

NBER WORKING PAPER SERIES

WHAT DO CONSUMERS BELIEVE ABOUT FUTURE GASOLINE PRICES?

Soren T. Anderson  
Ryan Kellogg  
James M. Sallee

Working Paper 16974  
<http://www.nber.org/papers/w16974>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
April 2011

We thank Richard Curtin for guidance in using the MSC data. For helpful comments, we also thank Hunt Allcott, Severin Borenstein, Lutz Kilian, Tim Vogelsang, and seminar participants at the ASSA meetings, University of California Energy Institute, University of California Berkeley, University of California Santa Barbara, University of California Davis, Iowa State University, and Michigan State University. Anderson and Sallee gratefully acknowledge financial support from the California Energy Commission. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Soren T. Anderson, Ryan Kellogg, and James M. Sallee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Do Consumers Believe About Future Gasoline Prices?

Soren T. Anderson, Ryan Kellogg, and James M. Sallee

NBER Working Paper No. 16974

April 2011

JEL No. D84,L62,Q40,Q41

**ABSTRACT**

Researchers estimating the demand for energy-using durable goods must specify consumers' beliefs about future energy prices. Policy-relevant inference hinges on this specification, yet there is little direct evidence on the nature of consumer beliefs. We provide such evidence by analyzing two decades of data on gasoline price expectations from the Michigan Survey of Consumers. We find that average consumer beliefs are indistinguishable from a no-change forecast. This finding has important implications for the literature on consumer valuation of energy efficiency, and it implies that researchers are likely justified in assuming a no-change forecast, as is common practice.

Soren T. Anderson  
Department of Economics  
Marshall - Adams Hall, Office 18C  
Michigan State University  
East Lansing, MI 48824-1038  
and NBER  
sta@msu.edu

James M. Sallee  
Harris School of Public Policy Studies  
University of Chicago  
1155 East 60th Street  
Chicago, IL 60637  
and NBER  
sallee@uchicago.edu

Ryan Kellogg  
University of Michigan  
Department of Economics  
238 Lorch Hall  
611 Tappan Street  
Ann Arbor, MI 48109-1220  
and NBER  
kelloggr@umich.edu

# 1 Introduction

When consumers buy energy-using durable goods, such as cars, air conditioners, and appliances, or when they replace home windows and insulation, they must forecast the future price of energy to determine their willingness to pay for energy efficiency. In turn, researchers modeling the demand for such goods must specify consumer beliefs, and econometric estimation hinges on this specification. How do consumers form their beliefs about future energy prices? Are their beliefs reasonable? How should researchers model consumer beliefs? In this paper, we seek answers to these questions for the important case of gasoline using high-quality survey data that directly elicits consumer beliefs.

Addressing these questions is central to understanding markets for energy-using durable goods and, by extension, for designing policies to curtail greenhouse gas emissions and other energy-related externalities. A large and growing literature tests empirically whether consumers fully value energy efficiency when purchasing durables.<sup>1</sup> These tests are important because if consumers undervalue energy efficiency—as this literature sometimes finds—then policies designed to raise the price of carbon-intensive fuels, such as a carbon tax or cap-and-trade program, may no longer be first-best. In this case, efficiency standards or subsidies may be justified as complements to policies that increase fuel prices (Fischer, Harrington and Parry 2007).

Research that attempts to estimate consumers' valuation of energy efficiency must explicitly model consumers' beliefs about future energy prices and may draw biased inferences if these beliefs are mis-specified. This issue is most relevant for studies using identification strategies that rely on time-series variation in energy prices to identify demand—a strategy that is particularly common in studies of automobile demand (Kahn 1986; Goldberg 1998; Kilian and Sims 2006; Busse et al. 2009; Li, Timmins and von Haefen 2009; Klier and Linn

---

<sup>1</sup>Important papers in this literature include Hausman (1979); Dubin and McFadden (1984); Kahn (1986); Goldberg (1998); Kilian and Sims (2006); Allcott and Wozny (2010); Busse, Knittel and Zettelmeyer (2009); Sallee, West and Fan (2009); Bento, Li and Roth (2010). For a recent survey, see Greene (2010). Note that the “full valuation” hypothesis refers only to the private gains from energy efficiency, not to reductions in externalities.

2010; Whitefoot, Fowlie and Skerlos 2010; Allcott and Wozny 2010; Sallee et al. 2009; Bento et al. 2010; Langer and Miller 2011; Linn and Klier Forthcoming). These studies frequently assume that consumers adopt no-change forecasts for future gasoline prices in real terms; that is, they assume that the expected future price is the current price.<sup>2</sup> If consumer beliefs deviate significantly from this assumption, then researchers may under-estimate or over-estimate consumers' valuation of fuel economy, depending on the direction of the deviation.<sup>3</sup> In addition, if consumer beliefs are “unreasonable,” in the sense that they are systematically biased or have low predictive accuracy in comparison to a benchmark, then these beliefs themselves may constitute a market failure that motivates government intervention.

In lieu of direct evidence, there is perhaps little reason to believe that consumer expectations will align conveniently with the no-change hypothesis favored by applied researchers. Future crude oil and gasoline prices are notoriously difficult to predict, and there is substantial controversy among academic and industry experts about what the future price of oil will be and how best to predict future prices (Hamilton 2009; Alquist and Kilian 2010; Alquist et al. 2010). The main goal of our paper is therefore to test the no-change belief assumption directly.

We conduct our analysis using high-frequency data on consumer beliefs about future gasoline prices from the Michigan Survey of Consumers (MSC). Every month, the MSC asks a nationally representative sample of about 500 respondents to report their beliefs about the current state of the economy and to forecast several economic variables. Since 1993, the MSC has regularly asked respondents to report whether they think gasoline prices will be higher or lower (or the same) in five year's time and then to forecast the exact price change. To the best of our knowledge, we are the first researchers to use this unique cache of information on gasoline price expectations, and we are not aware of a comparable dataset

---

<sup>2</sup>Equivalently, consumers are assumed to believe that gasoline prices follow a martingale process. Throughout the paper, we use the “no-change” terminology as it accords with the literature on oil price forecasting (see, for example Alquist, Kilian and Vigfusson (2010)). We do not use the term “random walk” as a random walk process further implies that the price innovations are iid.

<sup>3</sup>This issue is a specific instance of the broader empirical problem, discussed by Manski (2004), that preferences and expectations are generally not both identified from choice data alone.

or comparable analysis of consumer forecasts of energy prices elsewhere. Moreover, no prior work has attempted to isolate the role of energy price expectations in models testing for full valuation of energy efficiency, and virtually no existing work directly measures consumer beliefs about future energy prices in any context.<sup>4</sup>

Our analysis of these data indicates that in normal economic climates the average consumer expects the future real price of gasoline to equal the current price. That is, in our preferred specifications, consumer beliefs cannot be distinguished statistically from a no-change forecast. This finding is robust to a number of modeling alternatives, and it provides a justification for the common assumption in the automobile demand literature that beliefs follow a no-change forecast. This finding complements the analysis in Anderson, Kellogg, Sallee and Curtin (2011), which concludes that the MSC data have a forecast accuracy similar to that of a no-change benchmark. The main caveat to our conclusion that consumers have a no-change forecast is that beliefs did deviate substantially from a no-change forecast during the 2008 financial crisis, when consumers correctly predicted that gasoline prices would rebound following their sharp decline.

