1984 through 2008 in the Ward's Automotive Yearbook data. Between 15 and 30 MPG, which includes nearly all of our dataset, MSRPs decline from \$30,000 to \$11,000. This figure corroborates concerns about the traditional cross-sectional approaches to estimating γ : low-MPG and high-MPG vehicles are systematically different, and some of these differences may be unobservable. The figure also means that we must control for factors such as depreciation that differentially affect vehicles with higher price levels and might be correlated with gas prices.

Second, Figure 3 shows that vehicle prices have systematic seasonal patterns, and they tend to decrease within each year. Because this systematic depreciation is approximately a percentage of price and lower-MPG vehicles have higher prices, these patterns affect the price levels of low-MPG

vehicles more than high-MPG vehicles. Since gas price forecasts 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 falsely attribute seasonality to adjustments caused by gas price forecasts.

Third, there are other underlying shifts in seasonally-adjusted average prices that could affect low-MPG vs. high-MPG vehicles differently. If correlated with changes in gasoline prices, such shifts could bias our estimator. The double blue line on Figure 3 shows that after demeaning prices within vehicles and removing seasonal trends using month dummies, 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 mid- 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. We therefore document whether the estimates are sensitive to excluding the period before 2004.

Demeaned and seasonally-detrended 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. Similarly, we would certainly not want to include data from 2009, when the Cash for Clunkers program significantly changed used vehicle markets, differentially affecting prices of low-MPG vs. high-MPG vehicles.

4 Estimation

4.1 Empirical Strategy

We use vehicle fixed effects and look within the same vehicle over time as gasoline price forecasts change. The benefit of a panel is that the fixed effects soak up unobserved vehicle characteristics

that may be correlated with MPG: formally, our estimator is still unbiased even if $E[G\psi] \neq 0$.

We move from Equation (3) to our base specification estimating equation in three steps. First, we define an econometric error term that contains ξ and the market share term: $\varepsilon \equiv -\frac{1}{\eta} \ln s_{jat} + \xi_{jat}$. Second, we add time dummies τ_t , which absorb the outside option share s_{0t} and any shift in the overall market price level. Third, we define f^{50} as the median of the MPG distribution, and we include ν_{jat} , 11 indicator variables for month-of-year interacted with $1(f_{ja} > f^{50})$. As we shall discuss momentarily, these control for the differential seasonal trends in the prices of low- vs. high-MPG vehicles. Our base specification is:

$$p_{jat} = -\gamma G_{jat} + \tau_t + \nu_{jat} + \psi_{ja} + \varepsilon_{jat} \quad (5)$$

Equation (5) 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. If vehicle prices respond more than predicted to gasoline costs, we conclude that the market overvalues gasoline costs.¹⁴

Figure 4 illustrates our identifying variation by graphing the average G for an example month within age and fuel economy groups. Younger vehicles and lower-MPG vehicles have higher G . Equation (5) is identified by interacting this cross-sectional variation by the time-series variation in gasoline price forecasts illustrated in Figure 1. As gasoline price forecasts fluctuate over time, the bars in Figure 4 extend and contract proportionally. A given change in gas price forecasts has a larger effect on the level of G for younger and lower-MPG vehicles, and we test for whether relative prices move correspondingly.

Equation (5) is a consistent estimator of γ if $E[G\varepsilon|(\tau, \nu, \psi)] = 0$. Because ε contains both s and ξ , this actually contains two different economically-meaningful identifying assumptions:

Assumption 1: $E[Gs|(\tau, \nu, \psi)] = 0$

Assumption 2: $E[G\xi|(\tau, \nu, \psi)] = 0$

¹⁴Although we derive it differently, this is fundamentally the same estimating equation used by Kahn (1986) and Busse, Knittel, and Zettelmeyer (2012). The latter paper uses a slightly different configuration of fixed effects and controls. Sallee, West, and Fan (2009) allow G to vary at the transaction level, instead of at the jat level, based on the odometer reading of the specific auto being transacted. They can then insert jat -level fixed effects. This allows them to relax our two identifying assumptions and substitute an alternative assumption, which is that the within- jat relationship between odometer readings and transaction prices is not affected by unobserved factors that are correlated with gas prices.

In words, Assumption 1 is that gasoline costs are uncorrelated with market shares, conditional on time dummies and vehicle fixed effects. For new vehicles, this assumption would not hold: sales of low-fuel economy vehicles decrease relative to high-fuel economy vehicles when gas prices increase, as documented by Busse, Knittel, and Zettelmeyer (2012), Klier and Linn (2010), and Li, Timmins, and von Haefen (2009). This assumption might also not hold for older used vehicles, whose scrappage rates respond most to changes in gasoline prices (Li, Timmins, and von Haefen 2009). Therefore, we limit our sample to used vehicles less than or equal to 15 years old. In Appendix B, we show that weakening this assumption by allowing for endogenous market shares does not significantly affect the estimates, largely because we have limited our sample to the ages for which the conditional correlation between gasoline prices and market shares is insignificant.

Assumption 2 states that changes in discounted gasoline costs are not correlated with changes in preferences and unobserved characteristics that vary over the different model years within a generation of a vehicle. In Section 5, we discuss and test several potential violations of this assumption and document that they do not appear to substantially affect the results.

4.2 Measurement Error and the Grouping Estimator

Jerry Hausman's (2001) "Iron Law of Econometrics" is that the magnitude of a parameter estimate is usually smaller in absolute value than expected. This is due to attenuation bias that can occur even in seemingly high-quality data. In most applications in which the null hypothesis is that a parameter equals zero, this makes a test more conservative. By contrast, in our test of whether $\hat{\gamma} = 1$, this "Iron Law" would make us more likely to falsely reject, meaning that we could falsely conclude that consumers undervalue gasoline costs. One potential source of measurement error is the EPA fuel economy ratings, as discussed by Sallee (2010). As in any other application, the use of fixed effects exacerbates attenuation bias, as even small amounts of measurement error in G worsen the signal-to-noise ratio in $G|\psi$.

We address measurement error using the grouping estimator, which is just Wald's (1940) binary instrumental variables (IV) estimator generalized to the case with many group indicator variables. Wald's (1940) original objective in developing the estimator was to address measurement error, and Ashenfelter (1984), Angrist (1991), and others have used the grouping estimator for this same

purpose in various applications. In our base specification, we place each jat -level observation into one of 222 mutually exclusive and exhaustive groups: above- and below-median MPG vehicles for each of the 111 months in the sample. This grouping retains the fundamental identifying variation, which comes from how time-series changes in gas price forecasts affect G for high- vs. low-MPG vehicles. However, identifying γ off of group-level aggregates means that we average over any observation-level measurement error.

Formally, we define $Z_{jat}^u = 1(f_{ja} > f^{50}) \cdot 1(t = u)$ as an indicator variable for whether observation jat occurs at time u and has above-median MPG. We denote \mathbf{Z}_{jat} as the vector of 100 Z_{jat}^u variables for every month of the sample beginning with January 2000, and we instrument for G_{jat} with \mathbf{Z}_{jat} . Between instruments \mathbf{Z}_{jat} , time controls τ_t , and month-by-MPG group controls ν_{jat} , we have a full set of indicator variables in the "first stage" for each of the 222 groups, and the "second stage" fitted values of G are essentially group-level averages.

We also present results for groups at different levels of aggregation. Aggregating observations into fewer larger groups allows us to average over measurement error that is more severe or correlated across observations. However, using fewer groups effectively reduces the amount of identifying variation, which increases the IV standard errors.

We need not control for any variables that vary at a more disaggregated level than \mathbf{Z}_{jat} . For example, consider our controls ν_{jat} that soak up differential monthly price trends for above- vs. below-median MPG vehicles. The highest-MPG quartile also has different monthly price trends than the second-highest MPG quartile. However, we do not need to control for this because the fitted values of G do not vary between these two quartiles, and thus γ is not identified off of these differential price trends. In alternative specifications where we define \mathbf{Z}_{jat} by more disaggregated MPG groups, we also control for differential price trends for each separate MPG group.

5 Results

5.1 Graphical

Figure 5 illustrates our identification. On the vertical axis is the difference between the mean price across transactions 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. 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 .

Figure 6 is a plot of the price and gas cost residuals that identify γ . The solid blue line is the difference between average $G|(\nu, \psi)$, the seasonally detrended and demeaned 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 uses the martingale forecast instead of the futures forecast. The double black line is the difference between average $p|(\nu, \psi)$, the seasonally detrended and demeaned 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 at least somewhat responsive to relative gasoline costs. However, vehicle prices do not appear to respond very much to the high-frequency volatility in retail gasoline prices. To test whether prices are as responsive as theory predicts under the martingale and futures-based forecasts, we now turn to the formal estimates.

5.2 Base Specification

For our "base specification," we estimate Equation (5) using the fixed effects (demeaning) estimator while grouping \mathbf{Z}_{jat} at the level of above- and below-median MPG by time, assuming a six percent discount rate and using data from January 1999 through March 2008. To reflect the fact that p_{jat} is an average across varying numbers of transactions, we weight the jat -level observations by the number of transactions. We are careful to call this our "base specification," indicating that the alternative specifications modify individual assumptions from this base, but we do not call it our "preferred specification." This is because different readers may prefer different sets of alternative assumptions.

Table 3 reports the estimated coefficients on G from Equation (5), which correspond to $-\hat{\gamma}$. We denote the futures-based and martingale gas price forecasts as G^f and G^m , respectively. Columns 1 and 2 show that the estimated $\hat{\gamma}$ coefficients are 0.76 and 0.55. The smaller $\hat{\gamma}$ for the martingale forecast is consistent with our earlier discussion of the fact that the martingale forecast implies

higher levels of G and more monthly volatility compared to the futures forecast.

