

DISCUSSION PAPER SERIES

DP15168

**OIL PRICES, GASOLINE PRICES AND
INFLATION EXPECTATIONS: A NEW
MODEL AND NEW FACTS**

Lutz Kilian and Xiaoqing Zhou

INTERNATIONAL MACROECONOMICS AND FINANCE



OIL PRICES, GASOLINE PRICES AND INFLATION EXPECTATIONS: A NEW MODEL AND NEW FACTS

Lutz Kilian and Xiaoqing Zhou

Discussion Paper DP15168
Published 15 August 2020
Submitted 11 August 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Macroeconomics and Finance

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Lutz Kilian and Xiaoqing Zhou

OIL PRICES, GASOLINE PRICES AND INFLATION EXPECTATIONS: A NEW MODEL AND NEW FACTS

Abstract

The conventional wisdom that inflation expectations respond to the level of the price of oil (or the price of gasoline) is based on testing the null hypothesis of a zero slope coefficient in a static single-equation regression model fit to aggregate data. Given that the regressor in this model is not stationary, the null distribution of the t-test statistic is nonstandard, invalidating the use of the normal approximation. Once the critical values are adjusted, these regressions provide no support for the conventional wisdom. Using a new structural vector regression model, however, we demonstrate that gasoline price shocks may indeed drive one-year household inflation expectations. The model shows that there have been several such episodes since 1990. In particular, the rise in household inflation expectations between 2009 and 2013 is almost entirely explained by a large increase in gasoline prices. However, on average, gasoline price shocks account for only 39% of the variation in household inflation expectations since 1981.

JEL Classification: E31, E52, Q43

Keywords: inflation, Expectations, anchor, missing disinflation, oil price, gasoline price, Household survey

Lutz Kilian - lkilian2019@gmail.com
Federal Reserve Bank of Dallas and CEPR

Xiaoqing Zhou - xqzhou3@gmail.com
Federal Reserve Bank of Dallas

Acknowledgements

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of Federal Reserve Bank of Dallas or the Federal Reserve System. We thank Robert De Jong, Marc Giannoni, Atsushi Inoue and Joon Park for helpful comments. The authors have no conflict of interest to disclose.

Oil Prices, Gasoline Prices and Inflation Expectations: A New Model and New Facts

Lutz Kilian
Federal Reserve Bank of Dallas
CEPR

Xiaoqing Zhou
Federal Reserve Bank of Dallas

June 25, 2020. This version: August 11, 2020

Abstract: The conventional wisdom that inflation expectations respond to the level of the price of oil (or the price of gasoline) is based on testing the null hypothesis of a zero slope coefficient in a static single-equation regression model fit to aggregate data. Given that the regressor in this model is not stationary, the null distribution of the t-test statistic is nonstandard, invalidating the use of the normal approximation. Once the critical values are adjusted, these regressions provide no support for the conventional wisdom. Using a new structural vector regression model, however, we demonstrate that gasoline price shocks may indeed drive one-year household inflation expectations. The model shows that there have been several such episodes since 1990. In particular, the rise in household inflation expectations between 2009 and 2013 is almost entirely explained by a large increase in gasoline prices. However, on average, gasoline price shocks account for only 39% of the variation in household inflation expectations since 1981.

JEL code: E31, E52, Q43

Key words: Inflation, expectations, anchor, missing disinflation, oil price, gasoline price, household survey.

Acknowledgments: The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of Federal Reserve Bank of Dallas or the Federal Reserve System. We thank Robert De Jong, Marc Giannoni, Atsushi Inoue and Joon Park for helpful comments. The authors have no conflict of interest to disclose.

Correspondence to: Lutz Kilian, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. E-mail: lkilian2019@gmail.com.

Xiaoqing Zhou, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. Email: xqzhou3@gmail.com.

1. Introduction

It is well known that survey data of household inflation expectations may differ systematically from professional inflation forecasts. One of the explanations considered in the literature has been that households' expectations respond excessively to fluctuations in the price of crude oil. For example, Coibion and Gorodnichenko (2015) make the case that one-year mean household inflation forecasts, as measured by the Michigan Survey of Consumers (MSC), have tracked the price of oil closely with a contemporaneous correlation of 74% between January 2000 and March 2013. Almost all of the short-run volatility of household inflation expectations, according to their analysis, appears explained by changes in the level of the price of oil. They conclude that household inflation expectations at the one-year horizon are not fully anchored. This view has become part of the mainstream in recent years, has been elaborated on in numerous academic and policy studies, and has spurred a revival of work on the expectations-augmented Phillips curve using inflation expectations from household surveys.¹

In this paper, we reexamine this conventional wisdom. We first show that the stylized fact reported in Coibion and Gorodnichenko (2015) is highly sensitive to the estimation period. We then demonstrate that the type of single-equation regression that has been used to establish that increases in the price of oil are responsible for increases in household inflation expectations is problematic. Not only are there concerns about the choice of the regressor, but inference in regressions of $I(0)$ on $I(1)$ variables is nonstandard. We find that, when using the correct null distribution, none of the regression-based tests estimated on the full sample supports the conventional wisdom.

Next, we explain how to reconcile this result with the seemingly large correlation

¹ Examples include Anderson, Kellogg and Sallee (2013), Elliott, Jackson, Raczko and Roberts-Sklar (2015), Sussman and Zohar (2015), Wong (2015), Binder (2018), Hasenzagl, Pellegrino, Reichlin and Ricco (2018), Conflitti and Cristadoro (2018), and Coibion, Gorodnichenko, Kumar, and Pedemonte (2020).

between inflation expectations and the nominal price of oil observed during the period of 2000-13. We show that this high correlation vanishes when including earlier data in the estimation. We demonstrate by simulation that the temporal instability in the correlation estimates can be traced to the fact that the level of the nominal price of oil does not have a stable mean.

This evidence suggests that neither static regressions nor reduced-form correlations are the appropriate tool for understanding this empirical relationship. A more promising approach is to analyze this relationship using a structural model with multiple shocks. We propose a structural vector autoregressive model that sheds light on the interaction of actual inflation rates, household inflation expectations and the real price of gasoline. The model focuses on the price of gasoline rather than the price of oil, consistent with Coibion and Gorodnichenko's (2015) argument that households pay particular attention to gasoline prices when forming expectations of other prices. It is designed to quantify the cumulative effects of nominal gasoline price shocks on inflation expectations without strong restrictions on the dynamics of the relationship between inflation expectations and the price of oil. Inference in this model does not depend on whether the log real price of gasoline is $I(0)$ or $I(1)$.

Our identifying assumptions are motivated by the microeconomic literature on the determination of household inflation expectations. The model is estimated on monthly data back to mid-1981 to be consistent with Coibion and Gorodnichenko's original regression analysis. Virtually identical estimates are obtained when estimating the model on data starting in 1990, suggesting that the structural relationships of interest are stable.

Unlike the single-equation regression models discussed earlier, this structural model provides clear evidence of the transmission of nominal gasoline price shocks to household inflation expectations. Our estimates show that a shock that raises the nominal price of gasoline by 1% increases household inflation expectations by 0.05 percentage points on

impact, which is excessive relative to the gasoline share in consumer expenditures of near 3% on average. The response of inflation expectations declines rapidly over time, however, and is indistinguishable from zero after three months.

Next we construct a counterfactual for the evolution of one-year household inflation expectations in the absence of gasoline price shocks. The model implies that, since 1990, there have been six major episodes in which one-year inflation expectations have responded substantially to gasoline price shocks. In 1990 and in the mid-2000s, for example, household inflation expectations cumulatively rose by as much as one percentage point at annualized rates in response to gasoline price shocks, whereas in the late 1990s and from late 2014 to late 2017 nominal gasoline price shocks caused household inflation expectations to drop by as much as one percentage point. The sharp drop in gasoline prices in early 2020, at the end of our sample, appears to be the beginning of another episode like this. By April 2020, the drop in nominal gasoline prices already accounted for an almost one percentage point gap between actual and counterfactual household inflation expectations.

