Friend, Not Foe? Monetary Policy and Energy Prices *

Gökhan Ider† Alexander Kriwoluzky‡ Frederik Kurcz§ Ben Schumann ¶

May 31, 2025

Abstract

In this paper, we first document a new empirical finding by showing that the European Central Bank's (ECB) monetary policy decisions significantly influence global energy prices. Through Lucas-critique-robust counterfactual analysis, we then empirically study the implications of this result for the transmission of monetary policy. Our findings reveal that a central bank's ability to affect energy prices strengthens and accelerates the monetary transmission to inflation dynamics, and alleviates the inflation-output trade-off. We illustrate the relevance of these results by examining their role in the optimal policy response to an energy supply shock. Our estimates show that monetary policy's ability to affect global energy prices effectively halves the necessary tightening to stabilize inflation and the corresponding economic contraction, compared to a scenario in which energy prices are unaffected by monetary policy.

Keywords: inflation, energy prices, monetary policy transmission mechanism
JEL Codes: C32, E31, E52, Q43

*We thank Rüdiger Bachmann, Michael Bauer, Christiane Baumeister, Andrea Gazzani, Georgios Georgiadis, Refet Gürkaynak, Klodiana Istrefi, Yoon Joo Jo (discussant), Silvia Miranda-Agrippino, Gisle Natvik, Pascal Paul, Ricardo Reis, Giovanni Ricco, Fabian Seyrich, Christian Wolf and Leopold Zessner-Spitzenberg as well as seminar participants at the Kiel-CEPR Conference 2025 (*Monetary Policy After the Inflation Surge - What Have We Learned?*), EEA Annual Congress 2023, Oslo Macro Conference 2023, University of Vienna, the European Central Bank, the Humboldt University and DIW Berlin for helpful comments. We are especially grateful to Refet Gürkaynak and Ambrogio Cesa-Bianchi for sharing monetary policy surprise data with us. We also thank Antonia Vogel for excellent research assistance. This research has received financial support from the Leibniz Association through the project "Distributional effects of macroeconomic policies in Europe" and by the Deutsche Forschungsgemeinschaft (BE 5381/1-1).

†DIW Berlin, Mohrenstrasse 58, 10117 Berlin, gider@diw.de

‡Corresponding author: DIW Berlin, Mohrenstrasse 58, 10117 Berlin, Germany, and Freie Universität Berlin, Garystraße 21, 14195 Berlin, Germany, akriwoluzky@diw.de, +49 30 89789 430.

§DIW Berlin, Mohrenstrasse 58, 10117 Berlin, fkurcz@diw.de

¶DIW Berlin, Mohrenstrasse 58, 10117 Berlin, bschumann@diw.de

1 Introduction

After decades of low inflation, the global economy experienced a sharp surge in inflation during 2021–2023. It is widely recognized that energy prices were a major driver of this inflation episode (Bernanke and Blanchard, 2024), sparking renewed interest in the relationship between monetary policy and commodity markets—particularly energy prices. Among the economies most affected by the surge in energy prices was the euro area. As one of the world's largest economies and energy importers (see Figure 1), the euro area is highly exposed to fluctuations in commodity prices. This high degree of exposure and economic power raises the possibility that the European Central Bank's (ECB) monetary policy actions could influence global energy prices. This paper, therefore, first examines the ECB's impact on global energy prices and their role in the monetary transmission mechanism. Second, it explores how these dynamics shape the optimal conduct of monetary policy in response to an energy price shock.

First, using a high-frequency event study analysis and a Bayesian proxy structural vector autoregressive (BPSVAR) model, we show that ECB monetary policy transmits through energy prices. More precisely, a contractionary ECB monetary policy shock leads to a strong and persistent decline in both global oil prices and energy prices faced by consumers in the euro area. To flesh out the importance of this result, we conduct a Lucas-critique-robust counterfactual (McKay and Wolf (2023)), in which ECB policy is not able to influence global oil prices. Under this counterfactual scenario consumer price inflation and inflation expectations react considerably less to changes in the ECB's policy stance — the response is more than halved in the short term. Furthermore, we show using an approach akin to the "Phillips-Multiplier" of Barnichon and Mesters (2021), that the ability to affect energy prices equally alleviates the inflation-output trade-off faced by the ECB by approximately 50%. Consequently, the ability to affect relatively flexible energy prices provides monetary policy with a tighter grip on short- and even medium-term inflation dynamics.

Second, we investigate how the ECB should respond optimally to a supply-side energy price shock and what role its ability to influence these prices play within the empirical framework of McKay and Wolf (2023). We consider two different loss functions. We begin with a stylized, single-mandate loss function that mirrors the ECB's primary mandate of medium-term price stability. We estimate that the optimal response to an oil price shock is a small, front-loaded tightening that quickly curbs the inflationary impact while imposing only a mild additional contraction in economic activity. Absent the ability to affect energy prices the tightening required to achieve the same optimal stabilization is estimated to be more than twice as large - as is the corresponding induced contraction in output. In the case of a central bank loss function that balances output and inflation deviations, our results suggest that the optimal response of the ECB to an energy price shock is close to a "looking-through" strategy. When we compare this optimal response to the actual response of the ECB, we find that they are very similar. Crucially, we show in a counterfactual scenario that the relatively muted response to energy price shocks is only close to optimal because the ECB directly affects global oil prices. Thus, energy prices can at times be the ECB's friend rather than its foe.

Figure 1: Role of the euro area in the global economy and oil market

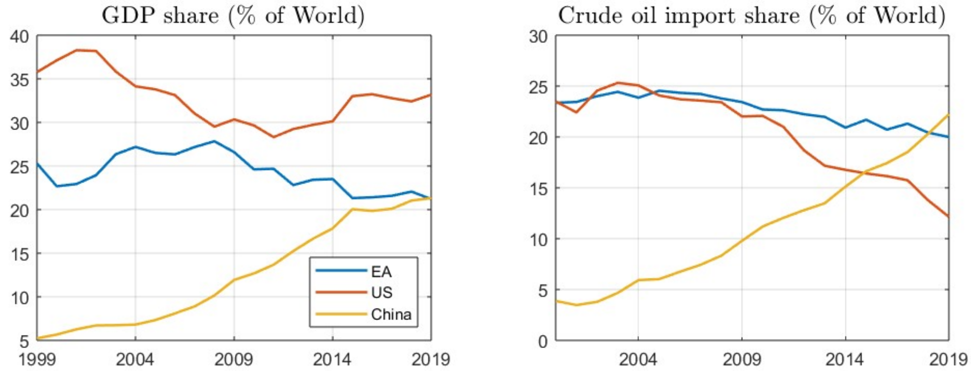

Notes: real GDP (in US dollars) and crude oil import shares of the euro area (EA), the United States (US) and China as a percentage of the world. GDP data from the IMF's World Economic Outlook database. Commodity import data are from UN Comtrade database.

In more detail, the paper first examines whether the ECB influences energy prices. As a motivating exercise, we follow the monetary policy event study literature (Gürkaynak et al. (2005); Altavilla et al. (2019); and many others) and use intra-day data to uncover the causal effects of changes in the ECB's monetary policy stance on the global oil price. Throughout the paper the Brent crude oil price acts as a stand-in for the price of energy goods traded on global financial markets – we refer to it as global energy price for short.[1] Our findings indicate that ECB policy decisions are rapidly transmitted to global energy prices, resulting in sizable and immediate effects. To further investigate these dynamics, we employ a BPSVAR model to analyze the business cycle and the dynamic effects of euro area monetary policy shocks on energy prices, as well as their role in monetary transmission. Importantly, our analysis shows that a contractionary monetary policy shock leads to a significant reduction in both globally traded energy prices and the energy prices faced by euro area consumers. This effect can be interpreted as follows: since the euro area is one of the world's largest oil importers (see Figure 1), changes in its monetary policy stance directly influence global oil and energy demand, which, in line with the theoretical work of Auclert et al. (2023) and Bayer et al. (2023), results in lower energy prices. Notably, these reductions in energy prices occur more quickly and are substantially more pronounced than the changes in headline consumer price indices. This is consistent with micro-data evidence showing that energy goods prices are updated much more frequently than those of other consumer goods (Aucremanne and Dhyne

[1]This assumption is not critical for our analysis. In particular, the Brent crude oil price is highly correlated with other major energy benchmarks, according to the IMF's Primary Commodity Prices database: its unconditional correlation with WTI crude is 0.99, with Dubai crude is 0.99, and even with Dutch TTF natural gas is 0.88. Also, the focus on oil prices in general is not critical, as we show in the online Appendix that, conditionally on a monetary policy shock, the same effects also materialize for gas prices (see Appendix B and E).

(2004)).

Having established that European monetary policy does, in fact, influence energy prices, we conduct a counterfactual exercise to assess the significance of this relationship for monetary policy transmission in the euro area. Our analysis is based on an empirical counterfactual scenario in which ECB decisions have no effect on global energy prices.[2] Specifically, in this counterfactual, the Organization of the Petroleum Exporting Countries (OPEC) adjusts supply to keep the global oil price at its preferred level, thereby neutralizing any impact of euro area monetary policy on global oil prices. We estimate this scenario using the method developed by McKay and Wolf (2023), which accounts for the anticipatory effects of such an OPEC policy rule change and is robust to the Lucas critique. To implement this approach, we draw on the literature on high-frequency identification of oil supply shocks (Känzig (2021)) to jointly identify both short-term and medium-term oil supply news shocks within our BPSVAR model. In the scenario where ECB decisions do not influence the global oil price, the response of energy prices faced by euro area consumers to a monetary tightening is substantially dampened. Crucially, this also results in a much weaker transmission of monetary policy to both inflation and inflation expectations. A comparison of the counterfactual and baseline responses highlights that, by influencing rapidly adjusting energy prices, monetary policy exerts significantly greater control over inflation dynamics—particularly in the short to medium term. Following an approach similar to the "Phillips-Multiplier" of Barnichon and Mesters (2021) we furthermore document that this also implies that the inflation-output trade-off faced by the ECB is substantially alleviated, by around 50%, even in the medium term.

Given the finding that the ECB's ability to influence energy prices plays a critical role in the transmission of monetary policy shocks, next, we study how this ability shapes the optimal conduct of monetary policy in response to a supply-side shock that induces an inflation-output trade-off. To this end, we employ the framework of McKay and Wolf (2023) to compute optimal policy. We implement the approach by identifying an oil supply shock as in Känzig (2021) and combine it with identified euro area monetary policy shocks. In line with the literature (cf. Barnichon and Mesters (2023), Barnichon and Mesters (2024)), we first define the optimal policy for the ECB as the policy that optimally stabilizes inflation only in the medium term and thereby achieves the primary mandate. While this reflects the institution's formal primary mandate, it is a stylized representation of actual policymaking. Therefore, we also consider a more realistic dual-mandate framework, where optimal policy assigns equal weight to stabilizing both inflation and output. For both loss functions, we

[2]This scenario aligns with the narrative often advanced by ECB officials. For example, at the press conference on February 3, 2022, Christine Lagarde responded to a question on this topic by stating: "If the ECB was to ... then raise interest rates in short order, do you think it would have any impact on energy prices? No, it is not in the ambit of monetary policy to decide the price of the barrel that is organized predominantly outside of Europe." (Lagarde (2022)) The assumption that ECB policy does not affect global energy prices is not only embedded in theoretical models used for policy analysis, but is also deeply rooted in the ECB's forecasting process (see Coenen et al. (2018) for an example and discussion).

compare the optimal policy response to an oil supply shock with the optimal response under the assumption that the ECB cannot influence global energy prices.

Under the ECB's primary mandate of medium-term inflation stabilization, the optimal policy response involves a modest, front-loaded tightening of monetary policy across both the short and long ends of the yield curve. This approach effectively curbs the rise in inflation, albeit at the cost of a slightly deeper but short-lived contraction in output. The underlying rationale is that a more contractionary policy stance swiftly counteracts the initial surge in oil prices, resulting in a substantially smaller increase in headline inflation and inflation expectations. Notably, only a marginal additional decline in output is required to minimize medium-term inflation deviations from target, as tighter ECB policy induces rapid and pronounced declines in the relatively flexible energy prices. Therefore, our findings suggest that the tightening necessary to optimally achieve the primary mandate is tightly linked to the ECB's ability to affect global energy prices. These findings underscore that the degree of monetary tightening necessary to optimally achieve the ECB's primary mandate is closely tied to its ability to influence global energy prices. When we consider a dual mandate framework—where the ECB assigns equal weight to stabilizing both inflation and output—the optimal policy response aligns with a "looking-through" strategy in the face of a supply-driven energy price shock. Interestingly, we find that the optimal response under the dual mandate closely mirrors the ECB's actual estimated response to an oil supply shock.

To further substantiate the importance of the ECB's ability to influence energy prices, we estimate the optimal policy response under the assumption that ECB policy decisions have no effect on global oil prices. In this case, the results of this counterfactual exercise reveal that, the ECB's optimal response to an oil supply shock would require a significantly stronger tightening of monetary policy, both at the short and long ends of the yield curve. As a consequence, the resulting decline in output would be substantially larger. This pattern holds under both the single-mandate framework—focused solely on inflation stabilization—and the dual-mandate framework. These findings highlight that the ECB's capacity to influence fast-moving energy prices is a critical factor in minimizing the economic costs associated with responding to energy price shocks.

The rest of the paper is structured as follows. In Section 2, we present the high-frequency event study analysis. Section 3 describes the empirical BPSVAR framework used throughout the paper. Section 4 examines if euro area monetary policy can affect energy prices. Section 5 studies the role of energy prices in the transmission of monetary policy. Section 6 investigates how the ECB's ability to affect energy prices shapes the optimal conduct of monetary policy. The last section concludes.

Related literature. Our paper contributes to the literature that studies how monetary policy transmits to the economy (Christiano et al. (1999); Gertler and Karadi (2015); Miranda-Agrippino and Ricco (2021); and many others). While the literature on monetary policy transmission is extensive, the specific role of energy prices has not been explored. In this paper, we fill that gap by examining how the response of energy prices to a monetary policy shock shapes its transmission to the economy. Importantly, our findings contribute to the literature studying the transmission pace of monetary policy and provide further evidence against the notion that monetary policy mostly transmits with long and variable lags (Buda et al. (2023)). The literature has documented that when monetary policy shocks are identified using instruments constructed from high-frequency financial data, monetary policy affects consumer prices already in the very short run (Miranda-Agrippino and Ricco (2021); Bauer and Swanson (2023)). We not only confirm this finding but also provide a more structural explanation for the quick response of consumer prices, which is tied to the ability of monetary policy to affect highly flexible energy prices.

Furthermore, our work speaks to the literature studying how monetary policy should react to an exogenous increase in energy prices. Previous studies have relied on dynamic stochastic general equilibrium (DSGE) models which are prone to model misspecification (Leduc and Sill (2004); Bodenstein et al. (2012); Natal (2012)) or have employed empirical methods vulnerable to the Lucas critique (Bernanke et al. (1997), Kilian and Lewis (2011)). We contribute to this literature by estimating the mandate-optimal monetary policy response empirically within a framework that is robust to the Lucas critique and by studying how a central bank's ability to impact crude oil prices is crucial for the optimal policy reaction to an oil supply shock. In this scope, Castelnuovo et al. (2024) comes closest to our analysis as they use the same Lucas-critique-robust empirical approach for a related question. However, while they focus on the role of the monetary policy response in the transmission of crude oil price and food price shocks, we study how the central bank's ability to impact the oil price shapes its mandate-optimal policy reaction to an oil supply shock.

Lastly, our paper is closely related to the literature that studies the effects of monetary policy on commodity prices. Existing work that utilizes the state-of-the-art approach of identifying monetary policy shocks in VARs using high-frequency monetary policy surprise series finds that contractionary US monetary policy shocks decrease commodity prices (Miranda-Agrippino and Rey (2020); Bauer and Swanson (2023)) and oil prices specifically (Degasperi et al. (2023); Miranda-Pinto et al. (2023)).[3,4] Meanwhile, Gazzani and Ferriani (2024) document a similar transmission of Chinese monetary policy to commodity prices. Building on these findings, we show that the oil price responds in a comparable manner to European

[3]Miranda-Agrippino and Ricco (2021) include the Commodity Research Bureau (CRB) commodity price index in their baseline VAR but do not report the IRFs. Therefore, using their replication files while keeping true to their baseline empirical specification, we produce the commodity price index IRFs and find that the commodity price index declines significantly in response to a contractionary US monetary policy shock (see Figure E.8 in the Appendix).

[4]A recent study that does not find this result is Gagliardone and Gertler (2023), who concludes that the real oil price does not respond to US monetary policy shocks. In Appendix C, we replicate their analysis and show that when the critique raised in Kilian (2024) regarding their aggregation of the monetary policy surprise series and the simultaneous use of the average-of-the-month crude oil price is addressed, the oil price does in fact strongly and significantly decline in response to a contractionary US monetary policy shock.

monetary policy as well. Crucially, and in contrast to the existing work, we show that this response has important implications for the transmission of monetary policy to inflation. Additionally, we find that these effects extend to inflation expectations, further connecting our paper to the interplay between energy prices and inflation expectations (Aastveit et al. (2023); Wehrhöfer (2023); Jo and Klopack (2024)).

2 Motivating evidence: Monetary policy & energy prices at high frequency

To start the analysis, we utilize the event study regression approach commonly employed in the literature to study the effects of monetary policy on asset prices. Using intraday data, we document that ECB monetary policy announcements impact global energy prices at high frequency.[5] For our baseline result we use the Brent oil price as our preferred measure of global energy prices, but we show in Appendix B that the results are robust to using natural gas prices instead. To put the results for the euro area into perspective, we compare them with the for the U.S. and the U.K., which are a large open economy and a small open economy, respectively.

To ensure comparability we measure unexpected changes in the interest rates — monetary policy surprises — using the intra-day changes in the three-month-ahead federal funds futures, the three-month overnight index swap (OIS) rate, and the three-month Libor rate in a narrow window around monetary policy announcements for the US, the euro area, and the UK, respectively.[6] We follow Jarociński and Karadi (2020) and purge monetary policy surprises from central bank information effects using changes in stock prices in the same window around the monetary policy announcement. Specifically, if stock prices and interest rates move in opposite directions, we label this a monetary policy shock. If not, we set the corresponding entry to zero. This corresponds to what Jarociński and Karadi (2020) call the "poor man's" identification approach. We use tick data from the Refinitiv Tick History database to compute the variation in the Brent crude oil price in the same narrow window around the monetary policy announcements. Precisely, we measure the price variation in the ICE Brent crude oil front-month futures (LCOc1), which is generally the benchmark global spot price quoted in financial news.

To study the effects of monetary policy on the global oil price, we estimate the following

[5]Gürkaynak et al. (2005), Beechey and Wright (2009), and others have shown that intraday data yield more precise point estimates of announcement effects than lower-frequency (daily) data. Rosa (2014) investigates this premise for crude oil futures prices and finds that oil prices respond to a broader range of news announcements than other U.S. asset prices, highlighting the importance of using intraday data for our analysis.

[6]Data sources are Gürkaynak et al. (2005), Altavilla et al. (2019), and Cesa-Bianchi et al. (2020). The choice of the interest rate maturity (three-month) is not only widely used in the literature (Jarociński and Karadi (2020),Cesa-Bianchi et al. (2020)) but also allows us to ensure comparability across countries given data availability. Results are similar when using other maturities.

high-frequency event study regression for the ECB, the Federal Reserve, and the Bank of England separately:

$$p_t^{oil} = \alpha_i + \beta_i mps_{i,t} + \epsilon_{i,t} \quad i \in [EA, US, UK]. \tag{1}$$

p_t^{oil} is the intraday percent variation in the Brent crude oil price (in US dollars) around the monetary policy announcement on day t, and $mps_{i,t}$ represents the corresponding monetary policy surprise of country i.[7]

Table 1: Results of the event study regression for the euro area, US, and UK

	EA	US	UK
$\widehat{\beta}^{std}$	-0.056^{**}	-0.078^{**}	0.019
	(0.026)	(0.037)	(0.037)
Sample	2002:1-2019:12	1996:1-2019:12	1997:6-2019:12
N	182	198	246
R^2 (%)	3.37	2.64	0.38

Notes: Coefficient estimates $\widehat{\beta}^{std}$ measure the percentage change in the front month future of the Brent crude oil price following a 1 standard deviation increase in the country-specific monetary policy surprise. Heteroskedasticity-robust standard errors are reported in parentheses. *, **, *** represent statistical significance levels at 10%, 5%, and 1%, respectively.

The intra-day responses of the Brent crude oil price to a one standard deviation – i.e. an average — contractionary monetary policy shock are presented in Table 1.[8] The results show that the Brent oil price declines immediately in response to an unexpected interest rate increase in both the euro area and the US, while it remains unaffected by a similar increase in the UK.[9]

This suggests that financial market participants update their expectations of the global oil market in light of surprise policy actions by the ECB and the Federal Reserve. In turn,

[7] Sample for the euro area event study regression starts from 2002 following the suggestions from Altavilla et al. (2019) and Andrade and Ferroni (2021) due to liquidity issues in the OIS during 1999-2001. Starting sample for the US regression is 1996 due to the availability of intraday Brent crude oil price data. Sample for the UK regression starts from June 1997 due to the availability of the UK monetary policy surprise series.

[8] Appendix B offers additional material and robustness along several dimensions such as sample period, instrument choice and choice of the energy price.

[9] Our findings for the US are similar to the results in the relevant literature that makes use of intraday data (Rosa (2014); Basistha and Kurov (2015)). However, there are two papers focusing on the euro area that have contradicting results to ours (Torro (2019); Soriano and Torró (2022)). These two papers do not control for central bank information effects, which is arguably a more prominent issue in the euro area rather than in the US (Jarociński and Karadi (2020)). Furthermore, Torro (2019) uses daily data against the recommendation of Rosa (2014) specific to crude oil prices.

this implies that, like the United States, the euro area functions as a large open economy in the global energy market. Such behavior is consistent with the fact that, for much of our sample period, the euro area was the world's largest oil importer —— surpassed only by the United States in earlier years and by China more recently (see Figure 1).

3 The empirical framework

The high-frequency event study shows that the ECB's monetary policy has an immediate and significant effect on energy prices. Motivated by this evidence, the rest of the paper studies the dynamics of this relationship. Therefore, this section presents our time series model of monetary policy and energy prices in the euro area. We first outline the general Bayesian proxy structural vector-autoregressive (BPSVAR) model framework of Arias et al. (2021) that allows us to identify dynamic causal effects with the use of instrumental variables. Next, we discuss our model specification and endogenous variables. Finally, we present our identifying assumptions in detail. Our empirical analysis in later sections requires the identification of up to two structural shocks simultaneously. Therefore, we discuss our approach to identifying both a single, as well as two shocks in the BPSVAR model. Note, however, that although the type and the number of structural shocks we identify varies according to the application, all shocks are identified in one consistent model with a constant set of endogenous variables.

