Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations

James M. Sallee
University of Chicago
and NBER

Sarah E. West
Macalester College

Wei Fan*
RVI Group

January 28, 2015

Abstract

Debate about the appropriate design of energy policy hinges critically on whether consumers might undervalue energy efficiency, due to myopia or some other manifestation of limited rationality. We contribute to this debate by measuring consumers' willingness to pay for fuel economy using a novel identification strategy and high quality microdata from wholesale used car auctions. We leverage differences in future fuel costs across otherwise identical vehicles that have different current mileage, and therefore different remaining lifetimes. By seeing how price differences across high and low mileage vehicles of different fuel economies change in response to shocks to the price of gasoline, we estimate the relationship between vehicle prices and future fuel costs. Our data suggest that used automobile prices move one for one with changes in present discounted future fuel costs, which implies that consumers fully value fuel economy.

---

1 Introduction

One of the great questions facing policy makers in the twenty-first century is whether and how to mitigate greenhouse gas emissions so as to limit climate change. Automobiles are a critical part of this policy problem—in the U.S., personal transportation accounts for 28% of greenhouse gas emissions (Environmental Protection Agency 2014). Gasoline consumption maps neatly into greenhouse gas emissions. This means that a Pigouvian tax on emissions is feasible (in the form of a gasoline tax). Such a tax can fully restore market efficiency, and alternative policies, such as fuel economy standards, will have inferior welfare properties, provided that the environmental externality is the only market failure leading to inefficiencies.[1]

However, many have argued that another market failure does exist, which is that consumers undervalue energy efficiency in a variety of choice situations, including automobile markets. The root of this hypothesis is the observation that engineering estimates of the cost of deploying fuel saving technologies often suggest that privately cost-effective technologies often go unadopted. Jaffe and Stavins (1994) call this the "energy paradox". If markets substantially undervalue energy efficiency, then the dominance of a gasoline tax over regulatory approaches may be broken because alternative policies may be better able to correct for inefficiencies from mis-valuation.[2]

Motivated by these policy implications, researchers have sought to use revealed preference data to determine whether consumers do in fact undervalue fuel economy. In this paper, we add to this literature by developing a unique identification strategy that utilizes two decades worth of microdata on used vehicle transactions to test whether used vehicle prices change by the amount predicted by a fully rational asset pricing model. We interpret our results as a test of whether or not consumers fully value fuel economy, and our results directly provide the parameters necessary for informed policymaking.

Intuitively, our approach is to first compare the prices of two used cars which are identical except in their current odometer readings—and therefore in the remaining future operating costs—and second to repeat this comparison when different gasoline prices prevail. We repeat this comparison

---

[1] For reviews of the design of policies to correct driving related externalities, see Parry, Walls, and Harrington (2007); Anderson, Parry, Sallee, and Fischer (2011) and Sallee (2011). The efficiency property of a gasoline tax is more complicated for local pollutants, as is explored in several papers, including Fullerton and West (2002, 2010) and Knittel and Sandler (2012).

[2] Fischer, Harrington, and Parry (2007), Allcott, Mullainathan, and Taubinsky (2014) and Heutel (2011) explore the implications of undervaluation for optimal policy design.

across many vehicle types and many months, during which changes in the price of gasoline drive changes in fuel costs, in order to estimate the relationship between vehicle prices and a measure of the present discounted fuel cost that we construct. For example, we calculate the price and fuel cost of a 2000 Ford Taurus SE six cylinder 3.0L vehicle with automatic transmission and front-wheel drive that has 50,000 miles in July 2005 to a different 2000 Ford Taurus SE six cylinder 3.0L vehicle with automatic transmission and front-wheel drive that has 60,000 miles in July 2005. We then calculate the price and fuel cost of two different cars with the exact same configuration and mileage in July 2006. Changes in the gasoline price between July 2005 and July 2006 will cause changes in the difference in expected fuel costs across the higher and lower mileage vehicles. We test whether the change in the price difference between the high and low mileage vehicle over time corresponds to the change in the cost difference.

This is conceptually similar to a difference-in-difference approach. The fact that our comparison is within vehicle type, and not across vehicles, allows us to provide an exceptionally rich set of controls, including time-period shocks and depreciation schedules that are unique for each vehicle type. Specifically, our preferred specification allows for a unique fixed effect in every month of our sample for every vehicle type, and it controls for a unique depreciation schedule for each vehicle type, where a vehicle type is very finely defined. To execute this research design, we employ used vehicle price data that include actual transaction prices, dates of sale, vehicle identification numbers (VINs), and odometer readings for a large sample of vehicles sold at wholesale auctions.

In our baseline specification, we find that vehicle prices do move one for one with future fuel costs. This conclusion is robust to a number of specification checks.[3] Given some simplifying assumptions about the structure of the used car market, this result implies that consumers do value fuel economy properly. This finding casts doubt on the idea that regulatory policies, such as fuel economy standards, might be more efficient than fuel taxation because they correct both the environmental externality and private mis-optimization due to limited rationality.

Our data come from wholesale auctions, but our interest is in what consumers pay in the retail market. Using an auxiliary data set from used car guide books, we demonstrate that price changes

---

[3]The main exception is that we find that for our highest mileage cars (those with over 100,000 miles when sold at auction), prices are significantly less responsive to fuel cost shocks. This may indicate that buyers of the oldest and least expensive used cars undervalue fuel economy, but it may also be due to a selection process by which only certain types of high mileage vehicles appear in wholesale auctions.

in the retail market appear to pass through one to one into retail prices. This is consistent with a competitive used car market, and it allows us to interpret our wholesale price results as directly reflecting consumer willingness to pay in the retail used car market.

We are not the first to ask whether or not consumers value fuel economy properly.[4] The most similar existing papers are Allcott and Wozny (2014), Busse, Knittel, and Zettelmeyer (2013) and Grigolon, Reynaert, and Verboven (2014). These papers use a panel identification strategy that leverages the fact that common gasoline price shocks translate into different fuel cost shocks for different vehicles based on their fuel economies.[5] Compared to these papers, we are able to relax a number of restrictive assumptions on the set of control variables. Specifically, because we leverage differences across vehicles of the same type in the same month by using variation in the odometer, we can control nonparametrically for time period shocks specific to each vehicle type, and we can control very flexibly for a depreciation schedule for each vehicle type. These papers find a range of estimates of consumer valuation across specifications that overlap with each other, with Allcott and Wozny (2014) emphasizing estimates that find modest undervaluation while Busse, Knittel, and Zettelmeyer (2013) and Grigolon, Reynaert, and Verboven (2014) emphasize that their results cannot consistently reject full valuation.

We interpret our estimates as consistent with their results. Moreover, we believe that our procedure presents a more difficult test because we identify consumer valuation off of variation in odometers within a set of identical vehicles, which may not be very salient to consumers, whereas other approaches identify valuation by looking across vehicles. If consumers have limited attention, in the sense of Sallee (2014), then we might expect them to ignore the type of within model variation in fuel costs that we leverage. That is, one could imagine consumers recognizing the fuel cost differences across categories of automobiles, but not "noticing" the difference in implied fuel costs across high and low mileage versions of the same model.

Our baseline model produces precise estimates consistent with full valuation. Our procedure yields statistical precision, and our results are robust across a number of dimensions. But, we emphasize that our procedure can be made to yield different results because they hinge on a number of assumptions about underlying parameters that are used to construct our estimate of the future

---

[4]We discuss the literature, and our relationship to it, more fully in the next section.

[5]Linn and Klier (2010) use the same strategy to study sales volumes, rather than prices. Li, Timmins, and von Haefen (2009) and Jacobsen and van Benthem (Forthcoming) use it to study vehicle scrappage decisions.

fuel cost of a vehicle, including consumer discount rates, expectations regarding future gasoline prices, perceived on road fuel economy, and typical patterns of vehicle utilization and scrappage. We have empirical support for each of the assumptions we use, but reasonable alternative parameter choices could shift our coefficient estimate in either direction. The same is true of other papers in the literature.

Thus, while the literature fails to consistently reject the null hypothesis of full valuation, the data cannot consistently rule out modest undervaluation, unless one takes a firm stand on underlying parameters that are themselves uncertain. What is clear from our results, in conjunction with the existing literature, is that the assumption that consumers place a *zero* value on fuel economy is indefensible. Nevertheless, this assumption is still employed, implicitly, in regulatory impact analyses that credit the entire fuel savings of consumers as a benefit of programs like the Corporate Average Fuel Economy (CAFE) standards. If consumers value fuel economy properly and automakers deploy fuel-saving technologies so long as their cost lies below consumer willingness to pay, then fuel economy improvements forced upon the market by regulation must be causing a trade-off in vehicle characteristics or market shares that lowers consumer surplus. This crediting of fuel savings as a program benefit is often pivotal to the cost-benefit analysis. For example, fuel savings represent nearly 80% of the total benefits of the 2017-2025 CAFE standards (Environmental Protection Agency 2012).

Our empirical evidence, combined with the previous literature, implies that consumers at worst undervalue fuel economy modestly. If there are energy efficient technologies that are not being deployed, then researchers and regulators should perhaps shift their attention to supply side explanations, like competitive failures, technological spill overs or other hold ups within the automobile industry.[6]

2   Conceptual approach

Our estimation strategy is based on exploiting variation in the expected future fuel cost of vehicles resulting from differences in the odometer reading and the price of fuel at the time a vehicle is sold.

---

[6]Sallee (2014) suggests one possible caveat. He argues that in a model with rational inattention it is possible that consumers are attentive to (fully value) the set of energy efficiency technologies that are deployed in equilibrium, but nevertheless there are cost-effective technologies that are not deployed because they are not salient and would be undervalued if they were deployed.