We also examine the view expressed in Coibion and Gorodnichenko (2015) that the missing disinflation in the U.S. economy between 2009 and 2013, can be explained almost entirely by the response of inflation expectations to the recovery of the price of oil. Our structural model shows that indeed the rise in household inflation expectations over this period was largely caused by an increase in gasoline prices. However, a variance decomposition based on the estimated model reveals that, on average, gasoline price shocks account for only 39% of the variation in household inflation expectations, with an additional 54% explained by idiosyncratic household expectations shocks and 7% by shocks to the core CPI.

We also link gasoline price shocks to the deviation between short-run and long-run inflation expectations which provides a measure of how anchored short-run inflation

expectations are. We conclude that monetary policymakers must be aware that gasoline price shocks may mask the underlying trends in one-year consumer inflation expectations. This point seems particularly relevant in the context of the recent sharp decline in gasoline prices that was triggered by Covid-19 epidemic. Finally, we discuss why consumers appear to overreact to gasoline price shocks and argue that this relationship need not persist going forward.

Our work relates to the literature on whether inflation expectations have been successfully anchored by the increased credibility of monetary policy in recent decades (see, e.g., Bernanke 2010; Jorgensen and Lansing 2019). It also relates to a growing literature on how household inflation expectations are determined (see, e.g., Madeira and Zafar 2015; Binder 2018; Angelico and Di Giacomo 2019; Binder and Makridis 2020). In addition, our analysis is related to earlier work on how oil and gasoline price shocks are transmitted to inflation (see, e.g., Kilian 2009; Clark and Terry 2010; Kilian and Lewis 2011; Wong 2015; Conflitti and Luciani 2019). Finally, our analysis sheds light on one of the premises of the recent literature on the expectations-augmented Phillips curve (see, e.g., Coibion and Gorodnichenko 2015; Coibion, Gorodnichenko, and Kamdar 2018; Hazennagl et al. 2018).

The remainder of the paper is organized as follows. Section 2 reviews the regression evidence reported in Coibion and Gorodnichenko (2015) and examines its robustness to the estimation period. In section 3, we draw attention to a number of econometric issues with this type of regression analysis and show that after suitable corrections, there is no statistically significant evidence that oil prices or gasoline prices are correlated with inflation expectations. Section 4 explains why the correlation between inflation expectations and the price of oil (or the price of gasoline) tends to be unstable over time. Section 5 presents empirical estimates of the joint determination of headline inflation rates, inflation expectations, and the real price of gasoline based on an alternative structural vector

autoregressive model that is stable over time. We quantify the extent to which inflation expectations have become unanchored in response to gasoline price shocks, and we put our estimates into perspective. The concluding remarks are in section 6.

2. How robust is the statistical relationship between the price of oil and inflation expectations?

Coibion and Gorodnichenko (2015) stress that historically the price of oil has been highly correlated with one-year household inflation expectations in the aggregate. The estimate reported in their paper, however, is based on regressing the difference between the one-year mean inflation expectation in the Michigan Survey of Consumers and the corresponding inflation forecast in the Survey of Professional forecasters, $\pi_t^{\text{exp}} - \pi_t^{\text{exp.SPF}}$, during the period of 1981Q3-2013Q1 on the level of the oil price, O_t .² They do not report estimates of regressions of π_t^{exp} on O_t . Regarding the latter relationship, they only report a 74% correlation for the period of 2000Q1-2013Q1 and a plot of these two time series from 1990Q1 to 2013Q1.

Table 1 provides a comprehensive overview of the key facts. Results that correspond to findings discussed in Coibion and Gorodnichenko (2015) are shown in bold. The first column is based on the estimation period used in Coibion and Gorodnichenko's (2015) Table 4. The second column is based on the time frame utilized in their time series plot. The third column focuses on the period for which Coibion and Gorodnichenko report the correlation between the oil price and inflation expectations. Finally, the last two columns correspond to the first two columns, except that the estimation period has been extended to early 2020.

The first panel of Table 1 shows that the correlation between $\pi_t^{\text{exp}} - \pi_t^{\text{exp.SPF}}$ and O_t is

² The oil price is measured by the price of West Texas Intermediate (WTI) crude oil, expressed in dollars per barrel.

robust across estimation periods. This evidence, however, does not address the question of how correlated π_t^{exp} and O_t are. The second panel shows that we can replicate the 74% correlation between π_t^{exp} and O_t highlighted by Coibion and Gorodnichenko for 2000Q1-2013Q1, but that correlation becomes much weaker when estimation starts in 1990Q1 and all but vanishes when estimation starts in 1981Q3. Moreover, the t-statistic drops to 1 and the R^2 to 1.3% using the same sample period as in their Table 4. The last panel of Table 1 shows that similar results are obtained when fitting the same regression on monthly data rather than quarterly data.

The near-zero and statistically insignificant coefficient in the regression of π_t^{exp} on O_t raises the question of what is driving the statistically significant slope coefficient of 0.024 in the regression of $\pi_t^{\text{exp}} - \pi_t^{\text{exp,SPF}}$ on O_t in Table 4 of Coibion and Gorodnichenko (2015). We find that this estimate is driven by the negative slope coefficient in the regression of $\pi_t^{\text{exp,SPF}}$ on O_t , which is -0.020 with a t-statistic of -4.2. This evidence is difficult to reconcile with the argument that households adjust their inflation forecasts more strongly in response to oil price changes than professional forecasters because of the salience of gasoline prices to consumers.

Table 1 illustrates that the stylized fact of a high correlation between the price of oil and inflation expectations is not robust. In no case is there evidence of a tight link between the price of oil and inflation expectations using data before 1990. Only the regressions of π_t^{exp} on O_t estimated on data since 1990 would have been considered statistically significant by Coibion and Gorodnichenko. In the next section, we examine this type of regression in more detail. To allow for the possibility of a structural break in this relationship before 1990, we focus on the monthly data starting in January 1990.

3. What do we learn from single-equation static regressions?

Figure 1 is an updated version of Figure 7 in Coibion and Gorodnichenko (2015), except that we focus on monthly data. At first glance, the relationship between the price of oil and consumer inflation expectations seems present in the updated data as well. As we discussed earlier, the static regression model

$$\pi_t^{\text{exp}} - \pi_t^{\text{exp},SPF} = \alpha + \beta O_t + \varepsilon_t$$

does not directly address this relationship, but the alternative regression specification

$$\pi_t^{\text{exp}} = \alpha + \beta O_t + \varepsilon_t. \quad (1)$$

does. Equation (1) is also in line with subsequent studies that have abstracted from the Survey of Professional Forecasters and have directly focused on the relationship between household inflation expectations and the level of the oil price (or the gasoline price) (see, e.g., Elliott et al. 2015; Sussman and Zuhar 2015; Wong 2015; Conflitti and Luciani 2019; Coibion et al. 2020). Likewise, policy discussions have focused on the direct link from the oil price to household inflation expectations (see, e.g., Hazennagl et al. 2018).