From a modeling perspective, our findings suggest that it is appropriate for researchers to assume that average consumers use a no-change price forecast during normal economic times. Following large shocks, however, beliefs about future gasoline prices may deviate substantially from current prices, consistent with MSC forecasts for other economic variables (Curtin 2007). Another caveat to our conclusion is that there is substantial heterogeneity in stated beliefs across consumers, the dispersion of which varies over time (Anderson et al. 2011). This heterogeneity is a potentially important source of variation in preferences for

---

<sup>4</sup>One recent exception is Allcott (2010), which estimates automobile demand using a specially designed survey instrument that asks consumers to report (among other things) their beliefs about future gasoline prices in real terms. Allcott finds that consumers expect a real price increase on average, whereas we find that consumers expect no price change in real terms. Allcott draws on a single cross section of data from October 2010, however, at which time the MSC series also predicts a small increase in real prices and an even larger increase in nominal prices that is roughly equivalent to the increase in Allcott's survey. Thus, the discrepant results are largely reconciled if respondents in the Allcott survey answer in nominal terms, contrary to instructions (our favored interpretation), or if respondents in the MSC series answer in real rather than nominal terms (Allcott's favored interpretation).

automobile fuel economy that is not fully captured in existing demand models.

Our findings also suggest that consumer beliefs themselves are unlikely to constitute a market failure. While a no-change gasoline price forecast is obviously not perfect, we believe it is a good benchmark for determining whether consumer forecasts are reasonable. A no-change forecast for crude oil is theoretically sensible because rapidly rising or falling prices would induce storage and extraction arbitrage (Hamilton 2009). In addition, no-change forecasts predict future crude oil prices as well as or better than forecasts based on futures markets and surveys of experts (Alquist and Kilian 2010; Alquist et al. 2010). We therefore interpret the statistical similarity between the MSC forecast and the no-change benchmark as evidence that consumers hold reasonable beliefs, implying that consumer beliefs themselves are unlikely to constitute a market failure.<sup>5</sup> The deviation between consumer beliefs and a no-change forecast during the financial crisis of late 2008 lends further credence to this conclusion. A no-change forecast is a good model during normal economic times, when price fluctuations reflect unexpected shifts in long-run supply and demand. When a large temporary shock hits the world economy, however, prices are likely to mean revert. Indeed, the deviation in the MSC data from a no-change forecast during the financial crisis, which closely paralleled a similar deviation in the futures market, turned out to be correct.

The paper proceeds as follows. In section 2 we discuss a model of consumer demand for fuel economy that highlights the importance of gasoline price expectations. In section 3 we describe the MSC data and detail our transformation of the raw data into aggregate measures. Section 4 provides graphical evidence regarding the relationship between current gasoline prices and consumer forecasts. We verify the graphical analysis with regression-based tests in section 5. Section 6 concludes.

---

<sup>5</sup>Our argument here is based on the literature on crude oil, not retail gasoline. Retail gasoline prices may behave differently on short time horizons, but they will be tethered to crude prices over a five-year horizon, which is our focus here. Likewise, retail prices may spike in specific locations due to refinery outages or supply disruptions, at which time it is reasonable to expect mean reversion in prices, but we believe such occurrences will be too rare to influence our average statistics.

## 2 Estimating the demand for automobile fuel economy

To motivate the importance of knowing what consumers actually believe about the future price of gasoline, consider the following standard expression for household utility that serves as the basis for many models of automobile demand:

$$u_{ijt} = -\alpha p_{jt} - \gamma E_t \left[ \sum_{s=0}^T (1+r)^{-s} g_{t+s} m_{ij,t+s} GPM_j \right] + \beta X_j + \xi_j + \varepsilon_{ijt}. \quad (1)$$

Here,  $u_{ijt}$  is the utility that household  $i$  derives from purchasing vehicle  $j$  at time  $t$ ;  $p_{jt}$  is the purchase price of this vehicle;  $E_t[\cdot]$  and its contents, detailed below, are expected fuel costs over the lifetime of the vehicle, in present-value terms;  $X_j$  is a vector of observable vehicle characteristics, such as interior volume and horsepower;  $\xi_j$  is unobservable (to the econometrician) vehicle quality; and  $\varepsilon_{ijt}$  is the idiosyncratic utility that an individual consumer derives from the vehicle.<sup>6</sup> Households are assumed to choose the vehicle model (if any) that gives them the highest utility, facilitating estimation of utility parameters using data on vehicle attributes and household choices. Similar utility models have been used in other energy-intensive durable goods settings, such as purchases of household appliances (Dubin and McFadden 1984).

In any given future time period  $t + s$ , fuel costs equal the number of miles  $m_{ij,t+s}$  the vehicle is driven, multiplied by the vehicle's fuel consumption rate in gallons per mile  $GPM_j$ , multiplied by the future real price of gasoline  $g_{t+s}$ . Discounting at rate  $r$  and summing over the full,  $T$ -period lifetime of the vehicle gives total lifetime fuel costs in brackets. The expectations operator is required because the vehicle's lifetime, future miles driven, and the future real price of gasoline (which embodies expectations about future gasoline prices and inflation) are not known with certainty at the time of purchase.<sup>7</sup> Thus, when trading off the

---

<sup>6</sup> $\varepsilon_{ijt}$  is typically modeled as iid logit or generalized extreme value. Random coefficients logit models that allow for heterogeneity in  $\gamma$  have generally not been used in the energy efficiency valuation literature, though a recent working paper (Bento et al. 2010) has begun to explore the implications of such an approach.

<sup>7</sup>Technically, the vehicle's future fuel consumption per mile (which varies with driving conditions and can degrade over time) and the real rate of discounting from one future period to the next are not known with

purchase price of a vehicle (and other vehicle attributes) against expected lifetime fuel costs, a consumer must consider the fuel efficiency of the vehicle, the number of miles she plans to drive, and the future price of gasoline in real terms.

In this model, testing whether consumers fully value the benefits of fuel economy is equivalent to testing the null hypothesis that  $\alpha = \gamma$ . Empirically implementing this test requires that a researcher populate the expected fuel costs term with each of its underlying components. Fuel consumption per mile for virtually every vehicle sold in the last several decades is readily available to consumers and researchers alike from the Environmental Protection Agency (EPA) based on standardized testing procedures.<sup>8</sup> Estimates for expected vehicle lifetimes (or rather, the probability that a vehicle survives a given number of years) and the number of miles that vehicles are driven are available directly from the National Highway Transportation Safety Administration, can be calculated from the National Household Travel Survey or other surveys, or can be obtained from state administrative datasets, as in Knittel and Sandler (2010). Lastly, discount rates for vehicle purchase decisions can be inferred from market interest rates, including rates on new and used car loans (after adjusting for expected inflation), which are available at the micro level in some vehicle transaction data sets and in aggregate from the Federal Reserve. In short,  $m$ ,  $GPM$ ,  $r$ , and  $T$  (or close approximations thereof) are all readily observable to researchers, if not for individual vehicles and consumers, then at least for broad classes of vehicles and consumers.

In contrast, expected *future* gasoline prices have not been directly observable to researchers in any form. In lieu of direct evidence, applied researchers frequently assume that consumers use a no-change forecast (Busse et al. 2009; Sallee et al. 2009). That is, researchers assume that the expected future real price of gasoline equals the current price, simply replacing future gasoline prices  $g_{t+s}$  in the expression above with the current price  $g_t$ . Less frequently, researchers estimate their own econometric forecast models to predict

---

certainty either. Moreover, miles driven in any future period may depend on the price of gasoline.

<sup>8</sup>These tests are the basis for fuel economy labels that auto dealers are required to place on the windows of new vehicles for sale. The tests are also the basis for verifying automaker compliance with the federal Corporate Average Fuel Economy (CAFE) standards program.



future gasoline prices as a function of current and lagged macroeconomic variables, sometimes specifying a probability distribution for the evolution of future prices (Kilian and Sims 2006). More recently, Allcott and Wozny (2010) assume that expected future gasoline prices equal the price of crude oil in futures markets, plus an add-on to account for refining costs, distribution, marketing, and taxes.