For example, suppose that two 2005 Toyota Camrys are sold in November 2009, one with 80,000 miles on the odometer and the other with 90,000 miles. The price difference between these two cars should reflect the difference in value of having a lower mileage vehicle (which is in better condition and has a longer expected remaining life) net of the larger operating costs, which are a function of fuel prices. Next, imagine that the price of gasoline changes between November 2009 and December 2009, and that in December two other 2005 Toyota Camrys, one with 80,000 miles and the other with 90,000 miles, are sold. Now, the price difference should reflect the same factors as before, and the difference-in-difference should reflect only the change in operating costs that resulted from the gasoline price change. There may be many factors that affect the level of these prices, but as long as these factors have the same effect on an 80,000 mile Camry and a 90,000 mile Camry, they can be accounted for with fixed effects.

Our final estimating equation has the intuitive flavor of this example, but it is not literally a difference-in-difference because we use a continuous measure of odometer readings. In the end, we regress used vehicle transaction prices on a rich set of mileage controls, time period fixed effects and a measure of the discounted expected future operating costs of vehicles. Our specification allows different models to have different depreciation schedules. Specifically, we allow the odometer polynomial and the time period fixed effects to vary for each vehicle type—where a type is defined by all observed characteristics contained in the "stub" of the Vehicle Identification Number (VIN), which includes model name, vintage, cylinders, displacement, and, sometimes additional information on transmission and trim levels. We are able to do this and still identify the parameter because we use individual transaction prices and use variation in fuel costs within a car type and time period by exploiting the odometer information.

To arrive at our estimation equation, we start with a simple model of the price of used cars which assumes that the used car market is competitive, the supply of used cars is inelastic, and that transaction costs are small.[7] The logic of the inelastic supply assumption is that, in any given period, the full stock of used cars (not necessarily the set listed for sale) of a given vintage and type is fully predetermined. Under these assumptions, the expected discounted price, P, of an individual vehicle i of type j at time t is equal to the expected discounted value of operating the

---

[7]With zero transaction costs, the market allocation will be the same as if people rented a vehicle each period. This abstracts from sorting concerns.

car $V(\cdot)$ over its remaining lifetime, minus the expected discounted values of fuel costs $C(\cdot)$, and maintenance costs $Z(\cdot)$:

$$P_{ijt} = V(O_{ijt}, X_j, r) - C(O_{ijt}, m_{ijt}, g_t, \text{MPG}_j, r) - Z(O_{ijt}, X_j, r), \tag{1}$$

where O is the car's mileage ("O" is for odometer) at the start of period j, X_j is a vector of vehicle attributes, m is per period miles driven, g is the price of gasoline, MPG is the vehicle's fuel economy in miles per gallon, and r is the discount rate.

The discounted value of fuel costs, $C(\cdot)$, is given by:

$$C_{ijt} = \mathbb{E}\left[\sum_{s=t}^{T} H(O_{ijt}, X_j)\left(\frac{1}{1+r}\right)^{(s-t)}\frac{m_{js}g_s}{\text{MPG}_j}\right], \tag{2}$$

where T is the final period at which time all vehicles are scrapped, and $H(\cdot)$ is the probability of survival of a vehicle as a function of its odometer reading and its attributes. We detail the construction of this variable in Section 3.1, below.

The fundamental question in the literature on the energy efficiency gap is whether or not consumers fully value fuel economy, which is typically formulated as the existence of a one to one mapping between price and future discounted fuel costs, in the used car market where supply is assumed to be perfectly inelastic. Thus, the goal is to find a way to regress prices P on future discounted fuel costs C that leaves variation in C but protects against omitted variable biases. Variation in fuel costs comes from several sources. It comes across vehicles because of differences in fuel economy. It comes over time because of differences in the price of gasoline. And, it comes across individual vehicles depending on their remaining lifetime.

Different articles in the literature have used different types of variation, and all approaches require some assumptions to achieve identification. The cross-sectional hedonic approach that was once prominent in the literature used cross vehicle variation in studies of new car prices.[8] Automobiles have many unobserved characteristics, however, so the cross-sectional approach lacks credibility as an identification strategy.

---

[8]This approach goes back at least to the seminal contributions in Hausman (1979) and Dubin and McFadden (1984) on household durables using cross-sectional data. The cross-sectional literature on automobiles includes Dreyfus and Viscusi (1995), Goldberg (1998) and Espey and Nair (2005). This literature found mixed evidence of consumer undervaluation. See Greene (2010) and Helfand and Wolverton (2009) for reviews.

The literature recently switched to a focus on the use of panel strategies that use changes in gasoline prices as a source of quasi-experimental variation (Kahn 1986; Kilian and Sims 2006; Allcott and Wozny 2014; Busse, Knittel, and Zettelmeyer 2013; Grigolon, Reynaert, and Verboven 2014). A common gasoline price shock creates different fuel cost changes across models, because of differences in fuel economy.[9] Thus, a single gasoline price time series generates panel variation that allows for price regressions that include fixed effects for each vehicle type, which will absorb any time invariant unobserved factors. When the unit of analysis is the same vintage of vehicle observed several times in the used car market, these fixed effects will capture all attributes of the vehicle.

This approach uses variation from the interaction of fuel prices and vehicle fuel economy. This does not require the use of micro data (though it can be used, as in Busse, Knittel, and Zettelmeyer (2013)) because it uses differences across car types over time, rather than making predictions about how different vehicles of the same type would differ in price within a time period. As such, this approach cannot control for vehicle type specific time period shocks; it rests on assumptions that some broader level of time period fixed effects are able to account for demand shocks that influence the price of models. If there are demand shocks in different time periods that have differential effects on different vehicles prices, and if these differential shocks are correlated with fuel economy, then this approach may still be biased. As suggested in Langer and Miller (2013), the prices of competing products may be one such source of correlated shocks.

Our approach focuses on vehicle mileage in order to relax those assumptions. We isolate the effect of fuel price changes over time on the differences in prices across vehicles of the same type in the same time period, which differ only in their remaining lifetime by virtue of having been driven more or less prior to that time period.[10] This approach inherently requires micro data.

Our approach may involve trade-offs, both practical and conceptual. Practically, our procedure eliminates the vast majority of variation in prices and fuel costs through our control variables,

---

[9]Grigolon, Reynaert, and Verboven (2014) differ from our approach and from the existing literature on the U.S. in focusing on the demand for diesel versus gasoline vehicles in Europe.

[10]Our conceptual framework assumes a risk neutral agent who cares only about the average future fuel cost. Risk averse agents will perceive some value in more fuel economic vehicles, which condense the variation in fuel costs that arise from shocks to the future gasoline price. Changes in the volatility of the price of gasoline may therefore influence the value of fuel economy over our sample period. But, our identification strategy limits this concern. Any increase in demand for a particular car type due to volatility considerations will be soaked up in our vehicle by time fixed effects.

7

and the variation that remains may be too small to yield statistical precision. That turns out not to be the case, both because the gasoline price moves a great deal during our sample period and because we have a large data set. Aside from practical concerns, isolating variation in fuel costs to within model differences due to odometer variation leads us to focus on a type of variation that may be particularly non-salient to consumers. Consumers might ignore differences in fuel costs due to odometer variation, while still accurately considering the difference in fuel costs across models.[11] That is, when gasoline prices rise, consumers may switch demand away from heavy trucks towards compact cars, but they might simultaneously fail to account for the differential implication for low versus high mileage compact cars. For this reason, we think that ours is a more stringent test of consumer valuation than the cross-vehicle approaches in the prevailing literature.

3 Data

Our used car price data come from a large sample of wholesale used car price auctions. The data include the transaction price, transaction date, odometer reading and truncated Vehicle Identification Number (VIN) of each vehicle sold in several large auction houses.[12] This market does not include individual end users. Automobile dealers, manufacturers and businesses and governments that own large fleets sell their vehicles at these auctions. The buyers are licensed used car dealers, who subsequently resell the vehicles to consumers. Used car dealers routinely use these auctions to optimize the stock of vehicles they have for sale to final consumers by both buying and selling vehicles.

Our data sample includes millions of transactions that took place between July 1993 and July 2008.[13] We match these vehicles to official EPA fuel economy ratings using all available information on model, model year, cylinders, displacement, body type, transmission and trim. In the estimation, we use the combined EPA fuel economy rating. For some early model years and for model years

---

[11]Allcott and Wozny (2014) make this same point in discussing our approach versus theirs. Sallee (2014) explores the same issue, arguing that, when perfect information is costly to acquire, consumers may be rationally inattentive to some types of fuel cost variation, while being attentive to others, due to differences in the consequences of ignoring each type of variation.

[12]One source of error may exist in the odometer readings for a small fraction of our data. Goh, Fischbeck, and Gerard (2007) show that many domestic vehicles built prior to the mid-1990s had five-digit odometers. This leads to rollover, where it is impossible to tell how many hundreds of thousands of miles a vehicle has been driven. Our estimates are similar when we look only at vehicles known to have six-digit odometers, such as Toyotas and Hondas.

[13]The raw data include some data before 1993, but the coverage is limited. We have access to data through part of 2009, but we limit our sample to the period before the financial crisis.

after 2007, we lack a complete VIN decoder and for vehicles made before 1978, there are no fuel economy ratings. Such vehicles are dropped from our sample. In the results reported here, we focus on a 10% random sample of our data for computational reasons.

We also focus primarily on vehicles sold at auction by dealers, rather than vehicles sold by auto manufacturers themselves or by fleet operators. We do so because our strategy relies on correctly specifying the remaining lifetime of vehicles, and we believe that vehicles sold by dealers, which were owned and operated by consumers, will be more likely to depreciate according to the average schedule that we use for estimation. That is, vehicles operated by fleet owners and manufacturers may have been used and maintained differently than those owned by consumers. We do report results from our preferred specifications on these alternative samples.

3.1 Estimating Remaining Cost

A critical step in our estimation is the construction of the future fuel cost variable. This construction requires assumptions, described here, about vehicle mileage and survival rates and about consumers' beliefs about future gasoline prices and their discount rates. Throughout the paper we emphasize our econometric analyses as tests of whether or not consumers fully value fuel economy, but it is critical to keep in mind that our tests (and those in the existing literature) depend on the accuracy of these assumptions, and that modifications of these assumptions will mechanically alter our estimated coefficients.

