DRAFT: PRELIMINARY AND INCOMPLETE
PLEASE CITE ONLY WITH PERMISSION

The Effect of Gasoline Prices on the Demand for Fuel Economy in Used Vehicles: Empirical Evidence and Policy Implications

James M. Sallee
The Harris School
University of Chicago
sallee@uchicago.edu

Sarah E. West
Macalester College
wests@macalester.edu

Wei Fan
RVI Group
wfan@rvigroup.com

Version 1.2

June 18, 2010

Key words: fuel economy, undervaluation, myopia, gasoline prices, vehicle prices, expectations

JEL codes: H23; Q48; D84

Acknowledgments: We thank Aleksander Azarnov, McLane Daniel, Pedro Bernal Lara, Alejandro Ome, Colleen O'Reilly and Sandya Swamy for excellent research assistance. We are also grateful for helpful comments and suggestions from Soren Anderson, Brian Cadena, Carolyn Fischer, Don Fullerton, David Gerard, David Greene, Jonathan Hughes, Mark Jacobsen, Gary Krueger, Raymond Robertson, Joel Slemrod, and participants at the 2009 Heartland Environmental Economics Workshop and the 2009 National Tax Association Meetings. Funding for this project was provided by the Energy Initiative at the University of Chicago and by a Keck Foundation Grant administered by Macalester College. Any errors or omissions are our own.

The Effect of Gasoline Prices on the Demand for Fuel Economy in Used Vehicles: Empirical Evidence and Policy Implications

Abstract

Economic analysis generally favors fuel taxation as an efficient policy for dealing with climate and environmental externalities related to fuel consumption and emissions in the personal transportation sector, while actual policy heavily leans on fuel economy regulation. One justification offered for the use of regulatory policy instead of direct fuel taxation is the hypothesis that consumers undervalue fuel economy. If consumers do not sufficiently value fuel economy when making the decision about what car to buy, alternatives such as fuel economy regulation or feebates may be welfare improving over a standard Pigouvian tax on fuel.

To determine whether or not consumers properly value fuel economy, we use microdata on used vehicle prices and a unique identification strategy based on micro-level variation in vehicle odometers to test whether used car prices change by the amount predicted by a fully rational asset pricing model. Our baseline results indicate that used car prices adjust by nearly 80% of the amount predicted by theory, and sensitivity analysis suggests that for a reasonable discount rate, they adjust by the full amount. These results contrast with recent literature and policy discussions that argue for tighter fuel economy standards on the basis of consumer undervaluation. Our results instead confirm the canonical finding that direct taxation of fuel is superior to fuel economy regulation or feebates.

1. Introduction

Fossil fuel consumption and greenhouse gas emissions from the personal transportation sector pose serious challenges to today's policy-makers. In terms of greenhouse gases, a tax on gasoline is as efficient as a tax on emissions, since the amount of carbon released by a gallon of gasoline is independent of the manner in which it is combusted.[1] As a result, the standard economic model suggests that a tax on gasoline is sufficient to fully correct for any climate-change externalities related to driving.[2] This conclusion rests on an economic model in which consumers make fully informed, rational decisions regarding both how many miles to drive and what vehicle to buy. There are reasons to think that these conditions may not hold. In particular, there is evidence that consumers may not sufficiently value fuel economy when making the decision about what car to buy.

Consumers may undervalue fuel economy because they do not fully and correctly calculate the present discounted value of the flow of future operating costs, which depend on driving behavior, fuel economy and the price of gasoline. This implies that fuel economy will be underprovided in the absence of a corrective policy. It also implies that consumers will under-react to a carbon pricing system that raises the value of fuel economy by increasing the price of gasoline. As such, an additional complementary policy, such as a Corporate Average Fuel Economy (CAFE) standard, or a feebate, which subsidizes the purchase of fuel efficient and taxes the purchase of inefficient vehicles, can be welfare improving, even when carbon emissions have been priced optimally via a gasoline tax, carbon tax, or cap-and-trade program

[1] See Parry, Walls, and Harrington (2007) for a thorough review of the relative efficiency of policies for the reduction of vehicle pollution.

[2] Since a gallon of gasoline corresponds to differing numbers of vehicle miles traveled, according to fuel economy, a gasoline tax is an imprecise instrument for addressing miles externalities, such as congestion.

(Fischer, Harrington, and Parry (2007), Greene, Patterson, Singh, and Li (2005), and Train, Davis, and Levine (1997)).

Despite the important policy implications of undervaluation of fuel economy, the literature on this topic is sparse. Existing studies that examine the demand for fuel economy rely on incomplete identification strategies, make unreasonably stringent restrictions on the way in which vehicle prices vary by vehicle type and age, or do not directly estimate the parameters necessary to design optimal complementary policies. In this paper, we use a unique identification strategy and 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. This is conceptually similar to a difference-in-difference approach. Importantly, the fact that the difference is *within* vehicle type allows us to provide an exceptionally rich set of controls, including time-period shocks and depreciation schedules that are unique for each vehicle type. 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.