Because fuel consumption per mile is highly correlated with a vehicle’s other attributes, such as engine size, weight, and horsepower, the unique variation in expected fuel costs needed to identify these models comes largely (and sometimes exclusively, to the extent that specific vehicle fixed effects are used) through time-series variation in expected gasoline prices. Thus, correct specification of consumer beliefs about future gasoline prices is crucial to identification of the ratio  $\gamma/\alpha$ . Suppose, for instance, that the researcher models consumers as having a no-change forecast. Under this assumption, whenever the current gasoline price increases by \$1, consumer beliefs about the future price will also increase by \$1. If, however, consumer beliefs about the future price actually increase by less than \$1, then the no-change assumption will lead to an estimate of  $\gamma$  that is biased downward toward zero: consumers will seem under-responsive to lifetime fuel costs. If, on the other hand, consumer beliefs increase by more than \$1, then conventional estimates of  $\gamma$  will be biased upward and consumers will seem over-responsive to lifetime fuel costs. It is this strong dependence of inferences about consumers’ valuation of fuel economy on the specification of their future price beliefs that motivates our study.

## 3 Data

### 3.1 Data sources

Our expectations data come from the Michigan Survey of Consumers (MSC), which every month asks a nationally representative random sample of respondents to state their beliefs about the current state of the economy and to forecast several economic variables. For

example, regarding unemployment expectations, the survey asks: “How about people out of work during the coming 12 months—do you think that there will be more unemployment than now, about the same, or less?” A subset of these questions are aggregated into a single measure known as the University of Michigan Consumer Sentiment Index, which is widely followed as a leading indicator of economic performance.

Each month’s survey is based on responses from approximately 500 households. The survey has a short panel component: about one-third of respondents each month are repeat respondents from six months earlier, another third are new respondents that will be surveyed again in six months, and the final third are new respondents that will never be surveyed again. A core set of questions appears in every survey, but the survey has added and discontinued and even restarted various questions over time, so not all information is available in every time period.

We are primarily interested in two questions related to expected future gasoline prices that appear in nearly every survey dating back to 1993:<sup>9</sup>

**Question:** “Do you think that the price of gasoline will go up during the next five years, will gasoline prices go down, or will they stay about the same as they are now?”

If respondents answer “stay about the same,” their expected price change is recorded as zero. If respondents answer “go up” or “go down,” they are asked a follow-up question:

**Question:** “About how many cents per gallon do you think gasoline prices will (increase/decrease) during the next five years compared to now?”

If consumers report a range of price changes, they are asked to pick a single number. If they refuse or are unable to pick a single number, then the median of their reported range is recorded instead. If consumers respond that they “don’t know” or refuse to respond at any stage of the questioning, then their non-response is noted as such, but only after

---

<sup>9</sup>There are several short gaps in the data availability: November 1993–February 1994, December 1999–February 2000, and January–April 2004.

being prompted several times to give a response. Less than 1% of respondents are coded as non-response. The survey has also asked an identical set of questions about expected twelve-month future gasoline prices since 2006 and occasionally during 1982-1992. We focus here on the five-year forecast because this time horizon is more relevant for automobile demand and because the data coverage is significantly better.

The survey was designed to elicit expectations about gasoline price changes in nominal terms, and there are several compelling reasons to believe that respondents answer in nominal rather than real dollars. First, experienced survey practitioners generally believe that respondents answer in nominal terms unless they are specifically coaxed into a real-price calculation (Curtin 2004). Second, because the questions about gasoline prices follow a series of questions about expected inflation and prices in general, we suspect that consumers are primed to answer in nominal terms. Third, the question asks for gasoline price changes in cents per gallon, so that answering in real terms would require the respondent to make an inflation adjustment calculation. Finally, should respondents ask for clarification, interviewers are instructed to tell respondents to answer in nominal values. Thus, we assume from here on that consumers respond in nominal terms.

In addition to the MSC data, we collected data on gasoline prices from the U.S. Energy Information Administration (EIA). These data record the monthly, sales-weighted average retail price of gasoline (including taxes) by regional Petroleum for Administration of Defense District (PADD) for all grades (regular, midgrade, and premium) and formulations (conventional, oxygenated, and reformulated) of gasoline.<sup>10</sup> We match MSC respondents by state to each of the seven PADD regions contained in the EIA data.

Because we believe that consumers are reporting expected future gasoline prices in nominal terms, we need to deflate these prices by a measure of expected inflation to facilitate comparison to current gasoline prices in real terms. Fortunately, the MSC asks a series of

---

<sup>10</sup>Preliminary analysis indicates that using state-level retail prices for the ten or so states for which these data are available (or for all states, with retail prices reconstructed by adding state gasoline and sales taxes to pre-tax retail prices from EIA) would have very little impact on our results.

questions that allow us to deflate each respondent’s gasoline price forecast using his or her stated beliefs about future inflation. The first question is: “What about the outlook for prices over the next 5 to 10 years? Do you think prices will be higher, about the same, or lower, 5 to 10 years from now?” If respondents answer “about the same,” their expected inflation rate is recorded as zero. If respondents answer “higher” or “lower,” then they are asked a follow-up question: “By about what percent per year do you expect prices to go (up/down) on the average, during the next 5 to 10 years?” (underlining in the original survey codebook). We use the responses to these questions to deflate nominal price forecasts by expected inflation, as described in section 3.2.

Lastly, we collected data on the Consumer Price Index (CPI) from the Bureau of Labor Statistics to put all prices into a common unit.<sup>11</sup> We have complete data on all of these variables—actual gasoline prices, inflation expectations, and CPI—for our study period of January 1993 to December 2009 (except for several short gaps due to missing MSC data).

### 3.2 Data procedures

We construct our variables of interest from these raw data in several steps. Let  $\tilde{C}_{it}^{60}$  be respondent  $i$ ’s expectation at time  $t$  for the change in nominal gasoline prices over the next 60 months (5 years), and let  $\tilde{P}_{it}$  be the nominal price of gasoline in respondent  $i$ ’s PADD. (Henceforth, tildes denote nominal variables.) The expected price change is reported directly in the MSC data, while the current price is given by the EIA retail price data. We use these data to construct respondent  $i$ ’s expectation at time  $t$  for the nominal gasoline price 60 months into the future:

$$\tilde{F}_{it}^{60} \equiv E_{it} \left[ \tilde{P}_{i,t+60} \right] = \tilde{P}_{it} + \tilde{C}_{it}^{60}, \quad (2)$$

which is the nominal price of gasoline plus the expected price change in nominal terms.

Now, let  $r_{it}$  be respondent  $i$ ’s expectation at time  $t$  for the average annual inflation rate

---

<sup>11</sup>We use series CUUR0000SA0LE, which is the non-seasonally adjusted index for all urban consumers, all items less energy.

over the next 60 months. We deflate the expected future nominal price by five years at this expected inflation rate and then deflate again by the realized CPI to construct the expectation at time  $t$  for the real price of gasoline 60 months into the future (in January 2010 dollars):

$$F_{it}^{60} \equiv E_{it}[P_{i,t+60}] = \tilde{F}_{it}^{60} \cdot (1 + r_{it})^{-5} \cdot CPI_{t,Jan2010}, \quad (3)$$

where  $CPI_{t,Jan2010}$  is the CPI inflation factor from time  $t$  to January 2010 (the lack of a tilde on  $F_{it}^{60}$  denotes real dollars). We also convert the current price of gasoline from nominal to real dollars:  $P_{it} \equiv \tilde{P}_{it} \cdot CPI_{t,Jan2010}$ . Deflating the price forecast by five years of expected inflation puts the forecast in time- $t$  dollars for an apples-to-apples comparison with the current price of gasoline at time  $t$ ; deflating both variables by realized inflation puts everything in January 2010 dollars for an apples-to-apples comparison across the many months of the survey.