A rejection of $H_0 : \beta = 0$ may be interpreted as a rejection of the hypothesis that these two series are mutually uncorrelated.³ Since the error term in the regression, under $H_0 : \beta = 0$ inherits the persistence of the dependent variable, inference is based on a one-sided t-test constructed from Newey-West standard errors. The first column of Table 2 shows estimates of equation (1) based on the data in Figure 1. The correlation between consumer inflation expectation and the level of the oil price is only 29%. The R^2 of the regression is only 8.6%. It is less clear what the value of the t-statistic is for testing $H_0 : \beta = 0$. The problem is that the t-statistic for $\hat{\beta}$ in Table 2 is sensitive to the choice of the truncation lag used in computing

³ Under the additional (economically questionable) assumption that the price of oil is strictly exogenous with respect to inflation expectations and not correlated with exogenous variation in other variables, β measures the causal effect of the price of oil on inflation expectations.

the variance of $\hat{\beta}$. Depending on the choice of this truncation lag, it ranges from 1.59 to 2.04. Since the dependent variable appears stationary, but is fairly persistent, as shown in Figure 1, the analysis calls for a large truncation lag. Andrews' (1991) data-based method for selecting the truncation lag suggests a t-statistic of 1.83, corresponding to a truncation lag of about 24.

3.1. Should the oil price be expressed in levels or in log-levels?

It is readily apparent that $\Delta o_t \approx (O_t - O_{t-1}) / O_{t-1}$ is I(0) and hence o_t is I(1). This means, however, that O_t cannot be I(1), calling into question regression (1). Even after differencing, O_t is not stationary because its variance changes over time in a manner that would not be captured by GARCH models of conditional heteroskedasticity, as illustrated in Figure 2. The log-differenced series has a much more stable variance than the differences of the price of oil. Given the evidence in Figure 2, a more conventional specification would be

$$\pi_t^{\text{exp}} = \alpha + \beta o_t + \varepsilon_t, \quad (1')$$

where the regressor is logged. It may seem that such a subtle specification change should not make much of a difference. However, the third column in Table 2 shows that the t-statistics are substantially smaller based on specification (1'). The R^2 of the regression drops from 9% to 4%. Using Andrews' (1991) estimator of the truncation lag, the t-statistic drops from 1.83 to 1.25. Based on the N(0,1) approximation employed by Coibion and Gorodnichenko (2015), this t-statistic would no longer allow us to reject the null at conventional significance levels.

3.2. The oil price is not a good proxy for the gasoline price

There is a deeper caveat about these results, however. The intuition underlying Coibion and Gorodnichenko's analysis is that the price of gasoline is particularly salient because consumers are confronted with this price daily, as they pass by gas stations. Their argument is that households are likely to pay particular attention to prices they see more often when

formulating their expectations of future inflation. This argument does not apply to the price of crude oil, however. Most consumers would be at a loss when asked about the current price of crude oil.

One possible rationale for including the price of crude oil in equations (1) and (1') is that the cost share of crude oil in producing gasoline has been remarkably constant at about 50% in recent years, implying that, for every percent change in the price of crude oil, the price of gasoline will change by half of that percentage. That relationship need not hold for older data, however. Moreover, the price of gasoline also depends on other cost components, on gasoline taxes, and on shifts in the demand for other refined products.⁴

It therefore makes more sense to evaluate the link between gas prices and consumer inflation expectations directly. This is also the approach taken by several more recent studies including Binder (2018) and Coibion, Gorodnichenko, Kumar and Pedemonte (2020). We employ the city average of the price for unleaded motor gasoline, as reported in the U.S. EIA's *Monthly Energy Review*. The results are robust to other measures of the gasoline price.

The second column in Table 2 shows that the evidence for model (1) becomes weaker, once we focus directly on the price of gasoline. The R^2 of the regression drops from 9% to 5% and the t-statistic computed using Andrews' (1991) estimator of the truncation lag drops from 1.83 to 1.26, at which point $\hat{\beta}$ is no longer statistically significant at conventional significance levels based on the $N(0,1)$ approximation. This result calls into question the original specification of model (1) based on the price of oil. Under the maintained assumption that the price of oil is a good proxy for the price of gasoline, the results for the price of oil and the price of gasoline should effectively be the same. There is no possible explanation for

⁴ In addition, when working with longer samples, it has to be kept in mind that the WTI price of crude oil was regulated until the early 1980s and not representative for the cost refiners paid for acquisitions of crude oil (see Alquist et al. 2013). For example, in the late 1970s, the WTI price was effectively constant for extended periods due to regulation, whereas the price of gasoline was not.

why the regression based on the price of oil should fit the data better. Working with the log-level of the price of oil, as in equation (1'), also substantially reduces the t-statistic and the R^2 . The last column in Table 2 evaluates model (1') using the log of the gasoline price rather than the log of the oil price. Combining the two corrections, reduces the R^2 to 2% and the t-statistic to 0.81. To summarize, the t-statistic for $\hat{\beta}$ is below conventional critical values whether working with the level or the log-level of the gasoline price.

3.3. The critical values are not standard

Regardless of whether we specify the gasoline price in levels or log-levels, however, the standard Gaussian critical values used in the existing literature are not appropriate in this context. Inflation expectations are plausibly I(0). When regressing an I(0) variable on an I(1) variable (or a nonlinear transformation of an I(1) variable), the distribution of the estimator is no longer Gaussian and the distribution of the t-statistic of $\hat{\beta}$ may be far from N(0,1). This is a particular concern when the dependent variable is positively serially correlated, as in the case of household inflation expectations (see Stewart 2011).⁵

To illustrate this important point, we consider a bivariate data generating process (DGP) for π_t^{exp} and the price of oil that embodies the restriction that $\beta = 0$. We will show that, under this DGP, the distribution of $t_{\hat{\beta}}$ is not well approximated by the N(0,1)

distribution. For simplicity, we postulate that π_t^{exp} follows an AR(1) process,

$\pi_t^{\text{exp}} = \alpha_0 + \alpha_1 \pi_{t-1}^{\text{exp}} + \varepsilon_t^{\text{exp}}$, where $\varepsilon_t^{\text{exp}} \sim N(0, \sigma_{\text{exp}}^2)$. The parameters of this process may be

recovered from the data. The estimated slope parameters is 0.84. The log-level of the price of oil, o_t , is independently generated by cumulating realizations of $\Delta o_t = \mu_o + \sigma_o \varepsilon_t^o$, where ε_t^o is

⁵ Under the alternative assumption that inflation expectations are better approximated by an I(1) process, this regression would be a classical spurious regression, $\hat{\beta}$ and R^2 would converge to a random variable and standard inference would be equally invalid.

a Student- t_4 innovation (standardized to have mean zero and variance 1) and μ_o and σ_o are calibrated to match the mean and the standard deviation of Δo_t in the data. The assumption that o_t approximately follows a random walk is consistent with Figure 2 (see also Alquist, Kilian and Vigfusson 2013). The importance of modeling ε_t^o as a fat-tailed distribution is also illustrated in Figure 2. The same type of DGP may also be used to generate realizations of $O_t = \exp(o_t)$. Realizations of the price of gasoline may be generated analogously, by replacing the price of oil in the DGP by the price of gasoline and recalibrating the parameters.

This DGP allows us to simulate the finite-sample distribution of the t-statistic under $H_0 : \beta = 0$ in repeated sampling for the sample size of interest. We generate synthetic data from the estimated process of the same length as the actual data. For each random draw, we regress the simulated time series of π_t^{exp} on an intercept and the log-level (or the level) of the simulated oil (or gasoline) price, allowing us to build up the empirical distribution of the t-statistic.

While there is no reason for this specific DGP to be the true process necessarily, the point of this exercise is to illustrate that the distribution of the t-statistic for $\hat{\beta}$ is far from the $N(0,1)$ distribution under plausible assumptions. Figure 3 shows that the null distribution under the log-level specification is centered on zero but much more fat-tailed. Very similar results hold for the level specification. Table 3 shows that the 10% and 5% critical values for testing $H_0 : \beta = 0$ against $H_1 : \beta > 0$ under these null distributions are considerably larger than those implied by the $N(0,1)$ distribution. The standard Gaussian critical values used in the existing literature are systematically too small, causing the test to reject too often.

