

WORKING PAPER

Inflation expectations and the pass-through of oil prices

NORGES BANK
RESEARCH

5 | 2020

KNUT ARE AASTVEIT,
HILDE C. BJØRNLAND
AND JAMIE L. CROSS



NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:

FacilityServices@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatterens regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail

FacilityServices@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8190 (online)

ISBN 978-82-8379-153-2 (online)

Inflation expectations and the pass-through of oil prices*

Knut Are Aastveit[†] Hilde C. Bjørnland[‡] Jamie L. Cross[§]

June 30, 2020

Abstract

Do inflation expectations and the associated pass-through of oil price shocks depend on demand and supply conditions underlying the global market for crude oil? We answer this question with a novel structural vector autoregressive model of the global oil market that jointly identifies transmissions of oil demand and supply shocks through the real price of oil to both expected and realized inflation. Our main insight is that US households form their expectations of inflation differently when faced with long sustained increases in the price of oil, such as the early millennium oil price surge of 2003 to 2008, as compared to short and sharp price fluctuations that characterized much of the twentieth century. We also find that oil demand and supply shocks can explain a large proportion of expected and realized inflation dynamics during multiple periods of economic significance, and resolve disagreements around the role of oil prices in explaining the missing deflation puzzle of the Great Recession.

JEL-codes: E31, D84, Q41, Q43

Keywords: Inflation expectations, inflation pass-through, oil prices

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of the Norges Bank. The authors would like to thank Benjamin Wong, Renee Fry-McKibbin, Warwick McKibbin, participants of the *International Economic Flows: Energy, Finance, Diplomacy and Market Structures* workshop in Sydney and seminar participants at BI Norwegian Business School for valuable comments. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School.

[†]Norges Bank & BI Norwegian Business School, Email: Knut-Are.Aastveit@Norges-Bank.no

[‡]BI Norwegian Business School & Norges Bank, Email: hilde.c.bjornland@bi.no

[§]BI Norwegian Business School, Email: jamie.cross@bi.no

1 Introduction

A central tenant of modern monetary policy is that the anchoring of inflation expectations is necessary for achieving stable prices.¹ Since such expectations are sensitive to increases in the price of oil, understanding the degree to which expectations facilitate the inflation pass-through of oil price shocks is an important policy question that has achieved much academic attention (see, e.g. [Bernanke et al. \(1997\)](#); [Hooker \(2002\)](#); [Harris et al. \(2009\)](#); [Coibion and Gorodnichenko \(2015b\)](#) and [Wong \(2015\)](#) among others). Despite this fact, there is currently no consensus on the empirical strength of this mechanism. For instance, while [Coibion and Gorodnichenko \(2015b\)](#) argue that the high sensitivity of household's inflation expectations to oil price shocks in the US can help explain the missing deflation puzzle of the Great Recession, other studies such as [Blanchard and Galí \(2007\)](#); [Blanchard and Riggi \(2013\)](#) and [Wong \(2015\)](#) suggest that this mechanism is weak at best, and may have altogether disappeared since the 1990s.

In this paper, we question whether inflation expectations and any associated oil price pass-through depends on demand and supply conditions underlying the global market for crude oil. Our question is motivated by the idea that households may form their expectations of inflation differently when faced with long sustained increases in the oil price, such as the early millennium oil price surge of 2003-2008, as compared to short and sharp price increases that characterized much of the twentieth century. If this hypothesis is true, then it not only adds to our understanding of how oil price shocks propagate throughout the economy, but may provide a resolution of the aforementioned debate about the empirical strength of the inflation expectations mechanism of oil price pass-through.

To model the relationship between oil prices, inflation expectations and realized inflation, we extend the structural vector autoregressive (SVAR) model of the global market for crude oil developed in [Baumeister and Hamilton \(2019a\)](#) to include monthly measures of expected and realized inflation in the US. Our Bayesian inference and identification has the benefit of being based on sign restrictions that are less restrictive than commonly used alternatives in the literature, such as e.g. [Kilian \(2009\)](#); [Kilian and Murphy \(2012, 2014\)](#), and further has the advantage of accounting for uncertainty about the identifying assumptions themselves. As highlighted by [Baumeister and Hamilton \(2019a\)](#), this is crucial for identifying demand and supply shocks to the oil market, and more generally for drawing structural conclusions in VARs, see also [Baumeister and Hamilton \(2015, 2018, 2019b\)](#).

Our results not only support existing evidence that inflation expectations are sensitive

¹[Coibion et al. \(2018a\)](#) survey the literature on how economic agents form their inflation expectations.

to oil price shocks (Harris et al., 2009; Coibion and Gorodnichenko, 2015a,b; Wong, 2015; Güntner and Linsbauer, 2018), but also provide the novel insight that the degree of sensitivity depends on the underlying source of oil market shock. In particular, we find that demand for oil associated with unexpected large global economic activity shocks, such as the early millennium oil price surge of 2003-08 (Aastveit et al., 2015), elicits an extremely persistent response in both expected and actual inflation. In contrast, when the economy is hit by shocks to oil supply, consumption demand, or inventory demand, we find that both expected and actual inflation initially increase but then revert back to zero within the following quarter. This suggests that the way in which households form their expectations differs depending on the type of oil price shock underlying the global market for crude oil.

Having shown that both expected and realized inflation are sensitive to oil price shocks, we then investigate their relative effects during four commonly studied periods of economic significance: the Iraqi invasion of Kuwait in 1990/91, the oil price surge of 2003-08, the oil price decline of the great recession and the collapse of the oil price in 2014. We find that while supply shocks can account for some of the increase in inflation expectations around the Iraqi invasion of Kuwait in 1990, they don't explain much of the inflation dynamics since then. Instead, much of the fluctuations in expected inflation is accounted for by unanticipated fluctuations in demand for crude oil. For instance, there is a close mapping between consumption demand and inflation expectations during the 1997/98 Asian Financial Crisis, while demand from economic activity played a key role in driving the persistent increases in expectations throughout the oil price surge of 2003-08 and the subsequent collapse in expectations in 2009. Finally, consumption and economic activity shocks jointly explain the persistent reduction in inflation expectations since the oil price drop of 2014.

A possible critique of our analysis, is that observing large movements in expected and actual inflation following an oil price shock is insufficient for concluding that expectations are an important driver of inflation pass-through. This is because pass-through arises through both *direct* and *indirect* mechanisms. On the one hand, the cost channel suggests that oil price increases directly feed through to prices via higher costs of production. On the other, higher inflation expectations indirectly pass-through the oil price increases via price setting and wage bargaining mechanisms. As a result, any dynamics from impulse response functions or historical decompositions will conflate these two mechanisms, and their isolation requires a counterfactual.

To address this issue, we build on the counter-factual analysis in Wong (2015), and approximate the direct mechanism by examining whether the inflationary response from an oil price increase differs under the assumption that expectations are insensitive to

each of the oil price shocks. Set in this manner, any difference between the actual and counterfactual inflationary effects must be due to a propagation by inflation expectations, i.e. the indirect effect. Given that our Bayesian analysis provides posterior distributions for both the actual and counter-factual impulse response functions, we then determine the relative strength of the two channels by testing whether the two sets of responses are drawn from the same underlying continuous population distribution using a one-sided two-sample Kolmogorov-Smirnov test ([Massey Jr, 1951](#); [Miller, 1956](#)).

The results from our test not only provide overwhelming evidence that the actual and counterfactual impulse response functions are drawn from distinct population distributions, but also that inflation tends to be smaller in the counterfactual compared to the actual responses. On average, holding expectations constant after each oil price shock will reduce the CPI level by 0.4-2.2 percent over the next 18 months, with the largest reduction stemming from economic activity shocks associated with changes in the global business cycle. This suggests that inflation expectations are an important channel in the pass-through of real oil price increases into realized inflation.

In light of these new insights, we use our framework to revisit the much debated question of whether high inflation expectations during the Great Recession can account for the missing deflation. We find that if the real price of oil was held constant throughout 2009-11, then expected and realized inflation would have respectively been 1.5 and 1 percent lower than observed. Given that observed expectations increased from around 2.5 in January 2009 to 4 percent in December 2011, our results suggest that oil price fluctuations can account for all of the rise in household inflation expectations during this period. Thus, in line with [Coibion and Gorodnichenko \(2015b\)](#) we conclude that if firms' inflation expectations track those of households, then the rise in their inflation expectations between 2009 and 2011 can explain the missing deflation during this period.

The paper is organized as follows. We detail the methodology in [Section 2](#), discuss the results in [Section 3](#), and conclude in [Section 4](#).