3.1 Bayesian proxy SVAR model

We lay out the BPSVAR model for the general case with $k \geq 1$ proxy variables and k structural shocks of interest. Following the notation of Rubio-Ramirez et al. (2010), the structural VAR model with one lag and without deterministic terms can be written as:

$$\boldsymbol{y}_t'\boldsymbol{A}_0 = \boldsymbol{y}_{t-1}'\boldsymbol{A}_1 + \boldsymbol{\epsilon}_t', \qquad \boldsymbol{\epsilon} \sim N(\boldsymbol{0}, \boldsymbol{I}_n), \tag{2}$$

where $\boldsymbol{y}_t$ is an $n \times 1$ vector of endogenous variables and $\boldsymbol{\epsilon}_t$ an $n \times 1$ vector of structural shocks. The BPSVAR framework builds on the following assumptions in order to identify k structural shocks of interest: There exists a $k \times 1$ vector of proxy variables $\boldsymbol{m}_t$ that are correlated with the k structural shocks of interest $\boldsymbol{\epsilon}_t^*$ and orthogonal to the remaining structural shocks $\boldsymbol{\epsilon}_t^o$. Formally, the identifying assumptions are

$$E[\boldsymbol{\epsilon}_t^* \boldsymbol{m}_t'] = \underset{(k\times k)}{\boldsymbol{V}}, \tag{3a}$$

$$E[\boldsymbol{\epsilon}_t^o \boldsymbol{m}_t'] = \underset{((n-k)\times k)}{\boldsymbol{0}}, \tag{3b}$$

and represent the relevance and the exogeneity condition, respectively. Below in section 3.3 we explicitly write out the identifying assumptions for two applications, with $k = 1$ and $k = 2$, respectively.

We estimate the BPSVAR model using the algorithm developed in Arias et al. (2021). In this algorithm the model in (2) is augmented by the equations for the proxy variables. More precisely, denote by $\tilde{\boldsymbol{y}}'_t \equiv (\boldsymbol{y}'_t, \boldsymbol{m}'_t)$, by $\tilde{\boldsymbol{A}}_\ell$ the corresponding $\tilde{n} \times \tilde{n}$ coefficient matrices with $\tilde{n} = n + k$, and by $\tilde{\boldsymbol{\epsilon}}' \equiv (\boldsymbol{\epsilon}'_t, \boldsymbol{\eta}'_t) \sim N(\mathbf{0}, \boldsymbol{I}_{n+k})$, where $\boldsymbol{\eta}_t$ is a $k \times 1$ vector of measurement errors. The augmented structural VAR model is then given by

$$\tilde{\boldsymbol{y}}'_t \tilde{\boldsymbol{A}}_0 = \tilde{\boldsymbol{y}}'_{t-1} \tilde{\boldsymbol{A}}_1 + \tilde{\boldsymbol{\epsilon}}'_t. \tag{4}$$

In the estimation of the model in (4), the algorithm by Arias et al. (2021) imposes the assumptions (3a) and (3b) to identify the structural shocks.

3.2 Data and model specification

Our baseline BPSVAR model for the euro area includes nine endogenous variables. Our starting point is a standard monetary model featuring the 1-year constant maturity yield on German Bunds as a monetary policy indicator, the industrial production index (excluding construction) as a measure of economic activity, the Harmonised Index of Consumer Prices (HICP) as a measure of the price level, and the BBB corporate bond spread to capture financial conditions (cf. Gertler and Karadi (2015)). To this setup, we add variables that are important for our analysis of the interplay of monetary policy with energy prices. Therefore, we add the energy component of the HICP as a measure of energy prices in the euro area, the Brent crude oil price, and one-year-ahead inflation forecasts from Consensus Economics to capture inflation expectations. Since the euro area is a major energy importer, we also add the EUR-USD exchange rate, as oil and other energy commodities are generally traded in US dollars. Finally, we add the 5-year constant maturity yield on German Bunds since our analysis is going to include the identification of forward guidance shocks later on. The sources and more details on the data can be found in Appendix A.

The variables are measured in monthly frequency. Furthermore, all variables except interest rates and credit spreads enter the SVAR in log levels ($\times 100$), so that the impulse responses can be interpreted as percentage deviations. The BPSVAR model is estimated on a sample from January 2002 to December 2019. As in our high-frequency event study, we exclude the period 1999 – 2001 due to liquidity issues in the OIS contracts, which, as discussed below, will serve as a proxy to identify monetary policy shocks. The model has 12 lags and includes a constant. We follow Arias et al., 2021 and use flat priors for estimating the BPSVAR parameters.[10] In addition, a relevance threshold is imposed to express the prior belief that the proxy is informative to identify monetary policy shocks. In particular, we assume that the identified structural monetary policy shocks account for at least 10% of the variance in the

[10] As in Born and Pfeifer (2021) and many other studies we impose the dogmatic prior that the SVAR is stable implying that, after being hit by an exogenous shock, the endogenous variables eventually converge back to their steady state.

proxy.[11]

3.3 Identifying assumptions

Our empirical strategy to identify structural shocks relies on an instrumental variables — or proxy — approach. In the following we first lay out how we identify a single structural shock with the use of an appropriate proxy. To conduct empirical counterfactuals, the method by McKay and Wolf (2023) requires the identification of as many structural shocks as possible to minimize approximation error. Therefore, we subsequently present our identification assumptions for the case of identifying two shocks simultaneously with the use of two proxy variables.

Define matrices containing all 12 lags of the endogenous variables and the proxy as $\boldsymbol{Y}'_{t-1} = (\boldsymbol{y}'_{t-1}, \dots, \boldsymbol{y}'_{t-12})$ and $\boldsymbol{M}'_{t-1} = (m_{t-1}, \dots, m_{t-12})$. For the identification of a single structural shock with a single proxy, we can derive the Equation governing the (scalar) proxy variable from Equation (4) (see Appendix D for details):

$$m_t = (\boldsymbol{Y}'_{t-1}, \boldsymbol{M}'_{t-1})\boldsymbol{B}_1 + v_{1,1}\epsilon^*_t + B_2\eta_t, \tag{5}$$

where $\boldsymbol{V} = v_{1,1} \neq 0$ by the relevance condition (3a). All remaining structural shocks, $\boldsymbol{\epsilon}^o_t$, are unrelated to m_t by means of the exogeneity condition (3b). Note that relative to the standard frequentist external instrument procedure as in Mertens and Ravn (2013), Equation (5) illustrates that here the proxy variable is allowed to be serially correlated, predictable, and affected by measurement error.

For the case of identifying two shocks with two proxies, define the structural shocks as $\boldsymbol{\epsilon}^*_t \equiv (\epsilon_{1,t}, \epsilon_{2,t})'$. In this case, we need additional identifying assumptions since the proxies are allowed to be correlated with both structural shocks of interest. Instead of imposing potentially contentious zero restrictions to disentangle the two shocks, we rely on relatively weak magnitude restrictions to obtain set-identification. Specifically, our identifying assumptions in the two-shock scenario are

$$m_{1,t} = (\boldsymbol{Y}'_{t-1}, \boldsymbol{M}'_{t-1})\boldsymbol{B}_1 + v_{1,1}\epsilon^*_{1,t} + v_{1,2}\epsilon^*_{2,t} + \boldsymbol{B}_2\boldsymbol{\eta}_t, \tag{6}$$

$$m_{2,t} = (\boldsymbol{Y}'_{t-1}, \boldsymbol{M}'_{t-1})\boldsymbol{B}_1 + v_{2,1}\epsilon^*_{1,t} + v_{2,2}\epsilon^*_{2,t} + \boldsymbol{B}_2\boldsymbol{\eta}_t, \tag{7}$$

$$v_{1,1} > v_{1,2}, \quad v_{2,2} > v_{2,1}. \tag{8}$$

As an example, we later identify two dimensions of monetary policy with high-frequency surprises in short- and long-maturity OIS contracts: a conventional, or contemporaneous

[11] This is a weak requirement compared to the 20% threshold of Arias et al. (2021) and the 'high-relevance' prior of Caldara and Herbst (2019). As shown in Figure E.5 in the Appendix, our results are robust to reducing the relevance condition to 0.

monetary policy shock, and a forward guidance shock. Then, in words, we only assume that the conventional monetary policy shock affects the short-maturity OIS contract proxy more strongly compared to the forward guidance shock ($v_{1,1} > v_{1,2}$), and vice versa.

Relative to the standard frequentist two-step estimation, the algorithm and the Bayesian approach in general have the following advantages. First, we refrain from imposing potentially contentious recursiveness assumptions between the endogenous variables when multiple structural shocks are identified. Second, the single-step estimation of the BPSVAR model is more efficient than the standard two-stage least squares estimation of proxy SVAR and facilitates coherent inference. In fact, the Bayesian set-up allows exact finite sample inference and does not require an explicit theory to accommodate potentially weak instruments. Third, the BPSVAR framework allows the proxy variables to be serially correlated, predictable, and affected by measurement error. Lastly, Bayesian inference is particularly convenient in the presence of set identification, which arises in our applications with two proxies.[12]

4 Monetary policy and energy prices: SVAR evidence

The event study in Section 2 demonstrated that the ECB's policy decisions impact energy prices at high frequency. We now utilize the framework established in Section 3 to investigate whether this effect also materializes at a monthly frequency and how it influences the economy. To achieve this, we construct an instrument for a monetary policy shock using high-frequency changes in interest rate futures around monetary policy announcements (similarly to Gertler and Karadi (2015); Jarociński and Karadi (2020); Miranda-Agrippino and Ricco (2021)). In this section, we analyze the euro area's economic response to such a shock, compare our findings with existing literature, and provide evidence supporting their robustness.

4.1 Dynamic response to monetary policy shock

We construct the monetary policy shock in three steps. First, we capture revisions in interest rate expectations at different interest-rate maturities. More precisely, we compute the first principal component of OIS surprises from the Altavilla et al. (2019) dataset with maturities from one month up to one year. This "generic" monetary policy shock (Nakamura and Steinsson (2018); Bauer and Swanson (2023)) has the advantage that is does not depend on one specific maturity. Second, we purge the resulting surprise series from central bank information effects (see Section 2). Third, to aggregate the surprises to monthly frequency, we employ

[12] We fully acknowledge the concerns that in the case of set identification, our uniform prior for the rotation matrix, which is embedded in the approach of Arias et al. (2021), may even asymptotically influence our results as forcefully raised by Baumeister and Hamilton (2019) and Giacomini and Kitagawa (2021). But recent contributions by Inoue and Kilian (2021) and Arias et al. (2023) called into question the empirical relevance of this concern in applied research with tightly identified sets as is the case in our applications. Therefore we conduct standard Bayesian inference along the lines of Rubio-Ramirez et al. (2010) and the subsequent literature.

the approach proposed by Kilian (2024), which takes the accounting relationship between daily and average monthly data into account. We denote this proxy by $m_{t,PC1}^{MP}$. It enters the framework in Section 3 Equation (5) in the following way: $m_t \equiv m_{t,PC1}^{MP}$ with $\epsilon_t^* \equiv \epsilon_{t,generic}^{MP}$ as the corresponding structural shock. The shock represents a linear combination of monetary policy shocks at different maturities which, in combination, move the term structure of interest rates (Inoue and Rossi (2021) and McKay and Wolf (2023)).

In Figure 2, we illustrate the impulse responses to a one standard deviation contractionary monetary policy shock. Both the 1-year and 5-year Bund yields increase by approximately 5 basis points on impact before reverting to zero. Industrial production experiences a slight decline initially, reaching its lowest point after 10 months. Similarly, the domestic headline consumer price index quickly declines, dropping to about 0.12% after a year and remaining depressed. Concurrently, the euro appreciates against the dollar by just under 0.4% in the short term, financial conditions tighten, and inflation expectations decline significantly and persistently. Overall, the estimated dynamics for the endogenous variables align with standard theory and previous findings in the literature.

The main result of this section is the substantial decline in global crude oil prices and the energy prices faced by euro area consumers (HICP energy). The oil price drops sharply on impact and remains subdued for over a year, with a trough response of 3%. Additionally, the HICP energy price index falls by 0.65%, a much larger decline than that observed in the overall HICP basket. Given that energy prices constitute about 10% of the overall HICP basket, a back-of-the-envelope calculation suggests that the majority of the decline in the overall HICP in the short- and medium-term can be attributed to the contractionary monetary policy shock's effect on oil prices and, subsequently, energy prices in the euro area.

Importantly, our finding that energy prices adjust more rapidly and significantly than other goods is supported by the micro-data literature on price changes. For instance, Aucremanne and Dhyne (2004) analyze the micro-data used to compute the Belgian HICP and find that the average price duration for energy goods is approximately one month. This contrasts sharply with the median price duration of all goods in the basket, which is around 14 months. From the perspective of a standard New-Keynesian Model, this implies that, all else being equal, the Phillips curve for energy goods is steeper compared to that for average goods. Consistent with this theoretical intuition, nearly all subcomponents of the HICP energy component exhibit a significant decline in their prices (see Figure E.7 in the Appendix). Among these, the price of fuels, which are more flexible, contracts the most.

4.2 Robustness and discussion of results

Our baseline empirical specification indicates that an exogenous increase in the one-year yield, scaled to 100 basis points (peak response), results in a 9.3% decline in industrial production (trough response). This represents a significantly larger industrial production elasticity to

Figure 2: Transmission of a generic EA monetary policy shock

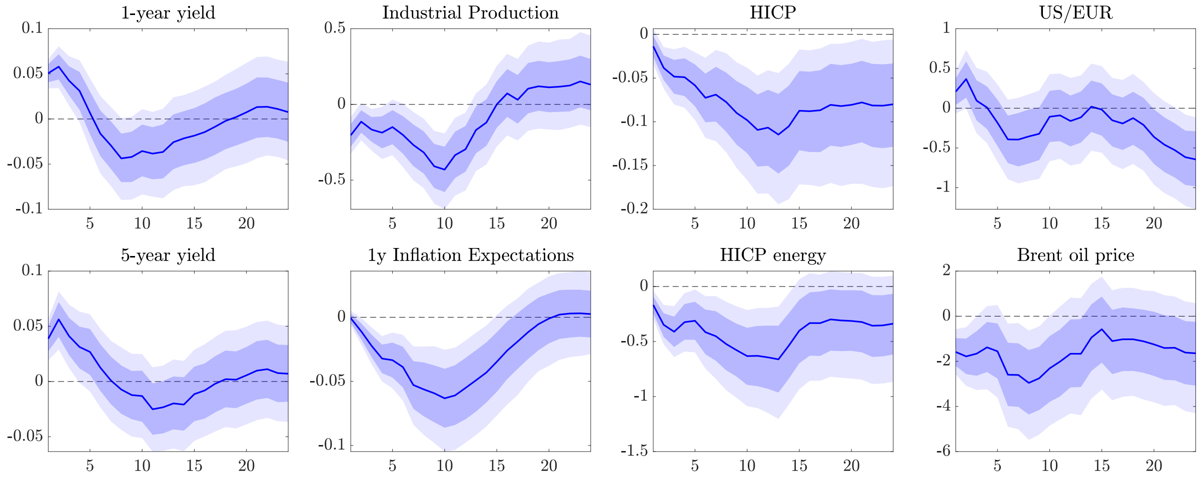

Notes: Impulse responses to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months. Impulse responses for variables that do not correspond to interest rates or inflation rates are expressed in percent. Impulse responses for inflation rates and interest rates are expressed in annualized percentage points. Response of the credit spread is omitted to save space (see Figure E.1).

the interest rate than typically found in well-known monetary policy proxy SVAR studies, which often focus on the US and use a considerably larger sample starting as early as 1973. In Table F.1 in Appendix E, we demonstrate that, once differences in the sample period are accounted for, our estimate aligns fully with the empirical evidence from existing studies for the US, UK, and the euro area (with an average elasticity of 11.3%). This survey of the literature suggests that, given state-of-the-art identification, a larger elasticity of output to surprise changes in the interest rate is a characteristic of modern data, particularly for the euro area.

The estimated impact of a contractionary monetary policy shock in our baseline specification results in not only a significant contraction in industrial output but also a pronounced decline in the Brent crude oil price, reflecting the severity of the overall economic downturn. In Appendix F, we elaborate on a mechanism that explains this substantial drop in energy prices. Key factors include the empirically relevant low elasticity of substitution for energy goods and the fixed energy supply in the short run, which leads to a relatively steep demand curve and a vertical supply curve. Consequently, for a given policy-induced fall in demand, the price reacts strongly to balance the energy market. We then simulate a monetary tightening of similar magnitude to our baseline results using a state-of-the-art general equilibrium model (Bayer et al. (2023)) that incorporates such a market structure, demonstrating that the energy price indeed drops by approximately the same amount. Finally, we show that, when measured in terms of the standard deviation of forecast errors, a 2% surprise fall in the oil price is less rare than a 5 basis point unanticipated increase in the one-year yield, consistent with theoretical arguments that energy prices are highly volatile (Figure F.2).

In Appendix E, we demonstrate that our results are robust to alternative specifications. To address concerns about the relevant information set regarding the global oil market (Baumeister and Hamilton (2019)), we extend the model to include global oil production and global industrial production (Figure E.2). Since the BPSVAR approach relies on the assumption of (partial) invertibility to identify the monetary policy shock, we follow the suggestion of Plagborg-Møller and Wolf (2021) and also report impulse responses based on their proposed "internal instrument" approach, which remains robust even in the case of non-invertibility. The results are very similar (Figure E.3). Furthermore, we document that gas prices also fall substantially after a monetary policy shock, indicating that our choice of the Brent oil price as a baseline measure of global energy prices does not significantly affect our results (Figure E.4). The results are consistent across all variables in the VAR. Lastly, we show that our findings remain largely unchanged when removing our prior on the relevance of the proxy variable (Figure E.5) or when incorporating the pandemic into the estimation using the Pandemic Priors approach of Cascaldi-Garcia (2022) (Figure E.6).

5 The role of energy prices in monetary transmission

The impulse response functions and back-of-the-envelope calculations from Section 4 suggest that energy prices play an important role in the monetary transmission mechanism. To further substantiate this point, we conduct an empirical counterfactual exercise in this section, in which the global oil price is assumed not to respond to ECB monetary policy shocks. We begin by outlining the general framework, then identify two additional structural shocks required to implement the counterfactual methodology, and finally apply this approach to examine how a monetary policy shock would propagate through the euro area economy and how the inflation-output trade-off would be altered if the ECB were unable to influence global oil prices.

5.1 Structural policy counterfactual: general framework

The approach to estimating impulse responses under the counterfactual OPEC policy rule builds on the recent insights of McKay and Wolf (2023, henceforth MW). In particular, MW develop an approach for constructing policy-rule counterfactuals empirically that is (i) robust to the Lucas critique and (ii) recovers the true policy-rule counterfactual for a wide range of underlying structural frameworks, including standard representative and heterogeneous-agent New Keynesian models. The key components in their counterfactual analysis are impulse responses to shocks to current and future policy. Specifically, they show that by combining the impulse response function to the structural shock of interest —estimated under the baseline policy rule— with a particular sequence of impulse responses to policy (news) shocks, one can uncover the impulse response functions to the structural shock under a counterfactual policy rule.

Formally, MW consider a linear, perfect-foresight, infinite-horizon economy in terms of deviations from the deterministic steady state for periods $t = 0, 1, 2, \ldots$. In sequence-space notation, this economy can be described by a set of equations

$$\boldsymbol{\mathcal{H}}_x \boldsymbol{x} + \boldsymbol{\mathcal{H}}_z \boldsymbol{z} + \boldsymbol{\mathcal{H}}_\epsilon \boldsymbol{\epsilon} = \mathbf{0}, \tag{9}$$

$$\boldsymbol{\mathcal{A}}_x \boldsymbol{x} + \boldsymbol{\mathcal{A}}_z \boldsymbol{z} + \boldsymbol{\nu} = \mathbf{0}, \tag{10}$$

where $\boldsymbol{x} \equiv (\boldsymbol{x}_1', \boldsymbol{x}_2', \ldots, \boldsymbol{x}_{n_x}')'$ stacks the time paths of the n_x endogenous variables over n_h periods, analogously $\boldsymbol{z}$ stacks the time path of the n_z policy instruments. The matrices $\boldsymbol{\mathcal{H}}$ summarize the behavior of agents in the non-policy block, while the matrices $\boldsymbol{\mathcal{A}}$ describe the baseline policy rule of interest. $\boldsymbol{\epsilon}$ represents the n_ϵ non-policy structural shocks and $\boldsymbol{\nu}$ the n_ν policy (news) shocks; the latter are deviations from the policy rule announced at date t but implemented only in some future period $t + i$, $i \geq 0$. The key assumption reflected in Equations (9) and (10) is that $\{\boldsymbol{\mathcal{H}}_x, \boldsymbol{\mathcal{H}}_z, \boldsymbol{\mathcal{H}}_\epsilon\}$ do not depend on the coefficients of the policy rule $\{\boldsymbol{\mathcal{A}}_x, \boldsymbol{\mathcal{A}}_z\}$, so that policy affects the non-policy block's decisions only through the path of the instrument $\boldsymbol{z}$, rather than through the policy rule *per se*. As shown in MW, this assumption holds true for a broad range of structural frameworks frequently used in counterfactual policy analysis such as standard representative and heterogeneous-agent New Keynesian models.

Under the assumption that the solution exists and is unique, the solution to Equations (9) and (10) can be written in impulse response space as

$$\begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{z} \end{pmatrix} = \boldsymbol{\Theta}_{\mathcal{A}} \times \begin{pmatrix} \boldsymbol{\epsilon} \\ \boldsymbol{\nu} \end{pmatrix}, \quad \boldsymbol{\Theta}_{\mathcal{A}} \equiv (\boldsymbol{\Theta}_{\epsilon,\mathcal{A}}, \boldsymbol{\Theta}_{\nu,\mathcal{A}}) \equiv \begin{pmatrix} \boldsymbol{\Theta}_{x,\epsilon,\mathcal{A}} & \boldsymbol{\Theta}_{x,\nu,\mathcal{A}} \\ \boldsymbol{\Theta}_{z,\epsilon,\mathcal{A}} & \boldsymbol{\Theta}_{z,\nu,\mathcal{A}} \end{pmatrix}. \tag{11}$$

where $\boldsymbol{\Theta}_{\mathcal{A}}$ collects the impulse responses of the policy instrument $\boldsymbol{z}$ and the non-policy variables $\boldsymbol{x}$ under the baseline policy rule summarized by $\boldsymbol{\mathcal{A}}$.

