The Time-Varying Effects of Oil News Shocks*

Ana María Herrera† Sandeep Kumar Rangaraju‡

Abstract

We investigate whether the transmission of oil news shocks to inflation and inflation expectations has changed over time. We employ a Bayesian time-varying parameter vector autoregressive model where we identify the structural shock of interest -an oil price shock- via a proxy variable. This framework allows us to explore what parameters drive the time variation and to inquire whether the reliability of the proxy is sensitive to alternative priors. We find evidence that there has been a considerable degree of time variation in the response of real oil prices, world economic activity, inflation, and, especially, inflation expectations to oil news shocks. Moreover, we trace back the sources of the time-varying effect of oil shock to fluctuations in the stochastic volatility of all variables–especially oil prices and inflation–and to time variation of the correlation between oil prices and inflation expectations. Our results suggest that, in times of higher and more volatile inflation, oil news shocks matter for the pass-through of shocks to realized inflation and inflation expectations.

JEL Classification: C32, E32
Keywords: Oil prices, TVP-VAR, inflation, inflation expectations.

*We thank Elena Pesavento and conference participants at the 2025 Society for Economic Measurement(SEM), the 2024 Midwest Econometrics Group Conference(MEG), and the 2025 SNDE Symposium for helpful suggestions and comments.

†Department of Economics, Gatton College of Business and Economics, University of Kentucky, Lexington 40206-0034; phone: (859) 257-1119; e-mail: amherrera@uky.edu

‡Department of Economics, Goddard School of Business & Economics, Weber State University, Ogden 84408-3801; e-mail: srangaraju@weber.edu

1 Introduction

Turmoil in the oil market, whether sparked by wars connected to oil interests, the post-COVID-19 pandemic, or uncertainty regarding future oil demand related to trade wars, can result in considerable fluctuations in oil prices. In turn, oil price increases, such as those observed after the recovery from the COVID-19 pandemic and Russia's invasion of Ukraine, can exert upward pressure on inflation and inflation expectations. Moreover, recent remarks about "the possibility that inflation could be more volatile going forward than during the inter-crisis period of the 2010s" expressed by the Chairman of the Federal Reserve (Powell, 2025), in conjunction with a rise in inflation expectations and looming concerns about slower economic growth, underscores the belief among policymakers and the press that the US economy is facing higher inflation risk and, perhaps, the possibility of stagflation.

In this context, how changes in expectations about future oil prices affect US inflation and inflation expectations continues to be an important topic of research. The possibility that a period of stagflation may recur raises the following question: Has the effect of oil price shocks on inflation changed over time? A priori, there are many reasons to think that this is the case. The fracking revolution has led to changes in the US oil industry: shale oil producers are more nimble than conventional oil producers and may respond faster to uncertainty about global economic activity, positive spillover effects linked to increased shale oil production appear to have altered the transmission of oil price shocks to labor markets (Bjørnland and Skretting (2018)), and reliance on imported oil has declined, with the US becoming a net exporter of petroleum in 2020. The US economy has also experienced profound structural changes in recent decades, with the rise of services and the transformation of the manufacturing sector. In addition, monetary policy and the tools employed by the Federal Reserve have changed significantly since the late 1970s (especially since the Great Recession), and the degree of anchoring of inflation expectations has changed over time (see Naggert *et al.* (2023)).

In this paper, we investigate whether the effect of oil news shocks on inflation and inflation expectations has changed over time. To do so, we employ a time-varying parameter structural vector autoregressive model where we identify the structural shock of interest -an oil price shock- via a proxy variable. To achieve identification, we follow Känzig (2021) in using OPEC announcements to construct an external instrument for oil news shocks. We depart from his setup in two ways. First, rather than imposing that the parameters in the Proxy-VAR are constant over time or imposing an ad hoc structural break (see Blanchard and Galí (2010) and Herrera and Pesavento (2009), and Ramey and Vine (2011)), we allow the parameters to evolve continuously over time. Second, we estimate the model using a Bayesian framework, which allows us to explore what parameters drive the time variation and inquire whether the reliability of the proxy is sensitive to alternative priors. This approach enables us not only to examine possible time variation in the responses of real oil prices, oil production, world economic activity, inflation, and inflation expectations to oil news shocks, but also in the contribution of these shocks to the forecast error variance of inflation and inflation expectations.

Four key findings are derived from our study. First, we uncover a significant degree of time variation in the response of the oil market variables to oil news shows. Specifically, while the response of oil prices has become more persistent over the years, the effect of the shock on global economic activity has declined. Second, the inflation impact response increased from the beginning of the sample until the Great Recession, exhibited a declining trend until the COVID-19 pandemic, and then increased in the last years of the sample. Similarly, inflation expectations have responded more to oil news shocks since the COVID-19 pandemic. Moreover, we find that the contribution of oil news shocks to the forecast error variance decomposition of inflation expectations has increased markedly since the pandemic. Finally, we trace back the sources of the changing effect of oil news shocks to fluctuations in the stochastic volatility of all variables –especially oil prices and inflation– and to time variation in the correlation between oil prices and inflation expectations.

This paper is related to two broad strands of literature. On the one hand, we contribute to the literature that studies the effects of oil price shocks on the macroeconomy. In particular, we expand on the work by Känzig (2021) who, borrowing from the monetary policy literature where an instrumental variable is constructed using unexpected changes in daily futures around FOMC announcements, constructs an instrument for oil supply shocks using surprises in daily oil price futures around OPEC announcements. He finds that news of future oil price increases leads to higher inflation expectations and a significant increase in CPI inflation. He estimates that news shocks that lead to a contemporaneous 10% increase in oil prices result in an immediate increase in inflation expectations of about a tenth of a percentage point and a contemporaneous increase in CPI inflation of about twice this magnitude. We depart from his work in two aspects. Känzig (2021) estimates a constant parameter Proxy-VAR and posits that the identified shocks represent oil supply disturbances. However, recent work by Degasperi (2021) and Kilian (2023) calls into question the interpretation of the shocks and suggests that such a model identifies a combination of storage and flow demand shocks. Our estimates support that view; hence, we refer the shock as an *oil news shock*. Because our purpose is not to separately identify supply- and demand-driven shocks, but to investigate whether the parameters of a Känzig-style model are time invariant, we use a Metropolis-within-Gibbs sampling algorithm (Mumtaz and Petrova, 2023) to approximate the posterior distributions of the Proxy TVP-VAR. To the best of our knowledge, ours is the first application of the estimation approach proposed by Mumtaz and Petrova (2023) in the context of studying the effect of oil news shocks on oil price fluctuations and inflation.

On the other hand, our paper is related to a vast literature that examines the time-varying effects of shocks on inflation using Bayesian TVP-VAR models. The TVP-VAR framework has been employed to investigate topics of interest to policy makers, such as the transmission of monetary policy (e.g., Cogley and Sargent (2005), Primiceri (2005), and Koop *et al.* (2009)), the effects of fiscal shocks (e.g., Klein and Linnemann (2020), Mumtaz and Petrova (2023)), and the changing impact of oil price shocks (e.g., Baumeister and Peersman (2013b), Baumeister and Peersman (2013a), Chang *et al.* (2023)). Our study contributes to this literature by inquiring whether the effect of oil news shocks on the oil market, inflation, and inflation expectation has varied over time and then investigating which parameters account for the time variation.

The paper is organized as follows. Section 2 outlines the data and Section 3 discusses the econometric framework, including the identification of oil news shocks, prior starting values, and reliability of the instrumental variable. Estimation results for the impact of oil news shocks on the oil market, inflation, and inflation expectations are reported in Section 4. Section 5 explores the sources of time variation. Section 6 presents several robustness checks before we conclude in Section 7.

2 Data

Our study employs five monthly time series that span from January 1978 to December 2022, encompassing three global oil market variables, US inflation expectations, and the inflation rate of the Consumer Price Index (CPI). Global oil production data, measured in thousands of barrels per day, is sourced from the Monthly Energy Review published by the Energy Information Administration (EIA). Global economic activity is quantified using the OECD + 6 industrial production index developed by Baumeister and Hamilton (2019) and employed by Känzig (2021) in the baseline model. The real price of oil is calculated by the WTI spot crude oil price deflated by the US Consumer Price Index (all items) obtained from the Federal Reserve Economic Data (FRED) database.

For empirical estimation purposes, global oil production, global real economic activity, and real oil price are incorporated as logarithmic growth rates. U.S inflation is computed as the annualized monthly rate of change in the US Consumer Price Index, retrieved from FRED. To capture forward-looking inflation dynamics, inflation expectations are measured as the median one-year-ahead inflation expectations obtained from the Michigan Survey of Consumer Inflation Expectations, which is widely recognized among academics and policy makers as a reliable measure of consumer sentiment regarding future prices.

As mentioned above, we identify the oil news shock using an instrumental variable. The construction of the instrument, z_t, follows Känzig (2021). Specifically, we gathered OPEC announcements from press releases spanning the period between 1983 and 2022. This results in a total of 150 announcements. Since 2002, press releases have been available on the official OPEC website.[1] Prior to 2002, we collected the dates of the announcements from official OPEC resolutions and press releases OPEC (1990), supplemented with information from Bloomberg's news service. To construct the oil news shock instrument, we calculate the change in the log settlement price of the WTI futures contract h months ahead for $h = 1, ..., 12$. Subsequently, the surprise series is computed as the first principal component of these changes. As is common in the literature, we assign each surprise to the month in which the corresponding OPEC meeting took place. If no meeting occurred in a month, the value of the surprise is set to zero. Figure 1 illustrates the evolution of the proxy and identifies notable historical episodes of oil market disruptions. Particularly relevant for our study are the large changes in oil price expectations around the OPEC meetings during the COVID-19 pandemic. Note how in March 2020, oil futures prices plunged to levels not seen since 1991 after OPEC's announced a cut in production as the COVID-19 outbreak led OPEC to revise downward its forecast of global oil demand.

3 The Time-Varying Proxy SVAR

3.1 Econometric Model

We consider the joint behavior of global oil production, world industrial production, the real price of oil, and inflation to be given by a Gaussian $VAR(p)$ model with time-varying parameters and stochastic volatility given as follows:

$$y_t = c_t + B_{1,t}y_{t-1} + B_{2,t}y_{t-2} + ... + B_{p,t}y_{t-p} + u_t = B_tX_t + u_t \tag{1}$$

[1] Collected from the OPEC press release archives.

where $y_t = \begin{bmatrix} \Delta prod_t & rea_t & rpo_t & \pi_t \end{bmatrix}'$, $X_t = \begin{bmatrix} y_{t-1} & y_{t-2} & ... & y_{t-p} \end{bmatrix}'$ is the 4×1 vector of endogenous variables, c_t is a 4×1 vector of time-varying intercepts.

The time-varying covariance matrix of the reduced form residuals u_t is given by:

$$\Sigma_t = (A_t q)(A_t q)'. \tag{2}$$

We define A_t as a lower triangular matrix with time-varying parameters and q as an element of orthogonal matrices of size 4 such that $q'q = I_4$. Note that the heteroskedastic reduced-form residuals, u_t, are a linear combination of the structural residuals, ε_t, given by $u_t = A_t q \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, I_4)$ and u_t has variance covariance matrix Σ_t. The 4×4 matrices of time-varying autoregressive parameters are denoted by $B_{i,t}$.

In the above time-varying model, we implement a variable substitution by rotating the fourth variable, i.e., π_t (inflation), with π_t^{ex} (inflation expectations). In other words, we first estimate the model with the three oil market variables and inflation and then reestimate the model replacing inflation with inflation expectations.[2] Estimation results reported in the online Appendix show that the estimated impulse response functions for the oil market variables are indistinguishable in these two model specifications.

