

Working Paper

Re-visiting the Relationship Between Oil Prices and Monetary Policy

Norges Bank Research

Authors:

Hilde C. Bjørnland

Jamie L. Cross

Jonas Hölz

Keywords:

Proxy-SVAR, Monetary Policy

Instrument, Oil Price Instrument,

Counterfactual Analysis, Monetary

Policy trade-offs

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles på e-post: servicesenter@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på <http://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 forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail: servicesenter@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 1502-8143 (online)

ISBN 978-82-8379-357-4 (online)

Re-visiting the Relationship Between Oil Prices and Monetary Policy*

Hilde C. Bjørnland[†] Jamie L. Cross[‡] Jonas Hölz[§]

March 15, 2025

Abstract

This paper examines how central banks respond to supply-side shocks and investigates the trade-offs they face in stabilizing inflation and output. To do so we develop a dual external instrument proxy structural vector autoregressive (SVAR) model to disentangle the macroeconomic effects of oil supply news and monetary policy shocks. Our identification strategy, which combines multiple external instruments with sign restrictions, enables a sharp distinction between structural shocks, allowing us to analyze their dynamic effects and construct policy counterfactuals for different central bank objectives. We find that both oil supply and monetary policy shocks significantly influence U.S. output and inflation. Moreover, while monetary policy can mitigate some of the output losses caused by oil price shocks, it cannot fully offset their inflationary effects. Finally, we estimate that the Federal Reserve's historical response aligns closely with a policy that places twice as much weight on inflation stabilization than on output stabilization.

JEL classification codes: C32, E31, E43, Q41, Q43

Keywords: Proxy-SVAR, Monetary Policy Instrument, Oil Price Instrument, Counterfactual Analysis, Monetary Policy trade-offs

*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 Norges Bank. We thank Knut Are Aastveit, Christiane Baumeister, Chenghan Hou, Silvia Miranda-Agrippino, Lorenzo Mori, Matthew Read, Leif Anders Thorsrud, Benjamin Wong, an anonymous referee for our Norges Bank Working Paper, and conference participants at the SNDE 2024 Symposium in Padova, the Dolomiti Macro Meeting 2024 in Kastelruth and the Midwest Econometrics Group Conference 2024 in Lexington for useful comments and suggestions. This paper is part of the research activities at the Centre for Applied Macro and commodity Prices (CAMP) at the BI Norwegian Business School.

[†]BI Norwegian Business School and Norges Bank. Email: hilde.c.bjornland@bi.no

[‡]Melbourne Business School, University of Melbourne. Email: j.cross@mbs.edu

[§]BI Norwegian Business School. Email: jonas.holz@bi.no

1 Introduction

Periods of commodity price volatility have long posed challenges for monetary policymakers, raising questions about the extent to which central banks should intervene to mitigate their economic consequences. The recent surge in inflation following the COVID-19 pandemic and the Russia-Ukraine war has reignited interest in this issue, drawing comparisons to the high-inflation episodes of the 1970s and 1980s. These historical parallels highlight an enduring macroeconomic question: what constitutes an optimal central bank response to supply shocks?

A central issue in this discussion is the trade-off between inflation stabilization and sustaining economic activity that central banks face when responding to supply shocks. This question has been at the core of macroeconomic research since the seminal work of Bernanke et al. (1997), who argued that the depth of oil-induced recessions is largely a consequence of the central bank's reaction rather than the shock itself. However, subsequent studies, including Hamilton and Herrera (2004), Herrera and Pesavento (2009), Kilian and Lewis (2011) and Bjørnland et al. (2018), have challenged this view, demonstrating that monetary policy effectiveness depends on the underlying model framework, the design of counterfactual exercises and broader economic conditions, including the prevailing policy regime.

Empirical challenges in disentangling the effects of monetary policy from other macroeconomic disturbances have long hindered progress in resolving this debate. However, recent advancements in multivariate time-series modeling have improved shock identification techniques, while the development of more robust counterfactual analysis offers the potential for a more precise assessment of how monetary policy influences economic outcomes following supply shocks.

Against this backdrop, our study employs state-of-the-art econometric techniques to reassess this long-standing question, offering new insights into the trade-offs central banks face in navigating supply-side shocks. Specifically, we employ a dual external instrument proxy structural vector autoregressive (SVAR) model that allows us to jointly identify monetary policy shocks and oil supply news shocks within the SVAR. This builds on the proxy SVAR literature (Stock and Watson, 2012; Mertens and Ravn, 2013) and recent extensions that allow for multiple instruments (Piffer and Podstawski, 2018; Arias et al., 2021; Braun and Brüggemann,

2022; Giacomini et al., 2022; Angelini et al., 2024).

Precise shock identification in structural analysis requires carefully chosen instruments. To construct the monetary policy instrument, we follow the methodology of Bauer and Swanson (2023a) with two adjustments: extending their analysis through the COVID-19 period and by adding oil prices into the set of macroeconomic variables used to orthogonalize the shock series. This latter enhances robustness, particularly in periods of strong oil-market-policy interaction. For oil supply news shocks, we adapt the instrument from Känzig (2021), which identifies shocks via oil futures price movements around OPEC meetings. To ensure robustness, we also test some recent alternative oil supply measures (Baumeister and Hamilton (2019); Degasperi (2023)).

While we expect the two external instruments to exhibit low correlation, Bruns et al. (2024) highlight that even weakly correlated instruments can introduce unintended correlations between identified shocks, distorting impulse response estimates. To mitigate this risk, we impose structured correlation restrictions that ensure each instrument is more strongly correlated with its intended shock than with the alternative, reinforcing identification. This approach enhances robustness in Proxy SVAR frameworks, where explicitly accounting for correlation is crucial.

A central innovation of our paper is the application of the McKay and Wolf (2023) framework for constructing monetary policy counterfactuals in a way that is robust to the Lucas critique. In particular, rather than relying on a parametric structural model, we use empirical estimates of policy shock transmission to infer alternative policy paths. This allows us to quantify the trade-offs policymakers face when responding to oil shocks, determining how different policy priorities—such as greater emphasis on inflation stabilization versus output stabilization—affect economic outcomes.

We have three main findings. First, we show that, consistent with many studies, both oil supply shocks and monetary policy shocks significantly influence inflation and economic output, with their relative importance shifting over time. Second, counterfactual analyses indicate that although monetary policy can reduce output losses, it cannot fully neutralize the inflationary pressures created by oil price shocks. This underscores the persistent influence of oil price movements on macroeconomic outcomes, even under optimal policy interventions, challenging

the conclusions of Bernanke et al. (1997). In particular, by highlighting the intrinsic stagflationary dynamics of oil supply news shocks, our analysis underscores the need for a nuanced understanding of how oil supply shocks interact with broader macroeconomic conditions, regardless of monetary policy intervention. Third, our analysis underscores the trade-offs central banks face when responding to oil price shocks. Prioritizing inflation control leads to tighter policy and greater output losses, while emphasizing output stabilization results in more accommodative policies but higher inflation. In fact, we show that baseline policy path closely mirrors a scenario where inflation is weighted 2.5 times more than output, reflecting a pragmatic balance between these competing objectives in the face of oil-driven disruptions.

Related Literature. This paper contributes to the empirical literature on the role of monetary policy in transmitting oil price shocks to the real economy, c.f. Bernanke et al. (1997) for a seminal paper showing that systematic monetary policy responses were a primary driver of recessions following major oil price shocks and Hamilton and Herrera (2004), Herrera and Pesavento (2009), Kilian and Lewis (2011) and Bjørnland et al. (2018) that challenge this view. More recent analysis, such as Miyamoto et al. (2024), have focused on the impact of oil supply shocks on the macroeconomy during the zero lower bound period, while Castelnuovo et al. (2024) study the transmission mechanism for various commodity market shocks. More related is the paper by Gagliardone and Gertler (2023), that analyze the trade-offs in monetary policy responses using a calibrated quantitative New-Keynesian model. We complement Gagliardone and Gertler (2023) by analysing the inflationary pressures of oil price shocks and the trade-offs in monetary policy responses using a proxy SVAR model. Our framework allows for an empirical assessment of the relative impacts of oil and monetary shocks over time, as well as the study of optimal monetary policy using an empirical model.

From a methodological viewpoint, our study relates to the literature that uses multiple external instruments to identify different structural shocks within SVAR frameworks. This line of work began with Mertens and Ravn (2013), who introduced a framework for multi-shock identification, assuming a recursive ordering of the two shocks identified. Piffer and Podstawski (2018) extend this approach by introducing the idea of multiple contaminated instruments, but worked in a

frequentist setting. Arias et al. (2021) and Giacomini et al. (2022) extend the multi-shock identification framework to Bayesian methods. Our study builds on Piffer and Podstawski (2018) by employing a frequentist approach, but augmented to allow for two instruments to distinctly identify oil supply and monetary policy shocks. Subsequently, the use of counterfactuals is also new to this paper.

Finally, in applying multiple high-frequency identified shocks as external instruments, this paper relates to the high-frequency identification literature, see e.g. Kuttner (2001), Gürkaynak et al. (2005), Gertler and Karadi (2015), Altavilla et al. (2019), Miranda-Agrippino and Ricco (2021), Acosta (2022) and Bauer and Swanson (2023a) for examples of how to identify monetary policy shocks and Känzig (2021), Degasperi (2023) and Miyamoto et al. (2024) for oil supply shocks. We contribute by extending the monetary policy instrument related to Bauer and Swanson’s framework to include data through the COVID period and by robustifying the instrument by accounting for co-movements between oil price changes and monetary policy shocks.

The paper is organized as follows. Section 2 outlines the econometric framework, including the identification and construction of the shocks. Section 3 presents our main empirical results regarding the effects of the oil and monetary policy shocks, while Section 4 studies optimal policy responses. Section 5 presents several robustness checks before we conclude in Section 6.

2 The Econometric Framework

The vector autoregression (VAR) model of order p can be expressed in its reduced-form representation as:

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}\mathbf{x}_t + \mathbf{u}_t, \tag{1}$$

where \mathbf{y}_t is an $n \times 1$ vector of endogenous variables, \mathbf{b} is an $n \times 1$ vector of intercepts, and $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_p)$ represents the coefficient matrices, with $\mathbf{B}_i = \mathbf{A}_0^{-1}\mathbf{A}_i$ for $i = 1, \dots, p$. Here, \mathbf{A}_0 is an invertible $n \times n$ matrix with positive diagonal elements, and \mathbf{A}_i are $n \times n$ matrices capturing the autoregressive dynamics

of the model. The vector $\mathbf{x}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$ aggregates the lagged values of the endogenous variables, and the reduced-form shocks \mathbf{u}_t are related to the structural shocks $\boldsymbol{\epsilon}_t$ via the linear mapping:

$$\mathbf{u}_t = \mathbf{S}\boldsymbol{\epsilon}_t. \quad (2)$$

The $n \times n$ impact matrix $\mathbf{S} = \mathbf{A}_0^{-1}$ is invertible and can be partitioned as $\mathbf{S} = [\mathbf{S}_1, \mathbf{S}_2]$, where \mathbf{S}_1 is an $n \times k$ matrix corresponding to the structural shocks of interest, and \mathbf{S}_2 is an $n \times (n - k)$ matrix corresponding to the remaining shocks. The structural shocks $\boldsymbol{\epsilon}_t$ are assumed to be uncorrelated, satisfying:

$$\mathbb{E}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t') = \mathbf{I}_n, \quad \mathbb{E}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_s') = \mathbf{0}_n \quad \text{for } t \neq s. \quad (3)$$

2.1 Identification via External Instruments

In the context of Proxy Structural VARs (Proxy SVARs), external instruments $\mathbf{m}_t = (m_{1t}, \dots, m_{kt})'$ are used to identify a subset of $k \leq n$ structural shocks, corresponding to the \mathbf{S}_1 matrix. These instruments must satisfy relevance and exogeneity conditions.