Using consumers' gas price forecasts elicited from the Michigan Survey of Consumers, Anderson, Kellogg, and Sallee (2011) show that the average consumer believes that real gasoline prices in five years will be the same as they are at the time of the survey. This suggests that the martingale-based $\hat{\gamma}$ will be very close to the MSC-based estimate. This is indeed the case: using the MSC-based gas price forecast provided by Anderson, Kellogg, and Sallee, $\hat{\gamma} = 0.51$ (SE=0.029), not statistically distinguishable from the martingale-based $\hat{\gamma}$.

One way to test whether it is more realistic to assume the futures-based or martingale gas price forecast is to test which fits the data better in a "horse race." Because G^m and G^f are highly correlated, we do not include both in a regression. Instead, we estimate Equation (5) using G^f and add $[G^m - G^f]$ as an additional explanatory variable. Column 3 presents the results of this regression. Neither the futures nor the martingale assumption appears to fully capture consumers' forecasts: conditional on G^f , $[G^m - G^f]$ is still associated with vehicle prices. (Mechanically, this also means that conditional on G^m , $[G^f - G^m]$ is still associated with vehicle prices.) Of course, these results are conditional on the model being correctly specified. Because there appears to be some evidence for both forecasts under this model, we present both sets of results in most of the robustness checks that follow.

One way of putting $\hat{\gamma}$ in context is to compare the observed vehicle price adjustments to the amounts predicted by theory. Consider a hypothetical set of used vehicles that have different fuel economy ratings but would all be driven 12,000 miles per year over a remaining lifetime of seven years. The double red line in Figure 7 reflects the change in vehicle prices that would be expected in response to a \$1 increase in the gas price forecast, relative to the price of a 25 MPG vehicle. The blue "Predicted Price Change" line illustrates the change in relative prices predicted by $\hat{\gamma} = 0.76$ from the futures-based gas forecast.

For example, $\hat{\gamma} = 0.76$ predicts that the price of a 20 MPG vehicle decreases by \$506 relative to a 25 MPG vehicle. In comparison, $\gamma = 1$ predicts that the relative price should decrease by \$670, which is \$164 more. As another example, $\gamma = 1$ predicts that the price of a 30 MPG vehicle should increase by \$2230 relative to a 15 MPG vehicle. By contrast, $\hat{\gamma} = 0.76$ suggests that the relative price adjusts by 72 percent of that amount, or \$1690. The dollar values of the apparent mispricing

are certainly not trivial.

5.3 Standard Errors

Throughout the paper, we report robust standard errors, clustered at the level of model j by age a . These standard errors are unbiased in the presence of serial correlation over time in the price of, for example, a three-year-old Honda Civic DX sedan. Clustering at the level of "nameplate" by age gives a standard error is 0.054. 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 0.086. Analogously, this is unbiased in the presence of serial correlation in the price of all Honda Civic DX sedans of any age.

In addition, all standard errors can be adjusted to 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. 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 $\hat{\gamma}$ for each draw. Using the base specification with G^f as an example, the clustered, robust IV standard error is 0.046. The standard deviation of the bootstrapped set of $\hat{\gamma}$ estimates is 0.023. Adding together these two variances and taking the square root gives an adjusted standard error of 0.052.

5.4 Alternative Specifications

5.4.1 Level of Grouping

Correcting for measurement error appears to be extremely important. Rows 1 through 3 of Table 4 present estimates of γ using increasingly disaggregated grouping estimators, while Row 4 is the ungrouped OLS estimator. As we disaggregate, the estimated values of γ drop toward the OLS estimate. This suggests that measurement error significantly biases the OLS estimates, and even some of the disaggregated grouping estimators. This pattern can arise either if the variance of the measurement error is large or if it is correlated across observations.

5.4.2 Alternative Time Periods

As discussed earlier, vehicle markets changed over the study period. How sensitive is $\hat{\gamma}$ to the choice of time periods? Rows 11 and 12 in Table 4 repeat the base specification for 2004 through March 2008 and for 1999 through the end of 2008, respectively. In both cases, $\hat{\gamma}$ is closer to one than the 1999-March 2008 base specification, although both are still statistically less than one for both G^f and G^m . In row 11, 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 bias the estimator. In row 12, including the gas price spike of 2008 increases $\hat{\gamma}$ more under G^f than under G^m . This is consistent with the fact that futures markets did not expect the high spot prices in 2008 to be sustained, which mechanically decreases G and thus increases $\hat{\gamma}$.

There are two explanations for the fact that including the latter part of 2008 moves $\hat{\gamma}$ statistically closer to one under the futures-based gas price forecast. First, changes in vehicle markets unrelated to changes in gasoline price forecasts could bias the estimated γ . Second, the true γ may have increased over this period as gas prices fluctuated substantially. This would be consistent with models of endogenous or "rational" inattention in which consumers pay more attention to more important attributes (Gabaix 2012, Sallee 2011). It would also be consistent with models in which large changes cause consumers to update beliefs between coarse categories, for example from gas costs being "inconsequential" to gas costs being "high" (Mullainathan 2002).

5.4.3 Alternative Discount Rates

Rows 21-25 of Table 4 show the sensitivity of $\hat{\gamma}$ to the assumed discount rate. At a discount rate just higher than 11 percent, we fail to reject $\gamma = 1$ with 90 percent confidence using the Murphy-Topel standard errors and the futures forecast. The "implied discount rate," the discount rate that rationalizes the data by giving $\gamma = 1$, is just under 15 percent for the futures forecast and just over 24 percent for the martingale.

5.4.4 Changes in Characteristics

The identifying assumption that $E[G\xi|(\tau, \nu, \psi)] = 0$ would be violated if a model's characteristics change in ways that are correlated with G . For example, since gasoline price forecasts rise on average during the study period, if characteristics improve more over model years for high-MPG models, we would misattribute these vehicles' increased desirability to changes in gasoline price forecasts. Our policy of redefining a new generation of a model as a different j addresses the bulk of these concerns. Even within a generation, however, there can be some variation in observable characteristics (Knittel 2011, figure 1).

An additional suggestive test is to add controls for observable characteristics to the estimation. Rows 31-35 re-estimate the base specification with progressively more controls for observable characteristics: horsepower, curb weight, wheelbase, anti-lock brakes, stability control, and traction control. The estimated $\hat{\gamma}$ remains very similar. 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). Therefore, this suggests that changes in unobservable characteristics might not bias our estimates.

5.4.5 Changes in Preferences

Aside from representing changes in vehicle characteristics, changes in ξ_{jat} can also represent changes in preferences for a vehicle with the same characteristics. Since gasoline price forecasts rise on average during the study period, differential trends in preferences for low- vs. high-MPG vehicles would violate the assumption that $E[G\xi|(\tau, \nu, \psi)] = 0$, which would bias $\hat{\gamma}$.

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 $\hat{\gamma}$ upward, as it would increase the prices of high-MPG vehicles over the study period as gasoline price forecasts rose. To test this, we exclude hybrids and the top 60 most green vehicles ranked by Yahoo (2009). Rows 41 and 42 of Table 4 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. This would bias $\hat{\gamma}$ 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 price forecasts rose. To test this, Rows 43 through 45 exclude SUVs, minivans, and all cars, respectively, from the estimation. While the point estimates change somewhat, the differences are not statistically significant, except for when excluding cars and using the martingale gas price forecast.

5.4.6 Age-by-Time Controls

Rows 51-53 of Table 4 include age-by-time controls to account for potential changes in depreciation patterns over time. These controls address bias that would occur if depreciation patterns changed as gas prices rose over the study period and the age composition of above- vs. below-median MPG groups changed. For example, older vehicles could have become systematically more valuable in later years, and there could be increasingly more below-median or above-median MPG vehicles on the road. The table shows that identifying $\hat{\gamma}$ only off of within-age variation in G does not significantly change the estimates.

5.5 Observation Weights and Heterogeneity by Vehicle Age

A natural question to ask is whether γ might be heterogeneous across consumer groups. For example, the types of consumers that buy newer or older vehicles might be more or less informed about or attentive to fuel economy. In fact, there is significant heterogeneity in $\hat{\gamma}$ by vehicle age. Column 1 of Table 5 displays estimates of Equation (5) for specific vehicle age groups. The $\hat{\gamma}$ is 0.93 for vehicles aged 1-3, but it drops to 0.26 for vehicles aged 11-15. All estimates in Table 5 use the futures-based gas price forecast, but the martingale forecast similarly shows that $\hat{\gamma}$ declines with vehicle age.

The differences in $\hat{\gamma}$ depend mechanically on the differences in G . Returning to Figure 4, we can see how G varies with age. In June 2004, the transaction-weighted average values of G for above-median and below-median MPG vehicles aged 11-15 were \$2620 and \$4150, respectively.

These are about 37 percent of the values for vehicles aged 1-5. These figures depend on underlying data for survival probabilities and VMT. According to the registration data, the average vehicle between 11 and 15 years old will survive for 6.1 more years. According to the National Household Travel Survey odometer readings, the average vehicle between 11 and 20 years old will be driven 8980 miles per year, conditional on survival.