As is common in the literature (Primiceri, 2005), we assume that the autoregressive parameters $b_t = \text{vec}(B_t')$ follow a random walk process given by:

$$b_t = b_{t-1} + Q_b^{1/2} \eta_t^b, \ \eta_t^b \sim \mathcal{N}(0, I_{4,4p+1}) \tag{3}$$

[2]This rotation is implemented primarily for computational efficiency. Time-varying parameter models are computationally intensive and require significantly longer estimation times compared to constant coefficient models due to the substantially larger number of parameters that must be estimated. Each time period introduces a new set of coefficients, exponentially increasing the dimensionality of the estimation problem. By substituting inflation expectations for inflation in this specific position of the variable ordering and reestimating the model, we reduce both the computational burden and the number of parameters to be estimated, while still maintaining the essential dynamics of the system.

and consider the decomposition of A_t

$$A_t = \tilde{A}_t H_t^{1/2} \tag{4}$$

where $\tilde{A}_t$ is a lower triangular matrix with ones along the main diagonal and H_t is a diagonal matrix. Let α_t, be the vector stacking –by rows– the non-zero and non-one elements of $\tilde{A}_t$ and h_t be the 4×1 vector stacking the diagonal elements of H_t. Then, we assume that the data generating processes for α_t and h_t are given by:

$$\alpha_t = \alpha_{t-1} + Q_\alpha^{1/2} \eta_t^\alpha, \ \eta_t^\alpha \sim \mathcal{N}(0, I_6) \tag{5}$$

and

$$lnh_t = lnh_{t-1} + Q_h^{1/2} \eta_t^h, \ \eta_t^h \sim \mathcal{N}(0, I_4). \tag{6}$$

Hence, as is common in the literature, we assume that the non-zero and non-one elements of $\tilde{A}_t$ follow independent random walks, the diagonal elements of H_t evolve as geometric random walks, and all the elements are independent of each other. This specification allows for time variation in both the contemporaneous relationships among variables and the variance of structural shocks, capturing potential changes in the transmission mechanism of oil news shocks over time.

3.2 Identification of the Oil News Shocks

Recall that we are interested in identifying only the shock to oil prices. Let the structural shocks of the TVP-VAR model, ε_t, be given by:

$$\varepsilon_t = A_{0,t}^{-1} u_t \tag{7}$$

where $A_{0,t} = A_t q = \tilde{A}_t H_t^{1/2} q$. To identify the effect of oil news shocks ε_{1t} –without necessarily distinguishing their specific source–, we employ a single external instrument that correlates with it. Let q_1 denote the first column of q, collect the remaining structural shocks in a $(n-1)\times 1$ vector $\varepsilon_{st} = [\varepsilon_{2t}, \varepsilon_{3t}, \varepsilon_{4t}]$, and let z_t denote the instrumental variable (proxy). We link the proxy to the structural shock of interest, ε_{1t}, via the following equation

$$z_t = \beta\varepsilon_{1t} + \sigma v_t \tag{8}$$

where v_t is assumed to be distributed independently and identically over time as $\mathcal{N}(0,1)$ and $E(v_t\varepsilon_t) = 0$.

Then, identification requires that the oil surprises employed as an instrument satisfy two key conditions:

- **Relevance**: $E[z_t\varepsilon_{1t}] = \beta \neq 0$
- **Exogeneity**: $E[z_t\varepsilon_{qt}] = 0$

The above conditions imply that the proxy, z_t, carries information about the oil news shock ε_{1t}, but is not driven by any other shock, ε_{st}.

To jointly model the interaction between the TVP-VAR and the proxy, we adopt the methodology proposed by Caldara and Herbst (2019) and adapted to the time-varying setup by Mumtaz and Petrova (2023). That is, the likelihood of the TVP-VAR is augmented with a measurement equation that relates the proxy to the structural shock of interest, and the model is estimated using Bayesian techniques that involve a Metropolis-with-Gibbs sampling algorithm. The advantages of such a methodology are twofold. First, as noted by Caldara and Herbst (2019) the estimation procedure incorporates all the sources of uncertainty, and hence the proxy becomes informative about the reduced- and structural-form parameters of the model. Second, by calculating the reliability statistic, we can empirically evaluate the relevance of the prior. Specifically, we calculate the reliability statistic proposed by Mertens and Ravn (2013), which is defined as follows:

$$\rho^2 = \frac{\beta^2}{\beta^2 + \sigma^2}. \tag{9}$$

As equation (9) illustrates, the reliability of the proxy depends on the signal-to-noise ratio, β/σ. The higher the correlation between the IV, z_t, and the shock of interest, ε_{1t}, the more informative the proxy is for identifying the shock.

We conclude this section by noting that, as evident from equation (8), we model the relationship between the proxy and the oil news shock via a constant coefficient model. We opt for this modeling choice to focus our attention on the time-varying transmission mechanism of the oil news shock, while imposing the prior that the relationship between the IV and the shock does not change over time. While we realize such an assumption is restrictive, it allows us to force the time variation to stem from the TVP-VAR parameters while assuming that the relevance of the instruments remains unchanged. We leave the study of changes in the relevance of the prior for future work.

3.3 Prior Distributions, Starting Values and Estimation Strategy

The selection of the priors follows the standard Bayesian TVP-VAR literature and, in particular, the work by Mumtaz and Petrova (2023) for the implementation of the identification via an external instrument.[3] While we provide a detailed description of all priors in Table 1, in this section we briefly discuss the prior distribution and starting values for some of the parameters.

The priors for b_t, α_t, h_t, and Q_b follow Primiceri (2005), Cogley and Sargent (2005), and Mumtaz and Petrova (2023). The priors for the initial states for the coefficients (b_t), the covariances (α_t), and $ln(h_t)$ are assumed to be normally distributed and independent of each other. The priors for b_t and Q_b are calibrated using the OLS estimates of the VAR model on a training sample of $T_0 = 120$ observations. Let the ordinary least squares(OLS) estimate of the VAR autoregressive parameters be denoted by B_{OLS} and the estimate of the variance-covariance matrix denoted by V_{OLS}. Then, the prior for Q_b is an inverse Wishart

[3] We are grateful to Haroon Mumtaz and Katerina Petrova for making their MATLAB code for the Proxy TVP-VAR estimation publicly available at `https://sites.google.com/site/hmumtaz77/research-papers`.

$IW(T_0 \times V_{0LS} \times 3.5 \times 10^{-4}, T_0)$. The initial states for b_t are modeled as $b_{0|0} \sim \mathcal{N}(B_{OLS}, V_{OLS})$.

As in Mumtaz and Petrova (2023) and Caldara and Herbst (2019), the prior for Σ_t is uniform and is taken from the QR decomposition of a standard normal distribution. The starting values for the parameters of Equation (8), are also calibrated using OLS on a training sample T_0. We employ the standard inverse-gamma distribution for σ^2 (i.e. $p(\sigma^2)$ is an inverse gamma density $IG(a, b)$, reparametrized in terms of the mean σ_0 and v_0). We evaluate the robustness of the results to different priors for the parameters of the IV equation in Section 6.

For details regarding the estimation and specifics of the Metropolis-within-Gibss algorithm, we refer the reader to Mumtaz and Petrova (2023) as well as to their technical appendix. We note here that, as is common in the literature, independent prior distributions are assumed so as to facilitate the selection of prior distributions for the different blocks of parameters. Hence, the algorithm consists of six steps that iterate through five conditional posterior distributions. The first step involves using a particle-Gibbs sampler to draw α_t and $\ln h_t$ (the parameters that determine the contemporaneous response) conditional on the data and the other parameters. The second step draws the conditional posterior distribution of b, the covariance between the reduced-form error of interest and the instrument, using the algorithm developed by Carter and Kohn (1994). An independent Metropolis step is used to sample q_1 as in Caldara and Herbst (2019) in the third step. The structural shock of interest is then derived after drawing β and σ. The fourth step deals with estimating the parameters in the instrumental regression in equation 8. We assume that the parameters in the instrumental regression are time invariant to focus our attention on the transmission mechanism of the oil shocks. However, future research work could relax these assumptions and let both coefficients vary over time.

The fifth step deals with drawing from the conditional posterior distributions of the volatility parameters, i.e. Q_b, Q_α, and Q_h. In this step, we estimate the variance parameters that govern the time variation in the model transition equations, i.e., equations (3), (5), and (6).

Steps 4 and 5 help us to estimate the structural parameters, i.e., β and σ_v, and the volatility parameters, i.e., Q_b, Q_α, and Q_h, that govern the dynamics of the state-space model. The final step addresses missing observations in the instrument variable, i.e., z_t, and it is executed only if there are missing observations in the instrument or proxy variable. Following Mumtaz and Petrova (2023), we run the MCMC algorithm for 100,000 iterations, discard the first 50,000 as burn-in, and use the remaining draws for inference.

3.4 Lag Selection and Reliability of the Instrumental Variable

Following Primiceri (2005), it has become common in the TVP-VAR literature to select a lag length of two; higher lag lengths increase the degree of complexity and impose computational constraints. However, we opt for a higher number of lags ($p = 4$) in order to capture longer-lasting effects of oil prices underlined by the empirical literature using constant-coefficient VARs (see e.g., Hamilton and Herrera (2004)). We believe that the slightly longer lag length in conjunction with the rotation of inflation and inflation expectations in a four-variable Proxy TVP-VAR is a good compromise to capture longer dynamics while reducing the dimensionality of the models to facilitate computation. However, a possible concern with estimating two different Proxy TVP-VAR models that differ only on the last variable in the system (i.e., with inflation or inflation expectations as the fourth variable) is that the reliability of the proxy hinges on the contemporaneous correlation with the reduced-form residuals of the endogenous variables included in the system. Thus, the correlation is affected by which variables are included or excluded. To address this concern, we examine the reliability of the proxy in both models.

Hence, before we examine the response to the oil news shock, we provide empirical evidence on the relevance of the instrument used to identify oil shocks. To do so, we compute the posterior median of the reliability statistics, ρ. The median estimate ρ is 0.234 and the credible region with 68% credibility is given by [0.187,0.305] when we rotate in the CPI inflation. Given that the credible region does not contain zero, we conclude that the proxy is informative for identifying the oil news shock. In other words, the correlation coefficient between the proxy and the structural shock is statistically different from zero.

As Figure 2 illustrates, the posterior distribution for the parameters that govern the reliability of the instrument in the model that uses inflation expectations has a slightly lower median correlation between the instrument and the shock of interest (0.241) than the model with inflation (0.271). In contrast, the median of the posterior distribution for the variance of the IV equation is higher in the model with inflation expectations (1.157 versus 1.138). As a result, the reliability statistic for the model with inflation expectation is slightly lower (0.203) than for the model with inflation (0.234), yet it is significant with 68% confidence region of [0.154, 0.252]. Despite these small differences, the reliability statistic suggests that the proxy is relevant for the identification of oil news shocks in both models.

4 The Changing Effect of Oil News Shocks

As outlined in the introduction, several factors could have led to changes in the transmission of oil news shocks to inflation and inflation expectations: OPEC's share of world oil production has varied during the period under analysis, the U.S. went from being a net importer to a net exporter of oil as shale oil production increased, the importance of supply and demand driven shocks has varied over time, and inflation has become more persistent in recent years. These changes raise the question of whether the effect of oil news shocks has changed over time. To empirically examine this question, we obtain time-varying estimates of the impulse responses and forecast error variance decomposition via the Bayesian Proxy TVP-VAR described in the previous sections.

Given that the proxy is relevant in both TVP-VAR models and because the estimates for the oil market variables are very similar and lead to the same conclusions, in what follows the results reported for the oil variables and the realized inflation correspond to the inflation model, whereas the results for inflation expectations correspond to the model where we include inflation expectations.[4]

4.1 The Response of the Oil Market Variables

Figures 3-7 plot the estimated time-varying responses to an oil news shock. The responses have been normalized so that the oil news shock increases the price of crude oil by ten percent. The top left panel in each figure shows the median response for horizons of $h = 0, 1, 2, ..., 20$ months. In the remaining panels of each figure, the solid blue line represents the median response for selected horizons of interest ($h = 0, 1, 2, 4, 6, 8$), and the shaded area denotes the 68% credible sets. The dotted red line indicates the response estimated from the constant-coefficient Proxy SVAR model.