Past research has provided several reasons to suspect that consumers may undervalue fuel economy. First, experimental evidence suggests that consumers are confused by the nonlinearity of cost savings in miles per gallon rated fuel economy (Larrick and Soll, 2008).[3] Second,

[3] Many, for example, do not realize that an increase in fuel economy from 15 to 16 miles per gallon implies a greater cost savings than an increase from 30 to 31 miles per gallon.

qualitative evidence from surveys of new car buyers suggests that most consumers are unable to articulate the key building blocks required for a present discounted value analysis, including typical mileage, fuel economy ratings of their current vehicles, and a discount rate (Turrentine and Kurani, 2007). Third, our own analysis of data from the Michigan Survey of Consumers produces several results that are difficult to explain in a model of infallible consumers. The survey asks a sample of households each month whether or not it is a good time to buy a vehicle and why. Even when controlling for the real price of gasoline, we find that consumers are more likely to cite gasoline prices and fuel economy as important when nominal gasoline prices are high and when prices have just changed (Sallee and West, 2008). They also appear to respond more strongly to gasoline price increases than decreases, which is consistent with findings in Kilian and Sims (2006).[4] These findings suggest that consumers may value fuel economy improperly, but they do not necessarily imply undervaluation as opposed to overvaluation, nor do they provide well-identified estimates of its degree.[5]

There is also a literature that, like us, attempts to directly test whether or not car prices fully adjust. Kahn (1986) develops a panel estimation procedure to test a model of asset valuation in the used car market, which is extended by Kilian and Sims (2006).[6] Allcott and Wozny (2010)

---

[4] Note that there may be justifications for some of these features depending on how consumers forecast gasoline prices.

[5] There is also an older literature that relies on cross-sectional variation for identification, which did not reach a consensus. The literature on energy intensive durables has found very large implied discount rates, which may be a symptom of myopia (Hausman, 1979, Dubin and McFadden, 1984), while Dreyfus and Viscusi's (1995) results imply discount rates that are roughly the same as those on used vehicle loans, suggesting that consumers in vehicle markets are not short-sighted.

[6] A recent body of literature also uses panel methods for identification, but does not estimate undervaluation parameters. Linn and Klier (2007) examines the effect of changes in national-level gas prices on the demand for fuel efficiency, but consider effects on quantities (as opposed to prices) of new cars. They find that a $1 increase in the price of gasoline results in a 0.5 miles per gallon increase in fuel economy, on average. This estimate fits neatly into the range of estimates in Li, Timmins, and von Haefen (2009), which examines the effects of changes in gasoline price on registrations of both used and new vehicles. Busse, Knittel, and Zettelmeyer (2009) examines the response

takes a different approach to arrive at the same location, identifying the effect of gasoline prices on the demand for fuel economy using a discrete choice framework that exploits consumer substitution between new and used vehicles and of varying fuel efficiencies. Like ours, these models develop an econometric framework in which there should be a one-to-one relationship between expected remaining operating costs and the price of a vehicle. If the relationship is less than one-to-one, then consumers undervalue fuel economy.

Our approach differs critically from this existing literature in the assumptions necessary to achieve identification. Kahn (1986) and Kilian and Sims (2006) implicitly assume that unobservable vehicle characteristics are constant within model across model year. That is, a two-year old 2007 Toyota Camry is assumed to have the same unobservable characteristics as a two-year old 2006 Toyota Camry. Allcott and Wozny (2010) relax that assumption but assume that changes in vehicle characteristics across model years are uncorrelated with gasoline price changes. This may be incorrect if automakers make vehicle modifications between model years, which is suggested in both Linn and Klier (2007) and Sallee and Slemrod (2009).

Kahn (1986) and Kilian and Sims (2006) also assume that all vehicles are driven the same number of miles each period and last the same number of years before falling apart completely. Allcott and Wozny (2010) allow miles driven per period to vary, but only by age and class. While their survival probabilities are specific to age, model year, class, and manufacturer, they do not vary with gas prices or with odometer reading. All three papers use aggregate vehicle prices rather than transaction-level prices.[7]

---

of new car transactions to gasoline prices on several margins, including how average prices change for vehicles in different quartiles of the fuel economy distribution.

[7] Allcott and Wozny aggregate transaction-level prices, while the other two papers use aggregates published in a leading used-car price guide.

