Pain at the Pump: The Effect of Gasoline Prices on New and Used Automobile Markets*

Meghan R. Busse
Northwestern University and NBER

Christopher R. Knittel
University of California, Davis and NBER

Florian Zettelmeyer
Northwestern University and NBER

July 2010
First version: October 2008

*We are grateful for helpful comments from Severin Borenstein, Tim Bresnahan, Igal Hendel, Ryan Kellogg, Aviv Nevo, Jorge Silva-Risso, Scott Stern and seminar participants at the ASSA, Iowa State, Ohio State, Triangle Resource and Environmental Economics seminar, Purdue, UC Berkeley, University of Michigan, Yale, MIT, University of British Columbia, University of Toronto, Cornell, Northwestern University, Chicago Federal Reserve Bank, University of Chicago, Texas A&M, and the University of Rochester. We also thank participants at the NBER 2008 Winter IO conference. We thank the University of California Energy Institute (UCEI) for financial help in acquiring data. Busse and Zettelmeyer gratefully acknowledge the support of NSF grants SES-0550508 and SES-0550911. Knittel thanks the Institute of Transportation Studies at UC Davis for support. Addresses for correspondence: E-mail: m-busse@kellogg.northwestern.edu, crknittel@ucdavis.edu, f-zettelmeyer@kellogg.northwestern.edu

Pain at the Pump: The Effect of Gasoline Prices on New and Used Automobile Markets

Abstract

We estimate the effect of gasoline prices on short-run equilibrium prices and sales of new and used cars of different fuel economies. We find that gasoline prices have larger effects on the prices of used cars than of new cars, but that they have large effects on market shares and sales of new cars. We use our findings to estimate a component of the effect of an environmental tax on gasoline, and to investigate whether consumers are myopic about future gasoline costs when they buy cars.

1 Introduction

According to EPA estimates, gasoline combustion by passenger cars and light-duty trucks is the source about fifteen percent of U.S. greenhouse gas emissions, "the largest share of any end-use economic sector."[1] As public concerns about climate change grow, policymakers are trying to figure out what instruments they can use to reduce carbon emissions from this source. Since carbon emissions are essentially proportional to the amount of gasoline used, changing emissions means either changing how much people drive or changing the cars—specifically the fuel economy of the cars—they drive. There are various instruments that might be employed to influence this. While Corporate Average Fuel Efficiency (CAFE) standards arguably have been the major policy instrument used so far to influence U.S. fuel consumption (Goldberg (1998), Jacobsen (2010)), some economists contend that changing the incentives to use gasoline—by increasing its price—would be a more efficient approach since it is the combustion of gasoline itself that is the source of the negative externality.

In considering such a policy, it is useful to know whether automobile buyers are responsive to the price of gasoline, and if so, how and how much they respond. In this paper, we investigate how the price of gasoline affects market outcomes in both new and used car markets. Specifically, we use data on individual transactions for new and used cars to estimate the effect of gasoline prices on equilibrium market shares and sales and equilibrium transaction prices for new and used cars of different fuel economies.

The results we find are of interest for three reasons. First, by estimating the effect of gasoline prices on market shares of new cars we learn about how much the fuel economy of the U.S. automobile fleet changes in response to gasoline prices. Today's car sales are cars that will continue to be driven for years to come. If high current gasoline prices lead to increased sales of high-MPG cars and lower sales of low-MPG cars, that will have a lasting effect on gasoline consumption over the life of those cars.

Second, by estimating the effect of gasoline prices on car prices we learn about how consumers trade off the up-front capital cost of a car and the ongoing usage cost of the car. Understanding this is useful for predicting how consumers might respond to a policy instrument that uses price incentives to encourage energy conservation. The more myopic customers are—meaning the less responsive they are to gasoline prices in the automobile context, or to other energy costs in other contexts—the less effective such policies are likely to be.

Third, by comparing our results for new car and used car markets (which turn out to be very

[1] EPA, Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2006, p. 3-8.

different), we learn something about the importance of the industrial organization of markets in determining the effect that price changes, or policy interventions that mimic price changes, have on the market outcomes that are the object of the policy.

In our empirical results, we find that changes in gasoline prices are associated with very large changes in relative prices of used cars of different fuel economies—a difference of $2839 in the relative price of the most fuel-efficient and least fuel-efficient quartile of cars. For new cars, the predicted relative price difference is much smaller—a $363 difference between the most and least fuel-efficient quartiles of cars. However, we find a large change in the market shares of new cars when gasoline prices change. A $1 increase in the gasoline price leads to a 20.5% increase in the market share of the most fuel-efficient quartile of cars and a 23.9% decrease in the market share of the least fuel-efficient quartile of cars. Using these estimates, we calculate that a gasoline tax or a carbon tax that increased the price of gasoline by $1 would lead to a one-year decrease in U.S. emissions of carbon dioxide of 2.1 million tons from the effect on new car purchases alone. We also find in our estimates little evidence for consumer myopia in valuing future gasoline costs when purchasing cars.

This paper proceeds as follows. In the next section, we position this paper within the related literature. In Section 3 we describe the data we use for the analysis in this paper. In Section 4 we investigate how gasoline prices affect the fuel economy of new cars coming into the U.S. vehicle fleet. In Section 5 we investigate how gasoline prices affect new car prices, and whether these patterns indicate that consumers are myopic. In Section 6 we compare how gasoline prices affect outcomes in new and used markets differently. Section 7 checks the robustness of our estimated results. Section 8 offers some concluding remarks.

2 Related literature

There is no single, simple answer to the question "How do gasoline prices affect gasoline usage?," and, consequently, no single, omnibus paper that answers the entire question. This is because there are many margins over which drivers, car buyers, and automobile manufacturers can adjust, each of which will ultimately affect gasoline usage. Some of these adjustments can be made quickly; others are much longer run adjustments.

For example, in the very short run, when gasoline prices change, drivers can very quickly begin to alter how much they drive. Donna (2010) and Goldberg (1998) investigate two different measures of driving responses to gasoline prices. Donna investigates how public transportation utilization is affected by gasoline prices, and Goldberg estimates the effect of gasoline prices on vehicle miles travelled.

At the other extreme, in the long run, automobile manufacturers can change the fuel economy of automobiles by changing the underlying characteristics—such as weight, power, and combustion technology—of the cars it sells or by changing fuel technologies to hybrid or electric vehicles. Gramlich (2009) investigates such manufacturer responses by relating year-to-year changes in the MPG of individual car models to gasoline prices.

This paper is most closely related to a set of papers that examine a question that could be classified as addressing an intermediate horizon question, relative to those above. This question is: How do gasoline prices affect the prices or sales of car models of different fuel economies? What this set of papers have in common is that they investigate the effect on some market outcome (prices, sales, market share) arising from the ownership choices drivers make among the set of cars currently available from manufacturers. Within this set of papers there are some papers that study the effect of gasoline prices on car quantities (these papers ultimately inform the fleet fuel economy question) and some that study the effect of gasoline prices on car prices (these papers relate ultimately to the question of consumer myopia).

2.1 Gasoline prices and car quantities

Two noteworthy papers that address the effect of gasoline prices on car quantities are Klier and Linn (forthcoming) and Li, Timmins, and von Haefen (2009). Although the two papers address similar questions, they use different data. Klier and Linn estimate the effect of national average gasoline prices on national sales of new cars by detailed car model. They find that increases in the price of gasoline reduce sales of low-MPG cars relative to high-MPG cars. Li, Timmins, and von Haefen also use data on new car sales, but to this they add data on vehicle registrations, which allows them to estimate the effect of gasoline price on the outflow from, as well as inflow to, the vehicle fleet. They find differential effects for cars of different fuel economies: a gasoline price increase increases the sales of fuel-efficient new cars and the survival probabilities of fuel-efficient used cars, while decreasing the sales of fuel-inefficient new cars and the survival probabilities of fuel-inefficient used cars.

2.2 Gasoline prices and car prices

There are several papers that investigate whether the relationship between car prices and gasoline prices indicates that car buyers are myopic about future usage costs when they make car buying decisions.

Kahn (1986) uses data from the 1970s to relate a used car's price to the discounted value of the expected future fuel costs of that car. He finds that used car prices do adjust to gasoline prices, by about one-third to one-half as much as they would need to in order to fully reflect the change in the gasoline cost. This, he concludes, indicates some degree of myopia. Kilian and Sims (2006) repeat Kahn's exercise, with a longer time series, more granular data, and a number of extensions. They conclude that buyers have asymmetric responses to gasoline price changes, responding nearly completely to gasoline price increases, but very little to gasoline price decreases.

Allcott and Wozny (2010) address this question using pooled data on both new and used cars. They also find that car buyers undervalue fuel costs. According to their estimates, a car's price changes by about 60% of the change in the discounted expected future gasoline costs for the car. These estimates imply less myopia than do those of Kahn (1986), although still not full adjustment.

Sallee, West, and Fan (2009) carry out a similar exercise as the papers above, also relating the price of used cars to a measure of discounted expected future gasoline costs. Their paper differs in that it controls very flexibly for odometer readings. This means that the identifying variation they use is differences between cars of the same make, model, model year, trim, and engine characteristics, but of different odometer readings. They find that car buyers adjust to 80-100% of the change in fuel costs, depending on the discount rate used.

Goldberg (1998) approaches the question of consumer myopia in a completely different way. While it is not the main focus of her paper, she calculates the elasticity of demand for a car with respect to its purchase price and with respect to its fuel cost. After adjusting the terms to be comparable, she finds that the two semi-elasticities are very similar, leading her to conclude that car buyers are not myopic.

2.3 Differences from the previous literature

Our paper differs from the papers described above in three ways. First, our paper uses data on individual new and used car transactions, rather than data from more aggregate sales figures, from registrations, or from surveys. Second, our data allow us to compare the effects of gasoline prices on both prices and quantities of cars, and in both used and new markets, in data from a single data source. Third, we choose to use more flexible functional forms than many (although not all) of the papers above.

2.3.1 Transactions data

As described in more detail in Section 3, we observe individual transactions, and observe a variety of characteristics about each transaction, such as location, timing, detailed car characteristics, and demographic characteristics of buyers. This allows us to use extensive controls in our regressions. For estimating the effect of gasoline price on market share, described in more detail in Section 4, we use a linear probability specification which allows us to incorporate at an individual transaction level the controls we observe, reducing the chances that our results arise from selection issues or aggregation over heterogeneous regions, time periods, or car models. In estimating prices, the controls are particularly important, since price negotiation means that individual cars of the same model can sell for very different prices depending on location, timing, buyer characteristics, and attributes of the individual vehicle. Again, being able to use these controls reduces the chances that our results are afflicted by selection or aggregation bias.

2.3.2 Single data source

Using transactions-based data means that we observe prices and quantities for new and used cars in a single data set. This enables us to do several things. First, we can investigate whether the finding of no myopia by Goldberg (1998) in *new* cars differs from the finding of at least some myopia in *used* cars by Kahn (1986), Kilian and Sims (2006), and Allcott and Wozny (2010) because the effect is actually different in new and used cars, or for some other reason. Second, this enables us to learn something about the industrial organization of new and used car markets, and specifically about how the differences in market structure might lead to differences in how a policy intervention would play out in the two markets.

2.3.3 Flexible specification

In addressing the question of myopia, researchers face a choice. The theoretical object to which customers should be responding is the present discounted value of the expected future gasoline cost for the particular car at hand. Creating this variable means having data on (or making assumptions about) how many miles the owner will drive in the future, the miles per gallon of the particular car, the driver's expectation about future gasoline prices, and the discount rate. Having constructed this variable, a researcher can estimate a single parameter that measures the extent of consumer myopia. The advantage of estimating a structural parameter such as this is that it can be used in policy simulations or counterfactual simulations (as Li, Timmins, and von Haefen (2009), Allcott and Wozny (2010), and Goldberg (1998) do). The disadvantage is that the specific assumptions the researcher has made are "baked into" the data, and thereby into the results. It is not possible to ask how the researcher's conclusions would differ if an assumption were different without having access to the researcher's data and redoing the analysis.

We choose to estimate more flexible, non-structural parameters. These cannot be used in the same way in policy simulations or counterfactuals. We argue that the results are nevertheless inputs into very informative back-of-the-envelope calculations, and furthermore, calculations that a reader of this paper can replicate for him- or herself with alternative assumptions about driving behavior, discount rates, or vehicle life.

3 Data

We combine several types of data for the analysis in this paper. Our main data contain information on automobile transactions from a sample of 20% of all dealerships in the U.S. from September 1, 1999 to June 30, 2008. The data were collected by a major market research firm, and include every new car and used car transaction within the time period that occurred at the dealers in the sample. For each transaction we observe the exact vehicle purchased, the price paid for the car, the dealer's cost of obtaining the car from the manufacturer, information on any vehicle that was traded in, and (census-based) demographic information on the customer. We discuss the variables used in each specification later in the paper.

We supplement these transaction data with data on car models' fuel consumption and data on gasoline prices. The fuel consumption data are from the Environmental Protection Agency (EPA). As the fuel consumption measure for each car model we use the "EPA Combined Fuel Economy" which is a weighted geometric average of the EPA Highway (45%) and City (55%) Vehicle Mileage. As shown in Figure 1, the average MPG of models available for sale in the United States shows a pattern of slow decline in the first part of our sample period, and some increase in the latter part.[2] Overall, however, the average MPG of available models (not sales-weighted) stays between about 21.5 and 23 miles per gallon for the entire decade.[3]

We also use gasoline price data from OPIS (Oil Price Information Service) which cover January 1997 to June 2008. OPIS obtains gasoline price information from credit card and fleet fuel card "swipes" at a station level. We purchased monthly station level data for stations in 15,000 ZIP codes. Ninety-eight percent of all new car purchases in our transaction data are made by buyers who reside in one of these ZIP codes.

Since our aim is to estimate the effect of gasoline prices on transactions, we need to match

[2]In 2008, the EPA changed how it calculates MPG. In this figure, the 2008 data point has been adjusted to be consistent with the EPA's previous MPG formula.

[3]While *vehicles* changed fairly little in terms of average fuel economy over this period, this does not mean that there was no improvement in technology to make *engines* more fuel-efficient. The average horsepower of available models increased substantially over the sample years, a trend that pushed toward higher fuel consumption, working against any improvements in fuel efficiency technology. See Knittel (2009) for a discussion of these issues and estimates of the rate of technological progress over this time period.

