

Ider, Gökhan; Kriwoluzky, Alexander; Kurcz, Frederik; Schumann, Ben

Working Paper

Friend, not foe: Energy prices and European monetary policy

DIW Discussion Papers, No. 2089

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: Ider, Gökhan; Kriwoluzky, Alexander; Kurcz, Frederik; Schumann, Ben (2024) : Friend, not foe: Energy prices and European monetary policy, DIW Discussion Papers, No. 2089, Deutsches Institut für Wirtschaftsforschung (DIW), Berlin

This Version is available at:

<https://hdl.handle.net/10419/300818>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

2089

Discussion
Papers

Friend, Not Foe – Energy Prices
and European Monetary Policy

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

DIW Berlin, 2024

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<https://www.diw.de>

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<https://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<https://ideas.repec.org/s/diw/diwwpp.html>
<https://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Friend, not foe - Energy prices and European monetary policy ^{*}

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

July 2, 2024

Abstract

This paper first shows that, contrary to conventional wisdom, the European Central Bank (ECB) can influence global energy prices. Second, through Lucas critique-robust counterfactual analysis, we uncover that the ECB's ability to affect fast-moving energy prices plays an important role in the transmission of monetary policy. Third, we empirically document that, to optimally fulfill its primary mandate, the ECB should swiftly tighten policy in response to an increase in energy prices. Crucially, the tightening required depends on the ECB's ability to influence global energy prices. Finally, we find this policy strategy could have largely prevented the post-pandemic inflation episode.

Keywords: inflation, energy prices, monetary policy transmission mechanism

JEL Codes: C22, E31, E52, Q43

^{*}We thank Rüdiger Bachmann, Michael Bauer, Christiane Baumeister, Georgios Georgiadis, Refet Gürkaynak, Kloadiana Istrefi, Silvia Miranda-Agrippino, Gisle Natvik, Pascal Paul, Ricardo Reis, Fabian Seyrich, Christian Wolf and Leopold Zessner-Spitzenberg as well as seminar participants at the EEA Annual Congress 2023, Oslo Macro Conference 2023, University of Vienna, the European Central Bank, the Humboldt University Berlin and DIW Berlin for helpful comments. We additionally thank Refet Gürkaynak and Ambrogio Cesa-Bianchi for sharing monetary policy surprise data with us. This paper combines the results of two papers Ider et al. (2023) and Schumann (2024). This research has received financial support by 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

Inflation in the euro area rose steadily in 2021 and reached unprecedented levels in 2022. This surge can be primarily attributed to a sharp increase in energy prices (Arce et al. (2024)). Despite inflation soaring to heights unseen in four decades, there has been controversy about the appropriate response of the ECB. One prevalent argument was that the ECB has limited capacity to combat inflation driven by energy prices, which are determined by global markets (Lagarde, 2022b). Moreover, even if it could influence energy prices, there were concerns about whether it should intervene, as the potential costs in terms of output and unemployment might outweigh the benefits of curbing inflation (Lagarde, 2022a). In this paper, we assess both of these conjectures by empirically examining the impact of European monetary policy on energy prices and evaluating the optimal policy response to an energy price shock.

First, using a high-frequency event study and a Bayesian proxy structural vector autoregressive model (BPSVAR) model we document that, both at the intra-day and business cycle frequency, policy decisions of the ECB quickly cause changes in the prices of energy goods. Second, we examine the importance of energy prices for the transmission of monetary policy using a counterfactual scenario computed along the lines of McKay and Wolf (2023). In particular, we base our assessment of this importance on a Lucas critique-robust empirical counterfactual, where the ECB's decisions do not affect global energy prices. We document that consumer prices and inflation expectations respond considerably less to changes in the ECB's policy stance when its decisions do not impact global oil prices. Hence, by affecting fast-moving, relatively flexible energy prices, monetary policy has a tighter grip on short- and medium-term inflation dynamics. Third, given the insight that the ECB can fight surges in inflation driven by energy prices, we investigate the ECB's optimal policy response to a shock in energy prices. Using the method to empirically conduct optimal policy counterfactuals proposed by McKay and Wolf (2023), we find that, under medium-term inflation targeting, the ECB should respond more aggressively to an oil price shock than it did historically. Doing so swiftly curbs the increase in inflation while only leading to a slightly deeper front-loaded contraction in output. Importantly, we document that the degree of additional tightening required to optimally stabilize medium-term inflation crucially depends on the ECB's ability to affect global energy prices.

In light of our findings, lastly, we analyze the ECB's response to the surge in inflation that occurred between 2021 and 2023. The optimal medium-term inflation targeting strategy would have entailed a strong initial tightening which would have prevented the observed inflation overshoot of 8% at the cost of approximately 3% lower output in 2022, with economic activity recovering quickly in 2023. Our results indicate that, if the ECB were not able to affect energy prices, the tightening required to tame the inflation surge would have been larger and more persistent. Therefore, if handled correctly, energy prices can be the ECB's friend, not its foe.

In more detail, the paper first examines whether the ECB can influence energy prices. As an initial 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. Our findings show that ECB’s policy decisions right away translate into changes in global energy prices. Subsequently, we proceed to employ a BPSVAR model to study the dynamic effects of euro area monetary policy shocks on energy prices and inflation at the business cycle frequency. Crucially, our analysis reveals that a contractionary monetary policy shock significantly reduces prices of energy goods traded on the global market as well as energy prices faced by consumers in the euro area. More precisely, and in line with the theoretical work by Auclert et al. (2023), contractionary monetary policy transmits to energy prices by affecting global energy prices due to a change in domestic and global demand, and by appreciating the exchange rate, which leads to a further fall in the prices of imported (local) energy goods. Importantly, the reductions in energy prices materialize quicker and are substantially more pronounced than the changes in the corresponding headline consumer price index measure. This underscores the pivotal role played by fluctuations of fast-moving energy prices in the transmission of monetary policy.