Our reliance on micro-level variation allows us to relax these assumptions. We expect for lifetime, mileage-age relationships, and scrappage rates to vary across vehicles and with gasoline prices, and for such variation to affect the estimated effect of changes in operating costs on used vehicle prices. For example, assuming constant mileage across vehicle ages ignores the fact that actual mileage declines with vehicle age. This implies that the calculated remaining lifetime operating costs are too low for all vehicles, and especially so for vehicles of lower fuel efficiency. Assuming that scrappage rates do not depend on odometer readings creates bias by underestimating the remaining mileage of older vehicles. And we expect for different vehicles to experience differing time-specific price changes.

Our baseline results indicate that used car prices adjust by nearly 80% of the amount predicted by theory, and sensitivity analysis suggests that for a reasonable discount rate, they adjust by the full amount. We interpret the results of this exercise as a test of whether or not consumers fully recognize the value of fuel economy, though we acknowledge that with our current specification, we can only jointly test assumptions about the discount rate, gasoline price expectations, and rationality. Our findings therefore suggest a limited role for complementary policy tools. These results contrast with Kahn (1986) and Kilian and Sims (2006), who report much lower adjustment estimates in most specifications.[8] Allcott and Wozny (2010) report a main estimate of 61% adjustment, which is sufficiently below our estimate to lead them to the conclusion that consumers undervalue fuel economy. We also find statistically significant differences in adjustment percentages across vehicle types and makes. For example, prices of pickups, SUVs, and minivans experience less adjustment than cars, while Toyotas experience more adjustment than other makes.

[8] Kahn (1986) concludes that the data do not reject full adjustment, but only after reporting a wide range of specifications that do reject it and finally settling on an instrumental variables specification that may be subject to bias.

The next section outlines our estimation strategy. Section 3 describes the data, while Section 4 explains the calculation of the expected remaining operating cost variable and the assumptions involved in this calculation. Section 5 explains how we implement our estimation strategy, presents results, and outlines steps for future research. Section 6 discusses the policy implications of our findings, and Section 7 concludes.

2. Estimation Strategy

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. For example, suppose that two 2005 Toyota Camrys are sold in November 2008, 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 2008 and December 2008, and that in December two other 2005 Toyota Camrys, one with 80,000 miles and the other with 90,000 miles, are sold. 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 differenced out.

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 an odometer polynomial, time period fixed effects and a measure of the discounted expected future operating costs of vehicles. Importantly, 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, including model, model year, cylinders, displacement, transmission and trim. 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. This use of individual data distinguishes our approach from related work and enables us to relax key assumptions.

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 and the supply of used cars is inelastic. Under these assumptions, the expected discounted price, P, of an individual observed vehicle i of type j at time t is equal to the expected discounted driving value $V(\cdot)$ over its remaining lifetime, minus the expected discounted values of fuel costs $C(\cdot)$, and maintenance costs $Z(\cdot)$:

$$P_{ijt} = V\left(O_{ijt}, X_j, r\right) - C\left(O_{ijt}, m_{ijt}, g_t, MPG_j, r\right) - Z\left(O_{ijt}, X_j, r\right) \tag{1}$$

where O is odometer reading, X 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:

$$C_{ijt} = E_t\left[\sum_{s=t}^{R} H(O_{ijt}, X_j)\frac{1}{(1+r)^{s-t}}\frac{m_{js}\, g_s}{MPG_j}\right] \tag{2}$$

where R is the final period in which 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 below in Section 4.

Since the functions $V(\cdot)$ and $Z(\cdot)$ do not depend on the price of gasoline $\left(\frac{\partial P_{ijt}}{\partial g_t} = -\frac{\partial C_{ijt}}{\partial g_t}\right)$, we could use a panel estimation strategy, where the variation in gasoline price interacts with fuel economy to enable identification at the vehicle level. But this requires that all vehicles experience the same time period shock and the same depreciation schedule. Rather than impose these rather stringent assumptions, instead we identify the effects of gasoline price changes on vehicle price using variation in odometer readings at the vehicle level, using the following estimation equation:

$$P_{ijt} = \beta C_{ijt} + \sum_{a=1}^{4} \alpha_a O_{ijt}^{a} + \delta_{jt} + \varepsilon_{ijt} \tag{3}$$

where the odometer terms O and their coefficients are specific to each vehicle type j, and the δ_{jt} are time-vehicle type fixed effects, where time is measured in months. The odometer polynomial is an approximation to the V and Z functions. Our null hypothesis is that β equals negative one, that there is a one-for-one relationship between the present discounted value of real remaining fuel costs and real vehicle price. This implies full adjustment.

3. Data Sources

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.[9] This market does not include individual end users. In addition to automotive dealers, businesses and

[9] One source of error may exist in the odometer readings. 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. Currently, we answer any concerns related to this by pointing out that our estimates are similar when we look only at vehicles known to have six-digit odometers, like Toyotas and Hondas.

governments that own large fleets sell vehicles on these auctions. The buyers are primarily used car dealers, who subsequently resell the vehicles to consumers. Many of the transactions in these markets are between dealers, who use these auctions as a way to reshuffle their inventories. Wholesale prices have a major advantage over retail prices because retail prices reflect idiosyncratic factors like bargaining ability and credit ratings that, from our standpoint, merely introduce noise.