Figure 1: Average MPG of available cars by model year

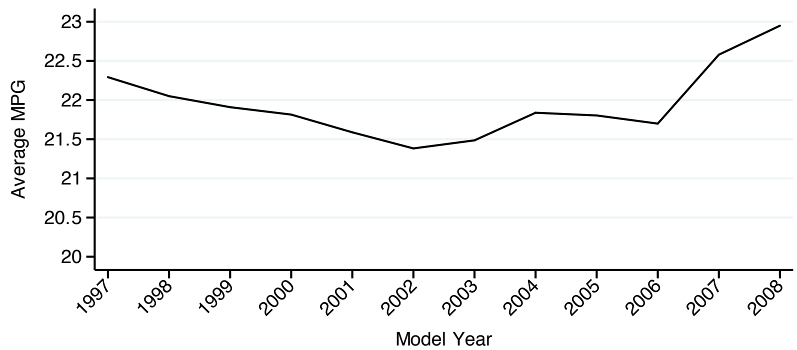

the gasoline price (which we observe at station level) to the location of a car buyer (which we observe at ZIP-code level). Althought we could aggregate station-level data to ZIP-codes, this may not be a good approach for two reasons. First, we only observe a small number of stations per ZIP-code, which may lead to measurement error.[4] Second, consumers are likely to react not only to the gasoline prices in their own ZIP-code but also to gasoline prices outside their immediate neighborhood. This is especially true if price changes that are specific to individual ZIP-codes are transitory in nature. As a result, we measure gasoline price by averaging the prices for basic grade over all stations in each local market (as defined by Nielsen Designated Market Areas, or "DMAs" for short). There are 210 DMAs. Examples are "San Francisco-Oakland-San Jose, CA," "Charlotte, NC," and "Ft. Myers-Naples, FL." Later we investigate the sensitivity of our results to different ways of aggregating gasoline prices (see section 7.3).

To get a sense of the variability of gasoline prices, we graph monthly national average gasoline prices. As shown in Figure 2 (left panel), there is substantial intertemporal variation in gasoline prices within our sample period. Between 1999 and 2008, average national gasoline prices were as low as \$1 and as high as \$4. While gasoline prices were generally trending up during this period there are certainly months where gasoline prices fall.

There is also substantial regional variation in gasoline prices. The right panel of Figure 2 illustrates this by comparing three DMAs: Corpus Christi, TX; Columbus, OH; and San Francisco-Oakland-San Jose, CA. California gasoline prices are substantially higher than prices in Ohio (which are close to the median) and Texas (which are low). While the three series generally track each other, in some months the series are closer together and in other months they are farther apart, reflecting the cross-sectional variation in the data.

[4] In our data, the median ZIP code reports data from 3 stations on average over the months of the year. More than 25% of ZIP-codes have only one station reporting.

Figure 2: Monthly average gasoline prices (national and by DMA)

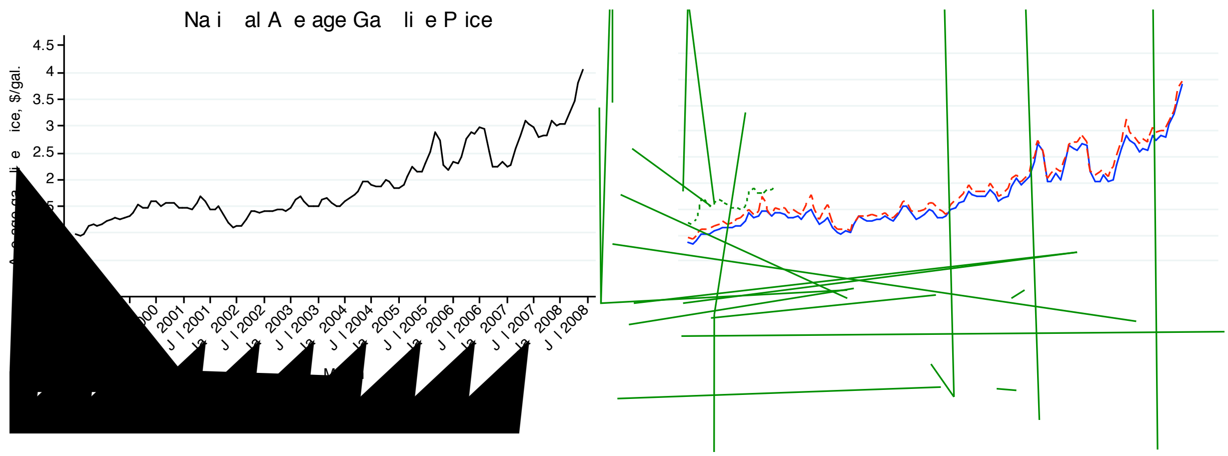

To create our final dataset, we draw a 10% random sample of all transactions.[5] After combining the three datasets this leaves us with a new car dataset of 1,866,366 observations and a used car dataset of 1,264,092 observations. Table 7 presents summary statistics for the two datasets.

4 Gasoline price and car sales

In this section, we address the first of the three main questions addressed in this paper, namely, what effect do gasoline prices have on the sales of cars of different fuel economies? In this section, we focus on new car sales, since these cars represent incremental additions to the U.S. vehicle fleet. While gasoline prices could also affect the pattern of used car sales, from an environmental perspective, these vehicles simply change hands with no effect on the composition of the vehicle fleet.[6] Nevertheless, for completeness, we briefly discuss the effect of gasoline prices on used car transaction shares.

4.1 Specification and variables

At the most basic level, our approach is to model the effect of covariates on short-run equilibrium quantity and (in the next section) price outcomes. For the car industry, the short-run horizon is measured in months, a time frame during which a manufacturer's offering of models is fixed, its model-specific production capacity is largely fixed, as are a number of input arrangements

[5] The 10% sample is necessary to accommodate large numbers of fixed effects, including fixed effect interactions, that we use later in the paper.

[6] Used cars could move from higher-mileage drivers to lower-mileage drivers, or vice versa, which would ultimately affect greenhouse gas emissions. Investigating this is beyond the scope—and the the informativeness of the data—of this paper.

(labor contracts, in particular). Over a longer horizon, say a year or two, some of these aspects become more flexible (models can be tweaked, some capacity can be altered). Only over a long-run horizon (several years), can a manufacturer introduce fundamentally different models into its product offering.

We use a reduced form approach. In completely generic terms, this will mean regressing observed quantities (Q), or some function of Q, on demand covariates (X^D) and supply covariates (X^S):

$$Q = \alpha_0 + \alpha_1 X^D + \alpha_2 X^S + \nu \qquad (1)$$

The estimated $\hat{\alpha}$'s we will obtain from this specification will estimate neither parameters of the demand curve nor of the supply curve, but will instead estimate the effect of each covariate on the equilibrium Q, once demand and supply responses are both taken into account.

We will estimate two variants of Equation 1. In the first variant, we will use the market shares of vehicles of different types as an outcome variable, rather than unit sales. There are two advantages to this approach. First, using market share controls for the substantial fluctuation in aggregate car sales over the year. Second, this approach enables us to control for transaction- and buyer-specific effects on car sales. The disadvantage is that if changes in gasoline prices affect total unit sales of new cars too much, changes in market share may not correspond to changes in unit sales. In light of this, we will later estimate a second variant of Equation 1 using several measures of unit sales.

In both cases, our demand covariates will be gasoline prices (the chief variable of interest), customer demographics, and variables describing the timing of the purchase, all described in greater detail below. We will also include region-specific year fixed effects and region-specific month-of-year fixed effects. Supply covariates should presumably reflect costs of production of new cars (raw materials, labor, energy, etc.). We suspect that these vary little within the region-specific year and within the region-specific month-of-year fixed effects that are already included in the specification. Furthermore, short- to medium-run manufacturing and pricing decisions for automobiles are not made on the basis of small changes to manufacturing costs. While we realize that almost any model of profit maximization an economist would write down would have pricing and production depend on costs, our interactions with executives responsible for these decisions at car manufacturers indicate that this is not the way short- to medium-run pricing and manufacturing decisions are made in practice.

This leaves us with the following specification for our market share regression. We estimate the effect of gasoline prices on market shares of models of different fuel economies using a series of linear probability models that can be written as:

$$I_{irt}(j \in K) = \gamma_0 + \gamma_1 \text{GasolinePrice}_{it} + \gamma_2 \text{Demog}_{it} + \gamma_3 \text{PurchaseTiming}_{jt} + \tau_{rt} + \mu_{rt} + \epsilon_{ijt} \quad (2)$$

$I_{irt}(j \in K)$ is an indicator that equals 1 if transaction i in region r on date t for car type j was for a car in class K.[7] A "car type" in our sample is the interaction of make, model, model year, trim level, doors, body type, displacement, cylinders, and transmission. (For example, one "car type" in our data is a 2003 Honda Accord EX 4-door sedan with a 4-cylinder 2.4-liter engine and automatic transmission.) We use quartiles of fuel economy to define the classes into which a car type falls.[8] (We define the quartile measure more precisely below.) The variable of primary interest is *GasolinePrice*, which is specific to the month in which the vehicle was purchased and to the DMA of the buyer.[9]

We use an extensive set of controls. First, we control for a wide range of demographic variables ($Demog_{it}$), namely the income, house value and ownership, household size, vehicles per household, education, occupation, average travel time to work, English proficiency, and race of buyers by using Census data from the 2000 Census at the level of "block groups," which, on average, contain about 1100 people.[10] We also control for a series of variables that describe purchase timing ($PurchaseTiming_{jt}$). These variables include: a dummy variable, *EndOfYear*, that equals 1 if the car was sold within the last 5 days of the year; a dummy variable, *EndOfMonth*, that equals 1 if the car was sold within the last 5 days of the month and a dummy variable, *WeekEnd*, that specifies whether the car was purchased on a Saturday or Sunday. If there are volume targets or sales on weekends or near the end of the month or the year, we will absorb their effects with these variables. $PurchaseTiming_{jt}$ also includes fixed effects for the difference between the model year of the car and the year in which the transaction occurs. This distinguishes between whether a car of the 2000 model year, for example, was sold in calendar 2000 or in calendar 2001. Finally, we include year, τ_{rt}, and month-of-year, μ_{rt}, fixed effects corresponding to when the purchase was made. Both

[7] Our results do not depend on the linear probability specification; we obtain nearly identical results with a multinomial logit model (see section 7.5).

[8] In previous versions we have used classes based on segments (e.g. compact, SUV, midsize) and subsegments (e.g., entry compact, premium compact, mini SUV, compact SUV). These results are available from the authors.

[9] Another approach would be to use a variable that represents gasoline price expectations, perhaps based on futures prices for crude oil. One might argue on theoretical grounds that this is the price customers should use in calculating the usage cost of a durable good. In practice, however, during the time period of our sample futures prices for crude oil are quite highly correlated with current gasoline prices, reducing the statistical power from such an exercise. As a result, we have not undertaken this approach.

[10] One might argue that our specification should not hold the demographics of buyers constant for the following reason: Any change in market shares of fuel-efficient vs. fuel-inefficient cars due to changes in demographics associated with fuel price changes can legitimately be considered to be part of the short-run equilibrium market share effect of changing gasoline prices. We have estimated all of our sales specifications without demographic covariates and find that our qualitative results are robust to the exclusion of the demographic variables.

year and month-of-year fixed effects are allowed to vary by the geographic region (29 throughout the U.S.) in which the car was sold. This takes into account that year-over-year and seasonal preferences for specific car classes may vary by region of the country. To examine the robustness of our results to which components of variation in the data are used to identify the effect of gasoline prices, we repeat our estimation with a series of different fixed effect specifications in Section 7.1.

Finally, note that our estimates should be interpreted as estimates of the short-run effects of gasoline prices. By "short-run" we mean effects on market shares and prices over the time horizon in which manufacturers would be unable to change the configurations of cars they offer in response to gasoline price changes. Defined this way, the short-run horizon is several years at least. Persistently higher gasoline prices would presumably cause manufacturers to change the kinds of vehicles they choose to produce, as U.S. manufacturers did in the 1970s at the time of the first oil price shock.[11] The nature of our data, its time span, and our empirical approach are all unsuited to estimating what the long-run effects of gasoline price would be on market shares and prices. The short-run estimates are nevertheless useful, we believe, both because short-run effects are important in the short-to-medium term (especially if financial solvency is an issue) and because they yield some insight into the size of the pressures to which manufacturers are responding as they move towards the long run.

4.2 New car market share results

We first consider the effect of gasoline prices on the market shares of new cars in different classes of fuel economy. Specifically, we classify all transactions in our sample by the fuel economy quartile (based on the EPA Combined Fuel Economy MPG rating for each model) into which the purchased car type falls. Quartiles are re-defined each year based on the distribution of all models *offered* (as opposed to the distributions of vehicles sold) in that year. Table 8 reports the quartile cutoffs and mean MPG within quartile for all years of the sample. (Note that the effect of a change in the EPA rating system can be seen in the abrupt change between 2007 and 2008. Our estimates include fixed year effects which capture level shifts in the EPA rating system.)

In order to estimate Equation 2 with car class defined by MPG quartile, we define four different dependent variables. The dependent variable in the first estimation is 1 if the purchased car is in fuel economy quartile 1, and 0 otherwise. The dependent variable in the second estimation is 1 if the purchased car is in fuel economy quartile 2, and 0 otherwise, and so on. To account for correlation in the errors due to either supply or demand factors, we cluster the standard errors at

[11] As gasoline prices began to fall in the early 1980s, CAFE standards also affected manufacturer offerings.

the DMA level.

The full estimation results are reported in Table A-1. The estimated gasoline price coefficients (γ_1) for each specification are presented below.[12] We also report the standard errors of the estimates, and the average market share of each MPG quartile in the sample period. (Note that the quartiles are based on the distribution of available models while the market share is sales-weighted, which is why market shares need not be 25% for each quartile.) Combining information in the first and third column, we report in the last column the percentage change in market share that the estimated coefficient implies would result from a $1 increase in gasoline prices.

Fuel Economy	Coefficient	SE	Mean market share	% Change in share
MPG Quartile 1 (least fuel-efficient)	-0.05**	(0.0049)	20.90%	-23.9%
MPG Quartile 2	-0.014**	(0.004)	21.20%	-6.60%
MPG Quartile 3	-0.0065*	(0.0029)	23.70%	-2.74%
MPG Quartile 4 (most fuel-efficient)	0.07**	(0.005)	34.20%	20.5%

These results suggest that a $1 increase in gasoline price decreases the market share of cars in the least fuel-efficient quartile by 5 percentage points, or 23.9%. Conversely, we find that a $1 increase in gasoline price increases the market share of cars in the most fuel-efficient quartile by 7 percentage points, or 20.5%. This provides evidence that higher gasoline prices are associated with the purchase of more fuel-efficient cars. Notice that these estimates do not simply reflect an overall trend of increasing gasoline prices and increasing fuel economy; since we control for region-specific year fixed effects, all estimates rely on within-year, within-region variation in gasoline prices and associated purchases.[13]

4.3 New car sales results

While the market share results allow us to investigate the effect of gasoline prices on automobile purchase choices while controlling for transaction– and buyer-specific characteristics, they do not allow us to draw inferences directly about changes in unit sales. The primary reason for this is that changes in gasoline prices may also be correlated with changes in the total number of vehicles sold. A higher market share of a smaller market could correspond to a unit decrease in sales, just as a smaller market share of a bigger market could correspond to a unit increase in sales. In this subsection, we re-estimate Equation 1 using two unit sales measures for Q.