Having established that monetary policy does affect energy prices, we set up a counterfactual exercise to examine the importance of this finding for the transmission of monetary policy in the euro area. We base our assessment on an empirical counterfactual in which the ECB’s decisions — as frequently claimed in its press conferences — do not affect global energy prices.¹ More precisely, in this counterfactual setting, the policy rule of the Organization of the Petroleum Exporting Countries (OPEC) is such that it aims to stabilize deviations of the global oil price from its preferred path, and consequently euro area monetary policy does not influence the global oil price. The estimation of this counterfactual is based on the method developed by McKay and Wolf (2023), which takes into account the anticipatory effects of such an OPEC policy rule change and is thus robust to the Lucas critique. McKay and Wolf (2023) prove that, for a large class of underlying structural models that generate the data, their approach alongside valid identification of the policy shock recovers the true counterfactual. To implement this approach, we borrow from the literature on high-frequency identification of oil supply shocks (Känzig (2021)) to jointly identify a short- and a long-run oil supply news shock as well as a euro area monetary policy shock in our BPSVAR model. When the ECB’s decisions do not impact the global oil price, the response of energy prices faced by euro area consumers to a monetary tightening is substantially muted and crucially, the transmission

¹For instance, at the press conference on the 3rd of February 2022 Christine Lagarde responded to a question on the 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 (2022b)) Presumably this argument is derived from the small open economy assumption. The assumption that the ECB does not affect global energy prices is not only incorporated into theoretical models used for policy analysis but also deeply embedded into their forecasting process (see Christoffel et al. (2008) for an example and a discussion).

to inflation and inflation expectations is considerably weaker. Comparing the counterfactual responses to the baseline reveals that, by affecting fast-moving energy prices which are known to be relatively flexible compared to the other goods in the consumption basket, monetary policy has a tighter grip on inflation dynamics especially in the short- to medium-term.

Given the finding that the ECB’s ability to influence energy prices plays a crucial role in the transmission of monetary policy shocks, next, we study how this effect shapes the optimal conduct of monetary policy in response to an energy price shock. In particular, we contrast the optimal response to an energy price shock in the case where the ECB can affect global energy prices, to the optimal allocation in a scenario where this is not the case. To this end, we integrate the framework of McKay and Wolf (2023) to compute optimal policy with their approach to estimate policy counterfactual impulse responses. This enables us to focus on how the ECB’s ability to affect energy prices shapes the optimal conduct of monetary policy for any given target criterion. In line with the literature that aims to empirically assess optimal monetary policy using a sufficient statistics approach, we define the optimal policy for the ECB as the policy that optimally achieves the primary mandate, i.e., stabilize inflation in the medium term (cf. Barnichon and Mesters (2023), Barnichon and Mesters (2024), McKay and Wolf (2023)). To implement this approach we jointly identify an oil price shock, a conventional monetary policy shock, and a forward guidance shock within our BPSVAR model.

Under the baseline policy rule, the estimated endogenous response of the ECB to an oil price shock indicates that, in line with their public statements, the ECB historically followed a “looking-through” strategy when faced with an energy price shock (Schnabel (2022)). On the contrary, the optimal policy response would entail a front-loaded tightening at the short and longer end of the yield curve. This would almost completely curb the rise in inflation at the cost of a slightly deeper yet short-lived contraction in output. The rationale behind this finding is that the more contractionary policy stance under optimal policy rapidly lowers global energy prices, consequently resulting in a substantially smaller increase in headline inflation and inflation expectations. Only a marginal additional decrease in output is sufficient to minimize medium-term inflation deviations from the target, as contractionary ECB policy causes quick and sizeable declines in relatively more flexible energy prices. Therefore, our findings suggest that the tightening necessary to optimally achieve the mandate is tightly linked to the ECB’s ability to affect global energy prices.

To further substantiate this notion we, as a next step, estimate the optimal policy response under the assumption that the ECB’s policy decisions do not affect global energy prices. The results of this thought experiment suggest that the optimal response of the ECB to an oil price shock would markedly differ if the ECB were not in a position to influence global energy prices. In this case, the optimal strategy would entail a substantially stronger tightening, especially at the longer end of the yield curve. Interestingly, even in this case, the optimal policy decision would not require engineering a major economic contraction to put downward pressure on the prices of domestically produced goods. Rather, it implies using additional

tightening at the longer end of the yield curve to strongly appreciate the euro, which in turn lowers fast-moving local energy prices. This finding further underscores the central role of energy prices in the conduct of monetary policy.

Finally, we apply our framework to the recent energy-price-driven surge in inflation in the euro area from 2021 to 2023. We first show that monetary policy decisions during this period did indeed contribute to the rise and fall of energy prices and inflation. We then use the recently developed method of Caravello et al. (2023) to estimate Lucas critique-robust historical counterfactuals, which allow us to quantify the additional tightening that would have been necessary to optimally stabilize inflation at the medium-term target. In particular, for the ECB to optimally achieve the 2% medium-term inflation target, the scenario would have required a strong initial tightening, which would have stabilized energy prices and brought inflation and inflation expectations in line with the target. This would have allowed the ECB to lower interest rates as early as 2023. By following this optimal strategy, we estimate that the ECB could have prevented inflation from exceeding the target by more than 8% at the cost of GDP being about 3% lower at its trough in 2022. Crucially, in the face of the recent energy price spike, the optimal strategy would have resulted in only a short-lived contraction, with output recovering quickly in 2023. According to our results, this is partly due to the fact that the ECB's monetary policy decisions affect energy prices, thereby directly addressing the source of the shock.

Related literature. Our paper contributes to the extensive literature that studies how monetary policy transmits to the economy (Christiano et al. (1999); Gertler and Karadi (2015); Miranda-Agrippino and Ricco (2021)). Relative to existing work on the transmission channels of monetary policy, we focus on and analyze how the response of energy prices to a monetary policy shock shapes the transmission of monetary policy to inflation and inflation expectations. Additionally, our findings connect to the literature studying the transmission pace of monetary policy and provide further evidence against the notion that monetary policy 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 output and inflation already in the very short run (Miranda-Agrippino and Ricco (2021), Miranda-Agrippino and Rey (2020), Bauer and Swanson (2023)). We not only confirm this finding using euro area data 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 that studies the transmission of energy price shocks conditional on the reaction of monetary policy. There exists a large literature analyzing the macroeconomic effects of oil price shocks (Baumeister and Hamilton (2019); Känzig (2021); and many others) and studying whether the response of monetary policy historically exacerbated the effects of oil price shocks (Bernanke et al. (1997); Leduc and

Sill (2004); Kilian and Lewis (2011)). Although our analysis provides insights about both questions, our novel contribution is the empirical analysis of the mandate-optimal monetary policy response to an oil price shock, which until recently could only be studied with theoretical models (Bodenstein et al. (2012); Natal (2012)).

In addition, our paper is related to the literature that studies the effects of monetary policy on commodity prices (Anzuini et al. (2012); Rosa (2014); Degasperis et al. (2023); Miranda-Pinto et al. (2023); Ca'Zorzi et al. (2023)). While much of the existing work focuses on the response of energy commodity prices to US monetary policy, we show that a similar response in sign and magnitude is present to European monetary policy as well. Crucially, and in contrast to the existing work, we employ a Lucas critique-proof approach to 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)).