2 Empirical Methodology

To model the effects that oil demand and supply shocks have on expected and realized inflation, we require a structural model that is capable of jointly identifying the underlying drivers in the world market for crude oil. To this end, we build on the recently developed SVAR model of the global market for crude oil developed in [Baumeister and Hamilton \(2019a\)](#). The primary difference between this model and earlier oil market SVARs, e.g. [Kilian \(2009\)](#); [Kilian and Murphy \(2012, 2014\)](#), is the use of explicit, as opposed to implicit, Bayesian priors in identifying the structural parameters and shocks. As highlighted

by [Baumeister and Hamilton \(2019a\)](#), this feature is crucial for identifying supply and demand shocks to the oil market, and more generally for drawing structural conclusions in VARs, see also [Baumeister and Hamilton \(2015, 2018, 2019b\)](#). This is because traditional identifying assumptions can be viewed as implicit priors, i.e. information that the analyst had about the economic structure before seeing the data. For instance, traditional point identification, such as placing zero restrictions on certain contemporaneous relationships, can be viewed as the researcher having a dogmatic prior. In general, however, we are often uncertain about the validity of such identifying assumptions. Rather than sweeping such assumption under the rug, we explicitly model this uncertainty by increasing the variance in those relevant prior distributions. Information from the data will then update our prior belief via a straight forward application of Bayes' Theorem and inference is drawn accordingly.

2.1 Structural VAR Model

Let $\mathbf{y}_t = (\mathbf{y}_t^o, \mathbf{y}_t^\pi)'$ denote an $(n^o + n^\pi) \times 1$ vector in which \mathbf{y}_t^o is an $n^o \times 1$ vector of variables associated with the global market for crude oil and \mathbf{y}_t^π is an $n^\pi \times 1$ vector of inflation variables. Following [Baumeister and Hamilton \(2019a\)](#), we set $\mathbf{y}_t^o = (q_t, y_t, p_t, i_t)'$ in which q_t is the percentage change in global crude oil production, y_t is the percentage change in global real economic activity, p_t is the percentage change in the global real price of oil, and i_t is the observable change in above-ground global crude oil inventories as a percent of the previous month's world production. Each variable is observed at a monthly frequency. Since the true change in above-ground global crude oil inventories, i_t^* , is not directly observable, i_t will be modeled via measurement error. Given these definitions, the global market for crude oil is then modeled through the following five equations

$$q_t = c_1 + \alpha_{qp}p_t + \mathbf{b}'_1\mathbf{x}_{t-1} + u_{1t}^*, \quad (1)$$

$$y_t = c_2 + \alpha_{yp}p_t + \mathbf{b}'_2\mathbf{x}_{t-1} + u_{2t}^*, \quad (2)$$

$$q_t = c_3 + \beta_{qp}p_t + \beta_{qy}y_t + i_t^* + \mathbf{b}'_3\mathbf{x}_{t-1} + u_{3t}^*, \quad (3)$$

$$i_t^* = c_4 + \psi_1^*q_t + \psi_2^*y_t + \psi_3^*p_t + \mathbf{b}'_4\mathbf{x}_{t-1} + u_{4t}^*, \quad (4)$$

$$i_t = \chi i_t^* + e_t, \quad (5)$$

where c_j , $j = 1, 2, 3, 4$ are intercept terms and $\mathbf{x}_{t-1} = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-12})'$ is a vector of lagged observations over the past year. A lag length of one year is designed to account for the long lags in the transmission of oil price shocks ([Hamilton and Herrera, 2004](#)), and was also used in both the SVARs of [Baumeister and Hamilton \(2019a\)](#) and [Wong \(2015\)](#). Equation (1) is the oil supply curve in which α_{qp} is the short-run price elasticity of supply and u_{1t}^* is an *oil supply shock*. Equation (2) models the determinants of global

real economic activity, in which α_{yp} captures any contemporaneous effects of oil prices and u_{2t}^* is an *economic activity shock* associated with changes in the global business cycle. Equation (3) is the oil demand curve in which β_{qp} is the short-run price elasticity of demand, while β_{qy} captures any contemporaneous effects of global real economic activity, i_t^* denotes the true change in global crude-oil inventories as a percent of the previous month's world production, and u_{3t}^* is a *consumption demand shock*. Equation (4) models the true dynamics in above-ground crude oil inventories, and u_{4t}^* an *inventory demand shock*. Finally, equation (5) is a measurement-error equation that accounts for the fact that observable data on global crude-oil inventories does not capture the true quantity of available inventories. Thus, i_t denotes the change in observable crude-oil inventories as a percent of the previous month's world production, $\chi < 1$ is a parameter that captures the fact that observable inventories are a proportion of the world total, and e_t is a measurement error which is assumed to be uncorrelated with the structural shocks.

Note that the specification in (1)-(5) assumes that the global oil market variables are predetermined to the US inflation block. This is motivated by the idea that monthly news about US changes in expected or realized US inflation is unlikely to impact the real price of crude oil, or the decisions of global oil producers.

To examine the effects that these fundamental oil market shocks have on inflation expectations, π^e and realized inflation π , we define $\mathbf{y}_t^\pi = (\pi^e, \pi)'$, and specify the following two equations

$$\pi^e = c_5 + \lambda_{\pi^e p} p_t + \mathbf{b}'_5 \mathbf{x}_{t-1} + u_{5t}^*, \quad (6)$$

$$\pi = c_6 + \gamma_{\pi p} p_t + \gamma_{\pi \pi^e} \pi_t^e + \gamma_{\pi y} y_t + \mathbf{b}'_6 \mathbf{x}_{t-1} + u_{6t}^*. \quad (7)$$

Equation (6) models inflation expectations dynamics which contemporaneously depends on the real price of oil through $\lambda_{\pi^e p}$, and an inflation expectations shock u_{5t}^* , but is only impacted by the remaining oil market variables and realized inflation through historical movements. Equation (7) models inflation dynamics, which responds contemporaneously to the real price of oil, inflation expectations, real economic activity and an idiosyncratic inflation shock u_{6t}^* .

We have made three assumptions in specifying the inflation block. In the first instance, we assume that expectations are not contemporaneously impacted by realized inflation. This reflects the time lag between the date at which households must submit their expectations to the survey and the release of actual inflation data. Second, we assume that expected inflation is contemporaneously independent of all oil market variables other than the real price of oil. This reflects the idea that households are most likely ignorant about the underlying dynamics in the global market for crude oil and only learn about the effects of shocks through the purchase of goods that are highly dependent on

the real price of oil, e.g. gasoline and automobile purchases (Hamilton, 2009; Coibion and Gorodnichenko, 2015b). Third, we assume that realized inflation is contemporaneously independent of oil production and oil inventories. In this way equation (7) can be interpreted as a Phillips curve relation.

For estimation purposes, it is easy to check that substituting (5) into (3) and (4) and combining (1)-(7) yields a system of the form

$$\tilde{\mathbf{A}}\mathbf{y}_t = \tilde{\mathbf{B}}\mathbf{x}_{t-1} + \tilde{\mathbf{u}}_t, \quad (8)$$

where $\mathbf{x}_{t-1} = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-12}, 1)'$,

$$\tilde{\mathbf{B}} = \begin{bmatrix} \mathbf{b}'_1 & c_1 \\ \mathbf{b}'_2 & c_2 \\ \mathbf{b}'_3 & c_3 \\ \mathbf{b}'_4 & c_4 \\ \mathbf{b}'_5 & c_5 \\ \mathbf{b}'_6 & c_6 \end{bmatrix}, \quad \tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 & 0 & 0 \\ 0 & 1 & -\alpha_{yp} & 0 & 0 & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} & 0 & 0 \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 & 0 & 0 \\ 0 & 0 & -\lambda_{\pi^e p} & 0 & 1 & 0 \\ 0 & -\gamma_{\pi y} & -\gamma_{\pi p} & 0 & -\gamma_{\pi \pi^e} & 1 \end{bmatrix}, \quad \tilde{\mathbf{u}}_t = \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ u_{3t}^* - \chi^{-1}e_t \\ \chi u_{4t}^* + e_t \\ u_{5t}^* \\ u_{6t}^* \end{bmatrix}.$$

This representation shows that u_{3t}^* and u_{4t}^* are correlated. To generate a representation in which the shocks are uncorrelated, the system in (8) can be pre-multiplied by

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \rho & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

in which $\rho = \frac{\chi^{-1}\sigma_e^2}{d_{33}^* + \chi^{-2}\sigma_e^2}$, and d_{33}^* comes from the representation

$$\mathbf{A}\mathbf{y}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u}_t, \quad \mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{D}), \quad (9)$$

in which $\mathbf{A} = \mathbf{\Gamma}\tilde{\mathbf{A}}$, $\mathbf{B} = \mathbf{\Gamma}\tilde{\mathbf{B}}$, $\mathbf{u}_t = \mathbf{\Gamma}\tilde{\mathbf{u}}_t$, and

$$\mathbf{D} = \begin{bmatrix} d_{11}^* & 0 & 0 & 0 & 0 & 0 \\ 0 & d_{22}^* & 0 & 0 & 0 & 0 \\ 0 & 0 & d_{33}^* + \chi^{-2}\sigma_e^2 & -\chi^{-2}\sigma_e^2 & 0 & 0 \\ 0 & 0 & -\chi^{-2}\sigma_e^2 & \chi^2 d_{44}^* + \sigma_e^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & d_{55}^* & 0 \\ 0 & 0 & 0 & 0 & 0 & d_{66}^* \end{bmatrix}.$$

Written in this manner, the SVAR with uncorrelated shocks in (9) can be estimated using the Bayesian methods discussed in Appendix B of [Baumeister and Hamilton \(2019a\)](#), for which we present priors in the next section. Structural inference can then be made by calculating the matrices in (8).