Our analysis of Figure 3 reveals two key findings regarding oil price dynamics. First, oil news shocks lead to a persistent increase in oil prices. Two years after the shock, the real oil price is still above its pre-shock level. Specifically, in the late 1980's, four quarters after the shock, the cumulative impact on the oil price had declined from the 10% increase on impact to about 5% eight quarters after the shock, suggesting a response half-life of around two years. In contrast, in the 2000s the cumulative response after eight quarters hovered around 9%, indicating a substantially more persistent effect of news shocks on oil prices. At most horizons, the constant-coefficient model suggests a slightly higher response. Yet, the fact that the cumulative estimate after eight quarters is greater in the 2020s than in earlier years provides evidence of increased persistence during the recent period of elevated inflation.

With respect to world oil production, Figure 4 shows a statistically insignificant response on impact and for most horizons. The time-varying impact response fluctuates around the value estimated for the constant-coefficient model. At most horizons, the 68% confidence regions contain zero. This general lack of significance for the time-varying responses is consistent with the estimates for a constant-coefficient Proxy-SVAR estimated over the full sample. However, it contrasts with the findings of Känzig (2021) derived from a constant-

[4] See Figures A.1–A.4 in the online Appendix for the responses of oil market variables using the inflation expectations proxy TVP-VAR model.

coefficient Proxy-VAR estimated on a shorter sample. He estimates that an oil news shock results in lower oil production: A shock that results in a 10% increase in oil prices leads to a 0.5% drop in world oil production 20 months after the shock. Estimation results reported in section 6, indicate that the differences stem from extending the sample and are not driven by the more parsimonious model specification. All in all, our estimation results are in line with the view that surprises derived from changes in oil futures prices around OPEC announcements do not necessarily reflect expectations of lower future oil production (see, e.g., Degasperi (2021), Kilian (2023)), but that the IV identifies precautionary or flow demand shocks.

The top left panel of Figure 5 illustrates the changes in the median response of the world industrial production index to time. In the earlier part of the sample, the decline in global economic activity is statistically and economically significant. An oil news shock that leads to a 10% increase in oil prices results in a persistent decline of about 1% in the world IP. On impact, this estimated response exceeds the 0.2% decline estimated with the constant-coefficient Proxy-SVAR estimated for the whole sample. However, starting in the 1990s, the impact effect became smaller, fluctuating between -0.5% in the early 1990s and turning insignificant in the late 2010s. This finding contrasts with the persistent decline in global economic activity estimated by Känzig (2021). The differences between estimates can be traced to two elements: the use of a sample that extends until 2022 and the use of a time-varying parameter model. Note that while the time-varying estimates indicate that oil news has a significant and persistent effect on oil prices throughout the sample, their impact on global economic activity has become muted.[5]

To summarize, we find evidence that the response of the oil market variables to oil news shocks has changed over time. Three insights are derived from the time-varying estimates.

[5]To allow for comparison, we also estimate a time-invariant Proxy VAR model using the same sample (i.e., January 1978 to December 2022) and with lag length set to $p = 4$. Figure A.20 shows impulse responses to the oil news shock, normalized to increase the real price of oil by 10%, with 68 and 95 percent confidence bands. A negative oil news shock immediately increases oil prices. Oil production exhibits a sharp increase after the shock, although it is statistically insignificant. Global industrial production shows a persistent and statistically significant decline.

First, the response of oil prices to oil news shocks has become more persistent over the years. Second, although we find some variation in the response of world oil production over time, the responses are estimated with a low degree of precision, suggesting the oil news shocks capture OPEC responses to future oil demand. Third, the effect of oil news shocks on world industrial production has become muted over time.

4.2 The Changing Impact of Oil News Shocks on Inflation and Inflation Expectations

Persistent inflation, not only in the US but globally, has prompted policymakers and academics to question whether structural changes post-COVID have made the inflation process more persistent (De Michelis *et al.*, 2024) and volatile (Powell, 2025). Such concerns underscore the importance of well-anchored inflation expectations and the need to understand the dynamic effect of changes in expectations about future oil prices on inflation and inflation expectations.

Figure 6 illustrates the time-varying response of inflation to oil news shocks. The top left panel of the figure shows that CPI inflation increases in response to an oil news shock, yet the effect is short-lived, lasting less than a year. The top right panel reveals substantial time variation in the contemporaneous effect of oil news shocks on inflation: the posterior median fluctuates from 0.6% in the late 1980s to 6.8% during the Great Recession, while a 2% increase is estimated with the constant-coefficient Proxy-SVAR. Furthermore, responsiveness tends to be higher during the expansionary period preceding the 2007-2008 financial crisis and the recovery period following the Covid-19 pandemic. It is interesting to note that we estimate a larger impact response than that obtained by Känzig (2021) with a constant-coefficient Proxy-VAR model. However, our time-varying estimates point to a less persistent effect on inflation.

Figure 7 shows that household inflation expectations increase in response to an oil news shock. Not surprisingly, the impact response of inflation expectations is considerably lower than that of inflation. The maximum increase in expectations reaches about 0.7% during the COVID-19 period versus 3% for realized inflation during the same period. Substantial time variation in the response of household inflation expectations is evident in Figure 7. It should be noted that while the constant coefficient estimate of the impact response lies between the values estimated with the TVP-VAR, the time-varying estimates for the four- and eight-quarter cumulative responses lie below the estimates from the constant coefficient model for the pre-COVID sample [6]

Finally, to further illustrate the time-varying effect of oil news shocks, we plot the responses for three specific historical episodes associated with unexpected increases in crude oil futures: (1) OPEC's agreement to cut production in November 1988 after a year of failed attempts to set quotas; (2) OPEC's agreement to reduce production in November 2016 (the first since 2008), and (3) OPEC+ agreement to maintain the reduced production levels established earlier in the year in April 2021.

Figure 8 depicts the impulse response functions for the Proxy TVP-VAR models where we rotate in the inflation rate (solid line) and inflation expectations (dashed line). Four results stand out. First, note that the impulse responses for the oil market variables are very similar in both models. The only notable differences across specifications are the faster decline in the real oil price in 2106M11 and 2011M04 and the positive response of oil production in 2021M4 when we rotate inflation expectations. Second, no statistically significant reduction in oil production is observed for any of the shocks. As mentioned above, this finding is in line with the view of Kilian (2024) and Degasperi (2021), who note that Känzig's proxy captures information about the economic outlook of OPEC and does not exclusively capture oil supply news. Moreover, the fact that only the oil news shock in 1988M11 leads to a decline in world industrial production reinforces this view. Third, we find a significant increase in the pass-through of oil news shocks to inflation in 2016M11 and 2021M04 relative

[6] In the time-invariant framework, both expected and realized inflation exhibit robust and persistent positive responses to the oil news shock, with the magnitude of the response of realized inflation exceeding that of inflation expectations in the short run (See the solid blue line response for the baseline time-invariant specification from the Online Appendix in Figure A.20).

to 1988M11. Finally, although the rise in inflation expectations is considerably smaller than that of inflation, we estimate a marked increase in the magnitude and persistence of the response during the COVID-19 pandemic.

4.3 Forecast Error Variance Decomposition

How important are oil news shocks to explaining the variation in inflation and inflation expectations over time? Has the contribution of oil news shocks to the forecast error variance decomposition changed? To answer these questions, we leverage the fact that the estimation method developed by Mumtaz and Petrova (2023) allows us to identify the shock scale and hence to estimate the time-varying forecast error variance decomposition.

The top panel of Figure 9a reveals two interesting facts. On impact, the contribution of oil news shocks to inflation volatility increased from around 10% at the beginning of the sample to above 50% in the early 2000s. Since then, it has fluctuated between 40% and 50%. A noticeable decline in the contribution of oil news shocks is evident since the beginning of the shale boom around 2014. Second, in the medium run (4 and 8 months), the contribution of oil news to CPI inflation exceeded the impact contribution at the beginning of the sample. Thereafter, the FEVD fluctuated between 35% and 55%.

Regarding the contribution of oil news shocks to volatility in inflation expectations, the bottom panel of Figure 9b shows a considerable degree of time variation. We estimate that the contribution fluctuated between 8 and 67 percent during the sample period. Large increases in the contribution of oil news shocks are observed during expansionary periods, especially in the mid-1990s and the late 2000s. Interestingly, we find that the contribution of oil news shocks to the variation in inflation expectations during the post-COVID period of high and persistent inflation exceeded the values estimated at any other points in the sample.

To conclude this section, we compare the time-varying FEVD with the estimates obtained with the constant-coefficient model (dotted line). At the beginning of the sample, the constant coefficient estimate falls outside the credible set 68% for the time-varying model, suggesting that the constant-coefficient model could lead to overstating the contribution of the shock during the earlier period. Regarding the contribution of the shock to inflation expectations, on impact, the constant-coefficient estimate often falls above the 68% credible set for the time-varying estimate; the only exception is the COVID-19 period, where the contribution of the time-varying model exceeds the constant-coefficient model estimate. At longer horizons, the constant-coefficient model can potentially lead to overstating the contribution of the oil news shock.

5 Digging deeper into the sources of time variation

As discussed above, we assume that the parameters of the IV regression are constant over time. Thus, we restrict the sources of the time variation to the evolution of the autoregressive parameters $B_{j,t}$, the non-zero and non-unity elements of the lower triangular matrix $A_{0,t}$, and the elements of H_t. Figure 10 reports the median coefficient estimates for the model where we rotate in the CPI inflation. The plots start in January 1988 as we use the earlier part of the sample to initialize the prior. Each panel reports the time-varying coefficients for the intercept c_t and the $B_{j,t}$ coefficients for each of the equations. The key takeaway is that there is very little time variation in the intercepts and lag coefficients. Some time variation is evident when we rotate in household inflation expectations (see Figure 11). In particular, we observe slight increasing or decreasing trends in the third/fourth lag coefficients that tend to be offset by an opposite trend in the shorter lags. However, the variation is limited; the only exception is the fourth autoregressive lag for inflation expectations, which shows a clear decline.

The six non-zero and non-unit elements of $A_{0,t}$ are plotted in Figure 12. The correlation pattern in the estimated covariance matrix of the shocks is related to these $\hat{\alpha}_{it}$. The panels labeled *Oil Price* $\rightarrow$ *Oil Production* and *Oil Price* $\rightarrow$ *World IP* depict correlations that are very close to zero, albeit negative and positive, respectively, for most of the sample. Of note is the positive correlation between the oil price and inflation (inflation expectations), as well as the negative correlation between oil production and inflation (inflation expectations). These observations are consistent with the notion that oil prices and inflation (inflation expectations) move in the same direction, while oil production and inflation (inflation expectations) move in opposite directions. Also of interest are: (1) the variation in the correlation between oil prices and inflation, which illustrates an increasing trend in the pass-through from oil prices to inflation expectations until the time of the fracking revolution, and (2) the rise in the correlation between oil prices and inflation expectations since the onset of the COVID-19 pandemic.

Figure 13 plots the coefficient estimates for the elements of H_t, which is a diagonal matrix. Fluctuations in these coefficients point to the importance of the volatility of the estimated errors in accounting for the time-variation in the impulse response functions. The overriding impression is that there is a significant degree of time variation in the stochastic volatility of all variables. The largest increases in stochastic volatility are estimated for oil prices, oil production, and world IP; considerable rises occur during recessionary periods, especially during the COVID-19 pandemic. For these three variables, the patterns observed for the two models (rotating in inflation or inflation expectations) are almost identical. Regarding inflation and inflation expectations, we note that the volatility of the former is an order of magnitude larger than that of the latter, and it exhibits a larger increase during the Financial Crisis than the COVID-19 pandemic. Overall, the most time variation is attributed to oil prices, followed by the inflation rate.