Lastly, our paper contributes to the recent surge in sufficient statistics approaches to macroeconomic policy evaluation (McKay and Wolf (2023), Barnichon and Mesters, 2023, Caravello et al. (2023)). We not only operationalize these approaches in a fully coherent Bayesian inference framework but also, to our knowledge, are the first to use them to estimate the contribution of oil prices to the transmission of euro area monetary policy and the optimal reaction of the ECB to an oil price shock. Furthermore, we embed different approximations to the solution of the McKay and Wolf (2023) approach into a joint framework and show how to use these to learn something about the role of expectations in the transmission of a shock. We also extend the approach of Caravello et al. (2023) to show how to decompose the difference between the actual and a counterfactual evolution of the economy into a part related to expectations about the future and a part that is due to changes in the propagation of incoming shocks.

The rest of the paper is structured around three main questions. Section 2 and Section 3 answer the question if euro area monetary policy can affect energy prices. Section 4 examines if these effects on energy prices matter for the transmission of exogenous monetary policy shocks to the euro area economy. Section 5 then provides an answer to the question if the ability of the ECB to affect energy prices matters for the optimal conduct of endogenous monetary policy. Section 6 studies the role of euro area monetary policy in the recent energy-price-driven inflation surge in 2021-2023. The final section concludes.

2 Monetary policy and oil prices - a high frequency analysis

In this section, we analyze if European monetary policy translates to changes in the global energy price at high frequency. Throughout the paper we use the Brent crude oil price as a

stand-in for the global energy price.² To put the results for the euro area into perspective, we compare them with the results for the U.S. and the U.K., which are arguably a large open economy (LOE) and a small open economy (SOE), respectively.

2.1 High-frequency data

A prerequisite for conducting a high-frequency event study to investigate the effects of monetary policy of different central banks is a reliable and comparable measure of the surprise component of the policy change. For the US, we utilize the commonly used intra-day changes in the three-month-ahead federal funds futures around Fed monetary policy announcements, and apply the “poor-man’s” approach to purge these monetary policy surprises from central bank information effects (Jarociński and Karadi (2020)).³ We construct a similar instrument for the euro area (EA) by replacing the change in the three-month-ahead federal funds futures with the change in the three-month OIS rate from the database of Altavilla et al. (2019). For the UK, we use the intraday variation in the three-month Libor rate around Bank of England (BoE) monetary policy announcements measured by Cesa-Bianchi et al. (2020).⁴

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 of the ECB, the Fed and the BoE. 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. We closely follow the methodology outlined in the online appendix of Altavilla et al. (2019) to measure a pre- and post-announcement price. Further details on the sources and transformations for data used in this section are provided in Table A.1 in Appendix A.

2.2 Estimation and Results

To study the effects of monetary policy on the global oil price, we estimate the following high-frequency event study regression for the ECB, the Fed, and the BoE separately:

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

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

²In Appendix B, we show that results are robust to using the natural gas price (Dutch TTF) instead of the Brent oil price.

³US monetary policy surprise data is from Gürkaynak et al. (2005). We employ the “poor-man’s” approach of Jarociński and Karadi (2020) because it is easily applicable for all three countries in our sample. As a robustness check, we also use the orthogonalized monetary policy surprises of Bauer and Swanson (2023) for the Fed, which are not available for the ECB nor the BOE. The results for the US are almost identical for both surprise measures, as can be seen in Appendix B.

⁴Ambrogio Cesa-Bianchi has kindly provided us an extended series of the BoE monetary policy surprises until March 2021.

policy surprise of country i .

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

	EA	US	UK
$\hat{\beta}^{100\text{bps}}$	-3.34** (1.54)	-2.15** (1.03)	0.36 (0.68)
$\hat{\beta}^{std}$	-0.056** (0.026)	-0.075** (0.036)	0.019 (0.037)
Sample	2002:1-2019:12	1996:1-2019:6	1997:6-2019:12
N	182	195	246
R^2 (%)	3.37	2.61	0.38

Notes: Coefficient estimates $\hat{\beta}^{100\text{bps}}$ ($\hat{\beta}^{std}$) measures the percentage change in the front month future of the Brent crude oil price following a 100bps (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 results are presented in Table 1, with additional material and robustness along several dimensions such as sample size, instrument choice and choice of the energy price being relegated to Appendix B. In the first row of Table 1, the coefficient $\hat{\beta}^{100\text{bps}}$ represents the estimated percentage change in oil prices in response to a 100 basis point contractionary monetary policy surprise.⁵ The first column displays our benchmark results for the euro area, which shows that global energy prices, as measured by the spot price of Brent crude oil, decline significantly and immediately after a euro area monetary tightening. This stands in stark contrast to the results for the UK, where the oil price does not significantly react to a policy tightening. This finding is in line with the notion that the UK is indeed a small open economy in the global energy market. Conversely, and in line with the assumption that the US is a LOE, we find that the oil price also significantly falls after a surprise Fed tightening. Thus, at least when viewed through the lens of this exercise, it seems that the euro area is by no means a small open economy in the global energy market. Instead, it seems to be a big player with significant market impact and a central bank that can steer energy prices on the world market. Indeed, the euro area is the second largest oil importer for the majority of our sample period only trumped by in early periods by the US and recently by China (see Figure

⁵For all countries, we stop the sample at the onset of the pandemic and choose the sample period in order to maximize the observations for each country. For the euro area this implies that we initiate the sample in 2002, consistent with Altavilla et al. (2019) and Andrade and Ferroni (2021), who exclude the 1999-2001 period due to liquidity issues on OIS contracts.

L.1).

When comparing the results for the euro area to those for the United States, the comparatively larger coefficient in front of ECB monetary policy surprises stands out. While in work subsequent to ours Miranda-Pinto et al. (2023) find a similar pattern, it is important to clarify that this does not imply that European monetary policy shocks of the same “size” have a stronger impact than those in the United States. It rather depends on how the “size” of the shock is measured. In particular, this result is primarily attributable to the fact that a 100 basis point surprise in the 3-month Overnight Index Swap (OIS) rate constitutes a historically much larger shock in the euro area compared to the United States. To address the relative sizes of these shocks, we conduct an analysis using standardized monetary policy surprises. In the second row of Table 1, we report the responses of the oil price to a one standard deviation – i.e. an average — contractionary monetary policy surprise, as measured by the coefficient $\hat{\beta}^{\text{std}}$. Although an average euro area monetary policy surprise still significantly moves the global oil price, it does so to a smaller extent than its US counterpart.

