



EI @ Haas WP 201R

Are Consumers Myopic? Evidence from New and Used Car Purchases

Meghan R. Busse, Christopher R. Knittel and Florian Zettelmeyer
December 2009

**Revised version published in
American Economic Review
103(1), February 2013.**

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board.

© 2009 by Meghan R. Busse, Christopher R. Knittel and Florian Zettelmeyer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

<http://ei.haas.berkeley.edu>

Are Consumers Myopic? Evidence from New and Used Car Purchases

By MEGHAN R. BUSSE, CHRISTOPHER R. KNITTEL AND FLORIAN
ZETTELMEYER*

We investigate whether car buyers are myopic about future fuel costs. We estimate the effect of gasoline prices on short-run equilibrium prices of cars of different fuel economies. We then compare the implied changes in willingness-to-pay to the associated changes in expected future gasoline costs for cars of different fuel economies in order to calculate implicit discount rates. Using different assumptions about annual mileage, survival rates, and demand elasticities, we calculate a range of implicit discount rates similar to the range of interest rates paid by car buyers who borrow. We interpret this as showing little evidence of consumer myopia.

According to EPA estimates, gasoline combustion by passenger cars and light-duty trucks is the source of 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, so does interest in designing policy instruments that will reduce carbon emissions from this source. In order to be effective, any such policy must reduce gasoline consumption, since carbon emissions are essentially proportional to the amount of gasoline used. The major policy instrument that has been used so far to influence gasoline consumption in the U.S. has been the Corporate Average Fuel Efficiency (CAFE) standards (Pinelopi K. Goldberg (1998), Mark R. Jacobsen (2010)). Some economists, however, contend that changing the incentives to use gasoline—by increasing its price—would be a preferable approach. This is because changing the price of gasoline has the potential to influence both what cars people buy and how much people drive.

* Busse: Northwestern University and NBER, Kellogg School of Management, 2001 Sheridan Rd, Evanston, IL 60208, m-busse@kellogg.northwestern.edu. Knittel: MIT Sloan and NBER, MIT Sloan School of Management, 100 Main Street, Cambridge, MA 02142-1347, knittel@mit.edu. Zettelmeyer: Northwestern University and NBER, Kellogg School of Management, 2001 Sheridan Rd, Evanston, IL 60208, f-zettelmeyer@kellogg.northwestern.edu. We are grateful for helpful comments from Hunt Allcott, Eric Anderson, John Asker, Max Auffhammer, Severin Borenstein, Tim Bresnahan, Igal Hendel, Ryan Kellogg, Aviv Nevo, Sergio Rebelo, Jorge Silva-Risso, Scott Stern, and three anonymous referees. We thank seminar participants at Brigham Young University, the Chicago Federal Reserve Bank, Cornell, Harvard, Illinois Institute of Technology, Iowa State, MIT, Northwestern, Ohio State, Purdue, Texas A&M, Triangle Resource and Environmental Economics seminar, UC Berkeley, UC Irvine, University of British Columbia, University of Chicago, University of Michigan, University of Rochester, University of Toronto, and Yale. We also thank participants at the ASSA, Milton Friedman Institute Price Dynamics Conference, NBER IO, EEE, and Price Dynamics conferences, and the National Tax Association. 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.

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

This paper addresses a question that is crucial for assessing whether a gasoline price related policy instrument (such as an increased gasoline tax or a carbon tax) could influence what cars people buy: How sensitive are consumers to expected future gasoline costs when they make new car purchases? More precisely, how much does an increase in the price of gasoline affect the willingness-to-pay of consumers for cars of different fuel economies? If consumers are very myopic, meaning that their willingness-to-pay for a car is little affected by changes in the expected future fuel costs of using that car, then a gasoline price instrument will not influence their choices very much and will not be sufficient to achieve the first-best outcome in the presence of an externality. This condition is not unique to the case of gasoline consumption. Jerry A. Hausman (1979) was the first to investigate whether consumers are myopic when purchasing durable goods that vary in energy costs. More generally, this is an example of the quite obvious point that a policy must influence something that consumers pay attention to in order to actually affect the choices consumers make.

Our analysis proceeds in two steps. First, we estimate 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 transaction prices, market shares, and sales for new and used cars of different fuel economies. We find that a \$1 change in the gasoline price is associated with a very large change in relative prices of *used* cars of different fuel economies—a difference of \$1,945 in the relative price of the highest fuel economy and lowest fuel economy quartile of cars. For *new* cars, the predicted relative price difference is much smaller—a \$354 difference between the highest and lowest fuel economy 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 21.1 percent increase in the market share of the highest fuel economy quartile of cars and a 27.1 percent decrease in the market share of the lowest fuel economy quartile of cars. These estimates become the building blocks for our next step.

In our second step, we use the estimated effect of gasoline prices on prices and quantities in new and used car markets to learn about how consumers trade off the up-front capital cost of a car and the ongoing usage cost of the car. We estimate a range of implicit discount rates under a range of assumptions about demand elasticities, vehicle miles travelled, and vehicle survival probabilities. We find little evidence that consumers “undervalue” future gasoline costs when purchasing cars. The implicit discount rates we calculate correspond reasonably closely to interest rates that customers pay when they finance their car purchases.

This paper proceeds as follows. In the next section, we position this paper within the related literature. In Section II we describe the data we use for the analysis in this paper. In Section III we estimate the effect of gasoline prices on equilibrium prices, market shares, and unit sales in new and used car markets. In Section IV we use the results estimated in Section III to investigate whether

consumers are myopic, meaning whether they undervalue expected future fuel costs relative to the up-front prices of cars of different fuel economics. Section V checks the robustness of our estimated results. Section VI offers some concluding remarks.

I. 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. Javier Donna (2011), Goldberg (1998), and Jonathan E. Hughes, Christopher R. Knittel and Daniel Sperling (2008) investigate three different measures of driving responses to gasoline prices. Donna investigates how public transportation utilization is affected by gasoline prices, Goldberg estimates the effect of gasoline prices on vehicle miles travelled, and Hughes *et al.* investigate monthly gasoline consumption.

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 they sell or by changing fuel technologies to hybrid or electric vehicles. Jacob Gramlich (2009) investigates such manufacturer responses by relating year-to-year changes in the MPG of individual car models to gasoline prices.

This paper belongs to a set of papers that examine a question with a time horizon in between this two extremes: 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 of gasoline prices taking as given 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 sales or market shares and some that study the effect of gasoline prices on car prices.²

A. Gasoline prices and car quantities

Two noteworthy papers that address the effect of gasoline prices on car quantities are Thomas Klier and Joshua Linn (2010) and Shanjun Li, Christopher Timmins and Roger von Haefen (2009). Although the two papers address similar

²There is a very large literature (reaching back almost half a century) that has investigated the effect of gasoline prices on car choices, the car industry, or vehicles miles travelled, and that has estimated the elasticity of demand for gasoline. In addition to the papers described in detail in the next section, other related papers include Ake G. Blomqvist and Walter Haessel (1978), Rodney L. Carlson (1978), Makoto Ohta and Zvi Griliches (1986), John S. Greenlees (1980), James W. Sawhill (2008), Asher Tishler (1982), and Sarah E. West (2007).

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 high fuel economy new cars and the survival probabilities of high fuel economy used cars, while decreasing the sales of low fuel economy new cars and the survival probabilities of low fuel economy used cars.

B. 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.

James 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 generally finds that used car prices do adjust to gasoline prices, by about one-third to one-half the amount that would fully reflect the change in the gasoline cost, although some specifications find full adjustment. This, he concludes, indicates some degree of myopia. Lutz Kilian and Eric R. 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.

Hunt Allcott and Nathan Wozny (2011) 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, consumers equally value a \$1 change in the purchase price of a vehicle and a 72-cent 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.

James M. Sallee, Sarah E. West and Wei 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 from others 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 percent of the change in fuel costs, depending on the discount rate used.

Frank Verboven (2002) implements a similar approach to the papers described above but using data on European consumers' choices to buy either a gasoline- or a diesel-powered car. This choice also involves a trade-off between the upfront price for a car and the car's future fuel cost, but with variation over different fuels rather than over time in the price of a single fuel. He estimates implicit

discount rates of approximately 11.5 percent, a value that is close to or slightly above contemporaneous interest rates.

Goldberg (1998) approaches the question of consumer myopia in a completely different way. 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.

C. 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 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 estimate reduced form parameters, which differentiates from some (although not all) of the papers above.

Transactions data: As described in more detail in Section II, we observe individual transactions, and observe a variety of characteristics about each transaction, such as location, purchase timing, detailed car characteristics, and demographic characteristics of buyers. This allows us to use extensive controls in our regressions, reducing the chances that our results arise from selection issues or aggregation over heterogeneous regions, time periods, or car models. We are also able to observe transactions prices for cars (rather than list prices) and we are able to subtract off manufacturer rebates and credits for trade-in cars.

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 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 (2011) because the effect is actually different for new and used cars, or for some other reason.

Reduced form 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 then 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 (2011), and Goldberg (1998) do).

We choose to estimate reduced form parameters. In order to interpret these parameters with respect to consumer myopia, we have to make assumptions similar

to what must be assumed in the structural approach; namely, how many miles the owner will drive each year, how long the car will last, and what the buyer's expectation of future gasoline price is. The advantage of this approach is that a reader of this paper can create his or her own estimate of consumer myopia using alternative assumptions about driving behavior, gasoline prices, or vehicle life. The disadvantage is that reduced form parameters cannot be used in policy simulations or counterfactuals the way structural parameters can.