[12]Two asterisks (**) signifies significance at the .01 level, * signifies significance at the .05 level and + at the .10 level. We do not restrict the γs to sum to zero; the sum equals -0.0001.

[13]Nor are the results due to seasonal correlations between gasoline prices and the types of cars purchased at different times of year, since the regressions control for region-specific month-of-year fixed effects.

The first measure we use aggregates our individual transaction data into unit sales by dealer, for each month, by MPG quartile.[14] Using this measure, we estimate:

$$Q_{dkrt} = \gamma_0 + \boldsymbol{\gamma_1}(\text{GasolinePrice}_{dt} \cdot \text{MPG Quartile}_k) + \boldsymbol{\gamma_2}\text{MPG Quartile}_k + \delta_d + \tau_{rt} + \mu_{rt} + \epsilon_{dkrt}. \quad (3)$$

Q_{dkrt} is the unit sales at dealer d located in region r for vehicles in MPG quartile k that occur in month t. The variable of primary interest is the *GasolinePrice* in month t in the DMA in which dealer d is located. This variable is interacted with an indicator variable which equals 1 if the observation is for cars in MPG quartile k. The coefficients of interest are the four coefficients in the vector $\boldsymbol{\gamma_1}$ which represent the effect of gasoline prices on the sales of cars in each of the four MPG quartiles. The interaction term allows us to estimate separate gasoline price effects for each of the four quartiles. We include fixed effects for each of the MPG quartiles and for individual dealers (δ_d). Finally, as in Equation 2, we include year, τ_{rt}, and month-of-year, μ_{rt}, fixed effects that are are allowed to vary by the geographic region of the dealer.

While this measure enables us to look at effects on unit sales (instead of market share) while still controlling for many local characteristics (via dealer fixed effects), the estimated coefficients will represent the effects on sales at an average dealer. In order to discuss the implication of the estimated effects for the U.S. as a whole, we need to use a measure that will indicate the effect of gasoline prices on national unit sales.

For national sales data, we use information from Ward's Auto Infobank.[15] Using these data, we estimate:

$$Q_{kt} = \gamma_0 + \boldsymbol{\gamma_1}(\text{GasolinePrice}_t \cdot \text{MPG Quartile}_k) + \boldsymbol{\gamma_2}\text{MPG Quartile}_k + \tau_t + \mu_t + \epsilon_{kt}. \quad (4)$$

Q_{kt} is the national unit sales for vehicles in MPG quartile k that occur in month t.[16] The variable of primary interest is again *GasolinePrice*, which is now measured as a the national average in month t. The coefficients of interest are the four coefficients in the vector $\boldsymbol{\gamma_1}$ which represent the effect of gasoline prices on the sales of cars in each of the four MPG quartiles. We include fixed effects for each of the MPG quartiles, and for year, τ_t, and month-of-year, μ_t.

Using a single national measure for sales means that we cannot utilize regional variation in our

[14] We aggregate from our full data set, not the 10% random sample that we use elsewhere in the paper.

[15] Our transaction data are from a representative sample of dealers, according to our data source. So one approach might be simply to use our data and multiply by the inverse of the sample percentage to get a national figure. Unfortunately, the sample percentage changes over time, and we don't know the year-to-year scaling factor.

[16] Ward's reports sales data for some cars by a more aggregate model designation than the EPA uses to report MPG's. We use the sales fractions in our transaction data to allocate models to which this issue applies in the Ward's data into MPG quartiles.

gasoline price data, or use the region-specific year and month-of-year controls that we have used so far in the paper. In our final sales specification, we use the information in our transaction data about the regional distribution of sales within an MPG quartile to divide the Ward's national sales into regional sales. Specifically, for each month in the sample, we calculate from the transaction data the fraction of sales in each MPG quartile that occurred in each region. We then designate that fraction of the Ward's sales in the corresponding MPG quartile to have occurred in the corresponding region. After doing so, we estimate:

$$Q_{krt} = \gamma_0 + \boldsymbol{\gamma_1}(\text{GasolinePrice}_{rt} \cdot \text{MPG Quartile}_k) + \boldsymbol{\gamma_2}\text{MPG Quartile}_k + \tau_{rt} + \mu_{rt} + \epsilon_{krt}. \quad (5)$$

Q_{krt} is the unit sales for vehicles in MPG quartile k that occur in month t that are imputed (on the basis of the transaction data) to have occurred in region r. The variable of primary interest is *GasolinePrice*, which is a regional average in month t. The coefficients of interest are the four coefficients in the vector $\boldsymbol{\gamma_1}$ which represent the effect of gasoline prices on the sales of cars in each of the four MPG quartiles. We include fixed effects for each of the MPG quartiles, and region-specific year, τ_{rt}, and month-of-year, μ_{rt}, fixed effects.

The coefficient estimates of Equations 3, 4, and 5 are reported in the tables below. For each specification, the tables report the estimated gasoline price coefficients for each of the four MPG quartiles, then the average unit sales, and the percentage change relative to the average implied by the coefficients for a $1 increase in the price of gasoline. The estimated results for dealer sales indicate that on average dealers sell 10.7 cars per month in the least fuel-efficient quartile of available cars, and that a $1 increase in gasoline prices reduces that number by 2.9 cars, or 27.2%. On average dealers sell 16.7 cars per month in the most fuel-efficient quartile of cars, and an increase in gasoline prices increases that number by 1.9 cars, or 11.4%. Notice that an increase in gasoline prices is predicted to reduce the total sales of new cars (adding up the predicted effects across quartiles), which makes the percentage changes in unit sales more negative quartile-by-quartile than the percentage changes in market share reported in the previous subsection.

Fuel Economy	Coefficient	SE	Average cars sold per month in dealer	% Change in sales
MPG Quartile 1 (least fuel-efficient)	-2.9**	(.089)	10.7	-27.20%
MPG Quartile 2	-0.76**	(.085)	10.6	-7.20%
MPG Quartile 3	-0.73**	(.087)	12.2	-5.97%
MPG Quartile 4 (most fuel-efficient)	1.9**	(.11)	16.7	11.38%

According to the estimates using the Ward's national sales data, reported in the next table, when gasoline prices increase by $1, there are 76,144 fewer cars per month sold in the least fuel-efficient quartile of cars. This is a 26.3% decrease relative to the 289,076 monthly average in this quartile. In the most fuel-efficient segment, a $1 increase in gasoline prices is associated with an increase in monthly sales of 34,791 cars, a 9% increase on the average monthly sales in this quartile of 384,990.

Fuel Economy	Coefficient	SE	Average cars sold per month nationally	% Change in sales
MPG Quartile 1(least fuel-efficient)	-76,144**	(9083)	289,076	-26.34%
MPG Quartile 2	-13,793	(9744)	263,214	-5.24%
MPG Quartile 3	-29,124**	(9197)	319,180	-9.12%
MPG Quartile 4 (most fuel-efficient)	34,791**	(11881)	384,990	9.04%

Using the regional breakdown of the Ward's data, reported in the last table, produces very similar estimates: in the least fuel-efficient quartile a unit sales decrease of 25.2% (2,158 cars relative to a mean of 8,555 cars) and in the most fuel-efficient quartile, an increase of 12.7% (1,435 cars relative to a mean of 11,257 cars).

Fuel Economy	Coefficient	SE	Average cars sold per month in region	% Change in sales
MPG Quartile 1(least fuel-efficient)	-2158**	(143)	8,555	-25.22%
MPG Quartile 2	-410**	(133)	7,715	-5.31%
MPG Quartile 3	-855**	(130)	8,960	-9.55%
MPG Quartile 4 (most fuel-efficient)	1435**	(157)	11,257	12.74%

Overall, the results we obtain using unit sales tell a consistent story whether they are measured at the dealer, national, or (imputed) regional level. They are also broadly consistent with the market share results estimated in the previous subsection, with the primary difference that the unit sales results reveal a reduction in total car purchases when gasoline prices increase that is masked by the market share results.

Implications for climate policy

We motivated this section by suggesting that knowing these results would be useful for understanding the possible effects of climate change policies that influenced the price of gasoline. It is beyond the scope of this paper to do a complete simulation of a policy such as a gasoline tax. (Bento, Goulder, Jacobsen, and von Haefen (2009) undertake precisely this task.) But we can do a back-of-the-envelope sketch that gives some context to our results. Our sketch should be interpreted with some care, as it represents the effect of a gasoline price change only as it affects new car sales, not driving behavior, scrappage, or other avenues through which a gasoline price change might have an effect.

Suppose that a tax raised the price of gasoline by $1. Our results suggest that that price increase would change the number of new cars of different fuel economies that are purchased. Suppose that the increase in the gasoline price affected one year's worth of new car sales at the rates predicted by the regional sales results above. How different would one year's worth of gasoline consumption be for the counterfactual set of new cars purchased compared to a baseline of the average number of new cars purchased in each segment? For example, on average, there are 3,363,040 cars sold per year that are in the least fuel-efficient quartile of cars. A $1 gasoline price increase predicts a 25.22% reduction in that number, to 2,514,782 cars. The average MPG of cars in this quartile is 16.2 MPG. If we assume that those cars are driven 10,984 miles per year[17], then on average, each uses 678 gallons of gasoline. Having 848,258 fewer of them on the road would reduce gasoline consumption by 575,139,799 gallons of gasoline. Doing the same exercise for the other three segments and adding up the results yields a reduction in gasoline consumption of 608,745,310 gallons per year.[18] A $1 increase in the price of gasoline is predicted to reduce the total number of new vehicles purchased by 789,384. Since those drivers will probably not stop driving entirely, but instead keep driving their current cars, we will count the gasoline usage of those cars against the gasoline usage reduction above. The average MPG of used cars in our sample is 22.1 MPG, which implies gasoline usage of 497 gallons per year, or 392,334,411 gallons per year in total. This makes the net change 216,410,900 fewer gallons per year. At 19.4 pounds of CO_2 emissions per gallon of gasoline, this is a reduction of 2.1 million tons of CO_2 per year.

4.4 Used car transaction share results

While we can easily estimate Equation 2 using our data on used car transactions, the estimates do not have the same interpretation as the estimates for new cars. As described in the introduction, changes in the market share of new cars measure how the incremental additions to the U.S. vehicle fleet change when gasoline prices change. The analogous estimates arising from the used car data would not measure changes in market share in this sense, but instead changes in "transaction share;" namely how gasoline price affects the share of used car transactions that are for cars in different quartiles. For completeness, we present these results briefly.

We estimate the same specifications as we used to estimate the new car results, namely Equation 2, but using data from used car transactions at the same dealerships at which we observe new car transactions. The full results of transaction share effects of gasoline prices by MPG quartiles

[17] This is the average odometer reading of one-year-old trade-ins that we observe in our data.

[18] The corresponding numbers for the other quartiles are the following. Quartile 2: 2,987,885 cars per year, going down by 5.31% to 2,829,198. At 19.5 average MPG this is 89,385,330 fewer gallons of gasoline per year. Quartile 3: 3,470,177 cars per year, going down by 9.55% to 3,138,945. At 22.7 average MPG this is 160,275,243 fewer gallons of gasoline per year. Quartile 4: 4,306,146 cars per year, going up by 12.74% to 4,854,938. At 27.9 average MPG, this is 216,055,062 more gallons of gasoline per year. We also note that because there is likely to be shifts toward fuel economy *within a quartile*, these calculations will understate the benefits.

are reported in Table A-3. The gasoline price coefficients are as follows:

Fuel Economy	Coefficient	SE	Mean share	% Change in share
MPG Quartile 1 (least fuel-efficient)	-0.016*	(0.0074)	24.10%	-6.64%
MPG Quartile 2	-0.019**	(0.006)	21.00%	-9.05%
MPG Quartile 3	0.026*	(0.012)	25.90%	10.04%
MPG Quartile 4 (most fuel-efficient)	0.01	(0.009)	28.90%	3.46%

The results at the extremes of the fuel-economy distribution are both smaller in magnitude and weaker in statistical significance than the analogous results for new cars. For new cars, market share changes were quite consistently related to gasoline price, with the most fuel-efficient quartile showing the largest increase (7 percentage points) and the least fuel-efficient quartile showing the largest decrease (5 percentage points) in conjunction with gasoline price increases. For used cars, the most fuel-efficient quartile shows no statistically significant effect of gasoline price changes on transaction share while the least fuel-efficient quartile shows a much smaller (1.6 percentage point) decline than for new cars.

Summary

Overall, we find both statistically and economically significant effects of gasoline prices on new car sales, measured either as market shares or by unit sales. This is particularly true for the "extremes," measured as the most fuel-efficient and least fuel-efficient quartiles, where market share shifts by more than 20% in response to a $1 increase in gasoline prices and where unit sales shift by more than 25% for the least fuel-inefficient quartiles and around 10% for the most fuel-efficient.

5 Gasoline price and car prices

In this section, we address the second of the three main questions investigated in this paper, namely what is the effect of gasoline price on the transaction prices of new and used cars? Answering this question will enable us to investigate whether there is evidence that consumers exhibit "myopia" about future fuel costs of different cars when they are considering the up-front purchase decision. We will begin by describing our empirical approach. Then we will present our estimates of the effect of gasoline prices on the transaction prices of new and used cars. Finally, we will use these estimates to investigate the question of myopia in each of the markets.

5.1 Empirical approach

The basic starting point for the consumer myopia literature is a simple idea: if the expected future usage cost of a durable good increases, all else equal, then the price of that durable good should fall by the same amount. In other words, consumers' total willingness-to-pay for the good should be unchanged, all else equal, so that if one component of the total cost rises, the other must fall if consumers are to keep purchasing the good. A direct approach to testing whether consumers "correctly" value future fuel costs would be to estimate a demand relationship in which expected future fuel costs were included as a covariate, and test whether the relevant coefficient has the value that would be implied by consumers correctly valuing fuel costs.

In the automotive setting, there are two difficulties to actually estimating this relationship. One is that, in the cross-section, differences between cars in fuel costs are often related to differences between those cars in other attributes that are valued by consumers as goods; for example, size, weight, power, or other unobservables. This can make the empirical cross-sectional relationship between price and fuel cost positive. Of course, adequate controls for characteristics, or detailed car fixed effects, could remedy this.