Table 4 reports the finite-sample p-values for the t-statistics in Table 2 associated with these null distributions. Using Andrews' (1991) estimator of the truncation lag, we are unable to reject the null for any of the specifications in Table 4. Neither model (1) nor model (1')

provides evidence against the null that the oil price affects inflation expectations. Under the log specification, the p-value increases. The p-values are higher for either specification when focusing on the gasoline price, which is the more relevant variable for this exercise. We conclude that caution is called for in interpreting the t-statistics reported in the existing literature. There is no systematic evidence that shifts in the level of the oil price or the gasoline price un-anchor consumers' one-year inflation expectations. Obviously, we could have done the same exercise with data starting in 1981, but the test results would have been even weaker.

4. The correlation evidence

One may be tempted to argue that the high correlation between the level of the nominal price of oil and survey inflation expectations found for 2001Q1-2013Q3 speaks for itself, even without formal inference. As Table 1 shows, however, this correlation is unstable over time. This is true, even when extending the estimation period back to 1960, when quarterly household inflation expectations first became available. For example, using quarterly data from 1960Q1 to 2013Q1, the correlation in question is 38%, but for 1960Q1-2020Q1 it drops to -5%.

This is no accident. Focusing on correlations can be deceiving, because the correlation of inflation expectations with variables that are not stationary (such as the price of oil) is sensitive to the estimation period. The reason is that computing the sample correlation requires the data to be demeaned. As the estimation period becomes short, the implied mean may be quite different from the mean for the full sample, which can render the correlation estimator erratic. A simple simulation study illustrates this point. Consider the sample period 1990.1-2013.3. Suppose that $O_t = \exp(o_t)$, $t = 1, \dots, 279$, where $\Delta o_t = \mu_o + \sigma_o \varepsilon_t^o$, μ_o and σ_o are calibrated to the data and $\varepsilon_t^o \sim NID(0,1)$, is generated independently of the observed π_t^{exp} .

Then, in repeated trials, the average absolute difference in the correlation of O_t with π_t^{exp} between the first and the second half of the sample is 56 percentage points.

Our analysis so far suggests that neither static regressions nor reduced-form correlations are the appropriate tool for understanding this empirical relationship. In the next section, we propose an alternative structural econometric model that reveals a systematic link between nominal gasoline price shocks and household inflation expectations using all the data available since July 1981.

5. A structural model

In this section, we introduce a structural vector autoregressive (VAR) model that further disentangles the relationship between shocks to the nominal price of gasoline and inflation expectations, drawing on economic insights from the microeconomic literature on inflation expectations. The model is estimated on monthly data starting in 1981.7 and ending in 2020.4, corresponding to the estimation period underlying Table 4 in Coibion and Gorodnichenko (2015), but expanded to include more recent data.

Our VAR approach is less restrictive than the static regression model discussed earlier in several dimensions. First, we work with the price of gasoline rather than the price of crude oil, consistent with the original motivation in Coibion and Gorodnichenko (2015). Second, we replace the implicit assumption that the nominal price of gasoline price is strictly exogenous by the weaker assumption that this price is predetermined with respect to U.S. inflation expectations. Third, we explicitly model the dynamics of the relationship between the gasoline price and inflation expectations, allowing us to quantify the cumulative effect of nominal gasoline price shocks on household inflation expectations. Fourth, our VAR model allows the log real price of gasoline to be I(0) or I(1).

Let $y_t = [rpgas_t, \pi_t, \pi_t^{\text{exp}}]'$, where $rpgas_t$ denotes the log-level of the real gasoline

price, π_t is the headline CPI inflation rate, and π_t^{exp} is the Michigan Survey of Consumers measure of household's one-year inflation expectations (see Figure 4). The inflation rates are not annualized. Consider the structural VAR process $B_0 y_t = B_1 y_{t-1} + \dots + B_p y_{t-p} + w_t$, where w_t denotes the mutually uncorrelated i.i.d. structural shocks and B_i , $i = 0, \dots, p$, represent 3×3 coefficient matrices. The intercept has been dropped for expository purposes. The reduced-form VAR representation is $y_t = A_1 y_{t-1} + \dots + A_p y_{t-p} + u_t$, where $A_i = B_0^{-1} B_i$, $i = 1, \dots, p$. We set the lag order to a conservative upper bound of 12 lags (see Kilian and Lütkepohl 2017).

The model explains variation in the data in terms of the structural shocks $w_t = [w_t^{\text{gas price}}, w_t^{\pi \text{ ex gas}}, w_t^{\text{exp}}]'$. The gasoline price shock captures innovations in the nominal price of gasoline that are presumably salient to households. In contrast, the shock to the "core" CPI excludes the gasoline price, but covers all other prices. Finally, we allow for idiosyncratic shocks to household's inflation expectations that are orthogonal to the first two structural shocks. This shock is designed to capture changes in households' perceptions of future inflation not captured by current prices. The importance of such idiosyncratic shocks has been emphasized in Madeira and Zafar's (2015) study of household-level inflation expectations data.

The identification of the structural model exploits a combination of sign and zero restrictions on the structural impact multiplier matrix B_0^{-1} , where $u_t = B_0^{-1} w_t$. A positive nominal gasoline price shock is assumed to raise the real price of gasoline on impact because the CPI responds more slowly than the nominal price of gasoline. It also is assumed to raise household inflation expectations, given the household-level evidence in Binder (2018). A positive shock to the core CPI (defined as consumer prices excluding gasoline price) raises consumer price inflation and inflation expectations, consistent with the evidence in Binder

(2018). It lowers the real price of gasoline on impact, given that the nominal gasoline price does not respond within the month to inflation shocks (see Kilian and Vega 2011). Finally, a positive shock to household inflation expectations not associated with the other structural shocks leaves the real price of gasoline and headline inflation unaffected on impact. The rationale for this restriction is that idiosyncratic expectations shocks that move actual consumer prices are already captured by the gasoline and core CPI shocks. Jointly these restrictions imply that

$$\begin{pmatrix} u_t^{rpgas} \\ u_t^\pi \\ u_t^{\pi^{exp}} \end{pmatrix} = \begin{bmatrix} + & - & 0 \\ + & + & 0 \\ + & + & + \end{bmatrix} \begin{pmatrix} w_t^{\text{nominal gasoline price}} \\ w_t^{\text{core CPI}} \\ w_t^{\text{idiosyncratic inflation expectation}} \end{pmatrix}.^6$$

The model is estimated by Bayesian methods using a uniform-Gaussian-inverse Wishart prior, as described in Arias, Rubio-Ramirez and Waggoner (2018). The reduced-form prior is a conventional Minnesota prior with zero mean for the slope parameters. In the appendix, we show that this prior is largely uninformative for the vector of structural impulse responses and is not driving our empirical results. Having simulated the posterior distribution of the structural impulse responses, we evaluate the joint impulse response distribution under absolute loss, as discussed in Inoue and Kilian (2020a).⁷

5.1. How do inflation expectations respond to nominal gasoline price shocks?

⁶ The log real price of gasoline in this model, in principle, could be replaced by the log of the real price of crude oil without affecting the identification. In that case, the first structural shock would be a nominal oil price shock. The reason we focus on the price of gasoline is that not all gasoline price changes reflect oil price changes, as discussed earlier.