2.2 Priors

The priors on parameters in the global market for crude oil are in line with those in Table 1 of [Baumeister and Hamilton \(2019a\)](#), p.1892). This includes the use of priors on (1) the determinant of $\tilde{\mathbf{A}}$ and (2) the equilibrium feedback effects in the impact matrix, $\tilde{\mathbf{A}}^{-1}$. This is permissible because both the determinant of our $\tilde{\mathbf{A}}$ matrix, and the equilibrium feedback effect of interest, are identical to those in their paper. We prove this result formally in Appendix 1. Intuitively, these results hold because of the identification assumption that expected and realized inflation have no contemporaneous impact on the oil market block. Given this result, we also follow [Baumeister and Hamilton \(2019a\)](#), p.1893) and restrict $\psi_2 = 0$ to identify the oil market shocks. Finally, since expected and realized inflation in the US are likely to have little impact on global oil market variables, e.g. global oil production or inventory accumulation decisions, we center the appropriate lagged coefficients on zero, with a variance of 10^{-2} .

The inflation block and subsequent parameters are new to this paper. We set the priors for structural variances and lagged structural coefficients in the same manner as the oil market block, with each of the lagged coefficients being centered on zero. In setting the priors affecting contemporaneous coefficients we emulate the approach proposed by [Baumeister and Hamilton \(2019a\)](#) and propose a set of priors for those coefficients based on different sources of information. As shown in Table 1, we use a set of Student's-t prior densities truncated to positive support. This captures the fact that both expected and realized inflation will react to an unexpected increase in the real price of oil in a non-negative manner. Similarly, higher inflation expectations should, if anything, result in a positive price change. When choosing the location of each distribution we appeal to the

Table 1: Student’s t prior distributions for structural parameters in the inflation block

Parameter	Definition	Location	Scale	Degrees of freedom	Sign restriction
$\lambda_{\pi^e p}$	Effect of real price of oil on inflation expectations	0.02	1	3	positive
$\gamma_{\pi p}$	Effect of real price of oil on realized inflation	0.04	1	3	positive
$\gamma_{\pi \pi^e}$	Effect of inflation expectations on realized inflation	1	1	3	positive
$\gamma_{\pi y}$	Effect of real economic activity on realized inflation	0.25	1	3	positive

existing studies. In the first instance, we use the estimates in [Coibion and Gorodnichenko \(2015b, p.226\)](#) and [Coibion et al. \(2018a, p.1483\)](#) to center the real oil price effect on expected and realized inflation on 0.02 and 0.04, respectively. Regression estimates from [Coibion et al. \(2018a, p.1464\)](#) are also used to center the effect of real economic activity on realized inflation at 0.25. Next, in line with the rational expectations hypothesis, we center the coefficient of inflation expectations in the inflation equation around unity ([Galí and Gertler, 1999](#); [Coibion et al., 2018a](#)). While [Coibion et al. \(2018a\)](#) present evidence to suggest that this value is around 1.36 for the Michigan survey, thereby deviating from the predictions of full information rational expectations, we think that this is a prudent starting point, and allow the data to inform us of any deviations. Finally, given the magnitude of these values, in each case we use a unit valued scale parameter with 3 degrees of freedom.

2.3 Data

Our sample consists of six monthly time series from January 1983 to December 2019. This includes observations on the four fundamental oil market variables, US inflation expectations and realized inflation. In line with previous studies that examine the link between expected and realized inflation (e.g., [Harris et al. \(2009\)](#); [Wong \(2015\)](#)), the starting point is chosen to restrict the analysis to the post-Volcker deflation regime and to account for the changes in oil price elasticities since the mid-1980’s, see [Baumeister and Peersman \(2013a,b\)](#).

The variables in the global oil market block are measured as in [Baumeister and Hamilton \(2019a\)](#). First, the percentage change in global crude oil production is measured using global crude oil production data as provided by the U.S. Energy Information Administration (EIA). Second, global real economic activity is measured using the OECD+6 industrial production index provided by [Baumeister and Hamilton \(2019a\)](#). Third, the percentage change in the global real price of oil is measured by the US refiners’ acquisition cost (IRAC) for imported crude oil, as reported by the EIA, and deflated by the US consumer price index (all items), which were obtained from the FRED database. Finally, the change in above-ground oil global crude oil inventories is measured by total US crude

oil inventories scaled by the ratio of OECD petroleum stocks over US petroleum stocks, all of which are obtained from the EIA. Note that the model based proxy, i_t , defined as the change in inventories as a proportion of last periods oil production, is obtained by converting the production data into millions of barrels per month by multiplying the million barrels of crude oil per day by 30.

The variables in the inflation block are measured as in [Coibion and Gorodnichenko \(2015b\)](#) and [Wong \(2015\)](#). First, inflation expectations are measured by the median one-year-ahead inflation expectations Michigan Survey of Consumer Inflation Expectations. Second, inflation is the annualized month-on-month rate of change in the US consumer price index (all items).

3 Results

3.1 Posterior estimates of structural parameters

The posterior distributions of the new contemporaneous coefficients in the inflation block are shown in [Figure 1](#). The solid red curves represent the prior information and the blue histogram is the posterior distribution. The posterior median of the real oil price effect on expected and realized inflation are respectively 0.01 and 0.14. The former estimate is slightly lower than the prior mean, but quite close to the regression estimate in [Coibion and Gorodnichenko \(2015b\)](#), while the latter is larger, but of a similar magnitude to the value in [Coibion et al. \(2018a\)](#). Next, the posterior mean of inflation expectations effect on inflation is 1.92, thereby providing strong support that the behavior of respondents in the Michigan survey are not in accordance with rational expectations. This is in line with evidence presented in [Coibion et al. \(2018a\)](#) which shows that survey based inflation expectations tend to deviate from what full information rational expectations theory would predict. Finally, the posterior mean of global real economic activity's effect on realized inflation is 0.21, which is very close to the 0.25 estimate obtained in [Coibion et al. \(2018a\)](#).

Since some of the posterior median values are of similar magnitudes to their associated prior mean, we also estimated the model with a set of Student's-t prior densities located at 0, with scale parameter 100, and 3 degrees of freedom priors and unbounded support. These values are chosen so that the priors are uninformative over the entire real line. The results are in line with those presented throughout the paper.




Figure 1: Prior (solid red curves) and posterior (blue histograms) distributions of structural parameters in the inflation block

Notes: $\lambda_{\pi^e p}$ denotes the effect of the real price of oil on inflation expectations, $\gamma_{\pi p}$ denotes the effect of the real price of oil on realized inflation, $\gamma_{\pi \pi^e}$ denotes the effect of inflation expectations on realized inflation, and $\gamma_{\pi y}$ denotes the effect of real economic activity on realized inflation.

3.2 Do oil price shocks effect expected and realized inflation?

Do expected and realized inflation respond differently to the underlying source of shock in the global market for crude oil? A quick glance at the impulse response functions in Figure 2, suggest that the answer is yes. For instance, an economic activity shock, that increases oil prices with 10 percent, elicits an extremely persistent response in both expected and actual inflation. In contrast, when the economy is hit by shocks to supply, consumption demand, or inventory demand, we find that both expected and actual inflation initially increase but typically dies out after a quarter. This suggests that households




Figure 2: Structural impulse response functions

Notes: The structural shocks have been normalized to elicit a 10% increase in the real price of oil. The posterior median is shown in boldface and the shaded area is the 95% joint credible set obtained from the posterior distribution of 100,000 structural models.

form their expectations of inflation differently when faced with economic activity shocks as compared to alternative oil price shocks.