Mileage and Scrappage

The cost function (equation 2), specified in discrete form, includes the annual miles driven of vehicles in all future periods. Additionally, since a vehicle might be scrapped, mileage is multiplied by a survival probability to generate expected miles driven.

We estimate mileage and scrappage by using the results reported in Lu (2006), which details the average mileage and scrappage probability for passenger cars and light-trucks (pick-ups, sport-utility vehicles, and vans) as a function of age. For our purposes, we believe the relevant measure of the lifecycle position of a vehicle is its odometer, not its age in months. Accordingly, we conduct a transformation of the results in Lu (2006) to respecify her results in terms of odometer, rather than age. Lu (2006) also does not provide separate results by make (e.g., Dodge), so we use our

9

auction data to introduce heterogeneity across makes in an intuitive way.

Our procedure is as follows. We generate a predicted odometer reading for all vehicles in our sample by class (car and light-duty truck) based on age, measured in annual increments, using the regression coefficients reported by Lu (2006). We then regress this predicted odometer reading on the actual odometer reading separately for each make as follows:

$$d_{ca} = \theta_{cm} O_{icm} + \varepsilon_{icm} \tag{3}$$

where c indexes class, a indexes age, and m indexes make. Note that the mileage estimates from Lu (2006) are all mileage conditional on survival. This procedure estimates a unique θ_{cm} for each make. We use these θs as a measure of how much each vehicle make is driven compared to the average across all makes. While coarse, this procedure enables us to introduce some heterogeneity in miles driven across makes.

Next, we transform the results of Lu (2006) to write mileage and scrappage as a function of odometer, rather than age. Lu (2006) reports mileage and scrappage as functions of a quartic in age. We invert these functions numerically by using the reported regression coefficients as a data generating process on a simulated data sample and then regressing the scrappage and mileage functions on a polynomial in odometer. We then use these functions to create a predicted scrappage function and mileage function as a function of odometer. The base mileage function is then multiplied by the estimated θ_{cm} for each make to allow for heterogeneity, but we currently impose the same scrappage function on all vehicles. Specification testing suggests that the introduction of this heterogeneity has little effect on our final estimates, but we prefer to include it in our baseline so as to allow some degree of data driven heterogeneity.

In constructing the cost measure, we assume that each vehicle lasts no more than 25 remaining periods from the time of observation. Since we are using a scrappage probability, expected mileage beyond 25 years in the future is negligible for all vehicles. In all results reported in the paper, we assume that the vehicle miles traveled and scrappage schedules are independent of the price of gasoline. There is evidence that the gasoline price affects miles traveled and scrappage, but these affects are modest. For small changes in the price of gasoline, an envelope theory argument suggests that any such responses will be second order. In addition, we experimented extensively

with introducing a mileage elasticity and make-specific scrappage elasticities based on results from Li, Timmins, and von Haefen (2009) and found that our results were insensitive to these additional considerations.[14]

Gasoline Price Expectations

The cost function in equation 2 also depends upon the price of gasoline in *future* periods. We assume that all consumers use a "no change" forecast—i.e., they expect that the future price of gasoline in all periods is equal to the current price. This is consistent with evidence on actual consumer beliefs reported in Anderson, Kellogg, and Sallee (2013). It is also the case that a no change forecast for oil prices preforms as well, or better than, alternative forecasts based on futures markets or expert surveys (Alquist and Kilian 2010; Alquist, Kilian, and Vigfusson 2013).

Fuel Economy

Several different fuel economy measures are available. In particular, the EPA reports both a city and a highway, along with a combined rating which is a harmonic average of the two. We use the combined rating. In 1986, the EPA adjusted ratings to account for the apparent discrepancy between the official rating and the actual fuel economy used. Starting in 2008 (which is the very end of our sample), a significant revision to the test procedures was initiated with the hopes of improving accuracy further. We use the fuel economy rating published at the time of the vehicle's sale, which is the official number still available in the EPA fuel economy guide, without attempting to adjust for these regime changes. Additionally, we assume that the fuel economy rating is accurate for the life of the vehicle. Some research exists quantifying the degree to which fuel economy may degrade as a vehicle ages. In principle, it would be possible to adjust for this in cost measure, but we suspect that consumers are largely unaware of this phenomenon and know only a vehicle's official fuel economy rating.

---

[14]Li, Timmins, and von Haefen (2009) and Jacobsen and van Benthem (Forthcoming) document how scrappage decisions respond to gasoline prices, and there is a large literature studying the mileage response, including Hughes, Knittel, and Sperling (2008); Gillingham (2011); Knittel and Sandler (2012). Kahn (1986) makes the envelope theory argument to justify a constant mileage assumption. Kilian and Sims (2006) make the same assumption based on the same reasoning. Allcott and Wozny (2014) and Busse, Knittel, and Zettelmeyer (2013) assume no mileage response in their baseline specifications, and, like us, they find that relaxing the assumption has little impact on their results.

The Discount Rate

The cost function also includes a discount rate. For our baseline estimates, we use a discount rate of 5%, which we intend to be a conservative (low) benchmark. Allcott and Wozny (2014) use 6%, based on the average interest rate on car loans according to the Survey of Consumer Finance. Alternatives rates could be justified by pointing to borrowing costs or other metrics of the opportunity cost of funds. We do not believe that it is possible to identify a single correct discount rate, so we select a low rate and discuss the sensitivity of our results to higher discount rates.

3.2 Summary Statistics

Before presenting our regression results, we briefly describe the summary statistics for our baseline sample in Table 1. Our baseline sample includes vehicles sold at auction by automobile dealers that have mileage between 10,000 and 100,000 miles. (Our reason for making this restriction is explained in the next section.) The average mileage in this sample is just over 58,000 miles. These vehicles are about four and one-half years old on average. Sixty-two percent are cars (as opposed to light-duty trucks).

Average transaction prices are about $10,300, as compared to remaining future fuel costs of $7,800, in present discounted value. (All dollar amounts are inflation adjusted using the CPI-U to 2008 dollars.) Gasoline prices are tax inclusive monthly retail prices from the Energy Information Administration. Our data range from the July 1993 to June 2008. During that period, there was substantial variation in the price of gasoline; the trough was less than one-third of the peak price during our sample. We also have substantial variation in fuel economy; the average fuel economy in the sample is 21.9 miles per gallon, with a standard deviation of 4.5.

Our estimation sample includes over 9,000 distinct types of vehicles, defined by the "VIN stub"— the portion of the Vehicle Identification Number that identifies the vehicle's attributes, excluding the serial number that identifies a unique vehicle. The VIN stub identifies the model, model year, and engine type of each vehicle. It typically also indicates the drive type and transmission, and it often indicates a unique trim or body style (such as an extended cab pickup truck). For example, an automatic transmission 4-cylinder 2.0L 2005 Toyota Camry will have a VIN stub that differs

Table 1: Summary Statistics

	Mean	Standard Deviation	Minimum	Maximum
Transaction Price (2008 USD)	10,294	7,424	1.03	131,277
Gasoline Price (2008 USD)	2.06	0.60	1.20	3.98
Present Value Remaining Fuel Cost (2008 USD)	7,787	3,333	1,511	32,175
Odometer (Miles)	58,205	25,213	10,000	99,999
Age (Years)	4.6	2.9	0.0	23.1
Fuel Economy (EPA Combined MPG)	21.9	4.51	10.0	46.0
Car Indicator Variable	0.62	0.49	0.0	1.0
Number of Observations		1,649,908		
Number of Unique VINs		9,076		

Summary statistics reported here are for the final estimation sample used in the baseline specification of vehicles sold by dealers.

from Camrys from different model years or with different engine sizes or transmission. We use the VIN stub as our definition of a vehicle type, and our specification allows each VIN stub to have unique month fixed effects and a unique depreciation schedule.

4 How Do Vehicles Depreciate with Mileage?

Our strategy isolates variation in fuel costs and vehicle price that is driven by variation in the remaining mileage on a particular vehicle, interacted with changes in the price of gasoline. Because we are isolating this variation, it is essential to accurately model the relationship between mileage and vehicle price. Any misspecification of this relationship could bias our results. Thus, before presenting our main results, we explore the relationship between price and odometer—that is, the way that vehicles depreciate with use—in our data.

To do so, we take a random sample of one million observations and explore the relationship between existing mileage at the time the vehicle is sold and price. We first plot the full sample, collapsed into 5,000 mile bins, in Figure 1. This shows that there is a strong relationship between vehicle price and odometer readings, but it also points to several anomalies.

First, at very high mileages, the relationship between mileage and price flattens out, and even begins to rise at the highest readings. This is likely the result of selection. Taking a vehicle to auction is costly; it takes time, and the vehicle must be transported. Sellers will not bear these costs for vehicles with a sufficiently low value. Thus, vehicles above 150,000 or 200,000 miles are