⁷ The identification of our model differs in several important dimensions from the VAR model employed in Wong (2015). First, unlike Wong, we allow inflation expectations to respond instantaneously to all nominal price shocks. This point is important because households can directly observe gasoline prices as well as many other prices such as food prices, allowing them to form inflation expectations without having to wait for CPI data releases. Second, Kilian and Vega (2011) only established that nominal gasoline prices are predetermined with respect to inflation news, allowing for the possibility that an unexpected increase in the inflation rate would lower the real price of gasoline. This argues against one of the exclusion restriction imposed in Wong (2015), which in our analysis is replaced by a sign restriction. Third, our model exploits additional information about the signs of three impact responses of inflation and inflation expectations not utilized by Wong (2015). Finally, we do not allow idiosyncratic expectations shocks to move headline inflation on impact because such shocks are already embodied in the nominal price shocks.

Figure 5 shows the set of impulse responses obtained by minimizing in expectation the joint loss function. It also shows the implied joint 68% credible sets.⁸ A positive nominal gasoline price shock causes a persistent appreciation of the real price of gasoline, as expected, and a sharp and precisely estimated increase in headline inflation that quickly dies out after three months. It also causes a small, but persistent increase in inflation expectations that is precisely estimated for the first three months. There is no evidence that nominal gasoline price shocks permanently affect one-year inflation expectations in that the longer-run response of inflation expectations is indistinguishable from zero. A positive shock to the core CPI has negligible effects on inflation expectations and the real price of gasoline, but raises headline inflation for about two months. Finally, a positive idiosyncratic shock to the survey inflation expectation raises inflation expectations persistently. It temporarily raises headline inflation for about one quarter, but that effect is not precisely estimated. The effect on the real price of gasoline is negligible.

The key result in Figure 5 is that a nominal gasoline price shock that raises the nominal price of gas by one percent on impact causes inflation expectations to increase by 0.054 percentage points on impact.⁹ After accounting for estimation uncertainty, this estimate may be as low as 0.042 and as high as 0.084, which is higher than would be expected based on consumers' average expenditure share of gasoline (which is 3% over this sample, when defined narrowly as gasoline and other motor fuel, compared with 5% using a broader definition of energy). This excess sensitivity is consistent with households overreacting to gasoline price shocks due to their salience, and consistent with the view that traditional

⁸ This approach may be viewed as a generalization of the traditional approach of reporting quantiles of the marginal distribution of each impulse response coefficient. Its advantage is that it preserves the dynamics of the impulse response functions as well as their comovement. It also accounts for the joint uncertainty in the response estimates.

⁹ This estimate is obtained as $\hat{\theta}_{31,0} / (\hat{\theta}_{11,0} + \hat{\theta}_{21,0})$, where $\theta_{jk,h}$ denotes the response of variable j to structural shock k at horizon h .

historical narratives about the role of oil price shocks in the 1970s and 1980s have become entrenched in households' psyche (see, e.g., Edelstein and Kilian 2009; Binder and Makridis 2020).

One may object that perhaps the relationship between gasoline price shocks and inflation expectations has evolved over time and has become stronger in recent years. If there had been a structural change, one would have expected this change to have happened in the early 1990s, which is when the correlation evidence is getting stronger (see also Wong 2015). We address this concern by re-estimating the structural VAR model on data for 1990.1-2020.4, corresponding to the sample period already used in section 3. The impulse response estimates are virtually identical. The impact response of inflation expectations to a gasoline price shock that raises the nominal price of gasoline by 1% is 0.07 percentage points, suggesting that there has been no important structural change in this relationship since the 1980s. Likewise, the remaining results reported below are robust to this change. We therefore proceed with the baseline VAR model for the remainder of the analysis, which allows us to study inflation expectations over a longer period.

5.2. How much of the evolution of inflation expectations must be attributed to nominal gasoline price shocks?

The evidence in Figure 5 speaks directly to the response of inflation expectations to a one-time nominal gasoline price shock, but it does not tell us how much of the evolution of household inflation expectations must be attributed to nominal gasoline price shocks. A better measure of the overall importance of nominal gasoline price shocks for inflation expectations is the historical decomposition in Figure 6, which provides clear evidence of inflation expectations increasing from the late 2002 to mid-2008 in response to the cumulative effects of nominal gasoline price shocks, followed by a sharp drop during the financial crisis and a gradual recovery until early 2011. There is also evidence of a temporary increase in inflation

expectations associated with the 1990 oil price spike and of sustained declines associated with the gasoline price drops in 2014-16 and in early 2020. Not all variation comes from gasoline price shocks. The upper panel of Figure 6 suggests that inflation expectations at times also increase for idiosyncratic reasons related to how households form expectations. For example, there was a sharp increase in idiosyncratic inflation expectations in mid-2008. Shocks to the core CPI, in contrast, tend to have only minor cumulative effects.

Figure 7 shows actual inflation expectations data, as in Figure 1, and the counterfactual series constructed by subtracting the cumulative effect of nominal gasoline price shocks shown in Figure 6 from the actual data. Figure 7 shows that one-year household inflation expectations may rise or fall for periods lasting more than one year at a time, when confronted with large and persistent gasoline price movements. In each case, however, these discrepancies ultimately vanish. One example is the period from late 2014 to late 2017, when falling gasoline prices masked a rise in inflation expectations. A similar event is about to unfold at the end of the sample in response to the Covid-19 epidemic. Another example is the late 1990s following the Asian crisis. Periods when inflation expectations were elevated by the cumulative effect of nominal gasoline price shocks, in contrast, include 1990, when Iraq invaded Kuwait, and the period of the Great Surge in the price of oil from 2004 to mid-2008.

What does the historical decomposition tell us about the missing disinflation from January 2009 to March 2013 discussed in Coibion and Gorodnichenko (2015)? Recall that, over this period, household survey inflation expectations increased by 1.5 percentage points (on an annualized basis). Figure 7 shows that in the absence of nominal gasoline price shocks, household inflation expectations would have cumulatively increased by only 0.1 percentage points over this period. Thus, the observed increase in inflation expectations is largely explained by gasoline price shocks. However, a variance decomposition based on the estimated model shows that, on average over the entire estimation period, gasoline price

shocks account for only 39% of the variation in household inflation expectations. The most of important determinant on average are idiosyncratic household expectations shocks, which account for 54% of the variation. Shocks to the core CPI explain 7%.

One metric for how anchored short-run inflation expectations are is the difference between the short-run and long-run expectation. Figure 8 illustrates that the gasoline price component of one-year household inflation expectations is clearly correlated with measures of the deviation between short-run and long-run household inflation expectations. We focus on the period since 1990.4, when the 5-year inflation expectation was first released by the Michigan Survey of Consumers. Notwithstanding the common component between these series, the overall correlation is only 52%, confirming that nominal gasoline price shocks are only one determinant of the spread between short- and long-run inflation expectations.

5.3. Why do consumers appear to overreact to nominal gasoline price shocks?

In the 1970s and 1980s, gasoline price increases were driven by oil price increases that reflected high global and domestic demand for goods and services (see, e.g., Barsky and Kilian 2002). In the 2000s, there was a similar demand boom, but domestic U.S. consumer price inflation remained subdued, possibly because at this point much of U.S. industrial production had been outsourced to emerging economies. The unanchoring of short-run inflation expectations observed in Figure 7 after 2002 is broadly consistent with households having associated higher gasoline prices with more general consumer price pressures, much like in the 1970s and early 1980s. Likewise, the disproportionate decline in household inflation expectations after 2014 and in early 2020 is consistent with this interpretation.