Our data sample includes over 90 million transactions that took place between 1990 and 2009. 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 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, due to computational limitations, we take a 10% sample of the vehicles with complete information, leaving us with around 8 million transactions in the estimation sample.

For the price of gasoline, we use the tax inclusive, monthly national retail price of gasoline from the Energy Information Administration (EIA). We adjust these and the vehicle prices using the CPI-U price series.

Table 1 presents summary statistics for our estimation sample.

4. Details of Cost Estimation and Implicit Assumptions

The most important, and most challenging, aspect of our estimation procedure is the proper construction of the future fuel cost function. Throughout, we emphasize the interpretation of our procedure as a test of the complete valuation of fuel economy by consumers, but it is important to acknowledge that the null hypothesis of our test relies on a number of assumptions, and the test is therefore a joint test of all of these factors. In particular, the cost measure requires an explication of the discount rate used by consumers and their expectations regarding the future path of gasoline prices. In this section, we describe in turn each of the assumptions necessary for the construction of a cost estimate and the sources of any auxiliary data used.

4.A Mileage and Scrappage

The cost function, 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 probability 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.[10] 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, so we use our auction data to introduce this heterogeneity 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

[10] Lu (2006) does not directly report average miles driven as a function of age, but instead reports coefficient estimates from a regression of odometer readings on a cubic in age. Lu's scrappage probabilities are calculated using coefficients from estimating two double exponential functions of age pieced together at 10 years for cars and 12 years for light trucks.

the regression coefficients reported by Lu (2006).[11] We then regress this predicted odometer reading on the actual odometer reading separately for each make as follows:

$$d_{ca} = \theta_{cm} O_{icm} + \varepsilon_{icm} \qquad (4)$$

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 theta for each make, which we report in Table 2. We use this theta 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 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. We invert Lu's 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 theta for each make to allow for heterogeneity, but we currently impose the same scrappage function on all vehicles.[12]

In constructing the cost measure, we assume that each vehicle lasts no more than 25 remaining years 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.

---

[11] To determine the age of vehicles in our price data, we assume that unless they appear in our data earlier than May of their model year, all vehicles were sold new in May of their model year, which we empirically determined to be the midpoint of the model year in recent years in a sample of new vehicle transactions. We assume that vehicles that appear in our data earlier than May of their model year were sold new in January of their model year.

[12] Li et al. (2009) use registration data to estimate scrappage functions, including responses to the price of gasoline. They have generously provided us with make-specific scrappage regressions, which will enable us to improve our current procedure. We also intend to use their estimates to adjust the scrappage probability for the level of gasoline prices.

Past research has simply assumed that all vehicles are driven the same number of miles in all periods, and that they eventually are scrapped at a common terminal mileage (Kahn (1986), Kilian and Sims (2006) and Allcott and Wozny (2010)). This is clearly not true, as it is well known that newer cars are driven more per year and that the average lifetime of cars and trucks varies significantly. Our procedure relaxes these assumptions and replaces them with a better approximation of actual driving behavior.

4.B Gasoline Price Expectations

Another key term in the cost function is the price of gasoline in future periods. To construct the cost measure, a value for the price of gasoline in all future periods must be specified, requiring that we impose an expectations formation process. We assume that all consumers use a random walk assumption, so that the future price of gasoline in all periods is equal to the current price. Support for this assumption comes from Anderson, Kellogg and Sallee (2010), which concludes that consumer reports of their expectations about future gasoline prices appear to follow a random walk in real terms.

We have also estimated ARIMA models of the gasoline price series, and we will use them to test whether or not our answer changes qualitatively when we use various expectations formation processes that conform to historical data. Due to computational time, we do not yet have these results to report. If consumers form expectations in some way that is inconsistent with observed data, we will be unable to distinguish this from an undervaluation. On one level, this is in line with the overall intention of our project because, if consumers are using a forecast that is out of line with observed data, then we may wish to consider this is a consumer failure in line with the problems of fuel economy undervaluation that might arise from other biases.

4.C Fuel Economy

Several different fuel economy measures are available. In particular, the EPA reports both a city and a highway MPG, along with a combined rating which is a harmonic average of the two. 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, a significant revision to the test procedures was initiated with the hopes of improving accuracy further. Our work here uses 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.

4.D The Discount Rate

The cost function also includes a discount rate. For our baseline estimates, we use a discount rate of 5%, and we conduct sensitivity analysis using rates of 10% and 15%. Just as with the gasoline expectations process, we are not able to separately identify the degree of consumer valuation and the discount rate from our data. Instead, we describe the range of estimates that follow from different discount rates. If vehicle price movements match our model expectations only at very high discount rates, we interpret this as evidence of incomplete adjustment.