These findings complement the work of Chang *et al.* (2023), who estimate a univariate endogenous regime switching model for crude oil prices and then employ the extracted mean and volatility factors to study the impact of oil price fluctuations on inflation expectations. They find that shocks to the mean factor result in a rise in inflation expectations and disagreement, whereas shocks to the volatility factor lead to lower inflation expectations.

Our framework differs from theirs in three key aspects. First, Chang *et al.* (2023) employs a univariate Markov switching model that embodies discrete changes in the latent process driving oil price fluctuations. Then, in a second step, they estimate a VAR for the mean and latent factors, inflation expectations, and disagreement about the inflation forecast. In contrast, we jointly model the interaction between oil prices, oil production, global economic activity, and inflation or inflation expectations. Second, we use a flexible framework with random time variation that can account for different patterns of time variation (stemming from the VAR parameters or stochastic volatility) instead of modeling discrete changes in the mean and volatility of real oil prices and then assuming that the coefficients in the second-step VAR are fixed. Third, we identify the structural shock of interest via a an external instrument whereas Chang *et al.* (2023) assume that the system is recursive. Despite all these differences, both analyzes underscore the importance of modeling time variation in the volatility of oil prices to understand the transmission of oil price shocks to inflation expectations.

6 Robustness Checks

We run a battery of robustness checks for the constant-coefficient and TVP models. For the sake of brevity, we relegate the figures to the online appendix and describe the results in this section.

Alternative priors on the relevance of the instrument in the Proxy TVP-SVAR models. As mentioned above, the reliability statistic allows us to gauge the proxy's relevance for identifying the oil price shock. Recall that, given the evidence that oil surprises are relevant for the identification of oil news shocks (Känzig, 2021), the priors for β and σ are set such that $\rho \approx 0.2$ in the baseline model. In both models (rotating inflation or inflation expectations), the posterior median for ρ exceeds 0.2 and the credible sets 95% exclude 0, thus suggesting that the proxy is reliable. Yet, the reader may consider that the reliability indicator is small, which could point to some degree of misspecification in the baseline models (see e.g., Caldara and Herbst (2019)). To illustrate the robustness of the results to the choice of the prior, we perform a series of robustness checks for the inflation model. The various scenarios are intended to force a higher (or lower) signal-to-noise ratio β/σ. Estimation results are reported in Table 2. For reference, the first row reports the results for the baseline model. The second and third row report the results when we impose a tighter or looser prior on the variance of the measurement error, ν_0, respectively. As expected, a tighter (looser) prior results in a higher (lower) posterior median for the reliability statistic. However, the impulse response functions for the tighter prior, reported in Figures A.5 - A.9 of the online appendix, are almost identical to the baseline responses. The last three lines of Table 2, report the posterior median and credible regions for ρ obtained when we impose a prior with a lower mean for the measurement error. As the table illustrates, this prior induces a lower posterior mean on ρ, yet the 95% credible sets exclude 0, which indicates the proxy continues to be informative. Furthermore, the impulse responses reported in Figures A.11- A.14 of the online appendix reveal responses that are indistinguishable from the baseline estimates. The only noticeable difference is the tighter (looser) confidence sets when the priors are more (less) informative. These results suggest that, for the variables included in the model, the TVP Proxy-SVARs are well identified; changing the prior distribution of the parameters that govern the signal-to-noise ratio does not alter the posterior[7].

Alternative constant-coefficient proxy-SVAR specifications estimated on the full sample. Results reported in Figure A.21 of the appendix show that our estimates are robust to employing a four-variable model that includes oil prices, oil production, world IP and rotates inflation (dotted line) or inflation expectations (dotted dashed line) as the fourth variable using the sample running from January 1978 to December 2022. The median posterior estimates fall within the 90% credible region for the baseline. Similarly, the results are robust to the inclusion of inventories in the Proxy-SVAR model (dashed line). While the

[7]See also Figures A.15- A.19 of the online appendix.

figure reveals some differences between the responses of real oil prices, world oil production, and world IP in the baseline and Känzig (2021) specification, the responses of inflation and inflation expectations (the main interest of this paper) are very similar. As noted above, in Känzig (2021), the response of the real oil prices is less persistent, and world oil production and industrial production exhibit a decline in response to the oil news shock.

Alternative constant-coefficient proxy-SVAR specifications estimated on the Känzig's sample. We compare our results to the baseline specification used by Känzig (2021), which includes the real price of oil, world oil production, world industrial production, world oil inventories, US industrial production and the US CPI log level (instead of the CPI inflation). As Figure A.22 in the online appendix illustrates, Känzig (2021) baseline specification estimated on his sample spanning January 1974 to December 2017 implies a less persistent response of oil prices, a decline in world oil production (consistent with the notion that oil news corresponds to a decline in oil supply), and a decline in world industrial production. Estimation results from the 5-variable specification that excludes oil inventories –consistent with our Proxy TVP-SVAR specification– fall within the 68% credible region for the baseline estimates.

7 Conclusions

In this paper, we examine the presence of time variation in the response of oil markets, US inflation, and inflation expectations to oil news shocks. We find evidence that the effect of oil news shocks that lead to a 10% contemporaneous increase in oil prices on the real oil price has become more persistent. However, such shocks continue to have a long-lasting and statistically significant effect on the real price of crude oil. Regarding other oil market variables, we find no statistically significant response of world oil production over the sample, and evidence that the responsiveness of world oil production has diminished over time.

Our results suggest that there has been a considerable degree of time variation in the response of US inflation and inflation expectations. Of relevance for the current debate on the possibility of a more volatile and persistent inflation is our finding of a significant increase in the sensitivity of inflation and, especially, inflation expectations to oil news shocks in the post-COVID period. Moreover, our estimates of the forecast error decomposition indicate that oil news shocks played a crucial role in explaining the volatility of inflation and, especially, inflation expectations in recent years. Finally, we traced most of the roots of the time variation to stochastic volatility.

The results presented in this paper demonstrate that, in the context of higher and more volatile inflation, oil news shocks matter for inflation expectations. Given that inflation expectations play a key role in determining the future course of inflation, policymakers should continue to monitor developments in the oil market.

References

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**, 1087–1109.

Baumeister, C. and Peersman, G. (2013b) Time-varying effects of oil supply shocks on the us economy, *American Economic Journal: Macroeconomics*, **5**, 1–28.

Bjørnland, H. C. and Skretting, J. (2018) The shale oil boom and the us economy: Spillovers and time-varying effects, *Journal of Applied Econometrics*.

Blanchard, O. and Galí, J. (2010) Labor markets and monetary policy: A new keynesian model with unemployment, *American economic journal: macroeconomics*, **2**, 1–30.

Caldara, D. and Herbst, E. (2019) Monetary policy, real activity, and credit spreads: Evidence from bayesian proxy svars, *American Economic Journal: Macroeconomics*, **11**, 157–192.

Carter, C. K. and Kohn, R. (1994) On gibbs sampling for state space models, *Biometrika*, **81**, 541–553.

Chang, Y., Herrera, A. M. and Pesavento, E. (2023) Oil prices uncertainty, endogenous regime switching, and inflation anchoring, *Journal of Applied Econometrics*, **38**, 820–839.

Cogley, T. and Sargent, T. J. (2005) Drifts and volatilities: monetary policies and outcomes in the post wwii us, *Review of Economic dynamics*, **8**, 262–302.

De Michelis, A., Lofstrom, G., McHenry, M., Orak, M., Queralto, A. and Scaramucci, M. (2024) Has the inflation process become more persistent? evidence from the major advanced economies.

Degasperi, R. (2021) Identification of expectational shocks in the oil market using opec announcements, *manuscript, University of Warwick*.

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*, pp. 265–286.

Herrera, A. M. and Pesavento, E. (2009) Oil price shocks, systematic monetary policy, and the "great moderation", *Macroeconomic Dynamics*, **13**, 107–137.

Känzig, D. R. (2021) The macroeconomic effects of oil supply news: Evidence from opec announcements, *American Economic Review*, **111**, 1092–1125.

Kilian, L. (2023) How to construct monthly var proxies based on daily futures market surprises.

Kilian, L. (2024) How to construct monthly var proxies based on daily surprises in futures markets, *Journal of Economic Dynamics and Control*, **168**, 104966.

Klein, M. and Linnemann, L. (2020) The time-varying effect of fiscal policy on inflation: evidence from historical us data, *Economics Letters*, **186**, 108823.

Koop, G., Leon-Gonzalez, R. and Strachan, R. W. (2009) On the evolution of the monetary policy transmission mechanism, *Journal of Economic Dynamics and Control*, **33**, 997–1017.

Mertens, K. and Ravn, M. O. (2013) The dynamic effects of personal and corporate income tax changes in the united states, *American economic review*, **103**, 1212–47.

Mumtaz, H. and Petrova, K. (2023) Changing impact of shocks: A time-varying proxy svar approach, *Journal of Money, Credit and Banking*, **55**, 635–654.

Naggert, K., Rich, R. W. and Tracy, J. (2023) The anchoring of us inflation expectations since 2012, *Economic Commentary*.

OPEC (1990) *OPEC: official resolutions and press releases 1960-1990*, Organization of the Petroleum Exporting Countries.

Powell, J. H. (2025) Opening remarks: Second thomas laubach research conference.

Primiceri, G. E. (2005) Time varying structural vector autoregressions and monetary policy, *The Review of Economic Studies*, **72**, 821–852.

Ramey, V. A. and Vine, D. J. (2011) Oil, automobiles, and the us economy: How much have things really changed?, *NBER macroeconomics annual*, **25**, 333–368.

Table 1: Prior Distributions and Starting Values for the Time-Varying Proxy VAR Model

Parameter	Description	Prior Distribution	Prior Scale Matrix/Value	Degrees of Freedom/Formula
b_t	VAR coefficients	$b_t = b_{t-1} + Q_b^{1/2}\eta_t^b$ $\eta_t^b \sim \mathcal{N}(0, I_{4,4p+1})$	$\beta_0 = vec(b_0)'$ $P_{0\|0} = 10 \times V_0$	Initial state based on OLS using training sample $T_0 = 120$
Q_b	Coefficient innovations	$IW(Q_0, T_0 + T)$	$Q_0 = V_0 \times T_0 \times 3.5 \times 10^{-4}$ $V_0 = kron(\Sigma_0, (X_0'X_0)^{-1})$	$T_0 = 120$
α_t	Non-zero, non-one elements of $\tilde{A}_t$ i.e.,$NNx = 6$ (lower triangular)	$\alpha_t = \alpha_{t-1} + Q_\alpha^{1/2}\eta_t^\alpha$ $\eta_t^\alpha \sim \mathcal{N}(0, I_6)$	Initial value from C_0 C_0 normalized and inverted	$C_0 = chol(\Sigma_0)$ $C_0 = (C_0^{-1})'$
Q_α	α_t innovations	$IW(D_0, TD_0 + T)$	D_0: diagonal with 10^{-4} elements	$TD_0 = NNx + 10 = 16$
$\ln h_t$	Log volatilities	$\ln h_t = \ln h_{t-1} + Q_h^{1/2}\eta_t^h$ $\eta_t^h \sim \mathcal{N}(0, I_4)$	$MU_0 = \ln(diag(\Sigma_0))$	Prior mean for the initial log-volatilities MU_0 and variance $SS_0 = 10$
Q_h	$\ln h_t$ innovations	$IW(g_0 I, TG_0 + T)$	$g_0 = 10^{-4}$	$TG_0 = N + 10 = 14$
q_1	First column of q	Uniform on the unit sphere	Metropolis step	
β	Instrument relevance	$\mathcal{N}(b_{0m}, v_{0m})$	b_{0m} from OLS regression of VAR residuals on instrument $v_{0m} = 0.1$	Estimated from regression of residuals on oil surprises
σ^2	Instrument noise variance	$\frac{1}{\sigma^2} \sim \mathcal{G}(a, b)$	$\sigma_0 = b_{0m}$(OLS estimate=0.354) $\nu_0 = 2(2 + \frac{\sigma_0^2}{v_0^2})$ $s_0 = 2\sigma_0(1 + \frac{\sigma_0^2}{v_0^2})$ $v_0^2 = 0.02$	$a = \frac{\nu_1}{2}$, $b = \frac{2}{s_1}$ $\nu_1 = \nu_0 + T$ $s_1 = s_0 + \hat{v}_t'\hat{v}_t$ where $\hat{v}_t = m_t - \hat{\beta}\varepsilon_{1t}$

Notes: The 4-variable VAR model includes real oil price, oil production, world industrial production, and inflation with $p = 4$ lags. Key model parameters: $N = 4$ (number of variables), $T_0 = 120$ (training sample size), T (remaining sample size), $NNx = N(N-1)/2 = 6$ (number of free elements in the impact matrix). Σ_0 is the variance-covariance matrix of reduced-form residuals from the training sample, $V_0 = kron(\Sigma_0, (X_0'X_0)^{-1})$ is the Kronecker product used for coefficient priors, and $\beta_0 = vec(b_0)'$ is the vectorized OLS coefficients. The external instrument is related to structural shocks through $m_t = \beta\varepsilon_{1t} + \sigma v_t$. MCMC uses 100,000 replications with 50,000 burn-in and every 10th draw is retained.