A second problem is that if intertemporal variation in gasoline prices is used to identify the relationship between a car's price and its future fuel cost, the "all else equal" condition is violated: a rise in the price of gasoline which increases the cost of operating one car will increase the cost of operating *all* gasoline-powered cars. This means that if consumers are sufficiently unwilling to substitute away from cars as a whole, a rise in the price of gasoline might well *increase* the price of relatively fuel-efficient cars even if their operating costs have actually gone up, because the operating cost would have decreased *relative* to that of a fuel-inefficient car.

To see how this latter point affects the estimation of the relationship between future fuel costs and car prices, consider a market with two vehicles, 1 and 2. Suppose that the price of vehicle i is given by p_i and that the present discounted value of the expected future gasoline cost for operating vehicle i over its lifetime is given by G_i. For simplicity, suppose that demand is linear, implying the demand for vehicle 1 can be written as:

$$q_1 = \alpha_1 + \beta_{11}(p_1 + G_1) + \beta_{12}(p_2 + G_2) \tag{6}$$

Solving this for price implies the following relationship:

$$p_1 = -G_1 + \frac{1}{\beta_{11}}q_1 - \frac{\alpha_1}{\beta_{11}} - \frac{\beta_{12}}{\beta_{11}}(p_2 + G_2) \tag{7}$$

The theoretical prediction for consumers who correctly value future fuel costs is that the coefficient on G_1 should equal -1. One could test whether consumers really do value upfront costs and discounted fuel costs equally by estimating a free parameter in front of G_1:

$$p_1 = -\gamma G_1 + \frac{1}{\beta_{11}} q_1 - \frac{\alpha_1}{\beta_{11}} - \frac{\beta_{12}}{\beta_{11}} (p_2 + G_2) \tag{8}$$

There are three difficulties in estimating this relationship in practice. First, there are many more than two cars available to consumers. A general model would have to specify the price of vehicle i as a function of the fuel cost of vehicle i and of the fuel costs of all other vehicles separately. Given the large number of vehicles offered in the U.S. market, this will be difficult to implement.[19] A second difficulty is that there may be endogeneity between q_i and p_i, arising from a supply relationship between the two variables. A third difficulty is that the econometrician must have a good model of how the marginal consumer constructs the present discounted value of the expected future gasoline cost. The inputs necessary for constructing this variable are expectations of future gasoline prices, expected future driving behavior (vehicle miles travelled), the discount rate, and the miles-per-gallon of the car. Of these, only the last is likely to be observable as data. This final difficulty is important because it means that attenuation bias in the key parameter is likely to exist, biasing the researcher toward concluding that consumers are myopic.

In this paper, we will take an alternative approach. Our approach will be to estimate the reduced form relationship between gasoline prices and equilibrium car prices. We will combine these estimates with estimates of elasticity of demand for new cars, and estimates of future gasoline prices, vehicle miles travelled, and discount rates in order to address the question of whether consumers are myopic with respect to future fuel costs. While our approach will lack the elegance of addressing the question in a single estimated parameter, it will be more amenable to examining the effect of a variety of assumptions about vehicle miles travelled, future gasoline prices, and discount rates. We will present a range of estimates; it will be fairly straightforward for readers to substitute their own assumptions as well.

5.2 Specification and variables

In this section we estimate the reduced form effect of gasoline prices on the equilibrium prices of new cars of different fuel economies (analogous to our estimate of the effect of gasoline prices on new car sales in section 4.1). We can solve Equation 1 for price instead of quantity, and rewrite

[19] An alternative approach, used by Allcott and Wozny (2010), is to specify a nested logit demand system and then to solve for equilibrium prices. The benefit of this approach is that in the logit model the usage cost of all other vehicles drops out of the estimating equation once the market share of each car is divided by the share of the outside good. The cost is that it imposes a specific functional form assumption on the data. If the model is not a good match for the data, the estimates could lead to erroneous inferences.

the result as:

$$P = \beta_0 + \beta_1 X^D + \beta_2 X^S + \eta \tag{9}$$

As with Equation 1, the estimated $\hat{\beta}$s will measure neither parameters of the demand curve, nor parameters of the supply curve, but instead the estimated short-run effects of the covariates on equilibrium prices. In practice, we will estimate the following equation, which contains the same controls as we used in Equation 2, with one addition:

$$\begin{aligned} P_{irjt} = \lambda_0 + \lambda_{\mathbf{1}}(\text{GasolinePrice}_{it} \cdot \text{MPG Quartile}_j) + \lambda_2 \text{Demog}_{it} + \\ \lambda_3 \text{PurchaseTiming}_{jt} + \delta_j + \tau_{rt} + \mu_{rt} + \epsilon_{ijt} \end{aligned} \tag{10}$$

The price variable recorded in our dataset is the pre-sales-tax price that the customer pays for the vehicle, including factory installed accessories and options, and including any dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.[20] Conceptually, our price variable should measure the customer's total wealth outlay for the car. In order to capture this, we make two modifications to the price variable from our dataset. First, we subtract off the manufacturer-supplied cash rebate to the customer if the car is purchased under a such a rebate, since the manufacturer pays that amount on the customer's behalf. Second, we subtract from the purchase price any profit the customer made on his or her trade-in (or add to the purchase price any loss made on the trade-in). The price the dealer pays for the trade-in vehicle minus the estimated wholesale value of the vehicle (as booked by the dealer) is called the *TradeInOverAllowance*. Dealers are willing to trade off profits made on the new vehicle transaction and profits made on the trade-in transaction, which is why the *TradeInOverAllowance* can be either positive or negative. When a customer loses money on the trade-in transaction, part of his or her payment for the new vehicle is an in-kind payment with the trade-in vehicle. By subtracting the *TradeInOverAllowance* we adjust the negotiated (cash) price to include this payment.

In Equation 10, P_{irjt} is the above-defined price for transaction i in region r on date t for car j, and the control variables are as described in section 4.1 (page 10). For the price specification, we also control for detailed characteristics of the vehicle purchased by including "car type" fixed effects (δ_j).[21]

We estimate how gasoline prices affect the transaction prices paid for cars of different fuel economies. One might think that since higher gasoline prices make car ownership more expensive, higher gasoline prices will lead to lower negotiated prices for all cars. However, this would ignore

[20]Dealer-installed accessories that contribute to the resale value include items such as upgraded tires or a sound system, but would exclude options such as undercoating or waxing.

[21]For a definition of "car type" see page 10.

the results of the previous section, which show that as gasoline prices increase, some cars experience sales increases and others decreases. It would thus not be surprising if the transaction prices of the most fuel-efficient cars were to increase as a result of a gasoline price increase. To capture this, we estimate separate coefficients for the *GasolinePrice* variable, depending on the fuel economy quartile into which car j falls; the quartiles are redefined each model year, as described in section 4.2.[22]

5.3 New car price results

The full results from estimating Equation 10 are presented in Table A-2. The gasoline price coefficients are as follows:

Variable	Coefficient	SE
GasolinePrice*MPG Quart 1 (least fuel-efficient)	-236**	(74)
GasolinePrice*MPG Quart 2	-74+	(40)
GasolinePrice*MPG Quart 3	6.9	(30)
GasolinePrice*MPG Quart 4 (most fuel-efficient)	127**	(43)

These estimates indicate that a \$1 increase in gasoline price is associated with a lower negotiated price of cars in the least fuel-efficient quartile (by \$236) but a higher price of cars in the most fuel-efficient quartile (by \$127), a relative price difference of \$363. Overall, the change in negotiated prices appears to be monotonically related to fuel economy. Note that this is an equilibrium price effect; it is the net effect of the manufacturer price response, any change in consumers' willingness to pay, and the change in the dealers' reservation price for the car.[23]

5.4 Used car price results

In this section, we estimate the effect of gasoline prices on the transaction prices of used cars. We do so by estimating the same specification we used for new car prices, namely Equation 10, using instead our used car transaction data. All the control variables are the same; in particular, we observe all the same car characteristics for used cars that we do for new cars, so the "car type" definition is the same. The definition of the dependent variable is almost the same as that used for new cars. A customer who is buying a used car can use a trade-in in the transaction, just as a buyer of a new car can, so the price definition subtracts the *TradeInOverAllowance* just as it does for new cars. However, used cars never have customer rebates offered, so there is no need to subtract that amount from the reported transaction price.

[22] We obtain similar results if we estimate four separate regressions, thereby relaxing the constraint that the parameters associated with the other covariates are equal across fuel economy quartiles.

[23] In previous versions, we estimated Equation 10 separately for each segment. These results are available from the authors.

As we did for new cars, we begin by estimating the effect of gasoline prices on used car prices separately by the MPG quartile of the used car being purchased. The full results are reported in Table A-4. The gasoline price coefficients are as follows:

Variable	Coefficient	SE
GasolinePrice*MPG Quart 1 (least fuel-efficient)	-1073**	(40)
GasolinePrice*MPG Quart 2	-900**	(58)
GasolinePrice*MPG Quart 3	118*	(53)
GasolinePrice*MPG Quart 4 (most fuel-efficient)	1766**	(51)

These estimates show a much larger effect on the equilibrium prices of used cars than was estimated for new cars. The estimates indicate that a $1 increase in gasoline price is associated with a lower negotiated price of cars in the least fuel-efficient quartile (by $1,073) but a higher price of cars in the most fuel-efficient quartile (by $1,766), a relative price difference of $2,839.[24]

5.5 Consumer myopia

In this subsection we address the question of whether consumers are myopic about future gasoline prices when they make car purchase decisions. Analyzing this means, in simple terms, comparing the predicted price effects of gasoline price changes on cars of different fuel economies to the changes in the discounted value of future gasoline costs that is implied by the gasoline price change and the fuel economy of the car. In practice, there are a few wrinkles.

First, to calculate the discounted value of expected future gasoline costs we need to know how many miles per year car owners drive. This requires not only how many miles a car is driven in a given year, conditional on the car surviving through that year, but also annual survival rates. We calculate miles driven, conditional on survival, three ways. We use NHTSA assumed values, which are used in a number of modeling efforts for both the NHTSA and DOT (Lu (2006)). These data report annual miles driven, separately for cars and light duty trucks, by vintage. Our other two measures come from within our data: we compute the average annual miles driven, by vintage, separately for cars and trucks, using the sample of vehicles sold as used vehicles in our transaction

[24]Our estimates suggest that changes in gasoline prices are associated with changes in the relative price of new vs. used cars. Specifically, we find that the price of a fuel-efficient new car falls relative to the price of a used car of the same fuel-economy when gasoline prices increase. In these circumstances we might expect, therefore, to see an increase in the share of purchases of fuel-efficient cars that are purchases of a new car. Conversely, we find that the price of a new fuel-*in*efficient car increases relative to the price of a fuel-inefficient used car as gasoline prices increase, suggesting a decrease in the share of purchases of fuel-inefficient cars that are new car purchases. In unreported results (available from the authors) we find that a $1 increase in gasoline prices is associated with an increase of 3.5 percentage points in the share of compact car transactions that are for new cars (a 5.4% change), 2.6 percentage points for sporty cars (a 4.7% change), and 1 percentage point for luxury cars (a 1.7% change). Conversely, the new car share falls by 2.1 percentage points for SUVs (a 3.3% change), and 2.7 percentage points for vans (a 4.5% change). (The results are not statistically significant for midsize cars or pickups). These effects are consistent with what we would expect given the predicted relative price change.

data and the analogous measure using the sample of all trade-ins we observe being used to purchase either new or used cars in our transaction data. If the typical new or used car purchased at our dealers is replacing the trade-in, one could argue that the calculations based on the miles driven of trade-ins most accurately reflect the driving patterns of those consumers in our data. We also use vehicle survival rates from NHTSA to calculate the expected miles driven for each year of the vehicle's life. Because the median used car is four years old at the time of purchase, we calculate miles driven beginning at the fourth year of life for used cars.

Second, we need to know how consumers form expectations of future gasoline prices. We adopt the simple model of expectations formation that consumers expect gasoline prices to follow a random walk. This has the convenient implication that the current gasoline price is the expected future gasoline price. One alternative is to assume that consumers are more sophisticated and use information on crude oil futures markets to make projections into the future. It turns out that during our sample period, the random walk is the more conservative assumption. But this we mean that assuming consumers view gasoline prices as a random walk biases us towards finding that consumers are myopic if the true model that consumers use is based on crude oil futures. Figure 3 plots both the spot crude price and the stream of expected prices in subsequent years for May of each year—the "forward curve." For the vast majority of time during our sample, the crude market was in backwardation; that is, the market expected crude prices to fall. This means that if consumers actually use crude futures prices to form expectations, and we assume instead that they use a random walk, then for any observed set of changes in willingness-to-pay for cars of different fuel economies, consumers would be more patient than our estimates would show.

Third, we need to know what discount rate customers use to discount future gasoline costs. We reserve this to be our free parameter. In other words, we will use estimates and make assumptions about the various other components of the calculation, and see what they imply for a discount rate.

Fourth, in order to address the question of myopia, what we would really like to observe is not the effects of gasoline prices on equilibrium transaction prices, but on consumers' willingness-to-pay for cars of different fuel economy. In the used car market, one might argue that a fixed supply curve is a reasonable assumption for used car supply. (See the discussion in Section 6 for the argument.) This means that the equilibrium price effect will be equal to the change in willingness-to-pay. (Figure 4 shows a representation of this for a hypothetical used car model.) However, in the new car market, one might well think that the supply relationship is more flexible and that auto manufacturers and car dealers likely have some scope to respond to changes in demand by altering prices, quantities, or both. (Again, see the discussion in Section 6.) This means that

Figure 3: Crude spot and futures prices during our sample

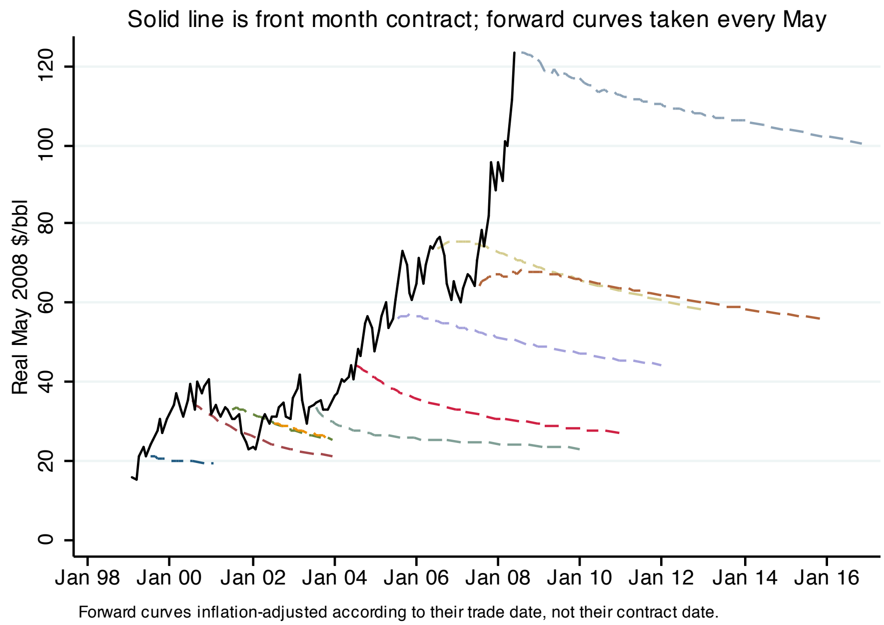

the equilibrium price effect will be less than the change in the willingness-to-pay, and that the difference between the two will be greater the more inelastic the demand curve is. (Figure 4 shows a representation of this for a hypothetical new car model.)

