What Do Consumers Believe About Future Gasoline Prices?*

Soren T. Anderson†

Michigan State University and NBER

Ryan Kellogg‡

University of Michigan and NBER

James M. Sallee§

University of Chicago and NBER

November 3, 2010

PRELIMINARY DRAFT: PLEASE DO NOT CITE WITHOUT PERMISSION

Abstract

How do consumers form beliefs about future energy prices? Do they form these beliefs in a "reasonable" way? The answer to this question is key to an ongoing debate regarding the "energy paradox" – which stipulates that consumers systematically undervalue energy efficiency – as well as having important implications for all models of consumer demand for energy-consuming durable goods. We study these questions directly by analyzing two decades of survey data from the Michigan Survey of Consumers, which asks a nationally representative sample of consumers about their expectations of future gasoline prices. Overall, we find that consumer beliefs follow a random walk, which we deem a reasonable forecast of gasoline prices, but we find a deviation from the random walk during the recent economic crisis of 2008-2009.

*The authors would like to thank Richard Curtin for guidance in using the data. For helpful comments and suggestions, we thank Lutz Kilian and seminar participants at the University of California Energy Institute, the Iowa State Bioenergy Camp and Michigan State University.

†Address: Soren T. Anderson, Department of Economics and Department of Agricultural, Food, and Resource Economics, Michigan State University. Email: sta@msu.edu. Web: http://www.msu.edu/∼sta.

‡Ryan Kellogg, Department of Economics, University of Michigan. Email: kelloggr@umich.edu; Web: http://www-personal.umich.edu/∼kelloggr.

§James M. Sallee, The Harris School, University of Chicago. Email: sallee@uchicago.edu; Web: http:/home.uchicago.edu/∼sallee.

1 Introduction

When consumers choose among energy-intensive durable goods, they have to forecast the future price of energy in order to determine their willingness to pay for energy efficiency. In turn, researchers modeling the demand for such durables must specify these consumer beliefs. How do consumers form these beliefs? Are their beliefs "reasonable"? How should researchers model consumer forecasts?

This paper directly answers these questions by analyzing two decades worth of high quality survey data on consumers stated beliefs about gasoline prices from the Michigan Survey of Consumers. The survey asks a nationally representative sample of households every month what they believe will happen to the price of gasoline in the future. To the best of our knowledge, we are the first to utilize this unique cache of information, and there is no comparable data available from other sources.

How consumers forecast future energy prices relates to an ongoing debate about the optimal policy response to externalities from the consumption of energy. If consumers fully value energy efficiency, then the optimal solution to energy-related externalities is a direct Pigouvian tax on emissions or fuels. If consumers over or undervalue fuel efficiency, however, then alternative policies like efficiency standards or durable taxation may be appropriate.

The proper valuation of energy efficiency requires consumers to calculate the present discounted value of the future flow of energy savings from improved efficiency. This calculation involves various components, including discounting, future energy prices and usage. The literature on consumer valuation has generally built econometric models that jointly tests hypotheses about discounting, future energy prices and usage (Hausman 1979; Dubin and McFadden 1984; Allcott and Wozny 2009; Busse, Knittel and Zettelmeyer 2009; Sallee, West and Fan 2009). No prior work has attempted to isolate the role of price expectations.

Moreover, no existing work measures consumer beliefs about future fuel prices, which is relevant well beyond the question of consumer valuation of efficiency. Any researcher modeling consumer choice among durable goods must specify consumers beliefs over future prices. A common assumption is that prices follow a random walk – this implies that today's prices are the best available estimate of future prices. Is this a good approximation of consumer beliefs?

Our analysis of the Michigan Survey of Consumer data on gasoline price forecasts indicates that the average consumer forecast looks very similar to a random walk. In our preferred specifications, consumer forecasts cannot be statistically distinguished from a random walk. In the paper, we employ a variety of robustness checks. The main caveat to the random walk result is that forecasts deviated significantly from a random walk during the 2008 financial crisis.

While a random walk is obviously not a perfectly accurate forecast, it historically has approximately equal predictive accuracy to futures prices (Alquist and Kilian 2010). Thus, our results indicate that consumer beliefs are "reasonable", making it unlikely that consumer beliefs about future gasoline prices are the source of fuel efficiency undervaluation or that these beliefs should be the target of policy. Moreover, these findings provide justification for the common practice of assuming random walk beliefs in models of consumer demand for durables.

2 Estimating the demand for automobile fuel economy

To highlight the importance of knowing what consumers actually believe about the future price of gasoline, consider the following model of consumer demand for automobiles:

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

where: u_{ijt} is the utility that consumer i derives from purchasing vehicle j at time t; p_{jt} is the purchase price of this vehicle; $E_t[\cdot]$ and its contents 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 or horsepower; ξ_j is unobservable (to the econometrician) vehicle quality; and ε_{ijt} is the idiosyncratic utility that an individual consumer derives from the vehicle.