II. Data

We combine several types of data for the analysis. Our main data contain information on automobile transactions from a sample of about 20 percent of all new car dealerships in the U.S. from January 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, 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. We measure each car model's fuel economy with the Environmental Protection Agency (EPA)'s "Combined Fuel Economy" which is a weighted geometric average of the EPA Highway (45 percent) and City (55 percent) Vehicle Mileage. As shown in Figure 1, the average MPG of models available for sale in the United States declined slowly in the first part of our sample period, then increased in the latter part.³ 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.⁴

We also used gasoline price data from OPIS (Oil Price Information Service) which cover the same time period. 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.

We aggregate the station-level data to obtain average prices for basic grade gasoline in each local market, which we define as Nielsen Designated Market Areas, or "DMAs" for short. There are 210 DMAs. Examples are "San Francisco-Oakland-

³In 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.

⁴While *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 Christopher R. Knittel (2011) 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

San Jose, CA,” “Charlotte, NC,” and “Ft. Myers-Naples, FL.” We aggregate station-level data to DMAs instead of to ZIP-codes for two reasons. First, we only observe a small number of stations per ZIP-code, which may make a ZIP-code average prone to measurement error.⁵ 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. Later we investigate the sensitivity of our results to different aggregations of gasoline prices (see section V.C).

Figure 2 gives a sense of the variation in the gasoline price data. The figure graphs monthly national average gasoline prices and shows substantial intertemporal variation 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. Figure 3 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.

To create our final dataset, we draw a 10 percent random sample of all transactions.⁶ After combining the three datasets this leaves us with a new car dataset of 1,863,403 observations and a used car dataset of 1,096,874 observations. Tables 1 and 2 present summary statistics for the two datasets.

⁵In our data, the median ZIP code reports data from 3 stations on average over the months of the year. More than 25 percent of ZIP-codes have only one station reporting.

⁶The 10 percent sample is necessary to allow for estimation of specifications with multiple sets of high-dimensional fixed effects, including fixed effect interactions, that we use later in the paper.




FIGURE 2. MONTHLY AVERAGE GASOLINE PRICES (NATIONAL)

III. Estimation and results

In this section we estimate the short-run equilibrium effects of changes in gasoline prices on the transaction prices, market shares, and unit sales of cars of different fuel economics. We separate our analysis by new and used markets. We will use the results estimated in this section to investigate, in Section IV, whether car buyers “undervalue” future fuel costs.

A. Specification and variables for car price results

At the most basic level, our approach is to model the effect of covariates on short-run equilibrium price and (in a later subsection) quantity outcomes. For the car industry, the short-run horizon is several months to a few years. During this time frame, a manufacturer can alter both price and production quantities, but its offering of models is predetermined, its model-specific capacity is largely fixed, and a number of input arrangements are fixed (labor contracts, in particular). While some of these aspects become more flexible over a year or two (models can be tweaked, some capacity can be altered), only over a long-run horizon (four years or more), can a manufacturer introduce fundamentally different models into its product offering.

We use a reduced form approach. In generic terms, this means regressing observed car prices (P) on demand covariates (X^D) and supply covariates (X^S):

$$(1) \quad P = \alpha_0 + \alpha_1 \mathbf{X}^D + \alpha_2 \mathbf{X}^S + \nu$$

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




FIGURE 3. MONTHLY AVERAGE GASOLINE PRICES (BY DMA)

both taken into account.

Our demand covariates are gasoline prices (the chief variable of interest), customer demographics, and variables describing the timing of the purchase, all described in greater detail below. We also include region-specific year fixed effects, region-specific month-of-year fixed effects, and detailed “car type” 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 region-specific month-of-year fixed effects that are already included in the specification. Furthermore, our interactions with executives responsible for short- to medium-run manufacturing and pricing decisions for automobiles indicate that, in practice, these decisions are not made on the basis of small changes to manufacturing costs.

We can write the specification we estimate more precisely as:

$$(2) \quad P_{irjt} = \lambda_0 + \lambda_1(\text{GasolinePrice}_{it} \cdot \text{MPG Quartile}_j) + \lambda_2 \mathbf{Demog}_{it} + \lambda_3 \mathbf{PurchaseTiming}_{jt} + \delta_j + \tau_{rt} + \mu_{rt} + \epsilon_{ijt}.$$

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.⁷

We make two adjustments in order to make P_{irjt} capture the customer’s total wealth outlay for the car. 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

⁷Dealer-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.

TABLE 1—SUMMARY STATISTICS (NEW CARS)

Variable	N	Mean	Median	SD	Min	Max
GasolinePrice	1,863,403	2	1.8	.67	.82	4.6
MPG	1,863,403	22	22	5.7	10	65
Price	1,863,403	25515	23,295	10,874	2,576	195,935
ModelYear	1,863,403	2004	2004	2.5	1997	2008
CarAge	1,863,403	.79	1	.46	0	3
TradeValue*	795,457	8,619	6,794	8,107	0	198,000
PctWhite	1,863,403	.72	.82	.26	0	1
PctBlack	1,863,403	.082	.024	.16	0	1
PctAsian	1,863,403	.05	.02	.087	0	1
PctHispanic	1,863,403	.12	.053	.18	0	1
PctLessHighSchool	1,863,403	.15	.12	.13	0	1
PctCollege	1,863,403	.38	.36	.19	0	1
PctManagment	1,863,403	.16	.15	.082	0	1
PctProfessional	1,863,403	.22	.22	.097	0	1
PctHeath	1,863,403	.016	.012	.018	0	1
PctProtective	1,863,403	.02	.016	.021	0	1
PctFood	1,863,403	.041	.035	.031	0	1
PctMaintenance	1,863,403	.028	.021	.029	0	1
PctHousework	1,863,403	.027	.024	.021	0	1
PctSales	1,863,403	.12	.12	.046	0	1
PctAdmin	1,863,403	.15	.15	.054	0	1
PctConstruction	1,863,403	.049	.042	.039	0	1
PctRepair	1,863,403	.036	.033	.027	0	1
PctProduction	1,863,403	.063	.049	.053	0	1
PctTransportation	1,863,403	.05	.044	.037	0	1
Income	1,863,403	58,110	53,188	26,274	0	200,001
MedianHHSIZE	1,863,403	2.7	2.7	.52	0	9
MedianHouseValue	1,863,403	178,306	144,700	131,956	0	1,000,001
VehPerHousehold	1,863,403	1.8	1.9	.39	0	7
PctOwned	1,863,403	.72	.8	.23	0	1
PctVacant	1,863,403	.063	.042	.078	0	1
TravelTime	1,863,403	27	27	6.8	0	200
PctUnemployed	1,863,403	.047	.037	.043	0	1
PctBadEnglish	1,863,403	.044	.016	.078	0	1
PctPoverty	1,863,403	.084	.057	.085	0	1
Weekend	1,863,403	.25	0	.44	0	1
EndOfMonth	1,863,403	.25	0	.43	0	1
EndOfYear	1,863,403	.022	0	.15	0	1

* This row summarizes the trade value for the subset of transactions that use trade-ins.

from the purchase price any profit or add to the purchase price any loss the customer made on his or her trade-in. Dealers are willing to trade off profits made on the new vehicle transaction and profits made on the trade-in transaction, including being willing to lose money on the trade-in.⁸ If 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 adding such a loss to the negotiated (contract) price we adjust the price to include the value of this in-kind payment. In Equation 2, P_{irjt} is the above-defined price for transaction i in region r on date t for car j .

⁸See Meghan R. Busse and Jorge Silva-Risso (2010) for further discussion of the correlation between dealers' profit margins on new cars vs. trade-ins.

TABLE 2—SUMMARY STATISTICS (USED CARS)

Variable	N	Mean	Median	SD	Min	Max
GasolinePrice	1,096,874	2	1.8	.69	.82	4.6
MPG	1,096,874	22	22	4.7	9.8	65
Price	1,096,874	15,637	14,495	8,281	1	173,000
ModelYear	1,096,874	2001	2001	3.5	1985	2008
CarAge	1,096,874	3.9	4	2.4	0	24
TradeValue*	435,813	5,233	3,000	5,992	0	150,000
PctWhite	1,096,874	.7	.81	.28	0	1
PctBlack	1,096,874	.11	.028	.2	0	1
PctAsian	1,096,874	.038	.013	.07	0	1
PctHispanic	1,096,874	.13	.05	.19	0	1
PctLessHighSchool	1,096,874	.18	.14	.13	0	1
PctCollege	1,096,874	.33	.29	.18	0	1
PctManagment	1,096,874	.14	.13	.074	0	1
PctProfessional	1,096,874	.2	.19	.092	0	1
PctHeath	1,096,874	.019	.014	.02	0	1
PctProtective	1,096,874	.021	.017	.021	0	1
PctFood	1,096,874	.046	.04	.033	0	1
PctMaintenance	1,096,874	.032	.025	.031	0	1
PctHousework	1,096,874	.028	.025	.022	0	1
PctSales	1,096,874	.12	.11	.044	0	1
PctAdmin	1,096,874	.16	.16	.054	0	1
PctConstruction	1,096,874	.056	.049	.041	0	1
PctRepair	1,096,874	.04	.037	.027	0	1
PctProduction	1,096,874	.075	.061	.059	0	1
PctTransportation	1,096,874	.059	.053	.039	0	1
Income	1,096,874	50,684	46,556	22,031	0	200,001
MedianHHSIZE	1,096,874	2.7	2.7	.51	0	8.5
MedianHouseValue	1,096,874	145,545	121,997	102,923	0	1,000,001
VehPerHousehold	1,096,874	1.8	1.8	.39	0	7
PctOwned	1,096,874	.69	.76	.24	0	1
PctVacant	1,096,874	.067	.048	.076	0	1
TravelTime	1,096,874	27	26	6.9	0	200
PctUnemployed	1,096,874	.053	.041	.046	0	1
PctBadEnglish	1,096,874	.045	.014	.079	0	1
PctPoverty	1,096,874	.1	.072	.095	0	1
Weekend	1,096,874	.25	0	.44	0	1
EndOfMonth	1,096,874	.21	0	.41	0	1
EndOfYear	1,096,874	.017	0	.13	0	1

* This row summarizes the trade value for the subset of transactions that use trade-ins.