Figure 1: Transaction Prices by Existing Mileage
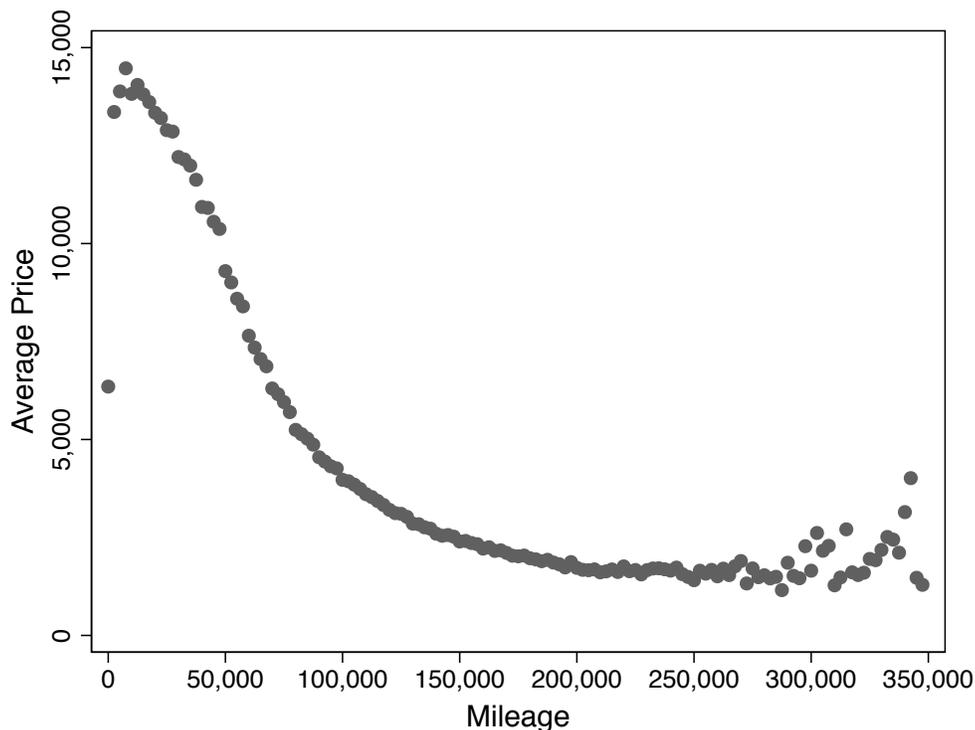
Figure plots average transaction prices by 5,000 mile bins.

unlikely to be brought to auction unless they are very valuable models or are in unusually good condition. Our empirical specification will account for the type of car and its initial value, but we have no way of accounting for the condition of an individual model (conditional on its odometer reading). This suggests that we truncate our data at lower odometer readings to avoid selection.

Figure 2 shows a histogram of our data by the same odometer categories. This shows that the data taper off quickly at the highest odometer ratings. The data also show a significant jump at 100,000 miles. There are discontinuously more vehicles in the auction data that have below 100,000 miles. To account for the selection involved, we restrict our main estimates to vehicles with fewer than 100,000 miles, and we also check for sensitivity when we expand our sample to include vehicles with up to 150,000 miles.

Second, Figure 1 shows that vehicles with very low mileage appear to be outliers; they have a much lower price than would be expected based on the other data. Further investigation reveals that the outliers primarily have mileage below 1,000 miles. It is unusual for very new vehicles

Figure 2: Histogram of Observations by Existing Mileage
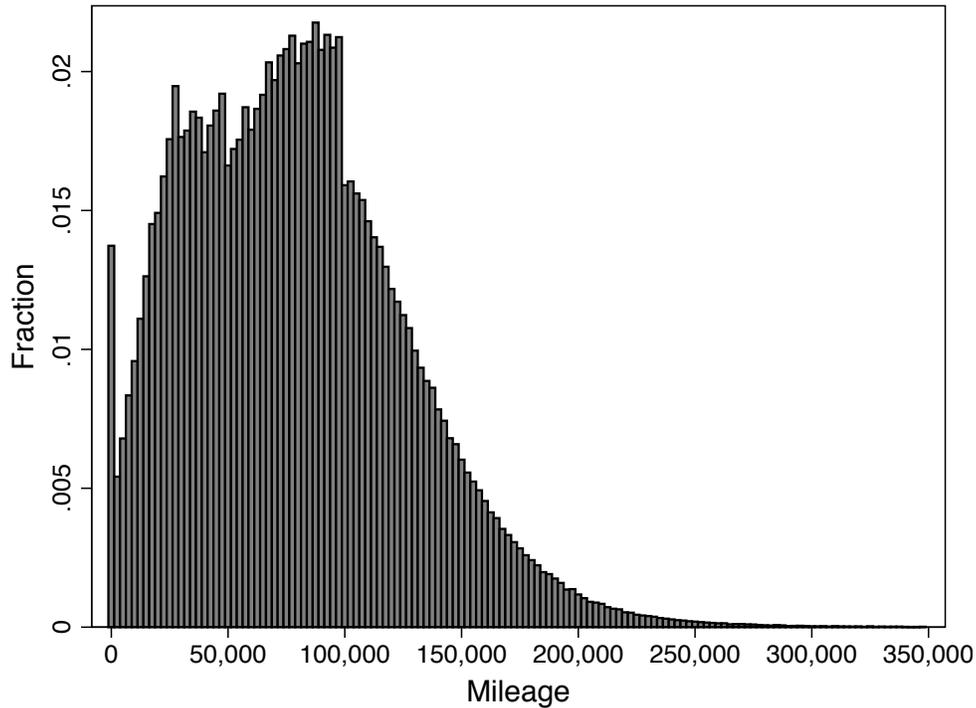
to appear in a wholesale auction because they were sold by consumers; yet Figure 2 shows that many such vehicles appear. We suspect that these may have been demonstration vehicles, vehicles that were returned by a customer under warranty due to some significant problem, vehicles that were in significant accidents, or they may even indicate coding errors. We thus drop all vehicles with odometer readings below 5,000 miles in the remaining analysis in this section, and, to be conservative, we limit our main sample to vehicles with at least 10,000.

We are interested in the relationship, in levels, between price and fuel costs, so we wish to estimate a model that has price levels on the left-hand side. But, it makes little sense to expect all vehicles to depreciate by the same amount *in price levels* (as opposed to percentages) over the mileage distribution. To demonstrate this, we calculate the average price over the entire sample period for each type of vehicle and, according to this average, define vehicles as either above or below average priced. Figure 3 plots the high and low priced vehicles separately to demonstrate that, in level terms, high priced vehicles depreciate faster. In general, different vehicles can be expected to depreciate according to a different schedule, and this is particularly true when measured in

Figure 3: Transaction Prices by Existing Mileage: High versus Low-Priced Models
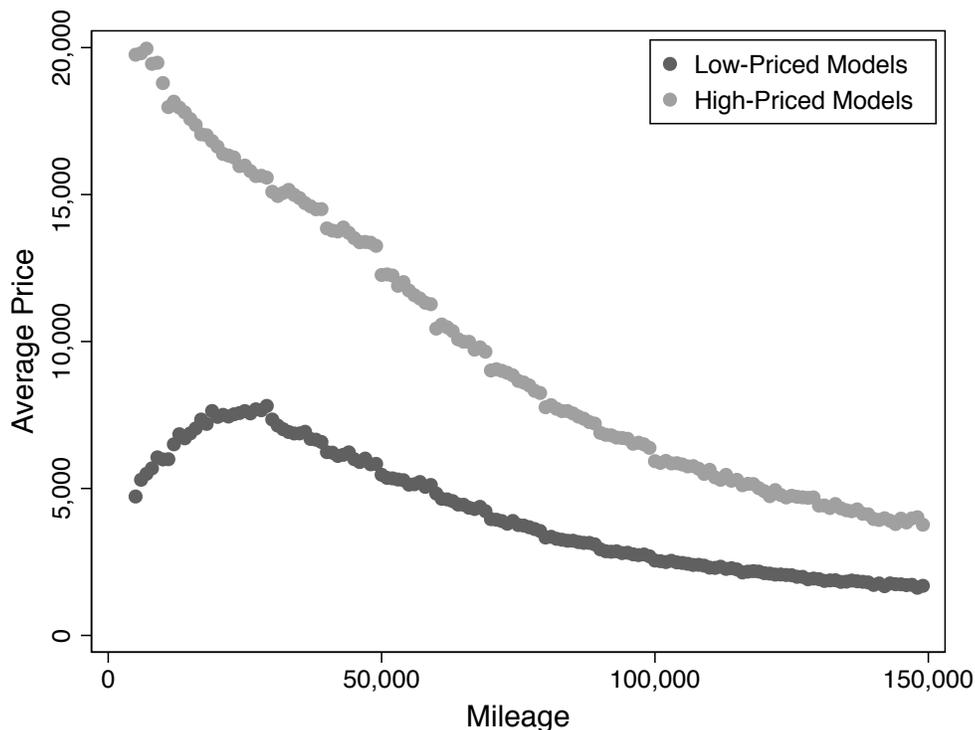
Figure plots average transaction prices by 5,000 mile bins, with sample divided into types of cars whose average price over the entire sample period is above or below the median sample average.

level dollars. Moreover, fuel economy is correlated with price (in our raw data this correlation is negative), which heightens concerns about bias if the depreciation schedule is not specified correctly for each model.

To account for this, we allow each type of vehicle (VIN stub) in our sample to depreciate according to a unique schedule. This differs substantially from existing literature, and it represents one of the key benefits of our microdata approach that leverages differences in cost and price across the same car that derive from differences in their odometer.[15]

Figure 1 suggests a fairly smooth relationship between odometer and price, and thus at first blush one might expect that a simple polynomial in mileage could account for depreciation. Figure 4, which plots the residuals—in 5,000 mile bins—from a regression of prices on a fifth-order polynomial

---

[15]Allcott and Wozny (2014) and Busse, Knittel, and Zettelmeyer (2013) both explore depreciation that varies by vehicle segments, but our approach is far less restrictive.