Since we estimate the equilibrium effects on prices and quantities, we could recover the implied effects of gasoline price changes on willingness-to-pay if we had an estimate of the elasticity of demand, as well as an assumed functional form for demand. While estimating an elasticity of demand is beyond the scope of this paper, there are a number of existing papers that have done just this. The relevant elasticity for our calculations is the demand elasticity for entire group of

Figure 4: Effects of gasoline price change on hypothetical used and new cars

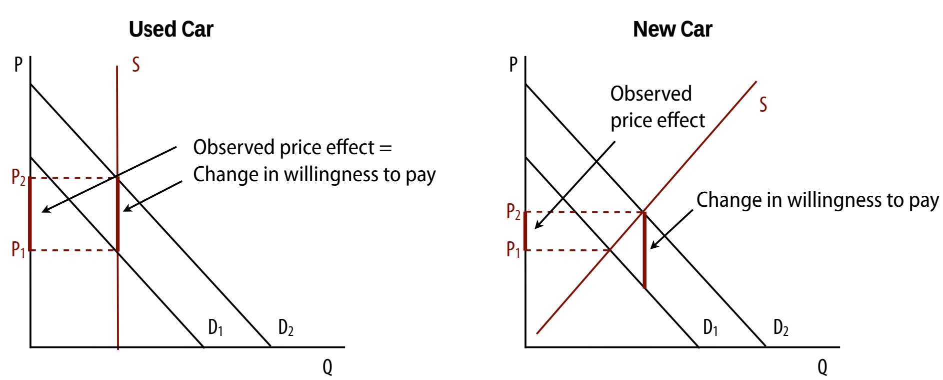

vehicles within a fuel economy quartile. Estimates of the residual demand elasticity for specific vehciles exist, and those estimates vary; Goldberg (1995) estimates elasticities of demand that are in the neighborhood of -2 to -4, while Berry, Levinsohn, and Pakes (1995) estimate elasticities in the -3 to -6 range.[25] We note that these estimates should be strong upper-bounds to the relevant demand elasticity for our purposes. Our estimates represent changes in the willingness-to-pay for a vehicle in a particular fuel-economy quartile; the estimates in Goldberg (1995) and Berry, Levinsohn, and Pakes (1995) are residual demand elasticity estimates for a *specific* vehicle, which are likely to be much higher than the demand elasticity for a much larger group of cars. Finally, we assume that demand has a constant elasticity functional form.[26]

In the table below, we present the results of our investigation into the question of whether consumers are myopic. The results of the table are in the last three columns. The numbers listed in these columns are the implied discount rates that would need to hold in order for the relative price differences between cars of different fuel economies to equal the relative differences in discounted expected future fuel costs between those cars. These are calculated for a variety of different assumptions.[27]

The top panel of Table 1 reports the implied discount rates when comparing the estimated price effects for the least fuel-efficient quartile of cars relative to the most fuel-efficient quartile. The middle panel reports for the least fuel-efficient quartile relative to the second most fuel-efficient; the bottom panel for the second least fuel-efficient relative to the most fuel-efficient. The top row of each panel reports the implied discount rates based on the relative price effects estimated for used cars. The next four rows report the implied discount rates based on the relative price effects estimated for new cars, adjusted to implied willingness-to-pay effects using elasticities of demand ranging from -2 to -5. Finally, the three columns use estimates of vehicle miles travelled from NHTSA, from the used car transaction in our data, and from the trade-ins in our data, respectively.

Overall, the implied discount rates seem to us to be quite small. In most cases, the estimates are in the single digits, with some combinations of assumptions actually implying negative discount rates. The more elastic demand is assumed to be, the smaller the implied change in willingness-to-pay is for a given relative price difference, and the higher is the implied discount rate necessary

[25] Goldberg (1995) reports average elasticites by vehicle segment and origin. The average elasticity across segments is -3.4, while the median is -3.5. Berry, Levinsohn, and Pakes (1995) report elasticity estimates for 13 specific vehicles. Assuming these are representative of the sample, the average elasticity is -5, while the median is -4.8.

[26] The assumption of a constant elasticity demand function has the benefit that, in order to make our calculations, it requires only percentage changes in equilibrium quantities. The calculations assuming a linear demand model where the slope and intercept are chosen such that the elasticity equals that of the constant elasticity demand curve at the average price and quantity are very similar to those reported here.

[27] The spreadsheet that makes this calculations—and could be used to show the influence of different assumptions from those presented here—is available from the authors.

Table 1: New and Used Cars: Implied Discount Rates†

	Market	Assumed Demand Elasticity	NHTSA VMT, NHTSA Survival Rates	VMT from Used Car Transactions, NHTSA Survival Rates	VMT from Tradeins, NHTSA Survival Rates
Q1 vs. Q4	Used	NA	2.4%	-2.7%	-1.7%
	New	2	-3.3%	-6.4%	-5.9%
	New	3	1.9%	-2.4%	-1.6%
	New	4	6.6%	0.0%	2.2%
	New	5	11.1%	4.1%	5.6%
Q1 vs. Q3	Used	NA	16.5%	7.3%	10.1%
	New	2	-3.0%	-6.2%	-5.7%
	New	3	2.2%	-2.2%	-1.4%
	New	4	7.0%	1.4%	1.0%
	New	5	11.5%	4.8%	6.2%
Q2 vs. Q4	Used	NA	-5.4%	-8.3%	-8.0%
	New	2	2.4%	-2.1%	-1.3%
	New	3	9.7%	3.4%	4.7%
	New	4	16.5%	8.3%	10.0%
	New	5	23.2%	13.1%	15.3%

to rationalize that with a given change in expected future gasoline costs. While some of the cells with the highest elasticity of demand have implied discount rates near, or even above, 20%, for the most part, all the estimates are within the range of typical car loan.

As a point of comparison, for those consumers financing their new car purchase through the dealer, the interquartile range of APRs is [4.0%, 8.3%], while the interquartile range for used car transactions is [7.0%, 15%]. When comparing the relative changes in willingness-to-pay for cars in Quartiles 1 and 4 in the new car market, only one of our calculated implied discount rates fall outside of this range. The same is true when comparing Quartiles 1 and 3. More of our calculated implied discount rates fall outside of this range when comparing Quartiles 2 and 4, but a substantial number remain within this range. For used cars, only one of our implied discount rate calculations falls outside of the used car interquartile range for APRs.

We conclude that there is little evidence that consumers undervalue changes in expected future fuel costs, and that the evidence from new and from used cars yield similar messages.

Summary

Overall, we see a modest effect of gasoline prices on new car transactions prices. The predicted effect of a $1 gasoline price increase is to increase the price difference between the most and least fuel-efficient quartiles of cars by $363. The estimated effects are much larger for used cars; in this market, the predicted effect is to increase the price difference between the most and least fuel-efficient quartiles by \$2,839. Despite this striking difference, once the new car equilibrium prices are transformed, using demand elasticity estimates, into willingness-to-pay differences, the two markets tell fairly consistent stories; namely, that there is little evidence in these data that consumers undervalue future fuel costs when they decide which cars to purchase.

6 Market structure and the effect of policy intervention

Two of the policy interventions that have been advocated by economists in the public discussion of climate change policy are an increased gasoline tax and a carbon tax. Both would reduce the incentive for drivers to use gasoline. Drivers could respond to this changed incentive by altering their driving habits, changing their vehicles, or both. In this section we argue that in order to predict the effect of such a policy intervention on vehicle markets, one has to consider the market structure of the vehicle market itself.[28]

We begin by noting that the new car and used car markets which we observe are very similar in most key aspects, including the cars, the consumers, and the retailers. All the transaction data in our sample come from new car dealers. The used cars sold at these dealerships are a positive selection of used cars, making them reasonably comparable to new cars. For example, on average, a used car in our sample sells for \$15,582 compared to \$25,515 for a new car. The consumers in our sample who buy used cars are also not dissimilar to those who buy new cars. For example, on average, a used car buyer comes from a Census block group with a median household income of \$50,826 instead of a median income of \$58,130 for new car buyers. Finally, new and used cars are sold at exactly same retailers in our data; at many of these dealerships, the same salesperson sells both new and used cars.

Where the two markets do differ is in the *supply* of new and used cars. New cars are supplied by the combined activities of auto manufacturers and dealers, while the supply of used cars arises ultimately from the cumulated new car purchases of past years. For new cars, auto manufacturers decide on the prices and quantities of each model they wish to sell. Even though manufacturers seldom adjust MSRP's, they can quickly adjust prices by using promotions (Busse, Silva-Risso, and Zettelmeyer 2006). They can adjust production quantities by adding or reducing shifts on assembly lines (Bresnahan and Ramey 1994), or in the case of some modern manufacturing plants, adjusting which kinds of vehicles are produced on a given line.[29] Car dealers, since they negotiate individual

[28] Busse and Keohane (2007) make a similar argument about the effect of sulfur dioxide regulation on the market for low-sulfur coal. They argue that in order to predict how changing the price of a "dirty" input affects the usage of that input, one must take into account the market structure of related industries; in their case, railroad transportation.

[29] For example, BMW builds their large X5 SUV, and their small Z4 roadster on the same assembly line in

prices with customers, can easily adjust the prices at which they are willing to sell vehicles, and can adjust quantities first by changing inventory holdings, and over several months by changing their orders to manufacturers. Both manufacturers and dealers are likely to have market power in the supply of a particular model, manufacturers because of the differentiation of individual car models, and dealers because of local market power, and so both have some discretion, when faced with a change in demand, to absorb the change in the model's price, in its quantity, or in a combination of the two. Profit maximization, taking into account the shape of each model's demand curve and the shape of its marginal cost curve, will determine what the optimal response is.

In the used car market, the stock of used cars is predetermined by the cumulation of past new car purchases. This stock is likely to respond very little to gasoline prices.[30] Many cars sold on the used market are fleet turnovers and lease returns. The timing for the entry of these cars into the used car market sill not be determined primarily by gasoline prices. To the extent that consumers' decisions to replace their existing cars is also based on timing choices unrelated to gasoline prices, the supply of a particular model available for sale on the used market at any point in time could be thought of as essentially fixed. If this is the case, then one might expect that the effect of a change in demand for that model would show up almost entirely in the equilibrium prices of used cars of different types, and could have very little effect on equilibrium unit sales or transactions shares.[31]

Rapid adjustment of prices in the used car market is also be aided by used car auctions. A large fraction of wholesale transactions for used cars go through independent auctions, unaffiliated with car manufacturers, which are ubiquitously available throughout the country. (For example, Man-

Spartanburg, SC. ("BMW subtracts to add flexibility in S.C.," Automotive News, June 5, 2006) In another example, Honda can build the compact Civic on the same assembly line that builds the Ridgeline pickup and the Acura MDX SUV. In 2008, the last year in our sample, the Civic was in highest fuel-economy quartile of cars while the Acura MDX was in the lowest fuel-economy quartile. ("Adaptability helps Honda weather industry changes," Automotive News, June 8, 2009)

[30] Davis and Kahn (forthcoming) suggest that some low-MPG vehicles may be more likely to be traded to Mexico when the U.S. price of gasoline deviates greatly from the prices set by PEMEX, the national petroleum company.

[31] One might argue instead that a car owner's decision to sell a used car should be influenced by gasoline prices; for example, one might believe that a commuter with a large SUV would be encouraged by high gasoline prices to sell or trade-in that car earlier than he or she otherwise would have in order to replace it with a higher-MPG car. We would argue that even if this is true, the effect of gasoline price changes is likely to show up primarily in used car prices. To see this, consider a potential seller and a potential buyer of a particular fuel-inefficient used car. If the gasoline price increases by some amount, then the *per-mile* cost of driving that particular car increases by the same amount for both drivers. If the two drivers have approximately the same driving habits, then one might expect the effect of the gasoline price increase on the buyer and the seller to be symmetric: for both the current owner and the potential buyer, the increased cost of usage for the current owner of that car will exactly equal the increased cost of usage for the potential buyer if she buys the car. Taking this logic one step further, if most drivers have similar driving habits (or if a large enough number of marginal buyers and sellers have similar driving habits), then the demand curve for a used car should shift inward by roughly the same amount that the supply curve for that car shifts outward (at least in the area of the demand and supply curves where equilibrium occurs).If this were the case, then the prices of a particular used car should adjust to reflect the value of the fuel expenditure disadvantage (or advantage) that car has, given the new gasoline prices.

heim, which is the largest operator of used car auctions in the U.S., has about 100 sites throughout the U.S.) Auctions are generally held every week, and 1,000-3,000 cars might be transacted in a typical week. This means that for car dealers, used cars are convertible into cash, and vice versa, at auction-determined prices on a weekly basis. One might expect that such a mechanism would reflect changes in equilibrium market conditions quite quickly, and would thereby help move the prices of used cars sold at car dealerships fairly quickly to a new, market-clearing equilibrium price that reflected changes in gasoline prices.

The results that we have estimated in the paper so far are consistent with these differences between new and used car supply. We find that the changes in equilibrium prices for used cars in response to changes in gasoline prices are much larger—roughly an order of magnitude larger—than the changes in prices estimated in the new car market. In the new car market, however, we estimate substantial effects of gasoline prices on the market shares and sales of new cars. This suggests that the new car supply chain—manufacturers and dealers combined—is choosing to respond to changes in demand by keeping price relatively unchanged, and absorbing the effect of the demand change in quantities—inventories or production levels—instead.

In this section, we present two pieces of supplementary evidence that are consistent with our previous findings of differences between new and used markets. First, we show that changes in inventory are consistent with the changes in sales we have presented already. Second, we show that the large effects we estimate for used car transaction prices are reflected also in dealers' internal valuations of used cars.