We estimate how gasoline prices affect the transaction prices paid for cars of different fuel economies. One might think that higher gasoline prices, by making car ownership more expensive, should lead to lower negotiated prices for all cars. Note, however, that cars do not increase uniformly in fuel cost: a compact car has lower fuel costs than an SUV at every gasoline price, but as gasoline price rises, its fuel cost advantage *relative* to the SUV actually *rises*. If enough people continue to want to own cars, even when gasoline prices increase, then higher gasoline prices may lead to increased demand for high fuel economy cars and decreased demand for low fuel economy cars, and consequently to the transaction price rising for the highest fuel economy cars and falling for the lowest fuel economy cars. To capture this, we estimate separate coefficients for the *GasolinePrice* variable depending on the fuel economy quartile into which car j falls. 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.⁹ 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 A1 in the online appendix reports the quartile cutoffs and mean MPG within quartile for all years of the sample.

We use an extensive set of controls. First, we control for a wide range of demographic variables ($Demog_{it}$) using data from the 2000 Census: income, house value and ownership, household size, vehicles per household, education, occupation, average travel time to work, English proficiency, and race of buyers.¹⁰ We use data at the level of “block groups,” which, on average, contain about 1100 people. We also control for a series of variables that describe purchase timing ($PurchaseTiming_{jt}$): $EndOfYear$ is a dummy variable that equals 1 if the car was sold within the last 5 days of the year; $EndOfMonth$ is a dummy variable that equals 1 if the car was sold within the last 5 days of the month; $WeekEnd$ is a dummy variable 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. For new cars, $PurchaseTiming_{jt}$ 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. For used cars, $PurchaseTiming_{jt}$ includes a flexible function of the car’s odometer, described in more detail below, which controls for depreciation over time.

We include year, τ_{rt} , and month-of-year, μ_{rt} , fixed effects corresponding to when the purchase was made. Both year and month-of-year fixed effects are allowed to vary by the geographic region (34 throughout the U.S.) in which the car was sold.¹¹ The identifying variation we use is therefore variation within a year and region that differs from the average pattern of seasonal variation within that region. 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 V.A. We also control for detailed characteristics of the vehicle purchased by including “car type” fixed effects (δ_j). 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.)

The coefficients of primary interest will be the coefficients on the monthly, DMA-level gasoline price measure. This variable contains both cross-sectional and intertemporal variation. Cross-sectional variation arises from factors such as

⁹We 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.

¹⁰Demographic variables do not change over time in our data.

¹¹See Table A13 in the online appendix for a list of regions and the DMAs within each region.

differences across locations in transportation costs (or transportation capacity), variation in the degree of market power, and differences in the costs of required gasoline formulations. Intertemporal variation in gasoline prices arises mostly from differences in the world price of oil. Because we use year and month-of-year fixed effects, both interacted with region, the component of the intertemporal variation that identifies our results will be within year variation in gasoline prices that differs from the typical seasonal pattern of variation for the region. The component of cross-sectional variation that will identify our results will be persistent differences among DMAs within a region in factors such as transportation costs or market power, as well as month-to-month fluctuations in the gasoline price differentials between DMAs or month-to-month fluctuations in the gasoline price differentials between regions that differs from the typical seasonal pattern.¹² By using a variable that contains both cross-sectional and intertemporal variation, our specification assumes that car buyers respond equally to both components of variation. In other words, we assume that intertemporal variation arising from changes in world oil prices and fluctuations in local market conditions both matter to car buyers in determining their forecasts of future gasoline prices, and in driving their decisions about what vehicles to buy. (In section V.C we consider specifications that use more geographically aggregated measures of gasoline price, one a national price series and another that varies by five regions of the country defined by Petroleum Administration for Defense Districts (PADDs).) A second, less obvious assumption implied by this specification is that vehicles are not traded across regions in response to gasoline price differentials.

Before describing the results, we note that our estimates should be interpreted as estimates of the short-run effects of gasoline prices, meaning effects on prices, market shares, or sales over the time horizon in which manufacturers would be unable to change the configurations of cars they offer in response to gasoline price changes, a period of several months to a few years. 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.¹³ 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 prices or sales. The short-run estimates are nevertheless useful, we believe, for two reasons. First, the short run effect is indeed the effect we want to estimate in order to investigate the question of consumer myopia. More generally, short-run effects are important for auto manufacturers in the short-to-medium term (especially if financial solvency is an issue) and because they yield

¹²The average price of gasoline in a DMA-month (our unit of observation) is \$1.91; the standard deviation is 0.68. The “within region-year” standard deviation is 0.21, a value that is 11 percent of the mean. The “between region-year” standard deviation is 0.72. (The “within” standard deviation is the standard deviation of $X_{DMA,month} - \bar{X}_{region,year} + \bar{X}$ where $\bar{X}_{region,year}$ is the average for the region-year and \bar{X} is the global mean. The between standard deviation is the standard deviation of $\bar{X}_{region,year}$.)

¹³As gasoline prices began to fall in the early 1980s, CAFE standards also affected manufacturer offerings.

some insight into the size of the pressures to which manufacturers are responding as they move towards the long run.

B. New car price results

We first estimate Equation 2 using data on new car transactions. The full results from estimating this specification are presented in Table A2 of the online appendix. The variable of primary interest is *GasolinePrice* in month t in the DMA in which customer i resides.¹⁴ 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 λ_1 which represent the effect of gasoline prices on the prices of cars in each of the four MPG quartiles; these coefficients and their standard errors are reported in Table 3.¹⁵ To account for correlation in the errors due to either supply or demand factors, we cluster the standard errors at the DMA level.

TABLE 3—GASOLINE PRICE COEFFICIENTS FROM NEW CAR PRICE SPECIFICATION

Variable	Coefficient	SE
GasolinePrice*MPG Quart 1 (lowest fuel economy)	-250**	(72)
GasolinePrice*MPG Quart 2	-96**	(37)
GasolinePrice*MPG Quart 3	-11	(26)
GasolinePrice*MPG Quart 4 (highest fuel economy)	104*	(47)

These estimates indicate that a \$1 increase in the price of gasoline is associated with a lower negotiated price of cars in the lowest fuel economy quartile (by \$250) but a higher price of cars in the highest fuel economy quartile (by \$104), a relative price difference of \$354. 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.

C. Used car price results

In this section, we estimate the effect of gasoline prices on the transaction prices of used cars by estimating Equation 2 (with some modifications) using the data on used car transactions. We observe all the same car characteristics for used cars that we do for new cars, enabling us to use all the covariates to estimate the used car price results that we used to estimate the results for new cars, including

¹⁴Another approach would be to use a variable that represents gasoline price expectations, perhaps based on futures prices for crude oil. In section V.B we explore such an approach.

¹⁵Two asterisks (**) signifies significance at the 0.01 level, * signifies significance at the 0.05 level and + at the 0.10 level.

identical “car type” fixed effects.¹⁶ However, there is one important difference between used cars and new cars. A new car of a given model-year can sell only during that model-year; a used car of a given model-year can sell in many different years. Over that time period, tastes may change, and individual vehicles will depreciate. To capture the effect of depreciation on used car transaction prices, we include a spline in odometer (*Odom*) when we estimate Equation 2 using the data on used car transactions.¹⁷ The spline has knots at 10,000-mile increments, allowing a different per mile rate of depreciation for each 10,000-mile range of mileage.¹⁸ We interact the spline with segment indicator variables to allow different types of cars to have different depreciation paths, and with indicators for five regions of the country defined by Petroleum Administration for Defense Districts (PADDs) to allow these paths to vary regionally.¹⁹ In addition, in order to allow for changes in tastes for different vehicles segments over time, we replace the year fixed effects in Equation 2 with segment-specific year fixed effects.²⁰ In the new car specification (Equation 2) we allowed the year fixed effects to differ by region. We also allow the segment-specific year fixed effects to vary by geography, however, to reduce the number of fixed effects we have to estimate, we now interact the segment-specific year fixed effects with PADD instead of region.²¹ This three way interaction controls for business cycle fluctuations that affect the entire car market, for year-to-year changes in tastes for different segments of cars (such as the increasing popularity of SUVs), and allows both of these effects to vary across the five PADD regions of the country. Taking into account these modifications, the specification we estimate for used cars is:

$$\begin{aligned}
 P_{irjt} = & \lambda_0 + \lambda_1(\text{GasolinePrice}_{it} \cdot \text{MPG Quartile}_j) \\
 (3) \quad & + f_{10,000}(\text{Odom}_i, \lambda_{2rj}) \cdot \text{Segment}_j \cdot \text{PADD}_r \\
 & + \lambda_3 \mathbf{Demog}_{it} + \lambda_4 \mathbf{PurchaseTiming}_{jt} + \delta_j + \tau_{rjt} + \mu_{rt} + \epsilon_{ijt},
 \end{aligned}$$

where τ_{rjt} is the year-segment-PADD fixed effect.

One could also consider allowing depreciation to vary by MPG quartile and region instead of by segment and region. (In other words, one could replace $f_{10,000}(\text{Odom}_i, \lambda_{2rj}) \cdot \text{Segment}_j \cdot \text{PADD}_r$ in equation 3 with $f_{10,000}(\text{Odom}_i, \lambda_{2rj}) \cdot$

¹⁶The definition of the price of the car is also the same. We subtract any profits (or add any losses) the customer makes trading in a car he or she currently owns in exchange for a different car. Used cars do not have any manufacturer rebate to subtract.

¹⁷In using odometer, our approach resembles Sallee, West and Fan (2009). We differ from Allcott and Wozny (2011), who use car age to measure depreciation. We use odometer for two reasons. First we find that adding car age does very little (in an R^2 sense) to explain depreciation once odometer is accounted for. Second, since odometer varies across individual vehicles, and does not move in lockstep with calendar time, odometer is less collinear with gasoline price than car age is. Using odometer thus increases our ability to identify a gasoline price effect in the data, if there is one.