Several factors could explain the differences across ages. First, we could systematically overstate G for the nationwide population of older vehicles by overestimating either VMT or survival probabilities. This appears to be unlikely. Because vehicle registrations must be renewed every year at some cost, it is unlikely that many vehicles would have their registrations renewed if not in use. Thus, it seems most accurate to estimate survival probabilities from vehicle registration data, and it is unlikely that this would overstate survival probabilities. However, we can also estimate survival probabilities using the cross-sectional age distribution observed in the National Household Travel Survey. There are fewer older vehicles observed in the NHTS than in the registration data, suggesting that the NHTS might under-sample households that own the oldest vehicles, even after weighting for national representativeness on other observables. As a result, using the NHTS to calculate survival probabilities should make G smaller for new vehicles, thereby increasing γ . Depending on the shape of the survival functions, this could change the estimated γ for older vehicles in either direction. Column 2 of Table 5 presents estimates of γ using the NHTS survival probabilities. All of the point estimates increase, but none of them change significantly, and the pattern of decreasing γ with vehicle age is unchanged.

Second, even if our G is appropriate for the US population of older vehicles, the population of vehicles going through auctions might have smaller G . This would happen if auctioned vehicles had lower future VMT or survival probabilities, perhaps because they are low quality or being sold for scrap. The data also do not support this. The average odometer reads for 11 to 15 year old vehicles differ by only a few percent between Manheim auctioned vehicles and vehicles in the NHTS, so auctioned vehicles have not been more heavily used. Furthermore, in constructing vehicle prices p_{jat} , we exclude transactions in which the auctioneers rated the auto as low-quality or scrap quality. As a robustness check, column 3 presents these same specifications using JDPa retail prices, which are for autos that are in good enough condition for retail sale. The significant decrease in $\hat{\gamma}$ for

older vehicles remains.

Third, the identifying assumption that $E[Gs|(\tau, \nu, \psi)] = 0$ might be increasingly violated for older vehicles. As gas prices rise, increased scrappage of older gas guzzlers would raise the prices of these vehicles both because demand is downward-sloping and because selective scrappage of the lowest-quality gas guzzlers would raise the average quality of remaining ones. However, while there is a strong negative correlation between G and market share for vehicles older than 15 years (which are excluded from our regressions), the correlation between G and s conditional on (τ, ν, ψ) for 11-15 year old vehicles is not statistically significant, and the point estimate is actually positive.

Fourth, buyers of older vehicles may have higher intertemporal opportunity costs of capital, i.e. higher discount rates r , compared to buyers of younger vehicles. If this were the case, our use of the average discount rate still gives an unbiased estimate of the market average γ . However, it would cause us to overstate γ for younger vehicles and overstate γ for older vehicles, and it would give the pattern of decreasing γ with age observed in Table 5. The SCF data do not support this hypothesis: auto loan interest rates are not statistically significantly different by vehicle age, and the standard errors are tight enough to rule out that the average interest rate for 11-15 year old vehicles is greater than 0.5 percentage points above the average rate for 1-5 year old vehicles. However, Adams, Einav, and Levin (2009) analyze subprime auto loans with interest rates of 25-30 percent, and consumers buying older vehicles may have very high intertemporal opportunity costs of capital because they cannot get any loan.

One implication of the differences in $\hat{\gamma}$ across ages is that regression weighting matters. In our base specifications, we weight jat -level observations by the number of observed transactions. Column 5 of Table 5 shows that most vehicles that pass through auctions are less than five years old, meaning that the jat -level observations for newer vehicles have more underlying transactions and thus are weighted more heavily. The difference in γ coefficients across ages means that the estimated γ for all ages should be smaller when weighting observations equally instead of by number of transactions. The first row of column 4 confirms this, showing that $\hat{\gamma} = 0.62$ for equal weights. Comparing columns 1 and 4 shows that when the regressions are restricted to particular age groups, the γ coefficient does not depend much on how observations are weighted.

6 Sticky Information

Mankiw and Reis (2002), Sims (2003), Woodford (2003), and others propose models of "sticky information," in which macroeconomic news diffuses slowly through the population of consumers and firms. This causes prices, savings, investment, and other decisions to respond to information with a delay. Sticky information could explain the visual evidence in Figure 6 that vehicle market prices do not respond to high-frequency fluctuations in retail gasoline prices.

Interestingly, neither of the two surveys that have elicited consumers' beliefs about gasoline prices were structured to provide evidence on sticky information. The Vehicle Ownership and Alternatives Survey analyzed by Allcott (2011) and Allcott (Forthcoming) asks consumers what current gasoline prices are in their area, and shows that responses are very close to state-level averages. However, the survey was carried out in October 2010, at a time when gas prices had not fluctuated by more than \$0.10 per gallon for eight months. If the survey were repeated at a time of more volatility, the sticky information model predicts that beliefs would lag current prices.

The Michigan Survey of Consumers (Anderson, Kellogg, and Sallee 2011) is a repeated cross section that covers multiple periods of gas price volatility. However, it does not ask consumers their beliefs about current gas prices. Instead, it only asks whether consumers believe that prices will increase or decrease relative to "current prices." We do not know what consumers believe are "current prices" in months when retail gasoline prices have changed significantly.

One prediction of a sticky information model is that γ should be attenuated when using the difference estimator instead of the demeaning estimator. This is because the former is identified only off of the variation at monthly frequency, while the latter is identified off of demeaned levels over the entire period that the vehicle is observed. Table 6 tests this by estimating Equation (5) using the difference estimator instead of the fixed effects (demeaning) estimator. In this table, unlike Tables 4 and 5, we report the coefficient on the G variable, which is $-\gamma$, instead of γ itself. Columns 1-3 use the futures-based forecast G^f , while columns 4-6 use the martingale forecast G^m . Columns 1 and 4 show that the contemporaneous differences in G are negatively associated with contemporaneous differences in vehicle prices, as would be expected. However, the estimates are significantly attenuated relative to the fixed effects estimates, especially for the martingale forecast.

Column 2, 3, 4, and 6 add lags of the change in G . These results show that changes in vehicle prices are strongly associated with changes in gasoline prices as far as three to four months in the past. The association becomes smaller and statistically insignificant by the fifth or sixth monthly lag.

Table 7 builds on these results by estimating Equation (5) using more aggregated units of time t instead of monthly time units used in the rest of the paper. For example, the "2-Month Units" row presents estimates with prices and gasoline costs aggregated to the bimonthly level. The fixed effects estimates of γ in the top panel do not change with this increased aggregation, as the demeaning estimator identifies γ off of lower-frequency variation. However, the estimates from the difference estimator move toward the fixed effects estimates as we aggregate over more months. This provides additional evidence that markets do not immediately adjust to high-frequency gasoline price variation.

Klenow and Willis (2007) use the data underlying the Consumer Price Index to document that US retailers set prices based on information that is up to one year old. Our results complement theirs by providing additional empirical support for the sticky information model in this particular market. Table 6 suggests that changes in gasoline price forecasts take four to six months to be full incorporated into vehicle market prices. This could either be because consumers are not aware that gasoline prices have changed, or because they somehow do not incorporate this knowledge into their willingness to pay when shopping for vehicles.

These results also highlight the distinction between models of rational inattention in macroeconomics and finance, in which agents do not immediately update beliefs about time-varying decision variables, and consumer choice models of static inattention, such as Gabaix and Laibson (2006) or Chetty, Looney, and Kroft (2009). Energy efficiency policies such as CAFE standards are motivated by the latter form of inattention or some similarly static misoptimization, not by inattention to high-frequency fluctuations in energy prices. While our strategy of exploiting time-series variation in energy costs allows us to sweep our unobserved product attributes using fixed effects, it is important to account for the fact that consumers appear to update gasoline price forecasts with some delay. Using quarterly or higher-frequency variation and a difference estimator could cause an analyst to misattribute inattention to high-frequency fluctuations in gasoline costs to static

undervaluation of fuel economy as a product attribute.

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. We propose a different and relatively plausible approach to estimating demand for energy efficiency, which exploits the interaction of time series variation in gasoline price forecasts with cross-sectional variation in fuel economy ratings of different vehicles. In our base specification assuming the futures-based gas price forecast, consumers appear to be indifferent between one dollar in discounted future gas costs and only 76 cents in vehicle purchase price. The parameter estimates are insensitive to many factors, including the potential endogeneity of market shares and vehicle-miles traveled and potential changes over time in consumer preferences or vehicle characteristics. However, assuming a lower discount rate or using the martingale gas price forecast strengthen the qualitative conclusion that $\gamma < 1$. On the other hand, other specifications move $\hat{\gamma}$ closer to one, including plausibly higher discount rates and adding the 2008 recession to the sample. Although $\hat{\gamma}$ is less than one in most of our alternative specifications, it is not hard to imagine plausible combinations of alternative assumptions that could generate this result. Importantly, the parameter estimates are very different for subsamples of younger vs. older vehicles.

Given the uncertainty in these estimates, it is worth returning to our calculation from the introduction. Even if γ is closer to one than our base specifications suggest, 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 potential 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 (March), pages 49-84.
- [3] Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2012). "Energy Policy with Externalities and Internalities." NBER Working Paper No. 17977 (September)
- [4] Allcott, Hunt (2011). "Consumers' Perceptions and Misperceptions of Energy Costs." *American Economic Review*, Vol. 101, No. 2 (May), pages 98-104.
- [5] Allcott, Hunt (Forthcoming). "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market." *American Economic Journal: Economic Policy*.
- [6] Alquist, Ron, and Lutz Kilian (2010), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, Vol. 25, No. 4, pages 539-573.
- [7] Alquist, Ron, Lutz Kilian, and Robert Vigfusson (2011). "Forecasting the Price of Oil." In Elliot, Graham, and Allan Timmerman, Eds., *Handbook of Economic Forecasting*, Vol. 2, forthcoming.
- [8] 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.
- [9] Anderson, Soren, Ryan Kellogg, and James Sallee (2011). "What Do Consumers Believe About Future Gasoline Prices." NBER Working Paper No. 16974 (April).
- [10] 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.
- [11] Arellano, Manuel (1987). "Computing Robust Standard Errors for Within-Groups Estimators." *Oxford Bulletin of Economics and Statistics*, Vol. 49, No. 4 (November), pages 431-434.
- [12] 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.
- [13] 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.
- [14] Austin, David (2008). "Effects of Gasoline Prices on Driving Behavior and Vehicle Markets." U.S. Congressional Budget Office Working Paper (January).
- [15] 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.
- [16] Baumeister, Christiane, and Lutz Kilian (2012). "Real-Time Forecasts of the Real Price of Oil." *Journal of Business and Economic Statistics*, forthcoming.
- [17] Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen (2009). "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes." *American Economic Review*, Vol. 99, pages 667-699.
- [18] Beresteanu, Ari, and Shanjun Li (2011). "Gasoline Prices, Government Support and the Demand for Hybrid Vehicles in the U.S." *International Economic Review*, Vol. 52, No. 1 (February), pages 161-182.
- [19] Berry, Steven (1994). "Estimating Discrete Choice Models of Product Differentiation." *RAND Journal of Economics*, Vol. 23, No. 2 (Summer), pages 242-262.