The structural shocks can be decomposed as:

$$\boldsymbol{\epsilon}_t = \begin{bmatrix} \boldsymbol{\epsilon}'_{1t} & \boldsymbol{\epsilon}'_{2t} \end{bmatrix}', \quad (4)$$

where $\boldsymbol{\epsilon}_{1t}$ is a k -dimensional vector representing the structural shocks of interest, and $\boldsymbol{\epsilon}_{2t}$ is an $(n - k)$ -dimensional vector representing the remaining shocks. Correspondingly, the $n \times n$ covariance matrix $\boldsymbol{\Phi}$ is decomposed as:

$$\boldsymbol{\Phi} = \begin{bmatrix} \boldsymbol{\Phi}_1 & \boldsymbol{\Phi}_2 \end{bmatrix}, \quad (5)$$

where $\boldsymbol{\Phi}_1$ is a $k \times k$ covariance matrix capturing the correlation between the k instrumental variables and the structural shocks of interest, and $\boldsymbol{\Phi}_2$ is an $(n - k) \times k$ matrix of correlations with the remaining shocks. The goal is to identify the structural shocks $\boldsymbol{\epsilon}_{1t}$ using the instrumental variables in \mathbf{m}_t . Given that $\mathbb{E}(\mathbf{m}_t \boldsymbol{\epsilon}_t') = \boldsymbol{\Phi}$, the combination of Equations (4) and (5) yields:

$$\mathbb{E}(\mathbf{m}_t \boldsymbol{\epsilon}'_{1t}) = \boldsymbol{\Phi}_1, \quad \mathbb{E}(\mathbf{m}_t \boldsymbol{\epsilon}'_{2t}) = \boldsymbol{\Phi}_2. \quad (6)$$

For an external instrument \mathbf{m}_t to serve as a valid proxy for the target shock, two key conditions must hold: relevance, which requires $\boldsymbol{\Phi}_1 \neq 0$, and exogeneity with respect to other structural shocks, which implies $\boldsymbol{\Phi}_2 = 0$. These conditions, as established by Stock and Watson (2012) and Mertens and Ravn (2013), enable the distinction between structural shocks correlated with the proxies and those uncorrelated. When identifying multiple shocks ($k > 1$), additional restrictions are necessary beyond these conditions.

2.2 Identification Restrictions

To identify both the monetary policy shock and the oil supply news shock simultaneously, we impose sign restrictions on the correlation structure between the structural shocks of interest and the external instruments. Specifically, we require that the proxy for the monetary policy shock is more strongly correlated with the monetary policy shock than with the oil supply news shock. Conversely, the proxy for the oil supply news shock must be more strongly correlated with the oil supply news shock than with the monetary policy shock. In doing so, we build on the broader Proxy-SVAR literature by demonstrating how this joint identification enables the construction of robust counterfactual scenarios, following, e.g., Sims and Zha (2006) and McKay and Wolf (2023). While similar sign and correlation restrictions have recently been used in other applications by Piffer and Podstawski (2018), Arias et al. (2021), and Giacomini et al. (2022), the joint identification of oil supply shocks and monetary policy shocks using external instruments, as well as the use of counterfactuals within Proxy-SVARs, is novel to this paper.

Let $\mathbf{m}_t = (m_t^o, m_t^m)'$ represent the proxy variables that are correlated with the k structural shocks of interest, $\boldsymbol{\epsilon}_{1t} = (\boldsymbol{\epsilon}_t^o, \boldsymbol{\epsilon}_t^m)'$. We can rewrite Equation (2) as:

$$\mathbf{u}_t = \mathbf{s}^o \boldsymbol{\epsilon}_t^o + \mathbf{s}^m \boldsymbol{\epsilon}_t^m + \tilde{\mathbf{S}} \tilde{\boldsymbol{\epsilon}}_t, \quad (7)$$

where \mathbf{s}^o represents the impulse vector associated with the oil supply shock, \mathbf{s}^m corresponds to the impulse vector related to the monetary policy shock, and $\tilde{\mathbf{S}}$

includes the impulse vectors of the remaining structural shocks.

The $k \times k$ covariance matrix Φ_1 , which contains the correlation coefficients¹ between the proxies and the first two structural shocks, is given by:

$$\begin{bmatrix} \mathbb{E}(\varepsilon_t^o m_t^o) & \mathbb{E}(\varepsilon_t^o m_t^m) \\ \mathbb{E}(\varepsilon_t^m m_t^o) & \mathbb{E}(\varepsilon_t^m m_t^m) \end{bmatrix} = \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix}. \quad (8)$$

We impose additional sign restrictions necessary for exact identification of the structural shocks on the subcomponents of Φ_1 , specifically:

$$\begin{aligned} \phi_{11} > 0 \quad , \quad \phi_{22} > 0 \\ \phi_{11} - \phi_{21} > \varphi \quad , \quad \phi_{22} - \phi_{12} > \varphi. \end{aligned} \quad (9)$$

The first set of restrictions ($\phi_{11} > 0$ and $\phi_{22} > 0$) ensures that the oil supply news proxy and the monetary policy proxy are positively correlated with their respective target shocks. We require these correlations to be statistically significant from zero at the 1% level. The second set of restrictions ensures that each proxy is more strongly correlated with its intended shock than with the alternative shock. To further reinforce this distinction, we impose a threshold φ , which must be exceeded. In our baseline calibration, φ is set at 0.10; however, we also assess the robustness of our results by varying φ within the range of 0 to 0.25.

While the external instruments used for identification exhibit a low but nonzero correlation our structured correlation restrictions ensure that each instrument is more strongly correlated with its intended shock than with the alternative, reinforcing identification.²

¹Since structural shocks are assumed to be mean zero and have unit variance, expressing restrictions in terms of correlations or covariances yields the same identification conditions. Additionally, we ensure that instruments are mean zero and standardized to unit variance for consistency.

²This idea aligns with recent findings by Bruns et al. (2024), who highlight that even weakly correlated instruments can lead to unintended correlation between identified shocks, distorting impulse response estimates.

2.3 Data

We collect monthly data from January 1984 to May 2023. The choice of 1984 as the starting point restricts the analysis to the post-Volcker disinflation period, avoiding the changes in oil price elasticities observed before this time, as discussed in Baumeister and Peersman (2013a, 2013b). Our model includes six macroeconomic variables in the Proxy-SVAR: the log of the real oil price³, the Federal Funds Rate, the log of U.S. Industrial Production, the log of the U.S. Consumer Price Index (CPI), the log of the Commodity Price Index, and the Excess Bond Premium following Gilchrist and Zakrajšek (2012). These variables are selected to closely match those used by Bernanke et al. (1997), ensuring comparability with their work. We assess the robustness of our results by including additional variables, such as wages, employment, and the S&P 500, among others.

A detailed description of the data and the respective sources can be found in Table 2 in Appendix A.

2.4 Monetary policy instrument

To construct the monetary policy surprise series, we follow the methodology of Bauer and Swanson (2023a) but extend their analysis by adding data from December 2019 to May 2023, and by orthogonalizing the shock of interest with respect to a larger set of macroeconomic variables. In total, we compile 351 FOMC announcements covering the period from 1988 to 2023. Various financial instruments have been used in the literature to construct monetary policy surprises. For instance, Kuttner (2001) uses the change in the current-month federal funds futures contract, while Gertler and Karadi (2015) and others use federal funds futures at different maturities. Gürkaynak et al. (2005) and Bauer and Swanson (2023a) use Eurodollar futures contracts. We follow the latter approach, using the first four quarterly Eurodollar futures contracts, ED1-ED4.

We consider a 30-minute event window around the FOMC announcements, starting 10 minutes before and ending 20 minutes after the press release, to identify exogenous variations in the prices of the Eurodollar futures contracts. Assuming

³The real oil price is calculated by deflating the WTI spot price with the U.S. Consumer Price Index.

that the risk premium does not change within this narrow window, the monetary policy surprise is given by

$$Surprise_t^h = \mathbb{E}_{t+20}(P_{t+h}) - \mathbb{E}_{t-10}(P_{t+h}) \quad (10)$$

where $Surprise_t^h$ represents the monetary policy surprise for announcement t for the $h = \{1, \dots, 4\}$ quarters ahead Eurodollar futures contract. The expected price of the futures contract h given the information 20 minutes after the announcement is defined as $\mathbb{E}_{t+20}(P_{t+h})$, while $\mathbb{E}_{t-10}(P_{t+h})$ is the expected price of the same futures contract given the information 10 minutes prior to the press release. In a subsequent step, we extract the first principal component of these surprises for each announcement t and rescale it so that a one-unit change in the principal component corresponds to a one percentage point change in the ED4 rate. Capturing fluctuations in the first four quarterly Eurodollar futures, this shock measure provides insight into both the current federal funds rate target and forward guidance regarding the future policy path.

Recent studies have documented that high-frequency monetary policy surprises contain a component that can be predicted to some extent using economic and financial data. Miranda-Agrippino and Ricco (2021) provide evidence of a Fed Information Effect by utilizing the Fed’s internal Greenbook forecasts to evaluate the predictability of these high-frequency surprises. Similarly, Bauer and Swanson (2023b) show that publicly available Blue Chip forecasts also contain predictable information regarding these monetary policy surprises. They further document this predictability using a set of macroeconomic and financial variables. Their findings suggest that this predictability is not solely a result of a Fed Information Effect but rather due to the correlation of these surprises with publicly available macroeconomic and financial data prior to FOMC announcements.

We follow Bauer and Swanson (2023a) and use a similar set of variables to construct our orthogonal measure of high-frequency monetary policy surprises. The variables⁴ used are: *Nonfarm payroll surprises* (calculated as the actual payroll growth minus the median expectation from the most recent Money Market Services

⁴We refer to Bauer and Swanson (2023b) and Bauer and Chernov (2024) for a detailed description and explanation of the construction and choice of the set of variables.

survey prior to the FOMC announcement), *Employment growth* (measured as the log change in nonfarm payroll employment from one year earlier), *S&P 500* (change in the stock price index over the previous 65 trading days), *Yield curve slope* (the change in the slope of the Treasury yield curve, measured as the second principal component of the one- to ten-year zero-coupon Treasury yields), *Commodity Price Index* (calculated as the log change in the Bloomberg Commodity Spot Price index during the previous 65 trading days) and *Treasury skewness* (following Bauer and Chernov (2024))⁵. In addition, we include *Oil Prices* (measuring the change in the log spot oil price over the previous 25 trading days) into the data set. This will strengthen the exogeneity of our monetary policy instrument by ensuring that high-frequency interest rate surprises are not systematically predictable based on oil price movements.⁶ The orthogonalized high-frequency instrument can then be constructed as

$$Surprise_t = \alpha + \beta'Y_t + \varepsilon_t. \quad (11)$$

The surprise component, calculated as outlined in Equation 10, represents our monetary policy instrument. Here, β denotes the estimated regression coefficients, while Y_t encompasses the set of predictors previously described. The residual of the regression, ε_t , serves as our monetary policy surprise instrument, which is orthogonal to the macroeconomic and financial variables considered.