In any given future time period $t + s$, fuel costs equal the number of miles $m_{ij,t+s}$ the vehicle is driven in that time period, multiplied by the vehicle's rate of fuel consumption in gallons per mile GPM_j, multiplied by the price of gasoline in that future time period 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 last step is to take expectations (at period t) over this entire bracketed term, since a vehicle's lifetime, future miles, 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 – nor technically are the vehicle's future rate of fuel consumption per mile (which varies with driving conditions and can degrade over time) or the real rate of discounting from one future period to the next. Thus, when trading off the 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 $\alpha = \gamma$. Of course, implementing this test empirically 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 EPA fuel economy testing procedures.[1] Estimates for expected vehicle lifetimes (or rather, the probability that the average vehicle of a particular class survives a given number of years) and the number of miles that vehicles are driven on average over their lifetimes are available directly from the National Highway Transportation Safety Administration, or can be calculated from surveys like the

[1]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.

National Household Travel Survey or from emissions test data, as in Knittel and Sandler (2010). Lastly, discount rates for vehicle purchase decisions can be inferred from interest rates, including rates on new and used car loans (after adjusting for expected inflation), which are available at the microlevel in some vehicle transaction data sets and in the aggregate from the Federal Reserve. In short, GPM_j, T, $m_{ij,t+s}$, and r — and close approximations of their expected values — are all readily observable to researchers, if not for individual vehicles and consumers, then at least for broad classes of vehicles and consumers.

Unfortunately, however, expected future gasoline prices, which are critical for calculating expected future fuel costs, are not directly observable to researchers. Thus, unable to observe these expectations directly, researchers in this literature almost always assume that consumers believe gasoline prices will follow a "random walk" (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 sometimes estimate their own econometric forecast models to predict future gasoline prices as a function of current and lagged macroeconomic variables. More recently, Allcott and Wozny (2009) assume that expected future gasoline prices equal the futures price of crude oil in commodity futures markets, plus an add-on to account for refining and distribution costs and taxes.

While fuel consumption per mile, vehicle lifetimes, future miles traveled, and discount rates are all directly observable or are sufficiently stable over time that they can be inferred accurately from historical data, this is not true of future gasoline prices. 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 oil prices (Hamilton 2009; Kilian 2009; Alquist and Kilian 2010). Thus, in lieu of direct evidence, there is perhaps little reason to believe that consumer expectations will conveniently align with the random-walk hypothesis favored by applied researchers in this literature. One goal of our paper is to test this assumption directly.

Whatever the method for constructing consumer expectations for the future price, they all suffer from the same criticism: they replace consumer expectations for the future price of gasoline, which are unobservable, with the researcher's own assumption about what consumers believe. Take for example the case of the random walk assumption. Even if most academic and industry experts believe that that gasoline prices approximately follow a random walk, and even if alternative econometric forecast models do not perform much better, this does not imply that consumers actually believe prices will follow a random walk. Thus, assuming that consumers expect a random walk could lead to biased estimates of automobile demand if consumer expectations do not comport with this assumption. For example, if when gasoline prices increase by \$1, consumer expectations of the future price increase by more than \$1, then the random walk assumption will lead to an estimate of γ that is biased upward, away from zero: consumers will appear more sensitive to gasoline prices than they really are. If, on the other hand, consumer expectations for the future price of gasoline increase by less than \$1 when the current price increases by \$1, then conventional estimates of γ will be biased downward, toward zero: consumers will appear less sensitive to gasoline prices than they really are.

3 Data source

Our main data source is the Michigan Survey of Consumers (MSC), which surveys a nationally representative random sample of consumers each month, asking how consumers view prospects for their own financial situation, prospects for the economy in general over the near term, and prospects for the economy in general over the long term. Questions include, for example, "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?" and "Speaking now of the automobile market — do you the next 12 months or so will be a good time or a bad time to buy a vehicle?" 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 6 months prior, another third are new respondents that will be surveyed again in 6 months, and the final third are new respondents that will never be surveyed again. A core set of questions appears in every survey, dating to the beginning of the survey in 1982. Individual researchers and other parties sometimes add questions for one or several rounds of the survey, however, and 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 a single question related to expected future gasoline prices, which appears in nearly every survey dating back to 1993 (with several short gaps) and then sporadically back to 1982: "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?" The end goal of the questioning is to elicit information about the expected price change over the next five years. If respondents answer "stay about the same," then their expected price change is recorded as zero. If respondents answer "go up" or "go down," then they are asked a follow-up question: "About how many cents per gallon do you think gasoline prices will (increase/decrease) during the next five years compared to now?" Their expected price changes are then recorded accordingly. 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. Less than 1% of respondents are coded as non-response. Although we focus on expectations about gasoline prices in five years, which is the most relevant for automobile purchases, the survey has asked an identical set of questions about expected gasoline prices during the next twelve months since 2006 and occasionally during 1982-1992.