If the appropriate discount rate reflects the opportunity cost of borrowing money in the car market, then interest rates on automobile loans may provide an objective measure. According to the Federal Reserve, the interest rate on 48-month new car loans originating at commercial banks has varied between 6 and 10% over the last decade, which would make our 5% baseline conservative.

4.E 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 (Kilian and Sims 2006) or attempted to account for the new car market influence structurally (Allcott and Wozny 2010). 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.

4.F Heterogeneity and Sorting

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 does not change with the price of gasoline.

4.G Heterogeneity and Interpretation

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.

This may be less of a problem if market forces cause this heterogeneity to collapse in its effect on market prices. We have not yet been able to develop a model that would tell us whether our procedure will reveal the average consumer adjustment, the marginal consumer adjustment, or some less coherent aggregation. In this regard, discrete choice approaches, such as the mixed logit approach of Sawhill (2008) offer some advantages over our own.

4.H Risk Aversion and Variability

Currently, we assume a risk neutral agent, who cares only about the expected future costs, not about other moments of the cost function. It might be the case that consumers care also about the variance in future costs because they are risk averse. In this case, a fuel efficient car is valuable as insurance, since it shrinks the dollar impact of a given gasoline price dispersion. Recent volatility in the price of gasoline may therefore influence the value of fuel economy. Again, our identification strategy limits this concern. Any resulting price increase in demand for a particular car type will be soaked up in our fixed effects.

5. Estimation

5.A Implementation

Estimation of equation (3) requires vehicle-specific odometer terms, which enter estimation as odometer terms interacted with vehicle indicators. As our unit of vehicle observation is at the level of the truncated VIN, running the estimation on even a subsample of our data hits computational memory constraints, as the number of independent variables is enormous.

We therefore invoke the Frisch-Waugh-Lovell theorem (Frisch and Waugh 1933, Lovell 1963), which allows us to break down the estimation procedure into steps that are less memory-intensive (albeit still very computation-time-intensive). To implement this procedure, first, for each VIN separately, we regress vehicle price on the odometer polynomial and time period (month) fixed effects and collect the residuals. We then regress, for each VIN separately, our cost variable on the odometer polynomial and time period fixed effects, and collect those residuals. Finally, we regress the residuals from the price regression on the residuals from the cost regression, clustering the standard errors on VIN and correcting for the degrees of freedom used in the prior stage. This yields a coefficient on the residuals from the cost equation that is numerically identical to the β in equation (3) and a standard error that is identical to the standard error of β.

5.B Baseline Results

Table 3 reports the results for the full sample and a selection of subsamples. For the full sample, the coefficient on cost is -0.788, implying that for every one dollar increase in fuel costs, used vehicle price falls by about 79 cents. While this estimate suggests that consumers' valuation of fuel economy is incomplete, it is substantially closer to negative one than the results of other recent studies. For example, over a broad spectrum of discount rates, Kilian and Sims (2006) obtain estimates in the range of -0.10 to -0.25. Assuming a 9 percent discount rate, Allcott and Wozny (2010) put this number at about -0.61.

Estimation on a set of subsamples reveals some intriguing heterogeneity in price adjustment across vehicle classes and makes, with vehicle price adjustment rates ranging from just over 50 percent for Ford cars to over 90 percent for Toyota cars. Overall, car prices adjust more fully than do light-duty truck prices, though the pattern is inconsistent within makes.[13]

5.C Sensitivity of Results with Respect to Changes in the Discount Rate

We also run the estimation on the full sample of vehicles imposing discount rates of 10% and 15%. Table 4 shows that if the discount rate is 10%, our coefficient of interest equals -1.01; consumers fully adjust to changes in the present discounted value of remaining operating costs.[14] Estimation imposing a discount rate of 15% implies overadjustment, underscoring the importance of careful construction of the cost term, and the determination of what constitutes a "reasonable" discount rate.

5.D Testing for the Presence of Asymmetric Responses to Changes in Gas Prices

Previous literature indicates that consumers may respond asymmetrically to gasoline price shocks (Hamilton (1996, 2003), Kahn (1986), Kilian and Sims (2006)). We therefore estimate a specification that includes the cost term and an interaction of the cost term with an indicator equal to one when monthly gas prices have increased. We report these results in Table 5. In contrast with previous findings, we find little evidence for such asymmetry, at least at the month-to-month level, as the estimated coefficient on the interaction term is very small.

[13] Sawhill (2008) finds substantial heterogeneity across consumers (where the unit of observation is actually a model year market share) but does not explore the dimensions along which this heterogeneity occurs.