2.5 Oil supply instrument

In contrast to the monetary policy instrument that we construct ourselves, the oil supply shock series is identical to Känzig (2021).⁷ It is constructed measuring unexpected fluctuations in oil futures prices that occur on the day of Organization of the Petroleum Exporting Countries (OPEC) meetings. The relevant event window for capturing these oil supply shocks is from the day of the announcement to the last trading day before the OPEC meeting. The sample period for the oil

⁵As this measure is only available until May 2021, we have also constructed an orthogonal measure without using this financial variable, and the results remain robust. However, we include it here to maintain comparability to the original monetary policy surprise measure.

⁶While the original correlation between the monetary policy and oil supply instruments is low (less than 0.1), inclusion of oil prices into the data set does reduce the correlation between the two instruments somewhat, ultimately bringing the correlation below 0.05.

⁷The updated oil supply news shock series was retrieved from the author's webpage.

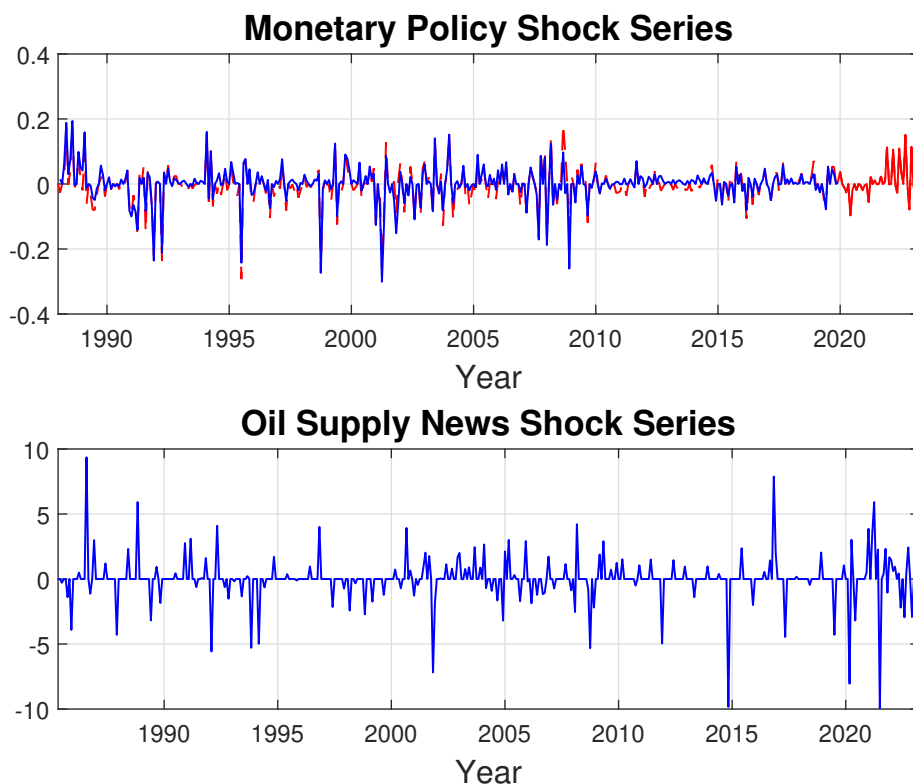


Figure 1: Historical Shock Series Plots

Notes: The upper panel plots two series, in blue the monetary policy shock series from Bauer and Swanson (2023a) and in red our extended and modified orthogonal monetary policy shock series. The lower panel displays the oil supply news series taken from Känzig (2021), where recent updates was retrieved from the author’s website.

supply news shock series is from January 1985 to May 2023, capturing a total of 150 OPEC press releases.⁸

Figure 1 display the two instruments. The upper panel plots the original monetary policy instrument from Bauer and Swanson (2023a) in blue and the modified and extended instrument we construct in red, while the oil supply instrument taken from Känzig (2021) is plotted in the lower panel.⁹ The monetary policy shock series demonstrates periods of higher volatility, particularly around major monetary

⁸For a more detailed description of how the instrument is constructed, see Känzig (2021).

⁹The correlation between the original and updated monetary policy shock is 0.98, confirming that our modifications preserve the core information content of the instrument.

policy events, while the extended and adjusted series in red aligns closely with the original instrument but provides additional information for recent periods.¹⁰ The lower panel illustrates the oil supply news shocks, capturing significant fluctuations during notable oil market events such as geopolitical tensions or global supply disruptions, c.f. Känzig (2021).

3 The effect of oil supply news and monetary policy shocks

Below we provide some details about the VAR estimation and instrument validity, before examining the effects of the oil supply news shocks and the monetary policy shocks through impulse responses and historical decomposition.

3.1 VAR estimation and instrument validity

The Proxy-SVAR is estimated at a monthly frequency over the sample period from March 1984 to April 2023. We include 12 lags and a constant in the estimation. Parameter uncertainty is simulated using the wild bootstrap method. To control for the sequence of extreme observations during the COVID-19 pandemic, especially throughout the first months of the pandemic, we follow Lenza and Primiceri (2022) and include dummies for March, April, and May 2020.¹¹

When external instruments are weakly correlated with the shocks they aim to capture, standard inference in SVARs may be unreliable, necessitating the use of weak-instrument robust inference methods. Instrument strength is typically evaluated using the F-statistic from the first-stage regression, where the residuals

¹⁰The instrument captures notable monetary policy events, such as the 2008 Financial Crisis, during which the Federal Reserve implemented rapid rate cuts and unconventional policies like quantitative easing, reflected in heightened volatility around 2008-2009 in the blue series. Similarly, the extended red series shows distinct spikes around 2020, corresponding to the Federal Reserve's emergency monetary responses to the COVID-19 pandemic.

¹¹While Lenza and Primiceri (2022) model heteroskedasticity in a Bayesian framework, our approach follows a frequentist estimation strategy, incorporating dummy variables for the initial months of the pandemic. This aligns with the logic of the Pandemic Prior approach but within a frequentist setting.

of the target variable from the VAR—in this case, the oil price and federal funds rate—are regressed on the external instrument.

We formally assess instrument strength in Appendix Table 1, which reports the first-stage regression results. Both the oil supply news and monetary policy instruments exhibit first-stage F-statistics exceeding 10 for their respective target variables, a threshold widely used in the literature to mitigate concerns about weak identification.¹² These results suggest that weak-instrument concerns are unlikely to be a significant issue in our setting.

3.2 Impulse responses to oil and monetary policy shocks

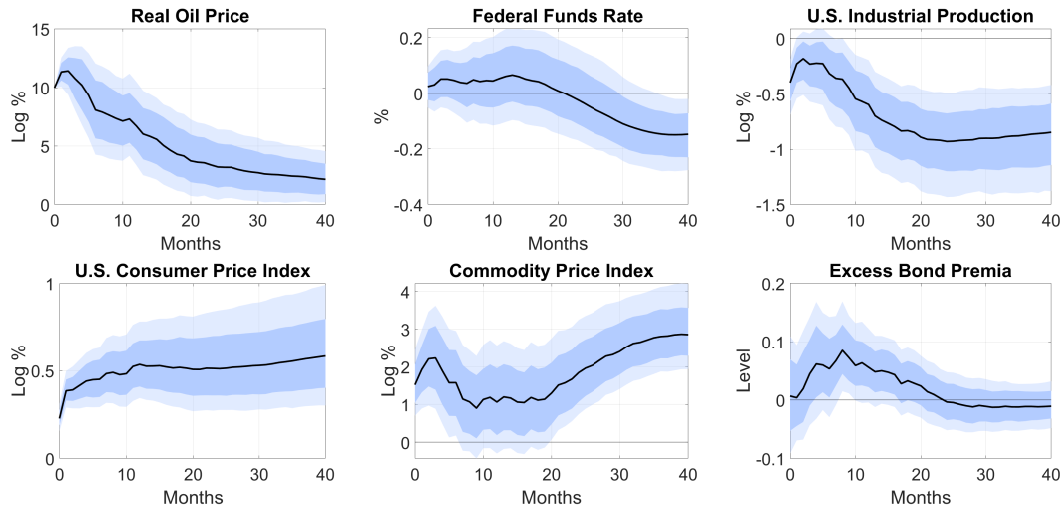
We analyse the effects of the shocks through impulse responses, before focusing on historical decompositions. The oil supply news shock is scaled to increase the real price of oil by 10%, while the monetary policy shock raises the federal funds rate by 25 basis points on impact. The corresponding impulse response functions (IRFs) and confidence intervals for both shocks are reported in Figure 2.

We first examine a contractionary oil supply news shock, which is constructed to increase the real price of oil. The impulse responses suggest a stagflationary effect on the economy. This shock leads to a gradual decline in industrial production by 1% at its peak, occurring around 20 months after the shock. Additionally, U.S. consumer prices increase significantly, with a peak increase of around 0.5% after 10 months. The federal funds rate rises over the first 20 months, indicating that the central bank tightens monetary policy in response to rising inflation. These findings align with Gagliardone and Gertler (2023), who report similar but stronger responses in the tightening of the federal funds rate following an oil supply shock. The increase in real oil prices leads to a persistent rise in the commodity price index over the whole horizon, peaking at around 3% before gradually reverting to the steady state. The IRFs are similar to those in Känzig (2021), with minor differences due to varying variable selections.¹³

¹²See Olea et al. (2020) and references therein for a discussion on the threshold level and the implications of weak instrument inference.

¹³Some recent studies, i.e., Bjørnland and Skretting (2024) and Hanson and Herrera (2024) have reported more positive responses in manufacturing production or job reallocation to oil shocks during the shale oil boom, in particular in oil producing states. Our findings suggest that the impact on aggregate U.S. industrial production during our full sample still remains negative.

(a) Oil Supply News Shock



(b) Monetary Policy Shock

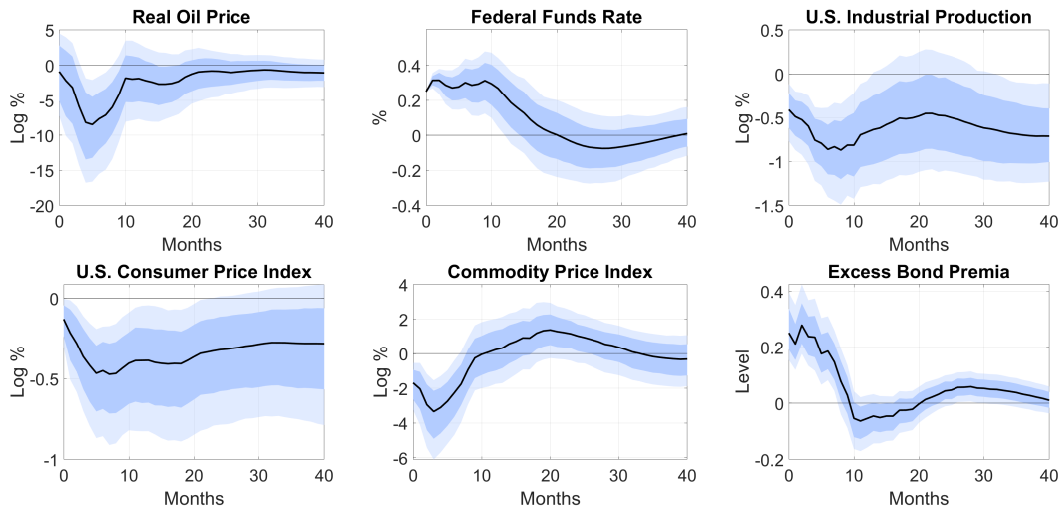


Figure 2: Impulse Responses to oil supply news and monetary policy shocks

Notes: The to oil supply news and monetary policy shocks, are normalized to increase oil prices by 10% and the federal funds rate by 25 bps, respectively. The solid line represents the median estimate, with dark and light-shaded areas indicating the 68% and 95% confidence intervals, respectively, obtained via the wild bootstrap method.