¹⁸We drop the 0.97 percent of the sample with odometer readings of 150,000 miles or greater.

¹⁹There are seven segments: Compact, Midsize, Luxury, Sporty, SUV, Pickup, and Van. The five PADDs are East Coast, Midwest, Gulf Coast, Rockies, and West Coast.

²⁰In the new car specification, changes in tastes are captured by the car type fixed effects since any particular car type sells as a new car only for one model-year.

²¹In unreported results we find that using year \times segment \times region fixed effects yields very similar results.

MPG Quartile_{*j*} · PADD_{*r*}.) A priori, we think that segment is a better categorization for vehicle depreciation than MPG quartile. Our belief is that SUVs are more likely to depreciate according to the same pattern as other SUVs, and luxury cars more like other luxury cars, than a midsize SUV and a high horsepower luxury car are to depreciate according to the same pattern just because they fall in the same MPG quartile. Additionally, allowing depreciation to vary by MPG quartile instead of segment divides vehicles into the same categorization for measuring gasoline price effects as for measuring depreciation effects. This will substantially increase the ability of our odometer measure to soak up any correlated gasoline price effect, and will make it difficult for us to identify whatever gasoline price effect is in the data. Nevertheless, we report results below that use this alternative interaction.

As we did for new cars, we estimate 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 column 1 of Table A3 in the online appendix. (Column 2 of Table A3 in the online appendix reports the results if depreciation is allowed to vary by MPG quartile instead of segment.) The gasoline price coefficients from columns 1 and 2 of Table A3 are reported in panels 1 and 2 of Table 4.

TABLE 4—GASOLINE PRICE COEFFICIENTS FROM USED CAR PRICE SPECIFICATION

Variable	(1)		(2)	
	Coefficient	SE	Coefficient	SE
GasolinePrice*MPG Quart 1 (lowest fuel economy)	-1182**	(42)	-783**	(49)
GasolinePrice*MPG Quart 2	-101	(62)	118*	(54)
GasolinePrice*MPG Quart 3	468**	(36)	369**	(33)
GasolinePrice*MPG Quart 4 (highest fuel economy)	763**	(44)	360**	(36)
Depreciation varies by	Segment × PADD		MPG Quartile × PADD	

These estimates show a much larger effect on the equilibrium prices of used cars than was estimated for new cars. The estimates in column 1 indicate that a \$1 increase in gasoline price is associated with a lower negotiated price of cars in the lowest fuel economy quartile (by \$1,182) but a higher price of cars in the highest fuel economy quartile (by \$763), a relative price difference of \$1,945, compared to a difference of \$354 for new cars.²²

²²The estimates in panel 2 of Table 4, which allows depreciation to vary by MPG quartile, imply that a \$1 increase in the price of gasoline would be predicted to increase the price of a car in the highest fuel economy quartile of cars relative to that in the lowest fuel economy by \$1,143. Note that the results in panel 2 are non-monotonic; they imply that an increase in the price of gasoline increases the price of an MPG quartile 3 used car by more than (statistically, by the same amount as) it increases the price of a quartile 4 car. Quartile 4 cars all have lower fuel costs per mile than quartile 3 cars, so one should be cautious about calculating implicit discount rates on the basis of this column.

D. *Specification and variables for car quantity results*

In this section we estimate the reduced form effect of gasoline prices on the equilibrium market shares and sales of new cars of different fuel economies. We can write an analog of Equation 1 that gives a reduced form expression for new car quantity, or some function of quantity, as a function of demand and supply covariates:

$$(4) \quad Q = \beta_0 + \beta_1 \mathbf{X}^D + \beta_2 \mathbf{X}^S + \eta$$

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 quantities.

We will estimate two variants of Equation 4. 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 purchases. 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 also estimate a second variant of Equation 4 using two different measures of unit sales.

In our market share regression we estimate the effect of gasoline prices on market shares of cars of different fuel economies using a set of linear probability models that can be written as:

$$(5) \quad I_{irt}(j \in K) = \gamma_0 + \gamma_1 \text{GasolinePrice}_{it} + \gamma_2 \mathbf{Demog}_{it} \\ + \gamma_3 \mathbf{PurchaseTiming}_{jt} + \tau_{rt} + \mu_{rt} + \epsilon_{ijt}.$$

$I_{irt}(j \in K)$ 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 .²³ We use quartiles of fuel economy to define the classes into which a car type falls. As described in Section III.A, quartiles are based on the distribution of fuel economies of car models for sale in a given year (i.e., the model-weighted, not sales-weighted, distribution).

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. We use the same demographic and purchase timing covariates and the same region-specific year and region-specific month-of-year fixed effects that we used to estimate the effect of gasoline prices on new car prices in Equation 2, although in estimating Equation 5 we cannot use the “car type” fixed effects that we used to estimate Equation 2 because “car type” would perfectly predict the fuel economy quartile of the transaction. We will estimate Equation 5 four times, once for each fuel

²³Our results do not depend on the linear probability specification; we obtain nearly identical results with a multinomial logit model (see section V.E).

economy quartile.

In order to estimate the effect of gasoline prices on unit sales, we use two different measures of unit sales. The first measure we use aggregates our individual transaction data into unit sales by dealer, for each month, by MPG quartile.²⁴ Using this measure, we estimate:

$$(6) \quad Q_{dkrt} = \gamma_0 + \gamma_1(\text{GasolinePrice}_{dt} \cdot \text{MPG Quartile}_k) \\ + \gamma_2 \text{MPG Quartile}_k + \delta_d + \tau_{rt} + \mu_{rt} + \epsilon_{dkrt}.$$

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. The coefficients of primary interest are γ_1 . These coefficients estimate the average effect of gasoline prices on new car sales within a fuel economy quartile. We include fixed effects for each of the MPG quartiles and for individual dealers (δ_d). Finally, as in Equation 5, we include year, τ_{rt} , and month-of-year, μ_{rt} , fixed effects that 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 our final specification, we measure sales at the national level using information from Ward's Auto Infobank.²⁵ Using these data, we estimate:

$$(7) \quad Q_{kt} = \gamma_0 + \gamma_1(\text{GasolinePrice}_t \cdot \text{MPG Quartile}_k) \\ + \gamma_2 \text{MPG Quartile}_k + \tau_t + \mu_t + \epsilon_{kt}.$$

Q_{kt} is the national unit sales for vehicles in MPG quartile k that occur in month t .²⁶ The variable of primary interest is again *GasolinePrice*, which is now measured as the national average in month t . The coefficients of interest are the four coefficients in the vector γ_1 which represent the effects 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 .²⁷

²⁴We aggregate from our full data set, not the 10 percent random sample that we use elsewhere in the paper.

²⁵Our 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 slightly over time, and we don't know the year-to-year scaling factor.

²⁶Ward's reports sales data for some cars by a more aggregate model designation than the EPA uses to report MPGs. 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.

²⁷In results available from the authors, we use a third unit sales measure. That third measure uses 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.

E. New car market share results

We first consider the effect of gasoline prices on the market shares of new cars in different quartiles of fuel economy. 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.

In order to estimate Equation 5, 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.

The full estimation results are reported in Table A4 in the online appendix. The estimated gasoline price coefficients (γ_1) for each specification are presented in Table 5. We also report the standard errors of the estimates, and the average market share of each MPG quartile in the sample period. (Since the quartiles are based on the distribution of available models, market shares need not be 25 percent 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.

TABLE 5—GASOLINE PRICE COEFFICIENTS FROM NEW CAR MARKET SHARE SPECIFICATION

Fuel Economy	Coefficient	SE	Mean market share	% Change in share
MPG Quartile 1 (lowest fuel economy)	-0.057**	(0.0048)	21.06%	-27.1%
MPG Quartile 2	-0.014**	(0.004)	20.95%	-6.7%
MPG Quartile 3	0.0002	(0.0027)	24.28%	0.1%
MPG Quartile 4 (highest fuel economy)	0.071**	(0.0058)	33.72%	21.1%

These results suggest that a \$1 increase in gasoline price decreases the market share of cars in the lowest fuel economy quartile by 5.7 percentage points, or 27.1 percent. Conversely, we find that a \$1 increase in gasoline price increases the market share of cars in the highest fuel economy quartile by 7.1 percentage points, or 21.1 percent. This provides evidence that higher gasoline prices are associated with the purchase of cars with higher fuel economy. 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 car purchases. 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.

F. 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. Changes in gasoline prices may be correlated, for macroeconomic reasons, 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 report the results of our two unit sales specifications, Equation 6 and Equation 7.

The coefficient estimates for these two specifications are reported in Tables 6 and 7. The tables report the estimated gasoline price coefficients for each of the four MPG quartiles, 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. On average, a dealer sells 11.2 cars per month in the lowest fuel economy quartile of available cars; a \$1 increase in gasoline prices is estimated to reduce that number by 3.1 cars, or 27.7 percent. On average, dealers sell 17.8 cars per month in the highest fuel economy quartile of cars; a \$1 increase in gasoline prices increases that number by 2.1 cars, or 11.8 percent. Adding up the predicted effects across quartiles shows that an increase in gasoline prices is predicted to reduce the total sales of new cars. Consistent with this, the percentage changes in unit sales are more negative quartile-by-quartile than the percentage changes in market share reported in the previous subsection.²⁸

TABLE 6—GASOLINE PRICE COEFFICIENTS FROM DEALER-LEVEL UNIT SALES SPECIFICATION

Fuel Economy	Coefficient	SE	Average cars sold per month in dealer	% Change in sales
MPG Quartile 1 (lowest fuel economy)	-3.1**	(0.091)	11.2	-27.7%
MPG Quartile 2	-0.83**	(0.087)	11.1	-7.5%
MPG Quartile 3	-0.71**	(0.088)	13.0	-5.5%
MPG Quartile 4 (highest fuel economy)	2.1**	(0.11)	17.8	11.8%

According to the estimates using the Ward's national sales data, reported in the next table, when gasoline prices increase by \$1, there are 79,169 fewer cars per month sold in the lowest fuel economy quartile of cars. This is a 27.2 percent decrease relative to the 291,533 monthly average in this quartile. In the highest fuel economy quartile, a \$1 increase in gasoline prices is associated with an increase in monthly sales of 40,116 cars, a 10.8 percent increase on the average monthly sales in this quartile of 372,998.