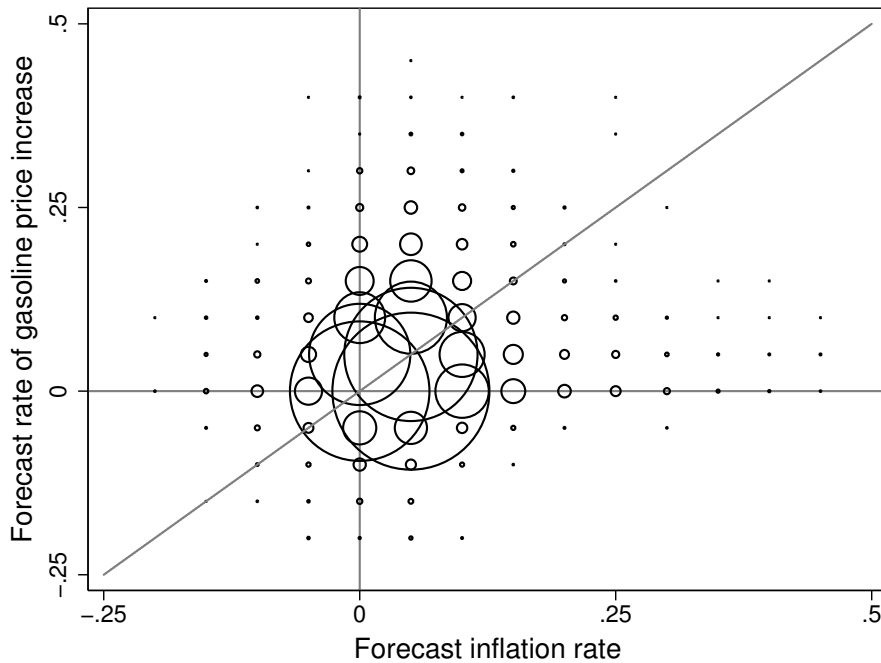
Having constructed both the real price forecast and the real current price of gasoline, we can then construct the expectation at time  $t$  for the real change in gasoline prices over the next 60 months:

$$C_{it}^{60} = F_{it}^{60} - P_{it}, \quad (4)$$

which is simply the real price forecast minus the current real price.

Figure 1 plots the joint pdf for the forecasted annual rate of change in gasoline prices and forecasted inflation rate across all respondents during 1993–2009; circle sizes are proportional to the number of respondents at each point in the distribution. Respondents above the 45-degree line forecast a real-price increase, while those below forecast a real-price decline. The vast majority of respondents are in the northeast quadrant, which indicates that they forecast *nominal* increases in both gasoline prices and overall price levels (inflation). These average expected increases are small: 3.3% per year for gasoline prices and 3.5% per year for inflation. The figure also shows that the distribution is roughly symmetric around the 45-degree line, which implies that the mean forecasted gasoline price change is roughly zero in real terms.

**Figure 1:** Joint pdf of gasoline prices and inflation expectations



Note: Figure plots the joint pdf for the forecasted annual rate of change in nominal gasoline prices and the forecasted inflation rate across all respondents during 1993–2009. The forecasted rate of gasoline price increase is given by  $(1 + \tilde{C}_{it}^{60} / \tilde{P}_{it})^{1/5} - 1$ , where  $\tilde{C}_{it}^{60}$  is the respondent’s forecasted change over five years and  $\tilde{P}_{it}$  is the current price, both in nominal terms. The forecasted inflation rate is reported directly by survey respondents. The figure was created by creating square bins of width 0.05; circles are centered on each bin with circle sizes proportional to the number of respondents within each bin. Respondents above the 45-degree line forecast a real-price increase, while those below forecast a real-price decline.

These descriptive statistics foreshadow our econometric results below: consumers forecast nominal gasoline price increases, but these increases are consistent with constant real prices.

Our main analysis consists of testing whether the average (mean or median) MSC respondent uses a no-change forecast. Our preferred approach is to calculate each individual’s real gasoline price forecast first, as described above, and then take the mean or median in the final step. This approach is superior to deflating average nominal price forecasts by average inflation rates, since the expectation of a ratio does not equal the ratio of expectations. In constructing these mean and median values, we use weights provided by the MSC that cor-

rect for survey sampling issues, such as ownership of multiple phone lines and non-response probabilities, so that our means and medians are representative of all U.S. households.<sup>12</sup>

## 4 Graphical analysis

We depict our price series graphically in figures 2 and 3. The top panel of figure 2 presents the mean current price of gasoline ( $\tilde{P}_t$ ), the mean forecast change over 5 years ( $\tilde{C}_t^{60}$ ), and the mean forecast level ( $\tilde{F}_t$ ) during our study period, all in nominal terms. The mean expected change always exceeds zero and rises with the increase in nominal gasoline prices over this time period. There is generally little month-to-month volatility in the forecast change, except in 2008, when gasoline prices shot up and then plummeted during the financial crisis. This figure suggests that we would reject a null hypothesis of a *nominal* no-change forecast: consumers consistently expect nominal gasoline prices to rise, and the expected change increases with the current nominal price.

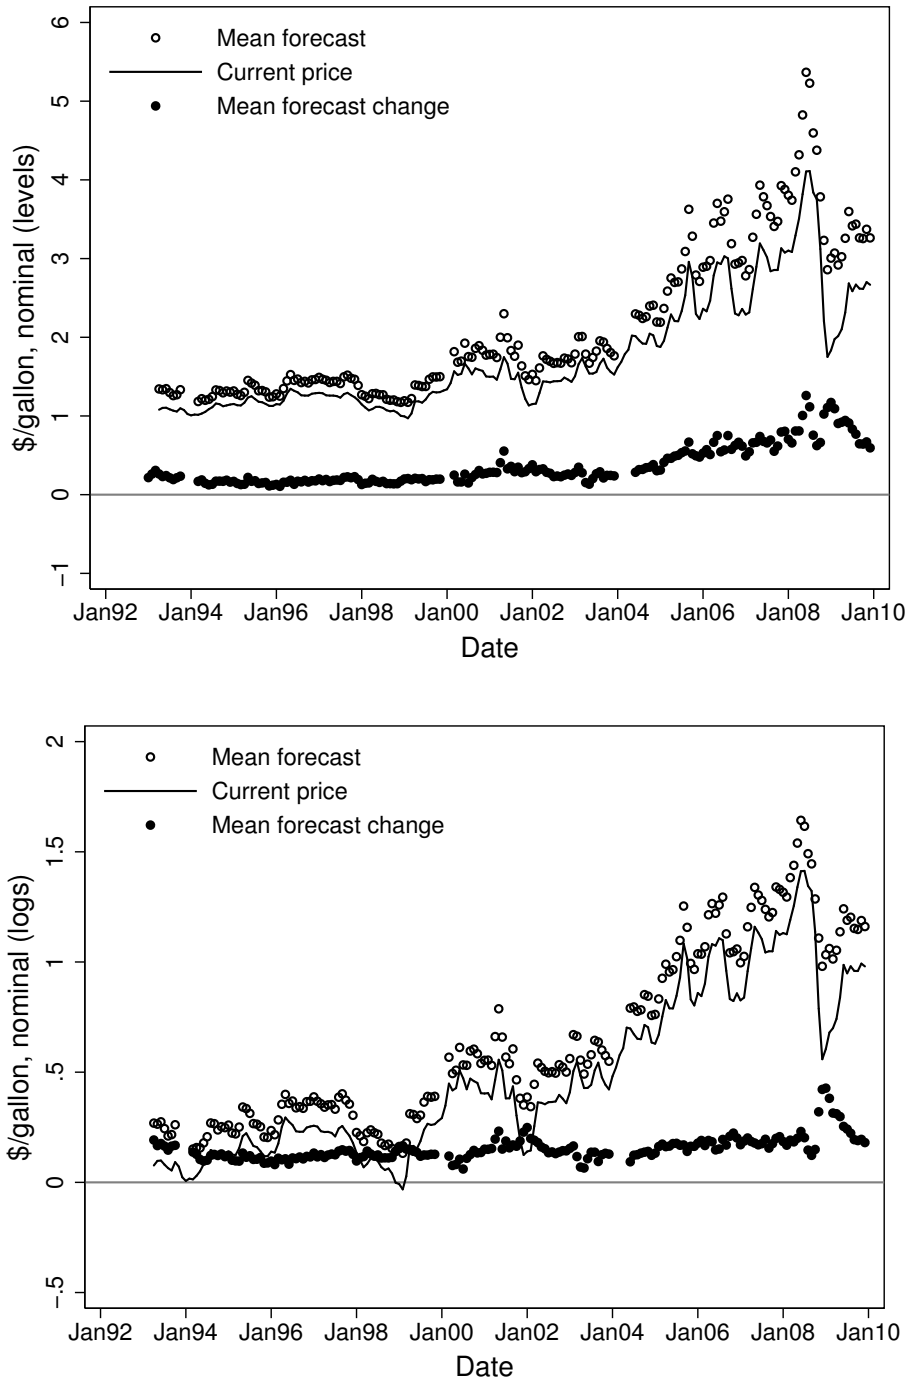
The bottom panel of figure 2 presents logged versions of these series: mean logged prices, mean logged forecasts, and the mean forecast change in logged prices, which we can interpret approximately as the forecast change in percentage terms. The mean forecast change holds steady at roughly 15% (about 3% per year) regardless of price levels, except during the economic crisis. This figure also suggests the data would reject a nominal no-change forecast.