Table 2: Robustness of Reliability Measure (ρ) to Alternative Priors - Median Estimates and Credible Regions

Prior mean σ_0	**Prior Variance** v_0	**Median**	**68% CR**	**95% CR**
0.354	0.020	**0.234**	[0.187, 0.279]	[0.157, 0.305]
0.354	0.010	**0.458**	[0.407, 0.507]	[0.371, 0.535]
0.354	0.100	**0.102**	[0.067, 0.135]	[0.045, 0.158]
0.177	0.020	**0.148**	[0.109, 0.184]	[0.083, 0.208]
0.177	0.015	**0.173**	[0.130, 0.213]	[0.102, 0.239]
0.177	0.025	**0.126**	[0.090, 0.162]	[0.063, 0.186]

Note: The estimates are obtained using Time Variant Proxy VAR for the Inflation Model. The priors for σ *are parameterised in terms of mean* σ_0 *and the variance* v_0*. Note that* $\frac{1}{\sigma^2} \sim \mathcal{G}(a, b)$ *where* $a = \frac{\nu_1}{2}$, $b = \frac{2}{s_1}$*. The parameters of this Gamma density are given by* $\nu_1 = \nu_0 + T$ *and* $s_1 = s_0 + \hat{v}_t' \hat{v}_t$ $\hat{v}_t = m_t - \hat{\beta}\varepsilon_{1t}$. $s_0 = 2\sigma_0 \left(1 + \frac{\sigma_0^2}{v_0^2}\right)$ *while* $\nu_0 = 2\left(2 + \frac{\sigma_0^2}{v_0^2}\right)$*. 68% confidence intervals correspond to the 16th and 84th percentiles, while 95% confidence intervals correspond to the 5th and 95th percentiles.*

Figure 1: Oil Surprises

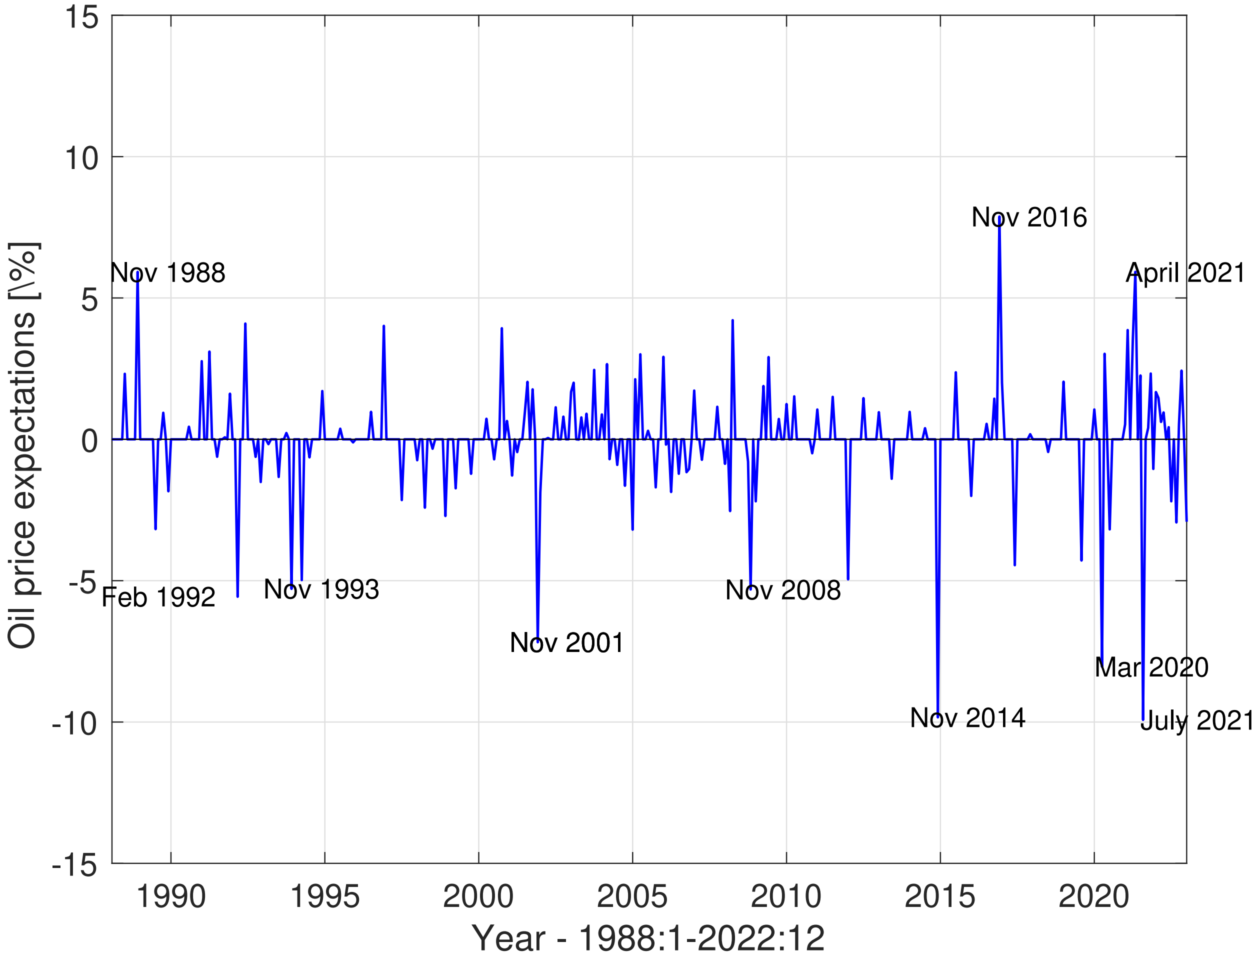

Figure 2: Proxy-VAR Posterior Parameter Distributions

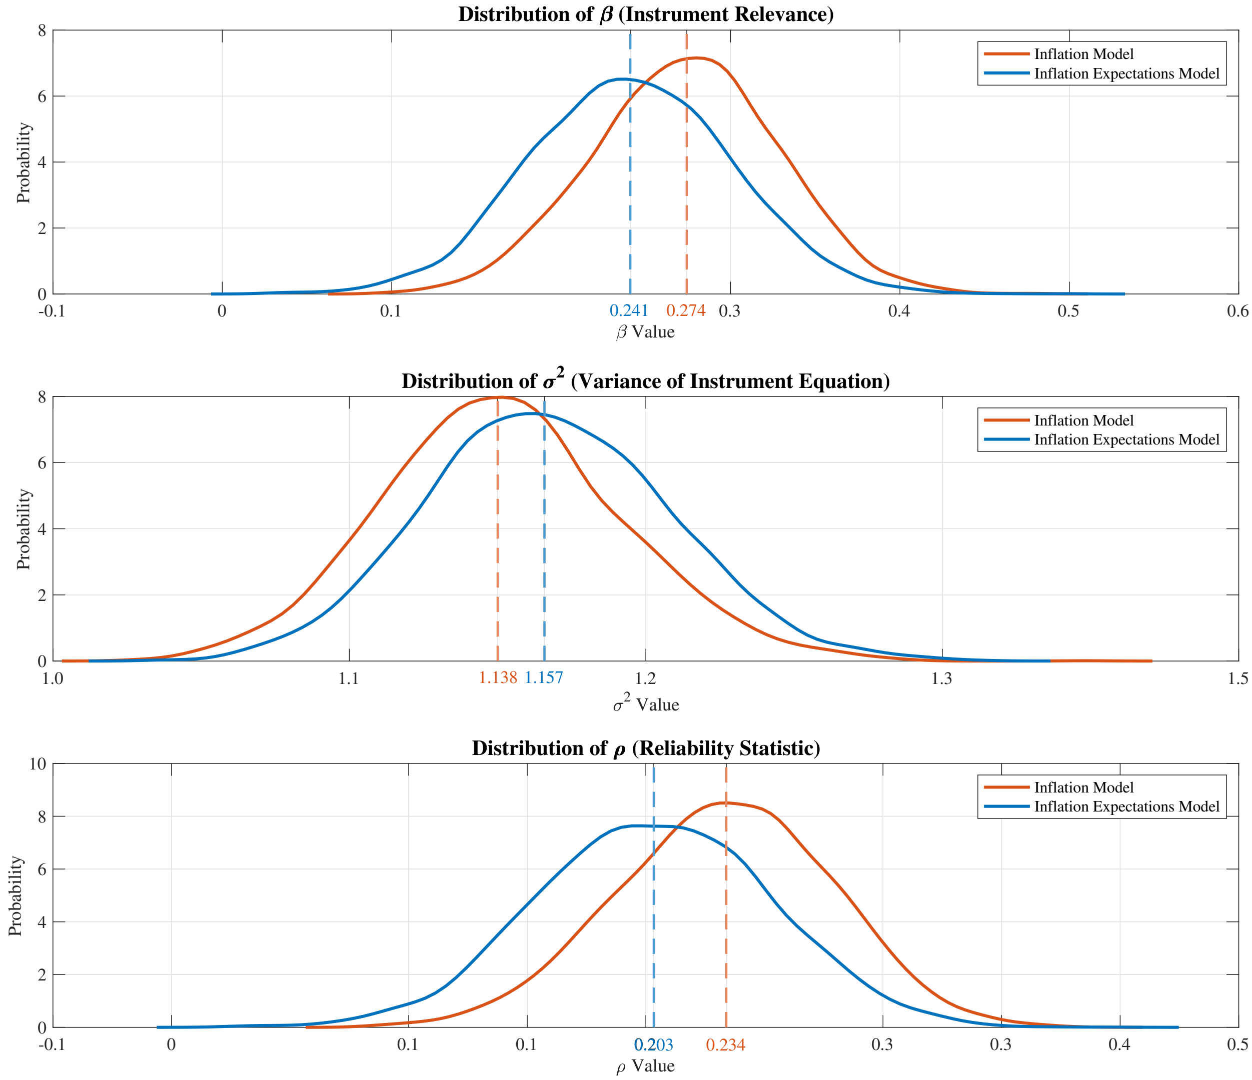

Notes: The figure shows the posterior distributions of three key parameters. Top panel: Distribution of β (instrument relevance) with median = 0.241 and 0.274. Middle panel: Distribution of σ^2 (variance of instrument equation) with median = 1.138 and 1.157. Bottom panel: Distribution of ρ (reliability statistic) with median = 0.203 and 0.234.