Figure 4: Transaction Price Residuals by Existing Mileage from Fifth-Order Polynomial
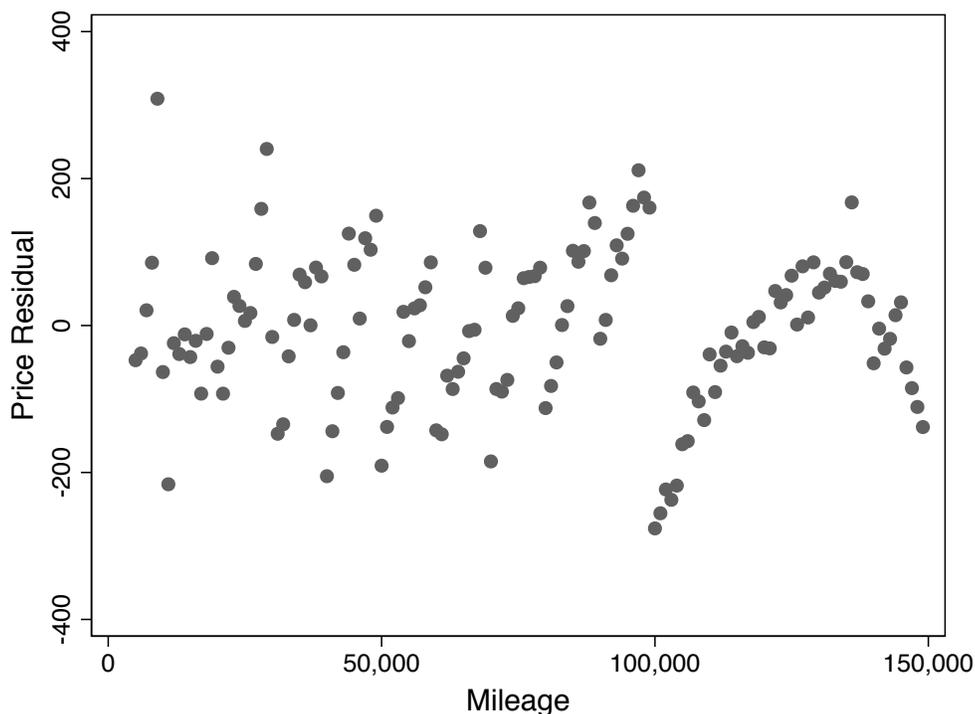
Figure plots average price residuals by 5,000 mile bins from a regression of prices on a fifth-order polynomial, omitting vehicles with fewer than 5,000 miles or more than 150,000 miles.

in mileage, shows that this approach suffers from at least two problems. First, the polynomial is a poor fit at higher odometer levels. In principle, this could be corrected by adding splines in addition to smooth polynomials or by simply increasing the order of the polynomial to very high orders. The second problem, however, is the "jigsaw" pattern of residuals at lower odometer levels. The residuals follow a pattern; they rise for several thousand dollars and then discretely jump down, and then rise again. This is due to discontinuities in vehicle prices at round numbers, which is documented by Lacetera, Pope, and Sydnor (2012), who interpret this as a sign of behavioral biases due to the salience of "left digits".

To fully account for this pattern, we follow Lacetera, Pope, and Sydnor (2012) and augment our polynomial by adding dummy variables for each 10,000 mile discrete category (e.g., there is a dummy equal to 1 for vehicles with 20,000 to 29,999 miles), and a variable that is a linear count of 1,000 mile bins (e.g., vehicle with between 20,000 and 20,999 will each have a thousand-mile count

17

Figure 5: Transaction Price Residuals by Existing Mileage from Seventh-Order Polynomial with Categorical Controls
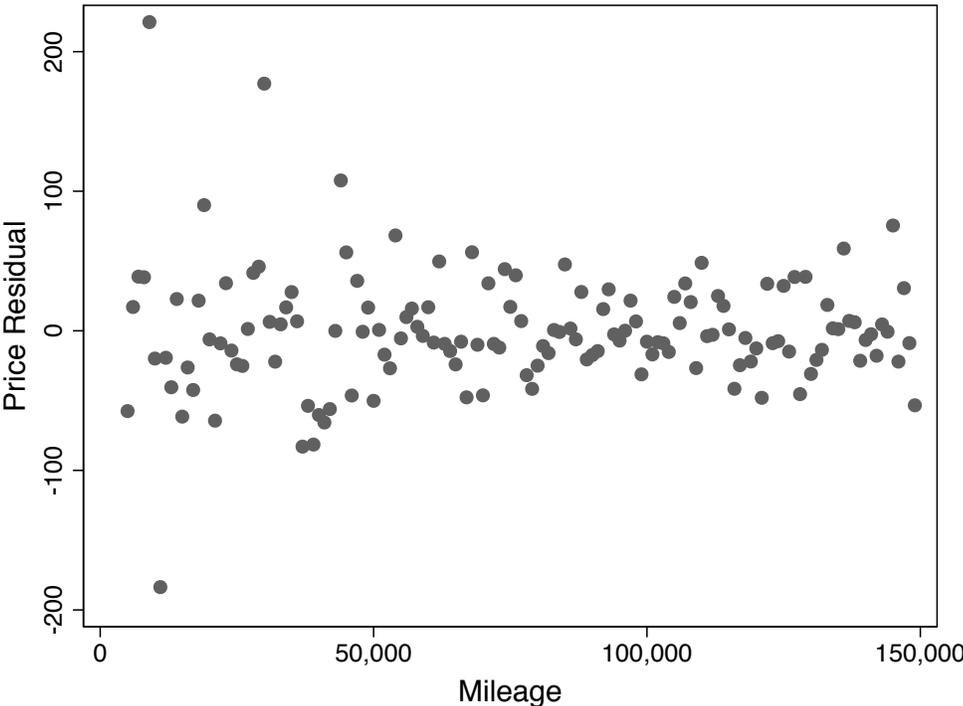
Figure plots average price residuals by 5,000 mile bins from a regression of prices on a seventh-order polynomial, a set of dummy variables for each 10,000 odometer category and a linear control for discrete 1,000 mile categories. The sample omits vehicles with fewer than 5,000 miles or more than 150,000 miles.

variable equal to 20, whereas a vehicles with 21,000 will have a value of 21). The addition of these terms helps the fit of the polynomial considerably. We also raise the order of the polynomial to seven in our baseline model.

Figure 5 shows the residuals from a regression of price on a seventh-order polynomial in mileage, with dummies for each 10,000 mile category, and the linear one thousand mile discrete category variable. This specification eliminates the patterns in the residuals, which leads us to use this specification as our baseline specification for modeling vehicle depreciation. Figure 5 also estimates a single (level) depreciation curve for all models, which hides the variation across types that is highlighted above in Figure 3. Our preferred estimator allows each model to depreciate in a unique pattern.

It is important to note that this set of odometer control variables accounts for the vast majority of the overall variation in vehicle prices and future fuel costs. We suspect that the remaining variation, which comes from differences in remaining lifetime within vehicles of the same type sold in the same month, is less salient to consumers and thus our specification presents a harsher test of the full valuation hypothesis than the approach taken in the related literature.

5   Results

Having arrived at a preferred specification through the examination of the relationship between price and odometer, we use our data to estimate the relationship between prices and present discounted value future fuel costs in the following estimating equation:

$$P_{ijt} = \beta C_{ijt} + \delta_{jt} + \sum_{j=1}^{J} \left(\sum_{a=1}^{7} \alpha_{ja} O_{ijt}^a + \sum_{k=1}^{10} \gamma_{jk} d_{ijt}^k + \rho_j l_{ijt} \right) + \varepsilon_{ijt} \tag{4}$$

$$\text{where } d_{ijt}^k = \begin{cases} 1 & \text{if } (k-1) \times 10{,}000 < O_{ijt} < k \times 10{,}000 \\ 0 & \text{otherwise} \end{cases}$$

$$l_{ijt} = \text{integer floor}(O_{ijt}/1000).$$

In equation 4, P_{ijt} is the real transaction price of vehicle i of VIN stub type j sold in month t, and C_{ijt} is the remaining future fuel cost for that vehicle. The additional controls include δ_{jt},

which is a vector of VIN stub by month of sample fixed effects (e.g., a fixed effect for an automatic 4-cylinder 2.0L 2005 Toyota Camry in January 2008), and a function of the vehicle's odometer reading, denoted O_{ijt}.

The odometer function includes a polynomial, denoted by the O_{ijt}^a terms, where a is an exponent ranging, in our baseline model, from one to seven. It also includes a set of dummy variables for each 10,000 odometer bin, denoted as γ_{jk}, and a linear control for the discrete thousand mile bin, denoted l_{ijt}, for each vehicle. As described above, these discrete odometer controls are used to account for the discontinuous relationship between price and mileage that is analyzed by Lacetera, Pope, and Sydnor (2012). Importantly, the odometer control function is estimated separately for each vehicle type j, which is indicated by the outside summation term. This allows each type of vehicle to depreciate in a unique way, which we believe is essential in a regression on price levels and in a sample with a diverse set of vehicles.

Our object of interest is β. If consumers fully value changes in fuel economy, we would expect β to be equal to negative one. That is, a one dollar increase in the present discounted value of the cost of operating a vehicle over its remaining life should correspond to a one dollar drop in consumer willingness to pay for the vehicle, which, under the assumptions about market structure that are outlined in our conceptual model (section 2), will translate into a one dollar drop in price. Along with β, the δ, α, γ and ρ terms are all parameters to be estimated.

Explaining how we actually estimate this equation may be useful for explaining its full flexibility. The equation includes many thousands of parameters. To estimate the model we make use of the Frisch-Waugh-Lovell theorem and first regress prices P_{ijt}, and then costs C_{ijt}, on time period fixed effects and the odometer control variables for each vehicle type j separately, one at a time. We then regress the residuals from these two sets of regressions on each other to recover $\hat{\beta}$, which returns numerically identical coefficients to the full model estimated in a single step. (We collect the number of parameters estimated and use them to make the proper degrees of freedom correction when calculating standard errors in the second step.) We exclude VIN stubs for which we have insufficient observations to estimate the first-step regression.

Table 2 reports our main results. Our baseline estimate yields a coefficient estimate of -0.98, with a standard error of 0.04.[16] This estimate is statistically indistinguishable from -1, which represents

---

[16] All standard errors are clustered at the VIN stub level.