[14] To ease comparison of our results with those in Allcott and Wozny (2010), we also ran the estimation assuming, like them, a 9 percent discount rate. Using data from our full sample period (1990-2009), with a 9 percent discount rate our coefficient on cost equals -0.963. Cutting our sample to match their sample's time period (1999-2008), with a 9 percent discount rate our coefficient on cost equals -0.934.

5.E Planned Future Work

In ongoing work, we are pursuing auxiliary data to improve the accuracy of our mileage estimates. We are also working to incorporate estimates of how mileage and scrappage schedules adjust to gasoline prices.

Finally, we also have a large number of robustness checks in process. For example, preliminary results indicate that our coefficient of interest is somewhat sensitive to the functional form of the odometer polynomial.[15] Given the very long computation time required for each iteration, we do not have these results to report or characterize yet. As such, our current estimates should be considered preliminary.

5.D Relationship to Previous Literature

Our results contrast with some recent work that attempts to estimate the same parameter in the car market. Kilian and Sims (2006) find significant underadjustment, with principle estimates hovering around 25 cents on the dollar. Allcott and Wozny (2010) report a main estimate of 61 cents on the dollar, which is closer to our findings but still far enough away from full adjustment that they consider this evidence of undervaluation. Both of these papers use a panel identification strategy that requires unobservable vehicle attributes to be common across model years. That is, a two-year old 2007 Toyota Camry is assumed to have the same unobservable characteristics as a two-year old 2006 Toyota Camry. This may be incorrect of automakers make vehicle

[15] As we were writing this draft, we obtained results suggesting that an order 5 polynomial (rather than the quartic we assume for this draft's results) may provide more appropriate controls in estimation of equation (3). Assuming a 5 percent discount rate, using an order 5 polynomial results in a coefficient on cost equal to -0.679, which is a bit closer to zero the main result reported here (-0.788). Estimates assuming 10 or 15 percent discount rates would similarly scale toward zero the estimates presented in this draft. Our overall conclusions remain unchanged—preliminary results suggest that consumers adjust rather fully to changes in gasoline prices. Estimation using polynomials of orders higher than 5 does not result in any further meaningful changes in the cost coefficient.

modifications between model years in response to fuel economy regulation or gasoline prices, as is suggested in both Linn and Klier (2007) and Sallee and Slemrod (2009). Further, since Allcott and Wozny (2010) use gasoline prices as instruments for the quantity sold, year-to-year modifications in vehicles that are related to gasoline prices may violate their exclusion restriction.

Both Kilian and Sims (2006) and Allcott and Wozny (2010) also make restrictions across vehicles regarding the depreciation schedules for vehicles and the time period effects. Importantly, when we restrict our estimator to impose the same odometer polynomial coefficients and the same time period fixed effects within a make, our estimates fall dramatically and become closer to zero than one. This suggests that even within a carmaker, the depreciation schedules and time period effects are different enough that imposing such restrictions causes a qualitative change in the analysis. Without a theoretical justification for imposing these restrictions, we think it is preferable to relax them, which we are able to do given our use of odometer variation.

Our results are broadly consistent with the findings of Busse, Knittel and Zettelmeyer (2009), who do not attempt to directly estimate the degree of adjustment, but provide a back of the envelope calculation comparing their estimates to a fully rational benchmark and conclude that price adjustments in the used car market are consistent with full adjustment.

6. Policy Implications

Economic analysis that assumes consumers are rational finds strong support for a gasoline tax over fuel economy regulation as a policy for reducing gasoline consumption (Austin and Dinan (2005) and West and Williams (2005)). If consumers systematically underestimate the value of fuel economy when making a vehicle choice, their mistakes represent an additional market failure that compounds externalities related to driving. In this setting, even an optimally determined carbon price or gasoline tax might be complemented by another policy, such as a feebate, or fuel economy standards.[16] Fischer et al. (2007), for example, presents a range of simulation estimates that show that the existence of myopia can change qualitative conclusions regarding the welfare impact of increasing CAFE; if consumers undervalue fuel economy enough, CAFE becomes welfare improving.

While consumer fallibility is probably a necessary condition for using feebates and CAFE versus relying on only a gasoline or carbon tax, it is not sufficient. Both feebates and CAFE, by improving fuel economy, reduce the cost of driving, and thereby lead to increases in vehicle miles driven. The negative externalities associated with driving, including congestion and accident costs, are substantial. In addition, since vehicle miles of travel and leisure are relative complements, the increase in miles traveled that occurs because of the reduction in the price per mile can exacerbate preexisting labor market distortions, which would further reduce the welfare improvements from correcting for myopia (West and Williams, 2005). In order to determine the net effect on welfare, one must include these considerations in a model that also accounts for consumer fallibility. For example, Fischer et al. (2007) conclude that undervaluation must be severe to justify tightening current fuel economy standards. Thus, while our estimates do not always fail to reject the null hypothesis of full valuation, our estimates do consistently indicate that consumers significantly value fuel economy, which makes the case for augmenting fuel taxation with regulation quite weak.