In the counterfactual analysis below, we are interested in analyzing impulse responses to a non-policy shock ϵ under a counterfactual policy rule $\{\tilde{\boldsymbol{\mathcal{A}}}_x, \tilde{\boldsymbol{\mathcal{A}}}_z\}$. The policy block with the counterfactual policy rule is then given by:

$$\tilde{\boldsymbol{\mathcal{A}}}_x \boldsymbol{x} + \tilde{\boldsymbol{\mathcal{A}}}_z \boldsymbol{z} = \mathbf{0}. \tag{12}$$

MW show that knowledge of the impulse responses $\boldsymbol{\Theta}_{\mathcal{A}}$ under the baseline policy rule is sufficient to determine the impulse responses to the structural shock of interest ϵ under any counterfactual policy rule even without knowing the true underlying structural model that generates the data. In particular, they prove that

$$\boldsymbol{x}_{\widetilde{\mathcal{A}}}(\boldsymbol{\epsilon}) = \boldsymbol{\Theta}_{x,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{x,\nu,\mathcal{A}} \times \widetilde{\boldsymbol{\nu}}, \qquad \boldsymbol{z}_{\widetilde{\mathcal{A}}}(\boldsymbol{\epsilon}) = \boldsymbol{\Theta}_{z,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{z,\nu,\mathcal{A}} \times \widetilde{\boldsymbol{\nu}}. \tag{13}$$

In words, the impulse response to the structural shock ϵ under the counterfactual policy rule $\boldsymbol{x}_{\widetilde{\mathcal{A}}}(\boldsymbol{\epsilon}) \equiv \boldsymbol{\Theta}_{x,\epsilon,\widetilde{\mathcal{A}}} \times \boldsymbol{\epsilon}$ is exactly equivalent to a combination of the corresponding impulse responses under the baseline policy rule $\boldsymbol{\Theta}_{x,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon}$ and the impulse responses to a specific sequence of policy news shocks $\widetilde{\boldsymbol{\nu}}$. Intuitively, as long as the decisions of the non-policy block depend on the (expected) path of the policy instrument rather than on the rule itself, it does not matter whether the path is due to the systematic conduct of policy or to policy news shocks. Consequently, the policy news shocks $\widetilde{\boldsymbol{\nu}}$ are chosen such that the counterfactual policy rule holds

$$\tilde{\boldsymbol{\mathcal{A}}}_x \left[\boldsymbol{\Theta}_{x,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{x,\nu,\mathcal{A}} \times \widetilde{\boldsymbol{\nu}}\right] + \tilde{\boldsymbol{\mathcal{A}}}_z \left[\boldsymbol{\Theta}_{z,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{z,\nu,\mathcal{A}} \times \widetilde{\boldsymbol{\nu}}\right] = \mathbf{0}. \tag{14}$$

What needs to be determined are the expressions $\boldsymbol{\Theta}_{x,\nu,\mathcal{A}}$ and $\boldsymbol{\Theta}_{z,\nu,\mathcal{A}}$ in Equation (13). Theoretically, this would require knowledge of impulse responses to news shocks that communicate changes in future policy over all possible n_h horizons. In practice, however, it is difficult, if not often impossible, to estimate impulse responses to policy news shocks $\boldsymbol{\nu}$. However, MW show that in practice, knowledge of a subset of the policy news shocks $\widetilde{\boldsymbol{s}} \subseteq \widetilde{\boldsymbol{\nu}}$ and their impulse responses $\boldsymbol{\Theta}_{\boldsymbol{s},\mathcal{A}}$ from the empirical literature often is sufficient as long as each entails a different future path of the policy instrument. In particular, they argue that using even only a small number of shocks $\boldsymbol{s}$ that solve

$$\underset{\widetilde{\boldsymbol{s}}}{min} \;\; ||\tilde{\boldsymbol{\mathcal{A}}}_x \left[\boldsymbol{\Theta}_{x,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{x,s,\mathcal{A}} \times \widetilde{\boldsymbol{s}}\,\right] + \tilde{\boldsymbol{\mathcal{A}}}_z \left[\boldsymbol{\Theta}_{z,\epsilon,\mathcal{A}} \times \boldsymbol{\epsilon} + \boldsymbol{\Theta}_{z,s,\mathcal{A}} \times \widetilde{\boldsymbol{s}}\,\right]||, \tag{15}$$

produces a reliable "best Lucas critique-robust approximation".[13]

5.2 Structural policy counterfactual: application details

To estimate $\boldsymbol{\Theta}_{z,s,\mathcal{A}}$ $\boldsymbol{\Theta}_{x,s,\mathcal{A}}$ we identify OPEC-related oil supply news shocks using the proxy variables constructed by Känzig (2021). These proxy variables capture high-frequency changes in oil price futures around OPEC meetings, making them a valid instrument for OPEC oil supply news shocks. Given that the accuracy of the approximations of Equation (14) depends on the number of policy news shocks identified, we diverge from Känzig (2021), who uses only the first principal component of changes in oil price futures at various horizons to identify a

[13] We compute this solution as follows: We first stack the system in Equation (14) across all the responses of all the $n = n_x + n_z$ endogenous variables $\boldsymbol{x}$ and the policy instrument $\boldsymbol{z}$ in Equation (14) and all horizons $n_H = n \times n_h$ in order to arrive at $\tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{\epsilon,A} \times \boldsymbol{\epsilon} + \tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{\nu,\mathcal{A}} \times \widetilde{\boldsymbol{s}} = 0$. Then we collect all impulse response functions to the n_s identified policy (news) shocks in $\boldsymbol{\Theta}_{s\mathcal{A}}$, and build $\boldsymbol{\Theta}_{\nu,\mathcal{A}} = [\boldsymbol{\Theta}_{s,\mathcal{A}}, \mathbf{0}_{(n \times n_h)\times(n_h - n_s)})]$. We then solve for $\widetilde{\boldsymbol{s}} = -\left(\tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}\right)^{\star} \tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{A,\epsilon} \times \boldsymbol{\epsilon}$ with $\left(\tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}\right)^{\star}$ as the Moore-Penrose inverse of $\tilde{\boldsymbol{\mathcal{A}}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}$.

single oil supply news shock. Instead, we use high-frequency changes in the 3-month ($m^{oil}_{t,3m}$) and 24-month ($m^{oil}_{t,24m}$) futures to identify short-term ($\nu^{oil}_{t,short}$) and medium-term ($\nu^{oil}_{t,medium}$) oil supply news shocks. In doing so, we incorporate recent suggestions in the literature to refine the identification of these oil supply news shocks (Degasperi et al. (2023), Kilian (2024)).[14] In the context of our general identifying assumptions framework in Equation (6), this implies that we set $\boldsymbol{\epsilon}^*_t \equiv (\epsilon^*_{1,t}, \epsilon^*_{2,t})' = (\nu^{oil}_{t,short}, \nu^{oil}_{t,medium})'$ and $\boldsymbol{m}_t \equiv (m_{1,t}, m_{2,t})' = (m^{oil}_{t,3m}, m^{oil}_{t,24m})'$.

Figure 3: Transmission of short-run (purple) and medium-run (magenta) oil supply news shocks

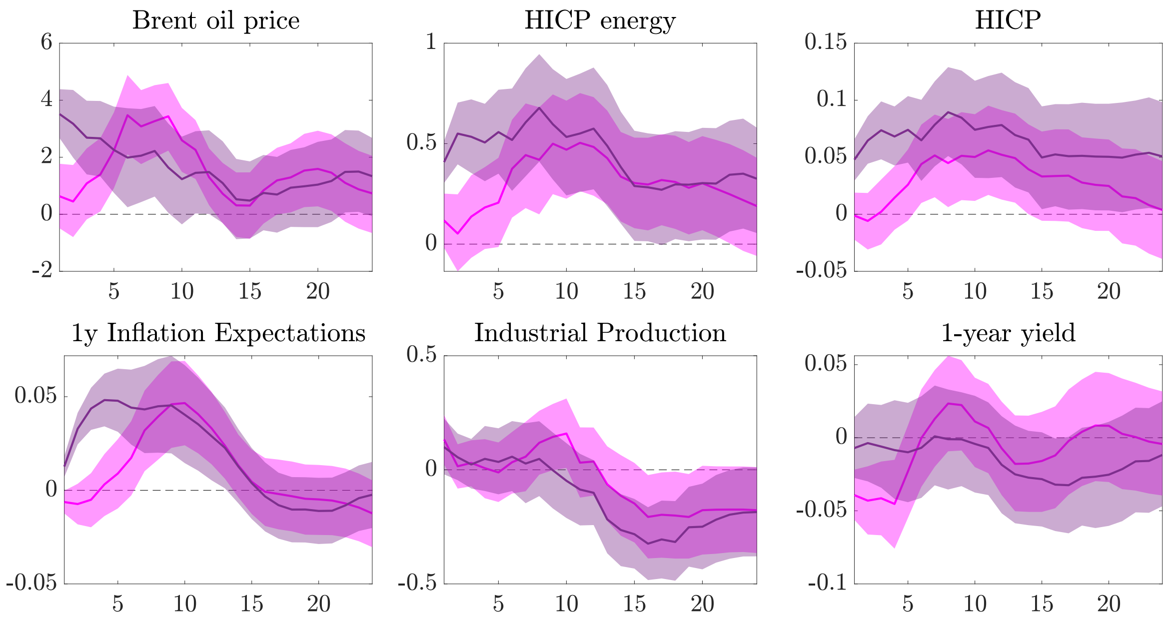

Notes: Impulse responses to the short-run oil supply news shock and corresponding 68% credible sets in magenta. Impulse responses to the medium-run oil supply news shock and corresponding 68% credible sets in purple. Horizon in months. We normalize the short-run (medium-run) oil supply news shock to increase the oil price on impact (after twelve months). Responses of the credit spread, exchange rate, and 5-year government bond yield are omitted to save space. See notes to Figure 2 for scaling of variables.

We present the impulse responses to these two oil supply news shocks in Figure 3.[15] In summary, both shocks increase the price of oil and lead to a contraction of output as well as a rise in consumer prices in the euro area. The response of the oil price and most of the other endogenous variables to the short-term oil supply news shock is strong and immediate, whereas the medium-term oil supply shock has more delayed effects on the oil price and the

[14] In particular, as we use average monthly data, we aggregate the surprises to monthly frequency using the approach proposed by Kilian (2024) instead of summing over them as in the original paper by Känzig (2021). Furthermore, we cleanse the surprises from possible oil demand shocks using the approach of Degasperi (2023).

[15] To maximize the number of observations we start the estimation directly in 1999 because the proxies proposed in Känzig (2021) do not suffer from the same liquidity issue as OIS contracts during the early period of the euro area.

broader of economy.

5.3 Structural policy counterfactual: What if OPEC stabilizes the oil price?

With the two estimated oil supply news shocks, we are now in the position to construct empirical policy rule counterfactuals to gauge the role of energy prices in monetary transmission in the euro area.

Specifically, we investigate how the euro area economy would respond if the ECB did not influence the oil price. The particular counterfactual we consider is a scenario in which OPEC aims to stabilize the global oil price by adjusting its oil supply accordingly. In other words, the counterfactual OPEC policy rule is such that it aims to stabilize the oil price at its steady-state level, i.e. $E_t[\widehat{p}_{t+s}^{oil}] = 0 \,\forall t, s \geq 0$.[16] This rule implies that in Equation (14), $\tilde{\boldsymbol{\mathcal{A}}}$ becomes a selection matrix that selects the entries in $\boldsymbol{\Theta}_{\epsilon,A}$ and $\boldsymbol{\Theta}_{\nu,A}$ corresponding to the oil price. The remaining components of the equation have been estimated: the impulse response functions to a generic monetary policy shock (Figure 2) and the two oil-price news shocks (Figure 3. We then combine these impulse responses to approximate the solution of Equation (14) that characterizes the counterfactual scenario where OPEC stabilizes the oil price.[17]

The results of this exercise are shown in Figure 4. In the counterfactual scenario depicted by the golden line, the reduced responsiveness of the oil price to an ECB monetary policy shock leads to a much smaller response of energy prices in the euro area, consumer price inflation, and inflation expectations. Notably, the transmission of monetary policy to consumer prices is more than halved in the absence of an oil price response. Conversely, industrial production declines more than in the baseline scenario, as the stabilizing effect of the fall in the oil price is absent in the counterfactual scenario. This suggests that the observed reduction in the policy-induced decline in the HICP in the counterfactual scenario is not due to changes in the prices of domestically produced goods, but rather due to the inability of the ECB to affect global oil prices.

5.4 The inflation-output trade-off in the counterfactual scenario

The counterfactual analysis suggests that the ability to influence global energy prices may play a crucial role in the inflation-output trade-off faced by a central bank. To quantify this effect,

[16] It is important to note that the existence of such a policy rule for OPEC is not a new assumption in the literature. As already discussed by Leeper et al. (1996), the assumption that OPEC-related shocks, such as those identified for instance by Känzig (2021), exist is equivalent to the assumption that there is a policy rule that characterizes the systematic part of the corresponding equation (see Caldara and Kamps (2017) for a discussion).

[17] This implies that we condition on the point-estimates in Figure 2, which is consistent with standard practice in the policy counterfactual literature, which tends to take initial point estimates as given (see, e.g., Rotemberg and Woodford (1997), Eberly et al. (2020), Wolf (2023), McKay and Wolf (2023)). Given the results in Plagborg-Møller and Wolf (2021), we could equivalently run local projections for each identified shock and then combine these as done in Broer et al. (2024).

Figure 4: Monetary transmission if EA MP can (blue) and cannot (gold) affect oil prices

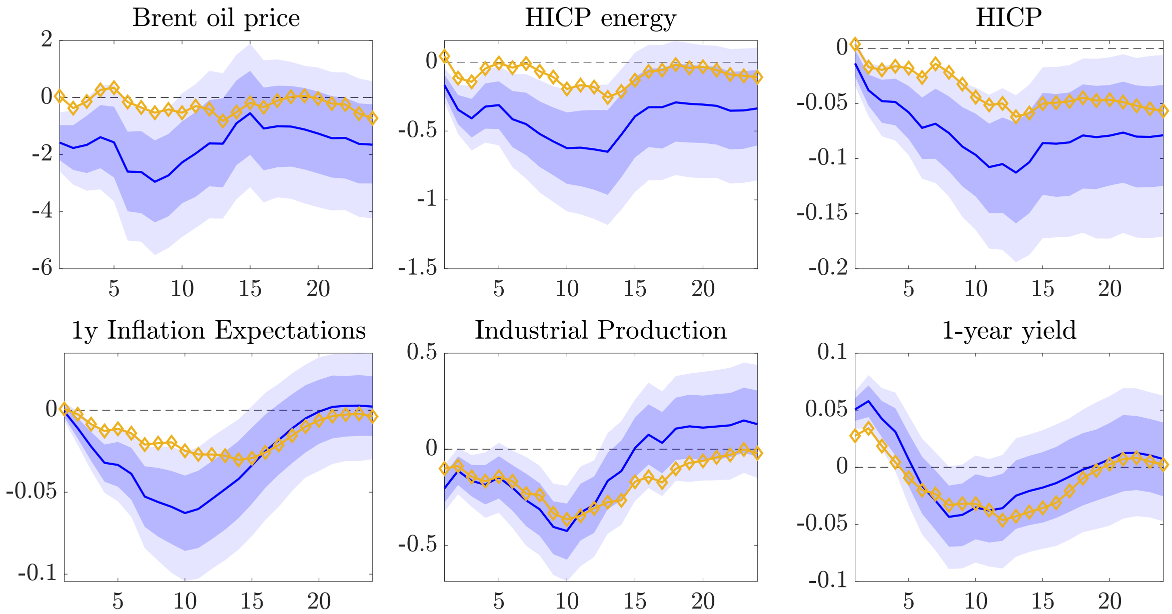

Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% point-wise credible sets in blue. Horizon in months. The golden line with diamonds shows the point-wise posterior means of the counterfactual where EA monetary policy does not affect the oil price. We approximate the solution to the counterfactual using the "best Lucas-Critique-robust approximation" of McKay and Wolf (2023), where we follow McKay and Wolf (2023) and condition on the point estimate to the monetary policy shock depicted in Figure 2. We report corresponding credible sets in Figure G.1 of the Appendix. Responses of the credit spread, exchange rate, and 5-year government bond yield are omitted to save space. See notes to Figure 2 for scaling of variables.

we follow a strategy very much akin to the non-parametric "Phillips-Multiplier" approach of Barnichon and Mesters (2021). We characterize this trade-off by measuring the average fall in inflation caused by a change in the monetary policy stance that lowers industrial production by 1% over the next h periods (see Appendix G.1 for details). In the spirit of Barnichon and Mesters (2021) we compute the "Output-Phillips-Multiplier" as

$$\mathcal{P}^h = \frac{\Theta^h_{\bar{\pi},\nu^{mp}}}{\Theta^h_{\bar{Y},\nu^{mp}}}, \tag{16}$$

where $\Theta^h_{\bar{Y},\nu^{mp}}$ ($\Theta^h_{\bar{\pi},\nu^{mp}}$) measures the horizon h impulse response of the average of industrial production (consumer price inflation) to a monetary policy shock ν^{mp}. We estimate this statistic for the baseline scenario and the counterfactual scenario. The results from this exercise are shown in Figure 5.

It becomes evident that the ECB's ability to influence energy prices alleviates the inflation-output trade-off substantially. For instance, engineering a 1% decline in industrial production

Figure 5: Inflation-Output Trade-off if EA MP can (blue) and cannot (gold) affect oil prices

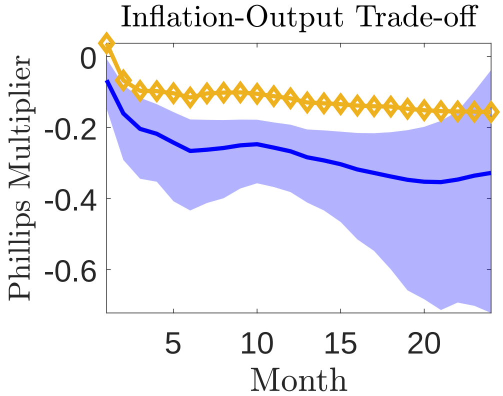

Notes: Point-wise median of the Output Phillips-Multiplier in blue. Counterfactual output Phillips-Multiplier under the assumption that EA monetary policy does not affect energy prices is depicted in gold. We only plot 68% credible sets to not distort the scale of the figure as the posterior distribution is very much skewed to the left.

over the next year is estimated to result in a 0.27% reduction in average inflation in the baseline scenario. This changes substantially in the counterfactual scenario, where the same policy-induced fall of industrial production only leads to a 0.12% reduction in average inflation. Thus, the ability of the ECB to affect global energy prices alleviates the inflation-output trade-off by approximately 55%.

Barnichon and Mesters (2021) demonstrates that if the true underlying model were characterized by a New-Keynesian Phillips Curve, this method would recover its slope with respect to changes in industrial production. Thus, when viewed through the lens of this model, our estimates suggest that this slope is much steeper when monetary policy can affect energy prices. Intuitively, and consistent with the microdata and the intuition outlined above, energy prices are updated much more frequently than other goods (implying a lower Calvo parameter), causing the slope of the aggregate Phillips curve to be steeper when monetary policy affects these goods.

6 Implications for the optimal policy response to supply shocks

The inflation-output trade-off faced by a central bank is particularly important when monetary policy faces a supply-side shock (Woodford (2003), Blanchard and Galí (2007), Galí (2015), Fornaro and Wolf (2023)). Our findings show that the ECB's ability to influence energy prices is a critical factor shaping this trade-off. To illustrate the significance of this result, we examine its implications for the optimal conduct of monetary policy. Specifically, we focus on a scenario in which the euro area faces an exogenous, supply-driven increase in oil prices. In this context, we consider two types of policy mandates: one focused solely on medium-term inflation stabilization, and another that targets the joint stabilization of both inflation and output. In both cases, we demonstrate how the ECB's mandate-optimal response depends on its capacity to affect energy prices.

Following MW, we approach this question empirically. First, we outline the general framework for computing the transmission of a shock under the mandate-optimal policy. Applying this framework to the question at hand necessitates the identification of a "generic" oil supply shock and its effects on the euro area economy, as well as mapping the ECB's mandate into a general loss function. These ingredients allow us to present the mandate-optimal policy responses under both mandate types and compare them to the corresponding responses in a counterfactual scenario where the ECB lacks the ability to influence energy prices.

6.1 Computing optimal policy counterfactuals

In this section, we outline the setup for computing the mandate-optimal policy responses to an exogenous shock. This includes the identification of the supply shock and two monetary policy shocks necessary for the analysis.

6.1.1 The framework to compute optimal policy counterfactuals

The approach of MW to estimating policy rule counterfactuals, discussed in Section 5.1, can be readily extended to compute impulse responses under the optimal policy. Specifically, in line with Barnichon and Mesters (2023), MW define the optimal policy response as the one that implements an allocation allowing the policymaker to optimally achieve its mandate. While this definition of optimality differs from the standard textbook definition, in which the policymaker seeks to maximize a measure of welfare —whose definition is inherently tied to a particular model and calibration— Barnichon and Mesters (2023) and McKay and Wolf (2022) convincingly argue that it is the "relevant objective-function for real-world central banks" (McKay and Wolf (2022, p.3)).

In particular, suppose the central bank minimizes the quadratic loss function of the form

$$\mathcal{L} = \frac{1}{2}\sum_{i=1}^{n_x} \lambda_i \boldsymbol{x}_i' W \boldsymbol{x}_i' = \frac{1}{2}\boldsymbol{x}'(\Lambda \otimes W)\boldsymbol{x} \tag{17}$$

where the $\boldsymbol{x}_i$ represents the time path of the endogenous variable i, λ_i describes the policy weights attached to that variable with $\Lambda = diag(\lambda_1, \lambda_2 .... \lambda_{n_x})$. The matrix $\boldsymbol{W}$ summarizes the effects of time discounting in the policymaker's preferences and can be (potentially) parameterized using a single discount factor β. MW show that the optimal policy problem can be stated in impulse-response space such that the loss function Equation (17) is minimized subject to Equation (11). In particular, the approach utilizes the observation that the implementable space of allocations for the endogenous variables $\boldsymbol{x}$ and for the policy instrument $\boldsymbol{z}$ is fully characterized by the impulse responses $\Theta_{\nu,\mathcal{A}}$ to the sequence of policy (news) shocks $\boldsymbol{\nu}$:

$$\begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{z} \end{bmatrix} = \Theta_{\nu,\mathcal{A}} \times \boldsymbol{\nu}. \tag{18}$$

The resulting optimality condition, which is substituted into the policy block of Equation (10), is given by:

$$\begin{aligned} \mathcal{A}^{\star}_{\boldsymbol{x}} &= (\lambda_1 \Theta'_{x_1,\nu,\mathcal{A}} W, \lambda_2 \Theta'_{x_2,\nu,\mathcal{A}} W, \ldots, \lambda_{n_x} \Theta'_{x_{n_x},\nu,\mathcal{A}} W), \\ \mathcal{A}^{*}_{\boldsymbol{z}} &= \mathbf{0}, \end{aligned} \tag{19}$$

where $\Theta_{x_i,\nu,\mathcal{A}}$ is the matrix of impulse responses of variable i to all shocks in $\boldsymbol{\nu}$ under the estimated policy rules ($\boldsymbol{\mathcal{A}}$). [18] Although numerically equivalent to previous approaches in the literature (see, e.g., Svensson (1997)), here, the optimal policy rule is fully characterized by impulse responses to policy (news) shocks, all of which can be, in theory, estimated from the data.