²⁸This is consistent with Christopher R. Knittel and Ryan Sandler (forthcoming) which finds that increases in gasoline prices reduces the scrappage rates of used vehicles, in aggregate.

TABLE 7—GASOLINE PRICE COEFFICIENTS FROM NATIONAL UNIT SALES SPECIFICATION

Fuel Economy	Coefficient	SE	Average cars sold per month nationally	% Change in sales
MPG Quartile 1 (lowest fuel economy)	-79,169**	(9,421)	291,533	-27.2%
MPG Quartile 2	-14,761	(9,994)	262,453	-5.6%
MPG Quartile 3	-30,029**	(9,609)	329,466	-9.1%
MPG Quartile 4 (highest fuel economy)	40,116**	(11,800)	372,998	10.8%

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

G. Used car transaction share results (an aside)

While we can easily estimate Equation 5 using our data on used car transactions, the estimates do not have the same interpretation as the estimates for new cars. 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 Equation 5 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 are reported in Table A5 in the online appendix. The gasoline price coefficients are reported in Table 8.

TABLE 8—GASOLINE PRICE COEFFICIENTS FROM USED CAR TRANSACTION SHARE SPECIFICATION

Fuel Economy	Coefficient	SE	Mean share	% Change in share
MPG Quartile 1 (lowest fuel economy)	0.00018	(0.0069)	24.19%	0.07%
MPG Quartile 2	-0.0077	(0.006)	20.89%	-3.7%
MPG Quartile 3	0.017	(0.011)	27.32%	6.2%
MPG Quartile 4 (highest fuel economy)	-0.009	(0.0074)	27.61%	-3.3%

The results are both smaller in magnitude and weaker in statistical significance than the analogous results for new cars.

SUMMARY OF RESULTS

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 highest and lowest fuel economy quartiles of new cars by \$354. The estimated effects are much larger for used cars; in this market, the predicted effect is to increase the price difference between the highest and lowest fuel economy quartiles by \$1,945.

We find both statistically and economically significant effects of gasoline prices on new car sales, measured either as market shares or as unit sales. This is particularly true for the highest fuel economy and lowest fuel economy quartiles, where market share shifts by more than 20 percent in response to a \$1 increase in gasoline prices, and where unit sales decrease by more than 25 percent for the lowest fuel economy quartile and rise by more than 10 percent for the highest fuel economy quartile.

IV. Consumer valuation of future fuel costs

In this section, we draw upon the estimates in the previous section to investigate whether 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.

A. Empirical approach

The basic starting point for the consumer myopia literature is a simple idea: an increase in the expected future usage cost of a durable good should not change consumers’ total willingness-to-pay for the good, all else equal. This means that if the usage cost component of the total cost rises, the up-front cost must fall by an equal amount if consumers (whose total willingness-to-pay is unchanged) 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, unobservable attributes. 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.²⁹

²⁹A recent example of a paper that takes this approach is Molly Espey and Santosh Nair (2005), who estimate a hedonic regression of list prices on a variety of attributes for a cross-sectional sample of 2001

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 cars with relatively high fuel economy even if their operating costs have actually gone up, because the operating cost would have decreased *relative* to that of a low fuel economy 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:

$$(8) \quad q_1 = \alpha_1 + \beta_{11}(p_1 + G_1) + \beta_{12}(p_2 + G_2)$$

Solving this for price implies the following relationship:

$$(9) \quad p_1 = -\gamma G_1 + \frac{1}{\beta_{11}} q_1 - \frac{\alpha_1}{\beta_{11}} - \frac{\beta_{12}}{\beta_{11}} (p_2 + G_2)$$

where $\gamma = -1$ is implied by consumers who correctly value future fuel costs. One could test whether consumers really do behave this way by estimating γ as a free parameter.

There are three difficulties in estimating this relationship in practice. First, 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 would be difficult to implement.³⁰ A second difficulty is that there may be endogeneity between q_i and p_i , arising from a supply relationship between the two variables.

In this paper, we will take an alternative approach. Our approach is to combine our reduced-form estimates of price and quantity effects with estimates of the elasticity of demand for new cars, and estimates of future gasoline prices, vehicle miles travelled, and vehicle survival rates in order to address the question of whether consumers are myopic with respect to future fuel costs. Note that these assumptions are very similar to the set of assumptions that must be made in

model year cars. They conclude that consumers use fairly low discount rates when valuing future fuel cost savings.

³⁰An alternative approach, used by Allcott and Wozny (2011), 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 structural approach. In this sense, the two approaches do not differ in how many assumptions must be imposed, but at what stage in the analysis they are imposed. The structural approach imposes them earlier and is able thereby to estimate a single parameter that captures the degree of consumer myopia and can be used in counterfactual simulations. The reduced form approach will be more amenable to examining the effect of a variety of assumptions about vehicle miles travelled, future gasoline prices, and vehicle survival rates. We will present a range of estimates; it will be fairly straightforward for readers to substitute their own assumptions as well.

B. Consumer myopia results

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 effects of gasoline price changes on buyers' willingness-to-pay for cars of different fuel economies to the changes in the discounted value of future gasoline costs that are 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 car owners drive in a given year, conditional on the car surviving through that year, and also annual survival rates. We calculate miles driven, conditional on survival, three ways. We use NHTSA-assumed values for annual miles driven, separately for cars and light duty trucks, by vintage. These data are used in a number of modeling efforts for both the NHTSA and DOT (S. Lu (2006)). Our other two measures come from within our data: we compute the average annual miles driven, by vintage, separately for cars and trucks, for vehicles in our used car transaction data and for 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.³¹ 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 model consumers' expectations of future gasoline prices as following a random walk for real gasoline prices. This has the convenient implication that the current gasoline price is the expected future real gasoline price. (Soren T. Anderson, Ryan Kellogg and James M. Sallee (2011), discussed in more detail in Section V.B, show empirical evidence that this is indeed the gasoline price

³¹One might worry that there is a survival bias toward lower mileage cars in the cars that we see as trade-ins and in used car transactions. In order to mitigate this, we use the NHTSA values for any vintage-vehicle class cell in which the VMT calculated from our data is lower than the NHTSA figure for the same cell.

expectation that consumers have on average.) One alternative is to assume that consumers are more sophisticated and use information on crude oil futures markets to make projections into the future.³² It turns out that for the vast majority of time during our sample, the crude market was in backwardation; that is, the market expected crude prices to fall. (See Figure 4 for a plot of both the spot crude price and the stream of expected prices in subsequent years for May of each year—the “forward curve.”) 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. In other words, our approach biases us toward finding myopia. (Our approach increases the chances of falsely concluding that consumers behave myopically.)³³




FIGURE 4. CRUDE SPOT AND FUTURES PRICES DURING OUR SAMPLE

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 use our estimates for some components of the calculation, we make assumptions about the other components, and see what the combination implies for the discount rate.

³²See section V.B for the results of such an approach.

³³A third justification for using current gasoline prices is that consumers may not be sophisticated in forming expectations, and may base their decisions on the most salient gasoline price they see—the one currently posted at gas stations nearby.

Fourth, in order to address the question of myopia, we need to observe the effects of gasoline prices on consumers' *willingness-to-pay* for cars of different fuel economies; what we have estimated so far is the effect of gasoline prices on *equilibrium transaction prices*. In order to translate a change in equilibrium price to a change in willingness-to-pay, we need to consider supply and demand in the new and used car markets. In the used car market, one might argue that a fixed supply curve is a reasonable assumption for used car supply. This is because the stock of used cars is predetermined by the cumulation of past new car purchases, and is likely to respond very little to gasoline prices.³⁴ Many cars sold on the used market are fleet turnovers and lease returns whose entry into the used car market will not be determined primarily by gasoline prices. If consumers are also driven to replace their existing cars by factors unrelated to gasoline prices, the supply of a particular used car model at any point in time could be thought of as essentially fixed. If this is the case, then the effect of a change in demand for that model ought to show up almost entirely in the equilibrium prices of used cars of different types.³⁵ This means that the equilibrium price effect will be equal to the change in willingness-to-pay. (Figure 5 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. Prices can be adjusted quickly by using promotions (Meghan R. Busse, Jorge Silva-Risso and Florian Zettelmeyer 2006). Production quantities can be adjusted by adding or reducing shifts on assembly lines (Timothy F. Bresnahan and Valerie A. Ramey 1994), or for some modern manufacturing plants by adjusting which kinds of vehicles are produced on a given line.³⁶ Car dealers can easily adjust the prices they negotiate with individual customers, and can adjust quantities by changing inventory holdings and orders to manufacturers. This means that 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 5 shows a representation of this for a hypothetical new car model.)

Since we estimate the equilibrium effects on prices and quantities, we could

³⁴Lucas W. Davis and Matthew Kahn (2010) 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.

³⁵The same equilibrium effect would obtain if an increase in the price of gasoline increases the demand for high fuel economy used vehicles, but leads to a commensurate reduction in the supply of such used vehicles (because the current owners choose to hold onto them for longer). Similarly, an increase in the price of gasoline might reduce the demand for low fuel economy used cars at the same time as the supply of low fuel economy vehicles increases (because the current owners wish to replace their current vehicles with a higher fuel economy vehicles).