From the wording of initial question, which asks whether consumers expect prices to "go up...or down," it is somewhat unclear whether consumers are reporting expected price changes in real or nominal terms. From the follow-up question, however, which asks by "how many cents" prices will change, it seems clear that consumers are being asked to report nominal price changes. Moreover, because the questions about gasoline prices follow several questions about expected inflation and prices in general, we suspect that consumers are primed to report gasoline prices in nominal terms. Thus, we assume from here on that consumers respond in nominal terms.

In addition to these MSC data on consumer expectations, we also collected data on contemporaneous gasoline prices from the U.S. Energy Information Administration (EIA). These data record the monthly, U.S. sales-weighted average retail price of gasoline (including taxes) for all grades of gasoline (regular, midgrade, and premium) and all formulations (conventional, oxygenated, and reformulated).[2] Ideally, we would like to be able to match individual respondents in the MSC data to state-level gasoline prices based on the respondents' state of residence. State-level gasoline prices are especially important when constructing individual expectations for future gasoline prices in real terms, as we describe below. At this point, however, we do not yet have access to state of residence for the respondents in the MSC data. We hope to obtain these data soon.

Because we assume that consumers are reporting expected future gasoline prices in nominal terms, we need to deflate these prices by some measure of expected inflation to facilitate comparison to the current price of gasoline. Fortunately, the questions about gasoline prices follow an earlier and analogous set of questions asking consumers directly about their inflation expectations: "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 per-

[2]These data are available at: http://tonto.eia.doe.gov/dnav/pet/pet_pri_gnd_dcus_nus_m.htm.

cent 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). To test the robustness of these MSC expected inflation measures, we also collect quarterly data on expected average inflation for the next 10 years from the Federal Reserve Bank of Philadelphia. These data report the median response from a survey of professional forecasters (SPF) conducted quarterly by the Bank. We linearly interpolate the quarterly data to the intervening months.

Lastly, we collect data on the Consumer Price Index (for all goods, urban consumers) from the Bureau of Labor Statistics. We have complete data on each of these variables for our study period of January 1993 to January 2010 (except for several small gaps due to missing MSC data).

4 Data procedures

Let $\tilde{C}_t^{60}$ be the expectation at time t for the change in gasoline prices over the next 60 months (5 years) and let $\tilde{P}_t$ be the current gasoline price, both in nominal terms (as denoted by the tildes above the price 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 the expectation at time t for the nominal gasoline price 60 months into the future:

$$\tilde{F}_t^{60} = E_t\left[\tilde{P}_{t+60}\right] = \tilde{P}_t + \tilde{C}_t^{60}, \tag{2}$$

which is simply the nominal price of gasoline plus the expected price change in nominal terms. Now, let r_t be the expectation at time t for the average annual inflation rate over the next 60 months (technically, 5-10 years). 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):

$$F_t^{60} = E_t[P_{t+60}] = \tilde{F}_t^{60} \cdot (1 + r_t)^{-5} \cdot CPI_{t,Jan2010}, \quad (3)$$

where $CPI_{t,Jan2010}$ is the CPI inflation factor from time t to January 2010 (and the lack of tilde denotes real dollars). We also convert the current price of gasoline from nominal to real dollars: P_t. Deflating the price forecast by five years of expected inflation puts the forecast in time t dollars for an apples-to-apples comparison to the actual price of gasoline at time t; deflating both by realized inflation puts everything in January 2010 dollars for an apples-to-apples comparison over time. Having constructed both the real price forecast and the real retail price of gasoline, we can then construct the expectation at time t for the real change in gasoline prices over the next 60 months:

$$C_t^{60} = F_t^{60} - P_t, \quad (4)$$

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

Our main analysis consists of testing whether the average (mean or median) MSC respondent believes that gasoline prices follow a random walk. We construct the "average" belief about future gasoline prices in several alternative ways. Our preferred approach is to calculate each individual's nominal and real gasoline price forecast, according to the equations above, based on individual MSC responses for expected price changes and expected inflation. Then, in the final step, we take the mean or median of these individual real price forecasts. The advantage of this approach is that it accounts for any correlation between individual expectations about future nominal price changes and expected inflation rates. One disadvantage of this approach is that we do not observe individual state of residence, which requires us to construct individual price forecasts based on national average retail gasoline prices. Not observing individual gasoline prices could be a problem in equation 3 above, since the individual-level measurement error enters nonlinearly, interacting with individual inflation expectations.

In addition to our preferred approach, we also construct the "average" belief about future gasoline prices by calculating the mean (or median) expected price change in the MSC responses first, adding this change to the current retail price to get the mean (or median) expected future nominal price, and then deflating this nominal price forecast by the mean (or median) expected inflation rate in the MSC responses. This approach does not account for possible correlation between individual nominal price expectations and individual expected inflation. Lastly, we construct the "average" belief using the mean (or median) expected price change in the MSC responses but deflating by the median expected inflation rate amongst professional forecasters based on a survey by the Philadelphia Fed.