Table 2: Effects of Future Fuel Cost Shocks on Transaction Prices

	(1)	(2)	(3)	(4)	(5)
Future fuel cost	-0.98	-1.01	-0.94	-0.98	-0.98
	(0.04)	(0.03)	(0.06)	(0.03)	(0.04)
Number of Observations	1,649,908	1,019,034	630,874	1,673,031	1,641,223
Number of Unique VINs	9,076	5,456	3,620	9,577	8,915
Polynomial Order	7	7	7	5	9
Class	All	Cars	Trucks	All	All
Minimum Mileage	10,000	10,000	10,000	10,000	10,000
Maximum Mileage	100,000	100,000	100,000	100,000	100,000

The dependent variable is the real transaction price. All specifications include month of sample by VIN stub fixed effects, a polynomial in odometer, a set of dummy variables for each 10,000 odometer bin, and a linear control for 1,000 odometer categories. Standard errors are clustered on VIN stub. Sample includes vehicles sold by dealers. Future fuel cost calculations use a 5% discount rate.

the full valuation null hypothesis. It is precise enough to rule out significant undervaluation or overvaluation. Our estimate changes little when we study only cars or only light-trucks (columns 2 and 3). Decreasing the order of the control polynomial to five, or increasing it to nine, has almost no effect on our final estimate (columns 4 and 5).[17]

Our baseline sample includes vehicles with odometer readings between 10,000 and 100,000 miles. This is because, as discussed in section 4, there is a discontinuity in the volume of vehicles sold at auction at 100,000 miles, and the lowest odometer vehicles are outliers. To see how this restriction affects our estimates, Table 3 reports coefficients from samples with alternative odometer ranges. Lowering the top odometer rating from 100,000 to 90,000 or 80,000 or 70,000 has little effect on our estimated coefficient.[18] Raising the odometer range to exceed 100,000, however, causes our coefficient to decline. The coefficient declines monotonically, and to an economically important degree, as we move from 100,000 up to 150,000. When we isolate the sample to include *only* vehicles above 100,000 miles (column 10), our coefficient estimate falls precipitously.

There are several possible interpretations of this final result. First, it may be that consumers who buy newer used vehicles properly value fuel economy, whereas consumers who buy high mileage vehicles substantially undervalue it. In the high-mileage segment of the used car market, consumers

---

[17]The number of observations changes slightly across specifications with different polynomial orders because adding (or subtracting) regressors in the first step changes whether or not the first-step regression has sufficient observations for a number of VIN stubs, each of which has a modest number of observations.

[18]In additional results (not shown), we find that our estimates are similarly insensitive to lowering the minimum odometer to either 5,000 or 1,000 miles.

Table 3: Effects of Future Fuel Cost Shocks on Transaction Prices

	(1)	(2)	(3)	(4)	(5)
Future fuel cost	-0.93	-1.02	-1.00	-0.98	-0.94
	(0.07)	(0.06)	(0.05)	(0.04)	(0.03)
Number of Observations	969,142	1,190,655	1,416,561	1,649,908	1,822,943
Number of Unique VINs	7,212	7,914	8,499	9,076	9,393
Polynomial Order	7	7	7	7	7
Minimum Mileage	10,000	10,000	10,000	10,000	10,000
Maximum Mileage	70,000	80,000	90,000	100,000	110,000
	(6)	(7)	(8)	(9)	(10)
Future fuel cost	-0.87	-0.79	-0.73	-0.69	-0.26
	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)
No. of Observations	1,975,025	2,101,451	2,201,968	2,278,386	499,199
No. of Unique VINs	9,622	9,791	9,906	9,976	4,232
Polynomial Order	7	7	7	7	7
Minimum Mileage	10,000	10,000	10,000	10,000	100,000
Maximum Mileage	120,000	130,000	140,000	150,000	150,000

The dependent variable is the real transaction price. All specifications include month of sample by VIN stub fixed effects, a polynomial in odometer, a set of dummy variables for each 10,000 odometer bin, and a linear control for 1,000 odometer categories. Standard errors are clustered on VIN stub. Sample includes vehicles sold by dealers.

might be myopic, or they may simply value other attributes more, which leads them to be inattentive to energy efficiency. A focus on other features of the car, such as its condition and accident history, may even be motivated by concerns related to asymmetric information.

Second, consumers who purchase older, lower priced vehicles may, on average, have substantially higher discount rates. Assuming that they are lower income on average, they are more likely to face liquidity constraints and will be forced to borrow at higher interest rates, which will (rationally) drive up the rate at which they trade off future discounted fuel costs against price.

Third, the result could be driven by selection. As argued above, the set of very high mileage vehicles that appear at wholesale auctions are a selected group. Among this selected group, our assumption about remaining lifetime mileage, which is based on the national averages reported in Lu (2006), may be biased. Critically, it could be that consumers who buy older vehicles do so because they intend to drive very few miles. Thus, we choose to emphasize our baseline sample, but the results for the high odometer vehicles is a notable caveat to our main results.

Our results align well with the existing literature. Allcott and Wozny (2014) find significant undervaluation among the oldest models in their sample, but full (or nearly full) valuation among

the newest models.[19] Busse, Knittel, and Zettelmeyer (2013), who intepret their findings as consistent with full valuation, do not report estimates separately by the age of vehicles in their sample, but their sample is drawn from vehicles sold at used car dealerships that also have a new car retail business. This leads their sample to be relatively low odometer (and high priced). Their sample is thus more similar to ours when we restrict to lower odometer vehicles. (Even then, our vehicles are somewhat older and less expensive on average.)

Finally, equation 4 includes a single cost coefficient estimated for all vehicle types, but our micro data approach enables us to calculate the coefficient on cost for each VIN stub separately— that is, we can estimate a vector β_j rather than a single scalar β. For our baseline sample and specification, we estimated these cost coefficients separately and plotted them in Figure 6. Because the coefficient for any individual VIN is estimated with considerable error, the point estimates show a large spread. Overall, however, we find it comforting that the distribution is centered around our unified estimate and appears to be roughly symmetric and, as suggested by econometric theory, normally distributed.

5.1 Vehicles Sold by Manufacturers and Fleets

Our data include vehicles sold by dealers, fleet operators and the manufacturers themselves. Our baseline results include only those vehicles sold by dealers because we expect that this sample of vehicles, which were previously owned by "normal" consumers will best match the depreciation and scrappage schedule established by Lu (2006), which is a critical component of our measure of future fuel cost.

Vehicles sold by fleet operators, such as large companies, rental car companies, or governments, are likely to have been used and maintained differently than the typical vehicle owned by a consumer household. Similarly, vehicles that were owned by the automobile manufacturers themselves for some time, for use by employees or for other purposes, may also differ from the average in their use and maintenance history. For example, in our sample, both fleet and manufacturer sold vehicles have much higher annual mileage than vehicles sold by dealers.[20] As a result, these vehicles may

---

[19]When Allcott and Wozny (2014) use mileage and scrappage schedules from Lu (2006), as we do, they get an estimate of -1.03 for the newest models (1 to 3 years old) and -0.28 for the oldest models (11 to 15 years old) in their sample (see their Table 5, column 2). This matches closely to our estimates of -0.98 for the baseline and -0.26 for the highest mileage vehicles.

[20]A full set of summary statistics for fleet and manufacturer sold vehicles are included in the appendix.

Figure 6: Histogram of Valuation Coefficient for Individual VIN Stubs
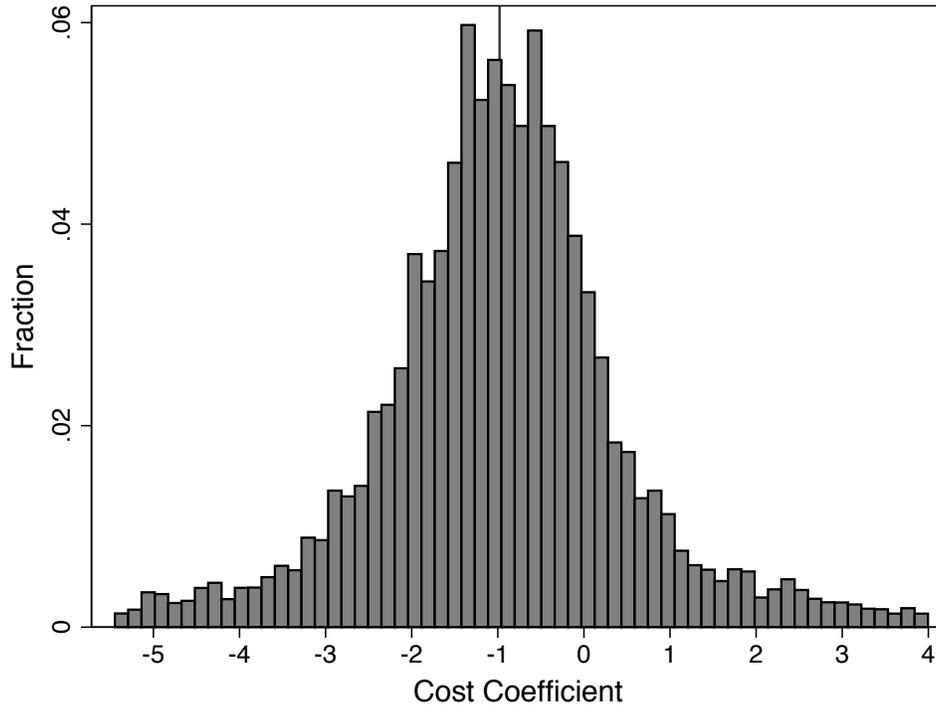
Each data point is the cost coefficient from equation 4 estimated for a single VIN stub. The figure truncates the highest and lowest 2.5% of coefficients. The vertical line is the point estimate from the joint regression from the same specification, which is found in Table 2, column 1.

have expected remaining lifetimes, conditional on current mileage, that differ from the national average.

Nevertheless, we present our baseline specification for vehicles sold by fleet operators and manufacturers in Table 4 for comparison. Column 1 repeats our main estimate from Table 2 for reference, and columns 3 and 5 show the corresponding result for fleet and manufacturer sold vehicles, respectively. Estimates from these samples suggest modest undervaluation, assuming a 5% discount rate. The point estimate for fleet operators is -0.87, while that for vehicles sold by manufacturers is -0.77. This could imply that these automobiles have indeed been depreciated differently than vehicles sold by dealers at auction.