³⁶For example, Honda can build the compact Civic on the same assembly line that builds the Ridgeline pickup and the Acura MDX SUV ("Adaptability helps Honda weather industry changes," *Automotive News*, June 8, 2009). 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.




FIGURE 5. EFFECTS OF GASOLINE PRICE CHANGE ON HYPOTHETICAL USED AND NEW CARS

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. Pinelopi K. Goldberg (1995) estimates residual demand elasticities of demand for specific vehicles that are in the neighborhood of -2 to -4 , while Steven Berry, James Levinsohn and Ariel Pakes (1995) estimate elasticities in the -3 to -6 range.³⁷ We note that these estimates should be strong upper bounds (in absolute value) to the relevant demand elasticity for our purposes. 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 higher than the demand elasticity for a vehicle in a particular fuel economy quartile, which is the relevant elasticity for us. Finally, we assume that demand has a constant elasticity functional form.³⁸

In the table below, we present the results of our investigation into the question of whether consumers are myopic. The entries in the last three columns are the implicit discount rates necessary to equate the relative price differences between vehicles of different fuel economies to the relative differences in discounted expected future fuel costs between those vehicles. The relative price differences we use are the estimates from Table 3 for new cars and from Table 4 for used cars. Because the expectation of future gasoline prices we have used is an expectation

³⁷Goldberg (1995) reports average elasticities 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 .

³⁸The 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.

about real, not nominal, gasoline prices, the implicit discount rates we calculate do not contain inflation expectations. Note that the table presents a range of possible estimates of the implicit discount rate that hold for a particular set of assumptions.³⁹ Changing those assumptions would, of course, change the implicit discount rates obtained.⁴⁰

The top panel of Table 9 reports the implicit discount rates when comparing the estimated price effects for the lowest fuel economy quartile of cars relative to the highest fuel economy quartile. The middle panel reports for the lowest fuel economy quartile relative to the quartile with second highest fuel economy; the bottom panel for the quartile with second lowest fuel economy relative to the highest fuel economy quartile. The top row of each panel reports the implicit discount rates based on the relative price effects estimated for used cars. The next four rows report the implicit 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 transactions in our data, and from the trade-ins in our data, respectively.

TABLE 9—NEW AND USED CARS: IMPLICIT 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	11.8%	4.4%	7.3%
	New	-2	-4.0%	-6.8%	-6.2%
	New	-3	1.0%	-3.0%	-1.9%
	New	-4	5.5%	0.5%	2.1%
	New	-5	9.8%	3.7%	5.8%
Q1 vs. Q3	Used	NA	5.9%	0.1%	1.9%
	New	-2	-3.6%	-6.6%	-5.9%
	New	-3	1.5%	-2.6%	-1.5%
	New	-4	6.1%	0.9%	2.5%
	New	-5	10.4%	4.2%	6.3%
Q2 vs. Q4	Used	NA	20.9%	11.0%	16.2%
	New	-2	0.3%	-3.5%	-2.5%
	New	-3	6.7%	1.4%	3.1%
	New	-4	12.6%	5.8%	8.3%
	New	-5	18.3%	10.0%	13.2%

Overall, the implicit discount rates range from moderate to quite small. In

³⁹The spreadsheet that makes this calculations—and could be used to show the influence of different assumptions from those presented here—is included in the online data appendix.

⁴⁰One plausible effect of gasoline prices that is not included in the assumptions underlying Table 9 is that vehicle miles travelled fall when gasoline prices increase. If this is the case, then the expected future fuel costs of cars of different fuel economies will be more similar than what we have assumed for the table, meaning that implicit discount rates will have to be smaller in order to reconcile the estimated change in willingness-to-pay with the change in expected future fuel costs.

most cases, the estimates are in the single digits, with some combinations of assumptions actually implying negative discount rates. The more elastic new vehicle 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 implicit discount rate necessary to rationalize the willingness-to-pay change with a given change in expected future gasoline costs. Since the largest elasticity for new vehicles we use in Table 9 was calculated for individual vehicles rather than for quartiles, whose demand is presumably less elastic, the estimates that use the lower values for elasticities may be the most relevant. The implicit discount rates arising from used car prices are generally higher than those implied by new car prices, although the ranges overlap.

Most of the estimates of implicit discount rates are near or below typical rates for car loans. In our sample, the 10th to 90th percentile range of APRs for consumers financing their car purchase through the dealer is [1.9 percent, 11.6 percent] for new car buyers and [5.5 percent, 19.9 percent] for used car buyers. These APRs are nominal interest rates. During our sample period, inflation rates were between 1.1 and 5 percent. We calculate “real APRs” by subtracting from each APR observation the annual inflation rate.⁴¹ The 10th to 90th percentile range for “real APRs” is [-0.9 percent, 9.0 percent] for new car purchases and [2.5 percent, 17.0 percent] for used cars.⁴² While some of the implicit discount rates fall outside this range, the evidence in Table 9 suggests that the discount rates people use to evaluate future fuel costs are generally comparable to interest rates they pay when they buy a car.

We conclude that there is little evidence that consumers dramatically undervalue changes in expected future fuel costs, and that the evidence from new and from used cars yield similar messages. Our findings on this are similar to Allcott and Wozny (2011) who calculate that their results correspond to a 16 percent implicit discount rate, and to Sallee, West and Fan (2009) who find somewhat less undervaluation of future fuel costs than do Allcott and Wozny (2011). It bolsters our confidence in the results of this entire set of papers that different configurations of identifying assumptions yield similar results. In our view, this lessens the worry readers should have that the results in any of these papers arise directly from a particular set of assumptions.

V. 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 investigate the effect of using an estimate of future gasoline prices (based on oil price futures)

⁴¹Inflation is generally considered a random walk, making the current inflation rate an appropriate measure of inflation expectations.

⁴²The new car APRs are negative when manufacturers subsidize interest rates to fall below market rates.

instead of current gasoline prices as our explanatory variable of interest. 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.

A. *Source of variation*

We now re-estimate the original specifications in this paper with a series of different fixed effect combinations. So far, all specifications have controlled for region-specific month-of-year fixed effects and either region-specific year fixed effects (new cars) or PADD \times segment-specific year fixed effects (used cars). This means that the estimated gasoline price effects have been identified by within-year, region-specific (or PADD \times segment-specific) variation in price, market share, or sales which deviates from region-specific seasonal effects.

In order to investigate the robustness of these estimates, we estimate eight additional specifications with different combinations of fixed effects. Five specifications are more parsimonious than our base specifications reported in Tables 3, 4, and 5, meaning that less of the variation in the left-hand-side variable is absorbed by fixed effects, and three specifications are richer, meaning that more of the variation is absorbed into fixed effects. Table 10 reports the results for new car prices, Table 11 for used car prices, and Table 12 for new car market shares. The most parsimonious specification includes only region fixed effects (no year or month-of-year fixed effects) and the richest specification uses month (not month-of-year) times region or PADD fixed effects. For ease of comparison, in all the tables, row 6 reports the results of our base specification.

In Tables 10 and 11, the most direct way to compare the price estimates across rows is to compare the implied change in the price of a car in quartile 1 relative to a car in quartile 4. This number is reported in the last column of both tables. For new cars (Table 10), the implied change in relative prices ranges from \$327 to \$355. This includes the results in row 9, which use monthly fixed effects. For used cars (Table 11), it ranges from \$1,917 to \$2,563, although the first seven rows vary only from \$1,917 to \$2,196. The fact that our estimates fall within a \$28 range for new cars and a \$279 range for most of the used car results seem to us to be quite stable results, especially considering the differences in the components of variation used to identify the effects in different rows. (Increasing the frequency of the time fixed effect, in rows 8 and 9, has bigger effect on the used car results than on the new car results. Using these estimates would reduce the implicit discount rates relative to what is reported in Table 9.)

Table 12 repeats the exercise for the market share specification. We find that the estimated effect of a \$1 increase in gasoline price on the market share of new cars in the lowest fuel economy quartile ranges from -4.8 percentage points to -7 percentage points. The effect on the market share of the highest fuel economy quartile ranges from 5.8 to 8.7 percentage points. The effect in quartile 2 ranges

TABLE 10—EFFECT OF TIME AND SEASONAL FIXED EFFECTS IN NEW CAR PRICE SPECIFICATION

	Region FE	Time FE	Seasonal FE	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4	Price change Quar 1 to 4
1	Region	–	–	-520** (82)	-373** (45)	-286** (31)	-165** (42)	\$355
2	Region	–	Month-of-year	-251** (73)	-98* (40)	-19 (33)	87+ (49)	\$338
3	Region	Year	Month-of-year	-245** (73)	-102** (37)	-17 (31)	82+ (49)	\$327
4	–	–	Month-of-year × Region	-255** (73)	-104* (40)	-26 (34)	84+ (50)	\$339
5	–	Year	Month-of-year × Region	-250** (73)	-108** (38)	-25 (32)	78 (50)	\$328
6 (Base)	–	Year × Region	Month-of-year × Region	-250** (72)	-96** (37)	-11 (26)	104* (47)	\$354
7	–	Year × Region Year × Trend	Month-of-year × Region	-205** (68)	-62+ (35)	18 (29)	127* (51)	\$332
8	–	Quarter × Region	Month-of-year × Region	-161* (68)	-16 (41)	66+ (35)	171** (58)	\$332
9	–	Month × Region		-313* (157)	-177 (152)	-90 (155)	14 (173)	\$327

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

from 2.9 percentage points to statistically zero, and the effects for quartile 3 are almost all statistically zero. The results seem to us to be again quite stable. The one exception to this is row 9, which estimates a fixed effect for each month of the sample separately for each region of the country; this approach taxes the data quite heavily, and the estimated effects, while not wildly different in magnitude from those in the other rows, are no longer statistically significant.