3 Monetary policy and energy prices in a dynamic setting

Section 2 shows that monetary policy decisions in the euro area have an instantaneous effect on global energy prices. In this section, we investigate how this high-frequency effect plays out dynamically at the business cycle frequency in order to measure the aggregate effect of the ECB’s policy decision on energy prices and the economy. To this end, we estimate a BPSVAR model.

3.1 The Bayesian Proxy SVAR model

In this section, we briefly lay out the BPSVAR model for the general case with $k \geq 1$ proxy variables and k structural shocks of interest. We keep the notation general because we simultaneously identify an oil supply news shock later in the paper. Following the notation of Rubio-Ramirez et al. (2010), the structural VAR model with one lag and without deterministic terms can be written as:

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

where \mathbf{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 \mathbf{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

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

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

and represent the relevance and the exogeneity condition, respectively.

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

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

The algorithm by Arias et al. (2021) imposes the assumptions (3a) and (3b) in the estimation of (4) to identify the structural shocks. We relegate the details of the algorithm to Appendix D where we also discuss the advantages of this Bayesian approach relative to the standard frequentist two-step proxy SVAR estimation.

3.2 Data and specification

Our baseline monetary VAR model for the euro area includes eight variables and, as a proxy, a high-frequency monetary policy surprise series to identify an ECB monetary policy shock. Specifically, our baseline version includes the 1-year constant maturity yield on German Bunds as a monetary policy indicator. Economic activity is measured by the euro area industrial production index (excluding construction). We use the Harmonised Index of Consumer Prices (HICP) as a measure of the overall (headline) Consumer Price Index (CPI) and the energy component of the HICP as a measure of local energy prices in the euro area. The one-year-ahead inflation forecast from Consensus Economics and the BBB corporate bond spread are included to capture inflation expectations and financial conditions, respectively. As in the high-frequency study, we use the Brent crude oil price as a measure of the global energy price and also add the EUR-USD exchange rate, as oil and other energy commodities are generally traded in US dollars.

As in Section 2 and following Jarociński and Karadi (2020), we utilize the high-frequency changes in the 3-month OIS rate over the monetary event window as our preferred proxy for the euro area monetary policy shocks. Again, we apply the poor-man's approach of Jarociński and Karadi (2020) to purge these surprises from central bank information effects based on the sign of the corresponding equity-price surprise. The variables are measured with a 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.

More details on the data can be found in Appendix A.

The BPSVAR model is estimated on a sample from January 1999 to December 2019. This omits the extraordinary volatility in the data caused by the COVID-19 pandemic.⁶ The model has 12 lags and includes a constant. Finally, we use flat priors for estimating the BPSVAR parameters.⁷ In addition, a relevance threshold is imposed to express the prior belief that the proxy is informative to identify monetary policy shocks. We set a prior $\gamma = 0.1$, imposing a threshold that the identified structural monetary policy shocks account for at least 10% of the variance in the proxy.⁸

3.3 Dynamic effects of a monetary policy shock

Figure 1 presents our baseline estimates of the effects of a one standard deviation contractionary monetary policy shock for the euro area. The 1-year Bund yield increases on impact and quickly reverts to zero, with an overall shape and magnitude that is very similar to Jarociński and Karadi (2020). Industrial production falls by approximately 0.35% on impact and remains depressed for about 1.5 years. Likewise, the fall in the domestic headline consumer price level is immediate, reaching a trough of about 0.07% after about 18 months. At the same time, the euro appreciates against the dollar by slightly less than 1% and remains significantly elevated for a year. Financial conditions tighten mildly, while inflation expectations decline quite strongly. In summary, all of the estimated dynamics for the endogenous variables described so far are in line with standard theory and in line with previous findings in the literature.

Our main result of this section is the substantial decline in the measures of local and global energy prices. The global oil price (in US dollars) falls sharply by about 3%. In addition, the local energy price index, as measured by the energy component of the HICP, falls by 0.5%, much more than the fall in the overall price of the HICP basket. Indeed, given that energy prices have a weight of about 10% in the overall HICP basket, a back-of-the-envelope calculation suggests that the overwhelming part of the decline in the overall HICP price level in the short and medium run can be attributed to the effect of the contractionary monetary policy shock on local energy prices.

Furthermore, our results corroborate the findings from the micro-data literature that shows at the consumer level that energy goods are by a large margin the goods with the highest frequency of price updating. For instance, when analyzing micro-data used for the computation of the Belgian HICP, Aucremanne and Dhyne (2004) show that the average


⁶We include and explicitly model this period along the lines of Cascaldi-Garcia (2022) in an extension in section 6. The results are robust to this, as shown in Figure C.5.

⁷As 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.

⁸This 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 C.7 in the Appendix, our results are robust to reducing the relevance condition to 0.

price duration for energy goods is roughly 1 month which stands in stark contrast to the median price duration of all goods in the basket, which amounts to approximately 14 months. Intuitively, when viewed through the lens of a standard New-Keynesian Model, this implies that all else equal the Phillips curve for energy goods has a steeper slope than for average goods. In line with this theoretical intuition, while virtually all subcomponents of the HICP energy component display a significant decline in their price (see Figure C.8 in the Appendix), the price of fuels, arguably the subcomponent with more flexible prices, contract the most.

Figure 1: Baseline Euro Area SVAR model



Notes: Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months.

In Appendix C we show that our results are robust to alternative specifications. To address concerns regarding the ZLB period we replace the 1-year yield with the 2-year yield to use a less constrained monetary policy indicator as in Swanson and Williams (2014) (Figure C.1). 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, oil inventories, and global industrial production (Figure C.2). Furthermore, we document that our results are robust to starting the sample in 2002 to address the concerns raised by Altavilla et al. (2019) and Andrade and Ferroni (2021) regarding the liquidity of the OIS contracts during the early years of the euro area. In this specification, the precision of the estimation even slightly increases (Figure C.3). As 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 is robust even in the case of non-invertibility (Figure C.4). The results are very similar; if anything, the oil price is estimated to

drop even more following a monetary policy shock of a similar size. Lastly, we show that our results are largely unchanged when incorporating the pandemic into the estimation alongside the Pandemic Priors approach of Cascaldi-Garcia (2022) (Figure C.5).