The even numbered columns in Table 4 show what happens if we recalculate our estimates assuming a 10% discount rate, instead of our baseline 5%. This change raises the coefficient

Table 4: Effects of Future Fuel Cost Shocks on Transaction Prices by Type of Seller

	(1)	(2)	(3)	(4)	(5)	(6)
	Dealers		Fleet Operators		Manufacturers	
Discount rate	5%	10%	5%	10%	5%	10%
Future fuel cost	-0.98	-1.38	-0.87	-1.22	-0.77	-1.11
	(0.04)	(0.05)	(0.03)	(0.04)	(0.07)	(0.11)
Number of Observations	1,649,908	1,649,908	2,107,541	2,107,541	1,146,855	1,146,855
Polynomial Order	7	7	7	7	7	7
Number of Unique VINs	9,076	9,076	9,302	9,302	4,464	4,464
Minimum Mileage	10,000	10,000	10,000	10,000	10,000	10,000
Maximum Mileage	100,000	100,000	100,000	100,000	100,000	100,000

The dependent variable is the real transaction price. All specifications include month of sample by VIN stub fixed effects, a polynomial in odometer, a set of dummy variables for each 10,000 odometer bin, and a linear control for 1,000 odometer categories. Standard errors are clustered on VIN stub.

estimates enough so that the full valuation null hypothesis cannot be rejected in any of the three samples. The point estimates actually suggest overvaluation of fuel economy.

5.2 What is the Value of Our Approach?

Our baseline model includes a very rich set of controls. This has two principal advantages. First, by allowing each VIN to have its own depreciation schedule, we eliminate any concerns about correlations between fuel economy, prices and misspecification of depreciation that could bias sparser specifications. Second, by isolating variation within VIN types within a time period, we are able to control for a very wide set of demand shocks nonparameterically. These features are very appealing ex ante, but we can check whether these flexibilities actually matter by running more limited specifications and comparing them to our results.

Table 5 shows our baseline sample estimated with sparser models. Column 1 repeats our baseline results for reference. Column 2 shows results from a regression that includes only VIN stub fixed effects, time period fixed effects (not interacted with VIN stubs), and a single odometer control function that applies to all models. This specification is a standard difference-in-difference estimator with a single odometer control function. It yields an estimate of -0.83, which suggests modest undervaluation.

In column 3 we use vehicle segment definitions to categorize our vehicles as falling into one of

Table 5: Effects of Future Fuel Cost Shocks on Transaction Prices: Alternative Specifications

	(1) Baseline	(2) Diff-in-Diff	(3) Richer Odo	(4) Richer Time
Future fuel cost	-0.98	-0.83	-0.97	-0.48
	(0.04)	(0.03)	(0.03)	(0.04)
Number of Observations	1,649,908	1,649,908	1,649,908	1,649,908
Time Fixed Effects	Interacted with VIN	Not Interacted	Not Interacted	Interacted with Segment
Odometer Controls	Interacted with VIN	Not Interacted	Interacted with Segment	Interacted with Segment
Minimum Mileage	10,000	10,000	10,000	10,000
Maximum Mileage	100,000	100,000	100,000	100,000

The dependent variable is the real transaction price. Standard errors are clustered on VIN stub.

eight segments.[21] We then interact our odometer control function with each segment. This allows each segment to have its own depreciation schedule. Our coefficient estimate rises to -0.97, which is almost exactly the same as our baseline.

When we add segment by time fixed effects in column 4, however, which allow the time period effects to have different affects for each segment, the coefficient falls to -0.48. This result is qualitatively different from our baseline and suggests substantial undervaluation. The sensitivity of the results to the structure of the controls suggests the value of our approach. The alternative specifications are all restrictions on our baseline model, which allows for VIN stub specification time period shocks and depreciation schedules. Thus, where the alternatives create disagreement, it is logical to prefer our approach because it relaxes restrictive assumptions.

In terms of the existing literature, the results in Busse, Knittel, and Zettelmeyer (2013) are also sensitive specifically to the inclusion of time period effects that are interacted with vehicle segment (see column 2 of Table 4 in that paper, where the inclusion of such controls causes a significant drop in the estimated effect of a gasoline price shock on relative vehicle prices, which translates to a sharp rise in the implied discount rate). That paper argues that alternative specifications are a priori preferred. Our analysis here provides direct evidence on the issue by relaxing the restrictions of the specification further.

---

[21]We use a standard industry source to determine classifications. The eight segments are Compact Cars, Midsize Cars, Fullsize Cars, Luxury Cars, Sports Cars, Pickup Trucks, SUVs and Vans.

5.3 How Do Wholesale Price Changes Influence Retail Prices?

Our interest is in how consumer prices change in response to future fuel cost changes, but our data are from wholesale auctions. Dealers purchase vehicles at these auctions and pass them on to consumers. To the extent that the used car market is competitive, we would expect wholesale price shocks to be passed on one for one into retail consumer prices. If so, then our regression coefficients can be directly interpreted as the effect of fuel cost shocks on consumer prices. Alternatively, if used car dealers add a proportional markup over the purchase price, then our coefficient estimates should be "scaled up". We find support for a one to one relationship between wholesale and retail prices using an auxiliary data set.

To examine the relationship between wholesale and retail prices, we use data from Kelley Blue Book. To measure wholesale prices, Kelley Blue Book collects data from auctions. To measure retail prices, they gather data on actual transactions from dealers and other market data sources. We gathered the wholesale and retail prices of all available cars and light-trucks in the July edition of the Kelley Blue Book guide from 2003 to 2008. We then regress a vehicle's retail price on its wholesale price to determine price pass through.

Table 6 reports our results. The first column presents simple OLS results. The unit of observation is a particular model and vintage sold in July of each year. We include year dummy variables and age dummy variables in all specifications. In the OLS specification, most of the variation in prices comes from differences across models. The OLS coefficient implies that when wholesale prices rise by $1, retail prices rise by $1.03. This is not statistically different from the one to one benchmark.[22]

We change the source of identification by adding model by vintage fixed effects in column 2. In this specification, pass through is identified only from changes in the price of particular models over time, controlling for year and age effects. These fixed effects soak up a majority of the variation, which is reflected in a much higher standard error. The coefficient falls significantly, though it is still not statistically different from 1. The reduction in the coefficient may reflect attenuation from measurement error, which is exacerbated by the panel strategy.

---

[22]This does not imply that retail and wholesale prices are the same. Rather, our data show that average markups of retail prices over wholesale prices are quite high, on the order of 30%. But, this may reflect fixed costs of shipping, repairing, holding and selling vehicles that do not vary with the wholesale price of the vehicle.

Table 6: Pass Through of Wholesale Price Changes to Retail Prices

	OLS	FE	IV
	(1)	(2)	(3)
Wholesale Price	1.03	0.64	1.05
	(0.06)	(0.21)	(0.05)
Number of Observations	29,318	29,318	29,318
Year Dummies	Yes	Yes	Yes
Age Dummies	Yes	Yes	Yes
VIN Stub Dummies	No	Yes	Yes
First-stage F-statistic			428

Dependent variable is the retail price of a vehicle, as reported by Kelley Blue Book. In column 3, the instrument is the interaction between each vehicle's fuel economy rating and the current real price of gasoline.

Both to overcome possible measurement error attenuation in the panel and to provide a more solid identification strategy, we use an instrumental variables approach in the third column. We instrument wholesale prices using the interaction of fuel economy and the price of gasoline in that year. This instrument isolates variation in wholesale prices due to price shocks associated with gasoline price changes. This is a valid instrument if the only way that fuel economy interacted with gasoline prices affects retail prices is "through wholesale prices", which is consistent with interpreting wholesale prices as the outside option for a dealer who is deciding whether to hold and sell a particular model or to take it to auction and swap it for another model to have on their lot. Our instrument is very powerful in the first stage.

The IV estimates are quite similar to the OLS. The point estimate implies a very small proportional markup, at 5%, of retail over wholesale prices. Taken at face value, this implies that our estimates from the auction data could be scaled up by 5%, which would change none of our qualitative conclusions. The standard error is not tight enough to rule out perfect pass through, however, and so we prefer to interpret these results as a failure to reject a null hypothesis of one to one pass through. This is consistent with the used car market being competitive. Under that interpretation, our main results can be interpreted directly as estimates of consumer valuation without any rescaling to account for retail markups.

5.4 Interpretation and Caveats

Our baseline estimates provide results that are precisely estimated and statistically indistinguishable from the full valuation null hypothesis. We interpret this as a failure to find evidence in support of consumer undervaluation. It is important to note, however, that the testing of this hypothesis requires a series of assumptions that are embedded in the construction of our future fuel cost variable. Just as in the other papers in this literature, our approach requires us to take a stand regarding consumer's beliefs about future fuel prices, about the schedule of vehicle mileage and scrappage, and about the discount rate. Our results would move, mechanically, if we changed any of these assumptions, and indeed the results in Table 4 suggest that the coefficients could suggest meaningful overvaluation at higher discount rates.

Additionally, our specification (again, like others in the literature) makes simplifying assumptions about the structure of the used car market. It is not obvious that any of these assumptions should create a bias that exaggerates our estimates, but it is important to understand the assumptions involved. Thus, we briefly discuss these implicit assumptions before concluding.

The New Car Market

Past literature has either assumed that supply responses in the new car market, which may be endogenous to the price of gasoline, do not affect used car prices (Kahn 1986; Busse, Knittel, and Zettelmeyer 2013) or attempted to account for the new car market using a vehicle choice model (Allcott and Wozny 2014). By including vehicle type by time fixed effects, we control nonparametrically for any effect that new car supply decisions have on used car prices, so long as the effect is common across vehicles with the range of odometer readings that appear in the sample.

Even so, the supply elasticity of the new car market could influence results. To see why, suppose that the supply for a particular vehicle in the used car market is perfectly inelastic. What we (and others) identify is the change in the willingness to pay for each model. But, the change in the willingness to pay for a model depends also on the prices of alternatives. (This point is made empirically in Langer and Miller (2013).) The elasticity of the new car market might cause a correlation in price shocks for competing products that could drive the willingness to pay to move more for some products than others in response to a gasoline price shock.