While impulse response functions are useful for analyzing the average response of each variable to shocks over the entire sample, they are silent on the relative importance of each shock during specific episodes of economic significance, such as the Great Recession. For this reason it is also useful to examine the historical decompositions of both expected and realized inflation to the various oil market shocks.

The results in Figures 3 and Figure 4 suggest that both expected and realized inflation have been quite sensitive to the oil market shocks over the sample period. For instance,




Figure 3: Historical Decomposition: Inflation Expectations

Notes: Actual changes in inflation expectations (red dotted lines) and median estimate of historical contribution of separate structural shocks (blue lines). Blue shaded regions indicate 95 percent posterior credibility regions from the posterior distribution of 100,000 structural models.

consumption demand shocks contributed towards both the high expected and realized inflation during the Iraqi invasion of Kuwait in 1990/91 and the subsequent decline during the 1997/98 Asian Financial Crisis. Such shocks also help explain the raising inflation expectations throughout the early millennium price surge of 2003-08, as well as low frequency inflation movements during this period. The results also suggest that economic activity shocks played a key role in shaping inflation expectations in the early millennium increase in inflation expectations. Combined together, consumption demand and economic activity shocks explain almost all of the sharp decline in both expected and realized inflation during the Great Recession, with the latter shock playing the primary role. Finally, while economic activity shocks explain the recent fall in inflation expectations since the oil price collapse of 2014, consumption demand shocks explain more of the inflation dynamics during this period.

Of course, this is not to say that oil price shocks can explain all of the movements in either expected or actual inflation over our sample period. For instance, there is no




Figure 4: Historical Decomposition: Inflation

Notes: Actual changes in inflation (red dotted lines) and median estimate of historical contribution of separate structural shocks (blue lines). Blue shaded regions indicate 95 percent posterior credibility regions from the posterior distribution of 100,000 structural models.

evidence that oil prices contributed towards inflation expectation dynamics during the mid-1990s or the sharp fall in 2002. Nonetheless, our results confirm that identifying the underlying oil market shock is extremely relevant for understanding inflation expectations both historically and in the present.

3.3 Counterfactual analysis

The evidence reported in the previous two sections showed that both expected and realized inflation are sensitive to oil price shocks. An important caveat noted by Wong (2015), however, is that a high degree of sensitivity does not necessarily imply that inflation expectations play a role in propagating the inflationary response of oil price shocks. This is because all of the inflationary response associated with an oil price shock may be due to the direct effects of higher production costs and not at all associated with the

indirect effects of price setting or wage bargaining mechanisms through which inflation expectations may impact inflation.

To address this potential concern, we consider a counterfactual in which inflation expectations are held constant during the impulse response analysis. This has the effect of isolating any increased inflation associated with shocks affecting the real oil price. Any observed difference between the inflationary effects in the actual and counterfactual impulse response functions must consequently be due to a propagation by inflation expectations. To ensure that the effects are comparable, we assume that each shock increases the real price of oil by 10 percent. Formal details of this exercise are deferred to Appendix B.

Figure 5 shows a comparison of the median counterfactual and actual IRFs. To conserve space we present the distributions of these responses to the Appendix. The main point to note are that the counterfactual IRFs contain more estimation uncertainty than the actual IRFs presented earlier in Figure 2. Nonetheless, comparing the median responses as in Figure 5 reveals that holding inflation expectations constant after an oil price shock will, on average, reduce the CPI level by 0.4-2.2 percent over the next 18 months, with the largest impacts stemming from economic activity shocks. This suggests that inflation expectations are an important channel in the pass-through of real oil price increases into realized inflation, and that failing to identify the underlying source of oil market shocks would underestimate the degree of pass-through.²

While this suggests that inflation expectations are an economically important channel in the pass-through of real oil price increases into realized inflation, observing that the averages of two distributions are different is insufficient for determining whether this difference is statistically significant. To formally test whether the counterfactual responses are statistically smaller than the actual responses, we consequently conduct a one-sided (two-sample) Kolmogorov-Smirnov test. The null hypothesis for the test is that the actual and counterfactual responses are generated from the same population distribution, and the one-sided alternative hypothesis is that the empirical cumulative distribution function of the actual IRFs is smaller than that of the counterfactual. Set in this manner, rejecting the null hypothesis implies that the observed IRFs tend to be larger than those in the counterfactual in a statistical sense.

²This is because averaging over the four oil market shocks would imply that holding inflation expectations constant after an oil price shock will, on average, reduce the CPI level by about 0.8 percent over the next 18 months.




Figure 5: Actual (blue solid line) and counterfactual (red dashed line) structural impulse response functions and their difference (magenta dotted line)

Notes: The structural shocks have been normalized to elicit a 10% increase in the real price of oil. The results presented are posterior medians obtained from the posterior distribution of 100,000 structural models.

Table 2 displays results from the one-sided (two-sample) Kolmogorov-Smirnov test.

Table 2: P-values for one-sided two-sample Kolmogorov-Smirnov test on actual and counterfactual inflation impulse response functions of CPI Inflation

Impulse horizon	Supply	Economic Activity Demand	Consumption Demand	Inventory Demand
1	0.00	0.00	0.00	0.00
2	0.00	0.00	0.32	0.67
3	0.00	0.00	0.00	0.00
4	0.35	0.00	1.00	1.00
5	0.00	0.00	1.00	0.00
6	0.00	0.00	0.00	0.00
7	1.00	0.00	0.00	1.00
8	1.00	0.00	0.00	1.00
9	0.00	0.00	0.00	1.00
10	0.00	0.00	0.00	0.00
11	0.00	0.00	0.00	0.00
12	0.00	0.00	0.00	0.00

Notes: The two-sample Kolmogorov-Smirnov test evaluates the difference between the empirical cumulative distribution functions (CDFs) for the actual and counterfactual impulse response functions (IRFs), i.e. $H_0 : \hat{F}_a = \hat{F}_c$, i.e. both samples come from a population with the same distribution. For the one-sided test: $H_1 : \hat{F}_a \leq \hat{F}_c$, i.e., the observed empirical CDF tends to be smaller than that of the counterfactual.

The table shows that the test often rejects the null hypothesis at the one percent level of statistical significance, thereby providing extremely strong evidence that inflation tends to be larger in the counterfactual. This suggests there is strong statistical support for the presence of indirect expectations effects in the inflation pass-through of oil price shocks. This is especially true when the real price of oil is responding to demand shocks associated with global economic activity, for which we reject the null hypothesis at each impulse horizon up to one year after the shock. In summary, if inflation expectations were held constant over our sample period, then, on average, inflation would have been significantly less than was actually observed.

3.4 Explaining the missing deflation during the Great Recession

In the year of February 2008 to 2009, US GDP contracted 4% while the unemployment rate rose by 3%. At the same time, despite falling almost 2% in November 2008, year-on-year CPI inflation was almost zero, and subsequent periods saw no deflation. This empirical fact called into question traditional backward-looking Phillips curve models of inflation, which suggest an inverse relationship between the price level and economic

output, and gave rise to the “missing deflation” puzzle (Ball and Mazumder, 2011; Hall, 2011). In an influential paper, Coibion and Gorodnichenko (2015b) argued that the puzzle can be explained by noticing that higher oil prices increased inflation expectations which prevented deflation. Despite this result, Wong (2015) found modest to nonexistent indirect effects of price setting or wage bargaining mechanisms, and suggested that the Coibion and Gorodnichenko (2015b) result was due to the fact that they confounded the direct and indirect effects.

In light of our results in the previous two sections, we here investigate whether our model can help to explain the differing conclusions reached in Coibion and Gorodnichenko (2015b) and Wong (2015). To this end, we generate counterfactual expected and realized inflation series holding the oil price constant after each of the oil market shocks between January 2009 and December 2011. Set in this manner, any deviation between the actual and counterfactual paths can be attributable to indirect effects of inflation pass-through during this period.

The results in Figure 6 reveal that if the real price of oil was held constant throughout 2009-11, then, on average, expected and realized would have respectively been around 1 percent lower than was observed. Given that observed expectations increased from 2.2 in January 2009 to 3.1 percent in December 2011, our results suggest that oil price fluctuations can account for all of the rise in household inflation expectations during this period. Thus, in line with Coibion and Gorodnichenko (2015b) we conclude that if firms’ inflation expectations track those of households, then the rise in their inflation expectations between 2009 and 2011 can explain the missing deflation during this period.³

Taken together, our results not only explain those in Coibion and Gorodnichenko (2015b), but also why Wong (2015) found results to the contrary. Specifically, the early millennium oil price surge of 2003-08 was due to an strong demand, primarily driven by unexpectedly high demand from Asia (Aastveit et al., 2015). This shock lead to an increase in inflation expectations which propagated the effects into higher actual inflation. If instead, oil prices were held constant, then our results suggest expected and actual inflation would have been 1 percent lower than observed. This explains the missing deflation result. Moreover, by averaging the effects of the underlying oil market shocks (see Section 3.2), Wong (2015) concludes that this mechanism is simply too weak to justify the results in Coibion and Gorodnichenko (2015b). Our results consequently highlight the importance of decomposing the underlying drivers of the oil market when studying

³While data limitations have made it impossible to examine the behavior of firm level inflation expectations in the US, Kumar et al. (2015) and Coibion et al. (2018b) report that inflation expectations of firms in New Zealand tend to co-move with gasoline prices to a similar extent of those in US households’ inflation expectations.