6.1 Inventories

In our data, we can observe an inventory-related measure called "days to turn." Days to turn counts the number of days that a specific vehicle was on a dealer's lot before it sold. Higher average days to turn for a particular class of cars indicates that the dealer is carrying higher inventory levels of that car.

In order to investigate inventory effects, we estimate the effect of gasoline prices on days to turn by MPG quartiles (the specification and results are reported in Table A-12). We find much larger changes in days to turn in response to gasoline price changes for new than for used cars. For new cars, the estimated coefficients imply that a \$1 increase in gasoline price is associated with a 12-day increase in days to turn for cars in the least fuel-efficient quartile, a 17.6% increase from the sample mean of 68.3 days. Conversely, we find that the same gasoline price increase reduces by 5.4 days the time that a car in the most fuel-efficient quartile remains on the lot, a 10.8% decrease relative to an average of 50.2 days. In contrast, for used cars, higher gasoline prices have no statistically significant effect on days to turn for either the least or the most fuel-efficient quartile. The only statistically significant change in days to turn occurs for used cars in the second most fuel-inefficient MPG quartile; days to turn increases by 1.5 days, or 3.2%.

6.2 Actual cash value of trade-ins

In our data, we also observe the amounts that dealers book as the "actual cash value" of trade-ins they receive. While a dealer might wish to manipulate the *price* paid to the customer for his or her trade in, the "actual cash value" is the dealer's internal assessment of the value of the vehicle. In this number, the dealer is trying to approximate the price for which he could have purchased—or could sell—the car at auction. We are interested in how the "actual cash value" of cars of different fuel economies varies with gasoline prices.

In regressions described and reported in Table A-13, we estimate the effect of gasoline prices on the "actual cash value" of trade-ins of different fuel economies. The estimated effects of gasoline prices on actual cash values are similar to the results obtained for the gasoline price effect on used car prices, reported on page 21. The first three quartiles of the actual cash value results are in almost all cases within \$100-300 of the used car price results.[32] This result is consistent with our argument that most of the adjustment to changes in demand in the used car market occur in prices, and that price are aided in their rapid adjustment by a well-functioning wholesale market.

Summary

One of the chief arguments for an increased gasoline tax or a carbon tax in car markets is that it would cause buyers to change the type of vehicle they buy, and thereby future gasoline consumption. The results we have estimated for the effect of gasoline prices themselves in new car markets suggests that this would indeed be the effect of such interventions. This section has pointed out, however, that it is important to be mindful of the market structure of the target market when such an intervention is contemplated. If the new car market were structured like the used car market (very inflexible supply, mechanisms to facilitate rapid adjustment to competitive market prices), then the effect on vehicle market shares or volumes likely would be much smaller.

[32]There is one interesting exception to actual cash values and used transaction car prices showing very similar adjustment, which is the most fuel-efficient quartile of trade-in cars. There the estimated effects of gasoline price on actual cash values are \$1275 (for trade-ins used to buy new cars) and \$778 (for trade-ins used to buy used cars) which are \$500-\$1000 less that the \$1766 estimated effect of a \$1 gasoline price increase on the price of the most fuel-efficient used car. One story that would explain this would be that when gasoline prices rise, customers are particularly interested in buying a good, fuel-efficient used car from dealers, and that dealers are able to mark up such cars in their retail transactions above what the actual market (auction) price is for such cars.

7 Robustness

In this section we explore the robustness of our results. First, we analyze whether our results are robust to changing the component of variation in the data that is used to identify the effect of gasoline prices. Second, we allow the response to gasoline price to differ by the level and prior direction of gasoline prices to see whether we are ignoring important response heterogeneity in our estimates. Third, we analyze the robustness of our findings to the aggregation of gasoline prices. Fourth we analyze whether we should treat gasoline prices as being endogenous. Fifth, we examine whether our results depend on our use of a linear probability model to estimate market share changes in response to gasoline prices.

7.1 Source of variation

We now re-estimate the original specifications in this paper with a series of different fixed effects. Recall that all specifications so far control for region-specific year and region-specific month-of-year fixed effects. This means that the estimated gasoline price effects are identified by within-year, region-specific market share or price changes which deviate from region-specific seasonal effects.

In order to investigate the robustness of theses estimates, we estimate a specification without any year or month-of-year (seasonal) fixed effects. (In our original specification (Equation 2) we captured regional effects by the fact that seasonal dummies were region-specific. To continue to account for regional variation we include region fixed effects for the 29 regions.) Next, we estimate a specification with region-specific seasonal effects but without year fixed effects. The results from these two more parsimonious specifications show whether we are missing an important source of variation due to the extensive set of fixed effects used in our base specification. Finally, we estimate a specification with *more* detailed time fixed effects, by replacing the region-specific year fixed effects with region-specific quarter fixed effects. These results allow us to increase our confidence that our estimated effects are not driven by the generally upward-trending gasoline prices in our sample period.

Table 2 shows the results of these three specification for new and used cars. For comparison, the table repeats the estimates of the original specification in Equation 2.

The coefficients on the gasoline price variable are remarkably robust to which fixed effects are included. In our original specification we found that a \$1 increase in gasoline prices decreased the market share of the least fuel-efficient new cars (MPG Quartile 1) by 5 percentage points. Omitting time and and then both time and seasonal fixed effects changes this estimate by only little, to 4.6 and then 4.8 percentage points, respectively. Similarly, we originally found that a \$1

Table 2: Effect of time and seasonal fixed effects in market share specification†

Specification	Time FE	Seasonal FE	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
New Cars						
Most Parsimonious (Region FE only)	–	–	-.048** (.0043)	-0.0035 (0.0046)	-0.0078** (0.003)	0.059** (0.0061)
More Parsimonious	–	Month-of-year × Region	-0.046** (0.0042)	-0.0022 (0.0042)	-0.0062* (0.0028)	0.054** (0.0055))
Base (original)	Year × Region	Month-of-year × Region	-0.05** (0.0049)	-0.014** (0.004)	-0.0065* (0.0029)	0.07** (0.005)
Richer	Quarter × Region	Month-of-year × Region	-0.065** (0.0067)	-0.024** (0.0063)	0.0093* (0.0042)	0.08** (0.0067)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

increase in gasoline prices increased the market share of the most fuel-efficient cars (MPG Quartile 4) by 7 percentage points. Omitting time and/or seasonal fixed effects decreases our estimate to no lower than 5.4 percentage points. Including more granular time fixed effects than in our original specification by using region-specific quarter fixed effects also has a modest effect; we find a market share decrease of 6.5 percentage points for the least fuel-efficient cars and a market share increase of 8 percentage points for the most fuel-efficient cars. Overall, we conclude that the results from our original specification seem robust to which component of time and seasonal variation in the data is used to identify the effect of gasoline prices on the market shares of cars of different fuel economy.

We investigate the robustness of our price results by estimating a comparable set of specifications. First, we eliminate all time and seasonal fixed effects from our original price specification (Equation 10), while retaining regional fixed effects. Next, we include region-specific seasonal effects but no year fixed effects. Finally, we include region-specific quarter effects in addition to region-specific seasonal effects. The estimation results are reported in Table 3.

The results in Table 3 show that our price estimates are also robust to changing how much variation in the data is absorbed by fixed effects versus used to identify the effect of gasoline price changes. This is especially true if what we compare across specifications is the estimated change in relative prices of the first and fourth quartile cars implied by a \$1 increase in gasoline prices. This estimate is reported in the last column of Table 3. Although the individual quartile coefficients are affected by changes in which fixed effects are used, if we focus on the last column, the estimated relative price effects are very similar across specifications. For new cars, a \$1 increase in gasoline price increases the price of the most fuel-efficient cars (Quartile 4) relative to the most fuel-inefficient cars (Quartile 1) by \$342 to \$366 across all specifications. For used cars, the equivalent numbers are \$2766 to \$2854 across all specifications. Hence, used cars are estimated to experience much

Table 3: Effect of time and seasonal fixed effects in price specification†

Specification	Time FE	Seasonal FE	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4	Price change Quar 1 to 4
New Cars							
Most Parsimonious (Region FE only)	–	–	-513** (82)	-359** (43)	-275** (33)	-147** (38)	$366
More Parsimonious	–	Month-of-year × Region	-245** (74)	-86* (41)	-14 (36)	103* (46)	$348
Base (original)	Year × Region	Month-of-year × Region	-236** (74)	-74+ (40)	6.9 (30)	127** (43)	$363
Richer	Quarter × Region	Month-of-year × Region	-124 (79)	30 (55)	108* (46)	218** (58)	$342
Used Cars							
Most Parsimonious (Region FE only)	–	–	-1117** (38)	-969** (58)	28 (51)	1649** (55)	$2766
More Parsimonious	–	Month-of-year × Region	-778** (41)	-639** (62)	365** (58)	1995** (60)	$2773
Base (original)	Year × Region	Month-of-year × Region	-1073** (40)	-900** (58)	118* (53)	1766** (51)	$2839
Richer	Quarter × Region	Month-of-year × Region	-1208** (59)	-1040** (73)	-25 (70)	1646** (65)	$2854

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

larger price adjustments than new cars when gasoline prices change, whichever set of fixed effects we use.

In summary, these estimates show that our findings are robust to which component of the variation in the data is used to estimate the effect of gasoline prices.

7.2 Heterogeneity in gasoline price response

So far we have used the gasoline price for the month of the transaction in the buyer's DMA as our measure of *GasolinePrice*. In doing so we are disallowing heterogeneity in the response to gasoline prices. In this subsection, we investigate two types of heterogeneity that might exist. First, we investigate whether the response to gasoline prices differs by the absolute price of gasoline. For example, does a \$1 increase in gasoline prices have a different effect on market shares and car prices when gasoline currently costs \$1.50 per gallon from when it costs \$3.50 per gallon? Second, we investigate whether the response to gasoline price differs by whether gasoline prices have been consistently rising or falling in prior periods. For example, is the effect of a \$1 gasoline price increase larger if gasoline prices have been increasing over the previous three months than if they have been flat or declining?

To answer the first question we repeat the main results in Sections 4 and 5 using *GasolinePrice* interacted with indicators for whether the gasoline price falls in the range "<\$1.50," "\$1.50-\$2.50,"

"$2.50-$3.50," or ">$3.50." The purpose is to see whether there is an inflection point of gasoline prices at which the effects suddenly kick in, or at which they grow much larger. News reports have posited that there is a gasoline price "threshold" above which consumers change their behavior more dramatically.[33] Summarizing over many results, we find that gasoline prices do have somewhat different effects at different price levels, but there is little evidence of a sudden inflection. See Tables A-8 and A-9 for a summary of these results; full results are available from the authors.

To answer the second question we also repeat the main results of this section interacting *GasolinePrice* with an indicator variable that records whether gasoline prices went up monotonically in the previous three months, went down monotonically in the previous three months, or had some kind of mixed pattern. These results also show some heterogeneity in the gasoline price response but do not have a consistent enough pattern to draw conclusions about systematic differences in effects under the three conditions. See Tables A-10 and A-11 for a summary of these results; full results are available from the authors.

7.3 Gasoline price aggregation

Next, we investigate the robustness of our findings to the aggregation of gasoline prices. Recall that, although we know the ZIP-code of each buyer, we chose to aggregate gasoline prices at the level of local markets (defined by DMAs). The advantage of using the higher level of aggregation is that we reduce the possibility of measurement error that could arise from our observing only a small number of stations per ZIP-code. The higher level of aggregation also allows consumers to react not only to the gasoline prices in their local ZIP-code but also to gasoline prices in a broader area. At the same time, however, we eliminate some of the cross-sectional variation that less aggregate data would allow us to use.

One could also make the argument that we should use a more aggregate measure of gasoline prices than the DMA-level prices we have used so far. This is because consumers may be more likely to notice gasoline price changes once the gasoline price changes have affected a large area and are thus reported in the media; alternatively, local price variation may contain more transitory price shocks that do not enter into the long-run forecasts of gasoline prices over the life of the car.

To investigate whether our conclusions depend of the level of aggregation of gasoline prices, we re-estimate our original MPG quartile specifications (Equations 2 and 10) using a less aggregated and a more aggregated measure of gasoline prices. We use 4-digit ZIP-code level gasoline price

[33]For example, an article in Automotive News on 5/22/08 entitled "Ford: $3.50 gasoline was tipping point for sales shift" states: "The segment shifts [away from SUVs and Pickups] 'really started to move' when gasoline prices hit $3.50 a gallon, [Ford CEO Alan] Mulally said. 'It seemed to us that we reached a tipping point where customers began shifting away from these vehicles at an accelerated rate,'...."

as our less aggregated measure. We use this instead of 5-digit ZIP-code level price because too many 5-digit ZIP-codes have too few gas stations to calculate a reliable average.[34] For our more aggregated measure, we average the prices for basic grade over all stations in each "Petroleum Administration for Defense District" (PADD). PADDs are the standard geographical classification used by the Energy Information Administration. A PADD's boundaries are defined such that they delineate a region in which supply is homogenous. There are five PADDs: East Coast, Midwest, Gulf Coast, Rockies, and West Coast. There remains substantial variation in gasoline prices across PADDs: not only are prices in some PADDs higher than in other PADDs, there is also variation in the magnitude of the difference (not reported).

The full results are reported in Tables A-5 and A-6. We find that the coefficients on gasoline prices in the 4-digit ZIP code aggregation are similar to those in our (original) DMA aggregation but somewhat smaller in magnitude. This is consistent with some measurement error occurring in the 4-digit ZIP code aggregation. If we aggregate gasoline prices at the PADD level, most coefficient estimates in the market share regression are unchanged. In the price regression we find some increases in magnitude for new cars.

Overall, we would reach many the same conclusions about the effects of gasoline price changes if we aggregated gasoline prices within 4-digit ZIP code or within PADD instead of within DMAs.

7.4 Endogeneity

So far we have assumed that gasoline prices are uncorrelated with the error term in the market share and price specifications. In this subsection, we relax that assumption.

It seems unlikely that such a correlation would arise due to reverse causality; this is because U.S. gasoline prices depend on world oil prices and refinery margins and these are unlikely to be influenced by the yearly sales of cars in the U.S. However, there are other potential sources of endogeneity which may taint our coefficient estimates. First, there could be local variations in economic conditions that are correlated with local variations in gasoline prices. If the changes in economic conditions change what cars people buy or how much they are willing to spend on them, then our gasoline price coefficients will capture (in part) cyclical effects on car sales and prices. Second, gasoline tax changes might be endogenous to economic conditions which also affect car sales and prices. Third, changes in gasoline prices could cause income shocks in local areas (say, areas with refineries or with car plants) and these income shocks may drive car sales and prices.