This behavioral pattern is not written in stone. For example, Madeira and Zafar (2015) and Binder and Makridis (2020) document that this pattern is more pronounced among households who personally experienced the 1970s than among younger households. As older households drop out of the survey and are replaced by younger households, one would

therefore expect less of a response of inflation expectations going forward. In addition, one should not underestimate the role of economists in shaping public perceptions either directly or through discussions in the financial press. For many years, the conventional wisdom has been that oil price (and gasoline price) increases are necessarily inflationary. The narrative of how the high inflation of the 1970s came about has changed only in recent years, starting with Barsky and Kilian (2002). As the textbook explanations of the 1970s are gradually replaced by more modern accounts, one would expect public perceptions of the effects of gasoline price shocks to gradually adjust.

The interpretation that households tend to associate rising gasoline prices with rising prices of other consumer goods is also supported by evidence from the missing disinflation period. It can be shown that the rising inflation expectations during the recovery of the price of gasoline between 2009 and 2011 were not driven by increases in expected gasoline price inflation, but by rising expected core inflation. Monthly data for one-year head gasoline price inflation expectations can be recovered from the Michigan Survey of Consumers, starting in February 2006, with the help of the EIA data on the price of motor gasoline. If we are willing to accept that $\pi_t^{\text{exp}} \approx 0.03\pi_t^{\text{gas,exp}} + 0.97\pi_t^{\text{core,exp}}$, we can infer $\pi_t^{\text{core,exp}}$ from the data for π_t^{exp} and $\pi_t^{\text{gas,exp}}$. As Figure 9 shows, at annualized rates, households' expectation of one-year gasoline price inflation fell from 25% in January 2009 to 5% in December 2011, while their expectation of one-year core inflation rose from 1.8% in January 2009 to 3.9% in December 2010. This fact is consistent with households expecting positive growth in gasoline prices to be slowly passed through to core prices, with households expecting core consumer prices to be driven up by a domestic economic recovery, or some combination thereof. If households view shocks to gasoline prices as being tied to broader domestic inflation pressures, the response of headline inflation expectations to nominal gasoline price shocks we documented in section 5.1 is not surprising.

A number of recent studies has argued that the increases in inflation expectations in the late 1990s and in the 2000s (as well as the decline in inflation expectations after 2014 and in early 2020) reflected the fact that the oil price during these episodes signaled broader changes in global demand (see, e.g., Sussman and Zohar 2015; Elliott, Jackson, Raczko and Roberts-Sklar 2015). These broader changes in global demand may have a direct effect on expected core inflation rates in the United States, as they stimulate the U.S. economy, beyond the direct effects operating through the price of gasoline. A recent example of this mechanism is the global financial crisis of 2008, when a collapse in the global demand for oil lowered U.S. gasoline prices. At the same time, this event reduced the ability of U.S. firms to raise prices, lowering actual inflation, and made households more pessimistic about the future economy, causing them to lower their inflation expectations. The view that household inflation expectations respond to gasoline price shocks because they tend to signal changes in broader economic conditions is consistent with our argument in this section that households extrapolate from the experience of the 1970s and 1980s, when oil and gasoline price fluctuations reflected global demand pressures (see Kilian 2008).

Some studies go as far as suggesting that households in forming their inflation expectations differentiate between oil price increases driven by global oil supply and oil price increases driven by global demand. This idea does not seem plausible. There is no doubt that differentiating between oil demand and oil supply shocks, building on the work of Kilian (2009), may help explain the excess sensitivity of one-year inflation expectations, but there are two obvious problems with this argument. First, this explanation is inherently anachronistic because the distinction between oil supply and oil demand shocks was only introduced in Barsky and Kilian (2002) and was only made operational in Kilian (2008, 2009). Thus, households could not possibly have been aware of the distinction between oil demand and oil supply shocks during many of the episodes of interest. Second, surely

professional inflation forecasters would have been able to exploit this distinction as well as households, yet there is no indication of a similar increase in the one-year SPF inflation forecast during 2009-13 (see Coibion and Gorodnichenko 2015).¹⁰ A model that assumes that households with no formal training in economics are better able to extract information about the state of the global economy than professional forecasters does not seem plausible.

6. Concluding Remarks

The conventional wisdom that inflation expectations respond to the level of the price of oil (or the price of gasoline) is based on estimates of static single-equation regressions fit to aggregate data. We identified a number of concerns about the specification of these regressions, each of which reduces the strength of this relationship in the data. We also drew attention to the fact that the t-statistics for testing the absence of feedback from the level of the price of oil have a nonstandard null distribution. We illustrated that under plausible assumptions the finite-sample critical values are considerably larger than the $N(0,1)$ critical values used in the literature. We showed that evidence in favor of a statistically significant impact of the level of oil prices on inflation expectations based on Gaussian critical values tends to be spurious. None of the model specifications we considered supported the conventional wisdom. We also cautioned against relying on estimates of the correlation between inflation expectations and the level of the price of oil (or of gasoline), which by construction tend to be erratic in small samples, given that these prices are not mean reverting, and unstable over time. The high correlations frequently referenced in the literature are only found in recent data, but vanish when using longer samples.

We then provided evidence based on an alternative multivariate structural regression model that gasoline price shocks may indeed cause one-year household inflation expectations

¹⁰ Likewise, there is no evidence that professional forecasters of real GDP growth understood the origins of the global demand boom in the 2000s (see Kilian and Hicks 2013).

to become temporarily unanchored. This conclusion is robust to only including data since 1990. We formally identified several episodes since 1990, when inflation expectations rose or fell substantially in response to nominal gasoline price shocks.

Our analysis confirms that gasoline price shocks may mask the underlying trends in one-year inflation expectations. It also supports the view that the rise in household inflation expectations between 2009 and 2013 can be almost entirely explained by the large increase in oil prices (and hence gasoline prices) over this period. However, on average, gasoline price shocks account for only 39% of the variation in household inflation expectations, with an additional 54% explained by idiosyncratic household expectations shocks and 7% by shocks to the core CPI. Gasoline price shocks can also be linked to deviations between short-run and long-run inflation expectations.

References

- Alquist, R., Kilian, L., and R.J. Vigfusson (2013), "Forecasting the price of oil," in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, 2013, 427-507.
- Anderson, S.T., Kellogg, R., and Sallee, J.M. (2013), "What do consumers believe about future gasoline prices?," *Journal of Environmental Economics and Management*, 66, 383-403.
- Andrews, D.W.K. (1991), "Heteroskedasticity and autocorrelation consistent covariance matrix estimation," *Econometrica*, 59, 817-858.
- Angelico, C., and F. Di Giacomo (2019), "Heterogeneity in inflation expectations and personal experience," manuscript, Bank of Italy,
- Arias, J.E., Rubio-Ramirez, J.J., and D.F. Waggoner (2018), "Inference based on structural vector autoregressions identified by sign and zero restrictions: Theory and applications," *Econometrica*, 86, 685-720.
- Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in B. Bernanke and K. Rogoff (eds.), *NBER Macroeconomics Annual 2001*, 137-183.
- Bernanke, B. (2010), "The economic outlook and monetary policy." Speech, Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 27. <http://www.federalreserve.gov/newsevents/speech/bernanke20100827a.htm>.
- Binder, C. (2018), "Inflation expectations and the price at the pump," *Journal of Macroeconomics*, 58, 1-18.
- Binder, C., and C. Makridis (2020), "Stuck in the seventies: Gas prices and macroeconomic expectations," forthcoming: *Review of Economics and Statistics*.
- Clark, T.E., and S.J. Terry (2010), "Time variation in the inflation passthrough of energy