In practice, researchers may not be able to identify the entire menu of policy shocks $\boldsymbol{\nu}$ but only a subset ($\widetilde{\boldsymbol{s}} \subset \boldsymbol{\nu}$). In this case, the set of hypothetical, feasible allocations can then no longer be described by Equation (18) but is rather given by

$$\boldsymbol{y} = \Theta_{\nu,\mathcal{A}} \times \widetilde{\boldsymbol{s}} \tag{20}$$

with $\boldsymbol{y} = (\boldsymbol{x}', \boldsymbol{z}')'$. In such a scenario, the monetary policy authority chooses the optimal policy rule (Equation (19)) and the corresponding allocation of $\boldsymbol{y}$ that minimizes the loss function in Equation (17) within the empirically identified space described by Equation (20).

6.1.2 Ingredient 1: two loss functions

To apply this framework to the question at hand, the first key component is to determine the ECB's relevant loss function. First, we follow McKay and Wolf (2022) and Barnichon and Mesters (2023) by conservatively deriving the loss function from the central bank's primary mandate. This approach provides a simple benchmark to evaluate the role of energy prices and their response to European monetary policy in the ECB's mandate-optimal conduct of monetary policy and illustrates the key mechanism. Second, we incorporate the secondary mandate into the analysis.

The ECB's primary mandate is to maintain price stability, which it defines as an inflation target of 2% over the medium term. Therefore, in our loss function, we aim to minimize the deviations of contemporaneous and future HICP inflation from the steady state. Although there is no clear definition of the medium-term horizon, there is good evidence based on

[18] Detailed derivations for this result can be found in Appendix H.1

the ECB's own projections that in practice, the relevant horizon corresponds to 6-8 quarters (Paloviita et al. (2021)). Therefore, we model the ECB's focus on the medium term by giving a higher weight to the inflation deviations that are present 6-8 quarters after the initial shock. Out of these considerations, the loss function takes the following form:

$$\mathcal{L} = \lambda_\pi \boldsymbol{\pi}' W \boldsymbol{\pi}, \tag{21}$$

with $\lambda_\pi = 1$. The weighting matrix is defined as $W = (\text{diag}(\beta^{24}, \ldots \beta^2, \beta, 1))$.[19] Additionally, $\boldsymbol{\pi} = \mathcal{D}\boldsymbol{P}^{\text{HICP}}$ represents the transformed impulse responses of the (log) level of the HICP, denoted as $\boldsymbol{P}^{\text{HICP}}$. The operator $\mathcal{D}$ appropriately converts these impulse responses to changes in year-on-year inflation rates. Furthermore, we set the discount factor β such that, in a standard New Keynesian model, the corresponding annualized real interest rate would be 2%.[20]

While instructive for fleshing out the key mechanism, the highly stylized loss function that is solely based on the ECB's primary mandate implies that the economic costs of the additional tightening are of no concern to the policymaker in the scenario. Therefore, as a next step, we incorporate the secondary mandate into the central banks' objective function. The secondary mandate loosely states that the ECB should " [...] support the general economic policies of the Union [...]" (*Treaty on the Functioning of the European Union* (2016, §127(1))). As one of these goals is "balanced economic growth" (*Treaty on European Union* (2012, §3(3))), an arguably simplified interpretation of the secondary mandate is that the ECB should also aim to stabilize economic activity. Albeit stylized, this allows us to integrate the inflation-output trade-off faced by the central bank into the loss function by adding an additional term containing the deviations of output as measured by deviations of GDP from trend $\boldsymbol{y}$.[21]

$$\mathcal{L} = \lambda_\pi \boldsymbol{\pi}' W \boldsymbol{\pi} + \lambda_y \boldsymbol{y}' W \boldsymbol{y}, \tag{22}$$

To best flesh out the differences between a dual and single-mandate loss function, we assume that in this scenario the ECB cares equally about output and inflation, implying that we set $\lambda_y = \lambda_\pi$.[22]

[19] This is only a linear approximation to the weighting problem, where the deviation at the last horizon (24 months) has the highest weight and the weight of deviations increases linearly. A quadratic approximation would not change the result significantly.

[20] We note that, given the focus on a single objective $\boldsymbol{\pi}$, the weighting matrix should not matter in theory, if in our application, the ECB were to operate in the fully unconstrained space of implementable allocations of Equation (18). Intuitively, if the ECB in our application had perfect knowledge of and perfect access to all 24 instruments (shocks) $\boldsymbol{\nu}$, it could perfectly stabilize the 24 targets that have a positive weight in the loss function. Since we do not fully identify the full menu of policy shocks $\boldsymbol{\nu}$, we restrict the set of possible allocations that can be implemented to the space of empirically identified policy shock paths, which implies that the weighting matrix matters because the central bank in our application lacks the tools to perfectly stabilize inflation at all horizons.

[21] We assume that the trend is unaffected by oil supply shocks and the monetary policy rule.

[22] To map industrial production deviations from our SVAR into GDP deviations, we scale the hypothetical equal weight that we want to give to GDP λ_{GDP} by the relative variance of GDP and industrial production

6.1.3 Ingredient 2: an identified oil supply shock

The second key component in estimating the ECB's mandate-optimal response to a supply shock is the estimated impulse response functions to such a shock. As in Section 5.3, we identify an oil supply shock following Känzig (2021). Specifically, we use high-frequency changes in oil price futures to identify OPEC-related changes in oil supply. This time, however, we identify only one "generic" oil supply news shock and do not distinguish between short- and medium-run oil supply news. See Inoue and Rossi (2021) and Caravello et al. (2023) for a similar interpretation of identified monetary policy shocks as a linear combination of underlying shocks. Just as we did for the monetary policy shock in Section 4, we employ the first principal component ($m^{OIL}_{t,PC1}$) of the changes in oil price futures from one month up to one year. In the language of our identifying assumptions framework in Equation (5) this implies that we set $\epsilon^*_t \equiv \epsilon^{oil}_{t,generic}$ and $m_t \equiv m^{OIL}_{t,PC1}$. We interpret this shock as news about oil supply policies at different maturities that, in combination, shift the current and expected future price of oil (Inoue and Rossi (2021), McKay and Wolf (2023)).

The resulting impulse responses, depicted in blue in Figure 7, illustrate the transmission of an average, one standard deviation oil supply news shock. Consistent with Känzig (2021), the oil supply shock leads to a long-lasting increase in the Brent oil price, resulting in higher consumer energy prices, inflation, and inflation expectations. Importantly, this rise in energy costs and inflation is not just a short-term phenomenon but extends into the medium term. Additionally, the shock creates an output-inflation trade-off by inducing a delayed yet significant economic contraction. In line with the conventional wisdom that "in the past, central banks have typically looked through energy shocks" (Schnabel (2022)), the estimated monetary policy response does not appear to counteract the inflationary pressures. Notably, it not only tolerates the increase in energy prices and inflation but even slightly lowers interest rates, possibly to mitigate the economic downturn.

6.1.4 Ingredient 3: two identified monetary policy news shocks

As the final ingredient, we need to identify two euro area monetary policy shocks to approximate the solution to the counterfactual in Equation (20). Specifically, we utilize high-frequency changes in the 3-month ($m^{MP}_{t,3m}$) and 24-month ($m^{MP}_{t,24m}$) futures to identify a short-term ($\nu^{MP}_{t,short}$) and a medium-term ($\nu^{MP}_{t,medium}$) monetary policy news shock. These shocks can be interpreted as shifting different segments of the yield curve.[23] Both shocks are incorporated into the framework in Section 3, Equation (5), as follows: we define $\boldsymbol{\epsilon}^*_t \equiv (\epsilon_{1,t}, \epsilon_{2,t})' = (\nu^{MP}_{t,short}, \nu^{MP}_{t,medium})'$ and $\boldsymbol{m}^*_t \equiv (m_{1,t}, m_{2,t})' = (m^{MP}_{t,3m}, m^{MP}_{t,24m})'$. The impulse responses to these

(IP) such that $\lambda_y = \lambda_{\text{IP}} = \lambda_{\text{GDP}} \times \frac{\sigma^2(\text{GDP})}{\sigma^2(\text{IP})}$.

[23] Given our identifying assumptions, the "generic" monetary policy shock that we identify in Section 4 should best be thought of as a linear combination of these two shocks. See Inoue and Rossi (2021) and Caravello et al. (2023) for a similar interpretation of identified monetary policy shocks as a linear combination of underlying shocks.

monetary policy shocks are presented in Figure 6, aligning with the results in the empirical literature. In summary, the short-term monetary policy shock primarily affects the short end of the yield curve, consistent with a conventional monetary policy shock. Conversely, the medium-term monetary policy shock influences both short-term (1-year) and long-term (5-year) yields, resembling a forward guidance type shock. Both shocks exhibit similar qualitative patterns, and their effects are consistent with existing empirical evidence for these types of shocks (Swanson (2024), Ricco et al. (2025), Miranda-Agrippino and Ricco (2023), Lakdawala (2019); see Rossi (2021) for a comprehensive literature review).

Figure 6: Response to a short-run (crimson red) and medium-run (sky blue) EA MP shock

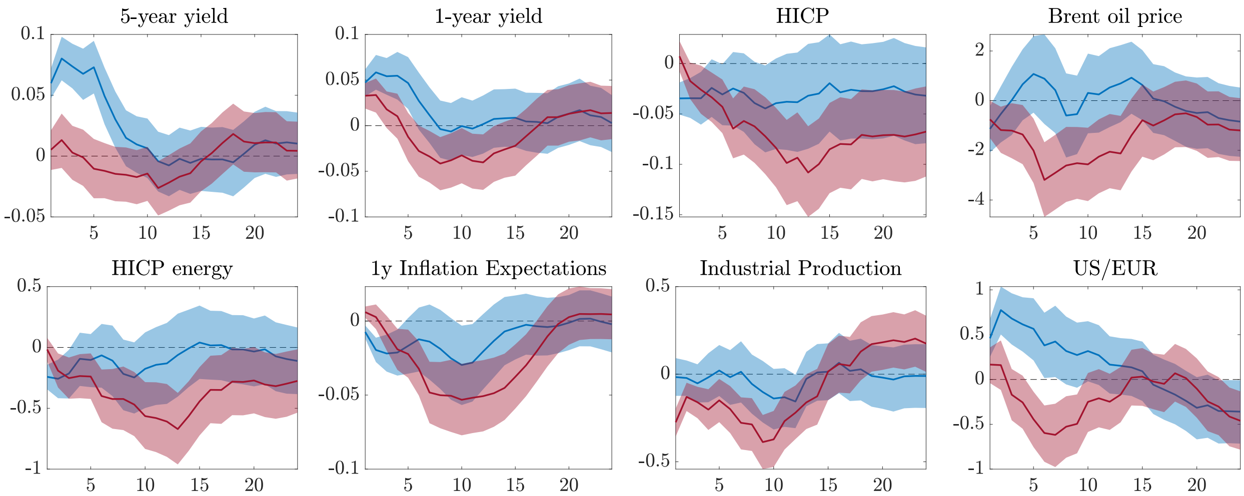

Notes: Impulse responses to a short-run (medium-run) EA monetary policy shock in crimson red (sky blue) alongside 68% credible sets. We normalize the responses such that the 1-year (5-year) yield increases on impact. Response of the credit spread is omitted to save space. See notes to Figure 2 for scaling of variables.

6.2 Oil supply shock transmission under primary-mandate-optimal policy

Equipped with the oil-supply shock $\epsilon_t = \epsilon^{oil}_{t,generic}$ and the two policy shocks $\widetilde{\boldsymbol{s}} = [\nu^{MP}_{t,short}, \nu^{MP}_{t,medium}]$ we compute the transmission of an oil supply shock under the assumption that the ECB aims to optimally achieve its primary mandate described by Equation (21) subject to the space of all possible allocation it can achieve as characterized by Equation (20). The results from this exercise are depicted by the black circled lines and bars in Figure 7.

Contrary to the observed empirical response (blue line), the primary-mandate-optimal response (black line) does not suggest that the ECB should lower short- and longer-term interest rates. Instead, to optimally stabilize inflation, the ECB under this scenario quickly raises interest rates to counteract the inflationary effects of the 3.5% increase in oil prices. It is important to emphasize that, despite the stylized loss function focusing solely on the primary mandate, our estimates do not support an excessive rate hike by the ECB following an oil supply shock. Nevertheless, this slight change in the policy stance stabilizes approximately 72% of the oil supply shock-induced deviations of inflation from target (first panel in the last row). The cost, relative to the scenario under the empirical monetary policy rule, is a front-loaded contraction in output. However, this strategy not only optimally stabilizes medium-term inflation and inflation expectations (see Figure H.1) but the initial additional output contraction is offset with higher output in the medium term. Therefore, on average, the economic contraction is only 10% higher under the primary-mandate optimal response (second panel in the last row).

One possible explanation for this relatively benign outcome is the quick and strong response of energy prices to a monetary contraction. As documented in previous sections, energy prices are comparatively flexible and react much faster and more strongly to changes in demand than other domestically produced goods in the HICP basket. As a result, to fulfill its primary mandate in the face of an exogenous increase in energy prices, the ECB does not need to tighten excessively and persistently, thereby possibly inducing a major output contraction. In fact, a large part of the adjustment is borne by relatively flexible oil/energy prices, which on average remain 30% below the corresponding path under the empirical policy rule (third panel in the last row). In the next section, we examine this hypothesis in more detail.

6.3 Primary-mandate-optimal policy response when the ECB does not influence energy prices

In this section, we substantiate the notion that the response of global energy prices to euro area monetary policy allows the ECB to optimally achieve its mandate with only limited increases in interest rates. To explore this, we conduct a thought experiment: What would the optimal monetary policy response be to the same exogenous increase in Brent oil prices, if the ECB could not influence global energy prices?

Specifically, we are interested in the optimal allocation $\boldsymbol{y}$ when the empirically identified, implementable space of possible allocations is not described by Equation (20), but rather by

$$\boldsymbol{y} = \Theta_{\nu,\widetilde{\mathcal{A}}} \times \widetilde{\boldsymbol{s}} \tag{23}$$

Here, the subscript $\widetilde{\mathcal{A}}$ indicates that the space is now characterized by counterfactual impulse responses. We construct these counterfactual impulse responses for the two identified monetary policy shocks $\widetilde{\boldsymbol{s}}$ in a manner analogous to our approach in Section 5.1.[24] The propagation of an oil supply shock under the resulting counterfactual optimal policy rule is depicted by

[24]A detailed step-by-step summary of our approach to estimating the optimal monetary policy response to an oil supply shock, under the assumption that the ECB's decisions do not affect global oil prices, can be found in Appendix H.2.

Figure 7: Oil supply shock transmission under different (optimal) monetary policy rules

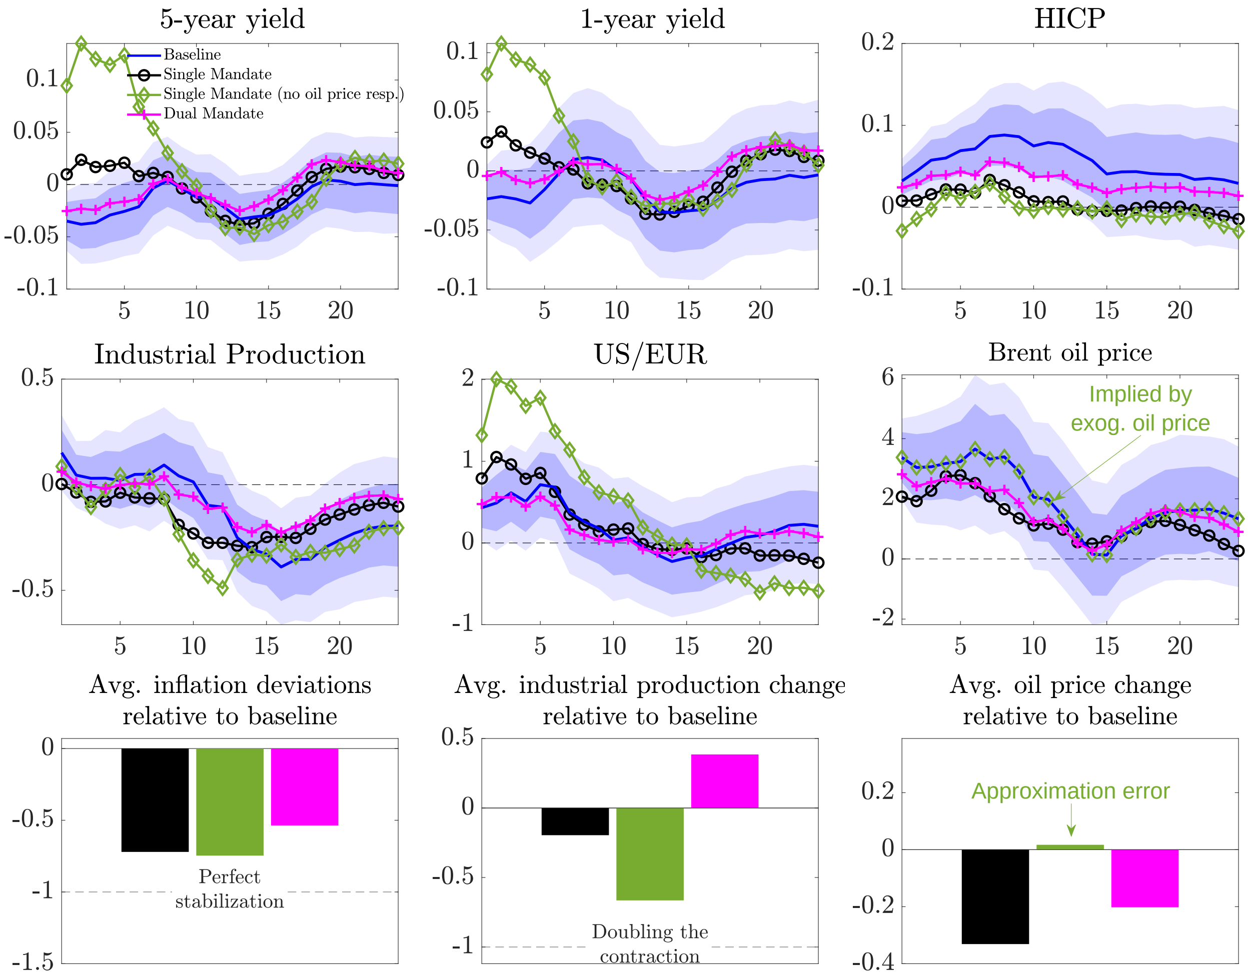

Notes: Impulse response functions to a one standard deviation oil supply shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. The black circled lines (magenta crossed lines) show the responses of the endogenous variables under optimal policy with a single mandate (dual mandate) loss function described in Equation (21) ((22)). The green lines with diamonds show the corresponding estimate for the single mandate loss function under the counterfactual assumption that the ECB's monetary policy decisions do not affect the oil price. This implies that in this scenario, the oil price path is exogenous and does not change in response to a change in the ECB's policy stance. Therefore, if it weren't for the approximation error, the green and blue lines for the oil price panel would coincide. To avoid cluttering, we report the time series of approximation errors in Figure H.2 of the Appendix. We measure the average policy implied changes for the variables in the bottom row as $\sum(|x_{t,\mathcal{A}^\star}| - |x_{t,\mathcal{A}}|)/\sum|x_{t,\mathcal{A}}|$, where x_t is the impulse response of the variable and the superscript $\mathcal{A}^\star$ indicates the counterfactual policy rule. Because the central bank aims to stabilize inflation, we define inflation deviations as absolute deviations from the target $|\pi_t|$ for ease of interpretation. Figure H.1 of the Appendix contains the responses of the remaining variables. See notes to Figure 2 for scaling of variables.

the green lines and bars in Figure 7.

Note that in this application, the oil price remains unaffected by changes in the conduct of monetary policy, so its path, up to approximation error, is the same as in the scenario under the empirical policy rule (third panel in the last row). When comparing the impulse response functions under primary-mandate-optimal policy from the last subsection (black) and those in the counterfactual scenario (green), it becomes clear that monetary policy needs to tighten significantly more to stabilize inflation when it does not influence global energy prices.[25] When cumulated over all periods, the additional tightening required to go from the observed interest rate path to the optimal interest rate path is more than three times as large. By implementing this policy, the ECB can again avoid approximately 72% of the induced deviations of inflation (first panel in the last row). However, the economic costs necessary to achieve this are much more severe in this scenario. Specifically, relative to the allocation under the empirical policy rule (blue line), the economic contraction is, on average, roughly 65% more severe. Most importantly, the costs in terms of economic activity that are required to optimally stabilize medium-term inflation in the face of an oil supply shock are more than three times as large when compared to the scenario where the ECB is able to affect the oil price and can therefore directly address the root cause of the shock (second panel in the last row).

6.4 Oil supply shock transmission under dual-mandate-optimal policy

To analyze the optimal policy response under a dual-mandate, we redo the previous analysis and compute the optimal policy response to an oil supply shock, taking into account the ECB's impact on energy prices. The results are depicted by the magenta lines and bars in Figure 7.

Two results are particularly noteworthy. First, by construction, under this loss function, the ECB aims to stabilize inflation and output (first and second panels in the last row). To do so, the optimal policy response does not entail raising interest rates quickly right away, but, when compared to the empirically observed response of the ECB, it rather only calls for slightly higher interest rates at the short end of the yield curve. As such, the observed response of the ECB is much closer to the estimated optimal response under the dual-mandate loss function than the single-mandate.

Second, we document that this benign result can again be traced back to the role of energy prices in the monetary transmission. As shown in Figure 8, even under the dual-mandate loss function, the additional tightening required to reach the optimal interest rate path is much higher. At the same time, the loss as measured by deviations of inflation and output from the target is significantly higher, since the ECB's inflation-output trade-off worsens. This underscores that the ability to influence energy prices is crucial not only for monetary transmission but is also of particular importance for policymaking in the face of an energy supply shock.

[25] In Figure H.3 in Appendix H, we show that this result holds even with a dual mandate loss function, where the ECB equally considers deviations of inflation and output.