B. Future vs. current gasoline prices

In the results we have presented so far, we have estimated the effect of current gasoline prices on the market outcomes from new and used cars. One might argue that since cars are durable goods, buyers should make decisions in response to their *expectations of future gasoline prices*, rather than current gasoline prices. There are several justifications for using current gasoline prices as our explanatory variable of interest. First, it may be the case that car buyers are not sophisticated in thinking about expectations, and that they instead respond to the price that they see posted prominently at gas stations and hear discussed in the news media. Second, if gasoline prices are a random walk, then the expected future gasoline price *is* the current gasoline price. With respect to this point, Anderson, Kellogg and Sallee (2011) use a set of questions on the Michigan Survey of Consumers that ask explicitly about consumers' gasoline price expectations to show that consumers' average forecast is for “no change in real gasoline prices.” This

TABLE 11—EFFECT OF TIME AND SEASONAL FIXED EFFECTS IN USED CAR PRICE SPECIFICATION

	Region FE	Time FE	Seasonal FE	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4	Price change Quar 1 to 4
1	Region	–	–	-3454** (70)	-2923** (98)	-2105** (68)	-1258** (43)	\$2196
2	Region	–	Month-of-year	-3773** (80)	-3240** (108)	-2428** (77)	-1589** (55)	\$2184
3	Region	Year × Segment	Month-of-year	-1197** (42)	-118+ (63)	436** (39)	720** (49)	\$1917
4	–	–	Month-of-year × Region	-3796** (82)	-3263** (111)	-2454** (79)	-1610** (58)	\$2186
5	–	Year × Segment	Month-of-year × Region	-1185** (41)	-106+ (59)	450** (36)	734** (45)	\$1919
6 (Base)	–	Year × Segment × PADD	Month-of-year × Region	-1182** (42)	-101 (62)	468** (36)	763** (44)	\$1945
7	–	Year × Segment × PADD, Year × Trend	Month-of-year × Region	-1264** (41)	-181** (59)	388** (33)	693** (41)	\$1957
8	–	Quarter × Segment × PADD	Month-of-year × Region	-1552** (59)	-355** (75)	377** (47)	781** (55)	\$2333
9	–	Month × Segment × PADD	Month-of-year × Region	-1759** (178)	-513** (194)	321+ (182)	804** (188)	\$2563

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

would suggest that our empirical approach, by using current prices, may indeed estimate the response to consumers' expectations of real gasoline prices. This also supports our interpretation of the consumer myopia results as measuring the implicit discount rate that rationalizes the current change in car prices with the change in future fuel costs implied by a change in the future real price of gasoline.

In this section, we take an alternative approach, which is to make use of the active futures market for crude oil to create a measure of consumers' expectations of future gasoline prices. While futures contracts for gasoline are listed on the NYMEX, futures in oil are actively traded in much larger volumes. Furthermore, gasoline prices are sufficiently closely correlated with oil prices that we suspect that a gasoline price forecast based on futures prices for oil will be a better measure of expectations of gasoline prices than using gasoline futures prices directly.

To be precise, we regress monthly, DMA-level gasoline prices on current, world prices for crude oil, and on DMA×month-of-year fixed effects. This allows us to translate the price of a barrel of oil into the price for a gallon of gasoline, allowing the “markup” between these two to vary by region, and by season differentially for each region. (This allows, for example, for refinery margins that vary geographically and seasonally.) We use the estimated coefficients from this regression to predict an expected future price for gasoline by using alternately the 6-month-ahead, 12-month-ahead, and 24-month-ahead futures price for crude oil in place of the current price of oil. We then use our prediction of the expected future price of gasoline in place of the current price of gasoline in our benchmark

TABLE 12—EFFECT OF TIME AND SEASONAL FIXED EFFECTS IN NEW CAR MARKET SHARE SPECIFICATION

	Region FE	Time FE	Seasonal FE	MPG Quartile 1	MPG Quartile 2	MPG Quartile 3	MPG Quartile 4
1	Region	–	–	-.05** (0.0044)	-.0062 (0.0047)	-.0055+ (0.0031)	.061** (0.0058)
2	Region	–	Month-of-year	-.049** (0.0043)	-.0056 (0.0046)	-.005+ (0.003)	.059** (0.0056)
3	Region	Year	Month-of-year	-.052** (0.0051)	-.016** (0.0058)	-.003 (0.0034)	.071** (0.0077)
4	–	–	Month-of-year × Region	-.048** (0.0042)	-.0053 (0.0043)	-.0041 (0.003)	.058** (0.0051)
5	–	Year	Month-of-year × Region	-.053** (0.0046)	-.016** (0.0051)	-.00038 (0.003)	.07** (0.006)
6 (Base)	–	Year × Region	Month-of-year × Region	-.057** (0.0048)	-.014** (0.004)	.0002 (0.0027)	.071** (0.0058)
7	–	Year × Region Year × Trend	Month-of-year × Region	-.049** (0.005)	-.025** (0.0034)	-.00084 (0.0034)	.075** (0.0043)
8	–	Quarter × Region	Month-of-year × Region	-.07** (0.0069)	-.029** (0.0056)	.012** (0.0047)	.087** (0.0065)
9	–	Month × Region		-.12+ (0.068)	-.013 (0.04)	.036 (0.03)	.097 (0.078)

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

specifications for new and used car prices and new car market shares. The results using our prediction of the expected future price of gasoline are reported in Tables A6 and A7 in the online appendix.

The estimated results for the new car market share in Table A7 are very similar whether we use current gasoline prices or our prediction of expected future gasoline prices. For the estimate of the effect of gasoline prices on car prices in Table A6, using our prediction of the expected future gasoline prices yields relative price effects that are 61 percent - 87 percent larger in magnitude for new cars and 17 percent - 19 percent larger in magnitude for used cars than using the current gasoline price. These results imply that consumers adjust their willingness-to-pay *more* in response to changes in future fuel costs than is reflected in the estimates used to produce Table 9. Using the estimates from Tables A6 and A7 would generate lower implicit discount rates than what is reported in Table 9, corresponding to less myopia among car buyers.

C. Gasoline price aggregation

Next, we investigate the robustness of our findings to the aggregation of gasoline prices to local markets (DMAs) rather than to ZIP-codes, which would be possible in our data. 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 for 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 DMA-level prices because consumers may notice gasoline price changes only once they have affected a large enough area to be reported in the media, or because local price variation contains transitory price shocks that do not enter into long-run forecasts of gasoline prices.

To investigate whether our conclusions depend of the level of aggregation of gasoline prices, we re-estimate our original car price and market share specifications (Equations 2, 3, and 5) using one less-aggregated and two more-aggregated measures of gasoline prices. We use 4-digit ZIP-code level gasoline price as our less aggregated measure.⁴³ 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, defined such that they delineate a region in which gasoline supply is homogenous. There are five PADDs: East Coast, Midwest, Gulf Coast, Rockies, and West Coast. There remains substantial variation in gasoline prices between PADDs and within PADDs over time. We also use national average gasoline prices as a still more aggregated measure.

The results are reported in Tables A8 and A9 in the online appendix. In both tables, we find that the coefficients on gasoline prices in the 4-digit ZIP code aggregation are slightly smaller than those in our (original) DMA aggregation, with the largest difference in the market share regression. This general finding of smaller coefficients is consistent with some measurement error occurring in the 4-digit ZIP code aggregation.

If we aggregate gasoline prices at the PADD or national level, most coefficient estimates in the market share regression are essentially unchanged. In the price regressions, the PADD-level estimates are distinctly larger in magnitude than the estimates using DMA-level prices, and the national-level estimates are larger still, although the differences are bigger for the new car price estimates than the used car price estimates.

Overall, we would reach many of the same conclusions about the effects of gasoline price changes if we aggregated gasoline prices within 4-digit ZIP code, within PADD, or nationally instead of within DMAs. We recalculated the implicit discount rates reported in Table 9 using the coefficients from the PADD and national price specifications, and from (unreported) unit sales regressions using the PADD- and national-level prices. We find that increasing the level of aggregation of gasoline prices decreases the estimated implicit discount rates. Loosely

⁴³We 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. 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.

speaking, the larger estimated price coefficients translate into larger changes in willingness-to-pay in response to changes in gasoline prices, which would mean that consumers are less myopic. The implicit discount rates reported in Table 9 fall mostly between -7 percent and the low positive teens. Using PADD-level prices yields implicit discount rates that are almost all mostly in the single digits, including some negative values. Using national average prices results in implicit discount rates that are mostly negative.⁴⁴

D. 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; U.S. gasoline prices are determined by world oil prices and refinery margins, and these are unlikely to be influenced by car transactions in the U.S. However, there are other potential sources of endogeneity which could taint our coefficient estimates. First, there could be local variations in economic conditions that are correlated with local variations in gasoline prices. If such 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 measure of gasoline price; this would make it less likely that local shocks lead to correlation between gasoline prices and the error term in the market share and price specifications. The specification using PADD- and national-level gasoline prices (described in the previous section and reported in Tables A8 and A9 of the online appendix) do 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 A10 and A11 of the online appendix.

We have already concluded that the OLS regression with PADD-level gasoline prices estimates similar market share effects but somewhat larger price effects compared to the original OLS regression with DMA-level gasoline prices. In Table

⁴⁴These results available on request from the authors.

A10, the PADD-level IV estimates of the effect of gasoline prices on market share are about 10 percent larger in magnitude than the PADD-level OLS estimates for the lowest fuel economy quartile, and about 15 percent smaller for the highest fuel economy quartile. We find 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. As a consequence, the estimated effect on the relative price difference between the highest and lowest fuel economy quartiles is larger by 32 percent for used car prices and 40 percent for new car prices in the IV specification than in the PADD-level OLS specification. This can be seen in Table A11.