- prices,” *Journal of Money, Credit and Banking*, 42, 1419-1433.
- Coibion, O., and Y. Gorodnichenko (2015), “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation,” *American Economic Journal: Macroeconomics*, 7, 197-232.
- Coibion, O., Gorodnichenko, Y., and R. Kamdar (2018), “The formation of expectations, inflation, and the Phillips curve.” *Journal of Economic Literature*, 56, 1447-1491.
- Coibion, O., Gorodnichenko, Y., Kumar, S., and M. Pedemonte (2020), “Inflation expectations as a policy tool?,” forthcoming: *Journal of International Economics*.
- Conflitti, C., and R. Cristadoro (2018), “Oil prices and inflation expectations”, Banca d’Italia Occasional Papers No. 423.
- Conflitti, C. and M. Luciani (2019). “Oil price pass-through into core inflation,” *Energy Journal*, 40, 221-247.
- Edelstein, P., and L. Kilian (2009), “How sensitive are consumer expenditures to retail energy prices?,” *Journal of Monetary Economics*, 56, 766-779.
- Elliott, D., Jackson, C., Raczko, M., and M. Roberts-Sklar (2015), “Does oil drive financial market measures of inflation expectations”, Bank of England.
- Hasenzagl, T, F Pellegrino, L Reichlin, and G Ricco (2018), “A model of the Fed’s view on inflation,” CEPR Discussion Paper No. 12564.
- Inoue, A., and L. Kilian (2020a), “Joint Bayesian inference about impulse responses in VAR models,” manuscript, Federal Reserve Bank of Dallas.
- Inoue, A., and L. Kilian (2020b), “The role of the prior in estimating VAR models with sign restrictions,” manuscript, Federal Reserve Bank of Dallas.
- Jorgensen, P.L., and K.J. Lansing (2019), “Anchored inflation expectations and the flatter Phillips curve,” Federal Reserve Bank of San Francisco Working Paper 2019-27.
- Kilian, L. (2008), “The economic effects of energy price shocks,” *Journal of Economic*

- Literature*, 46, 871-909.
- Kilian, L. (2009). "Comment on 'Causes and Consequences of the Oil Shock of 2007-08' by James D. Hamilton," *Brookings Papers on Economic Activity*, 1, Spring, 267-278.
- Kilian, L., and B. Hicks (2013), "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?," *Journal of Forecasting*, 32, 385-394.
- Kilian, L., and H. Lütkepohl (2017), *Structural Vector Autoregressive Analysis*, Cambridge University Press: Cambridge.
- Kilian, L., and L.T. Lewis (2011), "Does the Fed Respond to Oil Price Shocks?," *Economic Journal*, 121, 1047-1072.
- Kilian, L., and C. Vega (2011), "Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices," *Review of Economics and Statistics*, 93, 660- 671.
- Madeira, C., and B. Zafar (2015), "Heterogenous inflation expectations and learning," *Journal of Money, Credit and Banking*, 47, 867-896.
- Stewart, C. (2011), "A note on spurious significance in regressions involving I(0) and I(1) variables," *Empirical Economics*, 41, 565-571.
- Sussman, N., and O. Zuhar (2015), "Oil prices, inflation expectations, and monetary policy", Bank of Israel Discussion Paper No. 092015.
- Wong, B. (2015), "Do inflation expectations propagate the inflationary impact of real oil price shocks? Evidence from the Michigan Survey," *Journal of Money, Credit and Banking*, 47, 1673-1689.

Table 1: An overview of the regression evidence for the level of the price of oil

Dependent variable: $\pi_t^{\text{exp}} - \pi_t^{\text{exp},SPF}$	1981Q3- 2013Q1	1990Q1- 2013Q1	2000Q1- 2013Q1	1981Q3- 2020Q1	1990Q1- 2020Q1
$corr(\pi_t^{\text{exp}} - \pi_t^{\text{exp},SPF}, O_t)$	77.9%	85.1%	84.6%	78.6%	82.8%
$\hat{\beta}$	0.024	0.021	0.022	0.022	0.019
\hat{t}_{β}	12.91	15.71	10.94	10.83	11.80
R^2	60.7%	72.4%	71.6%	61.8%	68.5%
Dependent variable: π_t^{exp}	1981Q3- 2013Q1	1990Q1- 2013Q1	2000Q1- 2013Q1	1981Q3- 2020Q1	1990Q1- 2020Q1
$corr(\pi_t^{\text{exp}}, O_t)$	11.4%	35.0%	74.4%	3.7%	29.3%
$\hat{\beta}$	0.004	0.009	0.021	0.001	0.007
\hat{t}_{β}	1.01	2.16	5.38	0.29	1.68
R^2	1.3%	12.3%	55.3%	0.1%	8.6%
Dependent variable: π_t^{exp}	1981M7- 2013M3	1990M1- 2013M3	2000M1- 2013M3	1981M7- 2020M3	1990M1- 2020M3
$corr(\pi_t^{\text{exp}}, O_t)$	11.6%	33.9%	71.2%	4.6%	29.0%
$\hat{\beta}$	0.004	0.009	0.021	0.002	0.008
\hat{t}_{β}	1.08	2.24	5.39	0.39	1.80
R^2	1.3%	11.5%	50.7%	0.2%	8.4%

NOTES: Estimates based on regressions of $\pi_t^{\text{exp}} - \pi_t^{\text{exp},SPF}$ and π_t^{exp} , respectively, on an intercept and O_t for alternative estimation periods and data frequencies. \hat{t}_{β} based on Newey-West standard errors with a truncation lag of 8 for quarterly data and 24 for monthly data. Results that correspond to findings discussed in Coibion and Gorodnichenko (2015) are shown in bold. Our estimates are substantively similar or identical to theirs.

Table 2: Estimates of equations (1) and (1'), 1990.1-2020.4

	Level of oil price	Level of Gasoline price	Log-level of oil price	Log-level of gasoline price	
Correlation with π_t^{exp}	29.4%	22.1%	20.2%	14.6%	
R^2	8.6%	4.9%	4.1%	2.1%	
$\hat{\beta}$	0.008	0.190	0.242	0.256	
$t_{\hat{\beta}}$	NW(12)	2.04	1.48	1.44	0.96
	NW(24)	1.83	1.27	1.26	0.82
	NW(48)	1.59	1.06	1.08	0.68
	NW(Andrews)	1.83	1.26	1.25	0.81

NOTES: The standard errors underlying the t-statistics are computed based on Newey-West standard errors using alternative truncation lags of 12, 24, and 48 as well as the data-based estimator of the truncation lag proposed by Andrews (1991).

Table 3: Finite-sample critical values based on equations (1) and (1'), 1990.1-2020.4

		50%	90%	95%
Oil price	Level	0.01	1.86	2.45
	Log-level	0.01	1.77	2.33
Gasoline price	Level	0.00	1.81	2.38
	Log-level	0.00	1.78	2.33
N(0,1) critical values		0	1.28	1.65


NOTES: All results based on 100,000 Monte Carlo trials. The t-statistics are based on Newey-West standard errors with the truncation lag selected as in Andrews (1991). The DGP is described in the text.

Table 4: Finite-sample p-values based on equations (1) and (1'), 1990.1-2020.4

Level of oil price	0.102
Log-level of oil price	0.179
Level of gasoline price	0.180
Log-level of gasoline price	0.274


NOTES: See Table 3.

Figure 1: Household inflation expectations and the price of oil: 1990.1-2020.4




NOTES: The oil price is the spot price for WTI crude oil reported by the EIA. MSC denotes the Michigan Survey of Consumers mean one-year inflation expectations.

Figure 2: Alternative representations of the WTI price of crude oil, 1978.1-2020.4




NOTES: The oil price is the spot price for WTI crude oil reported by the EIA.