The monetary policy shock, normalized to increase the federal funds rate by 25 basis points on impact, leads to a decline in industrial production, with a peak

reduction of 0.8% after two years. The price level also decreases, reaching a maximum decline of 0.5% following the policy tightening, before gradually returning to the steady state. Real oil prices and commodity prices fall immediately after the shock, but they gradually revert to their initial levels over the first two years. Additionally, the excess bond premium rises by approximately 25 basis points on impact. Our findings broadly align with those of Bauer and Swanson (2023a), although slight differences arise due to our variable selection and sample period. We report somewhat larger effects, in particular for industrial production, than typically found in high-frequency monetary policy shock studies, such as Gertler and Karadi (2015), which rely on non-orthogonalized shock measures. Our larger estimates primarily reflect our use of orthogonalized monetary policy shocks, which remove biases associated with endogenous policy responses, thus yielding stronger estimated effects.¹⁴

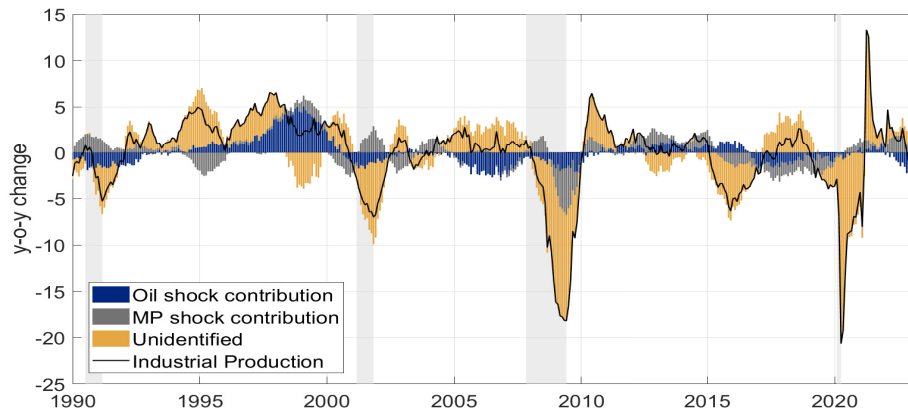
3.3 Historical Decomposition

To assess the influence of monetary policy and oil supply shocks on fluctuations in industrial production and consumer prices over the full sample period, we compute the historical decomposition, presented in Figure 3. Overall, the contribution of oil supply shocks to industrial production (upper panel) has remained relatively stable over time but has played a more prominent role in certain periods. Notably, oil shocks had a negative impact during the periods leading up to the global financial crisis and prior to the covid pandemic. We also see that periods of contractionary monetary policy, such as the rate hikes in the early 2000s and 2022, coincide with declines in industrial output, highlighting the sector’s sensitivity to policy adjustments.

Turning to inflation in the lower panel of Figure 3, we see that oil shocks consistently contribute to increase inflation over time. Notably, during the recent inflation surge, oil shocks accounted for a large share of the increase, peaking in early 2022. As oil prices declined afterward, the contribution of these shocks to inflation similarly receded, highlighting the responsiveness of prices to changes in

¹⁴Orthogonalization generally produces larger estimates by explicitly removing biases due to monetary policy responding endogenously to economic conditions (Bauer and Swanson (2023a), see Figure 3 and related discussion; see also Miranda-Agrippino and Ricco (2021)).

(a) Industrial Production



(b) Consumer Price Index

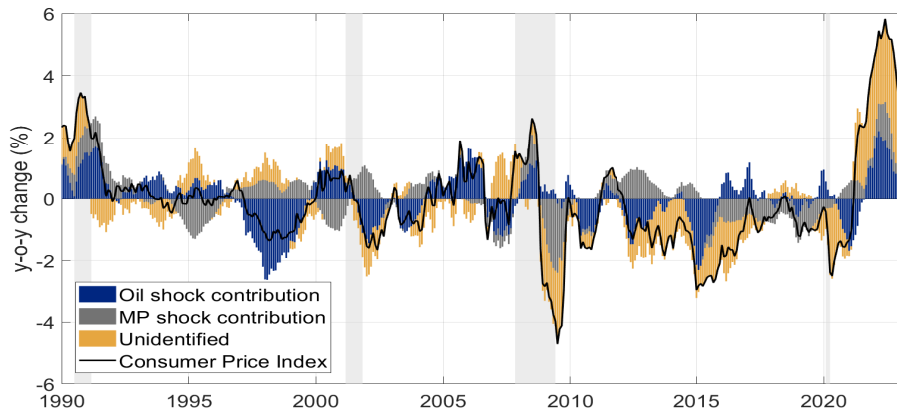


Figure 3: Historical decomposition of U.S. IP (top) and U.S. CPI (bottom)

Notes: The historical decompositions for Industrial Production and the Consumer Price Index due to monetary policy (MP) and oil shocks. The decomposition is computed from the stochastic component of the estimated model. The variables enter the model in logs, but for presentation, we report their year-over-year changes. The decompositions are expressed as deviations from their respective sample means and rescaled accordingly. Gray bars indicate NBER recession periods.

oil markets. Monetary policy shocks have also played a significant role in driving changes in consumer prices, particularly in the most recent period. The Federal Reserve's initially accommodative stance, characterized by stable interest rates, contributed positively to the rise in inflation. However, the subsequent interest

rate hikes from 2022 began exerting a negative influence on inflation from 2023.¹⁵

4 Monetary policy responses

Having examined the effects of oil supply news shocks and monetary policy shocks, we now turn to analyzing policy responses through counterfactual scenarios. Our approach builds on the methodologies of Sims and Zha (2006) and McKay and Wolf (2023). Finally, to assess optimal policy trade-offs, we build on the framework of McKay and Wolf (2023), focusing on the balance between inflation control and output stabilization in response to oil-driven disturbances.¹⁶

4.1 Counterfactual scenarios

We first conduct a counterfactual scenario analysis in which the Federal Funds Rate remains unchanged in response to oil price shocks, allowing us to isolate the direct economic effects of oil price shocks without monetary policy adjustments. To comprehensively assess the role of systematic monetary policy, we construct two counterfactual scenarios based on alternative methodologies: the first follows Sims and Zha (2006), and the second builds on the framework proposed by McKay and Wolf (2023).

The Sims and Zha (2006) approach introduces a sequence of monetary policy shocks that neutralize the effect of oil shocks on the Federal Funds Rate period-by-period.¹⁷ Because these adjustments are imposed sequentially and *ex post*, private-sector expectations are not aligned *ex ante* with the counterfactual policy regime, raising concerns highlighted by the Lucas Critique, *c.f.*, Lucas (1976).

In contrast, the methodology proposed by McKay and Wolf (2023) constructs policy counterfactuals through a linear combination of multiple monetary policy

¹⁵Bergholt et al. (2024) have argued that historical decomposition can be poorly identified as the deterministic component of VAR models tends to be imprecisely estimated. To check robustness to this, we follow one of their recommendations and use demeaned data and re-estimate our SVAR model. Our historical decomposition remains qualitatively robust, though monetary policy effects on inflation appear more pronounced, in particular post-GFC.

¹⁶A similar approach to assess optimal policy responses has been proposed by Barnichon and Mesters (2023), but our analysis is more closely aligned with the methodology of McKay and Wolf.

¹⁷A detailed description of this method is provided in Appendix B.

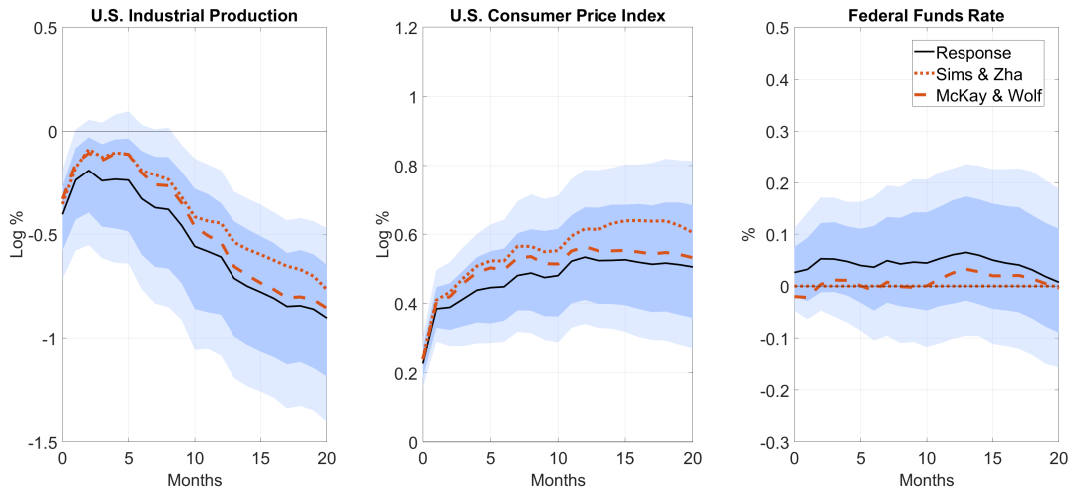


Figure 4: Counterfactual Impulse Responses

Notes: Counterfactual impulse responses to oil supply news shocks, under counterfactual Federal Funds Rate paths. The solid black line represents the median estimate, with dark and light-shaded areas indicating the 68% and 95% confidence intervals, respectively, obtained via the wild bootstrap method. The counterfactuals from the Sims & Zha and the McKay & Wolf methods are in the solid red line and dashed red line, respectively.

shocks implemented simultaneously at time 0, rather than sequentially across periods. The key rationale for employing multiple shocks is that systematic monetary policy responses typically require simultaneous adjustments of interest rates at multiple future horizons. A single monetary policy shock generally cannot achieve distinct adjustments at several future points in time. By combining monetary policy shocks with different persistence and dynamic response patterns, we can closely replicate, at the initial date, the full systematic monetary policy response across multiple horizons. Crucially, this simultaneous adjustment aligns private-sector expectations immediately with the counterfactual policy regime, directly addressing the concerns raised by the Lucas Critique, which arise when adjustments are implemented sequentially and ex post.

Specifically, we combine two distinct monetary policy shocks to approximate this systematic policy response. First, we use our orthogonalized monetary policy shock (described in Section 2.4), which captures immediate adjustments and forward-guidance components. Second, we employ the monetary policy shock from Romer and Romer (2004), which primarily captures short-term discretionary pol-

icy adjustments.¹⁸ By combining shocks that differ in persistence and dynamic implications, we can replicate the systematic monetary policy response closely at the initial date, aligning private-sector expectations immediately with the revised policy stance and thus addressing the concerns highlighted by the Lucas (1976).¹⁹

The counterfactual impulse responses are presented in Figure 4. While some minor differences arise between the two methodologies (Sims and Zha: red dotted; McKay and Wolf: red dashed), these differences are small and not statistically significant. Both approaches indicate that oil price shocks induce stagflationary effects even in the absence of a systematic monetary policy response. Specifically, the contraction in industrial production is slightly mitigated—by approximately 0.10 percentage points on average over the first 12 months—while inflationary pressures are amplified, with a peak increase of about 0.15 percentage points after 15 months relative to the baseline.