In addition, we also perform robustness exercises with respect to the proxy variable. Miranda-Agrippino and Ricco (2021), and Bauer and Swanson (2023) show that accounting for potential serial correlation in the proxy can considerably alter the results. Although our model already takes care of this issue at the (aggregated) monthly frequency, it is possible that there could be serial correlation present at the meeting frequency. Therefore, we estimate an AR(12) on the meeting-frequency data and use the residuals as our proxy instead. The results are robust to this specification (Figure C.6) and also remain unchanged when we remove our prior on the relevance of the proxy variable (Figure C.7).

Lastly in Appendix E, we address the concern that the implied 3% drop in the oil price seems quite large given that the one standard deviation monetary shock is estimated to increase the short-term interest rate by roughly 5 bps. To address these concerns, we first simulate a similarly sized shock in the state-of-the-art model of Bayer et al. (forthcoming) to analyze our results through the lens of a theoretical model that explicitly models the market for energy goods. Additionally, we also plot the impulse responses of these variables in terms of their respective standard deviations, to put the estimated magnitudes into a historical perspective. Both approaches indicate that our results are by no means excessively large.

In summary, the high-frequency analysis in Section 2 shows that a change in the ECB's policy stance has a significant impact on global energy prices. Combined with the results of the BPSVAR model analysis in this section, it is clear that monetary policy affects not only global but also local energy prices in the euro area.

4 The role of energy prices for the transmission of European monetary policy

Our results show that euro area monetary policy has a significant impact on global energy prices and, consequently, on the energy prices faced by euro area consumers. Moreover, the impulse response functions and back-of-the-envelope calculations suggest that energy prices play an important role in the monetary transmission mechanism. To further substantiate this notion, in this section we conduct an empirical counterfactual exercise in which the global oil price does not respond to a change in the ECB's monetary policy stance. In particular, we calculate the response to a euro area monetary policy shock under a scenario where OPEC's policy rule, counterfactually, is such that it aims to perfectly stabilize the oil price. We first present the general setup and the different ways to approximate the solution before presenting the results of the counterfactual analysis.

4.1 Computing structural (policy) counterfactuals

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 ingredients in their counterfactual analysis are impulse responses to shocks to current and future policy. In particular, they show that 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, uncovers 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 $t = 0, 1, 2, \dots$. In sequence-space notation, this economy can be described by a set of equations

$$\mathcal{H}_x \mathbf{x} + \mathcal{H}_z \mathbf{z} + \mathcal{H}_\epsilon \boldsymbol{\epsilon} = \mathbf{0}, \quad (5)$$

$$\mathcal{A}_x \mathbf{x} + \mathcal{A}_z \mathbf{z} + \boldsymbol{\nu} = \mathbf{0}, \quad (6)$$

where $\mathbf{x} \equiv (\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_{n_x})'$ stacks the time paths of the n_x endogenous variables over n_h periods, analogously \mathbf{z} stacks the time path of the n_z policy instruments. The matrices \mathcal{H} summarize the behavior of agents in the non-policy block, while the matrices \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 (5) and (6) is that $\{\mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_\epsilon\}$ do not depend on the coefficients of the policy rule $\{\mathcal{A}_x, \mathcal{A}_z\}$, so that policy affects the non-policy block's decisions only through the path of the instrument \mathbf{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 (5) and (6) can be written in impulse response space as

$$\begin{pmatrix} \mathbf{x} \\ \mathbf{z} \end{pmatrix} = \boldsymbol{\Theta}_{\mathcal{A}} \times \begin{pmatrix} \boldsymbol{\epsilon} \\ \boldsymbol{\nu} \end{pmatrix}, \quad \boldsymbol{\Theta}_{\mathcal{A}} \equiv (\boldsymbol{\Theta}_{\epsilon, \mathcal{A}}, \boldsymbol{\Theta}_{\nu, \mathcal{A}}) \equiv \begin{pmatrix} \boldsymbol{\Theta}_{x, \epsilon, \mathcal{A}} & \boldsymbol{\Theta}_{x, \nu, \mathcal{A}} \\ \boldsymbol{\Theta}_{z, \epsilon, \mathcal{A}} & \boldsymbol{\Theta}_{z, \nu, \mathcal{A}} \end{pmatrix}. \quad (7)$$

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

factual policy rule is given by:

$$\tilde{\mathcal{A}}_x \mathbf{x} + \tilde{\mathcal{A}}_z \mathbf{z} = \mathbf{0}, \quad (8)$$

where $\tilde{\mathcal{A}}_x$ and $\tilde{\mathcal{A}}_z$ contain the corresponding coefficients of the counterfactual rule. MW show that knowledge of the impulse responses $\Theta_{\mathcal{A}}$ 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

$$\mathbf{x}_{\tilde{\mathcal{A}}}(\epsilon) = \Theta_{x,\epsilon,\mathcal{A}} \times \epsilon + \Theta_{x,\nu,\mathcal{A}} \times \tilde{\nu}, \quad \mathbf{z}_{\tilde{\mathcal{A}}}(\epsilon) = \Theta_{z,\epsilon,\mathcal{A}} \times \epsilon + \Theta_{z,\nu,\mathcal{A}} \times \tilde{\nu}. \quad (9)$$

In words, the impulse response to the structural shock ϵ under the counterfactual policy rule $\mathbf{x}_{\tilde{\mathcal{A}}}(\epsilon) \equiv \Theta_{x,\epsilon,\tilde{\mathcal{A}}} \times \epsilon$ is exactly equivalent to a combination of the corresponding impulse responses under the baseline policy rule $\Theta_{x,\epsilon,\mathcal{A}} \times \epsilon$ and the impulse responses to a specific sequence of policy news shocks $\tilde{\nu}$. 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 $\tilde{\nu}$ are chosen such that the counterfactual policy rule holds

$$\tilde{\mathcal{A}}_x [\Theta_{x,\epsilon,\mathcal{A}} \times \epsilon + \Theta_{x,\nu,\mathcal{A}} \times \tilde{\nu}] + \tilde{\mathcal{A}}_z [\Theta_{z,\epsilon,\mathcal{A}} \times \epsilon + \Theta_{z,\nu,\mathcal{A}} \times \tilde{\nu}] = \mathbf{0}. \quad (10)$$