Figure 6: Great Recession Counterfactual: Actual time series (solid line) and counterfactuals (dashed line)

Notes: Actual percentage change in the price of oil, inflation expectations and month-on-month percentage change in inflation are in blue, while the counterfactuals are in red. The counterfactuals are made under the condition that the real price of oil remained flat since January 2009 using the method in Appendix B. The results presented are the median from a posterior distribution of 100,000 structural models.

the role of expectations in eliciting inflation pass-through of oil price shocks.

4 Conclusion

We have investigated whether the inflation pass-through of oil price shocks depends on the underlying drivers in the global market for crude oil. Methodologically, we extended a recently proposed structural vector autoregressive model of the global market for crude

oil to include an additional inflation block, which included monthly measures of US CPI inflation and the University of Michigan's survey of inflation expectations.

Using this new model, we confirmed existing evidence that expected and realized inflation are sensitive to oil prices, and provided the novel insight that the type of shock matters; with economic activity shocks having a significantly longer lasting effect than other types of real oil price shocks. Such shocks were able to explain almost all of the expected and realized inflation dynamics during the early part of the century, and also account for the reduction in inflation expectations since the oil price drop of 2014. This is not to say that the other shocks were not important. For instance, supply and consumption demand shocks were found to have contributed towards the high expected and realized inflation during the Iraqi invasion of Kuwait in 1990/91, and the latter shocks also played a key role in shaping households expected deflation during the 1997/98 Asian Financial Crisis.

In addition to these new insights, we resolved the debate on the role of oil price shocks in driving the missing deflation puzzle during the Great Recession. We found that the demand driven early millennium oil price surge of 2003-08 drove up inflation expectations and generated significant pass-through into actual inflation. If oil prices were instead held constant, expected and actual inflation would have been 1 percent lower than observed. Importantly, failing to distinguish between the different effects that demand and supply shocks have on actual and expected inflation would result in this effect being unidentified. This insight highlights the importance of decomposing the underlying drivers of the oil market when modeling the inflation pass-through of oil price shocks.

References

- Aastveit, K. A., Bjørnland, H. C., and Thorsrud, L. A. (2015). What drives oil prices? emerging versus developed economies. *Journal of Applied Econometrics*, 30(7).
- Ball, L. M. and Mazumder, S. (2011). Inflation dynamics and the great recession. Technical report, National Bureau of Economic Research.
- Baumeister, C. and Hamilton, J. D. (2015). Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica*, 83(5):1963–1999.
- Baumeister, C. and Hamilton, J. D. (2018). Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations. *Journal of Monetary Economics*, 100(C):48–65.
- Baumeister, C. and Hamilton, J. D. (2019a). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *American Economic Review*, 109(5):1873–1910.
- Baumeister, C. and Peersman, G. (2013a). The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics*, 28(7):1087–1109.
- Baumeister, C. and Peersman, G. (2013b). Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics*, 5(4):1–28.
- Baumeister, C. J. and Hamilton, J. D. (2019b). Drawing conclusions from structural vector autoregressions identified on the basis of sign restrictions. *Mimeo*.
- Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A., and Friedman, B. M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings papers on economic activity*, 1997(1):91–157.
- Blanchard, O. J. and Galí, J. (2007). The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s? In *International Dimensions of Monetary Policy*, NBER Chapters, pages 373–421. National Bureau of Economic Research, Inc.
- Blanchard, O. J. and Riggi, M. (2013). Why are the 2000s so different from the 1970s? A structural interpretation of changes in the macroeconomic effects of oil prices. *Journal of the European Economic Association*, 11(5):1032–1052.

- Caldara, D., Cavallo, M., and Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics*, 103:1 – 20.
- Coibion, O. and Gorodnichenko, Y. (2015a). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.
- Coibion, O. and Gorodnichenko, Y. (2015b). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.
- Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018a). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–1491.
- Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018b). How do firms form their expectations? new survey evidence. *American Economic Review*, 108(9):2671–2713.
- Galí, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44(2):195–222.
- Güntner, J. H. and Linsbauer, K. (2018). The effects of oil supply and demand shocks on us consumer sentiment. *Journal of Money, Credit and Banking*, 50(7):1617–1644.
- Hall, R. E. (2011). The long slump. *American Economic Review*, 101(2):431–69.
- Hamilton, J. D. (1985). Historical causes of postwar oil shocks and recessions. *The Energy Journal*, 6(1):97–116.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity*, page 215.
- Hamilton, J. D. and Herrera, A. M. (2004). Comment: oil shocks and aggregate macroeconomic behavior: the role of monetary policy. *Journal of Money, Credit and Banking*, pages 265–286.
- Harris, E. S., Kasman, B. C., Shapiro, M. D., and West, K. D. (2009). Oil and the macroeconomy: Lessons for monetary policy. In *US Monetary Policy Forum Report*, volume 23, page 2015.
- Hooker, M. A. (2002). Are oil shocks inflationary? asymmetric and nonlinear specifications versus changes in regime. *Journal of Money, Credit and Banking*, pages 540–561.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.

- Kilian, L. and Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: understanding the dynamics of oil market var models. *Journal of the European Economic Association*, 10(5):1166–1188.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Kumar, S., Afrouzi, H., Coibion, O., and Gorodnichenko, Y. (2015). Inflation targeting does not anchor inflation expectations: Evidence from firms in new zealand. *Brookings Papers on Economic Activity*, pages 151–208.
- Massey Jr, F. J. (1951). The kolmogorov-smirnov test for goodness of fit. *Journal of the American Statistical Association*, 46(253):68–78.
- Miller, L. H. (1956). Table of percentage points of kolmogorov statistics. *Journal of the American Statistical Association*, 51(273):111–121.
- Powell, P. D. (2011). Calculating determinants of block matrices. *arXiv preprint arXiv:1112.4379*.
- 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(8):1673–1689.

A Appendix: Proof

In this Appendix we formally prove our claims in Section 2.2 that

1. the determinant of $\tilde{\mathbf{A}}$, and
2. the equilibrium feedback effects in element (2,2) of the impact matrix $\tilde{\mathbf{A}}^{-1}$,

are equivalent to those in [Baumeister and Hamilton \(2019a\)](#), thereby allowing us to use the priors in their paper. To this end, we utilize the following well known theorem in linear algebra (see, e.g. [Powell \(2011\)](#) and references therein).

Theorem 1 *Let A be an $n \times n$ matrix that can be partitioned into*

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$$

where A_{11} is $l \times l$, A_{12} is $l \times k$, A_{21} is $k \times l$, A_{22} is $k \times k$, with $n = k + l$. Assume that A_{22} is nonsingular. Then,

1. The determinant of A is given by $\det(A_{11} - A_{12}A_{22}^{-1}A_{21}) \det(A_{22})$.
2. The matrix A is invertible if and only if the Shur complement $A_{11} - A_{12}A_{22}^{-1}A_{21}$ is invertible, and

$$A^{-1} = \begin{bmatrix} (A_{11} - A_{12}A_{22}^{-1}A_{21})^{-1} & - (A_{11} - A_{12}A_{22}^{-1}A_{21})^{-1} A_{12}A_{22}^{-1} \\ -A_{22}^{-1}A_{21} (A_{11} - A_{12}A_{22}^{-1}A_{21})^{-1} & A_{22}^{-1} \left(I_k + A_{21} (A_{11} - A_{12}A_{22}^{-1}A_{21})^{-1} A_{12}A_{22}^{-1} \right) \end{bmatrix}.$$

To see how $\tilde{\mathbf{A}}$ fits into Theorem 1, note that it can be partitioned as

$$\tilde{\mathbf{A}} = \begin{bmatrix} \tilde{\mathbf{A}}_{11} & \tilde{\mathbf{A}}_{12} \\ \tilde{\mathbf{A}}_{21} & \tilde{\mathbf{A}}_{22} \end{bmatrix}$$

in which

$$\tilde{\mathbf{A}}_{11} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 \end{bmatrix}, \tilde{\mathbf{A}}_{12} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix},$$

$$\tilde{\mathbf{A}}_{21} = \begin{bmatrix} 0 & 0 & -\lambda_{\pi^e p} & 0 \\ 0 & 0 & -\phi_{\pi p} & 0 \end{bmatrix}, \tilde{\mathbf{A}}_{22} = \begin{bmatrix} 1 & 0 \\ -\phi_{\pi \pi^e} & 1 \end{bmatrix}.$$

It is easy to see that the sizes of these matrices sums to the dimensions of $\tilde{\mathbf{A}}$ and that $\det(\tilde{\mathbf{A}}_{22}) = 1$.