These pictures change considerably after deflating by forecasted inflation. The top panel of figure 3 presents the mean price of gasoline ( $P_t$ ), the mean forecast level ( $F_t^{60}$ ), and the mean forecast change over 5 years ( $C_t^{60}$ ) during our study period, all in real terms. Note that the real forecast change hovers near zero for most of the study period, with large deviations only around September 11, 2001 and the large price swings during the financial crisis of 2008. Thus, this figure indicates that the mean MSC respondent forecasts the real price of

---

<sup>12</sup>Prior to aggregation, we omit a handful of observations for which the rates of increase exceed 50% and the rates of decrease exceed 33% annually, since such rates lead to implausibly high and low price forecasts for these respondents and an explosion in the variance of responses in a handful of months. Omitting these observations does not affect our main conclusions related to average price forecasts.

**Figure 2:** Nominal gasoline prices and forecasts.  
Top panel is in price levels, bottom panel is in logs.





gasoline in 5 years to equal the price at the time of the survey. That is, consumer forecasts appear to be consistent with a *real* no-change forecast model.

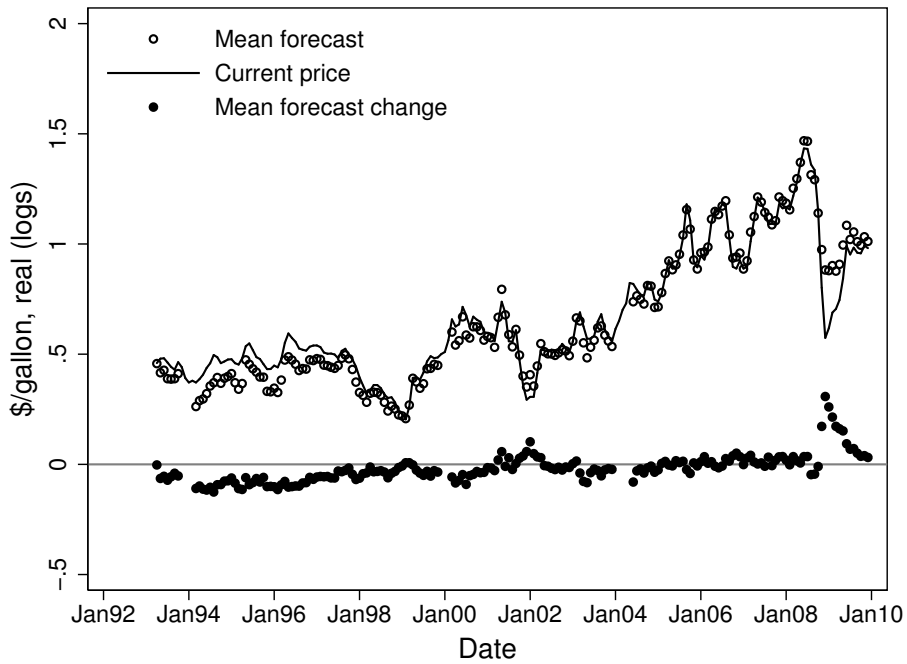
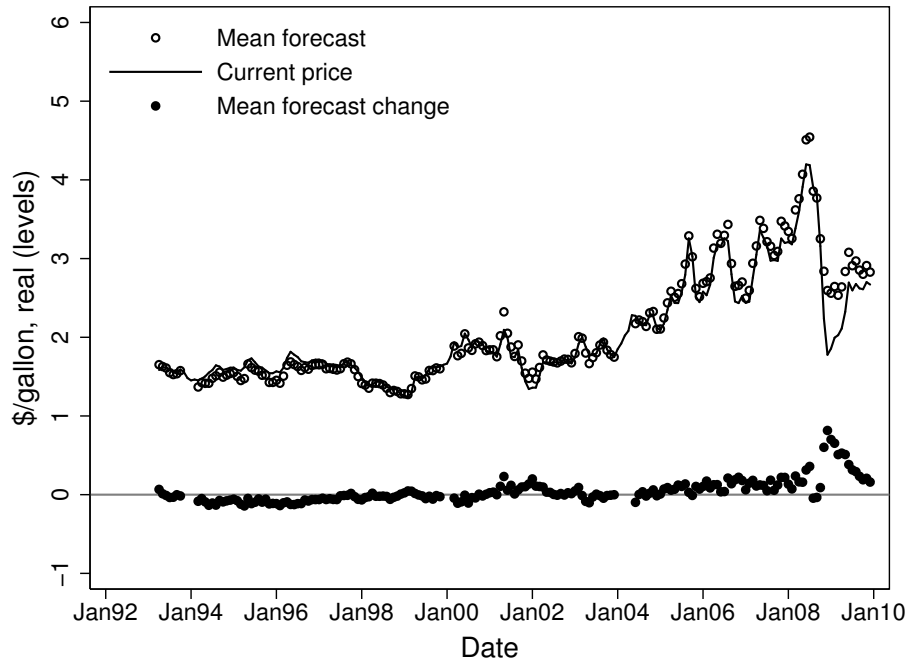
The bottom panel of figure 3 presents the logged versions of these series: mean logged prices, mean logged forecasts, and the mean forecast change in logged prices, which we can again interpret approximately as the forecast change in percentage terms. These logged data follow the same pattern as the real price data in levels: they appear to be consistent with a no-change real forecast. Both sets of figures (and these qualitative inferences) are nearly identical when we analyze sample medians instead of means.

The two characteristics of the nominal data that would lead to a rejection of a no-change forecast—the consistently positive forecast change and the correlation between the expected change and current prices in levels—are both eliminated when we account for expected inflation. Inflation alone will cause prices to rise over a five-year horizon, and a constant rate of inflation will have a larger impact on nominal prices in cents per gallon when the current price is higher.