[16] Train, Davis, and Levine (1997) and Greene, Patterson, Singh, and Li (2005) suggest that feebates might be justified by consumer myopia, but do not conduct formal analysis linking the two.

7. Conclusion

We use vehicle transaction data combined with EPA fuel economy ratings and gasoline prices to estimate the extent to which used vehicle prices adjust to changes in the present discounted value of operating costs. We identify the effect of gasoline prices not with time- and fuel-economy-based variation in operating costs across different vehicles, but instead with *odometer*-based variation in expected operating costs within tightly-defined vehicle types. That is, we compare the effect of a gasoline price shock on the price of a vehicle with many miles remaining to the price of an identical vehicle with few miles remaining.

We interpret the results of this exercise as a test of whether or not consumers fully recognize the value of fuel economy. With our current specification, we can only jointly test assumptions about the discount rate, gasoline price expectations, and rationality. That is, incomplete consumer valuation of fuel economy could be due to consumers employing a discount rate that is too high or to consumers underestimating the lifetime of a vehicle, having faulty gasoline price expectations, or being boundedly rational. Disentangling the effects of these potential sources of undervaluation remains a fruitful area for future research.

Preliminary baseline results, which assume a discount rate of 5%, indicate that for a dollar increase in operating costs, used vehicle prices fall by 79 cents. While this estimate is consistent with the general finding of underadjustment, it is closer to full adjustment than found by papers in the recent literature. And, our estimate assuming a 10% discount rate implies that adjustment is full. While our work is still progressing, these results, which contrast with some other recent findings, are important because they suggest that policies other than a gasoline tax are unlikely to be justified on the grounds of incomplete consumer valuation of fuel economy.

References

Allcott, Hunt, and Nathan Wozny, "Gasoline Prices and the Fuel Economy Discount Puzzle," March 29, 2010, Manuscript, MIT.

Anderson, Soren T., Ryan Kellogg and James M. Sallee, "What Do Consumers (Think They) Know About the Price of Gasoline?" May 26, 2010, Unpublished.

Austin, David and Terry Dinan, "Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Increased Gasoline Taxes," *Journal of Environmental Economics and Management*, November 2005, 50(3):562-582.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, 63 (4): 841-890.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer, "Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing," February 2009, Manuscript, University of California at Davis.

Dreyfus, Mark K. and W. Kip Viscusi, "Rates of Time Preference and Consumer Valuations of Automobile Safety and Fuel Efficiency," *Journal of Law and Economics*, 1995, 38: 79-105.

Dubin, Jeffrey A. and Daniel L. McFadden, "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption," *Econometrica*, March 1984, 52(2): 345-62.

Fischer, Carolyn, Winston Harrington, and Ian Parry, "Do Market Failures Justify Tightening Corporate Average Fuel Economy (CAFE) Standards?" *The Energy Journal*, 2007, 28 (4): 1-30.

Frisch, Ragnar and Frederick V. Waugh. Partial Time Regressions as Compared with Individual Trends" *Econometrica*, 1933, 1 (40): 387-401.

Goh, Vincent, Paul S. Fischbeck, and David Gerard, "Identifying and Correcting Errors with Odometer Readings from Inspection and Maintenance Data: Rollover Problem for Estimation of Emissions and Technical Change," *Transportation Research Record*, 2007, No. 2011: 87-97.

Greene, David L., Philip D. Patterson, Margaret Singh, and Jia Li, "Feebates, Rebates and Gas-Guzzler Taxes: A Study of Incentives for Increased Fuel Economy," *Energy Policy*, April 2005, 33 (6): 757-775.

Hamilton, James D. "This is What Happened to the Oil-Macroeconomy Relationship," *Journal of Monetary Economics*, 1996, 38(2): 215-220.

Hamilton, James D. "What is an Oil Shock?" *Journal of Econometrics*, 2003, 113: 363-398.

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

Kahn, James A., “Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach,” *Quarterly Journal of Economics*, May 1986, 101 (2): 323-340.

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

Larrick, Richard P. and Jack B. Soll, “The MPG Illusion," *Science*, June 2008, 320 (5883): 1593-4.

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

Linn, Joshua and Thomas Klier, “Gasoline Prices and the Demand for New Vehicles: Evidence from Monthly Sales Data," 2007, Manuscript, University of Illinois at Chicago.

Lovell, M., Seasonal adjustment of economic time series, *Journal of the American Statistical Association*, 1963, 58: 993-1010.

Lu, S., “Vehicle Survivability and Travel Mileage Schedules,” Technical Report DOT HS 809 952, NHTSA’s National Center for Statistics and Analysis, January 2006.