To prove our first claim, note that applying part 1 of Theorem 1 gives

$$\begin{aligned}\det(\tilde{\mathbf{A}}) &= \det(\tilde{\mathbf{A}}_{11} - \tilde{\mathbf{A}}_{12}\tilde{\mathbf{A}}_{22}^{-1}\tilde{\mathbf{A}}_{21}) \det(\tilde{\mathbf{A}}_{22}), \\ &= \det(\tilde{\mathbf{A}}_{11}).\end{aligned}$$

Since $\tilde{\mathbf{A}}_{11}^{-1}$, is equivalent to the impact matrix in Equation 35 of [Baumeister and Hamilton \(2019a, p.1889\)](#), this proves that the determinants of the two matrices are equivalent.

Next, to prove our second claim, note that the Shur complement is simply $\tilde{\mathbf{A}}_{11}$ which is invertible (with inverse given in Equation 40 of ([Baumeister and Hamilton, 2019a, p.1894](#))). Thus, it follows by part 2 of Theorem 1 that

$$\tilde{\mathbf{A}}^{-1} = \begin{bmatrix} \tilde{\mathbf{A}}_{11}^{-1} & -\tilde{\mathbf{A}}_{12} \\ \tilde{\mathbf{A}}_{21}\tilde{\mathbf{A}}_{11}^{-1} & \mathbf{I}_k \end{bmatrix}.$$

Since the upper 4×4 block of the impact matrix $\tilde{\mathbf{A}}^{-1}$ is $\tilde{\mathbf{A}}_{11}^{-1}$, it follows that the equilibrium feedback effects in element (2,2) are also equivalent.

B Appendix: Counterfactual

In this section we provide methodological details on the counterfactual exercise used in Section 3.3. To this end, first note that the reduced form VAR model variant of the SVAR in (8) is given by

$$\begin{aligned} \mathbf{y}_t &= \Phi \mathbf{x}_{t-1} + \boldsymbol{\epsilon}_t, \\ &= \mathbf{c} + \sum_{i=1}^{12} \Phi_i \mathbf{y}_{t-i} + \boldsymbol{\epsilon}_t, \end{aligned} \quad (10)$$

where $\Phi = \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{B}}$ and $\boldsymbol{\epsilon}_t = \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{u}}_t$. The companion form of (10) is then given by

$$\mathbf{Z}_t = \mathbf{F} \mathbf{Z}_{t-1} + \mathbf{v}_t, \quad (11)$$

in which

$$\mathbf{Z}_t = \begin{bmatrix} \mathbf{y}_t - \mathbf{c} \\ \vdots \\ \mathbf{y}_{t-11} - \mathbf{c} \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \Phi_1 & \Phi_2 & \dots & \Phi_{11} & \Phi_{12} \\ \mathbf{I}_n & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_n & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \dots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{I}_n & \mathbf{0} \end{bmatrix}, \quad \mathbf{v}_t = \begin{bmatrix} \boldsymbol{\epsilon}_t \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}.$$

Impulse response functions are then obtained by a straight forward application of the multivariable chain rule, i.e.

$$\frac{\partial \mathbf{Z}_{t+s}}{\partial \mathbf{u}'_t} = \frac{\partial \mathbf{Z}_{t+s}}{\partial \boldsymbol{\epsilon}'_t} \frac{\partial \boldsymbol{\epsilon}_t}{\partial \tilde{\mathbf{u}}'_t} = \Psi_s \tilde{\mathbf{A}}^{-1} = \Theta_s,$$

where $\Psi_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \boldsymbol{\epsilon}'_t}$ is the $(n \times n)$ non-orthogonalized impulse response function matrix at horizon s which is obtained with the first n rows of \mathbf{F}^s .⁴

Following a shock in the global market for crude oil, the counterfactual impulse response function uses a sequence of structural inflation expectations shocks to hold inflation expectations constant at each impulse horizon s . To ensure that the effects are comparable, we assume that each shock increases the real price of oil by 10 percent.⁵ To see how this sequence is constructed, define the selection vector \mathbf{e}_k to be an $n \times 1$ vector with 1 as the k -th element and 0 otherwise. Then, the impulse response function of variable i to a shock in variable j horizon s , denoted $\theta_{ij,s}$, is

$$\theta_{ij,s} = \mathbf{e}'_i \mathbf{F}^s \boldsymbol{\kappa}_j,$$

⁴Note that we have simplified notation by writing the impulse response function with respect to the correlated shock vector, \mathbf{u}_t , as opposed to the uncorrelated shocks \mathbf{u}_t^* . To get uncorrelated shocks we follow the procedure in Equation 39 of Baumeister and Hamilton (2019a, p.1890)

⁵The choice of a 10 percent price increase is in line with Wong (2015).

where $\boldsymbol{\kappa}_j = \frac{10\tilde{\mathbf{A}}^{-1}\mathbf{e}_j}{\mathbf{e}'_3\tilde{\mathbf{A}}^{-1}\mathbf{e}_j}$, and \mathbf{e}'_3 in the denominator occurs because the real price of oil is the third variable in the SVAR. Since inflation expectations is the fifth variable in the SVAR, the counterfactual exercise sets $\boldsymbol{\theta}_{5j,s} = 0$ for all s . For instance, if $s = 0$ then we require that

$$\hat{\boldsymbol{\theta}}_{5j,0} = \mathbf{e}'_5\boldsymbol{\kappa}_j + \mathbf{e}'_5\tilde{\mathbf{A}}^{-1}\mathbf{e}_5\hat{u}_0^{\pi^e} = 0,$$

where the simulated shock $\hat{u}_0^{\pi^e} = -\frac{\mathbf{e}'_5\boldsymbol{\kappa}_j}{\mathbf{e}'_5\tilde{\mathbf{A}}^{-1}\mathbf{e}_5}$. Letting S denote the maximum impulse horizon, the remaining sequence is then calculated recursively through

$$\hat{u}_s^{\pi^e} = \frac{\boldsymbol{\theta}_{5j,s} + \sum_{m=0}^{S-1} \mathbf{e}'_5\mathbf{F}^m\tilde{\mathbf{A}}^{-1}\mathbf{e}_5\hat{u}_m^{\pi^e}}{\mathbf{e}'_5\tilde{\mathbf{A}}^{-1}\mathbf{e}_5}.$$

Thus, the counterfactual impulse response function for variable i to shock j at horizon s , $\hat{\boldsymbol{\theta}}_{ij,s}$, is computed as

$$\hat{\boldsymbol{\theta}}_{ij,s} = \boldsymbol{\theta}_{ij,s} + \sum_{m=0}^{S-1} \mathbf{e}'_i\mathbf{F}^m\tilde{\mathbf{A}}^{-1}\mathbf{e}_5\hat{u}_m^{\pi^e}.$$

C Appendix: Additional Results

Full historical decompositions

In this Appendix we provide the full set of historical decompositions associated with Figures 3 and 4 in Section 3.2 of the paper.