The only large, sustained deviation from a real no-change forecast is during the financial crisis of 2008, during which the price of gasoline fell by half. Consumers expected prices to rebound quickly. As discussed in Anderson et al. (2011), futures markets predicted a similar rebound. This prediction turned out to be correct: prices had already risen by about one-third of the original decline within six months. Thus, in our sample, when consumers deviate substantially from a real no-change forecast, their deviation is accurate.

While the average respondent appears to use a real no-change forecast, there exists considerable cross-sectional heterogeneity in stated beliefs across consumers. This heterogeneity increases with gasoline prices when measured in levels and is roughly constant during 1993–2007 when measured in logs, although heterogeneity increased in both levels and logs during the financial crisis of 2008 (see Anderson et al. (2011)). During 1993–2007, the average cross-sectional standard deviation in the forecasted change in real prices (measured in logs, so that we can interpret the price changes in percentage terms) is 24%. We are exploring

**Figure 3:** Real gasoline prices and forecasts.  
Top panel is in price levels, bottom panel is in logs.



this cross-sectional heterogeneity and its implications in ongoing research.

## 5 Regression analysis

We now test formally the null hypothesis that the average MSC respondent expects future gasoline prices to equal the current price. Our econometric results will mirror the graphical analysis above, which suggests that the average consumer uses a no-change forecast. A simple way to implement a regression-based test is to regress the expected future price on the current price:

$$F_t^{60} = \beta_0 + \beta_1 P_t + \varepsilon_t \quad (5)$$

and then test the joint null hypothesis that  $\beta_0 = 0$  and  $\beta_1 = 1$ . A variant of this test is to impose that  $\beta_1 = 1$  and then regress the forecast change on a constant:

$$C_t \equiv F_t^{60} - P_t = \beta_0 + \varepsilon_t, \quad (6)$$

where the no-change null hypothesis implies that  $\beta_0 = 0$ . Either approach can be conducted in levels, as written, or in logs. We report results from simple regressions of  $C_t$  on a constant in table 1. Consistent with our graphical analysis, the results indicate that the mean real expected price change is positive when specified in levels (column 1), but this result is driven entirely by data from the financial crisis. When the sample is truncated at December 2007 (in column 3), this simple test fails to reject the no-change null hypothesis. When specified in logs (in columns 2 and 4), this regression indicates that respondents forecast a small percentage decrease in real prices, which appears to be driven by the early part of the sample.

We neither emphasize these results nor report additional specifications because the data demonstrate a high degree of persistence that limits the usefulness of such tests.  $F_t^{60}$  has a first-order autoregressive coefficient of 0.980 in levels and 0.984 in logs, while  $C_t^{60}$  has a

**Table 1:** Does the mean forecast change in gasoline prices equal zero on average?

<b>Panel A: Real gasoline prices and price forecasts</b>				
	Full sample: 1993–2009		Excluding crisis: 1993–2007	
	Levels	Logs	Levels	Logs
Constant	0.0457 (0.0332)	-0.0180 (0.0136)	0.0065 (0.0204)	-0.0320 (0.0103)

<b>Panel B: Nominal gasoline prices and price forecasts</b>				
	Full sample: 1993–2009		Excluding crisis: 1993–2007	
	Levels	Logs	Levels	Logs
Constant	0.3686 (0.0619)	0.1535 (0.0113)	0.2984 (0.0435)	0.1411 (0.0076)

Note: Standard errors (in parentheses) were estimated using Newey-West with 12 lags. Full sample for 1993–2009 includes 189 monthly observations; sample excluding crisis for 1993–2007 includes 165 monthly observations. See text for details.

first-order autoregressive coefficient of 0.880 in levels and 0.891 in logs. The current price of gasoline  $P_t$  has a first-order autoregressive coefficient of 0.973 in levels and 0.978 in logs. Augmented Dickey-Fuller (ADF) tests fail to reject unit roots on the real price and forecast series in nearly all cases.<sup>13</sup> As a result, we prefer estimates from a first-differenced model.<sup>14</sup>

The first-differenced model is also preferred because it is especially relevant for studies of automobile demand that rely on time-series variation in gasoline prices to identify consumers' valuation of fuel economy. Identification in such studies comes from observing how vehicle prices and quantities change in response to changes in gasoline prices (typically interacted with vehicle efficiency). Thus, it is the response of consumers' gasoline price forecasts to *changes* in the current gasoline price that is of primary interest in this paper.<sup>15</sup>

<sup>13</sup>This statement is true regardless of whether we focus on the full or pre-2008 sample, model prices in logs or in levels, analyze means or medians, or allow for a trend or not. We use a version of the ADF test that de-means and de-trends using GLS to increase power according to Elliott, Rothenberg and Stock (1996), and we determine the optimal number of lagged differences using the modified AIC of Ng and Perron (2001).

<sup>14</sup>We are able to reject a unit root in the mean forecast change in two cases: levels in the full sample, logs in the pre-2008 sample. If the price and forecast series are cointegrated, then we could estimate the dynamic response of expected prices to current prices using an error correction model. By imposing that expectations eventually equal current prices, however, this model *assumes* a long-run no-change forecast. We are not always able to reject a unit root in the difference between the two price series, however, supporting the use of a model in first differences.

<sup>15</sup>In this sense, we are testing a null hypothesis that is weaker than a no-change forecast assumption:

In our baseline specifications, we estimate the relationship between the monthly change in current prices and the monthly change in expected future prices using average prices in levels:

$$\Delta F_t^{60} = \beta_0 + \beta_1 \Delta P_t + \varepsilon_t, \quad (7)$$

and using average prices in logs:

$$\Delta f_t^{60} = \beta_0 + \beta_1 \Delta p_t + \varepsilon_t, \quad (8)$$

where the lower-case letters indicate logged price variables.

Table 2 presents our regression results. Panel A presents results with real variables. Results in levels based on the full 1993–2009 sample imply that when the mean current price of gasoline increases by \$1, the mean real-price forecast increases by about \$0.87. Results in logs based on the full sample imply that when the current price of gasoline increases by 1%, the mean real-price forecast increases by about 0.83%. We cannot statistically reject the null hypothesis of a real no-change forecast at a conventional 5% level, but this is partly because estimates using the full sample have sizable standard errors. The point estimates are economically significant. They suggest less-than-full adjustment, so that consumers anticipate mean-reverting gasoline prices.

This result (and much of the imprecision) is driven by the data from the financial crisis of 2008, which led to a large deviation between the current and expected future price. When we limit the sample to the 1993–2007 period, as reported in the right-hand side of the table, the regression coefficients are all tightly estimated and close to 1, consistent with a no-change forecast. Results in levels based on the limited 1993–2007 sample imply that when the current price of gasoline increases by \$1, the mean forecast price increases by \$0.99. Results

---

our null is that consumer forecasts change one-for-one with current gasoline prices. A strong test of the no-change hypothesis would also require that the forecast equals the current price. As shown in table 1 panel A, this hypothesis holds on average throughout the sample. When we look at the individual data one month at a time, however, there exist many months for which we can statistically reject equality between the current and mean forecasted price, which is not surprising given 500 individual observations per period. These deviations from equality are economically small in most periods (per figure 3).