Parry, Ian W. H., Margaret Walls and Winston Harrington, “Automobile Externalities and Policies, “*Journal of Economic Literature* XLV: 2007, 374-400

Sallee, James, and Joel Slemrod. “Car Notches,” November 2009, Paper prepared for the National Tax Association Meetings.

Sallee, James, and Sarah E. West. “The Effect of Gasoline Prices on the Demand for Fuel Economy,” November 2008, Paper prepared for the Association for Public Policy Analysis and Management Fall Conference.

Sawhill, James. “Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market,” April 2008, Manuscript, University of California at Berkeley.

Train, Kenneth E., William B. Davis, and Mark D. Levine, “Fees and Rebates on New Vehicles: Impacts on Fuel Efficiency, Carbon Dioxide Emissions, and Consumer Surplus,” *Transportation Research Part E: Logistics and Transportation Review*, March 1997, 33 (1): 1-13.

Turrentine, Thomas S. and Kenneth S. Kurani, “Car Buyers and Fuel Economy?" *Energy Policy*, February 2007, 35 (2): 1213-1223.

West, Sarah E. and Roberton C. Williams III, “The Costs of Reducing Gasoline Consumption.” *American Economic Review Papers and Proceedings*. May 2005, 99 (2): 294-299.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.
Nominal Vehicle Sale Price	9260.32	6529.70
Real Vehicle Sale Price ($1982-84)	5209.61	3657.68
Present Value Real Remaining Fuel Cost ($1982-84)	3076.19	1573.26
Nominal Gas Price ($)	1.78	0.75
Real Gas Price ($ 1982-84)	0.96	0.29
Odometer (miles)	57274	43726
Age (years)	3.88	3.31
Fuel Economy (EPA Combined MPG)	22.23	4.62
% Car	0.64	.
% Light-duty truck	0.36	.
Number of observations	8106030	

Table 2: Estimates of θ (By Make Vehicle Mileage, Relative to the Mean)

Make	Car Theta	Light-Duty Truck Theta
ACURA	0.892	0.920
ALFAROMEO	1.504	.
AMC	1.148	.
AUDI	0.963	0.677
BMW	0.950	1.000
BUICK	0.933	0.840
CADILLAC	0.990	0.850
CHEVROLET	0.838	0.874
CHRYSLER	0.895	0.886
DAEWOO	0.919	.
DAIHATSU	0.896	1.045
DODGE	0.801	0.902
EAGLE	0.900	.
FORD	0.869	0.916
GEO	0.879	1.138
GMC	.	0.892
HONDA	0.862	0.902
HUMMER	.	0.827
HYUNDAI	0.837	0.817
INFINITI	0.924	0.976
ISUZU	0.954	0.985
JAGUAR	1.123	.
JEEP	.	0.968
KIA	0.789	0.957
LANDROVER	.	1.087
LEXUS	0.916	0.963
LINCOLN	0.907	0.894
MAZDA	0.885	0.945
MERCEDESBENZ	0.989	0.992
MERCURY	0.936	0.949
MERKUR	1.112	.
MINI	0.968	.
MITSUBISHI	0.826	0.865
NISSAN	0.862	0.948
OLDSMOBILE	0.912	0.955

Make (cont.)	Car Theta	Light-Duty Truck Theta
PEUGEOT	0.981	.
PLYMOUTH	0.882	0.960
PONTIAC	0.837	0.875
PORSCHE	1.355	1.005
RENAULT	1.142	.
SAAB	0.959	1.067
SATURN	0.895	0.834
SCION	0.679	.
SUBARU	0.870	0.863
SUZUKI	0.759	1.030
TOYOTA	0.855	0.897
VOLKSWAGEN	0.936	1.067
VOLVO	0.906	0.955

Table 3: Baseline Estimation Results
Coefficients on the Present Value of Real Remaining Fuel Costs

	Full Sample	
	-0.788	
	(0.013)	
	Subsamples	
	Passenger Cars	Light Trucks
All	-0.851	-0.736
	(0.020)	(0.016)
Chevrolet	-0.599	-0.727
	(0.049)	(0.037)
Ford	-0.513	-0.735
	(0.032)	(0.037)
Honda	-0.787	-0.619
	(0.042)	(0.056)
Toyota	-0.914	-0.792
	(0.060)	(0.052)

Standard errors are reported in parentheses.
Estimates obtained assuming a 5 percent discount rate.

Table 4: Sensitivity of Results with Respect to the Discount Rate
(All Vehicles)

Discount Rate	5%	10%	15%
Cost Coefficient	-0.788	-1.01	-1.24
Standard Error	(0.013)	(.016)	(.020)

Table 5: Testing for Asymmetric Responses to Gas Price Changes
(All Vehicles)

Cost Coefficient	-0.792
Standard Error	(0.013)
Coefficient on (Cost*Gas Price Increase)	0.008
Standard Error	(0.002)

Estimates obtained assuming a 5% discount rate.