One way to address the potential endogeneity of gasoline prices would be to use a more aggregate

[34] In our data, the median 4-digit ZIP code reports data from 11.5 stations on average over the months of the year, up from 3 for 5-digit ZIP codes.

measure of gasoline price; this would make it less likely that local shocks leads to correlation between gasoline prices and the error term in the market share and price specifications. The specification using PADD-level gasoline prices (described in the previous section and reported in Tables A-5 and A-6) does exactly this.

A second approach we take is to use world oil price as an instrument for gasoline prices at the PADD level. Clearly, world oil prices are correlated with regional fuel prices. At the same time, it seems highly unlikely that local or regional variation in economic conditions, gasoline tax changes, or income shocks would have a meaningful effect on world oil prices. To allow for some variation by PADD in the correlation with world oil prices, we use as instruments world oil prices interacted with PADD dummies.

The results of these two approaches are reported in Tables 4 and 5. For easier comparison we also report our original DMA-level specification (which uses OLS).

Table 4: OLS and IV results in market share specification†

	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
New Cars				
DMA-level gas prices OLS (original)	-.05** (.0049)	-.014** (.004)	-.0065* (.0029)	.07** (.005)
PADD-level gas prices OLS	-.05** (.0048)	-.014** (.0049)	-.0078* (.0032)	.072** (.0051)
PADD-level gas prices IV	-.054** (.0063)	.0027 (.0088)	-.012* (.0048)	.063** (.0098)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

We have already concluded that the OLS regression with PADD-level gasoline prices yields similar market share estimates but somewhat larger price estimates compared to the original OLS regression with DMA-level gasoline prices. The PADD-level IV estimates of the effect of gasoline prices on *market share* are similar to the PADD-level OLS estimates. We find, however, that the estimates of the effect of gasoline prices on car *prices* are generally larger in the PADD-level IV specification than in the PADD-level OLS specification. This can be seen in Table 5.

In summary, controlling for endogeneity suggests that our original specification may have underestimated the magnitude of the gasoline price effect on car prices. The effect of gasoline prices on market shares is largely unaffected by our two approaches to control for endogeneity.

7.5 Alternative market share specification

As our last robustness check we address potential limitations of the linear probability model we have used to estimate the effect of gasoline prices on markets shares. One might be concerned that

Table 5: OLS and IV results in price specification†

	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
New Cars				
DMA-level gas prices OLS (original specif.)	-236** (74)	-74+ (40)	6.9 (30)	127** (43)
PADD-level gas prices OLS	-352** (74)	-70+ (37)	45 (30)	163** (34)
PADD-level gas prices IV	-466** (24)	-49* (23)	57* (23)	250** (20)
Used Cars				
DMA-level gas prices OLS (original specif.)	-1073** (40)	-900** (58)	118* (53)	1766** (51)
PADD-level gas prices OLS	-1108** (40)	-936** (58)	124* (49)	1830** (48)
PADD-level gas prices IV	-1302** (22)	-1084** (23)	97** (22)	1995** (22)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

the linear probability model does not constrain the estimates in the market share regressions to add up to 1.[35] To address this we reestimate our basic MPG quartile specification (Equation 2) with a multinomial logit ("mlogit" in Stata) which estimates the probability that, conditional on purchase, a car falls into MPG Quartile 1, 2, 3, or 4 (all variables and controls are the same as those specified in Equation 2). In Table 6 we compare the gasoline price coefficients estimated for the linear probability model with the marginal effects in probability associated with a $1 increase in gasoline prices as estimated by the multinomial logit. Full estimation results are reported in Table A-7.

Table 6: mlogit marginal effects in market share specification†

	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
New Cars				
LPM	-.05** (.0049)	-.014** (.004)	-.0065* (.0029)	.07** (.005)
mlogit	-.053** (.0047)	-.012** (.0038)	-.004 (.0027)	.069** (.005)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs in parentheses. The SEs in the LPM are robust and clustered at the DMA level. The SE of the marginal effects in mlogit are derived using the Delta-method in Stata's "margins" command.

The table shows that the estimated marginal effects for the multinomial logit are nearly identical to those estimated in the linear probability model. We conclude that our quartile results do not depend on our use of the linear probability model.

[35] In fact, the market shares predicted by the linear probability results in Section 4.2 come very close to summing to 1, despite no constraint to do so.

8 Concluding remarks

In this paper we have estimated the effect of gasoline prices on the short-run equilibrium prices and sales of new and used cars of different fuel economies. We have used these estimates to address three primary questions.

The first question we addressed is how gasoline prices affect the fuel economy of new car purchases. We estimated that a $1 increase in the price of gasoline increases the market share of cars in the most fuel-efficient quartile by 20.5% and decreases the market share of cars in the least fuel efficient segment by 23.9%. We also estimated the effect of a $1 increase in gasoline prices on unit sales of new cars and found that sales in the most fuel-efficient quartile increased by about 9-13%, while sales in the least fuel-efficient quartile fell by about 25-27%. Using these estimates, we made a back-of-the-envelope calculation of the one-year effect of a $1 gasoline price increase on carbon dioxide emissions. We calculated that the effect from altering the new vehicle mix alone—not accounting for changing driving behavior or other effects—would be to reduce carbon dioxide emissions by 2.1 million tons.

The second question we addressed is whether the changes in equilibrium prices for new and used cars associated with changes in gasoline prices show evidence that consumers undervalue future gasoline costs of cars with different fuel economies relative to the prices of those cars. We estimated the effect of gasoline prices on the equilibrium prices of new cars and found that a $1 increase in the price of gasoline is associated with an increase of $363 in the average price of the most fuel-efficient quartile of cars relative to that of the least fuel-efficient quartile. For used cars, the estimated relative price difference is $2,839. We used these estimates to calculate the implied discount rates that would rationalize these estimated price changes with the corresponding increase in expected fuel costs that would be associated with a $1 gasoline price increase. Using several different assumptions about vehicle miles travelled, a range of assumptions about the elasticity of demand, and comparing the relative price differences between different quartiles, we find little evidence of consumer myopia. Many of our implied discount rates are near zero, most are in the single digits, and almost all are less than 20%.

The third question we addressed is whether gasoline prices have similar effects in new and used markets. We find very different effects, with a gasoline price increase having a much larger effect on prices in the used car market than in the new car market, but having a substantial effect on sales in the new car market. We argued that these differences can be explained by differences in supply between new and used markets. A general implication of this point is that when designing policy, policymakers need to take into account the market structure of the markets whose outcome they wish to influence.

Forecasting the effect of policy interventions such as a carbon tax or a gasoline tax increase on greenhouse gas emissions from non-commercial vehicles is challenging because there are many possible margins of adjustment. We believe that our investigation of the effect of gasoline price on market outcomes in new and used car markets is useful for understanding some of them. While our paper is not the only one to address these issues, we believe our paper's particular advantages are that it uses transaction data; that the data on prices and quantities and on new and used markets are from the same source; and that the flexible specifications used allow us to estimate parameters whose interpretation is not dependent on a particular model of the data generating process and which can be combined with a range of assumptions about related parameters in order to answer policy-relevant questions.

References

Allcott, H., and N. Wozny (2010): "Gasoline Prices, Fuel Economy, and the Energy Paradox," Discussion paper, MIT.

Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. von Haefen (2009): "Distributional and Efficiency Impacts of Increased US Gasoline Taxes," *American Economic Review*, 99(3), 667–699.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

Bresnahan, T. F., and V. A. Ramey (1994): "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, 109(3), 593–624.

Busse, M. R., and N. O. Keohane (2007): "Market Effects of Environmental Regulation: Coal, Railroads, and the 1990 Clean Air Act," *RAND Journal of Economics*, 38(4), 1159–1179.

Busse, M. R., J. Silva-Risso, and F. Zettelmeyer

Li, S., C. Timmins, and R. von Haefen (2009): “How Do Gasoline Prices Affect Fleet Fuel Economy?,” *American Economic Journal: Economic Policy*, 1(2), 113–137.

Lu, S. (2006): “Vehicle Survivability and Travel Mileage Schedules,” Working Paper DOT HS 809 952, NHTSA Technical Report.

Sallee, J. M., S. E. West, and W. Fan (2009): “Consumer Valuation of Fuel Economy: A Microdata Approach,” Discussion paper, National Tax Association Conference Proceedings.

Table 7: Summary Statistics

	New Cars						**Used Cars**					
Variable	N	Mean	Median	SD	Min	Max	N	Mean	Median	SD	Min	Max
GasolinePrice	1866366	2	1.8	0.67	0.77	4.8	1264092	2.1	1.9	0.69	0.77	4.7
MPG	1866366	23	22	5.7	10	65	1264092	22	22	4.8	9.9	65
Price	1866366	25515	23295	10876	2576	195935	1264092	15582	14468	8504	1	181000
DaysToTurn	1801528	58	27	78	1	3859	1211535	47	25	74	1	6055
ModelYear	1866366	2004	2004	2.5	1997	2008	1264092	2001	2001	3.5	1985	2008
CarAge	1866366	0.79	1	0.46	0	3	1264092	4	4	2.6	0	24
TradeValue	796759	8619	6800	8107	-5350	198000	495083	5295	3000	6081	-3402	150000
PctWhite	1866366	0.72	0.82	0.26	0	1	1264092	0.7	0.81	0.28	0	1
PctBlack	1866366	0.082	0.024	0.16	0	1	1264092	0.11	0.028	0.2	0	1
PctAsian	1866366	0.05	0.02	0.087	0	1	1264092	0.038	0.013	0.07	0	1
PctHispanic	1866366	0.12	0.053	0.18	0	1	1264092	0.13	0.051	0.19	0	1
PctLessHighSchool	1866366	0.15	0.12	0.13	0	1	1264092	0.18	0.14	0.13	0	1
PctCollege	1866366	0.38	0.36	0.19	0	1	1264092	0.33	0.3	0.18	0	1
PctManagment	1866366	0.16	0.15	0.082	0	1	1264092	0.14	0.13	0.075	0	1
PctProfessional	1866366	0.22	0.22	0.097	0	1	1264092	0.2	0.19	0.092	0	1
PctHeath	1866366	0.016	0.012	0.018	0	1	1264092	0.019	0.014	0.02	0	1
PctProtective	1866366	0.02	0.016	0.021	0	1	1264092	0.021	0.017	0.021	0	1
PctFood	1866366	0.041	0.035	0.031	0	1	1264092	0.046	0.04	0.033	0	1
PctMaintenance	1866366	0.028	0.021	0.029	0	1	1264092	0.032	0.025	0.031	0	1
PctHousework	1866366	0.027	0.024	0.021	0	1	1264092	0.028	0.025	0.022	0	1
PctSales	1866366	0.12	0.12	0.046	0	1	1264092	0.12	0.11	0.045	0	1
PctAdmin	1866366	0.15	0.15	0.053	0	1	1264092	0.16	0.16	0.054	0	1
PctConstruction	1866366	0.049	0.042	0.039	0	1	1264092	0.056	0.049	0.041	0	1
PctRepaitn	1866366	0.036	0.033	0.027	0	1	1264092	0.04	0.037	0.027	0	1
PctProduction	1866366	0.063	0.049	0.053	0	1	1264092	0.075	0.061	0.059	0	1
PctTransportation	1866366	0.051	0.044	0.038	0	1	1264092	0.059	0.053	0.039	0	1
Income	1866366	58130	53199	26246	0	200001	1264092	50826	46580	22231	0	200001
MedianHHSize	1866366	2.7	2.7	0.52	0	9.4	1264092	2.7	2.7	0.51	0	8.5
MedianHouseValue	1866366	178431	144800	131866	0	1000001	1264092	145079	121674	102666	0	1000001
VehPerHousehold	1866366	1.8	1.9	0.38	0	7	1264092	1.8	1.8	0.39	0	7
PctOwned	1866366	0.72	0.8	0.23	0	1	1264092	0.7	0.77	0.24	0	1
PctVacant	1866366	0.062	0.042	0.076	0	1	1264092	0.067	0.047	0.075	0	1
TravelTime	1866366	27	27	6.7	0.91	200	1264092	27	26	6.8	1	200
PctUnemployed	1866366	0.047	0.037	0.043	0	1	1264092	0.053	0.041	0.046	0	1
PctBadEnglish	1866366	0.044	0.016	0.078	0	1	1264092	0.045	0.014	0.08	0	1
PctPoverty	1866366	0.084	0.057	0.085	0	1	1264092	0.1	0.072	0.095	0	1
Weekend	1866366	0.25	0	0.44	0	1	1264092	0.26	0	0.44	0	1
EndOfMonth	1866366	0.25	0	0.43	0	1	1264092	0.21	0	0.41	0	1
EndOfYear	1866366	0.022	0	0.15	0	1	1264092	0.017	0	0.13	0	1
TradeOdometer	632689	71181	64224	44632	1	250000	385625	93150	89903	48514	1	250000

Table 8: New car fuel economy quartile cutoffs

Modelyear	MPG Q1 Mean	25th Pctile	MPG Q2 Mean	50th Pctile	MPG Q3 Mean	75th Pctile	MPG Q4 Mean
1999	16.0	18.3	20.1	22.2	23.3	24.7	27.7
[redacted]					23.1	24.3	27.6
2001	16.0	17.7	17.7	17.7	17.7	17.7	2001

[redacted]

NOT FOR PUBLICATION

Appendix: Additional Tables

Table A-1: New Cars: Market share results, fuel economy quartiles†

	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
FuelPrice	-.05**	-.014**	-.0065*	.07**
	(.0049)	(.004)	(.0029)	(.005)
PctLessHighSchool	.034*	.025*	-.023+	-.036*
	(.015)	(.01)	(.012)	(.017)
PctCollege	-.056**	.018	.017	.021
	(.012)	(.011)	(.011)	(.017)
Income	2.8e-08	3.5e-07**	2.4e-07*	-6.1e-07**
	(9.0e-08)	(9.4e-08)	(1.0e-07)	(1.1e-07)
MedianHHSize	.016**	.0061*	-.006	-.016**
	(.0032)	(.0026)	(.0047)	(.006)
MedianHouseValue	7.3e-08*	3.1e-08+	1.2e-08	-1.2e-07**
	(3.0e-08)	(1.6e-08)	(9.3e-09)	(4.1e-08)
VehiclePerHH	.049**	.0033	-.029**	-.023
	(.014)	(.0036)	(.0057)	(.018)
TravelTime	-.000048	-.00029**	-.00027*	.00061*
	(.0002)	(.000098)	(.00013)	(.00025)
Weekend	-.019**	-.0036*	-.0012	.024**
	(.0018)	(.0015)	(.0016)	(.0021)
EndOfMonth	.0057**	.004**	.0035**	-.013**
	(.001)	(.0011)	(.001)	(.0013)
EndOfYear	-.0037	-.0056*	-.0016	.011**
	(.0026)	(.0023)	(.0027)	(.0037)
Observations	1866008	1866008	1866008	1866008
R-squared	0.033	0.010	0.009	0.039

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.
Not reported: Region × year, region × month-of-year, and car age fixed effects. We also don't report house ownership, occupation, english proficiency, and race of buyers.