**Table 2:** Does the mean forecast gasoline price change one-for-one with the current price?

**Panel A: Real gasoline prices and price forecasts**

Variable	Full sample: 1993–2009		Excluding crisis: 1993–2007	
	Levels	Logs	Levels	Logs
Current price	0.8742 (0.0895)	0.8250 (0.0933)	0.9938 (0.0430)	0.9624 (0.0340)
Constant	0.0017 (0.0029)	0.0011 (0.0011)	0.0021 (0.0020)	0.0012 (0.0010)

**Panel B: Nominal gasoline prices and price forecasts**

Variable	Full sample: 1993–2009		Excluding crisis: 1993–2007	
	Levels	Logs	Levels	Logs
Current price	1.0899 (0.1137)	0.8804 (0.0902)	1.2571 (0.0299)	1.0147 (0.0268)
Constant	0.0012 (0.0034)	0.0007 (0.0012)	0.0010 (0.0021)	0.0005 (0.0010)

Note: Standard errors (in parentheses) were estimated using Newey-West with 12 lags. Full sample for 1993–2009 includes 185 monthly observations; sample excluding crisis for 1993–2007 includes 161 monthly observations. See text for details.

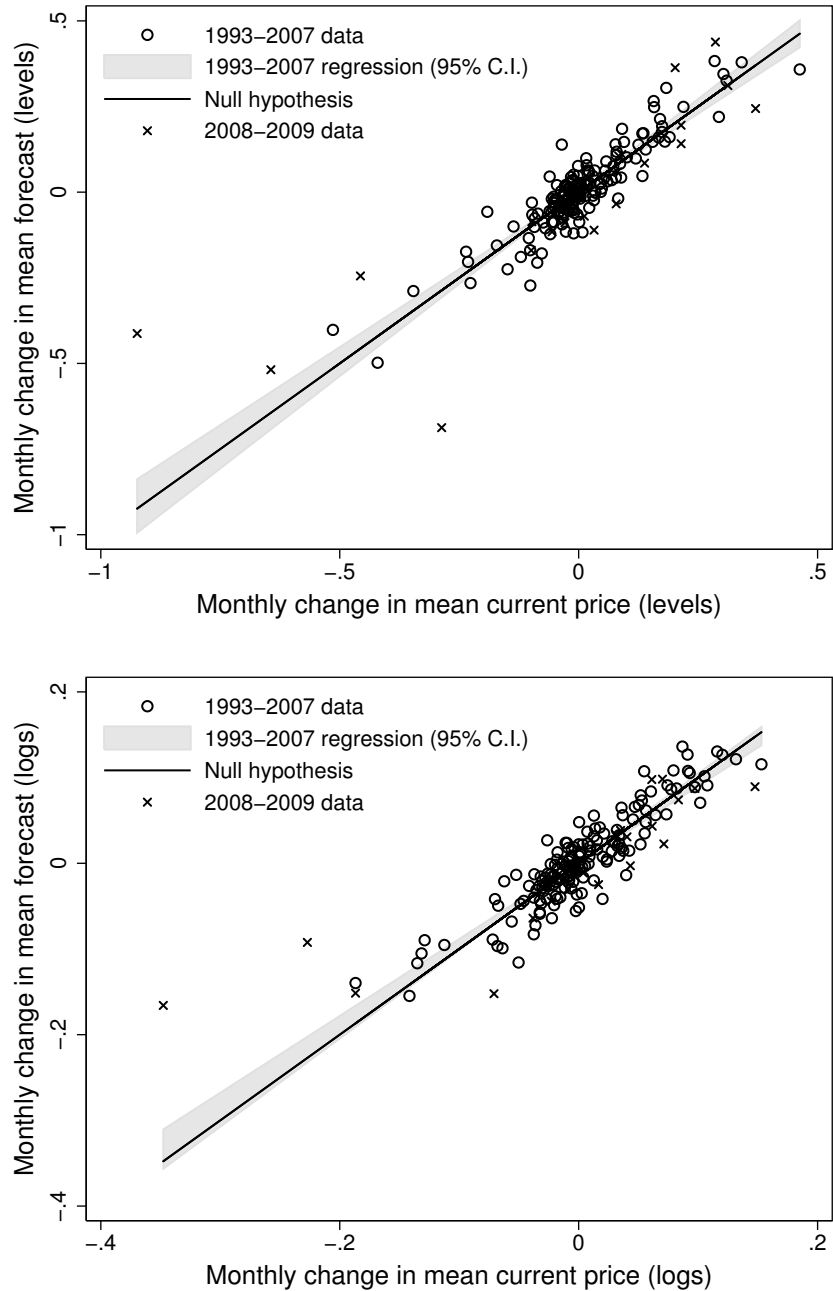
in logs based on the limited sample imply that when the current price increases by 1%, the mean forecast price increases by 0.96%. These coefficients are more precisely estimated than in the full sample, and we are still unable to reject no-change beliefs.<sup>16</sup>

Figure 4 presents these results graphically: forecasted gasoline prices increase one-for-one with current prices on average. To highlight the importance of the financial crisis, data from 2008 and 2009 are denoted with x's. These data clearly have the largest residuals. These figures also illustrate that while the slope is one, the correlation is not perfect, even during normal times. Thus, the current price is a somewhat noisy measure of average beliefs.

In table 2 panel B, we do not adjust for inflation, but rather regress the mean nominal forecast on the mean nominal price. In the full sample, a \$1 increase in the mean current price correlates with a \$1.09 increase in the mean forecast price. Given average expected

<sup>16</sup>Our findings are qualitatively unchanged if we use median instead of mean expectations or if we adjust for inflation using the Philadelphia Federal Reserve's survey of experts.

**Figure 4:** Does the mean forecast gasoline price change one-for-one with the current price?  
 Top panel is in price levels, bottom panel is in logs.



Note: Figure shows graphically the regression results in panel A of table 2 above. Top figure corresponds to the regression in levels in column (3); bottom figure corresponds to the regression in logs in column (4). See table and text for details.

inflation of about 3.5% per year, this finding implies that real prices only increase by about \$0.92, consistent with the estimates from panel A based on real prices. The estimates using logged nominal prices and forecasts are quite similar to those in panel A, consistent with the fact that average expected inflation is fairly constant over time and that multiplication by a constant has no effect on the coefficient estimate in a log-log model.

After excluding 2008–2009, the linear model suggests that a \$1 increase in the mean nominal current price leads to a \$1.26 increase in the mean nominal forecast price, consistent with a real increase of about \$1.06 after deflating for five years at 3.5% expected annual inflation. The model in logs suggests that a 1% increase in the mean current price leads to a 1.01% increase in the mean expected future price, a result that is not statistically different from 1% and consistent with the corresponding result from panel A.

The results in table 2 estimate the immediate response of expectations to changes in the current price of gasoline. It is possible that the long-run response to price changes differs from the short-run response, perhaps because consumers only update their expectations about future gasoline prices and inflation periodically. Thus, we estimate autoregressive distributed lag (ARDL) models that allow expectations to respond to both current and lagged changes in the gasoline price. That is, we estimate dynamic models of the form:

$$\Delta F_t^{60} = \beta_0 + \sum_{k=0}^q \beta_k \Delta P_{t-k} + \sum_{k=1}^q \gamma_k \Delta F_{t-k}^{60} + \varepsilon_t. \quad (9)$$

Table 3 presents the results of these regressions. Depending on the particular price series we used, we found that it was necessary to include up to 12 periods of lagged prices and forecasts to eliminate the serial correlation in the error term. Thus, all of our results are based on an ARDL model with 12 lags (i.e.,  $q = 12$  in the equation above). The table presents the long-run response of expectations to a permanent increase in the price of gasoline. As the table demonstrates, accounting for a delayed response narrows the gap between estimates based on the full and pre-crisis samples. Long-run responses estimated using the 1993–2007



**Table 3:** Does the mean forecast gasoline price change one-for-one with the current price over the long run?

<b>Panel A: Real gasoline prices and price forecasts</b>				
	Full sample: 1993–2009		Excluding crisis: 1993–2007	
Variable	Levels	Logs	Levels	Logs
Current price	0.9838 (0.0611)	0.8992 (0.0533)	1.0263 (0.0789)	0.9660 (0.0820)

<b>Panel B: Nominal gasoline prices and price forecasts</b>				
	Full sample: 1993–2009		Excluding crisis: 1993–2007	
Variable	Levels	Logs	Levels	Logs
Current price	1.2900 (0.1037)	0.9519 (0.0616)	1.2608 (0.0925)	0.9628 (0.0762)

Note: Table reports long-run response of mean forecast to a permanent increase in the current price of gasoline; standard errors are in parentheses. All regressions assume a lag structure of 12 months. Full sample for 1993–2009 includes 143 monthly observations; sample excluding crisis for 1993–2007 includes 119 monthly observations. See text for details.

pre-crisis sample are very similar to the immediate responses in the previous table, while responses estimated using the full 1993–2009 sample are now much closer to a no-change forecast.<sup>17</sup> These results lend further support to our inference that average consumers use a no-change forecast.