Figure 3: Posterior Median Response of the Oil Price to the Oil News Shock

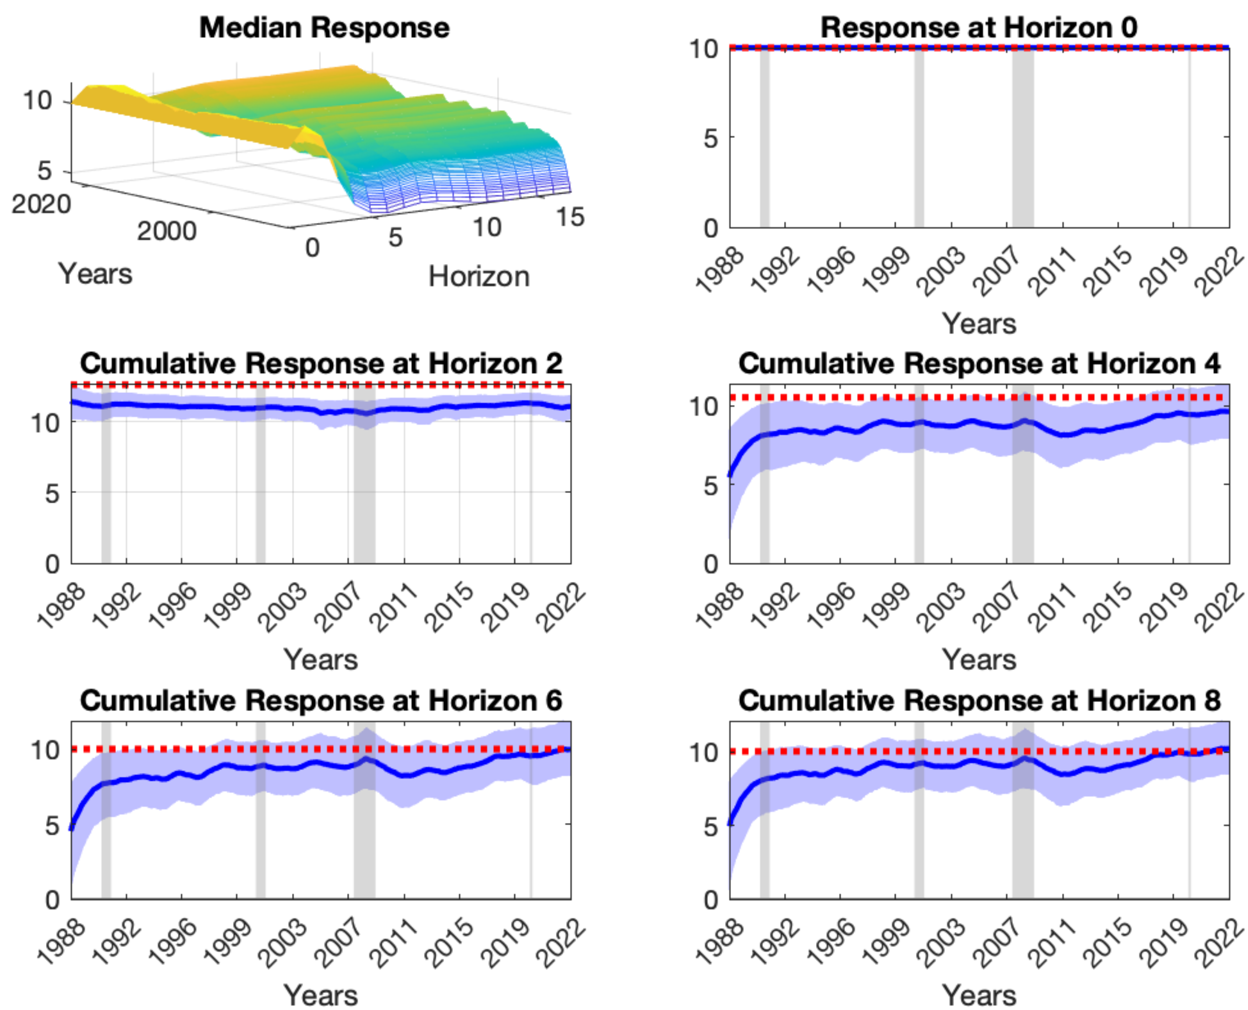

Note: The oil supply news shock normalized to a 10 percent increase in the real price of crude oil. The dotted red line is the response from a constant-coefficient Proxy SVAR model. The shaded area indicated the 68 percent posterior credible sets. The vertical bars represent the recession periods for the United States as identified by the NBER.

Figure 4: Posterior Median Response of World Oil Production to Oil News Shock

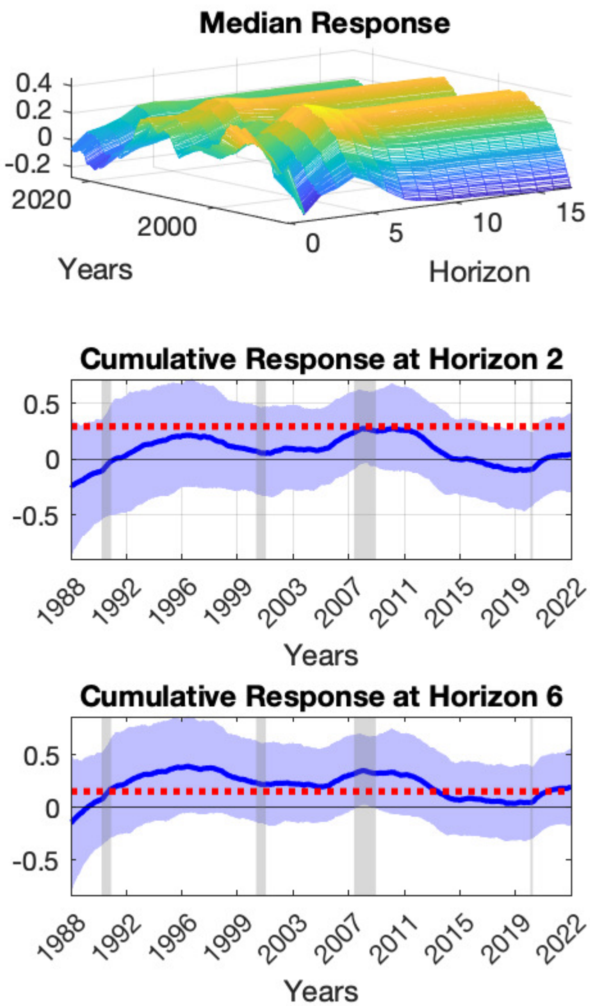

Note: The oil supply news shock normalized to a 10 percent increase in the real price of crude oil. The dotted red line is the response from a constant-coefficient Proxy-VAR model The shaded area indicated the 68 percent posterior credible sets. The vertical bars represent the recession periods for the United States as identified by the NBER.

Figure 5: Posterior Median Response of World Industrial Production to Oil News Shock

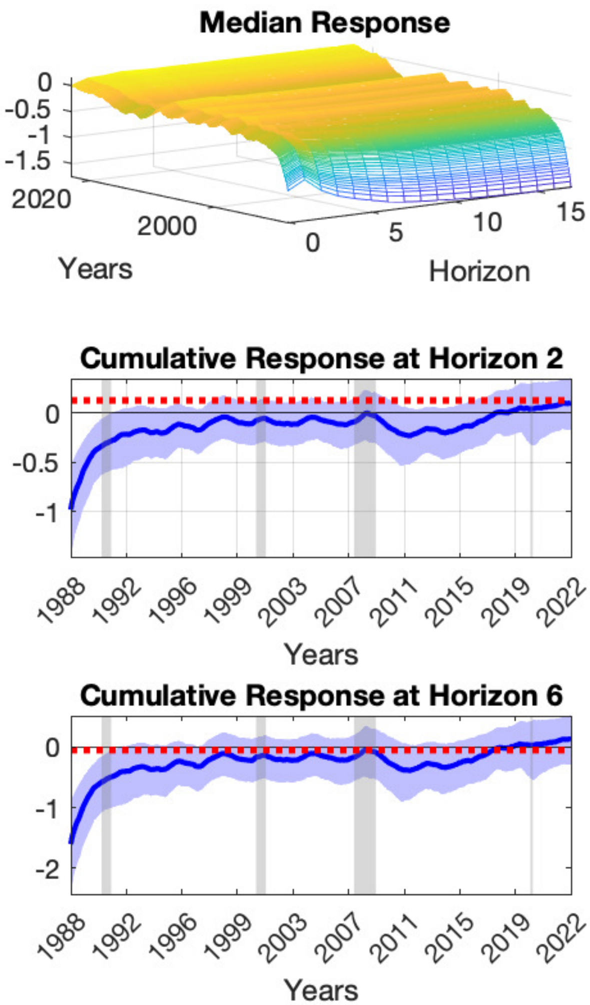

Note: The oil supply news shock normalized to a 1 percent increase in the real price of crude oil. The dotted red line is the response from a constant-coefficient Proxy-VAR model. The dashed lines indicated the 68 percent posterior credible sets. The vertical bars represent recession periods for the United States as identified by the NBER.

Figure 6: Posterior Median Response of Inflation to Oil News Shock

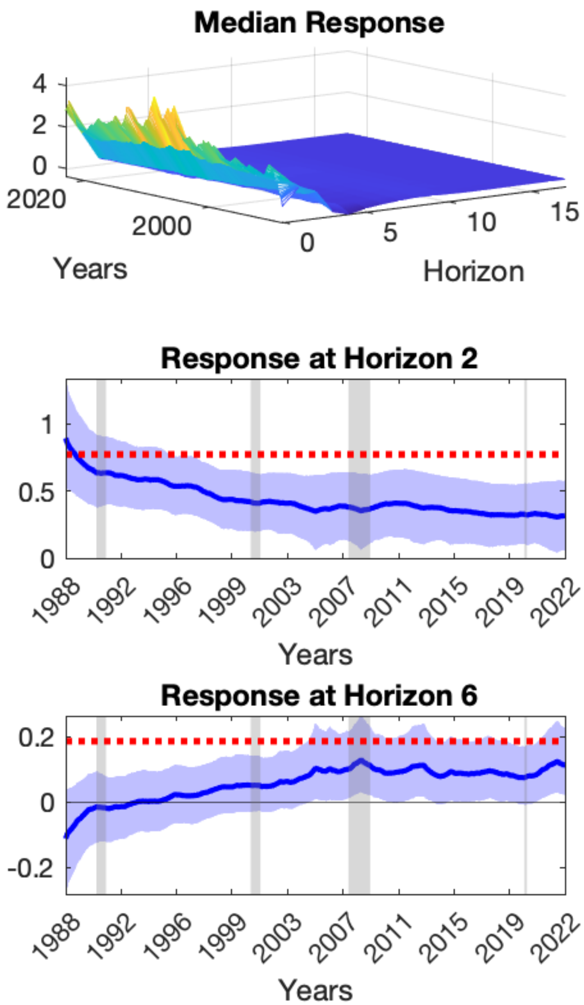

Note: The oil supply news shock normalized to a 1 percent increase in the real price of crude oil. The dotted red line is the response from a constant-coefficient Proxy-VAR model. The dashed lines indicated the 68 percent posterior credible sets. The vertical bars represent the recession periods for the United States as identified by the NBER.