- [20] Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium." *Econometrica*, Vol. 63, No. 4 (July), pages 841-890.
- [21] 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.
- [22] Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer (2012). "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review*, forthcoming.
- [23] 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.
- [24] Coibion, Olivier (2010). "Testing the Sticky Information Phillips Curve." *Review of Economics and Statistics*, Vol. 92, No. 1, pages 87-101.
- [25] Cragg, John, and Stephen Donald (1993). "Testing Identifiability and Specification in Instrumental Variable Models," *Econometric Theory*, Vol. 9, No. 2 (April), pages 222-240.
- [26] Davis, Lucas, and Lutz Kilian (2011a). "The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas." *Journal of Political Economy*, Vol. 119, No. 2 (April), pages 212-241.
- [27] Davis, Lucas, and Lutz Kilian (2011b). "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics*, Vol. 26, No. 7 (November-December), pages 1187-1214.
- [28] DellaVigna, Stefano (2009). "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, Vol. 17, No. 2, pages 315-372.
- [29] 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.
- [30] Dubin, Jeffrey (1992). "Market Barriers to Conservation: Are Implicit Discount Rates Too High?" Social Science Working Paper 802, California Institute of Technology (July).
- [31] 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.
- [32] Edelman, Paul, and Lutz Kilian (2009). "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, Vol. 56, pages 766-779.
- [33] Espey, Molly, and Santosh Nair (2005). "Automobile Fuel Economy: What Is It Worth?" *Contemporary Economic Policy*, Vol. 23, No. 3 (July), pages 317-323.
- [34] Finkelstein, Amy (2009). "E-ZTAX: Tax Salience and Tax Rates." *Quarterly Journal of Economics*, Vol. 124, No. 3 (August), pages 969-1010.
- [35] 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.
- [36] Gabaix, Xavier (2012). "A Sparsity-Based Model of Bounded Rationality." Working Paper, New York University (September).
- [37] 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.
- [38] Gillingham, Kenneth (2010). "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock in California." Working Paper, Stanford University (September).
- [39] Goldberg, Pinelopi (1998). "The Effects of the Corporate Average Fuel Economy Standards in the US." *Journal of Industrial Economics*, Vol. 46, pages 1-33.
- [40] 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.

- [41] 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.
- [42] Greenstone, Michael, Elizabeth Kopits, and Maryann Wolverton (2011). "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation." NBER Working Paper No. 16913 (March).
- [43] 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.
- [44] Hausman, Jerry (2001). "Mismeasured Variables in Econometric Analysis: Problems from the Right and Problems from the Left." *Journal of Economic Perspectives*, Vol. 15, No. 4 (Fall), pages 57-67.
- [45] 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.
- [46] 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.
- [47] 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).
- [48] Jaffe, Adam, and Robert Stavins (1994). "The Energy Paradox and the Diffusion of Conservation Technology." *Resource and Energy Economics*, Vol. 16, pages 91-122.
- [49] 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.
- [50] Khan, Hashmat, and Zhenhua Zhu (2006). "Estimates of the Sticky-Information Phillips Curve for the United States." *Journal of Money, Credit, and Banking*, Vol. 38, No. 1, pages 195-207.
- [51] Kilian, Lutz (2010). "Explaining Fluctuations in Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market," *The Energy Journal*, Vol. 31, No. 2 (April), pages 87-104
- [52] Kilian, Lutz, and Daniel Murphy (2011). "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." Working Paper, University of Michigan (May).
- [53] 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).
- [54] Klenow, Peter, and Jonathan Willis (2007). "Sticky Information and Sticky Prices." *Journal of Monetary Economics*, Vol. 54, pages 79-99.
- [55] Klier, Thomas, and Joshua Linn (2010). "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data." *American Economic Journal: Economic Policy*, Vol. 2, No. 3, pages 134-153.
- [56] Knittel, Christopher (2011). "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review*, Vol. 101, No. 7, pages 3368-3399.
- [57] Langer, Ashley, and Nathan Miller (2012). "Automakers' Short-Run Responses to Changing Gasoline Prices." *Review of Economics and Statistics*, forthcoming.
- [58] 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.
- [59] Manheim (2009). "Used Car Market Report." 14th Edition. Atlanta, GA: Manheim Consulting.
- [60] Mankiw, N. Gregory, and Ricardo Reis (2002). "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics*, Vol. 117, No. 4, pages 1295-1328.

- [61] Mankiw, N. Gregory, and Ricardo Reis (2003). "A Model of Monetary Non-Neutrality and Structural Slumps." in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (Eds.), Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps. Princeton, NJ: Princeton University Press.
- [62] Mullainathan, Sendhil (2002). "Thinking Through Categories." Working Paper, MIT (April).
- [63] 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.
- [64] 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).
- [65] 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.
- [66] Parry, Ian, Winston Harrington, and Margaret Walls (2007). "Automobile Externalities and Policies." *Journal of Economic Literature*, Vol. 45, pages 374-400.
- [67] Reis, Ricardo (2006). "Inattentive Consumers." *Journal of Monetary Economics*, Vol. 53, No. 8 (November), pages 1761-1800.
- [68] Rust, John (1985). "Stationary Equilibrium in a Market for Durable Assets." *Econometrica*, Vol. 53, No. 4 (July), pages 783-805.
- [69] Sallee, James (2010). "The Taxation of Fuel Economy." In Brown, Jeffrey, Ed., Tax Policy and the Economy, Vol. 25. Cambridge, MA: National Bureau of Economic Research.
- [70] Sallee, James (2011). "Rational Inattention and Energy Efficiency." Working Paper, University of Chicago (June).
- [71] Sallee, James, Sarah West, and Wei Fan (2009). "Consumer Valuation of Fuel Economy: A Microdata Approach." Working Paper, University of Chicago (September). Available from <http://www.ntanet.org/images/stories/pdf/proceedings/09/032.pdf>.
- [72] 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).
- [73] Sims, Christopher (2003). "Implications of Rational Inattention." *Journal of Monetary Economics*, Vol. 50, No. 3, pages 665-690.
- [74] Sims, Christopher (2010). "Rational Inattention and Monetary Economics." *Handbook of Monetary Economics*, Vol. 3, pages 155-181.
- [75] 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.
- [76] Small, Kenneth, and Harvey Rosen (1981). "Applied Welfare Economics with Discrete Choice Models." *Econometrica*, Vol. 49, No. 1 (January), pages 105-130.
- [77] Staiger, Douglas, and James Stock (1997). "Instrumental Variables Regression with Weak Instruments." *Econometrica*, Vol. 65, No. 3 (May), pages 557-586.
- [78] 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.

- [79] 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.
- [80] U.S. Bureau of Labor Statistics (BLS) (2007). "Consumer Expenditure Survey." <ftp://ftp.bls.gov/pub/special.requests/ce/standard/2007/quintile.txt>
- [81] U.S. Department of Transportation (2001). "National Household Travel Survey." <http://nhts.ornl.gov/download.shtml#2001>.
- [82] U.S. Energy Information Administration (2011). "Monthly Energy Review." <http://www.eia.doe.gov/totalenergy/data/monthly/>.
- [83] 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>.
- [84] Verboven, Frank (1999). "The Markets for Gasoline and Diesel Cars in Europe." Centre for Economic Policy Research Discussion Paper 2069 (February).
- [85] Wald, Abraham (1940). "The Fitting of Straight Lines if Both Variables are Subject to Error." *Annals of Mathematical Statistics*, Vol. 11, pages 284-300.
- [86] Woodford, Michael (2003). Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton, NJ: Princeton University Press.
- [87] Yahoo Autos (2009). "Top 100 Cars by Green Rating." http://autos.yahoo.com/green_center-top100/, accessed August 2009.

Tables

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Year	2003.5	2.9	1999	2008	1,068,459
Model Year	1996	5.2	1984	2007	1,068,459
Age	7.5	4.2	1	15	1,068,459
Price (\$)	7227	9845	0	358,289	1,068,459
Quantity	18,537	27,612	0	370,512	1,068,459
PDV of Gas Cost (\$)	7163	4147	686	41,576	1,068,459
Fuel Economy (MPG)	18.8	4	6.9	52.9	1,068,459
Horsepower	174	61	46	700	1,024,709
Weight (pounds)	3440	838	2.2	34,903	1,022,499
Wheelbase (inches)	110	13.9	79.9	972	1,025,023
Number of Transactions	37.8	122	1	9679	1,068,459

Notes: The number of observations differs from the base specification because this table includes observations between April and December 2008 and vehicles for which there is only one observation per fixed effect group. Unlike the regressions, this table is not weighted by number of transactions within an observation.