Figure 3: Finite-sample null distributions for the t-test of $H_0 : \beta = 0$



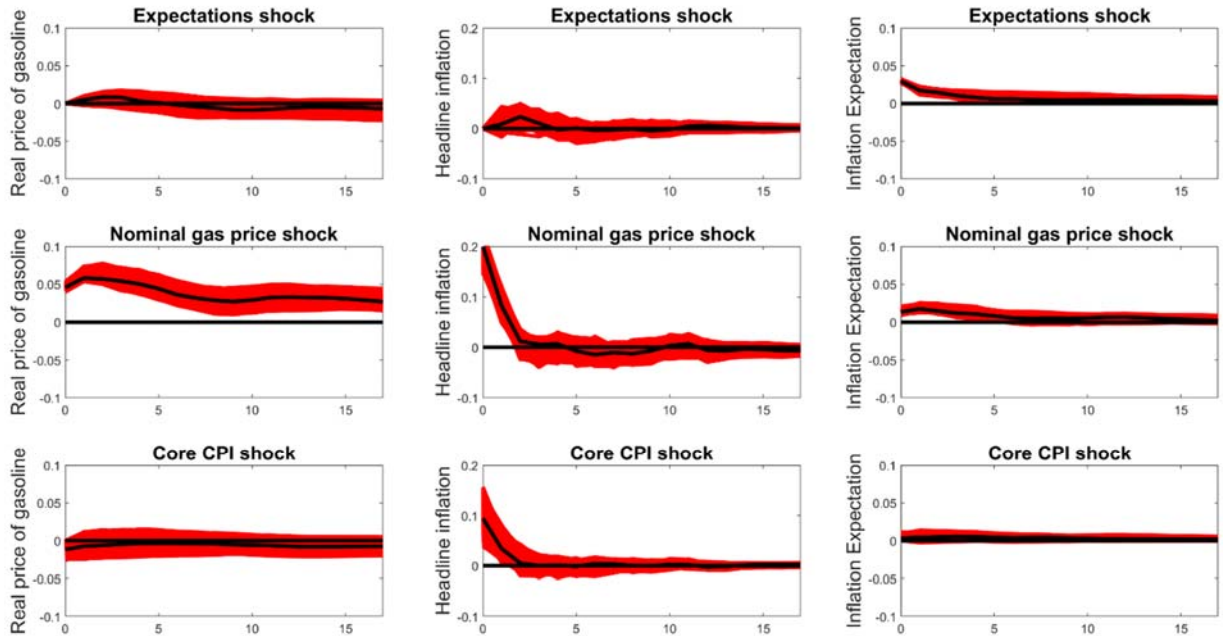
NOTES: All results based on NW(Andrews). Qualitatively similar results are obtained with fixed truncation lags. Based on 100,000 Monte Carlo trials from the data generating process described in the text.

Figure 4: Indicators used in the VAR analysis, 1981.7-2020.4




NOTES: The inflation rates are expressed in monthly percent changes. The real price of gasoline has been expressed in logs. All data have been demeaned.

Figure 5: Impulse response estimates and 68% joint credible sets, 1981.7-2020.4




NOTES: The set of impulse responses shown in black is obtained by minimizing the absolute loss function in expectation over the set of admissible structural models, as discussed in Inoue and Kilian (2020a). The responses in the corresponding joint credible set are shown in a lighter shade.

Figure 6: Historical decomposition of survey inflation expectations, 1990.1-2020.4




NOTES: The pre-1990 data have been discarded to reduce transient dynamics and to make the results compatible with the earlier regression evidence.

Figure 7: Actual mean one-year inflation expectation in *Michigan Survey of Consumers* and counterfactual series in the absence of nominal gasoline price shocks, 1990.1-2020.4




NOTES: The counterfactual time series is obtained by subtracting the cumulative effect of nominal gasoline prices shocks on household inflation expectations, as shown in Figure 6, from the actual data after rescaling the fitted data to represent annualized inflation rates.

Figure 8: Do nominal gasoline price shocks explain the unanchoring of one-year household inflation expectations?



NOTES: The cumulative effect of gasoline price on 1-year inflation expectations is from Figure 6, but expressed at annualized rates. The period shown is 1990.4-2020.4, given that the 5-year inflation expectations in the Michigan Survey of Consumers are only available starting in 1990.4.

Figure 9: Household core and gasoline price inflation expectations, 2009.1-2010.12

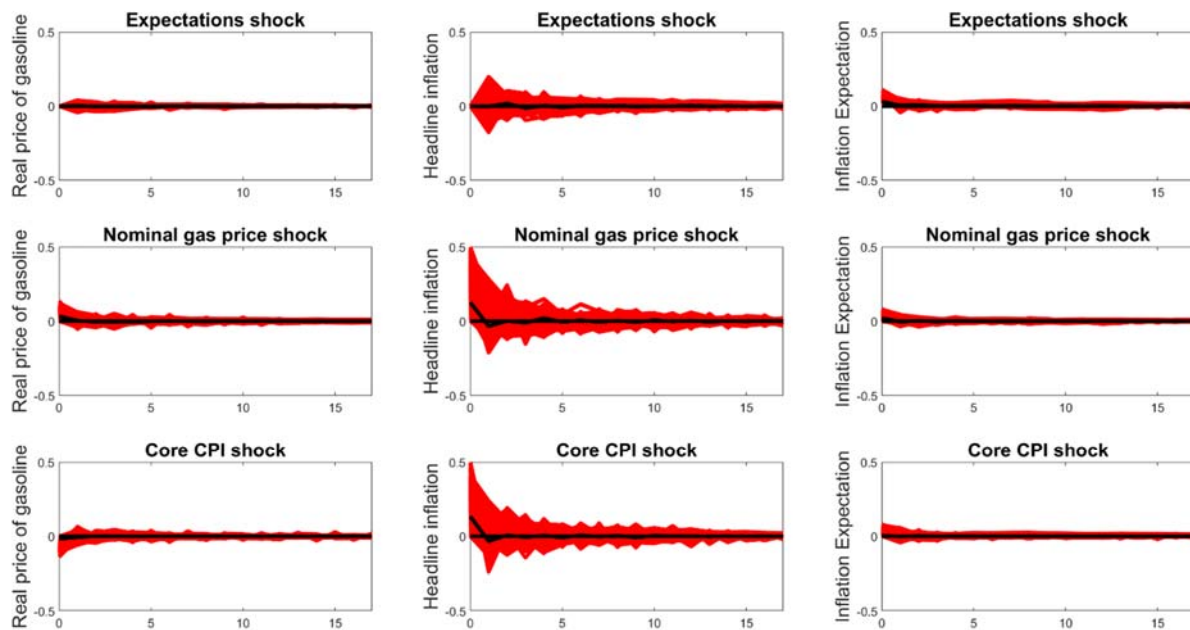


NOTES: The gasoline price inflation expectations are from the Michigan Survey of Consumers. The core inflation expectations are implied by the Michigan Survey of Consumer headline inflation expectations, taking account of the gasoline expenditure share of 3%.

Not-for-Publication Appendix:

This appendix illustrates that the uniform-Gaussian-inverse Wishart prior used in estimating the structural VAR model underlying Figure 5 is largely uninformative in that the response functions that minimize the absolute loss function under the prior are almost invariably flat, except where constrained by sign restrictions. In the rare cases, where these responses are not zero, the prior is not driving the posterior estimates shown in Figure 5. The differences between Figure A1 and Figure 5 are driven by the data.

Figure A1: Impulse response estimates and 68% joint credible sets simulated from the prior distribution



NOTES: The set of impulse responses shown in black is obtained by minimizing the absolute loss function in expectation over the set of admissible structural models, as discussed in Inoue and Kilian (2020a). The responses in the corresponding joint credible set are shown in a lighter shade.