What needs to be determined are the expressions $\Theta_{x,\nu,\mathcal{A}}$ and $\Theta_{z,\nu,\mathcal{A}}$ in Equation (9). 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 for all n_h periods. We approach the problem as follows. We start by stacking the system in Equation (10) across all the responses of all the $n = n_x + n_z$ endogenous variables \mathbf{x} and the policy instrument \mathbf{z} in Equation (10) and all horizons $n_H = n \times n_h$ in order to arrive at:

$$\tilde{\mathcal{A}}\Theta_{\epsilon,\mathcal{A}} \times \epsilon + \tilde{\mathcal{A}}\Theta_{\nu,\mathcal{A}} \times \tilde{\nu} = \mathbf{0}, \quad (11)$$

Afterwards, we approximate the solution to Equation (11) using three different techniques, which despite their seemingly different nature, are all nested in the stacked sequence space representation in Equation (11). In particular, we first employ the “best Lucas critique-robust approximation” (McKay and Wolf (2023, p.5)) of the problem, which amounts to choosing the counterfactual policy shocks $\tilde{\nu}$ so as to minimize the squared deviations from the counterfactual policy rule. Second, we use the procedure of Sims and Zha (2006), which has a long-standing tradition in the SVAR literature and which we show is nested in the representation in Equation (11). Finally, we use a “hybrid” version of the two approaches, which is briefly sketched in the online appendix of McKay and Wolf (2022b).

Least squares approximation: The first approach to computing an approximation to the solution of Equation (11) aims to have the counterfactual policy rule hold at point 0 and for all n_h periods ahead expectations. This is akin to a policy counterfactual in a fully structural (Dynamic Stochastic General Equilibrium (DSGE)) model, where agents are forward-looking and know exactly the structure of the economy and therefore form consistent, rational expectations about the path of the policy instrument.

MW discuss the estimation of the corresponding counterfactual impulse responses in the situation where the researcher has only a subset of the required policy news shocks ($\boldsymbol{\nu}^{\text{LSQ}} \in \boldsymbol{\nu}$) available. This constrains the space of allocations that can be implemented via policy changes, and thus requires an approximation to the solution of Equation (11), since the desired counterfactual may not lie in that space. In such a case, which is the relevant case in practice, MW advocate choosing a linear combination of the corresponding impulse responses that solves

$$\min_{\tilde{\boldsymbol{\nu}}^{\text{LSQ}}} \|\tilde{\mathbf{A}}\boldsymbol{\Theta}_{A,\epsilon} \times \boldsymbol{\epsilon} + \tilde{\mathbf{A}}\boldsymbol{\Theta}_{\nu^{\text{LSQ}},\mathcal{A}} \times \tilde{\boldsymbol{\nu}}^{\text{LSQ}}\|, \quad (12)$$

and thereby enforces the desired counterfactual rule as well as possible (in a least squares sense). This solution can be computed from Equation (11) by collecting all impulse response functions to the $n_{\nu^{\text{LSQ}}}$ identified policy (news) shocks in $\boldsymbol{\Theta}_{\nu^{\text{LSQ}},\mathcal{A}}$, by assuming that $\boldsymbol{\Theta}_{\nu,\mathcal{A}}^{\text{LSQ}} = [\boldsymbol{\Theta}_{\nu^{\text{LSQ}},\mathcal{A}}, \mathbf{0}_{(n \times n_h) \times (n_h - n_{\nu^{\text{LSQ}}})}]$, and by solving for $\tilde{\boldsymbol{\nu}}^{\text{LSQ}} = -(\tilde{\mathbf{A}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}^{\text{LSQ}})^* \tilde{\mathbf{A}}\boldsymbol{\Theta}_{A,\epsilon} \times \boldsymbol{\epsilon}$ with $(\tilde{\mathbf{A}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}^{\text{LSQ}})^*$ as the Moore-Penrose inverse of $\tilde{\mathbf{A}}\boldsymbol{\Theta}_{\nu,\mathcal{A}}^{\text{LSQ}}$.

While the approximation is fully robust to the Lucas critique because it only leverages policy news shocks announced at period 0 and does not make any assumptions on the expectations of agents, it only minimizes the squared deviations from the counterfactual policy rule described by $\tilde{\mathbf{A}}$ and thus does not perfectly enforce it. Although the error of approximation vanishes as the number of identified policy shocks $n_{\nu^{\text{LSQ}}}$ approaches the number of periods n_h , it can be large when $n_{\nu^{\text{LSQ}}} < n_h$.

Sims-Zha-approximation: The second approximation of the solution to Equation (11) is based on the procedure proposed by Sims and Zha (2006, henceforth SZ). Traditionally, this approximation is sketched as a hypothetical scenario where a policymaker has access to only a single (contemporaneous) policy shock (i.e. ν_0) and then, successively in each period, surprises agents by choosing the size of the policy shock so that the counterfactual policy rule holds perfectly after all shocks have played out. But this seemingly sequential procedure can be parsed into the stacked sequence space representation in Equation (12), which features sequences of news shocks $\boldsymbol{\nu}$ announced at horizon 0 instead of contemporaneous policy shocks announced at horizon $t = 0, 1, \dots$. In particular, the assumption underlying the SZ procedure is then that $\boldsymbol{\Theta}_{\nu,\mathcal{A},j}^{\text{SZ}}$ is a lower triangular matrix, which would be the case if agents in the underlying structural model are fully myopic with respect to announced changes

in the policy instrument. Consequently, under this assumption, the full matrix $\Theta_{\nu,\mathcal{A},j}^{\text{SZ}}$ can be constructed and thus the constraint can be perfectly enforced, using the knowledge of the impulse responses to a single contemporaneous policy shock $\nu_{0,t}$.

More formally, let $\Theta_{\nu,\mathcal{A}}^{\text{SZ}} = [\Theta'_{\nu,\mathcal{A},1}, \Theta'_{\nu,\mathcal{A},2} \dots \Theta'_{\nu,\mathcal{A},n}]'$. For each variable j the $\Theta_{\nu,\mathcal{A},j}$, has to be lower triangular and has to have a specific structure. In particular, each column c of matrix $\Theta_{j,\nu,\mathcal{A}}$, which describes the n_h responses of variables j to an announced change in the policy instrument c periods ahead, is given by $\Theta_{\nu,\mathcal{A},j,c} = [\mathbf{0}'_{1 \times (c-1)}, \Theta'_{\nu_{0,t},\mathcal{A},j,0:(n_h-c)}]'$ where the vector $\Theta_{\nu_{0,t},\mathcal{A},j,0:(n_h-c)}$ describes the impulse responses of variable j to the contemporaneous policy shock $\nu_{0,t}$ for periods 0 to $n_h - c$. The solution to the problem in Equation 11 is then given by $\tilde{\nu}^{\text{SZ}} = -\left(\tilde{\mathcal{A}}\Theta_{\nu,\mathcal{A}}^{\text{SZ}}\right)^{-1} \tilde{\mathcal{A}}\Theta_{A,\epsilon} \times \epsilon$.