Thus, we construct the "average" real price forecast in six different ways: three ways of deflating nominal price forecasts by expected inflation (each MSC respondent individually using his or her own expected inflation, in aggregate using the average MSC expected inflation, or in aggregate using the Philadelphia Feds survey of professional forecasters) by two ways of calculating the "average" response (mean or median).

5 Graphical analysis

Figures 1 – 3 show the price series we use in our analysis. Our presentation of these figures follows our step-by-step discussion above for how we construct our main price series. Since we have several alternative ways of constructing these price series, we focus on median MSC price expectations deflated by median MSC inflation expectations. These figures look similar if we use means instead of medians, deflate individual price expectations by individual inflation expectations and then calculate the median, or deflate the median MSC price expectations by the median inflation expectation from the Philadelphia Fed.

Figure 1a presents the nominal retail price of gasoline during our study period, along with the median MSC response for the expected nominal price change over 5 years. These series

Figure 1: Nominal retail gasoline prices and nominal MSC forecast

(a) Expected nominal price change

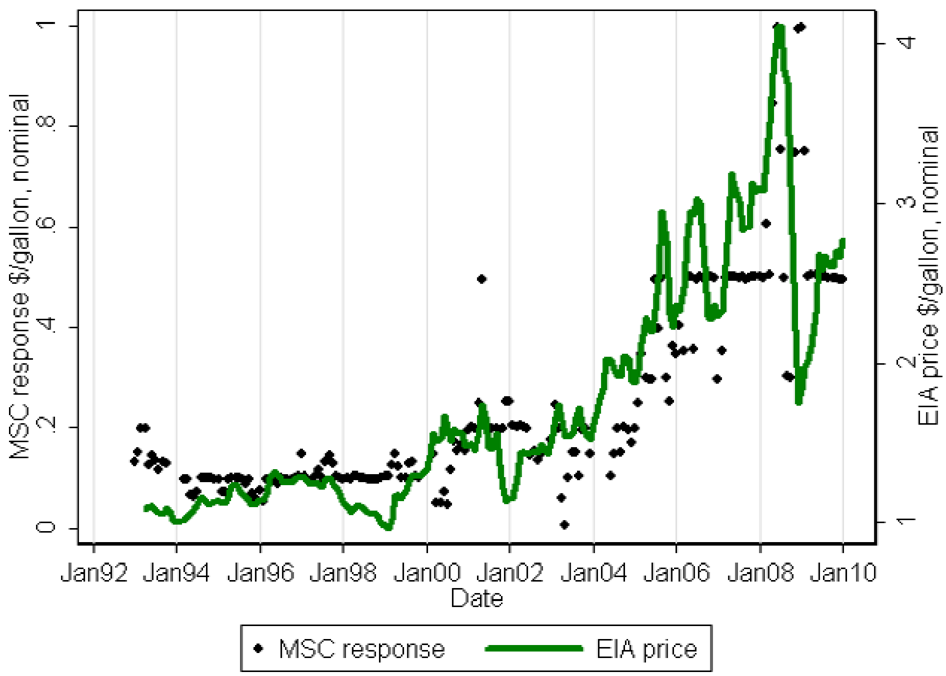

(b) Expected nominal price

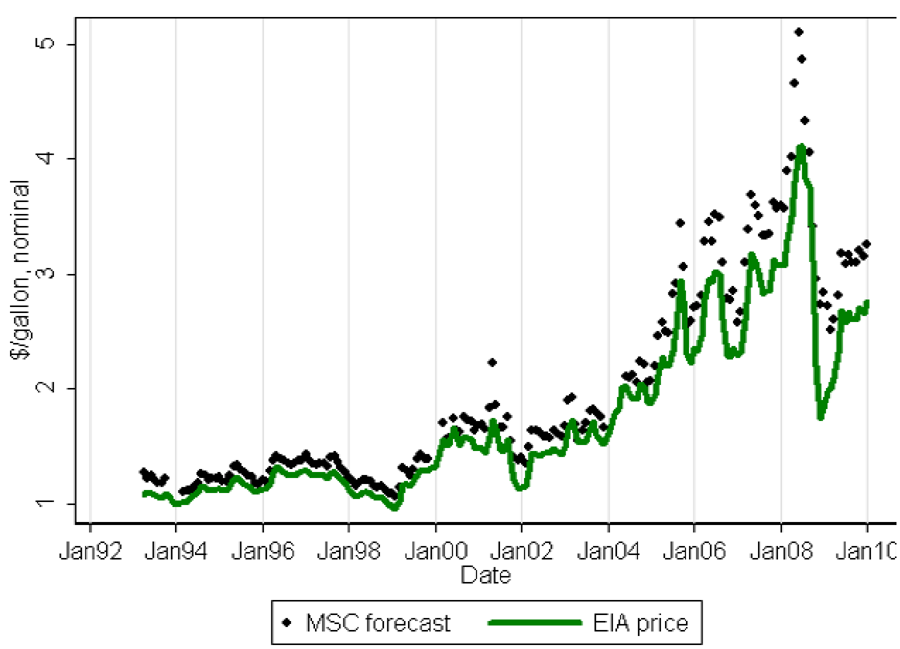

correspond to the nominal price of gasoline $\tilde{P}_t$ and the expected nominal price change $\tilde{C}_t^{60}$ in terms of the notation above. The median expected change is always strictly greater than zero and increases with the increase in nominal gasoline prices over the period. There is generally little month-to-month noise in this forecast, with the exception of the period in 2008 when gasoline prices spiked and then crashed during the financial crisis. The consistently positive expected change is also apparent in figure 1b, which plots the nominal MSC price forecast $\tilde{F}_t^{60}$ rather than just the MSC nominal price change.

Figure 2a presents the real price of gasoline during our study period, P_t, along with the median MSC response for the expected nominal price change over 5 years, C_t^{60}. The expected real price change hovers near zero for most of the study period, with large deviations only around September 11th, 2001 and the period of large gas price swings in 2008. Figure 2b plots the median MSC real forecast gasoline price F_t^{60} rather than the MSC price change and emphasizes that the median MSC respondent generally predicts that the real gasoline price in 5 years will be unchanged from the real price at the time of the survey. Of course, this equality does not hold exactly, and deviations on the order of 10 to 15 cents per gallon can persist for months or years. Still, this figure is qualitatively consistent with the median consumer believing that gasoline prices follow a random walk in real terms.

Figure 3 compares the median MSC expected inflation rate and the median expected inflation rate amongst professional forecasters based on the Philadelphia Fed's survey. The median rate amongst professional forecasters is consistent with MSC expectations during the first half of our study period, but expected inflation amongst MSC respondents increases quite substantially in the second half our sample, following the recent increase in gasoline prices. Apparently, MSC respondents perceive a connection between the current level of gasoline prices (or recent changes in gasoline prices) and the future rate of inflation in the economy.