Figure 8: The role of energy prices for monetary policy under a dual-mandate

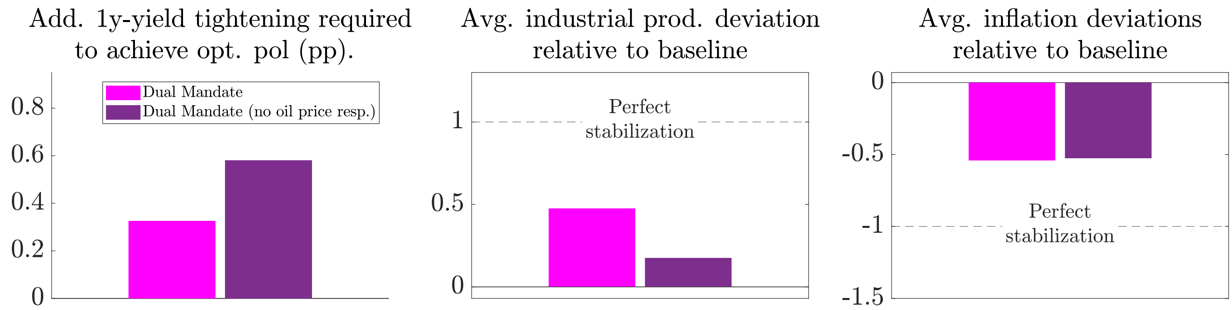

Notes: See notes to Figure 7. We define the additional tightening required as $\sum(i_{t,\mathcal{A}^\star} - i_{t,\mathcal{A}})$, where i_t is the impulse response of the interest rate and the superscript $\mathcal{A}^\star$ indicates the counterfactual policy rule. Note that we treat industrial production (output) as a policy target. Therefore, we define output deviations from the target as $|y|$ and treat them along the lines of inflation deviations in Figure 7. The impulse responses underlying these bar charts can be found in Figure H.3 in the Appendix.

7 Conclusion

This paper examines the influence of European monetary policy on energy prices and challenges the prevalent view that the ECB has limited capacity to combat energy-price-driven inflation. Employing a high-frequency event study we find that ECB policy decisions significantly affect global energy prices, such as the Brent oil price and the natural gas price. SVAR analysis corroborates a strong and persistent effect of monetary policy on oil prices as well as on consumer energy prices, inflation, and inflation expectations. Using Lucas critique-robust counterfactuals along the lines McKay and Wolf (2023), we establish that a substantial part of the impact of monetary policy on headline consumer prices is transmitted via energy prices. In particular, it is precisely the ECB's influence on relatively flexible energy prices that accounts for a major part of monetary policy's ability to affect inflation in the short- and medium-run. Furthermore, we estimate that the inflation-output trade-off faced by the ECB is alleviated by around 50% compared to a counterfactual in which energy prices are unaffected by the ECB's policy decisions.

We illustrate the importance of our findings by studying the optimal conduct of monetary policy in the face of a supply-side shock using the same empirical counterfactual method. Specifically, we compute the mandate-optimal response for the ECB to an oil supply shock both with the estimated effect of monetary policy on energy prices and a corresponding counterfactual scenario in which energy prices are unaffected by changes in monetary policy. Crucially, the mandate-optimal policy is much more contractionary once the ECB has no control over energy prices since the inflation-output trade-off is much more severe. In this scenario, the costs in terms of economic activity to stabilize inflation are on average more than twice as large, highlighting the key role played by energy prices in the conduct of monetary policy. It is in that sense, that energy prices can be considered a friend to and not a foe of central banks.

References

Aastveit, Knut Are, Hilde C Bjørnland, and Jamie L Cross (2023). "Inflation expectations and the pass-through of oil prices". *Review of Economics and Statistics 105* (3), 733–743.

Altavilla, Carlo, Luca Brugnolini, Refet S Gürkaynak, Roberto Motto, and Giuseppe Ragusa (2019). "Measuring euro area monetary policy". *Journal of Monetary Economics 108*, 162–179.

Andrade, Philippe and Filippo Ferroni (2021). "Delphic and odyssean monetary policy shocks: Evidence from the euro area". *Journal of Monetary Economics 117*, 816–832.

Anzuini, Alessio, Marco J Lombardi, and Patrizio Pagano (2012). "The Impact of Monetary Policy Shocks on Commodity Prices". *International Journal of Central Banking.*

Arias, Jonas, Juan F Rubio-Ramirez, and Daniel F Waggoner (2023). "Uniform priors for impulse responses".

Arias, Jonas E., Juan F. Rubio-Ramírez, and Daniel F. Waggoner (2021). "Inference in Bayesian Proxy-SVARs". *Journal of Econometrics 225* (1), 88–106.

Auclert, Adrien, Hugo Monnery, Matthew Rognlie, and Ludwig Straub (2023). "Managing an Energy Shock: Fiscal and Monetary Policy". *NBER Working Paper 31543.*

Aucremanne, Luc and Emmanuel Dhyne (2004). "How frequently do prices change? Evidence based on the micro data underlying the Belgian CPI". *ECB Working Paper Series No. 331.*

Bachmann, Ruediger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick (2022). "What if? The economic effects for Germany of a stop of energy imports from Russia". *ECONtribute Policy Brief 028.*

Badinger, Harald and Stefan Schiman (2023). "Measuring monetary policy in the Euro area using SVARs with residual restrictions". *American Economic Journal: Macroeconomics 15* (2), 279–305.

Barnichon, Regis and Geert Mesters (2021). "The phillips multiplier". *Journal of Monetary Economics 117*, 689–705.

Barnichon, Régis and Geert Mesters (2023). "A Sufficient Statistics Approach for Macro Policy". *American Economic Review 113* (11), 2809–2845.

——— (2024). "Evaluating policy institutions-150 years of US monetary policy". Tech. rep.

Basistha, Arabinda and Alexander Kurov (2015). "The impact of monetary policy surprises on energy prices". *Journal of Futures Markets 35* (1), 87–103.

Bauer, Michael D and Eric T Swanson (2023). "A reassessment of monetary policy surprises and high-frequency identification". *NBER Macroeconomics Annual 37* (1), 87–155.

Baumeister, Christiane and James D Hamilton (2019). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks". *American Economic Review 109* (5), 1873–1910.

Bayer, Christian, Alexander Kriwoluzky, Gernot J Müller, and Fabian Seyrich (2023). "Hicks in hank: Fiscal responses to an energy shock".

Beechey, Meredith J and Jonathan H Wright (2009). "The high-frequency impact of news on long-term yields and forward rates: Is it real?" *Journal of Monetary Economics 56* (4), 535–544.

Bernanke, Ben and Olivier Blanchard (2024). "An analysis of pandemic-era inflation in 11 economies". *Peterson Institute for International Economics Working Paper*, 24–11.

Bernanke, Ben S, Mark Gertler, Mark Watson, Christopher A Sims, and Benjamin M Friedman (1997). "Systematic monetary policy and the effects of oil price shocks". *Brookings papers on economic activity 1997* (1), 91–157.

Blanchard, Olivier and Jordi Galí (2007). "Real wage rigidities and the New Keynesian model". *Journal of money, credit and banking 39*, 35–65.

Bodenstein, Martin, Luca Guerrieri, and Lutz Kilian (2012). "Monetary policy responses to oil price fluctuations". *IMF Economic Review 60* (4), 470–504.

Born, Benjamin and Johannes Pfeifer (2021). "Uncertainty-driven business cycles: Assessing the markup channel". *Quantitative economics 12* (2), 587–623.

Braun, Robin, Silvia Miranda-Agrippino, and Tuli Saha (2024). "Measuring monetary policy in the UK: The UK monetary policy event-study database". *Journal of Monetary Economics*, 103645.

Broer, Tobias, John Kramer, and Kurt Mitman (2024). "The distributional effects of oil shocks".

Buda, Gergely, Vasco M Carvalho, Giancarlo Corsetti, João B Duarte, Stephen Hansen, Álvaro Ortiz, Tomasa Rodrigo, and José V Rodríguez Mora (2023). "Short and variable lags". *Robert Schuman Centre for Advanced Studies Research Paper* (22).

Caldara, Dario and Edward Herbst (2019). "Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs". *American Economic Journal: Macroeconomics 11* (1), 157–92.

Caldara, Dario and Christophe Kamps (2017). "The analytics of SVARs: a unified framework to measure fiscal multipliers". *The Review of Economic Studies 84* (3), 1015–1040.

Caravello, Tomás, Alisdair McKay, and Christian Wolf (2023). "Evaluating Policy Counterfactuals: A "VAR-Plus" Approach". *mimeo.*

Cascaldi-Garcia, Danilo (2022). "Pandemic priors". *International Finance Discussion Paper* (1352).

Castelnuovo, Efrem, Lorenzo Mori, and Gert Peersman (2024). "Commodity Price Shocks and Global Cycles: Monetary Policy Matters".

Cesa-Bianchi, Ambrogio, Gregory Thwaites, and Alejandro Vicondoa (2020). "Monetary policy transmission in the United Kingdom: A high frequency identification approach". *European Economic Review 123*, 103375.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans (1999). “Monetary policy shocks: What have we learned and to what end?” *Handbook of macroeconomics 1*, 65–148.

Coenen, Günter, Peter Karadi, Sebastian Schmidt, and Anders Warne (2018). “The New Area-Wide Model II: an extended version of the ECB’s micro-founded model for forecasting and policy analysis with a financial sector”.

Corsetti, Giancarlo, Keith Kuester, Gernot J Müller, and Sebastian Schmidt (2024). “Exchange rate insulation revisited”.

Degasperi, Riccardo (2023). “Identification of expectational shocks in the oil market using OPEC announcements”. *mimeo.*

Degasperi, Riccardo, Seokki S Hong, and Giovanni Ricco (2023). “The Global Transmission of US Monetary Policy”.

Eberly, Janice C, James H Stock, and Jonathan H Wright (2020). “The federal reserve’s current framework for monetary policy: A review and assessment”. *International Journal of Central Banking, 16*, 5—71.

Fornaro, Luca and Martin Wolf (2023). “The scars of supply shocks: Implications for monetary policy”. *Journal of Monetary Economics 140*, S18–S36.

Gagliardone, Luca and Mark Gertler (2023). “Oil Prices, Monetary Policy and Inflation Surges”. Tech. rep. National Bureau of Economic Research.

Galí, Jordi (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press.

Gazzani, Andrea Giovanni and Fabrizio Ferriani (2024). “The International Transmission of Chinese Monetary Policy and the Commodity Channel”. *Available at SSRN 4961268.*

Georgiadis, Georgios, Gernot J Müller, and Ben Schumann (2024). “Global risk and the dollar”. *Journal of Monetary Economics 144*, 103549.

Gertler, Mark and Peter Karadi (2015). “Monetary policy surprises, credit costs, and economic activity”. *American Economic Journal: Macroeconomics 7* (1), 44–76.

Giacomini, Raffaella and Toru Kitagawa (2021). “Robust Bayesian inference for set-identified models”. *Econometrica 89* (4), 1519–1556.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson (2005). “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements”. *International Journal of Central Banking 1* (1), 55–93.

Inoue, Atsushi and Lutz Kilian (2021). “The role of the prior in estimating VAR models with sign restrictions”. *Center for Financial Studies Working Paper* (660).

Inoue, Atsushi and Barbara Rossi (2021). “A new approach to measuring economic policy shocks, with an application to conventional and unconventional monetary policy”. *Quantitative Economics 12* (4), 1085–1138.

Jarociński, Marek and Peter Karadi (2020). “Deconstructing monetary policy surprises—the role of information shocks”. *American Economic Journal: Macroeconomics 12* (2), 1–43.

Jo, Yoon Joo and Ben Klopack (2024). "Fueling Expectations: The Causal Impact of Gas Prices on Inflation Expectations".

Känzig, Diego R (2021). "The macroeconomic effects of oil supply news: Evidence from OPEC announcements". *American Economic Review 111* (4), 1092–1125.

——— (2023). "The unequal economic consequences of carbon pricing". Tech. rep. National Bureau of Economic Research.

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

Kilian, Lutz and Logan T Lewis (2011). "Does the Fed respond to oil price shocks?" *The Economic Journal 121* (555), 1047–1072.

Lagarde, Christin (2022). *Transcript of the ECB Press conference 3 February.* 3 February, Frankfurt am Main. URL: \url{https://www.ecb.europa.eu/press/pressconf/2022/html/ecb.is220203~ca7001dec0.en.html#qa}.

Lakdawala, Aeimit (2019). "Decomposing the effects of monetary policy using an external instruments SVAR". *Journal of applied econometrics 34* (6), 934–950.

Leduc, Sylvain and Keith Sill (2004). "A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns". *Journal of Monetary Economics 51* (4), 781–808.

Leeper, Eric M, Christopher A Sims, Tao Zha, Robert E Hall, and Ben S Bernanke (1996). "What does monetary policy do?" *Brookings papers on economic activity 1996* (2), 1–78.

Mankiw, N Gregory (2001). "The inexorable and mysterious tradeoff between inflation and unemployment". *The Economic Journal 111* (471), 45–61.

McKay, Alisdair and Christian Wolf (2022). "Optimal policy rules in HANK". *Working Paper, FRB Minneapolis.*

——— (2023). "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica 91* (5), 1695–1725.

Mertens, Karel and Morten O Ravn (2013). "The dynamic effects of personal and corporate income tax changes in the United States". *American economic review 103* (4), 1212–47.

Miranda-Agrippino, Silvia and Tsvetelina Nenova (2022). "A tale of two global monetary policies". *Journal of International Economics 136*, 103606.

Miranda-Agrippino, Silvia and Hélene Rey (2020). "US monetary policy and the global financial cycle". *The Review of Economic Studies 87* (6), 2754–2776.

Miranda-Agrippino, Silvia and Giovanni Ricco (2021). "The transmission of monetary policy shocks". *American Economic Journal: Macroeconomics 13* (3), 74–107.

——— (2023). "Identification with external instruments in structural VARs". *Journal of Monetary Economics 135*, 1–19.

Miranda-Pinto, Jorge, Andreas Pescatori, Ervin Prifti, and Guillermo Verduzco-Bustos (2023). "Monetary Policy Transmission through Commodity Prices". *IMF Working Paper, 23/215.*

Nakamura, Emi and Jón Steinsson (2018). "High-frequency identification of monetary non-neutrality: the information effect". *The Quarterly Journal of Economics 133* (3), 1283–1330.

Natal, Jean-Marc (2012). "Monetary policy response to oil price shocks". *Journal of Money, Credit and Banking 44* (1), 53–101.

Paloviita, Maritta, Markus Haavio, Pirkka Jalasjoki, and Juha Kilponen (2021). "What does "below, but close to, 2 percent" mean? Assessing the ECB's reaction function with real-time data". *International Journal of Central Banking 17* (2), 125–169.

Plagborg-Møller, Mikkel and Christian K Wolf (2021). "Local projections and VARs estimate the same impulse responses". *Econometrica 89* (2), 955–980.

Ricco, Giovanni, Emanuele Savini, and Anshumaan Tuteja (2025). "Monetary Policy, Information and Country Risk Shocks in the Euro Area". *CEPR Discussion Paper 1967*.

Rosa, Carlo (2014). "The high-frequency response of energy prices to US monetary policy: Understanding the empirical evidence". *Energy Economics 45*, 295–303.

Rossi, Barbara (2021). "Identifying and estimating the effects of unconventional monetary policy: How to do it and what have we learned?" *The Econometrics Journal 24* (1), C1–C32.

Rotemberg, Julio J and Michael Woodford (1997). "An optimization-based econometric framework for the evaluation of monetary policy". *NBER macroeconomics annual 12*, 297–346.

Rubio-Ramirez, Juan F, Daniel F Waggoner, and Tao Zha (2010). "Structural vector autoregressions: Theory of identification and algorithms for inference". *The Review of Economic Studies 77* (2), 665–696.

Schnabel, Isabel (2022). *Looking through Higher Energy Prices? Monetary Policy and the Green Transition.* Speech at a panel on "Climate and the Financial System" at the American Finance Association 2022 Virtual Annual Meeting, 8 January 2022.

Soriano, Pilar and Hipolit Torró (2022). "The response of Brent crude oil to the European central bank monetary policy". *Finance Research Letters 46*, 102353.

Svensson, Lars EO (1997). "Inflation forecast targeting: Implementing and monitoring inflation targets". *European economic review 41* (6), 1111–1146.

Swanson, Eric T (2024). "The Macroeconomic Effects of the Federal Reserve's Conventional and Unconventional Monetary Policies". *IMF Economic Review*.

Ter Ellen, Saskia, Edvard Jansen, and Nina Larsson Midthjell (2020). "ECB Spillovers and domestic monetary policy effectiveness in small open economies". *European Economic Review 121*, 103338.

Torro, Hip Lit (2019). "The Response of European Energy Prices to ECB Monetary Policy". *International Journal of Energy Economics and Policy 9* (2), 1–9.

Treaty on European Union (2012). Article 3(3).

Treaty on the Functioning of the European Union (2016). Article 127(1).

Wehrhöfer, Nils (2023). “Energy Prices and Inflation Expectations: Evidence from Households and Firms”. *Mimeo.*

Wolf, Christian K (2023). “Fiscal stimulus and the systematic response of monetary policy”. *AEA Papers and Proceedings.* Vol. 113. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 388–393.

Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

A Data description

Table A.1: Detailed description of data used in the high-frequency event study regressions

Variable	Description	Notes	Source
Global oil price	Brent crude oil front-month futures (LCOc1) price (in US dollars)	Computed the percent price change around monetary policy announcements (ECB, Fed and BoE)	Refinitiv Tick History database
ECB monetary policy surprise	3-month OIS rate changes around ECB monetary policy announcements	Computed based on methodology of Jarociński and Karadi (2020)	EA-MPD from Altavilla et al. (2019)
Fed monetary policy surprise (baseline)	3-month-ahead federal funds future rate changes around FOMC announcements	Computed based on methodology of Jarociński and Karadi (2020)	Gürkaynak et al. (2005) and Marek Jarocinski's website
Fed monetary policy surprise (robustness)	The first principal component of the changes in ED1–ED4 around FOMC announcements	Orthogonalized monetary policy surprises uncorrelated with macroeconomic and financial data observed before FOMC announcements	Bauer and Swanson (2023)
Bank of England monetary policy surprise	3-month Libor rate changes around Bank of England monetary policy announcements	Computed based on methodology of Jarociński and Karadi (2020)	Cesa-Bianchi et al. (2020)
FTSE 100 index	FTSE 100 index price changes around Bank of England monetary policy announcements	Computed from tick data	Refinitiv Tick History database
Dutch TTF natural gas price	Daily (closing) price changes of 1-month and 1-year Dutch TTF futures around ECB monetary policy announcements		Bloomberg

Table A.2: Detailed description of data used in the VAR analysis

Variable	Description	Notes	Source
1-year yield	Germany Government 1 year yield	Monthly average of daily values	Macrobond Financial AB
5-year yield	Germany Government 5 year yield	Monthly average of daily values	Macrobond Financial AB
US/EUR	US-Dollar per Euro, spot rate	Monthly average of daily values	Macrobond Financial AB
Industrial Production	Euro Area Industrial Production excl. Construction		Eurostat
Brent oil price	Brent crude Europe Spot price FOB, US-Dollar per barrel	Monthly average of daily values	Energy Information Administration
CPI (headline)	Euro Area Harmonized Index of Consumer Prices	Seasonally adjusted using X13	Eurostat
HICP housing	Euro Area, HICP, Housing, Water & Electricity & Gas & Other Fuels	Seasonally adjusted using X13	Eurostat
HICP transport	Euro Area, HICP, Transport	Seasonally adjusted using X13	Eurostat
HICP heating	Euro Area, HICP, Housing, Water, Electricity, Fuel, Electricity, Gas	Seasonally adjusted using X13	Eurostat
HICP fuels	Euro Area, HICP, Fuels & Lubricants for Personal Transport Equipment	Seasonally adjusted using X13	Eurostat
HICP energy	Euro Area, HICP, Energy	Seasonally adjusted using X13	Eurostat
Credit spread	ICE BofA Euro High Yield Index Option-Adjusted Spread	Monthly average of daily values	FRED
Euro Area short-run monetary policy proxy	3 month (monetary event window) OIS surprise	Calculated based on data and methodology by Jarociński and Karadi (2020) ("poor-man" approach), aggregated to monthly frequency according to Kilian (2024)	Jarociński and Karadi (2020) and authors' calculations
Euro Area medium-run monetary policy proxy	2 year (monetary event window) OIS surprise	The same notes apply to all monetary policy proxies	Jarociński and Karadi (2020) and authors' calculations
Euro Area generic monetary policy proxy	First principle component of 1 month to 1 year (monetary event window) OIS surprises	The same notes apply to all monetary policy proxies	Jarociński and Karadi (2020) and authors' calculations
Global oil production	Global oil production (million barrels/day)		Baumeister and Hamilton (2019)
Global IP	Global industrial production		Baumeister and Hamilton (2019)
Consensus 1-year ahead inflation expectations	(GDP-) Weighted average of Germany, France, Italy, and Spain	We use the largest four euro area countries' data since the euro area aggregate data is only available starting from December 2002. The (monthly) 1-year ahead expectation is a weighted average of the "Current year" and "Next year" inflation forecasts, as in Miranda-Agrippino and Ricco (2021).	Consensus economics
Oil supply news proxy	Suprise in oil futures prices around OPEC announcements	Monthly sum of daily values	Känzig (2021)

As in Born and Pfeifer (2021), we demean the variables to avoid numerical problems arising from under/overflow during the posterior computations that involve the sum of squares.

B High-frequency event study robustness results

Table B.1: Additional results for the event study regression for the euro area, US and UK (Equation 1) for a 100 basis points monetary policy surprise

	EA	EA	US	US	UK	UK
	(1)	(2)	mps^{pm}_{FF4}	$mps^{\perp}$	(1)	(2)
$\widehat{\beta}^{100\text{bps}}$	-3.20^{**}	-3.34^{**}	-2.24^{**}	-2.23^{***}	0.37	0.36
	(1.31)	(1.54)	(1.04)	(0.83)	(0.67)	(0.68)
R^2 (%)	3.75	3.37	2.64	3.21	0.33	0.38
N	211	182	198	197	257	246
Sample	2002:1	2002:1	1996:1	1996:1	1997:6	1997:6
	2021:12	2019:12	2019:12	2019:12	2021:3	2019:12

Note: Coefficient estimates $\widehat{\beta}^{100\text{bps}}$ from the Brent crude oil price event study regression equation $p_t = \alpha + \beta mps_t + \epsilon_t$, where t indexes monetary policy announcements. Coefficient represents the percentage change in the Brent crude oil price in response to a 100 basis points increase in the country-specific monetary policy surprise measure. Each column represents the results for a different country-sample combination. Heteroskedasticity-consistent standard errors are reported in parentheses. *, **, *** represent statistical significance levels at 10%, 5% and 1% respectively.

Table B.2: Additional results for the event study regression for the euro area, US and UK (Equation 1) for a 1 standard deviation monetary policy surprise

	EA	EA	US	US	UK	UK
	(1)	(2)	mps^{pm}_{FF4}	$mps^{\perp}$	(1)	(2)
$\widehat{\beta}^{\text{std}}$	-0.054^{**}	-0.056^{**}	-0.078^{**}	-0.078^{***}	0.020	0.019
	(0.022)	(0.026)	(0.037)	(0.029)	(0.035)	(0.037)
R^2 (%)	3.75	3.37	2.64	3.21	0.33	0.38
N	211	182	198	197	257	246
Sample	2002:1	2002:1	1996:1	1996:1	1997:6	1997:6
	2021:12	2019:12	2019:12	2019:12	2021:3	2019:12

Note: Coefficient estimates $\widehat{\beta}^{\text{std}}$ from the Brent crude oil price event study regression equation $p_t = \alpha + \beta mps_t + \epsilon_t$, where t indexes monetary policy announcements. Coefficient represents the percentage change in the Brent crude oil price in response to a one standard deviation increase in the country-specific monetary policy surprise measure. Each column represents the results for a different country-sample combination. Heteroskedasticity-consistent standard errors are reported in parentheses. *, **, *** represent statistical significance levels at 10%, 5% and 1% respectively.