These findings underscore that both oil price and monetary policy shocks play a significant role in shaping inflation and economic activity. In line with Gagliardone and Gertler (2023), who examine the role of monetary policy in recent inflationary cycles using a DSGE framework, our results suggest that while monetary policy can dampen the stagflationary effects of oil supply shocks, it cannot fully eliminate inflationary pressures. This highlights the structural constraints faced by central banks in responding to oil-driven inflation, particularly when attempting to balance price stability with output stabilization.

4.2 Optimal Policy Response

The trade-off identified in our counterfactual analysis—where refraining from monetary policy intervention mitigates output declines but exacerbates inflationary pressures—underscores a key challenge for policymakers. Designing an optimal response requires weighing the relative importance of inflation stability versus output stabilization, particularly in the face of supply-side shocks. A stronger focus on inflation control necessitates a more contractionary stance, which push output fur-

¹⁸The extended Romer and Romer shock series was retrieved from Acosta (2022) and is available on the author’s website.

¹⁹See Appendix B.2 for additional methodological details, including the derivation of shock weights (Equation 28).

ther down, while prioritizing output stabilization calls for a more accommodative approach, prolonging inflationary pressures.

To construct an optimal policy response we follow McKay and Wolf (2023) who base the responses on policy counterfactuals. In particular, they demonstrate that the effects of alternative policy rules can be inferred directly from empirical estimates of policy shock transmission, without requiring a fully specified structural model. Applying their method, we identify the sequence of policy interventions that best aligns with a policymaker's objectives, ensuring the counterfactual response remains robust to the Lucas critique.

A central result in McKay and Wolf (2023) is that optimal policy can be derived by considering how different types of policy shocks influence the economy. Their framework highlights that, if the effects of both contemporaneous and anticipated policy shocks are known, it is possible to construct counterfactual policy paths that replicate the intended stabilization objectives. Following this logic, we characterize the optimal policy response to an oil supply shock by selecting the combination of policy interventions that achieves the best trade-off between inflation and output stabilization. Rather than estimating a rule-based policy function, we compute the required sequence of policy shocks that implement the desired outcome, ensuring consistency with the empirical evidence on policy transmission.²⁰ In practice, this method allows us to assess how different policy priorities shape the response to oil supply-side disturbances. A more contractionary stance, for example, reduces inflation persistence but amplifies output declines, whereas a more accommodative approach mitigates output losses at the cost of higher inflation.

Figure 5 illustrates the resulting optimal policy paths for a 10% oil price shock under three different weighting scenarios. The first assigns equal importance to output and inflation ("Equal weighted", dashed-dotted green), the second prioritizes inflation stabilization ("Inflation targeting", dashed red), and the third emphasizes output stabilization ("Output targeting", dotted red). The optimal policy response varies significantly depending on the relative weights assigned to output and inflation stabilization. First, when the policymaker assigns equal importance to output and inflation, the policy entails a moderate easing of the federal funds rate—on the order of 7 basis points—which helps cushion the decline in industrial

²⁰See Appendix B.3 for a detailed description of the method.

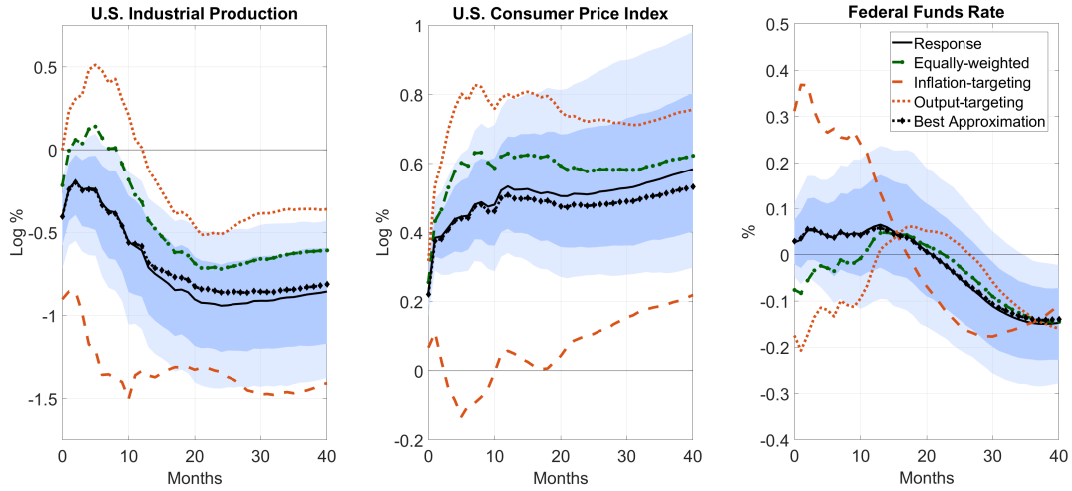


Figure 5: Optimal Policy Response Paths

Notes: The figure illustrates optimal policy responses to a 10% oil price shock under three loss function objectives: equally-weighted (dashed-dotted green), inflation-targeting (dashed red), and output-targeting (dotted red). The solid black line is the median estimate from the main model specification, with shaded areas indicating 68% (dark blue) and 95% (light blue) confidence intervals from the wild bootstrap method. The “Best Approximation” (dashed black line with markers) described in the text, identifies the weighting scenario that most closely replicates the baseline response.

production but allows for slightly higher inflation in the short run compared to the baseline.

In contrast, prioritizing inflation stabilization (inflation targeting) results in a more forceful tightening of over 35 basis points; although it nearly offsets the immediate inflationary effect of the oil shock, industrial production contracts by nearly twice as much as in the equal-weighted scenario. Finally, when the emphasis is placed on output stabilization (output targeting), the policy response features a sizable rate cut of about 20 basis points, mitigating the loss in industrial production but significantly amplifying inflationary pressures over a longer horizon.

These differences underscore the trade-offs central banks face when responding to oil price shocks. Prioritizing inflation stabilization entails deeper output losses, whereas emphasizing output stabilization intensifies inflation. Notably, the peak deviation in the policy rate across the scenarios exceeds 60 basis points, reflecting the range of outcomes driven by different policy priorities.

Finally, we construct the path of monetary policy that is closest to the actual monetary policy response. We denote this “Best Approximation” in Figure 5. Using a grid-based approach,²¹ we estimate that the Federal Reserve’s historical response aligns most closely with a policy placing 2.5 times more weight on inflation stabilization than on output stabilization.

5 Robustness Analysis

We perform three types of robustness tests, analysing sensitivity to (i) the sample period, (ii) the choice of instruments used to identify the oil shocks and (iii) the model specification.

First, we exclude the COVID-19 period from our analysis. Although our baseline model includes dummy variables to control for the extreme outliers during this period, we assess whether our results are influenced by these anomalies. As shown as the red line in Figure 7, the impulse responses to both shocks are closely aligned with our main results. The main difference is that the effect of a monetary policy shock is slightly more muted, and that the policy response begins to ease sooner following an oil supply news shock. Similarly, the counterfactual scenario in Figure 9 are closely aligned with our main results (see red versus black line). The main difference is that the responses decline slightly faster.

Second, we examine the sensitivity of our results to alternative measures of oil supply shocks. Specifically, we replace our baseline shock measure with the Degasper (2023) series, which refines the Känzig measure by filtering out demand-driven components using asset price comovements. Furthermore, we employ the Baumeister and Hamilton (2019) approach, which constructs oil supply shocks based on a Bayesian VAR with global production, oil prices, inventories, and other macroeconomic indicators. The results, depicted by the green and orange lines in Figures 7 and 9, confirm the robustness of our main conclusions. However, the Degasper (2023) measure suggests a stronger initial tightening of monetary policy in response to an oil price shock, leading to a sharper decline in industrial production from the oil shock, whereas the Baumeister and Hamilton (2019) approach

²¹A more detailed surface analysis illustrating how different configurations of these weights affect the interest rate response is provided in Figure 6 in Appendix B.3.

yields slightly muted responses.

Third, we assess the robustness of our results to modifications in model specifications, including alternative monetary policy indicators and additional macroeconomic controls. We find that replacing the Federal Funds Rate with the 1-year Treasury Yield Rate does not alter our conclusions (see the red line in Figure 8). Likewise, the inclusion of additional macroeconomic variables, such as the Unemployment Rate, the S&P 500 and the Exchange Rate (blue, orange and gray lines, respectively), leaves our main results unaffected. Finally, substituting the Consumer Price Index (CPI) with the Personal Consumption Expenditure (PCE) measure as our inflation indicator does not materially change our findings.

6 Conclusion

In this paper, we reexamine the relationship between oil price fluctuations and monetary policy, highlighting the nuanced roles of both oil shocks and monetary policy interventions in driving macroeconomic fluctuations. Our analysis shows that while monetary policy can help dampen some effects of oil price shocks, it is often insufficient to fully counter their inflationary pressures, especially during periods of heightened volatility such as experience during the Covid-19 pandemic. Using a dual-instrument proxy SVAR model, we identify distinct impacts of oil supply and monetary policy shocks, finding that their relative importance shifts over time. Our counterfactual analysis underscores a critical trade-off for central banks: prioritizing output stability can mitigate recessionary impacts but may amplify inflationary pressures, complicating policy decisions in the face of oil-driven shocks. Finally, we estimate that the Federal Reserve's historical response aligns closely with a policy that places twice as much weight on inflation stabilization than on output stabilization. These findings suggest that central banks must balance the need to stabilize output versus inflation objectives when responding to shocks in the oil market.

References

- Acosta, M. (2022). The perceived causes of monetary policy surprises. *Working Paper, Columbia University*.
- Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179.
- Angelini, G., Cavaliere, G., and Fanelli, L. (2024). An identification and testing strategy for proxy-svars with weak proxies. *Journal of Econometrics*, 238(2):105604.
- Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. (2021). Inference in Bayesian Proxy-SVARs. *Journal of Econometrics*, 225(1):88–106.
- Barnichon, R. and Mesters, G. (2023). A Sufficient Statistics Approach for Macro Policy. *American Economic Review*, 113(11):2809–45.
- Bauer, M. D. and Chernov, M. (2024). Interest Rate Skewness and Biased Beliefs. *The Journal of Finance*, 79(1):173–217.
- Bauer, M. D. and Swanson, E. T. (2023a). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*, 37(1):87–155.
- Bauer, M. D. and Swanson, E. T. (2023b). An Alternative Explanation for the ‘Fed Information Effect’. *American Economic Review*, 113(3):664–700.
- Baumeister, C. and Hamilton, J. D. (2019). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *The 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.
- Bergholt, D., Canova, F., Furlanetto, F., Maffei-Faccioli, N., and Ulvedal, P. (2024). What drives the recent surge in inflation? the historical decomposition roller coaster. *Norges Bank Working Paper 07/2024*.
- Bernanke, B., Gertler, M., and Watson, M. (1997). Systematic Monetary Policy and the Effects of Oil Price Shocks. *Brookings Papers on Economic Activity*, 28(1):91–157.
- Bjørnland, H. C., Larsen, V. H., and Maih, J. (2018). Oil and macroeconomic (in) stability. *American Economic Journal: Macroeconomics*, 10(4):128–151.
- Bjørnland, H. C. and Skretting, J. (2024). The Shale Oil Boom and The US Economy: A Time-Varying Analysis. *Journal of Applied Econometrics*, 39(6):1000–1020.
- Braun, R. and Brüggemann, R. (2022). Identification of SVAR Models by Combining Sign Restrictions With External Instruments. *Journal of Business Economic Statistics*, 41:1–38.
- Bruns, M., Lütkepohl, H., and McNeil, J. (2024). Avoiding Unintentionally Correlated Shocks in Proxy Vector Autoregressive Analysis. " Available at SSRN 4927320.
- Castelnuovo, E., Mori, L., and Peersman, G. (2024). Commodity Price Shocks and Global Cycles: Monetary Policy Matters. CAMA Working Papers 2024-36, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- Degasperi, R. (2023). Identification of Expectational Shocks in the Oil Market using OPEC Announcements. The Warwick Economics Research Paper Series (TWERPS) 1464, University of Warwick, Department of Economics.