Figure 2: Real retail gasoline prices and real MSC forecast

(a) Expected real price change

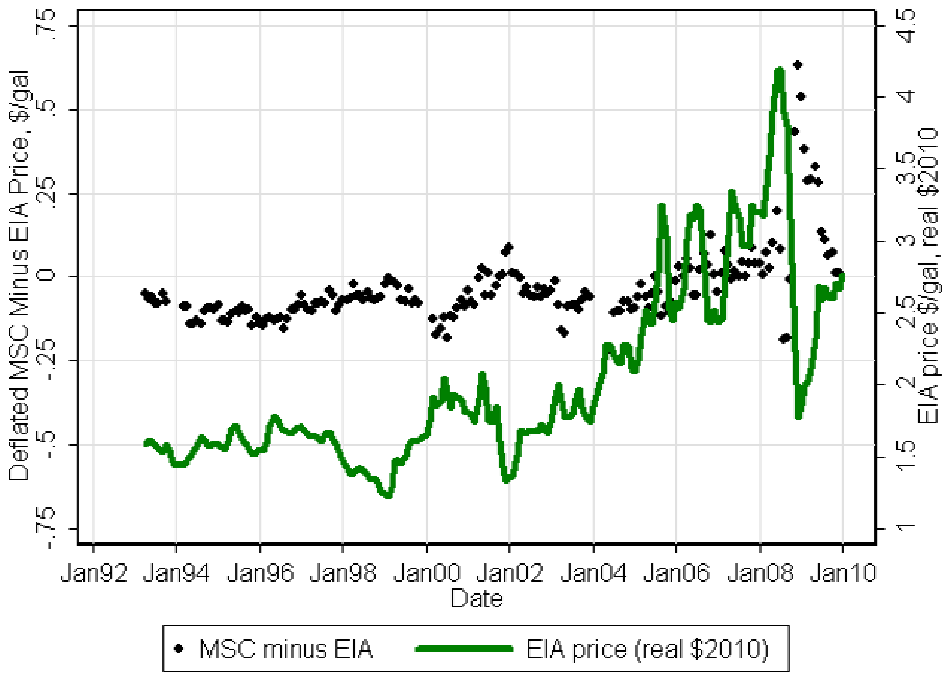

(b) Expected real price

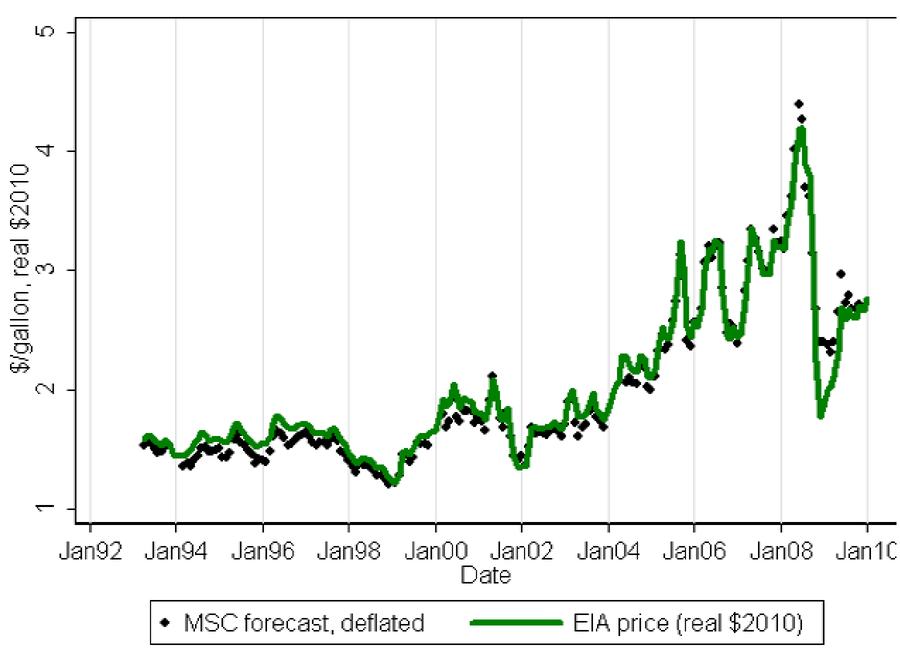

Figure 3: Retail gasoline prices (real), median MSC expected inflation, and median Philadelphia Fed expected inflation (survey of professional forecasters)

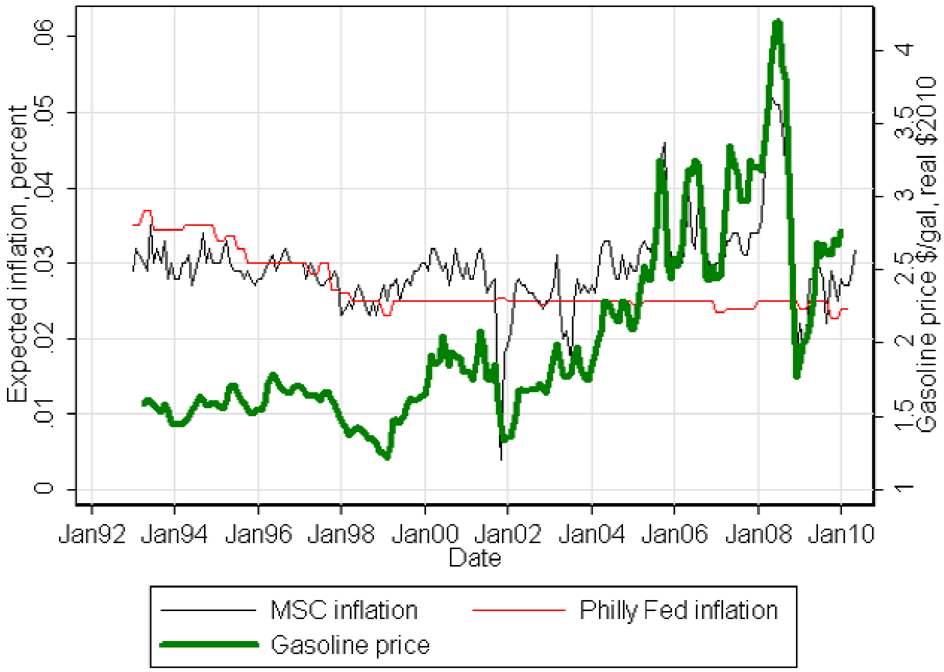

6 Regression analysis

We now test formally the null hypothesis that the "average" MSC respondent believes gasoline prices follow a random walk. Our econometric results will mirror the graphical analysis above, which suggests that the average consumer does forecast a real random walk. The naïve approach to implementing a regression-based test would consist of regressing the expected future price on the current price:

$$F_t^{60} = \beta_0 + \beta_1 P_t + \varepsilon_t \tag{5}$$

and then testing the joint null hypothesis that $\beta_0 = 0$ and $\beta_1 = 1$. This test could be conducted either in levels, as written, or in logs.