Although this approximation perfectly enforces the counterfactual policy rule ex-post (i.e. once all shocks have played out), the approximation only corresponds to the true solution and thereby the counterfactual, if the effects of the policy news shocks only materialize once the change in the policy instrument is actually enacted. Consequently, this approximation does not incorporate anticipation effects and thereby neglects the role of expectations about future policies. If agents are forward-looking and the policy communication is effective (i.e. they know about the announced change in the policy instrument in period s before it materializes in period t), this approximation deteriorates. Therefore, the accuracy of the approximation depends on the degree of myopia of agents in the non-policy block and/or the information set of these agents in the (unknown) structural model underlying the data.

“Hybrid”-approximation: The third approximation to the solution of Equation (11) combines the two approaches described above. While the first approximation tries to best enforce the counterfactual policy rule (in a least squares sense) at each time t and in expectations over all horizons n_h , the second approximation neglects the expectations component and rather enforces the counterfactual policy rule perfectly ex-post in each period. Since both are on the opposite spectrum regarding the degree of foresight of agents in the non-policy block of the true, unknown underlying model, a natural compromise would be the “hybrid” approach that combines both approaches.

In particular, MW in their online appendix sketch a sequential procedure for obtaining an approximation that enforces the policy rule both ex-post and in n_e period ahead of time using impulse responses to $\nu^{\text{hybrid}} \in \nu$ policy shocks with $n_{\nu^{\text{hybrid}}} = n_e + 1$. In their sequential procedure, the policymaker announces in each period that the counterfactual policy rule holds today and can be expected to hold for the next n_e periods.⁹ In the next period, agents are then surprised that she extends her promise for another period, causing agents to revise their expectations about this newly revealed path of the policy instrument for the next period, and thereby causing them to reconsider their decisions for the remaining periods. This revision

⁹This is reminiscent of an optimal policy problem, where the policymaker lacks a commitment device to convince agents of changes in the policy stance more than n_e periods ahead.

would imply that, despite their previous commitment, the path of the policy instrument would be different from the path announced one period ago. Therefore, the policymaker issues a set of well-specified policy news shocks that perfectly enforce the previously announced path of the policy rate and imply that the policy rule again holds at time t in the next n_e period ahead expectations. In Appendix F we provide a detailed description of how this sequential approach can be parsed into the stacked sequence space representation of Equation (11) and discuss which assumptions this approximation implies for the unknown, underlying structural model. In a nutshell, these assumptions imply that, in the unknown structural model that underlies the data, policy at time t can only manipulate expectations about the policy instrument for the next n_e periods. As in the previous cases, the assumptions behind this approach imply that the matrix $\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}$ has a certain structure, which can be fully recovered using only the impulse responses to the $n_{\nu, \text{hybrid}} = n_e + 1$ policy shocks, causing the matrix to be invertible. This implies that the hybrid solution to the problem in Equation (11) is given by $\hat{\nu}^{\text{hybrid}} = -\left(\tilde{\mathcal{A}}\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}\right)^{-1}\tilde{\mathcal{A}}\Theta_{\mathcal{A}, \epsilon} \times \epsilon$. Intuitively, given the assumed structure of $\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}$ implied by the “hybrid approximation” it becomes clear that the hybrid solution will correspond to the true solution of Equation (11) if expectations about the policy instrument more than n_e periods ahead are irrelevant to the decision problem of the agents in the underlying data generating structural model. This counterfactual approach might be a good approximation, for example, if agents in the underlying true model follow some form of exponential discounting so that the importance of expectations decays rapidly with the horizon.¹⁰ Note that, as with the least squares approximation, the approximation error vanishes as n_e approaches n_h , or in other words, as $n_{\nu, \text{hybrid}}$ approaches n_{ν} .

4.2 Counterfactual: What if OPEC stabilizes the oil price?

We employ the described approach to construct empirical policy rule counterfactuals to gauge the role of energy prices in monetary transmission in the euro area. In particular, we assume that OPEC, as stated in its statutes, aims to stabilize the global price of oil. As nicely summarized in Känzig (2021, p.6), “According to the statutes, OPEC’s mission is to stabilize global oil markets to secure an efficient, economic and regular supply of petroleum to consumers, a steady income to producers and a fair return on capital for those investing in the petroleum industry”. This implies that the counterfactual OPEC policy rule is such that it aims to stabilize the oil price at its steady-state level. The corresponding rule is $E_t[\hat{p}_{t+s}^{\text{oil}}] = 0 \forall t, s \geq 0$, which can be embedded in the system in Equation (11).¹¹ In fact, $\tilde{\mathcal{A}}$ simply becomes a se-

¹⁰Even for a full information rational expectations HANK model, the estimated counterfactual where the counterfactual policy rule is enforced only at point time t and in $n_e = 1$ period ahead expectations, this approximation already comes very close to the true counterfactual obtained from the underlying model (see McKay and Wolf (2022b)).

¹¹It 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

lection matrix that selects the entries in $\Theta_{\epsilon,A}$ and $\Theta_{\nu,A}$ corresponding to the oil price for all n_h horizons.

To implement the three different approaches to the solution of Equation (11) we identify OPEC-related oil supply news shocks using the proxy variables constructed by Känzig (2021). Specifically, Känzig (2021) argues that the high-frequency changes in oil price futures around OPEC meetings provide a valid instrument for OPEC oil supply news shocks. Because the accuracy of the approximations of Equation (11) depends on the number of policy news shocks identified, we depart from Känzig (2021), who uses only the first principal component of changes in the oil price futures at many horizons to identify a single oil supply news shock. Instead, we use high-frequency changes in the 1-month ($m_{t,1m}^{oil}$) and 12-month ($m_{t,12m}^{oil}$) futures to identify a short-term ($\nu_{t,short}^{oil}$) and a medium-term ($\nu_{t,medium}^{oil}$) oil supply news shock. For our baseline specification, we directly follow the procedure of MW and condition on the point estimate of the impulse responses for the separately estimated initial shock ϵ . In this application, this implies that we condition on the point-estimates in Figure 1, 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)). For reasons of consistency and robustness, when identifying the two oil supply news shocks, we nevertheless also use high-frequency surprises short-run interest rate futures to identify— and thereby account for the presence of— a euro area monetary policy shock. We report the resulting impulse response functions in Figure G.3 in Appendix G which shows that results for the effects of a euro area monetary policy shock are robust to identifying the monetary policy shock jointly with the oil supply news shocks.