Figure 7: Posterior Median Response of Inflation Expectations to Oil News Shock

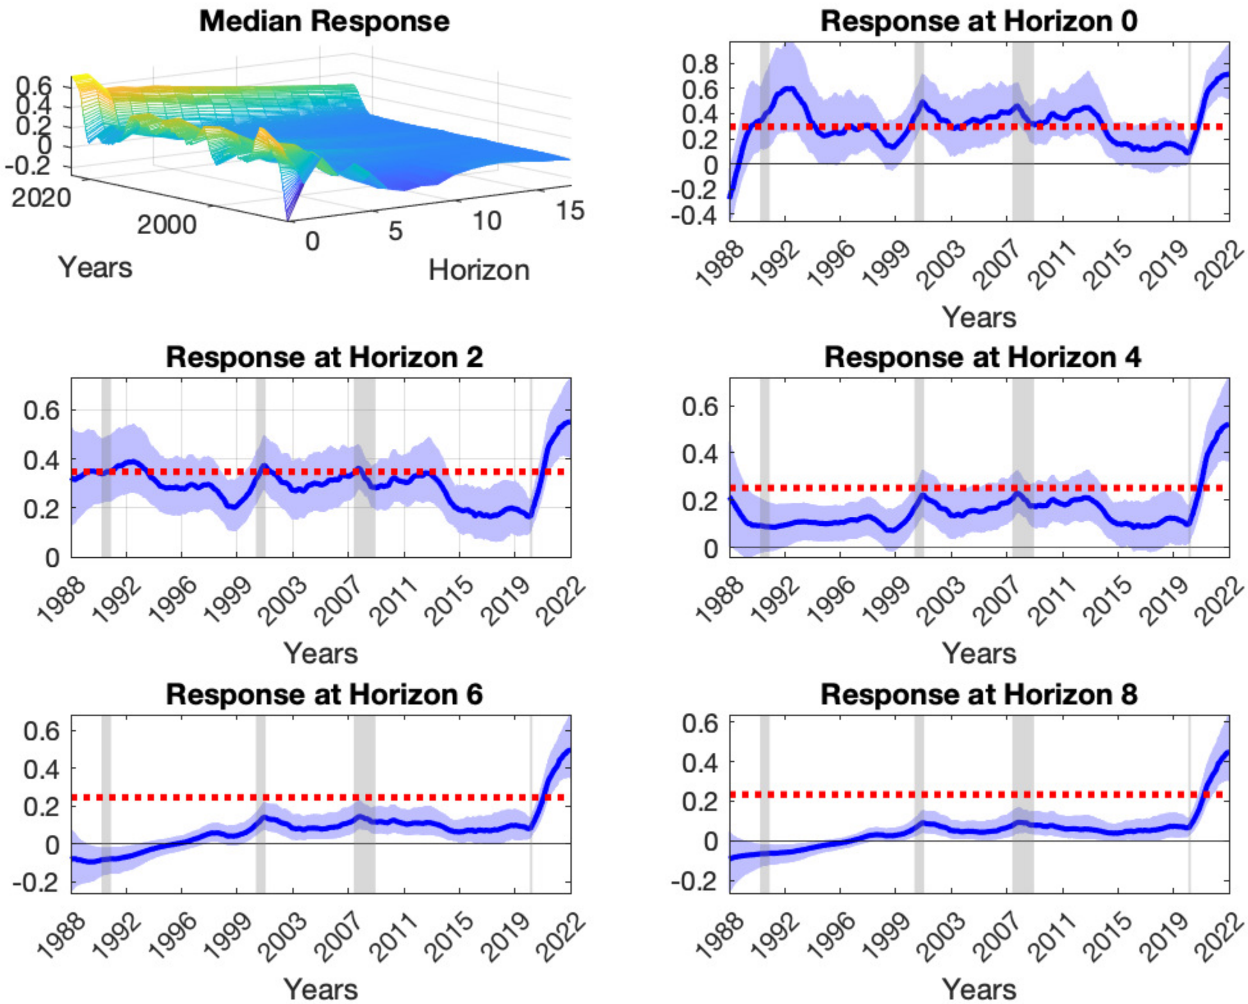

Note: The oil supply news shock normalized to a 1 percent increase in the real price of crude oil. The dotted red line is the response from a constant-coefficient Proxy-VAR model. The dashed lines indicated the 68 percent posterior credible sets. The vertical bars represent the recession periods for the United States as identified by the NBER.

Figure 8: Impulse Response Functions at Three Points in Time

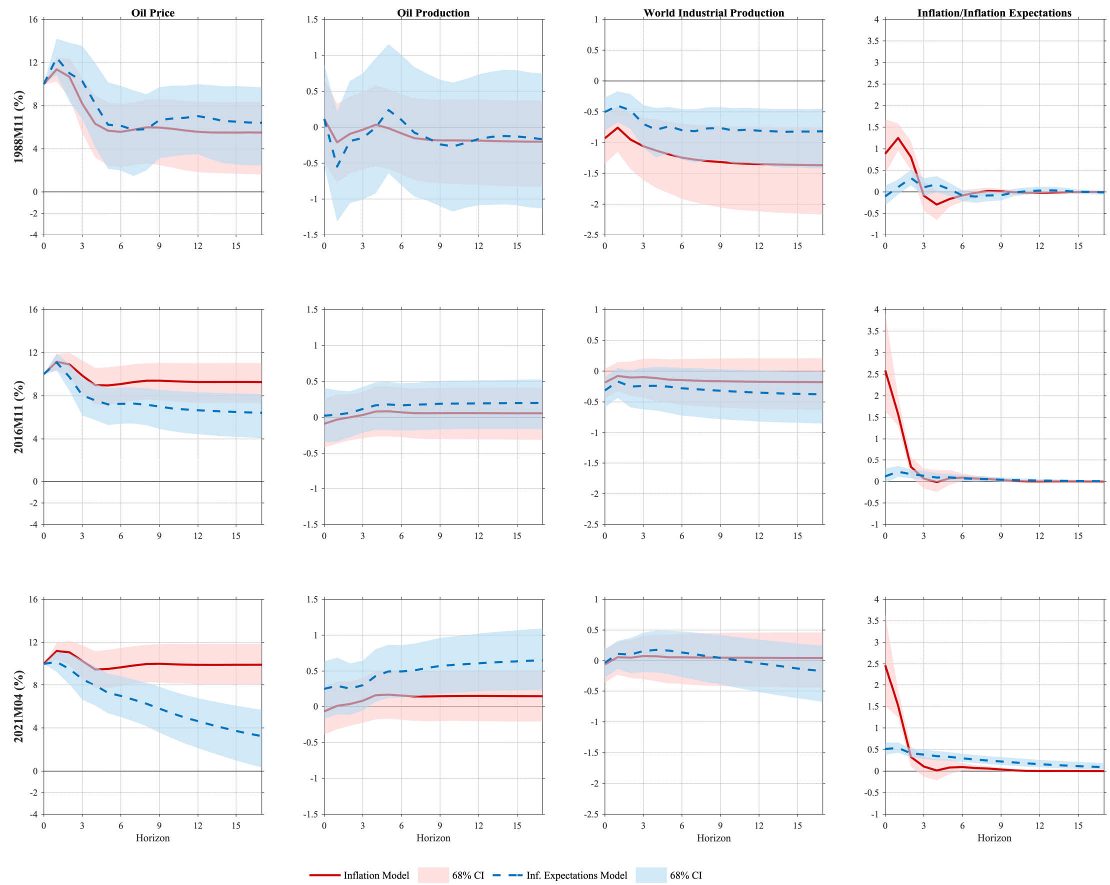

Note: This figure shows the impulse response functions of oil price, oil production, world industrial production, and inflation/inflation expectations (last column) for three episodes: 1988M11 (top), 2016M11 (middle), and 2021M04 (bottom). Red solid lines represent the Inflation Model, and blue dashed lines represent the Inflation Expectations Model. Shaded areas indicate 68% confidence bands.

Figure 9: Contribution of the Oil News Shock to the Forecast Error Variance on Inflation and Inflation Expectations in percent

(a) Inflation Rate

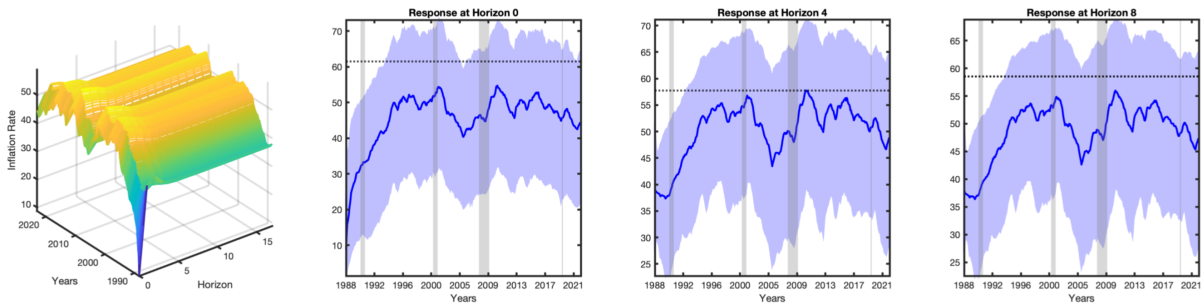

(b) Inflation Expectations

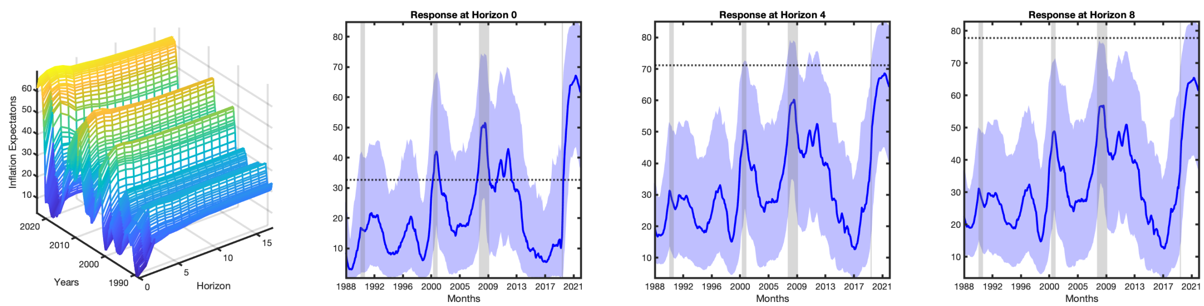

Note: The solid line indicates the median contribution in the TVP Proxy VAR model. The dotted line is the median contribution from the time-invariant model. The shaded areas indicate the 68 percent posterior credible sets. The vertical bars represent the recession periods for the United States as identified by the NBER.

Figure 10: Autoregressive Coefficient Median Estimates (1988-2022) for the Inflation Model.

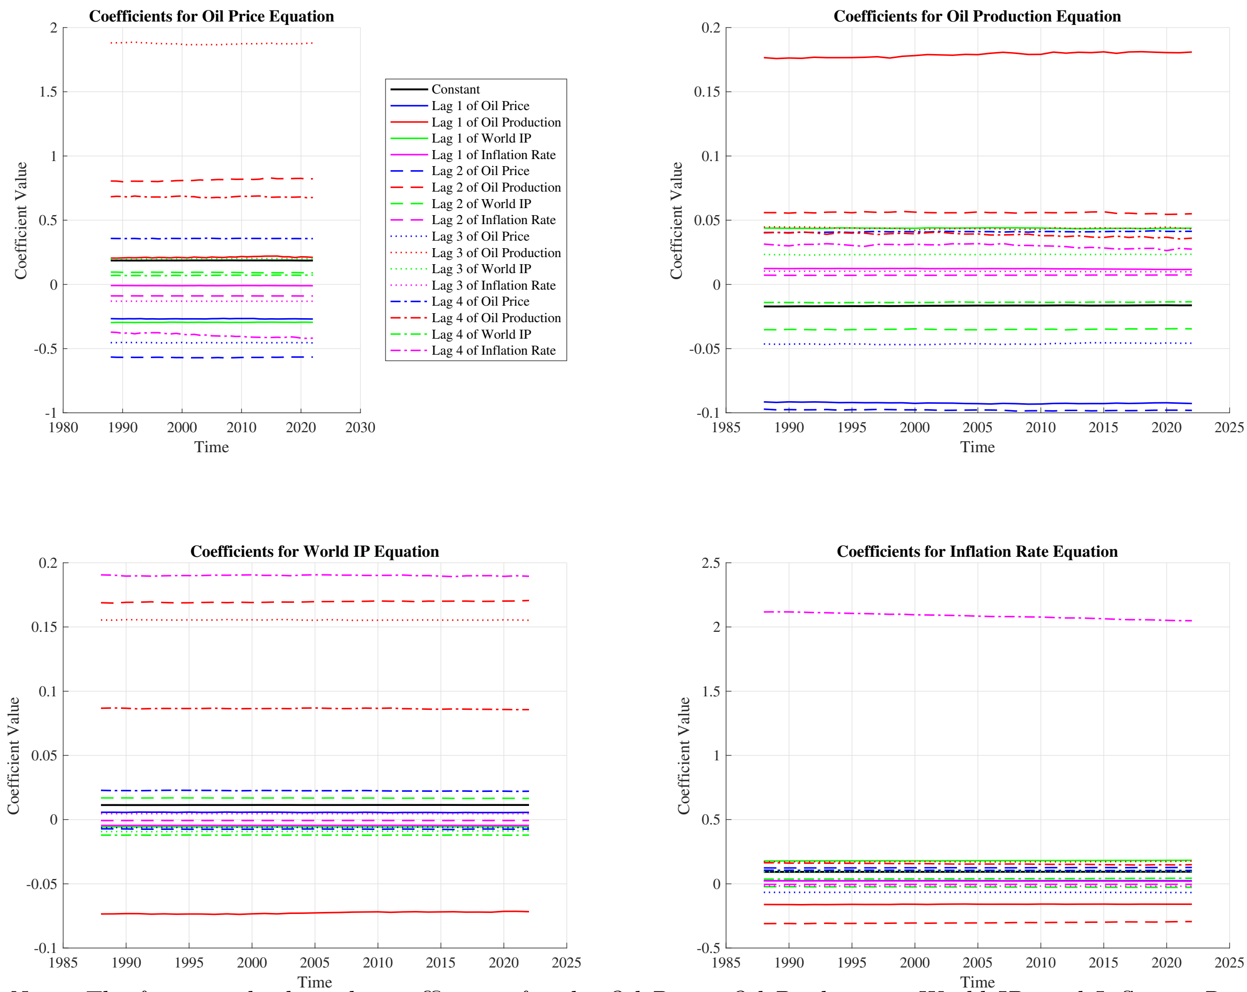

Note: The four panels show the coefficients for the Oil Price, Oil Production, World IP, and Inflation Rate equations.