Table B.3: Coefficient estimates $\widehat{\beta}^{100\text{bps}}$ from the natural gas price (Dutch TTF) event study regressions.

	1-month TTF	1-year TTF	1-month TTF	1-year TTF
$\widehat{\beta}^{100\text{bps}}$	-17.42^{***}	-12.32^{***}	-13.85^{***}	-13.41^{***}
	(4.50)	(3.12)	(3.92)	(3.23)
R^2 (%)	2.68	2.61	1.39	2.69
Sample	2007:10-2019:12	2007:10-2019:12	2007:10-2021:12	2007:10-2021:12
N	127	127	143	143

Note: Event study regressions are of the form $p_t = \alpha + \beta mps_t + \epsilon_t$ for the ECB, where t indexes ECB policy announcements, p_t is the daily change of the relevant futures price, computed as the difference between the closing price of the ECB policy announcement day and the closing price of the previous day. Coefficient represents the percentage change in the Dutch TTF natural gas price in response to a 100 basis points increase in the country-specific monetary policy surprise measure. Each column presents the event study regression for the combination of a different TTF maturity and a different sample period. mps_t is the high frequency change in the three month Overnight Index Swap (OIS) rate with poor man's sign restrictions as in Jarociński and Karadi (2020). Daily Dutch TTF price data is available from October 2007. Heteroskedasticity-consistent standard errors are reported in parentheses.

Table B.4: Coefficient estimates $\widehat{\beta}^{\text{std}}$ from the natural gas price (Dutch TTF) event study regressions.

	1-month TTF	1-year TTF	1-month TTF	1-year TTF
$\widehat{\beta}^{\text{std}}$	-0.33^{***}	-0.24^{***}	-0.25^{***}	-0.24^{***}
	(0.09)	(0.06)	(0.07)	(0.06)
R^2 (%)	2.68	2.61	1.39	2.69
Sample	2007:10-2019:12	2007:10-2019:12	2007:10-2021:12	2007:10-2021:12
N	127	127	143	143

Note: Event study regressions are of the form $p_t = \alpha + \beta mps_t + \epsilon_t$ for the ECB, where t indexes ECB policy announcements, p_t is the daily change of the relevant futures price, computed as the difference between the closing price of the ECB policy announcement day and the closing price of the previous day. Coefficient represents the percentage change in the Dutch TTF natural gas price in response to a one standard deviation increase increase in the country-specific monetary policy surprise measure. Each column presents the event study regression for the combination of a different TTF maturity and a different sample period. mps_t is the high frequency change in the three month Overnight Index Swap (OIS) rate with poor man's sign restrictions as in Jarociński and Karadi (2020). Daily Dutch TTF price data is available from October 2007. Heteroskedasticity-consistent standard errors are reported in parentheses.

C Revisiting Gagliardone and Gertler (2023)

In a VAR using high-frequency identification of monetary policy shocks, Gagliardone and Gertler (2023) do not find that the real oil price declines in response to a contractionary US monetary policy shock. This contradicts our findings and much of the related literature showing that contractionary US monetary policy shocks decrease commodity and/or oil prices (Anzuini et al. (2012); Miranda-Agrippino and Ricco (2021); Bauer and Swanson (2023); Degasperi et al. (2023); Miranda-Pinto et al. (2023)).[26]

In order to understand the source of their different result regarding the crude oil price impulse response, we replicate the VAR in Gagliardone and Gertler (2023). There are two limitations of this replication exercise. First, the authors are not explicit about whether they run separate VARs for the monetary policy shock and the oil supply shock. Therefore, since there is no explanation of a joint identification procedure, we assume they estimate two separate VAR models. Second, the authors also measure "surprises around non-FOMC dates on which the Federal Reserve revealed information", but they neither specify these dates nor grant access to the corresponding data. As a result, we cannot precisely replicate their approach to identifying shocks across the full event set.

Through our replication analysis, we find that the different result reported by Gagliardone and Gertler (2023) may be driven by certain empirical choices that can affect the validity of the estimates. First, the authors aggregate monetary policy surprises from a higher frequency to a monthly frequency by summing the surprises in a given month, yet, in the VAR, they use average-of-period monthly data for variables available at a daily frequency. Kilian (2024) shows that such practice can bias the impulse response estimates. Instead, he proposes using end-of-period data, if one wishes to simply sum over the high-frequency surprises. In our replication of their monetary policy VAR, we find that using end-of-month oil prices (see Figure C.1), or the average price over the last three or five trading days of each month (see Figure C.2), leads to a significant decline in real crude oil prices following a contractionary monetary policy shock. Second, if instead one prefers average-of-period monthly prices, it is crucial to construct the monthly surprises in line with the methodology in Kilian (2024); adopting that approach in the VAR replication of Gagliardone and Gertler (2023) again produces a decline in the real oil price in line with the results from our VAR framework (see Figure C.3).[27]

Moreover, the oil price series used by the authors (FRED code: WTISPLC) is not contin-

[26] Miranda-Agrippino and Ricco (2021) include the Commodity Research Bureau (CRB) commodity price index in their baseline VAR but do not report the IRFs. Therefore, using their replication files while keeping true to their baseline empirical specification, we produce the commodity price index IRFs and find that the commodity price index declines significantly in response to a contractionary US monetary policy shock (see Figure E.8).

[27] The BPSVAR framework we employ allows the proxy variable to be serially correlated and predictable. Therefore, any serial correlation or predictability arising from the aggregation scheme does not pose a problem in our empirical framework.

uous during the first six years of their sample, remaining constant over prolonged intervals.

Figure C.1: Replication of Gagliardone and Gertler (2023) with end-of-month real oil price

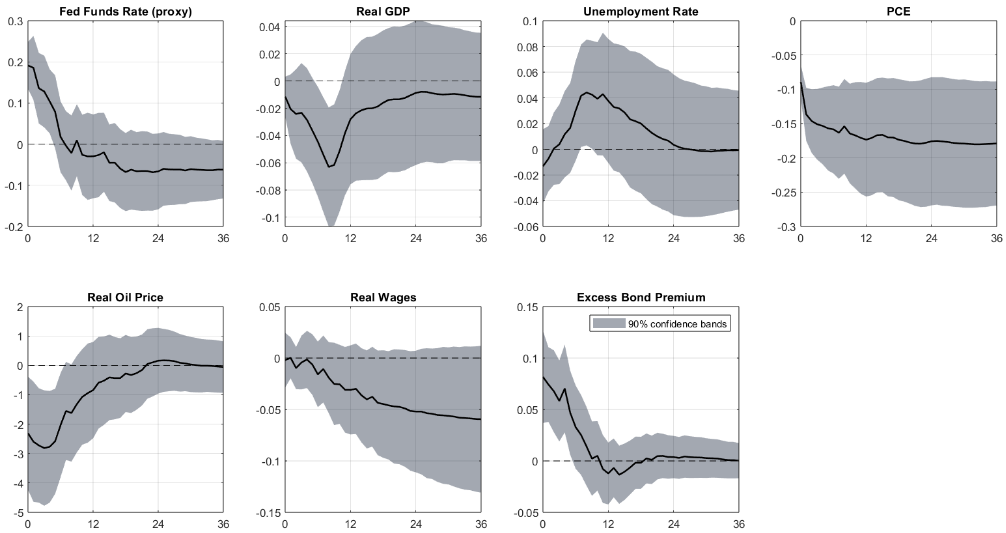

Notes: Baseline seven-variable VAR from Gagliardone and Gertler (2023). Sample is 1973M1–2019M12. The solid line is the point estimate and the shaded areas are the 90 percent confidence bands, computed using the wild bootstrap.

Figure C.2: Replication of Gagliardone and Gertler (2023) with real oil price as the average of the last 5 days in a month

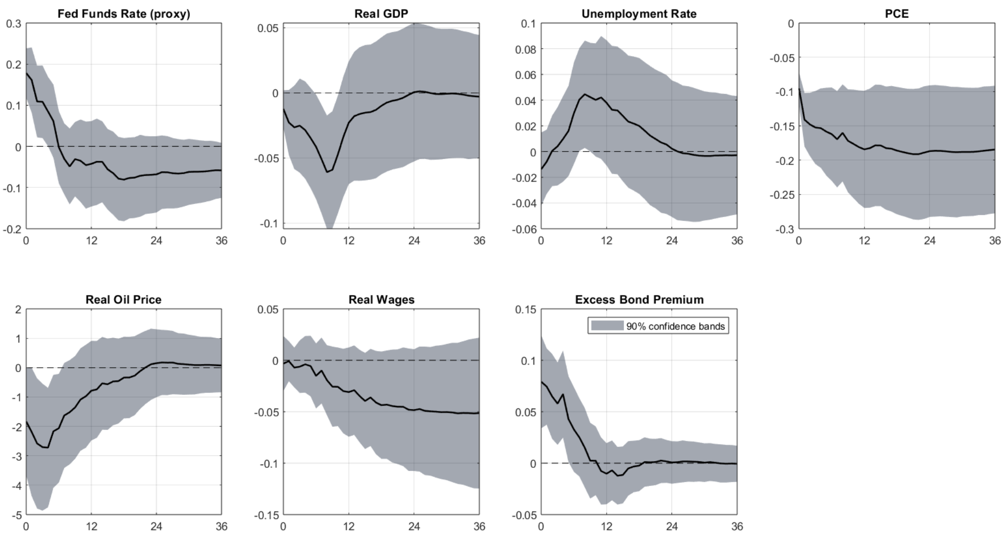

Notes: Baseline seven-variable VAR from Gagliardone and Gertler (2023). Sample is 1973M1–2019M12. The solid line is the point estimate and the shaded areas are the 90 percent confidence bands, computed using the wild bootstrap.

Figure C.3: Replication of Gagliardone and Gertler (2023) average-of-month real oil price and monetary policy surprises aggregated following Kilian (2024)

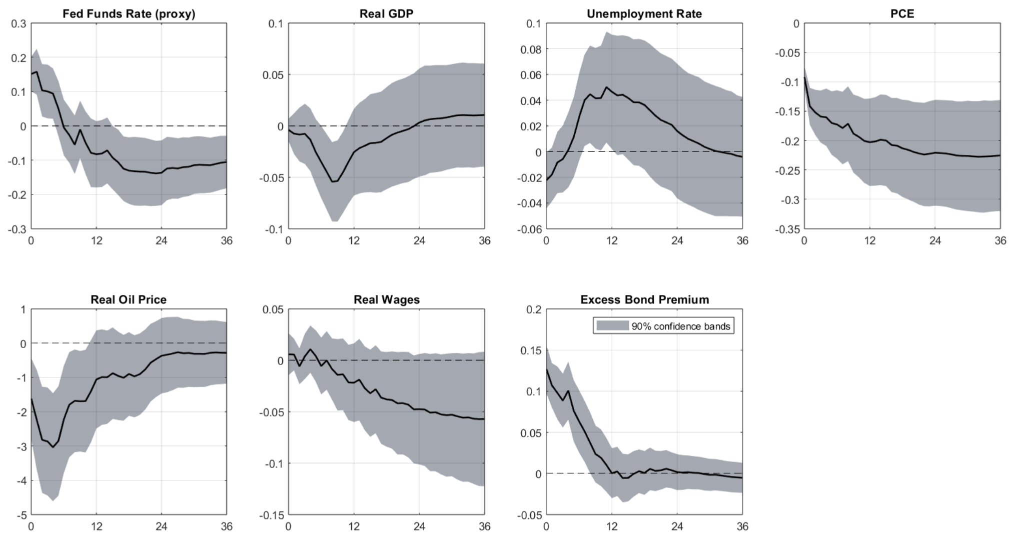

Notes: Baseline seven-variable VAR from Gagliardone and Gertler (2023). Sample is 1973M1–2019M12. The solid line is the point estimate and the shaded areas are the 90 percent confidence bands, computed using the wild bootstrap.

D Details on the Bayesian Proxy SVAR model

In this appendix we give more details on the implementation of the algorithm of Arias et al. (2021) and derive equations (5) and (6). For convenience, we reproduce the VAR Equation (4), augmented with proxies:

$$\tilde{\boldsymbol{y}}_t' \tilde{\boldsymbol{A}}_0 = \tilde{\boldsymbol{y}}_{t-1}' \tilde{\boldsymbol{A}}_1 + \tilde{\boldsymbol{\epsilon}}_t'. \tag{D.1}$$

To ensure that the augmentation with equations for the proxy variables does not affect the dynamics of the endogenous variables, the coefficient matrices $\tilde{\boldsymbol{A}}_\ell$ are specified as

$$\tilde{\boldsymbol{A}}_\ell = \begin{pmatrix} \underset{(n\times n)}{\boldsymbol{A}_\ell} & \underset{(n\times k)}{\boldsymbol{\Gamma}_{\ell,1}} \\ \underset{(k\times n)}{\boldsymbol{0}} & \underset{(k\times k)}{\boldsymbol{\Gamma}_{\ell,2}} \end{pmatrix}, \qquad \ell = 0, 1. \tag{D.2}$$

The zero restrictions on the lower left-hand side block imply that the proxy variables do not enter the equations of the endogenous variables. The reduced form of the model is

$$\tilde{\boldsymbol{y}}_t' = \tilde{\boldsymbol{y}}_{t-1}' \tilde{\boldsymbol{A}}_1 \tilde{\boldsymbol{A}}_0^{-1} + \tilde{\boldsymbol{\epsilon}}_t{}' \tilde{\boldsymbol{A}}_0^{-1}. \tag{D.3}$$

Because the inverse of $\tilde{\boldsymbol{A}}_0$ in Equation (D.2) is given by

$$\tilde{\boldsymbol{A}}_0^{-1} = \begin{pmatrix} \boldsymbol{A}_0^{-1} & -\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1} \\ 0 & \boldsymbol{\Gamma}_{0,2}^{-1} \end{pmatrix}, \tag{D.4}$$

the last k equations of the reduced form of the VAR model in Equation (D.3) read as

$$\boldsymbol{m}_t' = \tilde{\boldsymbol{y}}_{t-1}' \tilde{\boldsymbol{A}}_1 \begin{pmatrix} -\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1} \\ \boldsymbol{\Gamma}_{0,2}^{-1} \end{pmatrix} - \boldsymbol{\epsilon}_t' \boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1} + \boldsymbol{\eta}_t' \boldsymbol{\Gamma}_{0,2}^{-1}, \tag{D.5}$$

which is the generalization for k proxies of equations (5) and (6), with

$$\boldsymbol{B}_1 = \begin{pmatrix} -\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1} \\ \boldsymbol{\Gamma}_{0,2}^{-1} \end{pmatrix}, \quad \boldsymbol{B}_2 = \boldsymbol{\Gamma}_{0,2}^{-1}.$$

To see that only $\boldsymbol{\epsilon}_t^{*\prime}\boldsymbol{V}$ enters the equation, we order the structural shocks so that $\boldsymbol{\epsilon}_t = (\boldsymbol{\epsilon}_t^{o\prime}, \boldsymbol{\epsilon}_t^{*\prime})'$, which yields

$$E\left[\boldsymbol{\epsilon}_t \boldsymbol{m}_t'\right] = -\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1} = \begin{pmatrix} \underset{((n-k)\times k)}{\boldsymbol{0}} \\ \underset{(k\times k)}{\boldsymbol{V}} \end{pmatrix}. \tag{D.6}$$

The first equality is obtained using Equation (D.5) and because the structural shocks $\boldsymbol{\epsilon}_t$ are by assumption orthogonal to $\boldsymbol{y}_{t-1}$ and $\boldsymbol{\eta}_t$. The second equality is due to the exogeneity and relevance conditions in Equations (3a) and (3b). Equation (D.6) shows that the identifying assumptions imply restrictions on the last k columns of the contemporaneous structural impact coefficients in $\tilde{\boldsymbol{A}}_0^{-1}$. In particular, if the exogeneity condition in Equation (3b) holds, the first $n-k$ rows of the upper right-hand side sub-matrix $\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1}$ of $\tilde{\boldsymbol{A}}_0^{-1}$ in Equation (D.4) are zero. From the reduced form in Equation (D.3) it can be seen that this implies that the first $n-k$ structural shocks do not impact contemporaneously the proxy variables. In turn, if the relevance condition in Equation (3a) holds, the last k rows of the upper right-hand side sub-matrix $\boldsymbol{A}_0^{-1}\boldsymbol{\Gamma}_{0,1}\boldsymbol{\Gamma}_{0,2}^{-1}$ of $\tilde{\boldsymbol{A}}_0^{-1}$ are different from zero. From the reduced form in Equation (D.3) it can be seen that this implies that the last k structural shocks impact the proxy variables contemporaneously. The Bayesian estimation algorithm of Arias et al. (2021) determines the estimates of $\boldsymbol{A}_0$ and $\boldsymbol{\Gamma}_{0,\ell}$ such that the restrictions on $\tilde{\boldsymbol{A}}_0^{-1}$ implied by Equations (3a) and (3b) as well as on $\tilde{\boldsymbol{A}}_\ell$ in Equation (D.2) are simultaneously satisfied, and hence the estimation identifies the structural shocks $\boldsymbol{\epsilon}_t^*$.

E BPSVAR robustness results

Figure E.1: Baseline Euro Area SVAR model: Response of the credit spread

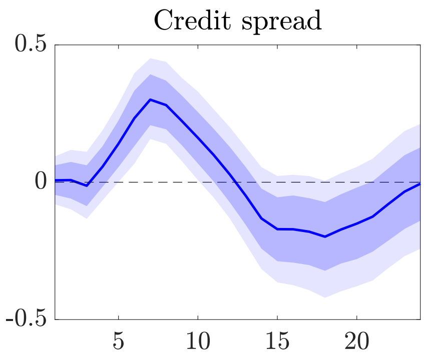

Notes: Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Response of the BBB corporate bond spread in percentage points.

Figure E.2: Euro Area SVAR model, including global oil market variables

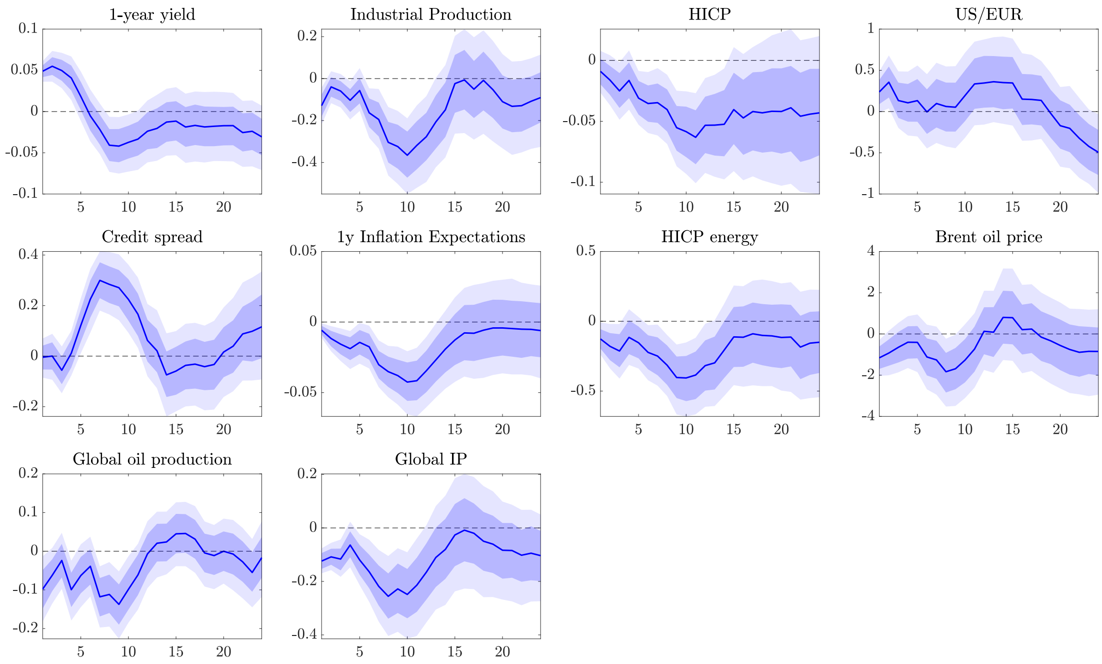

Notes: Euro Area model including additional variables that are typically used in models of the oil market. Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Horizon in months.

Figure E.3: Euro Area SVAR model, external vs. internal instrument identification

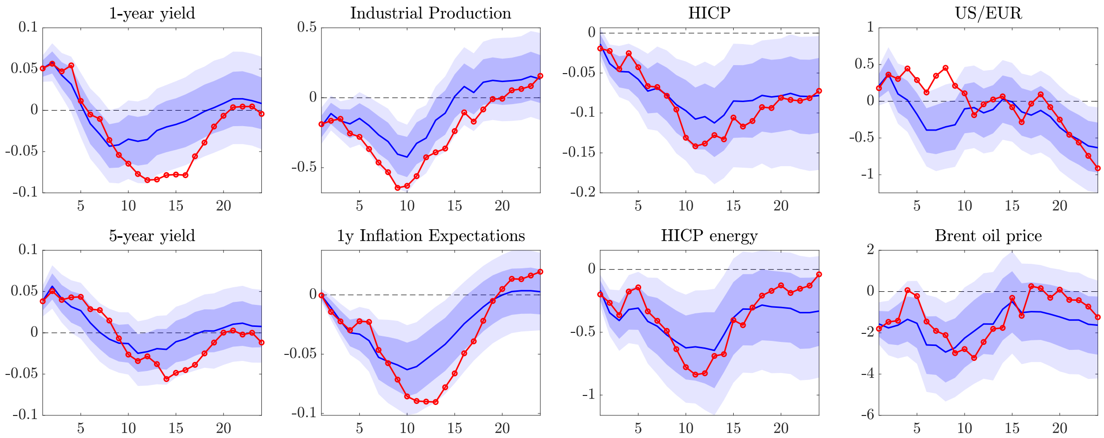

Notes: Posterior means of Impulse response functions to a one standard deviation monetary policy shock using the external-instrument BPSVAR identification alongside in blue 68% and 90% point-wise probability bands. Posterior mean impulse response functions to monetary policy shock identified using the internal-instrument approach of Plagborg-Møller and Wolf (2021) are depicted in red. To make the estimation of the IRFs using the internal instrument approach comparable to the BPSVAR approach we use a version of the conjugate normal-inverse-wishart prior. The impulse responses for the internal instrument approach are scaled such that they induce the same impact effect for the 1-year yield.