- Gagliardone, L. and Gertler, M. (2023). Oil Prices, Monetary Policy and Inflation Surges. NBER Working Papers 31263, National Bureau of Economic Research, Inc.
- Gertler, M. and Karadi, P. (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.
- Giacomini, R., Kitagawa, T., and Read, M. (2022). Robust Bayesian inference in proxy SVARs. *Journal of Econometrics*, 228(1):107–126.
- Gilchrist, S. and Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4):1692–1720.
- Gürkaynak, R. S., Sack, B., and Swanson, E. (2005). The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models. *American Economic Review*, 95(1):425–436.
- Hamilton, J. D. and Herrera, A. M. (2004). Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment. *Journal of Money, Credit and Banking*, 36(2):265–286.
- Hanson, R. and Herrera, A. M. (2024). The Effect of Oil News Shocks on Job Creation and Destruction. mimeo, University of Kentucky.
- Herrera, A. M. and Pesavento, E. (2009). Oil price shocks, systematic monetary policy, and the “Great Moderation”. *Macroeconomic Dynamics*, 13(1):107–137.
- Kilian, L. and Lewis, L. T. (2011). Does the Fed Respond to Oil Price Shocks? *Economic Journal*, 121(555):1047–1072.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.
- Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review*, 111(4):1092–1125.
- Lenza, M. and Primiceri, G. E. (2022). How to estimate a vector autoregression after March 2020. *Journal of Applied Econometrics*, 37(4):688–699.

- Lucas, R. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1(1):19–46.
- McKay, A. and Wolf, C. K. (2023). What Can Time-Series Regressions Tell Us About Policy Counterfactuals? *Econometrica*, 91(5):1695–1725.
- 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(4):1212–47.
- Miranda-Agrippino, S. and Ricco, G. (2021). The Transmission of Monetary Policy Shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107.
- Miyamoto, W., Nguyen, T., and Sergeyev, D. (2024). How Oil Shocks Propagate: Evidence on the Monetary Policy Channel. CEPR Discussion Paper 18755, CEPR Press, Paris & London.
- Olea, J., Stock, J., and Watson, M. (2020). Inference in Structural Vector Autoregressions identified with an external instrument. *Journal of Econometrics*, 225.
- Piffer, M. and Podstawski, M. (2018). Identifying Uncertainty Shocks Using the Price of Gold. *The Economic Journal*, 128(616):3266–3284.
- Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, 94(4):1055–1084.
- Sims, C. and Zha, T. (2006). Does monetary policy generate recessions? *Macroeconomic Dynamics*, 10(2):231–272.
- Stock, J. H. and Watson, M. (2012). Disentangling the Channels of the 2007-09 Recession. *Brookings Papers on Economic Activity*, 43(1 (Spring)):81–156.

—Appendix—

A Data

A.1 Instrument Strength

	Real Oil Price	Fed Funds Rate	Industrial Production	Consumer Price Index	Commodity Price Index	Excess Bond Premia
Oil Price Instrument						
β	0.009***	0.008*	-0.001	0.001***	0.001**	-0.011*
F	17.349	3.181	0.433	10.469	4.684	2.867
R ²	0.036	0.007	0.001	0.022	0.0101	0.006
Monetary Policy Instrument						
β	0.017	0.422***	-0.008	-0.001	-0.022	0.386**
F	0.056	10.216	1.469	0.808	1.309	4.577
R ²	0.001	0.022	0.003	0.002	0.003	0.009

Notes: The models estimated are $\hat{u}_t = \alpha + \beta_i m_t^i + \eta_t$, where \hat{u}_t is the residual from the VAR equation corresponding to each variable in the table, and m_t^i is the proxy for either the oil supply news shock or the monetary policy shock. The null hypothesis tests whether $\beta_i = 0$. Statistical significance is based on the asymptotic distribution of the OLS estimator.

Table 1: Tests on the Strength of the Instruments

A.2 Data Description

Variable	Description	Source
Baseline variables		
Real Oil Price	WTI spot crude oil price (WTISPLC) deflated by U.S. CPI (CPIAUCSL)	FRED
Fed Funds Rate	Federal Funds Effective Rate (FEDFUNDS)	FRED
U.S. Industrial Production	U.S. Industrial Production: Total Index (INDPRO)	FRED
U.S. CPI	U.S. Consumer Price Index for All Urban Consumers (CPIAUCSL)	FRED
Commodity Price Index	U.S. Spot market price index: All commodities	Commodity Research Bureau (CRB)
Excess Bond Premia	Excess bond premia from Gilchrist and Zakrajšek (2012)	Federal Reserve

Table 2: Data Description and Data Sources

B Methodology

B.1 Counterfactual construction following Sims & Zha

To analyze the role of systematic monetary policy in the transmission of oil supply news shocks, we construct a counterfactual scenario in which the Federal Funds Rate (FFR) remains unchanged following an oil supply news shock. This approach follows the methodology of Sims and Zha (2006), which neutralizes the monetary policy response by applying a sequence of exogenous monetary policy shocks at each horizon.

The reduced-form **VAR representation** of our model in Equation 1 can be stated as:

$$\mathbf{y}_t = \mathbf{b} + \sum_{i=1}^p \mathbf{B}_i \mathbf{y}_{t-i} + \mathbf{u}_t, \quad (12)$$

while the **companion form** of the VAR in Equation (12) is given by:

$$\mathbf{Z}_t = \mathbf{F} \mathbf{Z}_{t-1} + \mathbf{v}_t, \quad (13)$$

where:

$$\mathbf{Z}_t = \begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \vdots \\ \mathbf{y}_{t-p+1} \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{B}_2 & \dots & \mathbf{B}_{p-1} & \mathbf{B}_p \\ \mathbf{I} & 0 & \dots & 0 & 0 \\ 0 & \mathbf{I} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \mathbf{I} & 0 \end{bmatrix}, \quad \mathbf{v}_t = \begin{bmatrix} \mathbf{u}_t \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

The **impulse response functions (IRFs)** to a structural shock are computed recursively as:

$$\frac{\partial \mathbf{Z}_{t+s}}{\partial \boldsymbol{\varepsilon}'_t} = \frac{\partial \mathbf{Z}_{t+s}}{\partial \mathbf{u}'_t} \cdot \frac{\partial \mathbf{u}_t}{\partial \boldsymbol{\varepsilon}'_t} = \Psi_s \mathbf{A}_0^{-1} = \Theta_s, \quad (14)$$

where Ψ_s is the (non-orthogonalized) IRF matrix at horizon s , obtained using the first n rows of \mathbf{F}^s , with $\Psi_s = \mathbf{e}'_1 \mathbf{F}^s \mathbf{e}_1$.

Following an oil supply news shock, the counterfactual impulse response function must neutralize the impact on the Federal Funds Rate at each impulse horizon s . Given the variable ordering in our model, where oil prices are ordered first and the Federal Funds Rate is the second variable, we define a selection vector e_{FFR} , which extracts the interest rate component:

$$\theta_{FFR,j,s} = e'_{FFR} \Theta_s \boldsymbol{\varepsilon}_t. \quad (15)$$

To construct the counterfactual, we introduce a sequence of monetary policy shocks $\boldsymbol{\varepsilon}_t^{m,cf}$ that offsets the impact of the oil supply news shock on interest rates:

$$e'_{FFR} \left(\Theta_s \boldsymbol{\varepsilon}_t + \Theta_s \boldsymbol{\varepsilon}_t^{m,cf} \right) = 0. \quad (16)$$

Solving for $\boldsymbol{\varepsilon}_t^{m,cf}$, we obtain:

$$\boldsymbol{\varepsilon}_t^{m,cf} = - (e'_{FFR} \Theta_s e_{FFR})^{-1} e'_{FFR} \Theta_s \boldsymbol{\varepsilon}_t. \quad (17)$$

The remaining sequence of counterfactual shocks is computed recursively:

$$\boldsymbol{\varepsilon}_{t+s}^{m,cf} = \frac{\theta_{FFR,j,s} + \sum_{m=0}^{S-1} e'_{FFR} \mathbf{F}^m \mathbf{A}_0^{-1} e_{FFR} \boldsymbol{\varepsilon}_{t+m}^{m,cf}}{e'_{FFR} \mathbf{A}_0^{-1} e_{FFR}}. \quad (18)$$

Thus, the counterfactual impulse response function for any variable i to the oil supply news shock is:

$$\theta_{i,j,s}^{cf} = \theta_{i,j,s} + \sum_{m=0}^{S-1} e'_i \mathbf{F}^m \mathbf{A}_0^{-1} e_{FFR} \boldsymbol{\varepsilon}_{t+m}^{m,cf}. \quad (19)$$

B.2 Counterfactual construction following McKay & Wolf

The second method we use to construct counterfactuals follows the approach of McKay and Wolf (2023). Generally, assume that the economy can be described by the following system of equations:

$$\mathbf{H}_w \mathbf{w} + \mathbf{H}_x \mathbf{x} + \mathbf{H}_z \mathbf{z} + \mathbf{H}_\varepsilon \varepsilon = 0 \quad (20)$$

$$\mathbf{A}_x \mathbf{x} + \mathbf{A}_z \mathbf{z} + \boldsymbol{\nu} = 0 \quad (21)$$

where \mathbf{w} and \mathbf{x} are vectors of endogenous variables, with the crucial distinction that the variables in \mathbf{x} are observable to the econometrician—representing the macroeconomic variables of interest—while \mathbf{w} is unobservable. The vector \mathbf{z} contains the policy instrument, in our case the Federal Funds Rate, while ε represents the exogenous structural non-policy shock, specifically the oil supply shock in our model. The linear mappings $\{\mathbf{H}_w, \mathbf{H}_x, \mathbf{H}_z, \mathbf{H}_\varepsilon\}$ describe the non-policy block of the economy, independent of the baseline policy rule in Equation 21, where $\{\mathbf{A}_x, \mathbf{A}_z\}$ are the coefficients of the policy rule and $\boldsymbol{\nu}$ contains the full set of policy-shocks to the prevailing rule across all horizons. To derive and construct counterfactual impulse response functions of the economy under alternative policy functions, it is necessary that the policy rule in Equation 21 induces a unique equilibrium. The counterfactual policy rule under an alternative policy specification can then be expressed as:

$$\tilde{\mathbf{A}}_x \mathbf{x}_{\tilde{\mathbf{A}}} + \tilde{\mathbf{A}}_z \mathbf{z}_{\tilde{\mathbf{A}}} = 0, \quad (22)$$

where $\tilde{\mathbf{A}}_x$ and $\tilde{\mathbf{A}}_z$ are the coefficients associated with the alternative policy rule, and $\mathbf{x}_{\tilde{\mathbf{A}}}(\varepsilon)$ and $\mathbf{z}_{\tilde{\mathbf{A}}}(\varepsilon)$ represent the dynamic responses under this new rule to the non-policy shock ε , which we seek to construct. It is also required that the alternative policy rule in Equation 22 induces a unique equilibrium.