In summary, controlling for endogeneity suggests that our original specification may have underestimated the magnitude of the gasoline price effect on car prices. Using the PADD-level IV estimates in our myopia calculations would lead to smaller implicit discount rates—implying consumers who value the future more—than what is reported in Table 9.

E. 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 market shares. We reestimate our basic market share specification (Equation 5) 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 5.) Full estimation results are reported in Table A12 of the online appendix.

The coefficients reported in Table A12 correspond to predicted marginal effects of a \$1 increase in gasoline prices that are slightly larger than the effects predicted by the linear probability model. Specifically, the predicted marginal effects are -0.064** vs. -0.057** for MPG quartile 1, -0.014** vs. -0.014** for MPG quartile 2, -0.004 vs. 0.0002 for MPG quartile 3, and 0.075** vs. 0.071** for MPG quartile 4, where ** indicates estimates that are statistically significant at the 1 percent level. We conclude that our market share results do not depend on our use of the linear probability model.

VI. Concluding remarks

In this paper we have estimated the effect of gasoline prices on the short-run equilibrium prices, market shares, and sales of new and used cars of different fuel economies. We have used these estimates to address a question that is important for understanding the ability of a policy intervention such as a gasoline tax or a carbon tax to influence what cars people buy, which is one avenue through which such an intervention can affect greenhouse gas emissions.⁴⁵

⁴⁵Other potential avenues include changing vehicle miles travelled, car designs, fuel technologies, or urban design.

We estimated that a \$1 increase in the price of gasoline increases the market share of cars in the highest fuel economy quartile by 21.1 percent and decreases the market share of cars in the lowest fuel economy quartile by 27.1 percent. We also estimated the effect of a \$1 increase in gasoline prices on unit sales of new cars and found that sales in the highest fuel economy quartile increased by 10–12 percent, while sales in the lowest fuel economy quartile fell by 27–28 percent. 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 \$354 in the average price of the highest fuel economy quartile of cars relative to that of the lowest fuel economy quartile. For used cars, the estimated relative price difference is \$1,945.

We used these estimates to investigate 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. This could be thought of as a necessary condition for effective policy: the more car buyers discount future fuel costs, the less effective a gasoline tax or carbon tax will be in influencing vehicle choice. 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 implicit discount rates are near zero, most are less than 20 percent.⁴⁶

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 these margins. 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 reduced form specifications we use estimate parameters which can be combined by later researchers with a range of assumptions about related parameters in order to answer policy-relevant questions.

⁴⁶The alternative specifications we investigated in Section V generally led to larger relative price effect estimates, which would reduce the estimated implicit discount rates compared to what is reported in Table 9.

REFERENCES

- Allcott, Hunt, and Nathan Wozny.** 2011. "Gasoline Prices, Fuel Economy, and the Energy Paradox." New York University.
- Anderson, Soren T., Ryan Kellogg, and James M. Sallee.** 2011. "What Do Consumers Believe about Future Gasoline Prices?" National Bureau of Economic Research 16974.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.
- Blomqvist, Ake G., and Walter Haessel.** 1978. "Small Cars, Large Cars, and the Price of Gasoline." *The Canadian Journal of Economics/Revue canadienne d'Economique*, 11(3): 470–489.
- Bresnahan, Timothy F., and Valerie A. Ramey.** 1994. "Output Fluctuations at the Plant Level." *Quarterly Journal of Economics*, 109(3): 593–624.
- Busse, Meghan R., and Jorge Silva-Risso.** 2010. "'One Discriminatory Rent' or 'Double Jeopardy': Multi-componnet Negotiation for New Car Purchases." *American Economic Review*, 100(2): 470–474.
- Busse, Meghan R., Jorge Silva-Risso, and Florian Zettelmeyer.** 2006. "\$1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions." *American Economic Review*, 96(4): 1253–1270.
- Carlson, Rodney L.** 1978. "Seemingly Unrelated Regression and the Demand for Automobiles of Different Sizes, 1965–1975: A Disaggregate Approach." *The Journal of Business*, 51(2): 243–262.
- Davis, Lucas W., and Matthew Kahn.** 2010. "International Trade in Used Vehicles: The Environmental Consequences of NAFTA." *American Economic Journal: Economic Policy*, 2(4): 58–82.
- Donna, Javier.** 2011. "Estimating Switching Costs in Urban Travel Demand in Chicago." Northwestern University.
- Espey, Molly, and Santosh Nair.** 2005. "Automobile Fuel Economy: What is it Worth?" *Contemporary Economic Policy*, 23(3): 317–323.
- Goldberg, Pinelopi K.** 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica*, 63(4): 891–951.
- Goldberg, Pinelopi K.** 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the US." *Journal of Industrial Economics*, 46(1): 1–33.
- Gramlich, Jacob.** 2009. "Gas Prices, Fuel Efficiency, and Endogenous Product Choice in the U.S. Automobile Industry." Georgetown University.
- Greenlees, John S.** 1980. "Gasoline Prices and the Purchases of New Automobiles." *Southern Economic Journal*, 47(1): 167–178.
- Hausman, Jerry A.** 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *The Bell Journal of Economics*,

10(1): 33–54.

- Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling.** 2008. “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand.” *Energy Journal*, 29(1).
- Jacobsen, Mark R.** 2010. “Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity.” University of California, San Diego.
- Kahn, James.** 1986. “Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach.” *Quarterly Journal of Economics*, 101(2): 323–340.
- Kilian, Lutz, and Eric R. Sims.** 2006. “The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data.” University of Michigan.
- Klier, Thomas, and Joshua Linn.** 2010. “The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data.” *American Economic Journal: Economic Policy*, 2(3): 134–153.
- Knittel, Christopher R.** 2011. “Automobiles on Steroids: Product Attribute Trade-offs and Technological Progress in the Automobile Sector.” *American Economic Review*, 101(7): 3368–3399.
- Knittel, Christopher R., and Ryan Sandler.** forthcoming. “Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins.” In *The Design and Implementation of U.S. Climate Policy*, ed. Don Fullerton and Catherine Wolfram. University of Chicago Press.
- Li, Shanjun, Christopher Timmins, and Roger 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.” NHTSA Technical Report Working Paper DOT HS 809 952.
- Ohta, Makoto, and Zvi Griliches.** 1986. “Automobile Prices and Quality: Did the Gasoline Price Increases Change Consumer Tastes in the U.S.?” *Journal of Business and Economic Statistics*, 4(2): 187–198.
- Sallee, James M., Sarah E. West, and Wei Fan.** 2009. “Consumer Valuation of Fuel Economy: A Microdata Approach.” National Tax Association Conference Proceedings.
- Sawhill, James W.** 2008. “Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market.” University of California, Berkeley Working Paper.
- Tishler, Asher.** 1982. “The Demand for Cars and the Price of Gasoline: The User Cost Approach.” *The Review of Economics and Statistics*, 64(2): 184–190.
- Verboven, Frank.** 2002. “Quality-based Price Discrimination and Tax Incidence - The Market for Gasoline and Diesel Cars in Europe.” *RAND Journal of Eco-*

nomics, 33(2): 275–297.

West, Sarah E. 2007. “The Effect of Gasoline Prices on the Demand for Sport Utility Vehicles.” Macalester College Working Paper.

CALCULATION OF IMPLICIT DISCOUNT RATES

In this appendix, we provide more detail on how we calculate the implicit discount rates reported in Table 9. The general idea is to calculate the discount rate consistent with the *relative* shift in the demand for two vehicles in two different fuel economy quartiles. For any given discount rate, one can calculate the increase in the discounted fuel costs resulting from a \$1 change in gasoline prices for the average vehicle in each quartile, given assumptions on miles driven and fuel economy. Our spreadsheet calculations search for the implicit discount rate that equates the relative change in these discounted fuel costs between the average cars in two different quartile with the estimated change in relative willingness-to-pay between average cars in the two quartiles.

As described in section IV.B, for used vehicles we use the estimated change in prices as the measure of the shift in the willingness-to-pay. For new vehicles, we use the estimated change in both prices and quantities, plus an assumption about the elasticity of demand, to measure the shift in the demand curves. Figure 5 illustrates how this is done in the case of a linear demand curve, but our calculations assume a constant elasticity demand curve.

The remaining ingredients are the annual expected mileage of the vehicles, accounting for survival, as well as their fuel economies. As described in section IV.B, we report results from three different sets of estimates of mileage. The first are NHTSA estimates of mileage, which are reported separately for cars and trucks, by vintage of the vehicle. The second estimate of annual mileage comes from our data; we calculate the difference between the average odometer of used cars of adjacent vintages that we observe in our used car transactions. We calculate this measure for each vintage, separately by cars and trucks. Finally, our third estimate is the average change in odometer readings, by vintage, for trade-ins in our data. We also calculate this separately for cars and trucks. The last two measures tend to be smaller than the NHTSA estimates. This leads to lower implicit discount rates since, for a given discount rate, the relative price changes appear larger in relation to the change in discounted fuel costs. In light of this, we make one adjustment to the two mileage estimates that based on our data. If ever our observed change in the average odometer falls below the minimum observed in the NHTSA data, we replace the mileage with this minimum; these are 6,131 miles for cars and 6,648 miles for trucks.

We use the same mileage assumption for each of the four quartiles. The implicit assumption is that the modeled consumer does not change her driving patterns when she moves from a vehicle in quartile X to a vehicle in quartile Y. Consistent with this, we take the weighted average of the car and truck mileage using the share of trucks sold in new and used vehicle transaction data. Finally, because the used vehicles in the data do not begin their driving patterns at year 1, we use the estimated mileage patterns beginning in year 4. This is the median age of used vehicle sold in our data is 4 years old; the average is 3.98 years old.

The final ingredient is the fuel economy rating for each of the quartiles. We

use the average fuel economy rating of offered vehicles within each of the four quartiles. The results are very similar if we use the average fuel economy rating of purchased vehicles.