Table 2: Calculating the Average Discount Rate

Payment Method	Share of Vehicles	Discount Rate
Financed	37%	6.9%
Leased	0%	N/A
Cash	63%	5.8%
Weighted Average		6.2%

Notes: Share of Vehicles and the Finance Discount Rate 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, inclusive.

Table 3: Base Specification Results

	(1) Futures Forecast	(2) Martingale Forecast	(3) Both
G^f	-0.76 (0.046)		-0.73 (0.046)
G^m		-0.55 (0.032)	
$(G^m - G^f)$			-0.14 (0.037)
Time Dummies	Yes	Yes	Yes
Month-by-MPG Group Controls	Yes	Yes	Yes
Vehicle Fixed Effects	Yes	Yes	Yes
Observations	931,888	931,888	931,888

Notes: This table presents estimates of Equation (5) assuming the futures-based and martingale gas price forecasts, which are denoted G^f and G^m , respectively. All specifications use data from 1999 through March 2008 and use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Observations are weighted by number of transactions. Robust standard errors, clustered by vehicle, are in parenthesis.

Table 4: Alternative Assumptions

Row	Gas Price Forecast: Specification	Futures		Martingale	
		$\hat{\gamma}$	$SE(\hat{\gamma})$	$\hat{\gamma}$	$SE(\hat{\gamma})$
<i>Base</i>					
0	2 MPG Quantiles	0.76	0.046	0.55	0.032
<i>Grouping</i>					
1	5 MPG Quantiles	0.64	0.037	0.49	0.026
2	10 MPG Quantiles	0.64	0.034	0.49	0.024
3	20 MPG Quantiles	0.58	0.034	0.47	0.024
4	OLS (No grouping)	0.60	0.028	0.46	0.020
<i>Time Periods</i>					
11	2004-March 2008	0.80	0.037	0.61	0.027
12	1999-End 2008	0.85	0.035	0.57	0.025
<i>Discount Rate</i>					
21	r=0%	0.59	0.046	0.42	0.024
22	r=3%	0.67	0.035	0.48	0.028
23	r=10%	0.87	0.040	0.65	0.037
24	r=11%	0.90	0.053	0.67	0.039
25	r=15%	1.01	0.060	0.77	0.045
<i>Characteristics</i>					
31	Sample with characteristics	0.76	0.049	0.56	0.034
32	Control for HP	0.76	0.049	0.57	0.034
33	Control for HP, weight	0.76	0.049	0.57	0.035
34	Control for HP, weight, wheelbase	0.76	0.049	0.56	0.034
35	Control for all characteristics	0.75	0.048	0.55	0.034
<i>Preferences</i>					
41	Exclude hybrid vehicles	0.76	0.046	0.55	0.032
42	Exclude "green" vehicles	0.73	0.048	0.54	0.033
43	Exclude SUVs	0.71	0.065	0.52	0.046
44	Exclude all cars	0.87	0.087	0.68	0.061
45	Exclude minivans	0.70	0.046	0.51	0.032
<i>Age-by-Time Controls</i>					
51	One year per group	0.79	0.041	0.57	0.029
52	Two years per group	0.79	0.043	0.57	0.030
53	Four years per group	0.78	0.044	0.57	0.031

Notes: This table presents estimates of Equation (5) under alternative assumptions. Unless otherwise stated, all specifications include 1999 through March 2008. For these specifications, N=931,888, and there are 37,677 *ja* fixed effect groups. N=433,815 for row 11, N=1,010,667 for row 12, N=614,217 for rows 31-35, N=931,618 for row 41, N=886,300 for row 42, N=826,180 for row 43, N=319,361 for row 44, and N=874,497 for row 45. All specifications include a full set of month-by-year time dummies and MPG group by month-of-year controls. Observations are weighted by number of transactions. Unless otherwise stated, all specifications use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Robust standard errors, clustered by vehicle, are in parenthesis.

Table 5: Age Subgroups

Ages	(1) Transaction- Weighted $\hat{\gamma}$	(2) NHTS Survival Probabilities $\hat{\gamma}$	(3) JDPA Retail Prices $\hat{\gamma}$	(4) Equally- Weighted $\hat{\gamma}$	(5) Transactions per Observation
All	0.76 (0.046)	0.84 (0.050)	0.60 (0.041)	0.62 (0.027)	37.8
1-3	0.93 (0.074)	1.03 (0.082)	0.66 (0.057)	0.97 (0.072)	89.2
4-6	0.64 (0.054)	0.70 (0.060)	0.52 (0.055)	0.56 (0.049)	43.9
7-10	0.45 (0.032)	0.50 (0.036)	0.37 (0.040)	0.48 (0.030)	21.3
11-15	0.26 (0.017)	0.28 (0.019)	0.18 (0.032)	0.24 (0.015)	7.7

Notes: This table presents estimates of Equation (5) for different age groups using the futures-based gas price forecast. All specifications use data from 1999 through March 2008, include a full set of month-by-year time dummies and MPG group by month-of-year controls, and use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Observations are weighted by number of transactions, except for in column 4. Robust standard errors, clustered by vehicle, are in parenthesis.

Table 6: Difference Estimator

	(1) Futures Forecast	(2)	(3)	(4) Martingale Forecast	(5)	(6)
ΔG_t	-0.33 (0.036)	-0.30 (0.041)	-0.27 (0.044)	-0.16 (0.009)	-0.16 (0.011)	-0.20 (0.014)
ΔG_{t-1}		-0.43 (0.029)	-0.49 (0.033)		-0.14 (0.009)	-0.19 (0.012)
ΔG_{t-2}		-0.40 (0.030)	-0.34 (0.032)		-0.13 (0.012)	-0.13 (0.015)
ΔG_{t-3}			-0.26 (0.028)			-0.13 (0.013)
ΔG_{t-4}			-0.08 (0.037)			-0.07 (0.016)
ΔG_{t-5}			0.03 (0.035)			-0.09 (0.013)
ΔG_{t-6}			-0.02 (0.040)			-0.03 (0.014)
Number of Obs.	835,837	719,936	573,961	835,837	719,936	573,961

Notes: This table presents estimates of Equation (5) with the difference estimator instead of the fixed effects (demeaning) estimator. The dependent variable is $p_{jat} - p_{ja,t-1}$. The table reports the coefficient on $\Delta G_t = G_{jat} - G_{ja,t-1}$, which corresponds to $-\gamma$. All specifications use data from 1999 through March 2008, include a full set of month-by-year time dummies and MPG group by month-of-year controls, and use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Observations are weighted by the number of observed transactions in each pair of differenced months. Robust standard errors, clustered by vehicle, are in parenthesis.

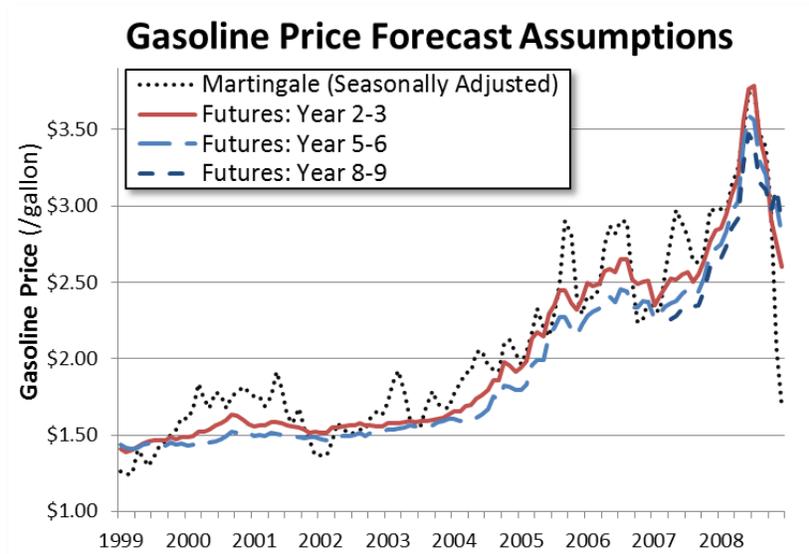
Table 7: Time Aggregation

	(1) Futures Forecast $\hat{\gamma}$	(2) Martingale Forecast $\hat{\gamma}$	Observations
Fixed Effects			
1-Month Units (Base)	0.76 (0.046)	0.55 (0.032)	931,888
2-Month Units	0.76 (0.045)	0.57 (0.032)	503,146
3-Month Units	0.75 (0.045)	0.57 (0.034)	342,551
4-Month Units	0.74 (0.046)	0.59 (0.034)	263,573
6-Month Units	0.72 (0.046)	0.58 (0.038)	182,300
Differences			
1-Month Units (Base)	0.33 (0.036)	0.16 (0.009)	835,837
2-Month Units	0.49 (0.035)	0.21 (0.010)	446,100
3-Month Units	0.65 (0.038)	0.24 (0.015)	294,862
4-Month Units	0.82 (0.038)	0.39 (0.017)	219,650
6-Month Units	0.86 (0.036)	0.57 (0.026)	141,295

Notes: This table presents estimates of Equation (5) with time units at different levels of aggregation. All specifications use data from 1999 through March 2008, include a full set of time dummies and MPG group by season controls, and use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Observations weighted by number of observed transactions. Robust standard errors, clustered by vehicle, are in parenthesis.