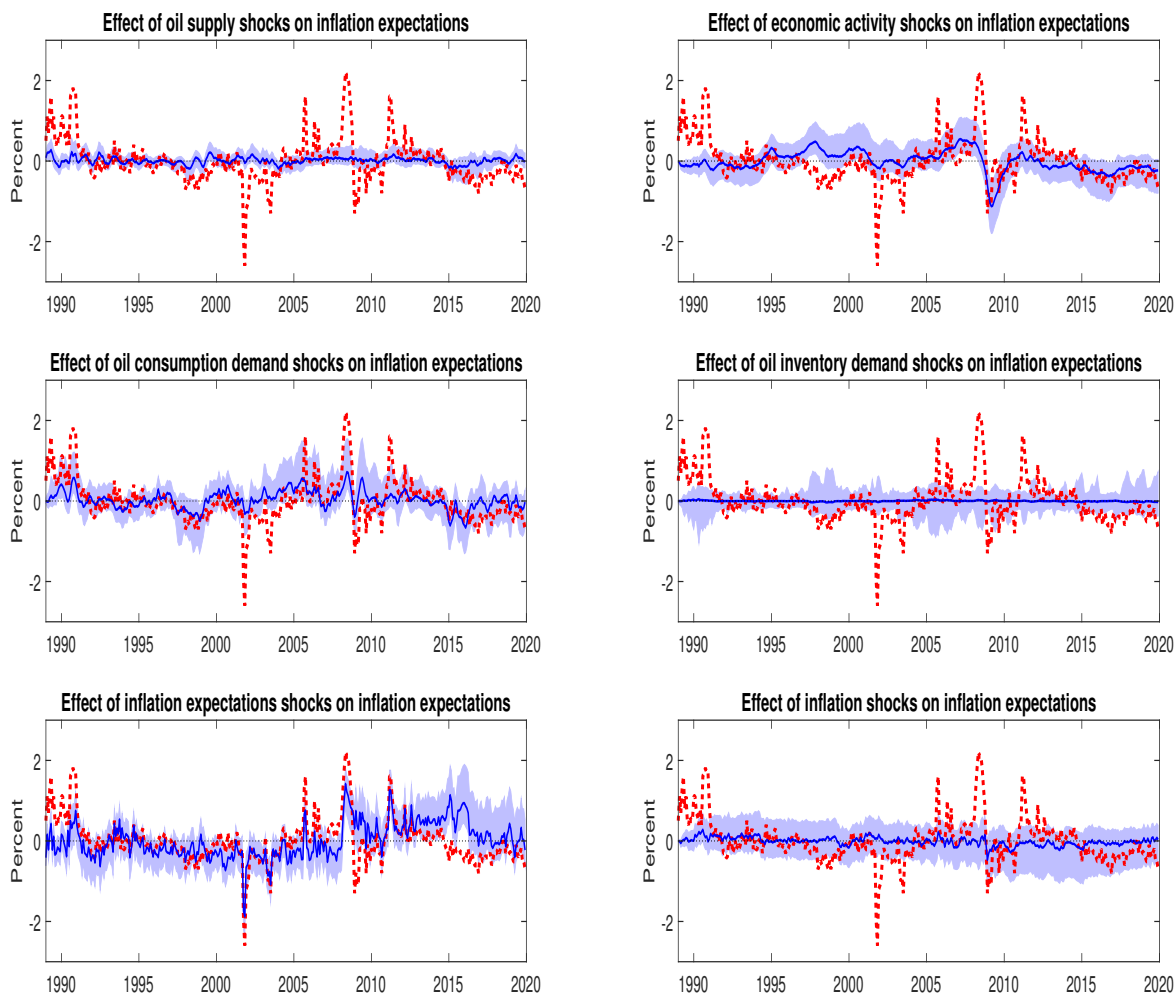


Figure 7: Historical Decomposition: Inflation Expectations

Notes: Actual changes in inflation expectations (red dotted lines) and median estimate of historical contribution of separate structural shocks (blue lines). Blue shaded regions indicate 95 percent posterior credibility regions from the posterior distribution of 100,000 structural models.

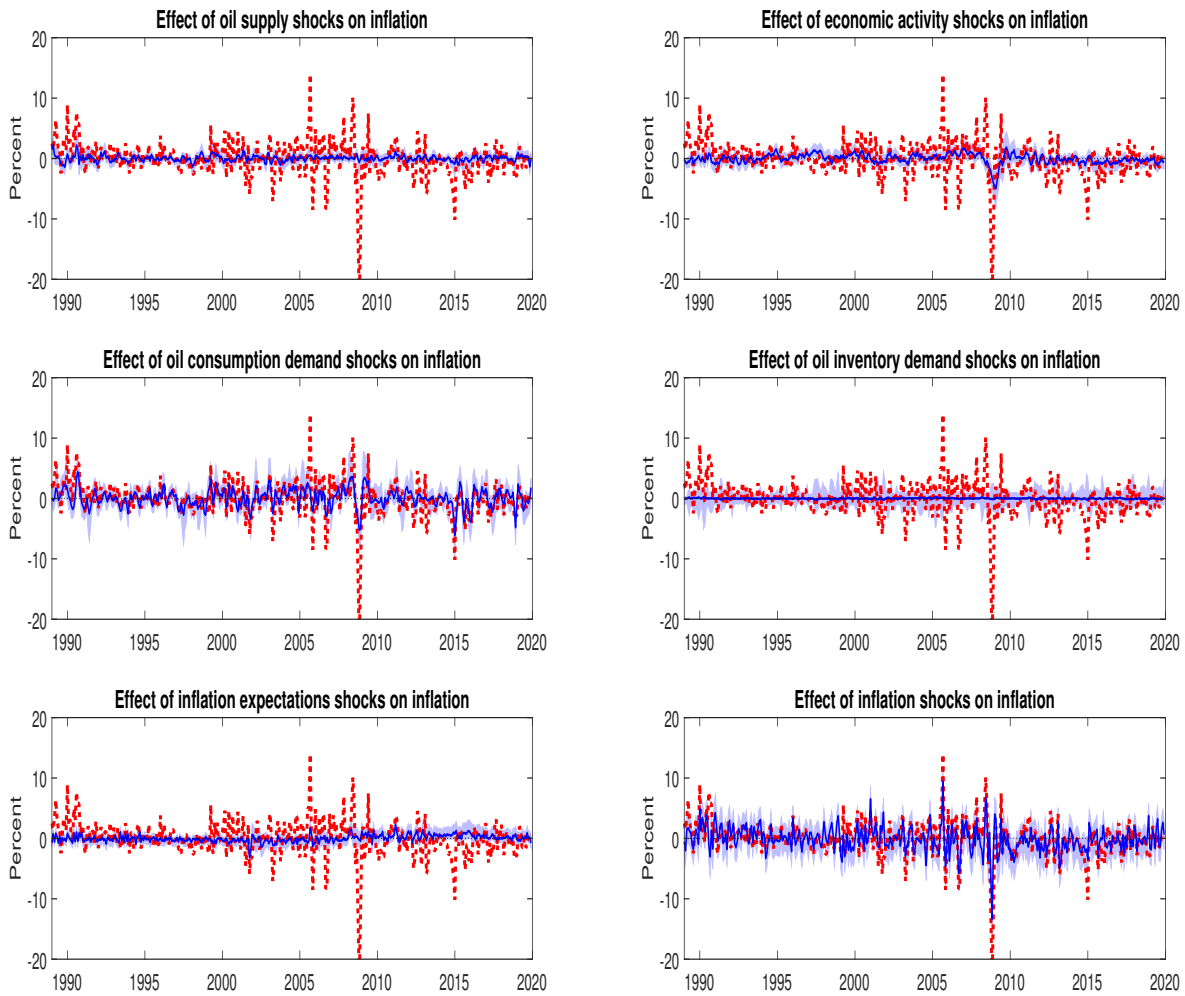


Figure 8: Historical Decomposition: Inflation

Notes: Actual changes in inflation (red dotted lines) and median estimate of historical contribution of separate structural shocks (blue lines). Blue shaded regions indicate 95 percent posterior credibility regions from the posterior distribution of 100,000 structural models.

Counterfactual IRFs

Figure 9 graphs the counterfactual structural impulse response functions with a 95% joint credible set associated with results in Section 3.3 of the paper. As discussed there, the main observation is that the counterfactual IRFs contain quite a bit more estimation uncertainty surrounding relative to the actual IRFs presented in Figure 2.