Table A-2: New Cars: Price results, fuel economy quartiles†

Variable	Coefficient/SE
FuelPrice*MGP Quart 1	-236** (74)
FuelPrice*MGP Quart 2	-74+ (40)
FuelPrice*MGP Quart 3	6.9 (30)
FuelPrice*MGP Quart 4	127** (43)
PctLessHighSchool	196** (75)
PctCollege	46 (53)
Income	0.0011** (0.00035)
MedianHHSize	25* (11)
MedianHouseValue	0.00017* (0.000078)
VehiclePerHH	-121** (37)
TravelTime	-0.27 (0.9)
Weekend	-12+ (6)
EndOfMonth	-135** (4.4)
EndOfYear	-79** (17)
Observations	1866008
R-squared	0.054

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses. Not reported: Region × year, region × month-of-year fixed effects, car type, and car age fixed effects. We also don't report house ownership, occupation, english proficiency, and race of buyers.

Table A-3: Used Cars: Market share results, fuel economy quartiles†

	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
FuelPrice	-0.016*	-0.019**	0.026*	0.01
	(0.0074)	(0.006)	(0.012)	(0.009)
PctLessHighSchool	0.011	0.039**	-0.029	-0.021
	(0.015)	(0.012)	(0.027)	(0.024)
PctCollege	-0.041*	0.0096	-0.012	0.044*
	(0.017)	(0.016)	(0.021)	(0.022)
Income	-5.2e-07**	3.7e-07**	6.0e-07**	-4.5e-07**
	(1.3e-07)	(7.2e-08)	(1.5e-07)	(1.4e-07)
MedianHHSize	0.014**	-0.002	-0.0057+	-0.0066
	(0.004)	(0.0031)	(0.0031)	(0.004)
MedianHouseValue	4.5e-08*	9.8e-08**	-4.2e-08*	-1.0e-07**
	(1.9e-08)	(1.1e-08)	(1.7e-08)	(2.9e-08)
VehiclePerHH	0.048**	-0.0089+	-0.037**	-0.0027
	(0.01)	(0.0052)	(0.0061)	(0.014)
TravelTime	0.00014	-0.0003*	-0.00036+	0.00052
	(0.00024)	(0.00014)	(0.00021)	(0.00033)
Weekend	-0.0053*	-0.0086**	0.0045	0.0095**
	(0.0021)	(0.0025)	(0.0028)	(0.0028)
EndOfMonth	0.0036	-0.00073	0.00088	-0.0037+
	(0.0024)	(0.0013)	(0.0023)	(0.0022)
EndOfYear	-0.012**	0.004	0.004	0.0045
	(0.0039)	(0.0041)	(0.0042)	(0.0045)
Observations	1263940	1263940	1263940	1263940
R-squared	0.030	0.013	0.015	0.023

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

Not reported: Region × year, region × month-of-year, and car age fixed effects. We also don't report house ownership, occupation, english proficiency, and race of buyers.

Table A-4: Used Cars: Price results, fuel economy quartiles†

Variable	Coefficient/SE
FuelPrice*MGP Quart 1	-1073**
	(40)
FuelPrice*MGP Quart 2	-900**
	(58)
FuelPrice*MGP Quart 3	118*
	(53)
FuelPrice*MGP Quart 4	1766**
	(51)
PctLessHighSchool	120
	(95)
PctCollege	89
	(80)
Income	.0026**
	(.00074)
MedianHHSize	-38
	(25)
MedianHouseValue	.00065**
	(.00016)
VehiclePerHH	-178**
	(28)
TravelTime	-1.8+
	(1.1)
Weekend	122**
	(11)
EndOfMonth	-106**
	(7.1)
EndOfYear	-23
	(23)
Observations	1263857
R-squared	0.631

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses. Not reported: Region × year, region × month-of-year fixed effects, car type, and car age fixed effects. We also don't report house ownership, occupation, english proficiency, and race of buyers.

Table A-5: Effect of gasoline price aggregation in market share regression†

New Cars				
Gas Price Aggregation	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
4-digit ZIP	-.038**	-.016**	-.0041	.058**
	(.0081)	(.0039)	(.0035)	(.0084)
DMA	-.05**	-.014**	-.0065*	.07**
(original specification)	(.0049)	(.004)	(.0029)	(.005)
PADD	-.05**	-.014**	-.0078*	.072**
	(.0048)	(.0049)	(.0032)	(.0051)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

Table A-6: Effect of gasoline price aggregation in price regression†

New Cars				
Gas Price Aggregation	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
4-digit ZIP	-219**	-71*	5	137**
	(69)	(33)	(30)	(43)
DMA	-236**	-74+	6.9	127**
(Base Case)	(74)	(40)	(30)	(43)
PADD	-352**	-70+	45	163**
	(74)	(37)	(30)	(34)
Used Cars				
Gas Price Aggregation	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
4-digit ZIP	-1080**	-921**	83	1615**
	(39)	(59)	(74)	(59)
DMA	-1073**	-900**	118*	1766**
(Base Case)	(40)	(58)	(53)	(51)
PADD	-1108**	-936**	124*	1830**
	(40)	(58)	(49)	(48)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

Table A-7: mlogit estimates (baseline is MPG Quartile 4) †

	New Cars		
	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3
FuelPrice	-.48**	-.27**	-.23**
	(.036)	(.029)	(.018)
PctLessHighSchool	.25*	.23**	-.036
	(.11)	(.084)	(.076)
PctCollege	-.38**	-.0079	-.046
	(.11)	(.1)	(.062)
Income	3.1e-06**	4.2e-06**	3.7e-06**
	(7.9e-07)	(6.6e-07)	(5.4e-07)
MedianHHSize	.13**	.071**	.015
	(.029)	(.024)	(.034)
MedianHouseValue	7.6e-07*	5.4e-07*	4.2e-07*
	(3.1e-07)	(2.1e-07)	(1.7e-07)
VehiclePerHH	.32**	.089	-.054
	(.12)	(.065)	(.068)
TravelTime	-.0029+	-.0031**	-.0027**
	(.0017)	(.001)	(.00098)
Weekend	-.17**	-.089**	-.076**
	(.014)	(.012)	(.0096)
EndOfMonth	.069**	.059**	.055**
	(.0077)	(.008)	(.0064)
EndOfYear	-.052*	-.061**	-.041*
	(.021)	(.02)	(.019)
Observations	1866008		
Log pseudolikelihood	-2482493.4		

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (clustered at the DMA level) in parentheses.

Table A-8: New Cars: Market share (quartile) results by gasoline price levels†

	New Cars			
	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
GasolinePrice (<1.5 dollar)	-.049**	-.012*	-.011*	.071**
	(.0067)	(.0056)	(.0046)	(.0058)
GasolinePrice (1.5-2.5 dollars)	-.05**	-.0092+	-.0087*	.068**
	(.0062)	(.0049)	(.0041)	(.0054)
GasolinePrice (2.5-3.5 dollars)	-.05**	-.011*	-.0067+	.067**
	(.0055)	(.0043)	(.0034)	(.0047)
GasolinePrice (>3.5 dollars)	-.051**	-.013*	-.014**	.078**
	(.0059)	(.0058)	(.0035)	(.0075)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.
This table only reports the coefficients on gasoline prices.

Table A-9: New and Used Cars: Price results by gasoline price levels†

	New Cars, MPG Quartiles	Used Cars, MPG Quartiles
GasolinePrice(< 1.5)*MPG Quart 1	-266** (88)	-1366** (75)
GasolinePrice(1.5-2.5)*MPG Quart 1	-223* (100)	-1201** (62)
GasolinePrice(2.5-3.5)*MPG Quart 1	-225** (82)	-1112** (52)
GasolinePrice(> 3.5)*MPG Quart 1	-384** (75)	-1350** (65)
GasolinePrice(< 1.5)*MPG Quart 2	-130+ (66)	-1237** (107)
GasolinePrice(1.5-2.5)*MPG Quart 2	-110+ (64)	-1156** (85)
GasolinePrice(2.5-3.5)*MPG Quart 2	-110* (53)	-1009** (78)
GasolinePrice(> 3.5)*MPG Quart 2	-81 (49)	-954** (89)
GasolinePrice(< 1.5)*MPG Quart 3	-42 (41)	476** (92)
GasolinePrice(1.5-2.5)*MPG Quart 3	-57 (41)	417** (87)
GasolinePrice(2.5-3.5)*MPG Quart 3	-36 (34)	315** (74)
GasolinePrice(> 3.5)*MPG Quart 3	.22 (33)	227* (91)
GasolinePrice(< 1.5)*MPG Quart 4	82 (64)	2654** (167)
GasolinePrice(1.5-2.5)*MPG Quart 4	28 (70)	2454** (158)
GasolinePrice(2.5-3.5)*MPG Quart 4	58 (56)	2217** (109)
GasolinePrice(> 3.5)*MPG Quart 4	136** (47)	1903** (88)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.
This table only reports the coefficients on gasoline prices.

Table A-10: New Cars: Market share (quartile) results by gasoline price trends†

New Cars Results	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
GasolinePrice (3 months up)	-.052** (.0046)	-.016** (.0039)	-.0015 (.0028)	.07** (.0046)
GasolinePrice (3 months mixed)	-.051** (.005)	-.016** (.0042)	.00048 (.003)	.067** (.0049)
GasolinePrice (3 months down)	-.054** (.0054)	-.019** (.0047)	.004 (.0034)	.069** (.0053)

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.
This table only reports the coefficients on gasoline prices.

Table A-11: New and Used Cars: Price results by gasoline price trends[†]

	New Cars, MPG Quartiles	Used Cars, MPG Quartiles
GasolinePrice(3 mo up)*MPG Quart 1	-290** (80)	-1025** (42)
GasolinePrice(3 mo mixed)*MPG Quart 1	-312** (90)	-1094** (44)
GasolinePrice(3 mo down)*MPG Quart 1	-365** (99)	-1202** (51)
GasolinePrice(3 mo up)*MPG Quart 2	-110* (46)	-855** (61)
GasolinePrice(3 mo mixed)*MPG Quart 2	-135** (51)	-920** (63)
GasolinePrice(3 mo down)*MPG Quart 2	-143* (56)	-1003** (70)
GasolinePrice(3 mo up)*MPG Quart 3	-34 (37)	146** (56)
GasolinePrice(3 mo mixed)*MPG Quart 3	-50 (41)	161** (59)
GasolinePrice(3 mo down)*MPG Quart 3	-77+ (46)	168** (60)
GasolinePrice(3 mo up)*MPG Quart 4	98+ (52)	1768** (52)
GasolinePrice(3 mo mixed)*MPG Quart 4	86 (60)	1945** (60)
GasolinePrice(3 mo down)*MPG Quart 4	82 (64)	2081** (64)

[*] significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.
This table only reports the coefficients on gasoline prices.

Table A-12: New and Used Cars: Inventory results[†]

	New Cars			Used Cars		
Variable	Coefficient (SE)	DTT sample mean	% Change in DTT	Coefficient (SE)	DTT sample mean	% Change in DTT
GasolinePrice * Quart. 1 (least fuel-efficient)	12** (2.3)	68.3	17.57%	.71 (.63)	47.8	1.49%
GasolinePrice * Quart. 2	2.3** (.89)	61.4	3.75%	1.5** (.57)	47.3	3.17%
GasolinePrice * Quart. 3	.56 (.9)	57.2	0.98%	.12 (.63)	49.1	0.24%
GasolinePrice * Quart. 4 (most fuel-efficient)	-5.4** (.88)	50.2	-10.76%	-.9 (.6)	45.4	-1.98%

[†] This table only reports the coefficients on gasoline prices. The full specification for both new and used cars is:

$$DTT_{irdjt} = \omega_0 + \omega_{\mathbf{1}}(\text{GasolinePrice}_{it} \cdot \text{MPG Quartile}_j) + \omega_2\text{Demog}_{it} + \omega_3\text{PurchaseTiming}_{jt} + \delta_{dj} + \tau_{rt} + \mu_{rt} + \nu_{ijt}$$

where DTT_{irdjt} measures days to turn for transaction i in region r at dealer d on date t for car j. We use the same extensive set of controls we have used in the market share specification (see page 10) with one addition. To control for the fact that different dealerships may have different inventory policies we include car type $\times$ dealer fixed effects (δ_{dj}).
[*] significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.

Table A-13: New and Used Cars: Actual cash value of trade-in†

	New Car Trade-in Actual Cash Value	Used Car Trade-in Actual Cash Value	Used Car Transaction Prices
GasolinePrice*MPG Quart 1	-1177**	-995**	-1073**
	(56)	(27)	(40)
GasolinePrice*MPG Quart 2	-887**	-588**	-900**
	(39)	(45)	(58)
GasolinePrice*MPG Quart 3	174**	205**	118*
	(46)	(43)	(53)
GasolinePrice*MPG Quart 4	1275**	778**	1766**
	(37)	(39)	(51)

† This table only report the coefficients on gasoline prices. The full specification for both new and used cars in columns 1 and 2 is:

$$\text{ACV}_{ilrt} = \beta_0+\beta_1\text{GasolinePrice}_{it}\cdot\text{MPG Quartile}_l+\beta_2\text{Odometer}_{ilt}+\beta_3\text{Demog}_{it}+\beta_4\text{PurchaseTiming}_{jt}+\delta_l+\tau_{rt}+\mu_{rt}+\xi_{ilrt}$$

where ACV_{ilrt} is the actual cash value booked in transaction i for trade-in car l in region r on date t. We add a new control variable to this specification, which is the odometer reading of the trade-in car; cars with higher odometer readings have experienced greater depreciation and should be booked at lower actual cash values, all else equal. In the specification, we include the demographic characteristics of the buyer; these should not have a direct effect on the average cash value, but may be correlated with unobservable quality characteristics ("wear and tear") of the trade-in car. We also include the purchase timing of the transaction, in case cars are assigned different actual cash values on, for example, weekend days, when there is typically higher transaction volume. Finally, we include detailed "car type" fixed effects for the trade-in, as well as our region-specific year and region-specific month-of-year fixed effects.

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the DMA level) in parentheses.