The current price of gasoline is highly persistent, however, as is the expected future price, regardless of whether we focus on mean or median expectations or how we choose to

Table 1: Regression tests for "immediate" random-walk beliefs

Panel A: Individual MSC price forecast deflated by individual inflation forecasts

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.879	0.8614	0.8847	0.8246	0.991	1.0144	1.0075	0.9897
(0.0957)	(0.1027)	(0.0358)	(0.0394)	(0.0242)	(0.0272)	(0.0426)	(0.0408)

Panel B: Aggregate MSC price forecast deflated by aggregate inflation forecasts

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.7019	0.7043	0.6374	0.6427	0.8529	0.8851	0.7561	0.773
(0.1115)	(0.1223)	(0.0688)	(0.0798)	(0.0552)	(0.0620)	(0.0432)	(0.0454)

Panel C: Individual MSC price forecast deflated by Philadelphia Fed inflation forecast

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.9021	0.8566	0.9265	0.8473	1.0446	1.0313	1.0395	0.9691
(0.1165)	(0.0384)	(0.0843)	(0.0801)	(0.0431)	(0.0496)	(0.0285)	(0.0316)

Panel D: Aggregate MSC price forecast and gasoline prices in nominal dollars

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
1.0028	0.8604	1.0373	0.8518	1.1817	1.0389	1.1807	0.9767
(0.1384)	(0.1196)	(0.1005)	(0.0816)	(0.0467)	(0.0500)	(0.0319)	(0.0322)

Note: All standard errors are Newey-West with 4 lags.

account for expected inflation. F_t^{60}, when measured by the median of the individual real price forecasts, has a first-order autoregressive coefficient of 0.979 in levels and 0.984 in logs. The current price of gasoline P_t has a first-order autoregressive coefficient is 0.975 in levels and 0.979 in logs. Given these coefficients, augmented Dickey-Fuller tests fail to reject unit roots on the real price and forecast series with p-values greater than 0.5 in all cases.[3] Thus, we estimate the model in first differences, using either price levels or logged prices.[4]

[3]This statement is true regardless of how inflation expectations are treated, whether logs or levels are used, whether a trend is allowed for, or whether zero or 12 lagged differences are included.

[4]We are able to reject a unit root in the difference between the MSC forecast and the current price when both are measured in logs, implying that the logged price series are cointegrated and that we could estimate the dynamic response of expected prices to current prices using an error correction model. By imposing that external prices eventually return to equaling current prices, however, this model assumes that consumers forecast are "no change" in the long run. We are not able to reject a unit root in the difference between the

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

$$\Delta F_t^{60} = \beta_0 + \beta_1 \Delta P_t + \varepsilon_t, \tag{6}$$

and using logged prices:

$$\Delta f_t^{60} = \beta_0 + \beta_1 \Delta p_t + \varepsilon_t, \tag{7}$$

where the lower-case letters indicate logged price variables. We estimate these models using several different versions of the "average" expected future price. That is, we estimate the model based on both mean and median expectations, as well as based on three ways of accounting for expected inflation, for a total of six different versions.

Table 1 presents our regression results. Panel A of the table presents regression results based on individual price forecasts deflated by individual inflation expectations, prior to aggregation. This is our preferred approach, since it accounts for correlation between individual price forecast and individual inflation forecasts. Results in levels based on the full 1993-2010 sample imply that when the current price of gasoline increases by \$1, the "average" forecast of the future price increases by about \$0.88. Results in logs based on the full sample imply that when the current price of gasoline increases by 1%, the "average" forecast of the future price increases by about 0.82%-0.86%. Results based on the mean and median responses are very similar in magnitude, but the results based on the mean response are more precisely estimated.Thus, while we cannot reject statistically a random walk belief based on the median responses, we soundly reject a random walk belief based on the mean responses.

It appears from these results that the "average" consumer does not believe that gasoline

two price series in levels, however, supporting the use of a model in differences. Persistence in the difference between the MSC real forecast and the current price of gasoline is evident in figure 2a.

prices will follow a random walk. Increases in the current price of gasoline do not translate to a one-for-one increase in the expected future price. These results appear to be driven entirely by the financial crisis of 2008, which led to a large deviation between the current and expected future price. When we exclude 2008 and subsequent years, as reported in the right-hand side of the table, the regression coefficients are all consistent with a random walk belief. Results in levels based on the limited 1993-2007 sample imply that when the current price of gasoline increases by $1, the "average" expected future price increases by $0.99. Results in logs based on the limited sample imply that when the current price increases by 1%, the average expected future price increases by 1.01%. These coefficients are more precisely estimated than in the full sample, and we are unable to reject a random walk belief in any case.

Panels B and C of table 1 replicate the results in panel A but use alternative methods for accounting for expected inflation. In panel B, we deflate the aggregate MSC nominal price forecasts by the aggregate MSC inflation forecast. In panel C, we deflate the aggregate MSC nominal price forecasts by the median forecast by professional forecasters based on quarterly survey by the Philadelphia Fed. The panel C results are similar to the reference case results from panel A: the "average" consumer predicts that future gasoline prices will not be substantially different from the current price. Panel B, however, indicates that future expectations about the real price of gasoline respond to changes in the current price in a way that is less than one-for-one, even after we exclude 2008. These results differ from those of panel A because panel A uses a dependent variable that is a mean (or median) of the change in a ratio, whereas panel B uses a change in the ratio of means (or medians). Whenever the two random variables in the ratio (here, the forecasts of the nominal gas price and of inflation) are not perfectly correlated, the ratio of means will be smaller than the mean of the ratio. This result is therefore a mechanical byproduct of the fact that the mean of the individual real price forecasts is consistent with random walk beliefs.