Figure E.4: IRFs of natural gas prices (in dollars) and EA consumer gas prices (in euros)

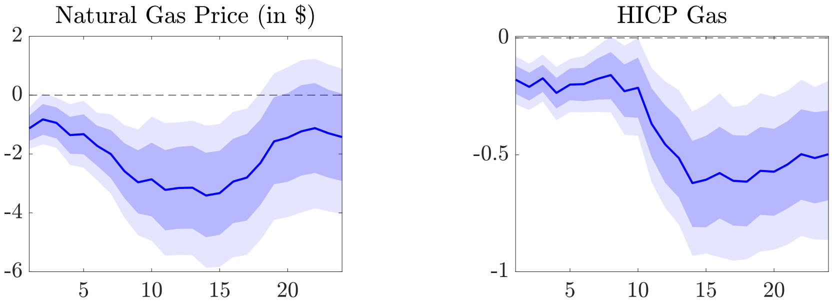

Impulse responses from the baseline BPSVAR model with (i) Worlds Bank's average, European natural gas primary commodity price (traded and quoted in US-$) and (ii) the natural gas component of the HICP (quoted in euros) as additional endogenous variables. Note that, even leaving aside differences in the currency in which the prices are quoted, these are not the same object. In particular, the HICP Gas series refers to gas prices faced by consumers (which naturally are more sticky) while the World Bank's average European natural gas price commodity price is the price traded on financial markets and faced by firms and consumers.

Figure E.5: Euro Area SVAR model, zero proxy relevance prior threshold

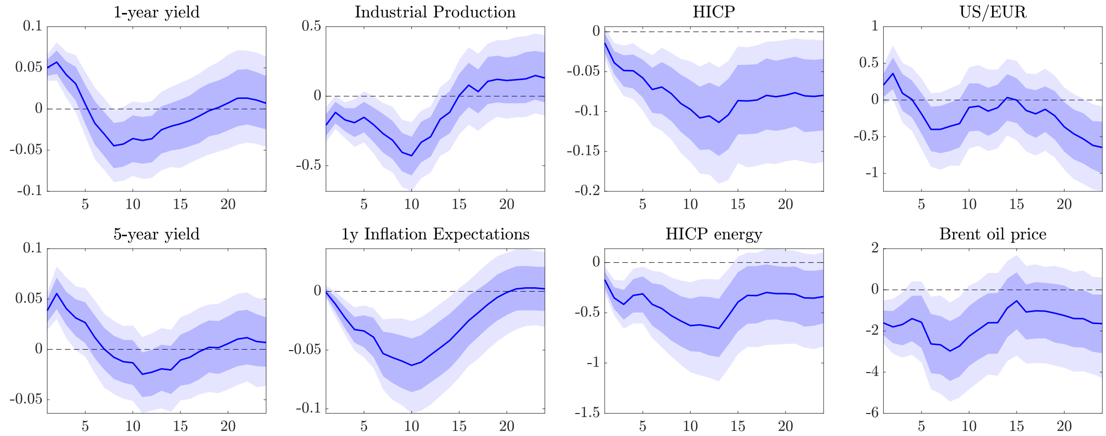

Notes: Impulse responses to a euro area monetary policy shock when the prior on the relevance of the shock for the proxy set to 0%. Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Horizon in months.

Figure E.6: Euro Area SVAR model including the Pandemic

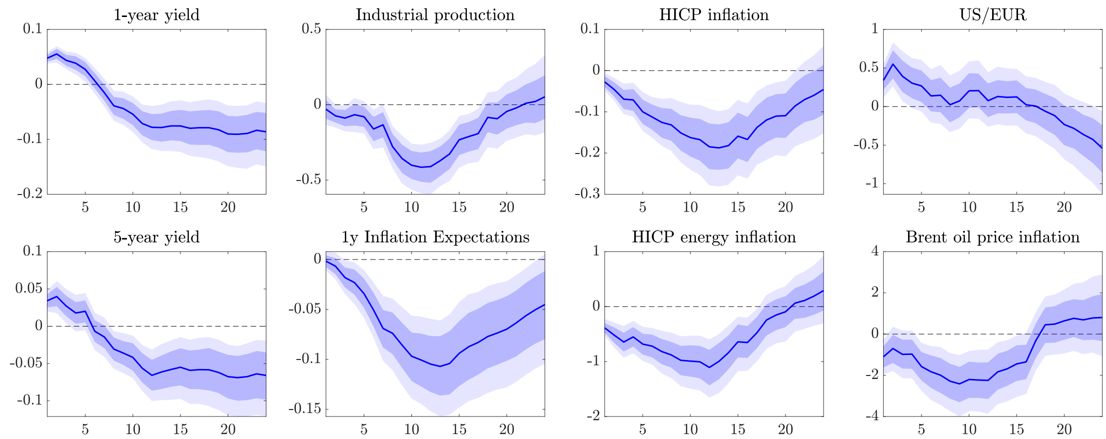

Notes: Impulse response functions to a one standard deviation monetary policy shock from the estimated BPVSAR model when including the pandemic period into the estimation sample. We model the pandemic using the "Pandemic-Priors" approach of Cascaldi-Garcia (2022) and transform prices to inflation rates to preserve stationarity. Due to data availability, we follow Känzig, 2023 and set to zero the values of the proxy starting in 2020. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months.

Figure E.7: Euro Area SVAR model including different subcomponents of HICP energy

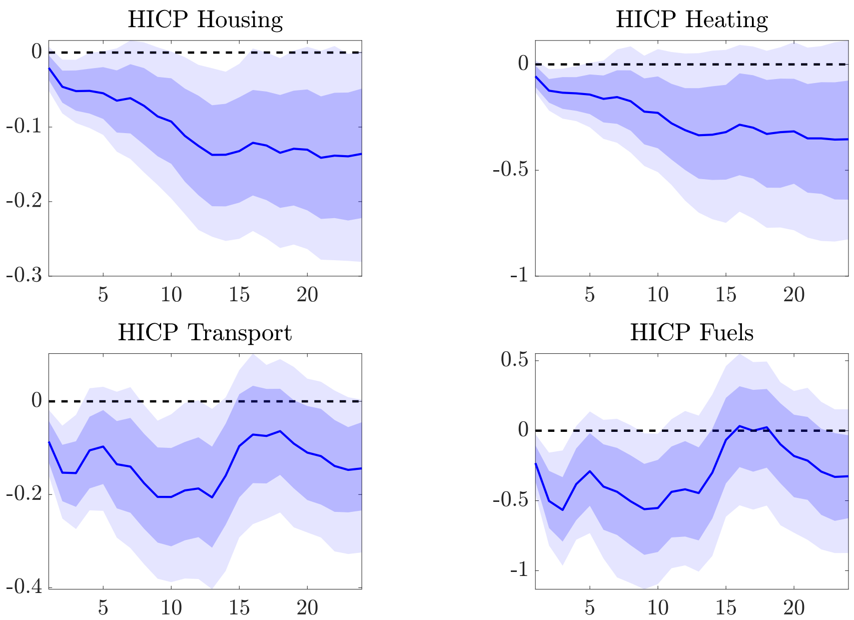

IRFs from the baseline BPSVAR model with the individual subcomponents of the HICP as additional endogenous variables. The official Eurostat categories are called "Housing, Water, Electricity, Gas & Other Fuels", "Transport", "Housing, Water, Electricity, Fuel, Electricity, Gas & Other Fuels", and "Fuels & Lubricants for Personal Transport Equipment". Their weights in the headline HICP are, in percent, 16.5, 15.4, 5.9, and 4.3, respectively (2019 values).

Table E.1: Literature review: Trough industrial production response to 100 bps tightening

Model	Sample period	Trough response
A. US Models		
Bauer and Swanson (2023) (Original sample)	1973:1 – 2020:2	−5.1%
Bauer and Swanson (2023) (Starting 2002)[a]	2002:1 – 2020:2	−10.5%
Miranda-Agrippino and Ricco (2021) (Original sample)	1979:1 – 2014:12	−1.7%
Miranda-Agrippino and Ricco (2021) (Starting 1999)	1999:1 – 2018:12	−9.2%
Jarociński and Karadi (2020) (Original Sample)	1984:2 – 2016:12	−3.6%
Jarociński and Karadi (2020) (Starting 1999)	1999:1 – 2016:12	−8.7%
IKKS US SVAR (starting 1990)	1990:1 – 2019:12	−1.3%
IKKS US SVAR (starting 1999)	1999:1 – 2019:12	−6.9%
B. EA Models		
Jarociński and Karadi (2020)	1999:1 – 2016:12	−17.4%
Corsetti et al. (2024)	1999:1 – 2021:12	−32%
Badinger and Schiman (2023)	1999:1 – 2019:12	−6.6%
This paper (main specification)	2002:1 – 2019:12	−9.3%
This paper (from 1999)	1999:1 – 2019:12	−6.1%
C. UK Models		
Braun et al. (2024)	1997:1 – 2019:12	≈ −6.6%
D. Summary		
Model average	Start < 1999	-3.7%
Model average	Start ≥ 1999	-11.9%

Notes: We report the trough Industrial Production response to a peak of one percentage point increase in the 1-year government bond yield of the respective country. As Braun et al. (2024) use monthly GDP instead of industrial production, we rescale their estimate to industrial production by taking into account the relative volatilities of the series (see Georgiadis et al., 2024 for a similar approach). For the euro area, the cited studies use the 1-year Bund yield. IKKS US SVAR refers to the baseline SVAR model presented in Section 3 of this paper, where we replace the euro area data with the corresponding time series for the US. See Figures E.9 and E.10.

[a] The yield fails to increase in response to a contractionary monetary policy shock when starting the sample from 1999. Therefore, we chose to start the sample in 2002.

Figure E.8: Baseline VAR from Miranda-Agrippino and Ricco (2021)

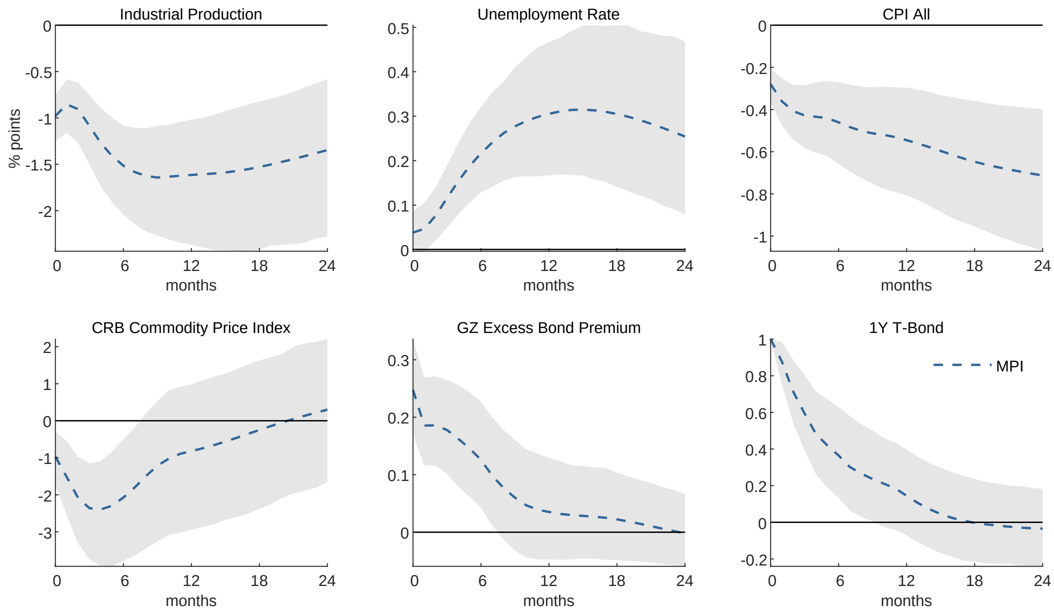

Notes: Baseline six-variable VAR from Miranda-Agrippino and Ricco (2021). MPI stands for the informationally robust monetary policy surprise series the authors construct. The shock is normalized to induce a 100 basis point increase in the 1-year rate. Sample is 1979M1–2014M12. Shaded areas are 90 percent posterior coverage bands.

Figure E.9: US SVAR model, starting 1990

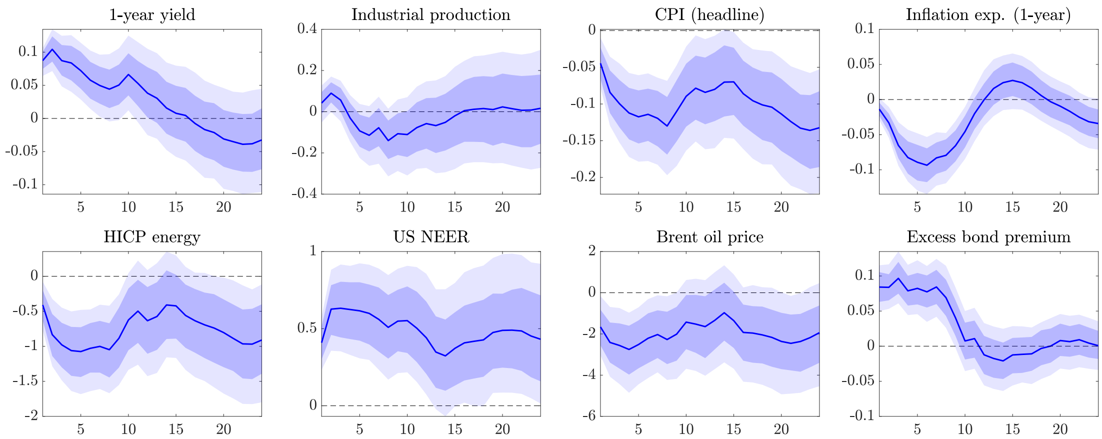

Notes: Impulse responses to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months. Impulse responses for variables that do not correspond to interest rates or inflation rates are expressed in percent. Impulse responses for inflation rates and interest rates are expressed in annualized percentage points.

Figure E.10: US SVAR model, starting 1999

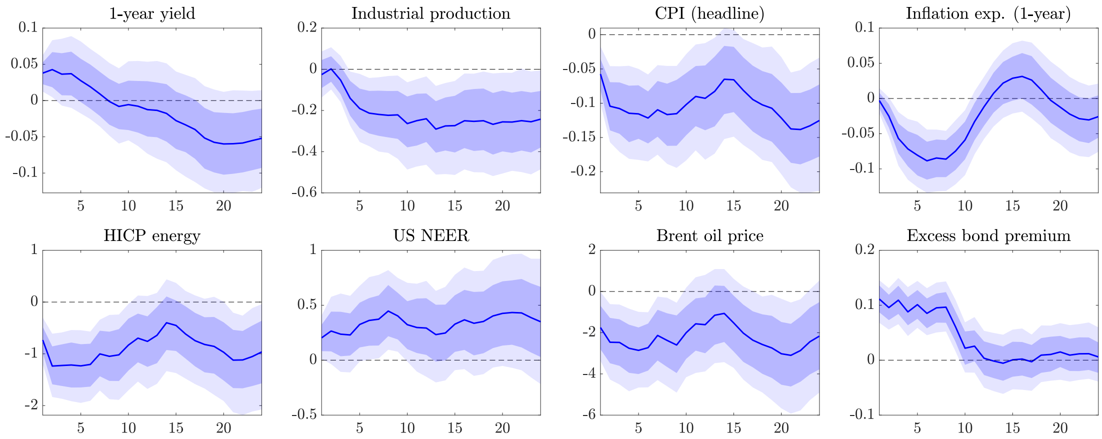

Notes: Impulse responses to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months. Impulse responses for variables that do not correspond to interest rates or inflation rates are expressed in percent. Impulse responses for inflation rates and interest rates are expressed in annualized percentage points.

F Discussion of the magnitude of the oil price response

F.1 A simple theoretical model

To rationalize why euro area monetary policy affects global energy prices and why energy prices are particularly sensitive to monetary policy induced changes in demand, we build a partial equilibrium model of the global market for energy goods. In particular, we assume that energy goods are traded globally, energy prices are flexible and energy supply is fixed at $\bar{E}$ in the short run as in Bayer et al. (2023).[28] Furthermore, we assume that world demand (Y_t^W) is allocated according to a standard CES aggregate of energy goods and non-energy goods, with α denoting the weight of energy goods in the aggregate basket.[29] Under these assumptions we can write the demand ($Y_{E,t}^D$) for and supply of energy goods ($Y_{E,t}^S$) as the following system of equations

$$Y_{E,t}^D = \alpha\Big(\frac{P_t^E}{P_t^W}\Big)^{-\sigma} Y_t^W, \qquad\qquad Y_{E,t}^S = \bar{E}, \tag{F.1}$$

with σ as the elasticity of substitution between energy and non-energy goods and P_t^E/P_t^W as the relative price of energy goods with respect to the world aggregate. Imposing market clearing ($Y_{E,t}^S = Y_{E,t}^D$) and log-linearizing, we can write the equilibrium relation as

$$\widehat{p}_{t,r}^E = \frac{\widehat{y}_t^W}{\sigma} \tag{F.2}$$

with $\widehat{y}_t^w$ denoting deviations of global demand from its steady state and ($\widehat{p}_{t,r}^E$) as the corresponding deviation of the relative price of energy goods.

Given that the euro area constitutes approximately 12% of global GDP and that the ECB's monetary policy decisions have sizable spillovers to other countries (Miranda-Agrippino and Nenova (2022), Ter Ellen et al. (2020)), an increase in the ECB's policy rate affects global demand $\widehat{y}_t^W$, which in turn affects the relative price of energy. In particular, recent estimates on the elasticity of substitution between energy and non-energy goods imply $\sigma \in [0.1, 0.2]$ (Bachmann et al. (2022) and Bayer et al. (2023)), i.e. an ECB-induced 1% decrease in demand causes the relative price of energy goods to fall by 5% to 10%. The intuition is that, given a low elasticity of substitution, the demand curve is very steep. When confronted with a vertical short-run supply curve, a change in demand will result in strong price adjustments in order to reach an equilibrium between supply and demand. This stylized mechanism not

[28] At the intraday frequency the assumption of flexible energy prices follows immediately from our high-frequency event study and at the monthly frequency this assumption is also in line with the micro-data underlying the computation of the HICP as for instance Aucremanne and Dhyne (2004) show that the prices of energy goods are on average updated every month.

[29] The aggregator is given by $Y_t^W = [\alpha^{\frac{1}{\sigma}} Y_{E,t}^{\frac{\sigma-1}{\sigma}} + (1-\alpha)^{\frac{1}{\sigma}} Y_{NE,t}^{\frac{\sigma-1}{\sigma}}]$ with $Y_{NE,t}$ as non energy goods.

only rationalizes why ECB policy decisions affect energy prices but the assumptions on the energy market structure also offer one possible interpretation for why energy prices are more responsive to changes in monetary policy than prices of other goods. In the next section we discuss how this this partial equilibrium model of the energy market is embedded in the fully-fledged, state-of-the-art general equilibrium model of Bayer et al. (2023) and show that in this model, an increase in the ECB's policy rate of similar magnitude than the one estimated in our BPSVAR, indeed implies a similar drop in the price of energy goods.

F.2 The mechanism in general equilibrium

Figure 2 reveals that a standard-deviation monetary policy shock, which increases the short-term interest rate by roughly 5 basis points, leads to an immediate fall in the oil price by approximately 2%. To shed some light on the plausibility on the magnitudes we use a representative agent version of the HANK model of Bayer et al. (2023) to gauge the plausibility of out results through the lens of a state-of-the-art model that features an energy market along the lines sketched above.[30]

The model of Bayer et al. (2023) is an arguably standard model of a monetary union with two countries and nominal frictions in terms of price and wage setting. The crucial ingredient is that firms (households) in these countries use (consume) energy goods. Energy goods are assumed to be in fixed supply and, crucially, as the model does not feature a small open economy assumption, a change in the demand for energy from households and firms in the monetary union, will affect the price of energy. Therefore the model features a market for energy goods along the lines of the one that we sketch in Section F.1. To use the model for our purposes we add a monetary policy shock to the Taylor Rule of the monetary authority in the monetary union, while keeping the calibration and all other model features exactly as in Bayer et al. (2023).

Figure F.1 illustrates that in this state-of-the-art model, a monetary shock that leads to roughly the same interest rate response as in our empirical model, causes energy prices to fall by even more than what we find empirically. The intuition for the large volatility of energy prices is that energy goods are in fixed supply and the elasticity of energy- and non-energy goods is assumed to be non-zero but relatively low in line with Auclert et al. (2023) and Bachmann et al. (2022). Therefore, as the market for energy goods has to clear, the energy price has to move a lot to realign the demand with the supply of energy goods. Intuitively, all else equal, a monetary shock induces a fall in demand for all goods. Given the low elasticity of household demand to a change in the energy price, the price of these energy goods has to fall a lot to ensure that households ultimately buy the fixed amount of energy goods supplied.

[30]We thank Fabian Seyrich for sharing the code with us.

Figure F.1: IRF of the energy price to a monetary shock in the model of Bayer et al. (2023)

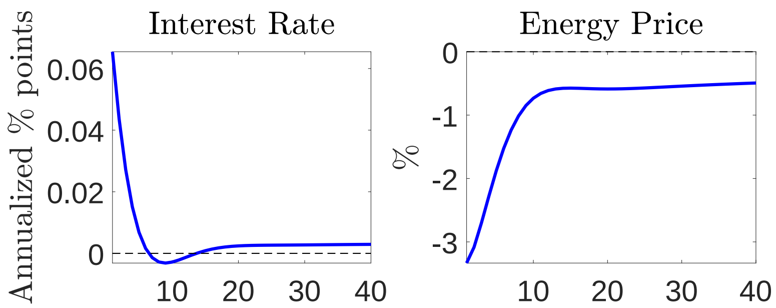

F.3 Relative volatilities of the interest rate and oil price

An arguably simpler argument for why the "elasticity" of oil prices to a monetary shock that we find empirically is not excessively large can be made by taking into account the relative volatilities of these variables. In line with the intuition sketched above, the oil price is very volatile compared to the short-term interest rate.

Figure F.2 illustrates this by plotting the impulse responses of the interest rate and the oil price in terms of their unconditional and conditional standard deviations (i.e. the standard deviation of their one-step ahead forecast error). It becomes apparent that, when measured in terms of the standard deviation of the respective forecast error, the average monetary policy shock causes the interest rate to increase by approximately half a standard deviation and the oil prices to fall by roughly a quarter of the standard deviation. Comparing this to the responses in levels of 5 basis points and 2% it becomes apparent that the 2% fall in the oil price is by no means excessively large. Thus, under the assumption that the forecast error is normally distributed around zero, an unexpected 5 basis point surprise in the interest rate is less likely than an unexpected 2% fall in the oil prices, as the latter only corresponds to a quarter of a standard deviation while the former corresponds to a surprise of half a standard deviation. The second row, which plots the IRFs in terms of their unconditional standard deviations, underscores the fact, that even unconditionally, the two magnitudes are more than comparable. In line with the intuition that we derived above, the oil price is just an arguably much more volatile object.