The construction of counterfactual scenarios under alternative policy rules requires two key inputs:

- The dynamic causal effects of the *non-policy* shock ε under the baseline rule, denoted as $\{\mathbf{x}_A(\varepsilon), \mathbf{z}_A(\varepsilon)\}$
- The dynamic causal effects of the *policy* shocks $\boldsymbol{\nu}$ on the policy instrument

\mathbf{z} and the macroeconomic observables of interest \mathbf{x}

The impulse responses under the baseline policy rule can be expressed as:

$$\Theta_A \equiv \begin{bmatrix} \Omega_{\mathbf{x},\varepsilon,A} & \Omega_{\mathbf{x},\nu,A} \\ \Omega_{\mathbf{z},\varepsilon,A} & \Omega_{\mathbf{z},\nu,A} \end{bmatrix}, \quad (23)$$

where $\Omega_{\mathbf{x},\varepsilon,A}$ and $\Omega_{\mathbf{z},\varepsilon,A}$ contain the impulse responses of macroeconomic variables of interest and the policy instrument to the non-policy shock ε , while $\Omega_{\mathbf{x},\nu,A}$ and $\Omega_{\mathbf{z},\nu,A}$ contain the responses to the policy shocks ν .

Given that the econometrician knows the policy shock causal effects $\{\Omega_{\mathbf{x},\nu,A}, \Omega_{\mathbf{z},\nu,A}\}$ and the response of the variables $\{\mathbf{x}_A(\varepsilon), \mathbf{z}_A(\varepsilon)\}$ to the shock ε , the counterfactual policy responses $\{\mathbf{x}_{\tilde{A}}(\varepsilon), \mathbf{z}_{\tilde{A}}(\varepsilon)\}$ under a counterfactual policy rule $\{\tilde{A}_x, \tilde{A}_z\}$ can be recovered as:

$$\mathbf{x}_{\tilde{A}}(\varepsilon) = \mathbf{x}_A(\varepsilon, \tilde{\nu}) \equiv \mathbf{x}_A(\varepsilon) + \Omega_{\mathbf{x},\nu,A} \times \tilde{\nu} \quad (24)$$

$$\mathbf{z}_{\tilde{A}}(\varepsilon) = \mathbf{z}_A(\varepsilon, \tilde{\nu}) \equiv \mathbf{z}_A(\varepsilon) + \Omega_{\mathbf{z},\nu,A} \times \tilde{\nu} \quad (25)$$

where $\tilde{\nu}$ solves the system:

$$\tilde{A}_x[\mathbf{x}_A(\varepsilon) + \Omega_{\mathbf{x},\nu,A} \times \tilde{\nu}] + \tilde{A}_z[\mathbf{z}_A(\varepsilon) + \Omega_{\mathbf{z},\nu,A} \times \tilde{\nu}] = 0. \quad (26)$$

Thus, given the complete set of policy shock perturbations, it is always possible to construct a date-0 policy shock vector ν that replicates the equilibrium path of \mathbf{z} under the counterfactual policy rule. Since the non-policy block of the economy is influenced by monetary policy solely through the expected path of the policy instrument, the resulting equilibrium allocations under the counterfactual rule align exactly with those obtained under the appropriately adjusted baseline rule.

However, in practice, an exact replication of the equilibrium path under the counterfactual rule is not always feasible. The lower-dimensional structure of the identified policy shocks typically renders the system of equations governing the counterfactual adjustment underdetermined or infeasible. As a result, researchers must approximate the counterfactual rule as closely as possible by selecting optimal weights ν_w on the date-0 identified policy shocks in ν . The natural approach is to solve an optimization problem that minimizes the deviation from the desired

counterfactual path while remaining within the span of the empirically identified policy shocks:

$$\min_{\nu_w} \left\| \tilde{A}_x(\mathbf{x}_A(\epsilon) + \Omega_{x,\nu,A} \times \nu_w) + \tilde{A}_z(\mathbf{z}_A(\epsilon) + \Omega_{z,\nu,A} \times \nu_w) \right\|. \quad (27)$$

The solution of the minimization problem is given as:

$$\begin{aligned} \nu_w = & - \left[(\tilde{A}_x \Omega_{x,\nu,A} + \tilde{A}_z \Omega_{z,\nu,A})' (\tilde{A}_x \Omega_{x,\nu,A} + \tilde{A}_z \Omega_{z,\nu,A}) \right]^{-1} \\ & \times \left[(\tilde{A}_x \Omega_{x,\nu,A} + \tilde{A}_z \Omega_{z,\nu,A})' (\tilde{A}_x \mathbf{x}_A(\epsilon) + \tilde{A}_z \mathbf{z}_A(\epsilon)) \right]. \end{aligned} \quad (28)$$

This solution allows for the construction of counterfactual impulse responses that approximate the intended alternative policy counterfactual as closely as possible, given the available set of identified policy shock paths. Since all shocks are dated at $t = 0$, the constructed counterfactual paths remain robust to the Lucas critique. In general, the accuracy of the counterfactual approximation improves as the number of identified policy paths increases. As we are interested in offsetting the Federal Funds Rate response, we implement a rule²² that can be thought of as similar to rate targeting. In our empirical application, \mathbf{x} includes the variables of interest, U.S. CPI and U.S. industrial production, while \mathbf{z} represents the Federal Funds Rate as the policy instrument. The non-policy shock, ϵ , corresponds to the oil supply news shock, see Section 2.5, whereas ν consists of the two monetary policy shocks described in Section 4.1.

²²This depends on the application the econometrician has in mind; other options include the implementation of a Taylor-type rule, output targeting, or other policy rules.

B.3 Optimal policy construction following McKay & Wolf

In this section, we outline the methodology used to derive optimal policy responses, following the notation and approach of McKay and Wolf (2023). While Barnichon and Mesters (2023) develop a similar framework for optimal policy assessment, our analysis is more closely aligned with McKay and Wolf (2023), particularly in focusing on the trade-off between inflation and output stabilization. The central idea is to determine the policy rule that minimizes a given policymaker's loss function, subject to the constraints imposed by the economic system. Assume that the policymaker seeks to minimize the following quadratic loss function:

$$\mathcal{L} = \frac{1}{2} \sum_{i=1}^{n_x} \lambda_i \mathbf{x}_i' W \mathbf{x}_i \quad (29)$$

where λ_i represents the weights assigned to different policy objectives, \mathbf{x}_i contains the macroeconomic variables of interest and $W = \text{diag}(1, \beta, \beta^2, \dots)$ allows for discounting. A necessary condition for solving the policymaker's optimization problem is that, given any non-policy exogenous shock $\boldsymbol{\varepsilon}$, there exists a unique solution for choosing the policy variable \mathbf{z} to minimize the loss function, subject to the non-policy constraint in Equation 20.

With this condition in place, the goal is to derive the dynamic response paths $\mathbf{x}_A(\boldsymbol{\varepsilon}), \mathbf{z}_A(\boldsymbol{\varepsilon})$ under the optimal policy rule $\mathbf{A}_x, \mathbf{A}_z$, which takes the form:

$$\mathbf{A}_x^* = (\lambda_1 \boldsymbol{\Omega}'_{x_1, \nu, A} \mathbf{W}, \lambda_2 \boldsymbol{\Omega}'_{x_2, \nu, A} \mathbf{W}, \dots, \lambda_{n_x} \boldsymbol{\Omega}'_{x_{n_x}, \nu, A} \mathbf{W}) \quad (30)$$

$$\mathbf{A}_z^* = \mathbf{0} \quad (31)$$

With these policy coefficients, the counterfactual impulse response paths under the optimal policy rule can be characterized as in Proposition 1²³, following the structure of Equation 24:

$$\mathbf{x} = \mathbf{x}(\boldsymbol{\varepsilon}) + \boldsymbol{\Omega}_{x, \nu, A} \times \boldsymbol{\nu}^* \quad (32)$$

where the minimization problem of the loss function in Equation 29 leads to

²³For the proof of equivalence, see McKay and Wolf (2023).

the optimality condition:

$$\sum_{i=1}^{n_x} \lambda_i \Omega_{x_i, \nu, A} \mathbf{W} \mathbf{x}_i = 0 \quad (33)$$

Thus, the optimal policy counterfactual is given by:

$$\boldsymbol{\nu}^* = - \left[\Omega'_{x, \nu, A} \boldsymbol{\lambda} \mathbf{W} \Omega_{x, \nu, A} \right]^{-1} \times \left[\Omega'_{x, \nu, A} \boldsymbol{\lambda} \mathbf{W} \mathbf{x}_A(\boldsymbol{\varepsilon}) \right]. \quad (34)$$

The same information that enables the construction of valid counterfactuals for arbitrary policy rules also allows for the characterization of optimal policy rules. The intuition remains unchanged: since we have knowledge of the causal effects of all possible policy perturbations $\boldsymbol{\nu}$ on the policymaker's target variables \mathbf{x} , we can determine the full set of outcomes that can be achieved through policy actions. Optimal policy corresponds to the point within this implementable set that minimizes the policymaker's loss function. As before, whether this optimum is reached via a systematic policy rule or through shocks to an alternative rule is inconsequential.

B.4 Optimal policy - surface analysis

This section presents a surface analysis of optimal policy responses based on a grid search over the policymaker’s preferences for inflation and output stabilization. The procedure uses impulse response functions (IRFs) estimated from the baseline model as inputs, capturing the historical dynamics of a 10% oil price shock and 25 bps increase in the federal funds rate. The optimal policy is computed by minimizing a weighted loss function, where weights on inflation and output deviations are systematically varied. The three-dimensional surfaces in Figure 6 illustrate how U.S. Consumer Prices (left), U.S. Industrial Production (center), and the Federal Funds Rate (right) evolve over time (horizontal axis) as preferences shift from inflation prioritization to output prioritization (front-to-back axis). Warmer colors indicate larger responses, while cooler colors signal smaller responses.