Figure 11: Autoregressive Coefficient Median Estimates (1988-2022) for the Inflation Expectations Model.

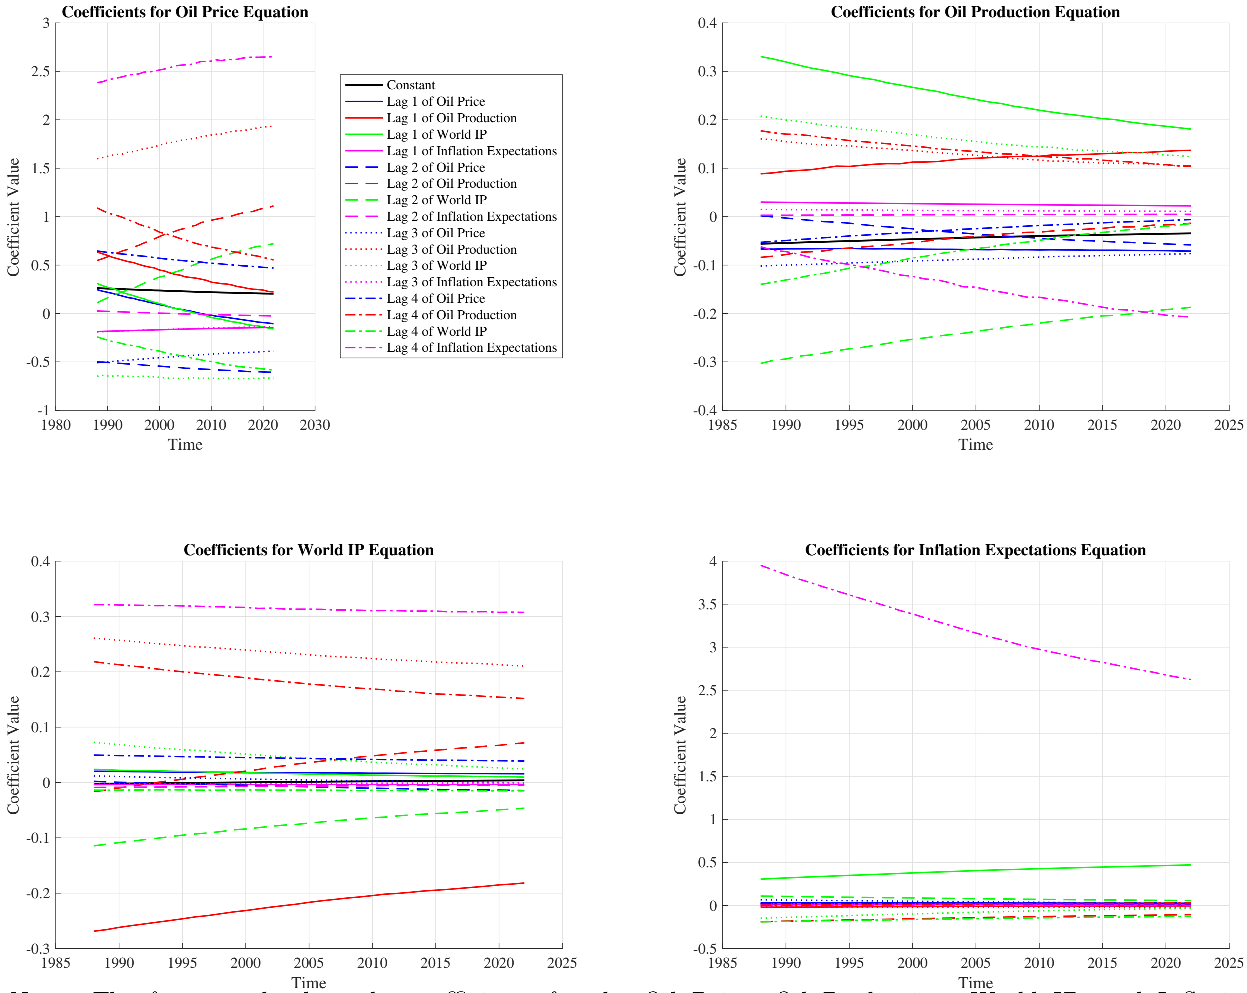

Note: The four panels show the coefficients for the Oil Price, Oil Production, World IP, and Inflation Expectations equations.

Figure 12: Time-Varying Contemporaneous Relationships (α_t)

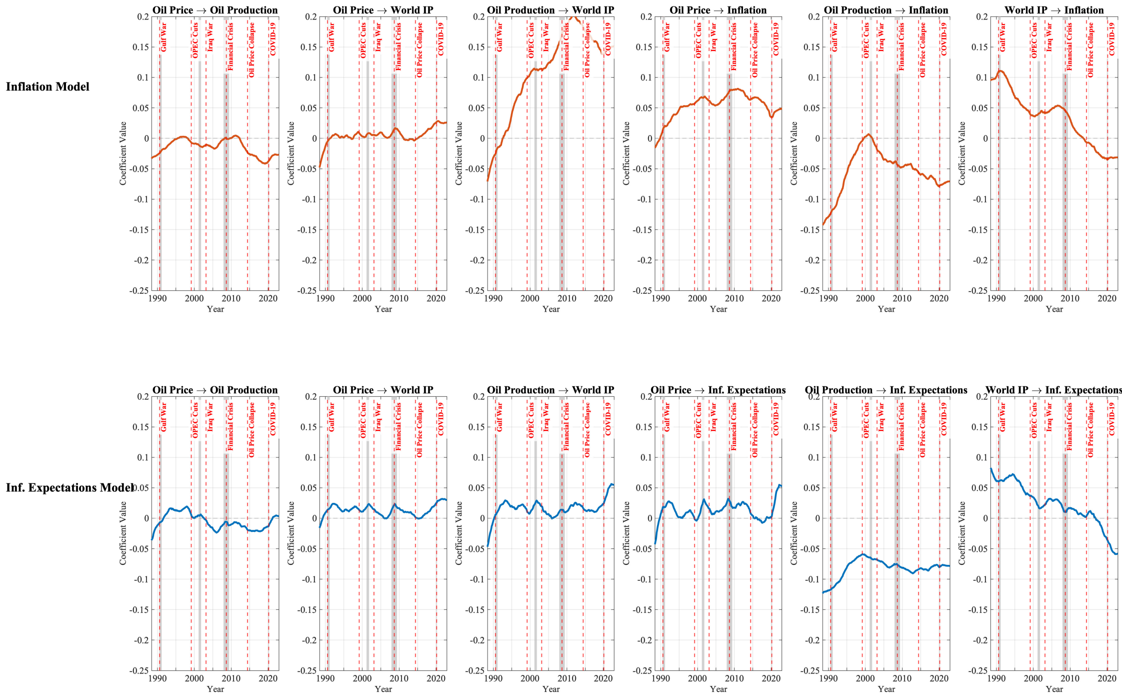

Note: This figure displays the evolution of contemporaneous relationships between oil market variables from 1988 to 2022. Red dashed lines mark major economic and geopolitical events. Gray shaded areas indicate NBER recession periods.

Figure 13: Time-Varying Stochastic Volatilities (Standard Deviations)

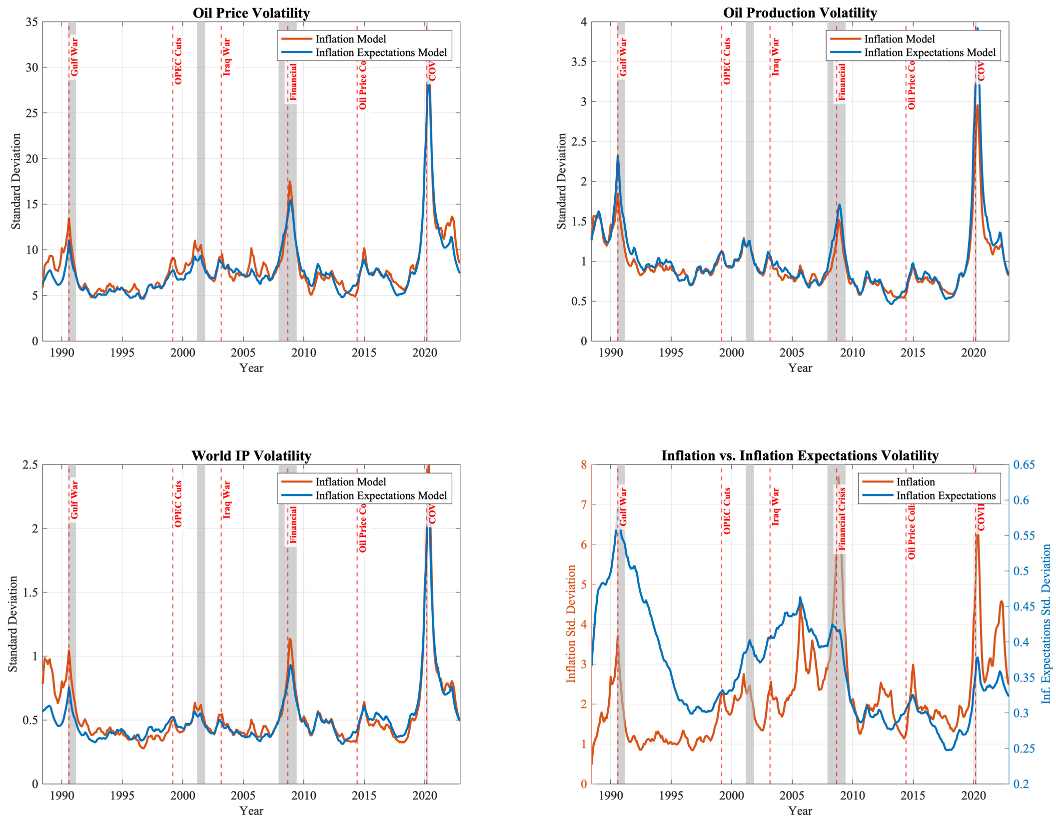

Notes: The figure shows the estimated stochastic volatilities for Oil Price (orange), Oil Production (blue), World IP (yellow), and Inflation (purple) from 1988 to 2022. Volatility is measured as the standard deviation of monthly percentage changes. The COVID-19 period exhibits unprecedented volatility in oil prices, reaching approximately 30 percentage points.