Figure F.2: IRF of the oil price to a monetary policy shock in standard deviation units

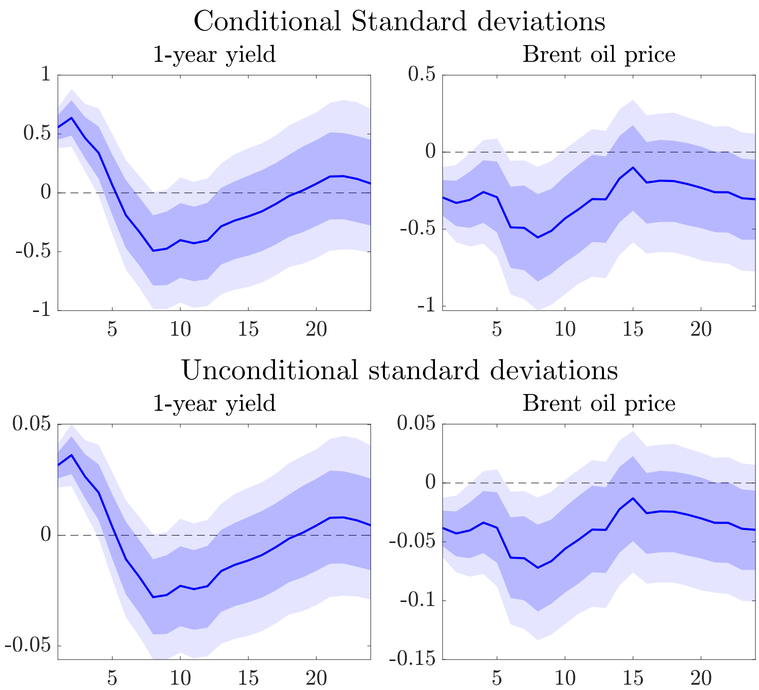

G Additional material on the OPEC counterfactual

G.1 Quantifying the role of global energy price for monetary policy trade-offs

Our framework not only allows us to quantify how the ability of the central bank to influence energy prices shapes the transmission of monetary policy to domestic inflation and inflation expectations but also enables us to speak to the role that this ability plays in the inflation-unemployment/inflation-output trade-off that central banks typically face. We follow Mankiw (2001) and Barnichon and Mesters (2021) and aim to measure the central banks's average trade-off, meaning we aim to measure the average fall in inflation caused by a change in policy that increases the unemployment rate by 1ppt or output by 1%.

As shown by Barnichon and Mesters (2021) the inflation-unemployment trade-off can be quantified using standard semi-structural methods such as our BPSVAR and can be estimated by a statistic that the authors coin "Phillips-Multiplier". In particular, the authors suggest to compute the sequence of "Phillips-Multipliers" ($\mathcal{P}_{\mathcal{A}}$) as

$$\mathcal{P}^h_{\mathcal{A}} = \frac{\Theta^h_{\bar{\pi},\nu^{mp},\mathcal{A}}}{\Theta^h_{\bar{U},\nu^{mp},\mathcal{A}}}. \tag{G.1}$$

Figure G.1: What if EA monetary policy shocks do not affect oil prices (including credible sets)

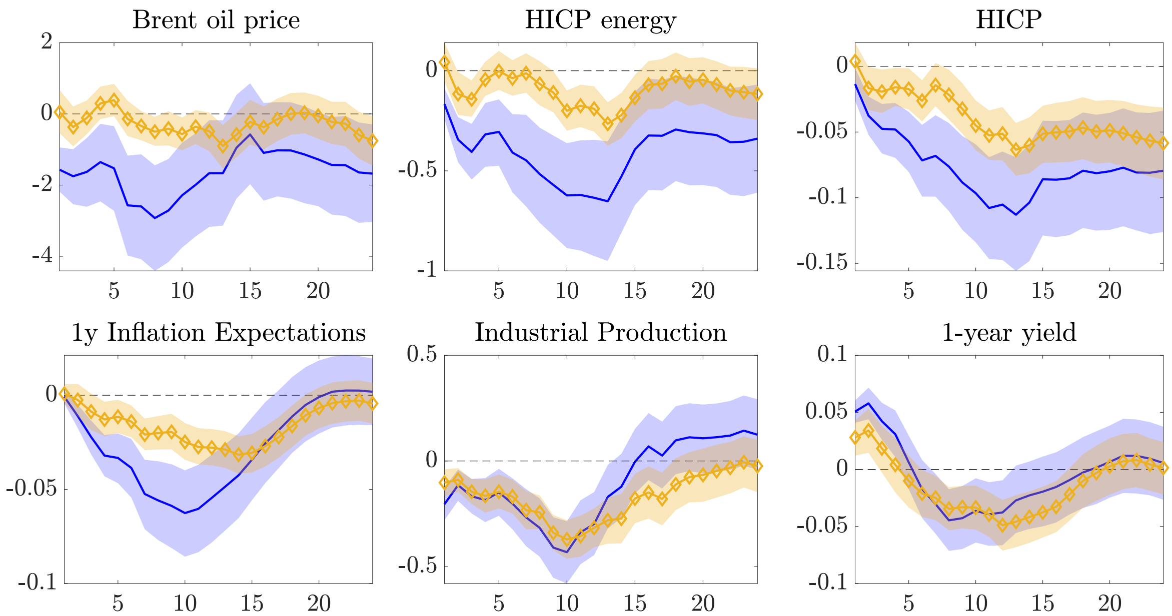

Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% point-wise credible sets in blue. Horizon in months. The golden line with diamonds shows the point-wise posterior means of the counterfactual where EA monetary policy does not affect the oil price. We approximate the solution to the counterfactual using the "best Lucas-Critique-robust approximation" of McKay and Wolf (2023), where we follow McKay and Wolf (2023) and condition on the point estimate to the monetary policy shock depicted in Figure 2. We also plot the 68% point-wise credible sets of this approximation in yellow.

where $\Theta^h_{\bar{x},\nu^{mp},\mathcal{A}}$ measures the horizon h impulse response of the average of variable x to a unit monetary policy shock ν^{mp} under the OPEC policy rule rule described by $\mathcal{A}$. For each variable x and horizon h this quantity can be computed by the average of the running cumulative sum of impulse responses $\Theta^h_{\bar{x},\nu^{mp},\mathcal{A}} = \frac{1}{h}\sum_{j=0}^{h}\Theta^j_{x,\nu^{mp},\mathcal{A}}$.

The "Phillips-Multiplier" at period h measures how the average rate of inflation changes if monetary policy would engineer a 1ppt increase in the unemployment rate over the next h periods. In other words, this statistic measures how the expectation (forecast) of inflation changes if monetary policy announces at period t that it will engineer an average increase of unemployment by 1ppt. over the next h horizon. Intuitively, in the textbook three equation New Keynesian model of Galí (2015), the "Phillips-Multiplier" is constant across periods and recovers the slope of the Phillips-curve with respect to unemployment (see Barnichon and Mesters (2021))

Given that this statistic can be computed solely based on impulse responses, we can also ask, how this measure of the inflation-unemployment trade-off would change if the ECB would not affect global energy prices as in Section 5. In particular, as in Section 5 we assume that OPEC follows a counterfactual policy rule $\widetilde{\mathcal{A}}$ aims to stabilize the oil price and calculate

$$\mathcal{P}^h_{\widetilde{\mathcal{A}}} = \frac{\Theta^h_{\bar{\pi},\nu^{mp},\widetilde{\mathcal{A}}}}{\Theta^h_{\bar{U},\nu^{mp},\widetilde{\mathcal{A}}}}. \tag{G.2}$$

To calculate this statistic we transform IRFs of industrial production to IRFs of unemployment by (i) transforming IRFs of industrial production to its GDP counterpart by taking into account the relative volatilities (3.3) and then compute the implied IRF of unemployment using Okun's law with a Coefficient of 2. Figure G.2 compares the baseline sequence of the Phillips-Multipliers as depicted by the blue lines and the counterfactual sequence of Phillips-Multipliers in depicted in gold. By comparing the blue and the golden line in Figure G.2 it becomes apparent that the ability to affect global energy prices plays a crucial role for the inflation output trade-off that the ECB faces. For instance, if the ECB were to engineer an increase in unemployment by 1ppt over the next year ($\mathcal{P}^{12}_{\mathcal{A}}$), this is estimated to yield a fall of average inflation of roughly 1.8% in the case where the ECB can affect energy prices. This value lies in the ballpark of the estimates of Barnichon and Mesters (2021). But this statistic changes dramatically in the counterfactual scenario, where this only brings about a fall in average inflation of about 0.8%. Thus, the ability of the ECB to affect global energy prices alleviates the unemployment-inflation trade-off by approximately 55% ((1.8-0.8)/1.8). When viewed through the lens of the textbook New-Keynesian model, this implies that the slope of the Phillips-curve is steeper, when monetary policy can affect fast-moving energy prices.

As quantifying the inflation-unemployment trade-off while sticking to our baseline model we have to take a stand on how a monetary policy-induced change in industrial production translates into a change in unemployment. Therefore we also compute the same statistic but replace the impulse response of average unemployment with the response of average industrial production, which is what report in the main text. To be more precise, we define the "Output-Phillips-Multiplier" as

$$\mathcal{P}^h_{\widetilde{\mathcal{A}}} = \frac{\Theta^h_{\bar{\pi},\nu^{mp},\widetilde{\mathcal{A}}}}{\Theta^h_{\bar{Y},\nu^{mp},\widetilde{\mathcal{A}}}}, \tag{G.3}$$

where $\Theta^h_{\bar{Y},\nu^{mp},\widetilde{\mathcal{A}}}$ measures the response of average industrial production to a monetary policy shock. This allows us to quantify the inflation-output trade-off, without having to make additional assumptions on the underlying mapping of output to unemployment. The results are shown in Figure 5 in the main text.

Figure G.2: Inflation-Unemployment trade-off under Baseline and Counterfactual OPEC rule

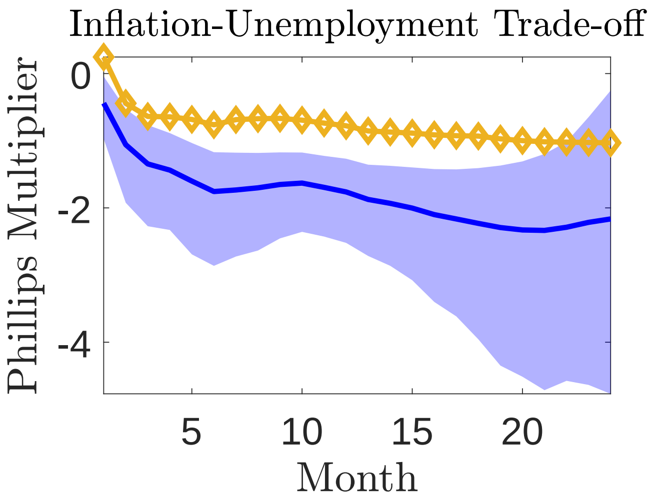

Notes: Point-wise median of the Phillips-Multiplier under the baseline policy rule alongside 50% credible sets in blue. Phillips-Multiplier estimated using the point-estimate of the Least-squares approximation of the counterfactual impulse responses depicted in in gold. To calculate this statistic we transform IRFs of industrial production to IRFs of unemployment by (i) transforming IRFs of industrial production to its GDP counterpart by taking into account the relative volatilities (3.3) and then compute the implied IRF of unemployment using Okun's law with a Coefficient of 2. We plot 68% credible sets to not distort the scale of the figure as the posterior is very much skewed to the left.

H Further material for the optimal policy counterfactuals

H.1 Deriving the optimal policy rule

Focusing on a single variable $\boldsymbol{x}_i$, Equation (18) implies that the space of possible allocations that the policymaker can achieve for this variable is given by

$$\boldsymbol{x}_i = \sum_{j=1}^{n_\nu} \Theta_{x_i,\nu_j,\mathcal{A}} \times \nu_j. \tag{H.1}$$

Plugging this expression into Equation (17) and taking the first-order conditions with respect to each ν_j, one arrives at the condition

$$\sum_{i=1}^{n_x} \lambda_i \Theta'_{x_1,\nu,\mathcal{A}} W \times \boldsymbol{x_i} = \mathbf{0}. \tag{H.2}$$

For each $\boldsymbol{x}_i$ the term in front of the sum describes how a change in the policy instruments $\boldsymbol{\nu}$ would translate into a change in the endogenous variable $\boldsymbol{x}_i$ and weights these changes over time using the time discount matrix W. All the implied changes are then summed over all variables x_i using the policy weight λ_i, which translates them into changes in the loss function of Equation (17). This rule then implies that the (weighted) sum of changes in the objective function resulting from a change in the policy instruments $\boldsymbol{\nu}$ has to equal zero. In other words, the gradient of the loss function with respect to the policy instruments has to be set to zero at the optimum.

This condition can be encapsulated into the matrices $\boldsymbol{\mathcal{A}_x}, \boldsymbol{\mathcal{A}_z}$ of the sequence-space representation of the model in Equation (10) by noting that the optimality condition in Equation (H.2) can be written as

$$\sum_{i=1}^{n_x} \lambda_i \Theta'_{x_i,\nu,\mathcal{A}} W \boldsymbol{x_i} = \mathcal{A}^*_{\boldsymbol{x}} \boldsymbol{x} = \mathbf{0}. \tag{H.3}$$

H.2 Estimating the impulse responses under counterfactual optimal policy

The procedure for the counterfactual optimal policy is very similar to the one sketched in the main text for the baseline optimal policy exercise but involves two additional steps.

First, we estimate impulse responses to a generic identified oil supply shock. We use the same endogenous variables as in our baseline BPSVAR model and include the 5-year German Bund yield. All variables enter the estimation in log levels if they are not already expressed in percentage terms.

Second, we identify the euro area conventional monetary policy and forward guidance shocks by combining the high-frequency proxies with the magnitude and sign restrictions described in the text. Again we use the same variables and transformations as in step 1.

Third, we use the same endogenous variables and sample as in the first step to estimate the impulse responses to a short- and long-run oil supply news shock in line with the description in Section 5.3.

Fourth, we compute the posterior distribution of each of the counterfactual impulse responses, where the euro area monetary policy shocks from the second step do not affect the by applying the procedure of McKay and Wolf, 2023 to each draw from the posterior distribution of the second and third step.

Fifth, we condition the impulse responses from the first step and compute the optimal policy counterfactual for each draw from the posterior distribution of the fourth step.

Lastly, we plot the point-wise mean which can be interpreted as summarizing the posterior distribution of impulse responses under the optimal (counterfactual) policy response conditional on the data and the impulse responses from step 1.

H.3 Additional figures for the optimal policy exercise

Figure H.1: Approximation error for the optimal policy exercise under exogenous oil prices

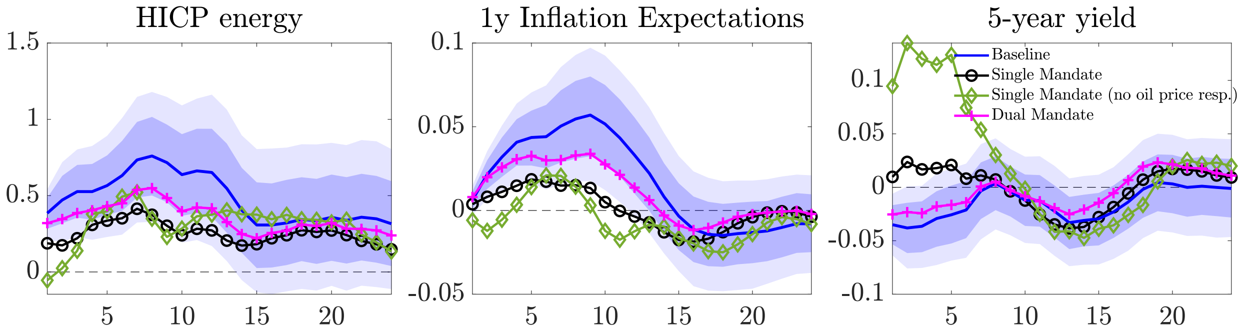

Notes: See notes to Figure 7.

Figure H.2: Approximation error for the optimal policy exercise under exogenous oil prices

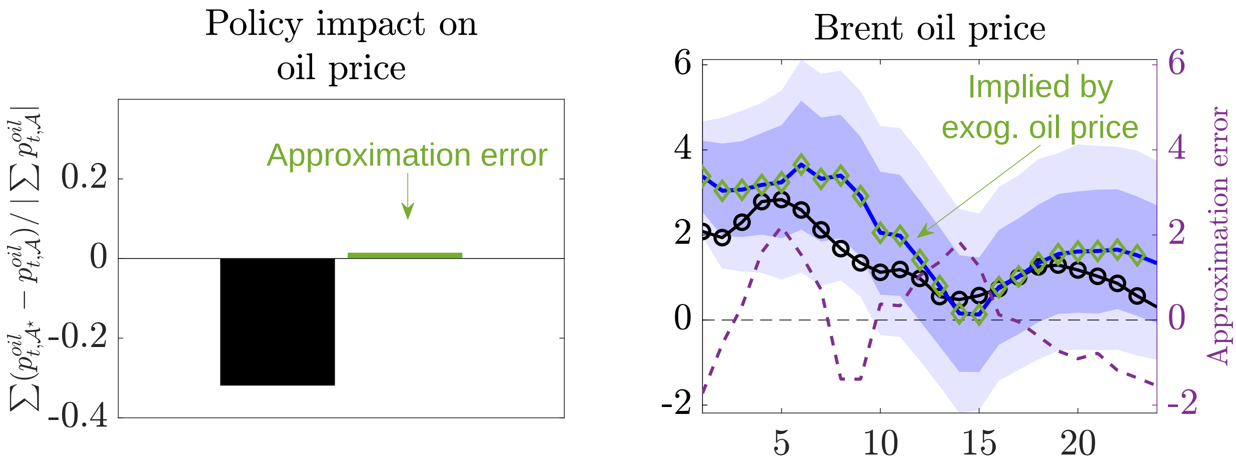

Notes: The left panel reports the impulse responses to the oil supply shock as well the transmission under the optimal policy when the ECB can (black) and cannot (green) affect the oil prices. The purple line shows the difference between the oil price response in the baseline case (blue) and the path of the oil price under the counterfactual optimal policy. As these two should coincide in theory, the difference can be traced to the approximation error from the "best Lucas-Critique-robust approximation". The right panel shows that, although the approximation may perform poorly for some horizons, the implied average changes for the oil price closely correspond to the baseline case, giving an average approximation error of close to zero.

Figure H.3: Impulse responses to an oil supply shock (blue) under Dual-Mandate optimal monetary policy when euro area monetary policy can (magenta) and cannot (purple) affect the Brent oil price

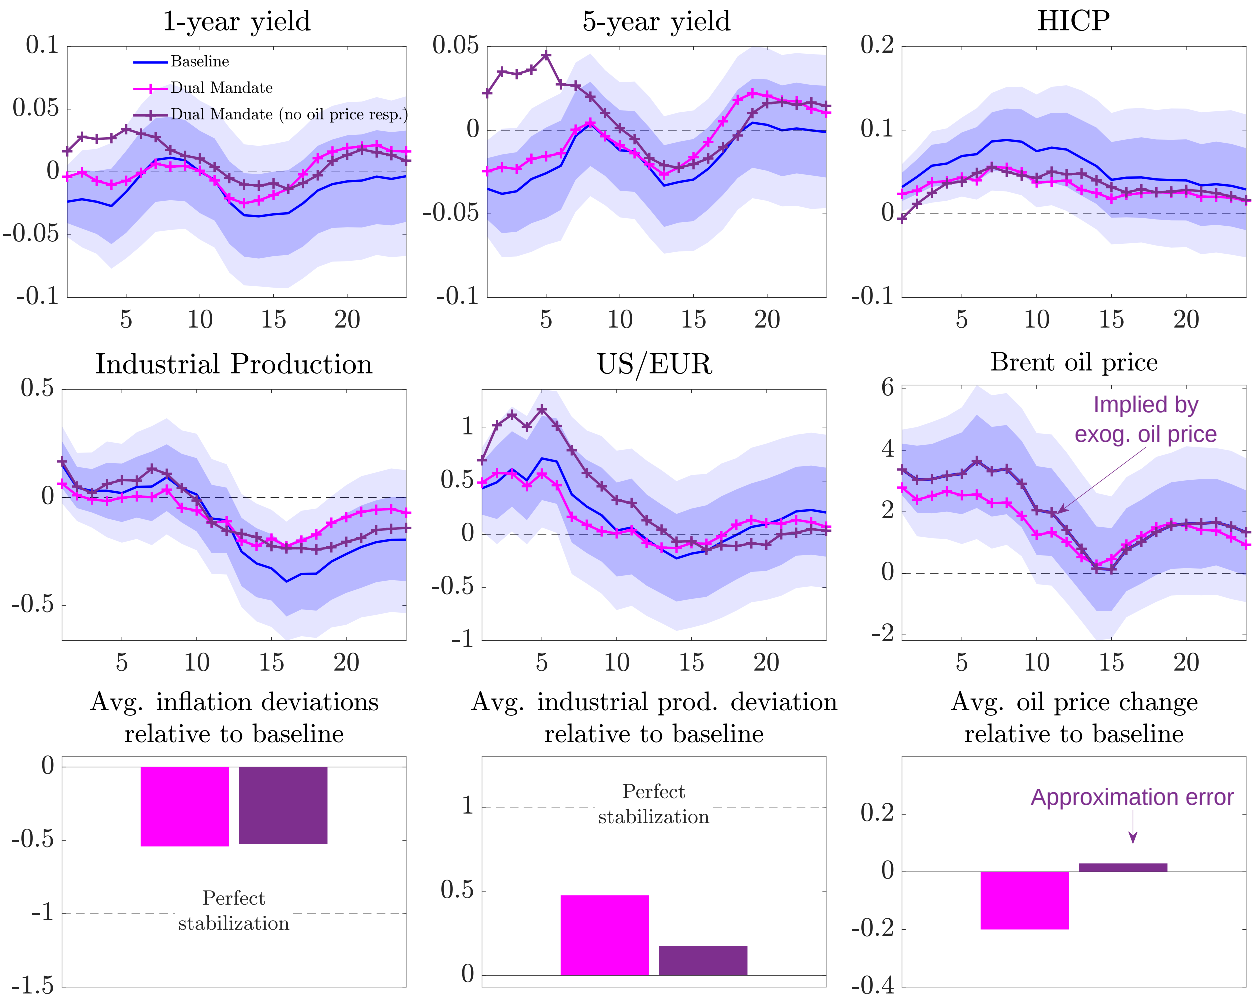

Notes: See notes to Figure 7. Under a dual mandate, we specify a loss function that gives a weight of $\lambda = 1$ to y-o-y inflation and deviations of GDP from the steady state. To map industrial production deviations into GDP deviations, we scale the hypothetical equal weight of 1 that we want to give to GDP by the relative variance of GDP and industrial production ($\approx 1/3.3$). Note that we treat industrial production as a policy target in the plot in the last row and therefore define IP deviations from the target as $|y|$. This is different from the scale used in Figure 7, where stabilizing output was not a target of the policymaker.