In panel D, we do not adjust for inflation at all, but rather regress the aggregate MSC nominal price forecast on the nominal price of gasoline. A \$1 increase in the current price of gasoline is estimated to be associated with a \$1 increase in the expected nominal future price when the entire sample is used. Given expected inflation of about 2%-3% per year, this finding implies that real prices only increase by about \$0.85-\$0.90, consistent with the estimates from panel A that used real prices and forecasts. The estimates using logged nominal prices and forecasts are very similar to those in panel A. This finding is consistent with the fact that 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, the linear model suggests that a \$1 increase in the current price leads to a \$1.18 dollar increase in the expected future price, consistent with 2%-3% expected annual inflation, or a random-walk belief in real terms. Even more tellingly, the model in logs suggests that a 1% increase in the current price leads to a 0.98%-1.04% increase in the expected future price, which is not statistically different from 1%. Assuming that expected future inflation does not change dramatically from one month to the next, these results are further evidence that the average consumer believes gasoline prices to follow a real random walk.

The results in table 1 estimate the immediate response of expectations to changes in the current price of gasoline. It is possible that consumer expectations respond with a delay. Thus, we estimate autoregressive distributed lag (ARDL) models that allow expectations to respond to changes in the current price with a delay. That is, we estimate dynamic models of the form:

$$\Delta F_t^{60} = \beta_0 + \sum_{k=0}^{q} \beta_k \Delta P_{t-k} + \sum_{k=1}^{q} \gamma_k \Delta F_{t-k}^{60} + \varepsilon_t. \tag{8}$$

These models allow changes in expectations to depend on changes in the current price of gasoline as well as on lagged changes in the price of gasoline and lagged changes in expectations from previous months.

Table 2: Regression tests for "long-run" random-walk beliefs

Panel A: Individual MSC price forecast deflated by individual inflation forecasts

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.8993	0.8557	0.8971	0.8077	0.9653	0.9497	0.9770	0.9141
(0.0382)	(0.0431)	(0.0419)	(0.0482)	(0.0337)	(0.0365)	(0.0437)	(0.0486)

Panel B: Aggregate MSC price forecast deflated by aggregate inflation forecasts

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.7308	0.7194	0.6247	0.6425	0.8582	0.8475	0.7915	0.7845
(0.0337)	(0.0417)	(0.0365)	(0.0459)	(0.0457)	(0.0521)	(0.0467)	(0.0560)

Panel C: Individual MSC price forecast deflated by Philadelphia Fed inflation forecast

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
0.9585	0.8669	0.9210	0.8298	0.9981	0.9384	1.0525	0.9367
(0.0326)	(0.0396)	(0.0381)	(0.0408)	(0.0455)	(0.0457)	(0.0463)	(0.0430)

Panel D: Aggregate MSC price forecast and gasoline prices in nominal dollars

Full sample: 1993-2010				Excluding crisis: 1993-2007			
Median Levels	Median Logs	Mean Levels	Mean Logs	Median Levels	Median Logs	Mean Levels	Mean Logs
1.0904	0.8756	1.0478	0.8412	1.1568	0.9586	1.2282	0.9594
(0.0368)	(0.0412)	(0.0480)	(0.0445)	(0.0499)	(0.0474)	(0.0534)	(0.0462)

Table 2 presents the results of these regressions. Depending on the particular price series we used, we found that it was necessary to include up to 7 periods of lagged prices and expectations to eliminate the serial correlation in the error term. Thus, all of the results in the table are based on ARDL model with 7 lags (i.e., $q = 7$ in the equation above). The table presents the "long-run" responses of expectations to a "permanent" increase in the price of gasoline. 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 almost immediately: 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 the table are all very similar to the coefficients in table 1, which measure the impacts on gasoline price forecasts that happen immediately in response to a change in the current price.

7 Conclusion

Do consumers exhibit a "reasonable" forecast of future energy prices? Our analysis suggests that they do. Using two decades of high quality survey data from the Michigan Survey of Consumers, we find that consumers, on average, report random walk beliefs regarding the future price of gasoline. This is a reasonable forecast in the sense that the predictive power of a random walk forecast is comparable to the predictive power of futures prices, and there is no obviously superior forecast. This finding suggests that if there is an "energy paradox," it likely does not stem from consumers having systematic bias in their beliefs about energy prices and that researchers are justified in using a random walk assumption when modeling consumer demand for durables.

References

Allcott, Hunt and Nathan Wozny, "Gasoline Prices and the Fuel Economy Discount Puzzle," September 2009. Manuscript: Princeton University.

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.

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," September 2009. Manuscript, University of California, Davis.

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

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.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, June 2009, *99* (3), 1053–1069.

Knittel, Christopher R. and Ryan Sandler, "Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins," October 2010. NBER Working Paper no. 16482.

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.