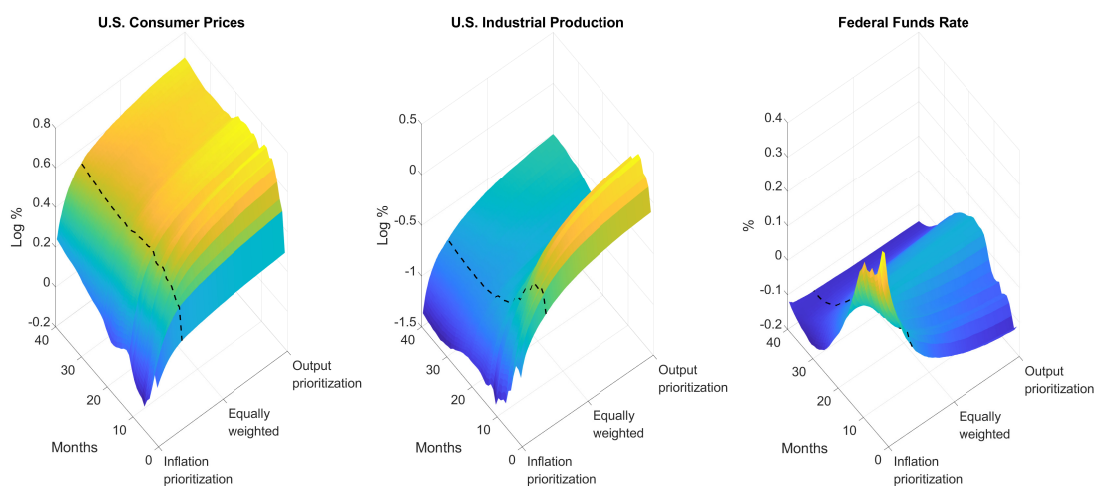
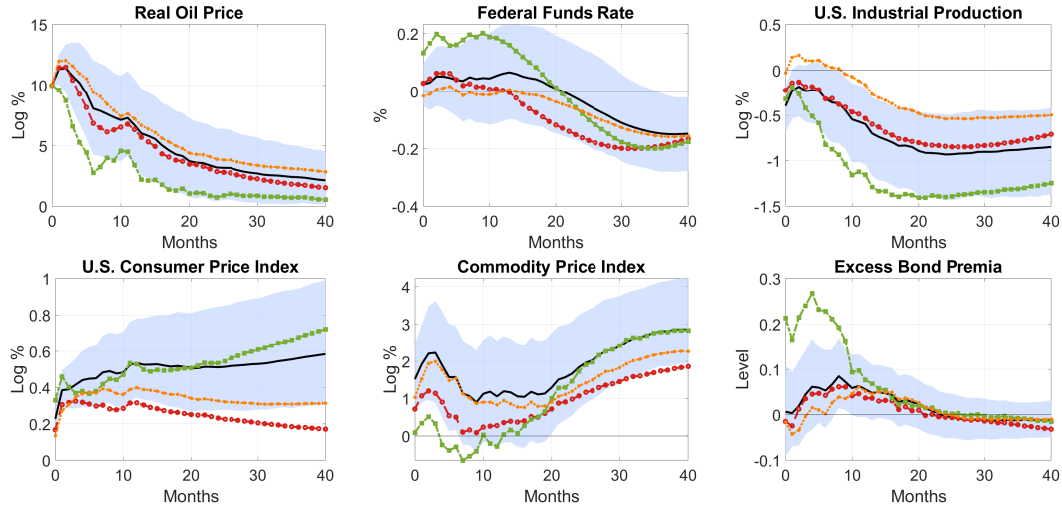


Figure 6: Surface Plot: Optimal Policy Response Paths

Notes: These three-dimensional surfaces show how the dynamic responses of U.S. Consumer Prices (left), U.S. Industrial Production (center), and the Federal Funds Rate (right) evolve over time (horizontal axis) as the policy-maker’s preferences shift from inflation prioritization to output prioritization (front-to-back axis). Warmer colors (i.e., yellow) indicate larger responses, while cooler colors (i.e., blue) signal smaller responses. The black-dashed line corresponds to the “best approximation” path discussed in the main text, reflecting the weighting of inflation and output stabilization that most closely replicates the baseline response.

C Robustness Analysis

(a) Oil Supply News Shock



(b) Monetary Policy Shock

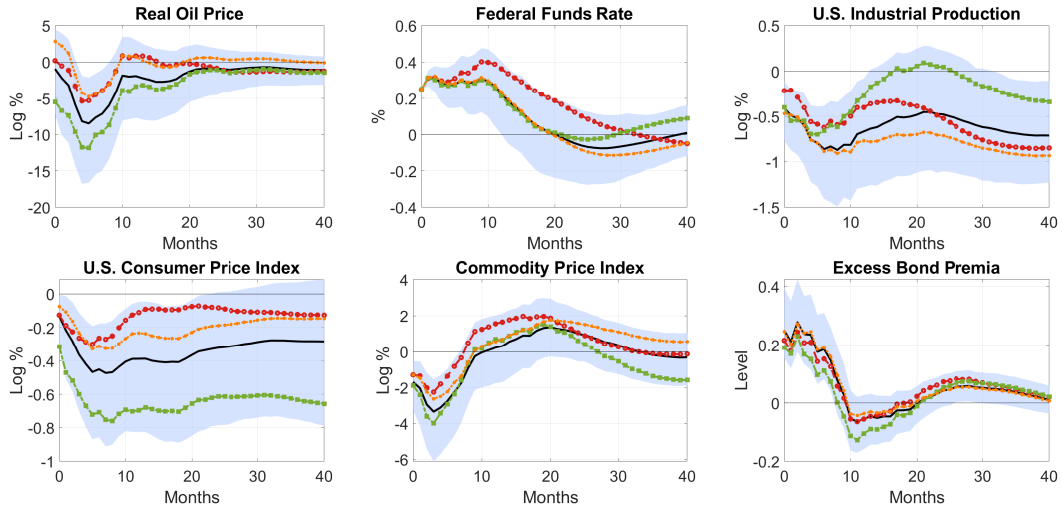
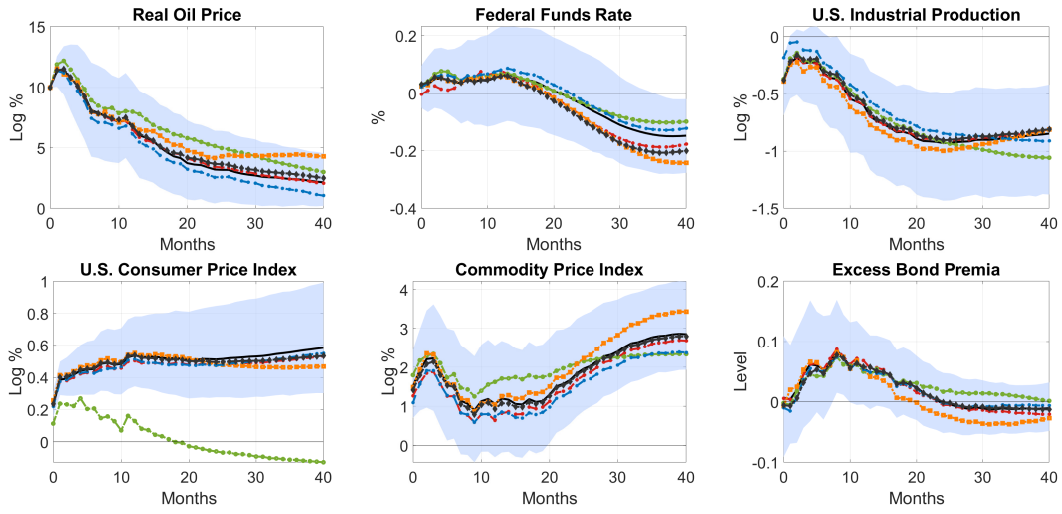


Figure 7: Impulse Responses to oil supply news and monetary policy shocks

Notes: The median estimates and the 95% confidence intervals from the main model specification are represented by the solid (black) and the shaded (blue) area, respectively. The results from the robustness analyses are shown as follows: (1) excluding COVID (red dashed with circles); (2) using the Degaspari oil supply measure (green dash-dotted with squares); and (3) using the Baumeister & Hamilton supply measure (orange dotted with crosses).

(a) Oil Supply News Shock



(b) Monetary Policy Shock

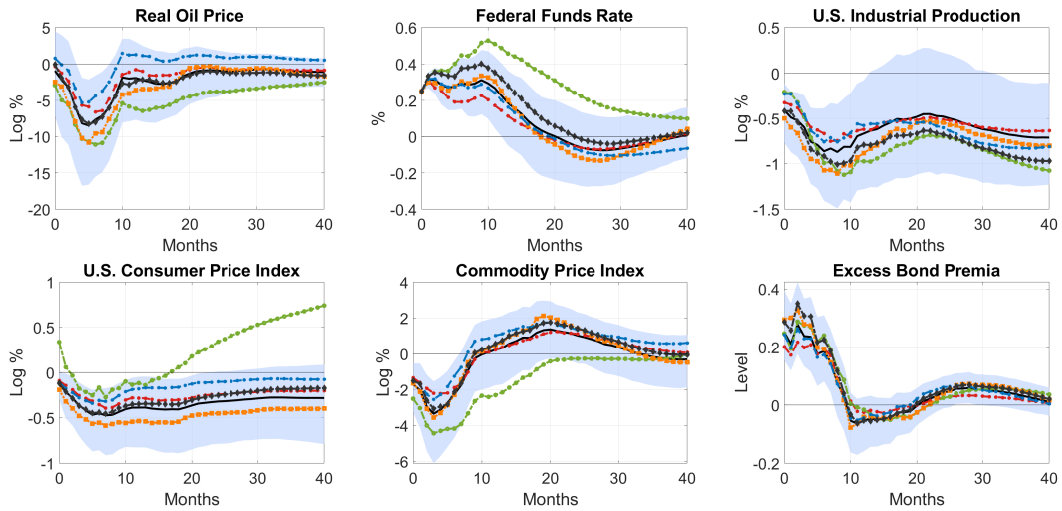


Figure 8: Impulse Responses to oil supply news and monetary policy shocks

Notes: The median estimates and the 95% confidence intervals from the main model specification are represented by the solid (black) and the shaded (blue) area, respectively. The results from the robustness analyses including different and additional series are shown as follows: (1) using the 1-year Treasury Yield instead of the Federal Funds Rate (red dashed with stars); (2) using the PCE instead if the CPI (green dashed-dotted with circles); (3) additionally including the S&P500 (orange dotted with squares), (4) including the Unemployment Rate (blue dashed with crosses) and (5) including the Exchange Rate (gray dashed-dotted with diamonds). Note, the labeling of the plots has not been changed for robustness test (1) and (2), to remain consistent throughout the analysis.

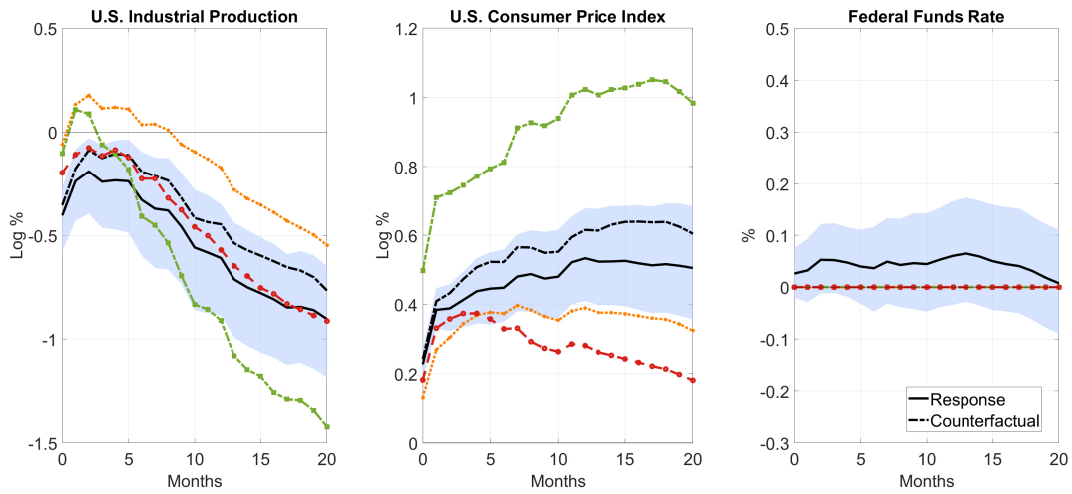


Figure 9: Counterfactual Impulse Responses

Notes: Counterfactual impulse responses to oil supply news shocks, under counterfactual Federal Funds Rate paths. The median estimates and the 95% confidence intervals from the main model specification are represented by the solid (black) and the shaded (blue) area, respectively. The results from the robustness analyses are shown as follows: (1) excluding COVID (red dashed with circles); (2) using the Degaspari oil supply measure (green dash-dotted with squares); and (3) using the Baumeister & Hamilton supply measure (orange dotted with crosses).