The intuitive case is that this should bias us against finding full valuation, which mitigates this concern given our results. To see why, consider, for example, a rise in the price of gasoline. This will cause a contraction in the sales of new cars that are inefficient and a (relative) expansion in the sales of new cars that are efficient. This supply effect will dampen the change in demand for efficient used cars relative to the change in demand for inefficient used cars, which works against the main effect we mean to isolate. We suspect that, in practice, the new and used car markets are sufficiently segregated that, while they undoubtedly affect each other even in the short run, the effects are being picked up sufficiently by our fixed effects. This would be true so long as new car market considerations did not differentially affect the same type of vehicle with different current mileage.

Heterogeneity

We do not explicitly model heterogeneity among consumers in their driving behavior, but of course some consumers drive more than others. If the scrappage decision, as a function of odometer, is common across individuals, then heterogeneity in miles driven per period will generate differences in valuation only through intertemporal shifts in mileage that influence the cost function through discounting. But, our data are wholesale prices. This means that dealers buying the vehicle do not know if they will be selling to a high or low mileage customer, so this variation should not actually influence our estimates, provided that the distribution of mileage across individuals buying a given type do not change with the price of gasoline.

Consumers may be heterogeneous in a variety of ways. It is surely the case that consumers have some variation in their discount rates, their degree of myopia, and their gasoline price expectations, but we model the problem as if there is a single representative agent. Bento, Li, and Roth (2012) argue that ignoring heterogeneity in fuel economy valuation is likely to cause downward bias. If the bias is in fact downward, that would work against our finding of full valuation. Note, however, that Grigolon, Reynaert, and Verboven (2014) explicitly model heterogeneity and conclude that, when allowing heterogeneity in both fuel economy valuation and in other coefficients, for their data set, this bias turns out to be minimal.

6   Conclusion

We use a new approach to answer the question of whether or not consumers properly value fuel economy in the market for automobiles. Our approach isolates a type of variation in fuel costs that represents a very small fraction of the overall variation. We do so for reasons of research design, but we note also that this method could be argued to isolate variation in fuel costs that are particularly non-salient to consumers, especially compared to the cross quartile design of Busse, Knittel, and Zettelmeyer (2013). Thus, ex ante, we think that we have presented a more stringent test for consumer rationality.

Our baseline estimates are consistent with full valuation at low discount rates. These results are robust to several specification alternatives. The conclusion that consumers fully value fuel economy is, however, subject to several caveats. First, to estimate the remaining lifetime of vehicles, which is critical for our estimation, we rely on mileage and vehicle scrappage schedules from a single government study. This prevents us from directly measuring differences in lifetime mileage across models. Second, we do find evidence of modest undervaluation (equivalently, of higher discount rates) in several specifications, especially among vehicles that have higher odometer readings.

We emphasize, however, that given our isolation of a limited source of fuel cost variation that is likely especially difficult for consumers to understand and calculate, even these coefficients are strong evidence that consumers are responsive to fuel costs. Even the most pessimistic reading of our evidence firmly rejects the assumption used in regulatory impact analyses that consumers are completely inattentive to fuel costs.

References

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. "Energy Policy with Externalities and Internalities." *Journal of Public Economics* 112:72–88.

Allcott, Hunt and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Review of Economics and Statistics* 96 (5):779–795.

Alquist, Ron and Lutz Kilian. 2010. "What Do We Learn From the Price of Crude Oil Futures?" *Journal of Applied Econometrics* 25 (4):539–573.

Alquist, Ron, Lutz Kilian, and Robert J. Vigfusson. 2013. "Forecasting the Price of Oil." In *Handbook of Economic Forecasting*, vol. 2, edited by Graham Elliot and Allan Timmerman. Amsterdam: North-Holland, 427–507.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe About Future Gasoline Prices?" *Journal of Environmental Economics and Management* 66 (3):383–403.

Anderson, Soren T., Ian W.H. Parry, James M. Sallee, and Carolyn Fischer. 2011. "Automobile Fuel Economy Standards: Impacts, Efficiency and Alternatives." *Review of Environmental Economics and Policy* 5 (1):89–108.

Bento, Antonio M., Shanjun Li, and Kevin Roth. 2012. "Is There an Energy Paradox in Fuel Economy? A Note on the Role of Consumer Heterogeneity and Sorting Bias." *Economics Letters* 115 (1):44–48.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1):220–256.

Dreyfus, Mark and Kip Viscusi. 1995. "Rate of Time Preference and Consumer Valuations of Automobile Safety and Fuel Efficiency." *Journal of Law and Economics* 38 (1):79–98.

Dubin, Jeffrey A. and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52 (2):345–362.

Environmental Protection Agency. 2012. "Final Rulemaking for 2017-2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards." EPA-420-R-12-016.

———. 2014. "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2012." EPA 430-R-14-003.

Espey, Molly and Santosh Nair. 2005. "Automobile Fuel Economy: What Is it Worth?" *Contemporary Economic Policy* 23 (3):317–323.

Fischer, Carolyn, Winston Harrington, and Ian W.H. Parry. 2007. "Should Automobile Fuel Economy Standards Be Tightened?" *The Energy Journal* 28 (4):1–29.

Fullerton, Don and Sarah E. West. 2002. "Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions?" *Journal of Environmental Economics and Management* 42 (1):135–157.

———. 2010. "Tax and Subsidy Combinations for the Control of Car Pollution." *The B.E. Journal of Economic Analysis & Policy (Advances)* 10 (1).

Gillingham, Kenneth. 2011. "How Do Consumers Respond to Gasoline Price Shocks? Heterogeneity in Vehicle Choice and Driving Behavior." Manuscript: Yale University.

Goh, Vincent, Paul S. Fischbeck, and David Gerard. 2007. "Identifying and Correcting Errors with Odometer Readings from Inspection and Maintenance Data: Rollover Problems for Estimation of Emisions and Technical Change." *Transportation Research Record* (2011):87–97.

Goldberg, Pinelopi Koujianou. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S." *The Journal of Industrial Economics* 46 (1):1–33.

Greene, David L. 2010. "How Consumers Value Fuel Economy: A Literature Review." U.S. Environmental Protection Agency report EPA-420-R-10-008.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2014. "Consumer Valuation of Fuel Costs and teh Effectiveness of Tax Policy: Evidence from the European Car Market." Manuscript: University of Leuven.

Hausman, Jerry A. 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *The Bell Journal of Economics* 10 (1):33–54.

Helfand, Gloria and Ann Wolverton. 2009. "Evaluating the Consumer Response to Fuel Economy: A Review of the Literature." National Center for Environmental Economics Working Paper 09-04.

Heutel, Garth. 2011. "Optimal Policy Instruments for Externality-Producing Durable Goods Under Time Inconsistency." NBER Working Paper 17083.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *The Energy Journal* 29 (1).

Jacobsen, Mark R. and Arthur A. van Benthem. Forthcoming. "Vehicle Scrappage and Gasoline Policy." *American Economic Review* .

Jaffe, Adam and Robert Stavins. 1994. "The Energy Paradox and the Diffusion of Conservation Technology." *Resource and Energy Economics* 16:91–122.

Kahn, James A. 1986. "Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach." *The Quarterly Journal of Economics* 101 (2):323–339.

Kilian, Lutz and Eric R. Sims. 2006. "The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data." Manuscript: University of Michigan.

Knittel, Christopher R. and Ryan Sandler. 2012. "Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins." In *The Design and Implementation of U.S. Climate Policy*, edited by Don Fullerton and Catherine Wolfram. NBER, 287–299.

Lacetera, Nicola, Devin Pope, and Justin Sydnor. 2012. "Heuristic Thinking and Limited Attention in the Car Market." *American Economic Review* 102 (5):2206–2236.

Langer, Ashley and Nathan H. Miller. 2013. "Automakers' Short-Run Responses to Changing Gasoline Prices and the Implications for Energy Policy." *Review of Economics and Statistics* 95 (4):1198–1211.

Li, Shanjun, Christopher Timmins, and Roger H. von Haefen. 2009. "How Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy* 1 (2):113–137.

Linn, Joshua and Thomas Klier. 2010. "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data." *American Economic Journal: Economic Policy* 2 (3):134–153.

Lu, S. 2006. "Vehicle Survivability and Travel Mileage Schedules." Technical Report DOT HS 809 952, National Highway Traffic Safety Administration.

Parry, Ian W.H., Margaret Walls, and Winston Harrington. 2007. "Automobile Externalities and Policies." *Journal of Economic Literature* 45 (2):374–400.

Sallee, James M. 2011. "The Taxation of Fuel Economy." *Tax Policy and The Economy* 25:1–37.

———. 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics* 57 (3):781–820.

Appendix A   Additional Tables

Table 7: Summary Statistics: Fleet Transactions

	Mean	Standard Deviation	Minimum	Maximum
Transaction Price (2008 USD)	12,977	7,149	1.05	161,697
Gasoline Price (2008 USD)	2.04	0.60	1.20	3.98
Present Value Remaining Fuel Cost (2008 USD)	7,108	3,290	980	34,701
Odometer (Miles)	43,797	21,765	10,000	99,999
Age (Years)	2.8	1.7	0.0	20.2
Fuel Economy (EPA Combined MPG)	22.0	4.49	13.0	46.0
Car Indicator Variable	0.62	0.49	0.0	1.0
Number of Observations		2,107,872		
Number of Unique VINs		9,312		

Summary statistics reported here are for the final estimation sample of vehicles sold by fleet operators.

Table 8: Summary Statistics: Manufacturer Transactions

	Mean	Standard Deviation	Minimum	Maximum
Transaction Price (2008 USD)	15,940	5,612	55.02	79,511
Gasoline Price (2008 USD)	1.91	0.54	1.20	3.98
Present Value Remaining Fuel Cost (2008 USD)	5,227	2,672	398	26,922
Odometer (Miles)	22,874	10,255	10,000	99,993
Age (Years)	1.0	0.9	0.0	13.25
Fuel Economy (EPA Combined MPG)	22.7	3.65	13.0	40.6
Car Indicator Variable	0.75	0.43	0.0	1.0
Number of Observations		1,146,925		
Number of Unique VINs		4,467		

Summary statistics reported here are for the final estimation sample of vehicles sold by manufacturers.