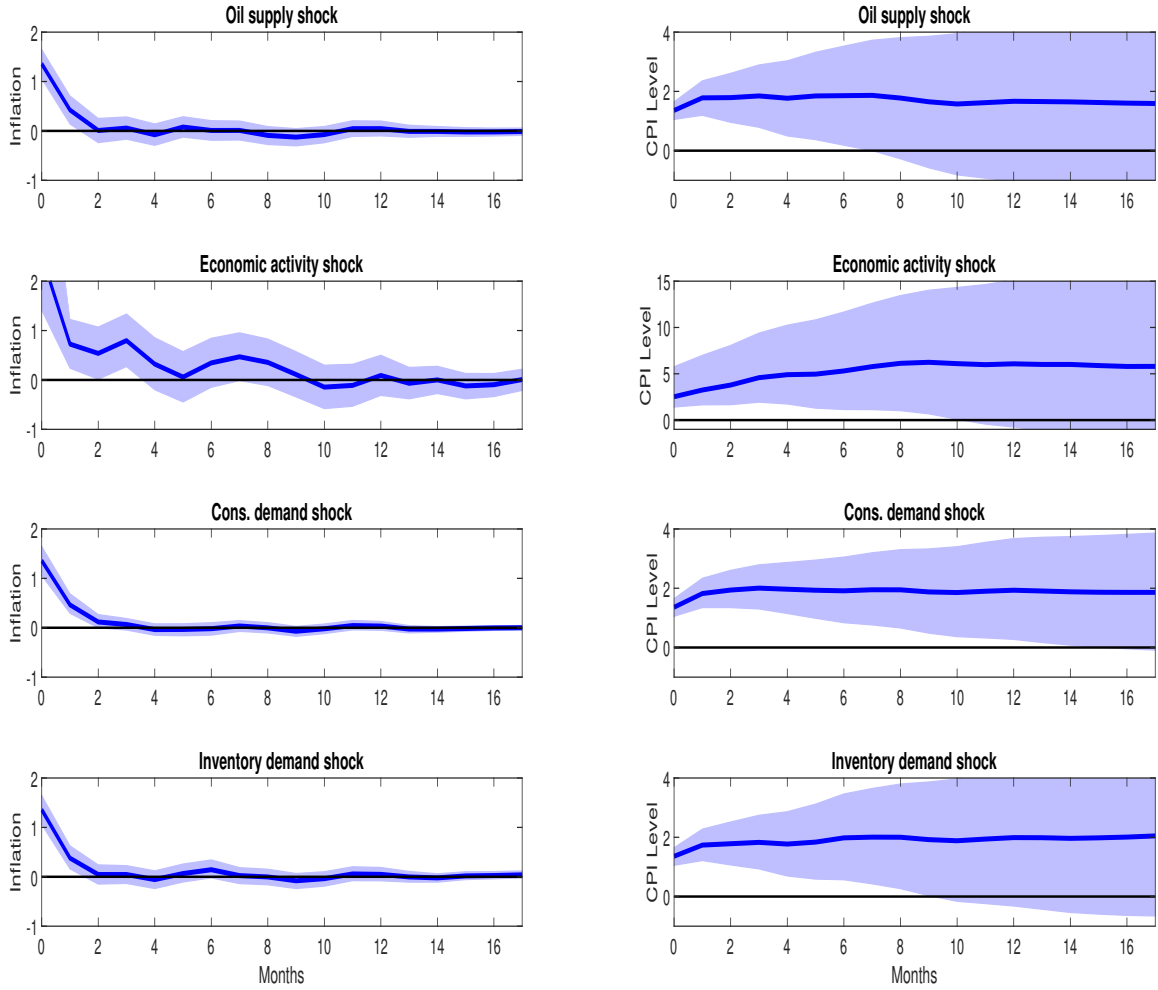


Figure 9: Counterfactual structural impulse response functions

Notes: The structural shocks have been normalized to elicit a 10% increase in the real price of oil. The posterior median is shown in boldface and the shaded area is the 95% joint credible set obtained from the posterior distribution of 100,000 structural models.

Oil market block

We here examine the impulse response functions and historical decompositions associated with the global oil market block of our SVAR model, and compare them to those in [Baumeister and Hamilton \(2019a\)](#). This has the effect of determining whether the use of a shorter sample size and additional inflation block, changes the main results in

Baumeister and Hamilton (2019a).⁶

The posterior distributions of the contemporaneous coefficients in the oil block are shown in Figure 1. The solid red curves represent the prior information and the blue histogram is the posterior distribution. Each of the posterior distributions exhibits similar characteristics to those in Baumeister and Hamilton (2019a). One exception is that the posterior median of the short-run price elasticity of oil supply, α_{qp} , is 0.04, which is smaller than the 0.15 estimated in Baumeister and Hamilton (2019a) and about half of the 0.11 value obtained in Caldara et al. (2019). In contrast, our estimate of the posterior median of the short-run price elasticity of oil demand, β_{qp} , is -0.38, which is quite close to the value of -0.35 obtained in Baumeister and Hamilton (2019a). Why is the supply elasticity is lower, while the demand elasticity is almost the same? The answer is that our shorter sample excludes periods in which supply shocks played a key role and economic activity shocks didn't matter. Since our sample starts in 1983, as compared to 1958, we have omitted the pre-1973 period in which the Texas Railroad Commission (TRC) controlled the price of petroleum (Hamilton, 1985), thus forcing oil supply shocks to have a smaller role.

The impulse response functions in Figure 11 are mostly identical to those in Baumeister and Hamilton (2019a, p.1900). For instance, an unanticipated increase in consumption demand for oil (see row 3) has a negligible impact on oil production and above-ground inventories, while slightly increasing world industrial production and pushing up the real oil price. That being said, there are two notable differences that occur from using a smaller sample. First, we find that the economic activity shock have a relatively larger impact on the real price of oil. Second, we find that supply shocks have a relatively smaller impact on world industrial production and the real price of oil. This finding is in line with our above estimates of the supply elasticity. Moreover, given that demand due to economic activity played a much larger role in shaping dynamics of the global market for crude oil since the turn of the century, shortening the sample increases the average effect.

⁶Recall that our sample size is 1983M1-2019M12, compared to 1958M1-2016M12 in Baumeister and Hamilton (2019a). We also re-estimated their SVAR model on our sample with their replication files. The results (which are provided in Appendix 2) are identical to those presented here.




Figure 10: Prior (solid red curves) and posterior (blue histograms) distributions of structural parameters in the oil block




Figure 11: Structural impulse response functions: oil market variables

Notes: The impulse response functions are in response to a one standard deviation structural shock. The posterior median is shown in boldface and the shaded area is the 95% joint credible set obtained from the posterior distribution of 100,000 structural models.

The historical decompositions in Figure 12, are identical to those in [Baumeister and Hamilton \(2019a, p.1901\)](#). For instance, we observe that demand and supply shocks to were equally important during the First Persian Gulf War of 1990/91, while consumption and economic activity explain much of the early millennium oil price surge of 2003-2008, along with the subsequent collapse and rebound during the Great Recession. The consumption demand shock also accounts for much of the oil price collapse in 2014-2016, and subsequent recovery in 2016. We highlight that this result is identical, while the IRFs differ, because IRFs are sample averages, while historical decompositions analyze the effects of shocks at each point in time.

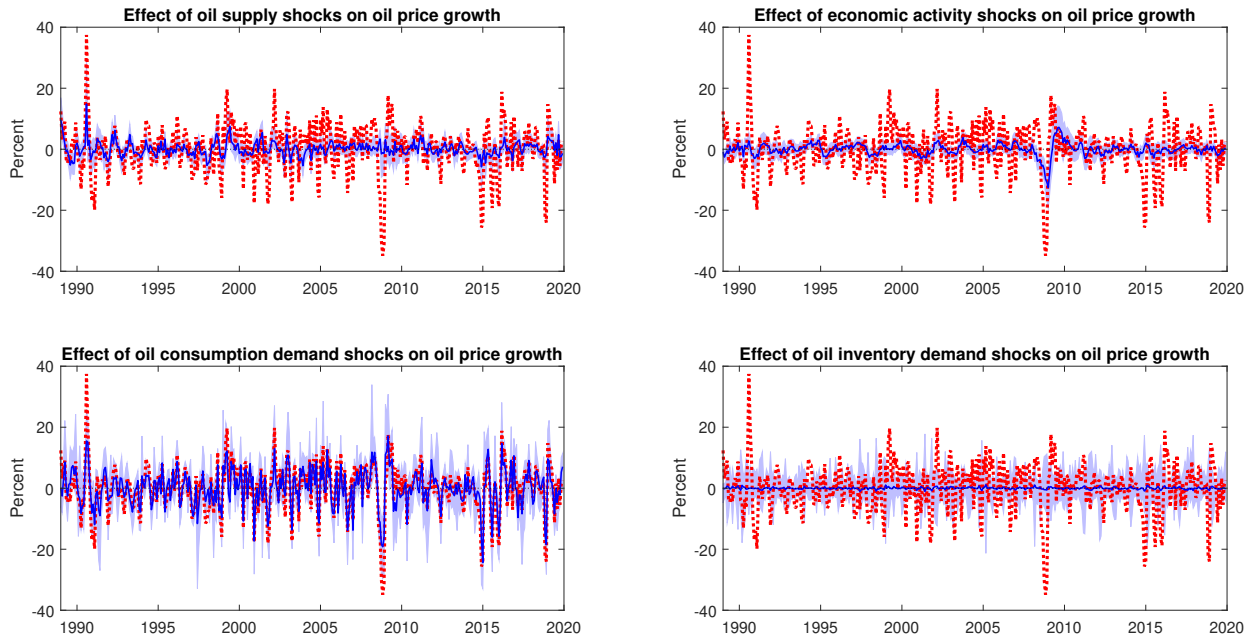


Figure 12: Structural historical decompositions: oil market variables

Notes: Actual changes in the real price of oil (red dotted lines) and median estimate of historical contribution of separate structural shocks (blue lines). Blue shaded regions indicate 95 percent posterior credibility regions obtained from the posterior distribution of 100,000 structural models.