## 6 Conclusion

Do consumers exhibit a reasonable forecast of future energy prices? Our analysis suggests that they do, at least in the important case of gasoline. Using two decades of high-quality survey data from the Michigan Survey of Consumers, we find that consumers, on average, report forecasts that are consistent with a real no-change model of future gasoline prices. This finding suggests that consumers have reasonable beliefs about future prices. Thus, if consumers do undervalue fuel economy, the undervaluation likely does not stem from consumers

<sup>17</sup>We also examined the impulse response functions associated with a permanent increase in the current price of gasoline. We found, however, that the long-run responses occurred quite quickly: in most cases, the short-run effect was statistically indistinguishable from the long-run effect upon impact. This finding is consistent with the fact that the long-run coefficients in table 3 are fairly similar to the coefficients in table 2, which measure the immediate impacts of changes in the current price on expected future prices.

having systematic bias in their beliefs about energy prices. Moreover, this finding implies that applied researchers are likely justified in assuming that average consumers employ a no-change forecast when modeling consumer demand for energy-using durables.

There are two important caveats to this conclusion. First, as we have shown, expectations deviated substantially from a no-change forecast during the financial crisis of 2008 and beyond. Researchers should use caution when estimating demand for automobiles as a function of current gasoline prices during this time period. Second, while we have not explored this issue in significant detail here, our data reveal substantial heterogeneity in beliefs across respondents. This heterogeneity may in turn cause consumers to form different valuations of fuel economy, an issue that has not been addressed in most of the fuel economy valuation literature (Bento et al. (2010), which uses a random coefficients demand model, is an exception). In addition, the cross-sectional dispersion in price forecasts correlates positively (in levels) with the current price of gasoline and sharply increased (even in logs) during the financial crisis. Typical implementations of random coefficient models do not, however, parameterize the dispersion of preference coefficients as functions of covariates in any setting. To further explore these issues, in ongoing work we are examining how variation in individual beliefs about future gasoline prices correlates with stated and revealed preferences for efficient versus inefficient cars and how these correlations vary over time.

## References

- Allcott, Hunt**, “Beliefs and Consumer Choice,” December 2010. Manuscript.
- **and Nathan Wozny**, “Gasoline Prices, Fuel Economy, and the Energy Paradox,” March 2010. Manuscript.
- Alquist, Ron and Lutz Kilian**, “What Do We Learn From the Price of Crude Oil Futures?,” *Journal of Applied Econometrics*, June/July 2010, 25 (4), 539–573.

– , – , and **Robert J. Vigfusson**, “Forecasting the Price of Oil,” in Graham Elliot and Allan Timmerman, eds., *Handbook of Economic Forecasting*, Vol. Forthcoming, Elsevier, 2010.

**Anderson, Soren T., Ryan Kellogg, James M. Sallee, and Richard T. Curtin**, “Forecasting Gasoline Prices Using Consumer Surveys,” *American Economic Review Papers & Proceedings*, 2011, *Forthcoming*.

**Bento, Antonio M., Shanjun Li, and Kevin Roth**, “Is There an Energy Paradox in Fuel Economy? A Note on the Role of Consumer Heterogeneity and Sorting Bias,” November 2010. Resources for the Future Discussion Paper 10-56.

**Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer**, “Pain at the Pump: The Differential Effect of Gasoline Prices on New and Used Automobile Markets,” December 2009. NBER Working Paper 15590.

**Curtin, Richard T.**, “Psychology and Macroeconomics,” in Robert L. Kahn, F. Thomas Juster, James S. House, Howard Schuman, and Eleanor Singer, eds., *A Telescope on Society: Survey Research and Social Science at the University of Michigan & Beyond*, Ann Arbor: University of Michigan Press, 2004.

– , “Consumer Sentiment Surveys: Worldwide Review and Assessment,” *Journal of Business Cylce Measurement and Analysis*, July 2007, 4 (2), 1–37.

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

**Elliott, Graham, Thomas J. Rothenberg, and James H. Stock**, “Efficient Tests for an Autoregressive Unit Root,” *Econometrica*, July 1996, 64 (4), 813–836.

- Fischer, Carolyn, Winston Harrington, and Ian W.H. Parry**, “Should Automobile Fuel Economy Standards Be Tightened?,” *The Energy Journal*, 2007, 28 (4), 1–29.
- Goldberg, Pinelopi Koujianou**, “The Effects of the Corporate Average Fuel Efficiency Standards in the U.S.,” *The Journal of Industrial Economics*, March 1998, 46 (1), 1–33.
- Greene, David L.**, “How Consumers Value Fuel Economy: A Literature Review,” March 2010. U.S. Environmental Protection Agency report EPA-420-R-10-008.
- Hamilton, James D.**, “Understanding Crude Oil Prices,” *The Energy Journal*, 2009, 30 (2), 179–206.
- Hausman, Jerry A.**, “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables,” *The Bell Journal of Economics*, Spring 1979, 10 (1), 33–54.
- Kahn, James A.**, “Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach,” *The Quarterly Journal of Economics*, 1986, 101 (2), 323–339.
- Kilian, Lutz and Eric R. Sims**, “The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data,” April 2006. Manuscript.
- Klier, Thomas and Joshua Linn**, “New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard,” December 2010. Resources for the Future Discussion Paper 10-50.
- Knittel, Christopher R. and Ryan Sandler**, “Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins,” October 2010. NBER Working Paper 16482.
- Langer, Ashley and Nathan H. Miller**, “Automakers’ Short-Run Responses to Changing Gasoline Prices and the Implications for Energy Policy,” March 2011. Manuscript.

- Li, Shanjun, Christopher Timmins, and Roger H. von Haefen**, “How Do Gasoline Prices Affect Fleet Fuel Economy?,” *American Economic Journal: Economic Policy*, August 2009, *1* (2), 113–137.
- Linn, Joshua and Thomas Klier**, “The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data,” *American Economic Journal: Economic Policy*, Forthcoming.
- Manski, Charles F.**, “Measuring Expectations,” *Econometrica*, September 2004, *72* (5), 1329–1376.
- Ng, Serena and Pierre Perron**, “Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power,” *Econometrica*, November 2001, *69* (6), 1519–1554.
- Sallee, James M., Sarah E. West, and Wei Fan**, “Consumer Valuation of Fuel Economy: A Microdata Approach,” *Proceedings of the National Tax Association Annual Conference on Taxation*, 2009.
- Whitefoot, Katie, Meredith Fowlie, and Steven Skerlos**, “Evaluating U.S. Reformed Corporate Average Fuel Economy Standards using an Engineering Attribute Selection Model,” October 2010. Manuscript.