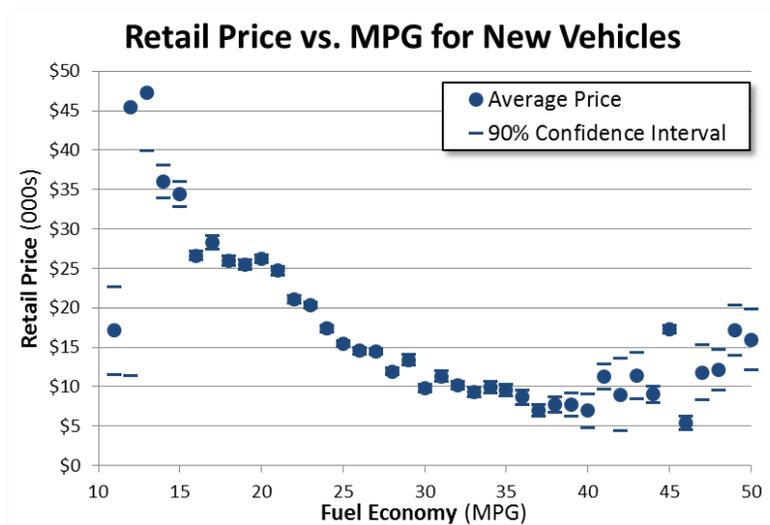
Figures

Figure 1: Retail Gasoline and Futures Prices



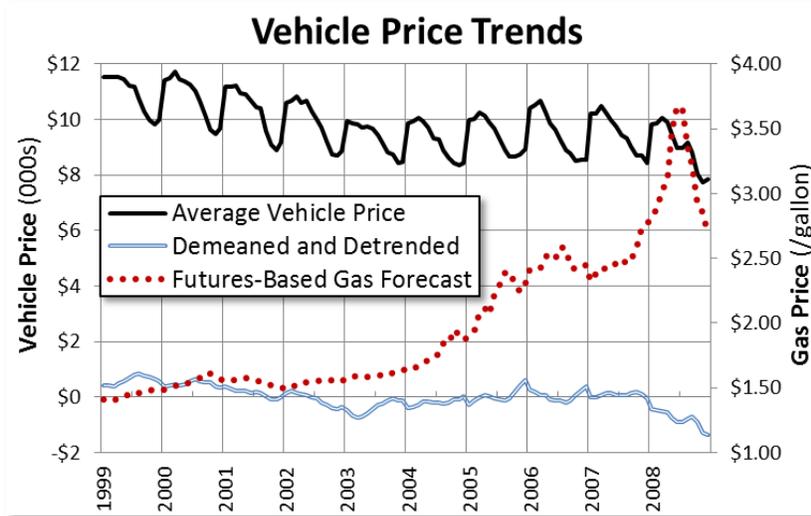
Notes: This figure shows seasonally adjusted retail gasoline prices and the gasoline prices implied by oil futures at different settlement horizons.

Figure 2: Manufacturer's Suggested Retail Price vs. Fuel Economy



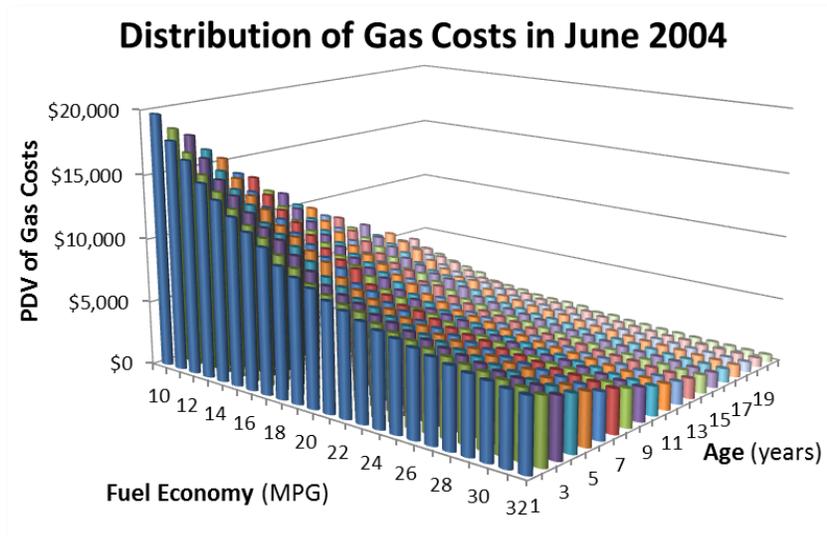
Notes: This figure shows the mean Manufacturer's Suggested Retail Price and the standard error of the estimated mean for each integer MPG rating. Data include all new vehicles from model year 1984 through 2008.

Figure 3: Vehicle Price Trends



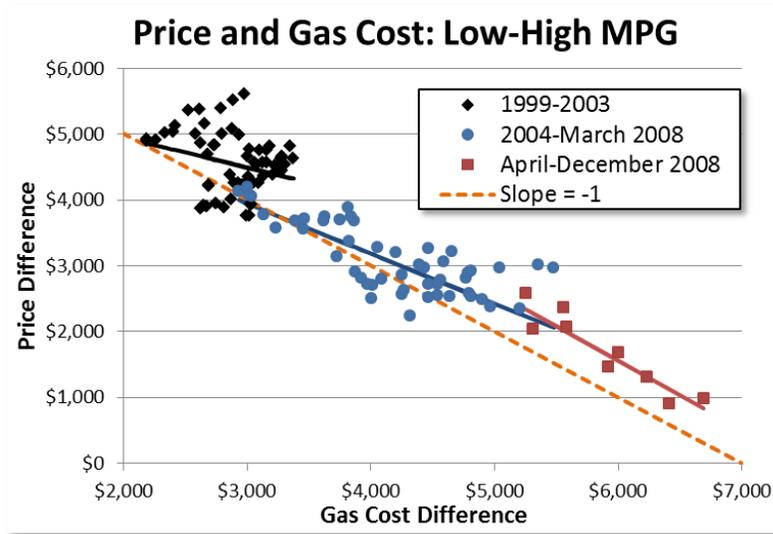
Notes: The black line shows the mean vehicle price across all transactions in our dataset. The double blue line is the average price conditional on vehicle fixed effects and month dummies. The Futures-Based Gas Forecast series is the average of oil futures prices for all future years, transformed into dollars per gallon of gasoline.

Figure 4: Identifying Variation in Gas Cost



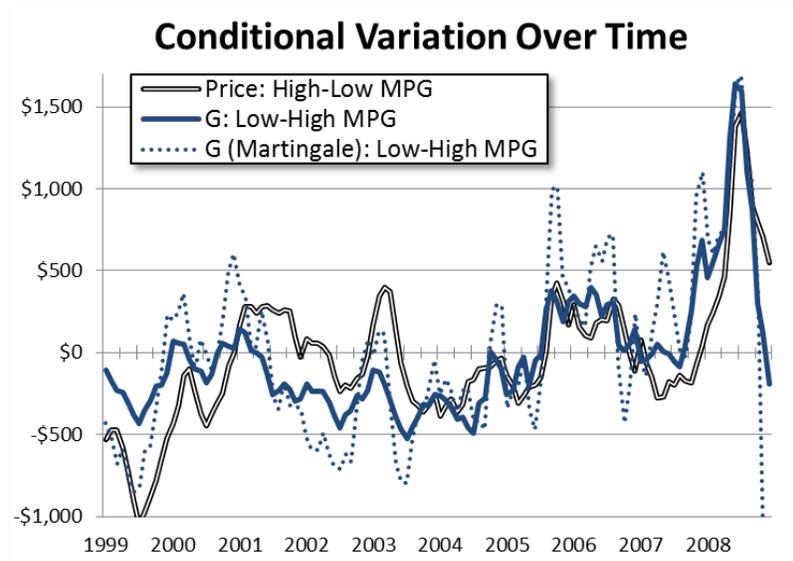
Notes: This figure shows the present discounted value of fuel costs under the futures-based gas price forecast, for vehicles of different MPG ratings and ages in June 2004.

Figure 5: Raw Data Scatterplot



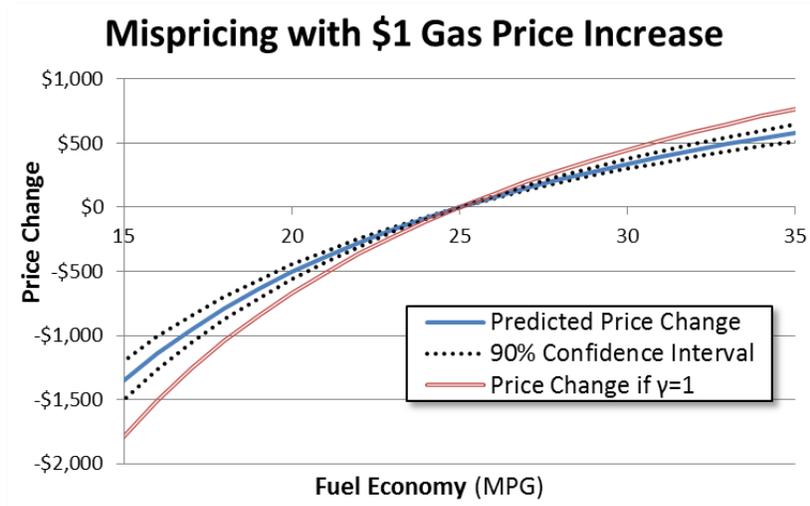
Notes: This figure plots the difference in transaction-level average prices for below-median minus above median-MPG vehicles against the difference in the PDV of gasoline costs, using the futures-based gas price forecast. There is one observation for each month.

Figure 6: Conditional Variation Over Time



Notes: The double black line is the difference in transaction-level average prices for above-median minus below-median MPG vehicles, conditional on vehicle fixed effects and MPG group by month-of-year dummies. The solid blue line is the conditional difference in the PDV of gasoline costs using the futures forecast. The dotted blue line is the conditional difference in the PDV of gasoline costs using the martingale forecast.

Figure 7: Predicted Mispricing



Notes: The double red line shows the price changes if $\gamma = 1$ in response to a \$1 increase in the gas price forecast for vehicles that would be driven 12,000 miles per year over a remaining lifetime of seven years, relative to a 25 MPG vehicle. The single blue line shows the predicted relative price changes for the base specification $\hat{\gamma} = 0.76$. The equations for the double red line and the blue line are, respectively, (Price Change if $\gamma = 1$) = $(\frac{1}{25} - \frac{1}{f}) \cdot \$1 \cdot 12,000 \cdot \sum_{y=1}^7 \delta^y$ and (Predicted Price Change) = $(1 - 0.76) \cdot (\frac{1}{25} - \frac{1}{f}) \cdot \$1 \cdot 12,000 \cdot \sum_{y=1}^7 \delta^y$.

Appendices: Not for Publication

Gasoline Prices, Fuel Economy, and the Energy Paradox

Hunt Allcott and Nathan Wozny

A Derivation from Dynamic Model

In Section 2, we derive the market share equation from a static discrete choice model. In this appendix, we show how that static model can in turn be derived from a simple dynamic model. This is also helpful in clarifying the nature of our econometric assumptions.

In each period t , each consumer chooses one vehicle from a set of new and used vehicles, including the vehicle that he or she already owns. Consumers expect to own the vehicle for a holding period of h years, after which time it will be resold at price $p_{ja,t+h}$. The budget constraint w_i should now be interpreted as the budget constraint over that holding period. Individuals discount utility and cash flows in future periods by discount factor $\delta = \frac{1}{1+r}$.

The PDV of gas costs, G_{jat} , can be divided into two parts: $G_{jat} = G_{jat}^o + \delta^h G_{ja,t+h}$. The variable G_{jat}^o 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 variable $\tilde{\psi}_{jat}^o$ captures the average utility across consumers from owning and using vehicle ja over the holding period.

The indirect utility function in this model is:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G_{jat}^o + \delta^h p_{ja,t+h}) + \tilde{\psi}_{jat}^o + \epsilon_{ijat} \quad (6)$$

The variable $\tilde{\psi}_{jat}$ now captures average usage utility over the holding period $\tilde{\psi}_{jat}^o$ plus the present discounted utility value of resale price plus remaining gasoline costs:

$$\tilde{\psi}_{jat} = \tilde{\psi}_{jat}^o + \eta \delta^h (p_{ja,t+h} + \gamma G_{ja,t+h}) \quad (7)$$

In our estimation, we assume that $\tilde{\psi}_{ja}$ 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 $\tilde{\psi}$ defined by Equation (7). The first element, the utility $\tilde{\psi}^o$ 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 $\tilde{\psi}$ 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, such as Rust (1985) and Stolyarov (2002). Under this assumption, 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 price forecasts change, the vehicle’s resale price will change by the difference in discounted remaining fuel costs $G_{ja,t+h}$ multiplied by the valuation weight γ . This demonstrates 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 price forecasts increase, this will affect both his own gasoline expenditures G^o and also the resale value.

B Endogenous Quantities

In deriving our base estimating equation in the body of the paper, we assume that the PDV of gasoline costs G is uncorrelated with market shares. In this appendix, we relax this assumption.

B.1 Empirical Strategy

Our estimating equation for the endogenous quantity specifications simply adds time dummies τ_t and model year dummies μ_{at} to Equation (3):

$$p_{jat} = -\gamma G_{jat} - \frac{1}{\eta} \ln s_{jat} + \tau_t + \mu_{at} + \psi_{ja} + \xi_{jat} \quad (8)$$

To control for seasonal depreciation patterns that might be correlated with MPG, we define fixed effects ψ_{ja} at the level of model by age (in months), instead of model by age (in years). These disaggregated controls are necessary instead of MPG group by month-of-year dummies ν because

we now identify off of all the variation in MPG. By contrast, the grouping estimator identifies γ using only variation in p and G between the two MPG groups.

Even conditional on our fixed effects, Equation (8) would suffer 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 for used vehicle market shares is the predicted lifetime gasoline costs of model j at the time when it was new, which we denote $G_{0,jat}$. 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 futures-based 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.

The 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. Appendix Figure A1 illustrates the identifying variation. 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.

This allows us to relax Assumption 1 from Section 4: we can now allow the market share of vehicle ja to be correlated both with gas costs G and unobservable characteristics ξ . In its place, we substitute the exclusion restriction: $E[G_0\xi|(\tau, \mu, \psi)] = 0$. We also note that for other empirical studies of used vehicle markets during periods when there is variation in gas price forecasts, our 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.

In theory, it might be possible to estimate Equation (8) using a grouping estimator, while instrumenting with group dummies (primarily to predict G) and with group-level averages of G_0 (primarily to predict $\ln s$). However, the grouping estimator involves a large number of instruments, of which only G_0 has effectively any correlation with $\ln s$. This causes a many instruments problem in fitting $\ln s$. Thus, in this appendix, we cannot use the grouping estimator, and we therefore must maintain the assumption that G is measured without error.

As Appendix Figure A1 suggests, most of the variation in the instrument affects model years 2004 and later. Thus, the 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 restrict our sample to transactions from 2004-March 2008.

B.2 Results

Column 1 of Appendix Table A2 presents estimates of Equation (5) for January 2004-March 2008 under the futures-based gasoline price forecast. Column 2 presents the "reduced form" estimates of Equation (8), without instrumenting for s . The coefficient on $\ln s$ is positive. This is the usual apparently upward-sloping demand symptomatic of simultaneity bias, driven by the correlation of equilibrium market share with the unobserved demand shifter ξ . Column 3 instruments for s with G_0 .

In the first stage of the IV estimates, 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 0.0669 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. The difference in G_0 changed by \$1966 between the two years, and as a result, the Civic market share is predicted to decrease

by 13 percent relative to the F-150. The first stage Angrist-Pischke F-statistic is 11.76, meaning that the estimates for this sample of years do not suffer from a weak instruments problem (Stock and Yogo 2005).

Now consider the estimated IV coefficients on $\ln s$. 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 \$361. Conversely, this also implies that if the price of a given vehicle were increased by \$1000, its market share would decrease by 28 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.

Because these specifications do not use the grouping estimator, the estimated $\hat{\gamma}$ is attenuated toward zero due to measurement error. Therefore, the key point of this table is not the absolute magnitudes of the $\hat{\gamma}$ coefficients. Instead, it is how the estimated coefficients compare across columns. Regardless of whether we instrument, using the logit model does not make any difference compared to the estimates of Equation (5) in column 1. This implies that our base specification assumption that $E[Gs|(\tau, \nu, \psi)] = 0$ does not generate statistically significant bias.

The basic reason for this is straightforward. Omitted variables bias requires both G and s as well as p and s to be correlated conditional on the other controls. Appendix Table A2 shows that p and s are indeed somewhat correlated: demand is downward sloping when we instrument properly. However, G and s are not highly correlated for our sample of used vehicles between ages 1 and 15, because once used vehicles have been sold, their stock is fixed until their value drops enough for some to be scrapped. In fact, when we estimate Equation (5) in our sample but use $\ln s$ as the left-hand-side variable instead of price, we cannot reject that $E[Gs|(\tau, \nu, \psi)] = 0$: the t-statistic is -1.597. This is why we restrict the sample to ages 1-15. If we were to include new vehicles or very old vehicles in our specifications, this assumption would not hold. For example, when we estimate Equation (5) with $\ln s$ as the left-hand-side variable for vehicles aged 15-25 years, we resoundingly reject that $E[Gs|(\tau, \nu, \psi)] = 0$: the t-statistic is -4.584.

In summary, it would ideally be possible to both allow endogenous market shares and address measurement error in the same specification, but we cannot do this because the grouping estimator requires a large number of instruments that only weakly predict $\ln(s)$. Appendix Table A2 shows

that endogenizing market shares does not significantly change the estimated γ , while we know from the body of the paper that measurement error causes significant attenuation bias. Therefore, we focus on the grouping estimator in the body of the paper and can realistically maintain the assumption that $E[Gs|(\tau, \nu, \psi)] = 0$.

C Endogenous Vehicle-Miles Traveled

C.1 Model

Our base specification assumes that vehicle-miles traveled (VMT) 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 the Envelope Theorem approximation might not hold. This section presents an alternative specification that endogenizes VMT.

As illustrated in Appendix Figure A2, the change in VMT resulting from a change in gasoline prices affects utility through two channels. First, changes in VMT affect gasoline expenditures G_{jat} . The NHTS survey that measures VMT was carried out in 2001; at 2001 gas price g_{2001} , the cost to drive vehicle ja one mile is $g_{2001}/f_{ja,2001}$, and the consumer chooses VMT $m_{ja,2001}(g_{2001})$. In future year y with higher gasoline price g_y , the consumer reduces VMT to $m_{jay}(g_y)$. The annual gasoline cost is now the shaded blue rectangle bounded by the g_y/f_{jay} and m_{jay} .

The second effect is that usage utility $\tilde{\psi}_{jat}$ 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 A2, the consumer's total willingness to pay for vehicle use is the area under the VMT demand curve. As the gasoline price increases from g_{2001} to g_y and the consumer's utility-maximizing VMT decreases, this total willingness to pay decreases by the solid green area.

Of course, for a marginal change in m_{jay} from $m_{ja,2001}$, the change in usage utility equals the change in gasoline expenditures, because consumers' observed choices must be such that marginal cost of driving equals the marginal willingness to pay. This is the Envelope Theorem logic that we used in the body of the paper to help justify the assumption that m_{jat} is constant over t .

To formalize these two effects of a non-marginal change, explicitly denote $G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t), \cdot)$ and

$\tilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))$ 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 at time t . The utility function in Equation (1) can be written as:

$$u_{ijat} = \eta(w_i - p_{jat} - \gamma G(\mathbf{m}_{jat}(\mathbf{g}_t), \cdot)) + \tilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_{2001})) + \int_{\mathbf{g}_{2001}}^{\mathbf{g}_t} \frac{\partial \tilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))}{\partial \mathbf{g}_t} d\mathbf{g}_t + \epsilon_{ijat} \quad (9)$$

The per-period dollar value of this new term with the integral is the solid green area in Appendix Figure A2. It can be calculated under an assumed functional form of VMT demand. Here, we derive this under the assumption of constant elasticity, denoted ζ . The demand function intercept κ_{ja} is pinned down by the vehicle's fitted VMT from the 2001 NHTS data and the gas prices from that year.¹⁵ The dollar value of the usage utility changes over the vehicle lifetime is denoted $I_{jat}(\mathbf{g}_t)$:

$$I_{jat}(\mathbf{g}_t) \equiv \frac{1}{\eta} \int_{\mathbf{g}_{2001}}^{\mathbf{g}_t} \frac{\partial \tilde{\psi}_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))}{\partial \mathbf{g}_t} d\mathbf{g}_t \quad (12a)$$

$$= \sum_{y=0}^L -\delta^y \cdot \left\{ \left(m_{jay}^{1+1/\zeta} - m_{ja,2001}^{1+1/\zeta} \right) \cdot \frac{\exp\left(\frac{-\kappa_{ja}}{\zeta}\right)}{1 + 1/\zeta} \right\} \cdot f_{jay}^{-1} \cdot \phi_{jay} \quad (12b)$$

C.2 Empirical Strategy

Using the same steps as in the derivation of the base specification, we carry the $I_{jat}(\mathbf{g}_t)$ term through from the utility function to an alternative estimating equation. Because I_{jat} and G_{jat} are

¹⁵This can also be stated mathematically. Under constant elasticity of demand ζ , we have VMT at any gasoline price:

$$\ln(m_{jas}) = \zeta \cdot \ln(g_s) + \kappa_{ja} \quad (10)$$

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:

$$\kappa_{ja} = \ln(m_{ja,2001}) - \zeta \cdot \ln(g_{2001}) \quad (11)$$

highly correlated, we move I_{jat} to the left hand side. The estimating equation is:

$$p_{jat} - I_{jat}(\mathbf{g}_t) = -\gamma G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t)) + \tau_t + \psi_{ja}(\mathbf{m}_{ja}(\mathbf{g}_{2001})) + \xi_{jat} \quad (13)$$

In this equation, $G_{jat}(\mathbf{m}_{jat}(\mathbf{g}_t))$ is calculated as before in Equation (4), except with VMT as a constant elasticity function of gas price. The term $\psi_{ja}(\mathbf{m}_{ja}(\mathbf{g}_{2001}))$ is constant, and it thus can represent a vehicle fixed effect, as in the base specification.

C.3 Results

Appendix Table A3 presents the results of the endogenous quantity specifications under the futures-based gasoline price forecast. Row 0 of Appendix Table A3 repeats column 1 from Table 3, our base specification estimates of Equation (5) using G^f , for 1999-March 2008. Rows 1-4 estimate Equation (13) for that same time period. We use both the linear and constant elasticity functional forms. 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, which is even larger than the empirical estimates of Bento *et al.* (2009), Kilian and Murphy (2011), and Davis and Kilian (2011b), $\hat{\gamma}$ is still very close to the base estimates.

Row 10-14 and 20-24 present the same set of estimates for the late period (2004-March 2008) and for the entire sample (1999-End 2008). As discussed in the body of the paper, the estimated γ from Equation (5) is sensitive to the time period. However, endogenizing vehicle-miles traveled does not change the estimates in any time period. For this reason, the body of the paper focuses on the simpler specification that assumes that VMT is exogenous.

Appendix Tables

Appendix Table A1: Gas Prices and Futures-Based Gas Price Forecasts

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.

Appendix Table A2: Results with Endogenous Quantities

Model:	(1)	(2)	(3)
	Base	Logit Non-IV	Logit IV
-G	0.44 (0.033)	0.44 (0.033)	0.44 (0.048)
$\ln(s)$		61 29	-3,607 (1,380)
First Stage			
$G_0/10^6$			-66.9 19.5)
$G/10^6$			32.3 11.5)
Angrist-Pischke F Stat			11.76

Notes: The first column of this table presents estimates of Equation (5), while the second two columns present estimates of Equation (8). All specifications include 2004 through March 2008. $N = 314,519$, and there are 120,917 fixed effect groups. Observations are weighted by number of observed transactions. All specifications include a full set of month-by-year time dummies and model year controls. Standard errors are in parenthesis; they are robust, clustered at the level of model by age, and use bootstrapping to account for the variance in first step estimates of G .

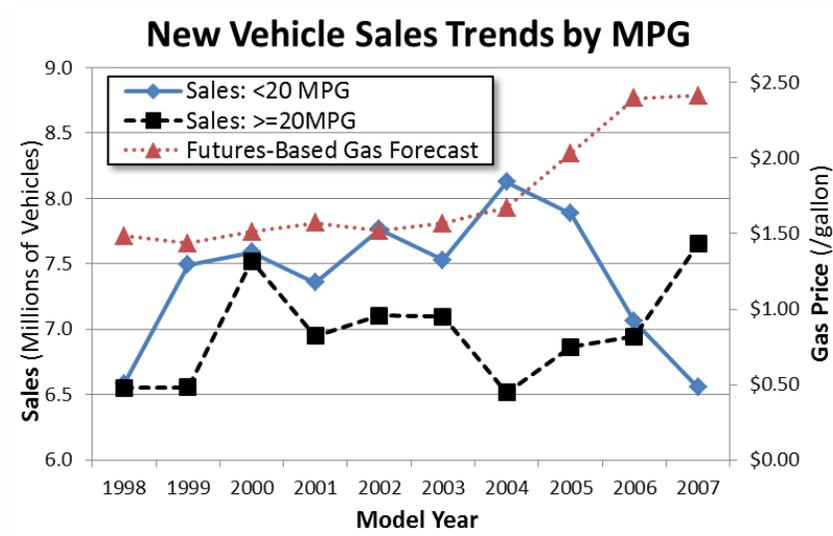
Appendix Table A3: Results with Endogenous Utilization

Row	Specification	$\hat{\gamma}$	$SE(\hat{\gamma})$
<i>Time Period: 1999-March 2008</i>			
0	Base	0.76	0.052
1	Constant elasticity = 0.2	0.76	0.063
2	Linear elasticity = 0.2	0.76	0.060
3	Constant elasticity = 0.5	0.77	0.097
4	Linear elasticity = 0.5	0.76	0.072
<i>Time Period: 2004-March 2008</i>			
10	Base	0.80	0.044
11	Constant elasticity = 0.2	0.82	0.053
12	Linear elasticity = 0.2	0.81	0.052
13	Constant elasticity = 0.5	0.88	0.083
14	Linear elasticity = 0.5	0.82	0.063
<i>Time Period: 1999-End 2008</i>			
20	Base	0.85	0.042
21	Constant elasticity = 0.2	0.90	0.053
22	Linear elasticity = 0.2	0.88	0.052
23	Constant elasticity = 0.5	1.03	0.086
24	Linear elasticity = 0.5	0.90	0.067

Notes: Row 0 presents estimates of Equation (5), while all other rows present estimates of Equation (13). For rows 0-4, $N = 931,888$, and there are 37,677 fixed effect groups. For rows 10-14, $N = 433,815$, and there are 21,839 fixed effect groups. For rows 20-24, $N = 1,010,667$, and there are 37,922 fixed effect groups. All specifications include a full set of month-by-year time dummies and MPG group by month-of-year controls, and all specifications use the grouping estimator with groups at the level of time by MPG group, with two MPG groups. Observations are weighted by number of transactions. Robust standard errors, clustered by vehicle, are in parenthesis.

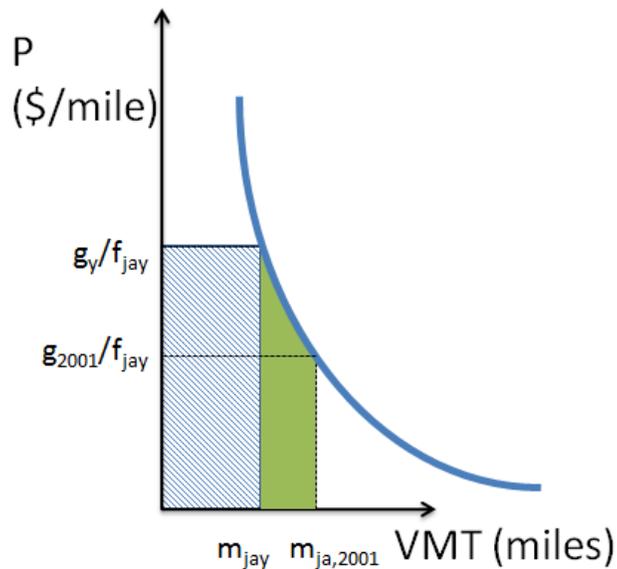
Appendix Figures

Appendix Figure A1: New Vehicle Sales by MPG Rating



Notes: This figure shows the total sales by model year for vehicles rated less than versus greater than 20 MPG. The Futures-Based Gas Forecast series is the average of oil futures prices for all future years, transformed into dollars per gallon of gasoline.

Appendix Figure A2: VMT Demand



Notes: This figure illustrates how utility and gasoline expenditures change when gas prices rise from g_{2001} to g_y .