In the notation of Equation (3a) and (3b), this implies that we extend the vector of proxies $\mathbf{m}_t = [m_{t,1m}^{oil'}, m_{t,12m}^{oil'}, m_t^{ir'}]$ and the vector of structural shocks of interest $\epsilon_t^* = [\nu_{t,short}^{oil'}, \nu_{t,medium}^{oil'}, \epsilon_t^{MP'}]$. To disentangle the three structural shocks identified using the three instruments we postulate

$$E[\mathbf{m}_t, \epsilon_t^{o'}] = \mathbf{0}, \quad E[\mathbf{m}_t, \epsilon_t^{*'}] = \begin{bmatrix} v_{1,1} & v_{1,2} & v_{1,3} \\ v_{2,1} & v_{2,2} & v_{2,3} \\ 0 & 0 & v_{3,3} \end{bmatrix}, \quad v_{1,1} > v_{1,2}, \quad v_{2,2} > v_{2,1} \quad (13)$$

Thus, we assume first that all three instruments are uncorrelated with the remaining structural shocks ϵ_t^o , second that the high-frequency changes in the interest rate m_t^{ir} are unaffected by OPEC oil supply news shocks, and third that the medium-term (short-term) oil supply news shock has a larger effect on the high-frequency change in 12 months (1 month) oil price futures. Furthermore, we require that a contractionary medium-term (short-term) oil supply news shock increases the oil price at horizon 12 (1).¹² The impulse responses to these two

that characterizes the systematic part of the corresponding equation (see Caldara and Kamps (2017) for a discussion).

¹²A natural additional assumption would be that the high-frequency changes in the oil price futures around

shocks are depicted in Figures G.1 and G.2 in Appendix G. In a nutshell, both shocks increase the price of oil, the energy component of the HICP, and lead to a contraction of output in the euro area. The response of the oil price to the short-term oil supply news shock is strong and immediate, while the medium-term oil supply shock tends to increase the oil price at longer horizons.

We then use the estimated impulse responses to these shocks to compute the three approximations to the solution of Equation (11) that characterize the true counterfactual, in which OPEC stabilizes the oil price. The results of this exercise are shown in Figure 2. In particular, the golden line in Figure 2 corresponds to the least squares approximation. In the counterfactual scenario, the reduced responsiveness of the oil price to an ECB monetary policy shock translates directly into a notably smaller response of local energy prices and thereby inflation and inflation expectations, especially at shorter horizons. This is the case even though the transmission to economic activity is little changed, which further supports the intuition that the reduction in the policy-induced decline in the CPI in the counterfactual is not due to changes in the prices of domestically produced goods. It is important to note that in this application the number of identified OPEC oil supply news shocks is smaller than the envisioned horizon $n_h = 24$ and therefore the oil price is not exactly zero all the time. In particular, it still falls after an ECB monetary tightening, although much less than in the baseline. This could lead us to underestimate the importance of oil prices in the monetary transmission.

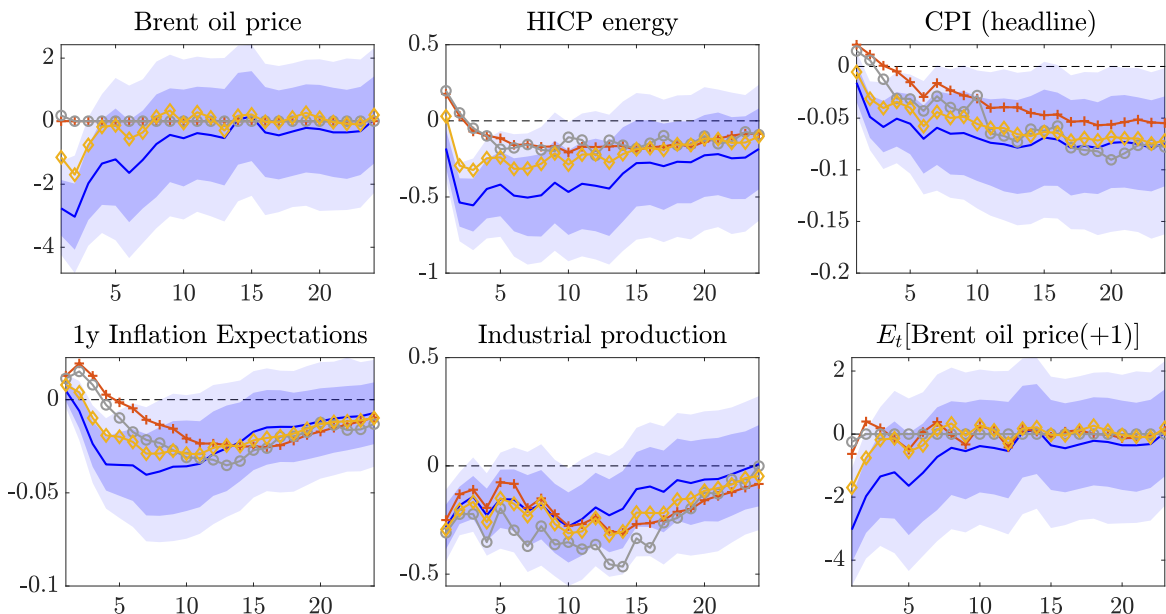
An alternative approximation is the Sims-Zha approximation, which is depicted in bronze. Recall that in this approximation the policy rule is perfectly enforced at the cost of assuming that agents in the underlying model are myopic, which implies that (i) the oil price is perfectly stabilized and (ii) the expectational component about the future path of the oil price that feeds into agents' decisions in the non-policy block is neglected.¹³ In other words, it is assumed that what matters for the behavior of agents in the non-OPEC block of the system is the spot price of oil (i.e., roughly the price at the gas station) and not their expectations about the future path of the oil price. As shown in Figure 2, the decline in the HICP is further reduced when using this approximation. Thus, the importance of global energy prices in the domestic transmission of euro area monetary policy is estimated to be larger under the Sims-Zha approximation than under its least squares counterpart.

This could be the case for two reasons. The first possible explanation is that, as shown in the upper left panel Figure 2, the spot price of oil price is perfectly stabilized under the Sims-Zha approximation which is not the case under the least squares approximation. Therefore,

OPEC announcements are unaffected by monetary policy shocks (i.e. $v_{1,3}, v_{2,3} = 0$). This would imply an overidentified system that cannot be handled by the algorithm of Arias et al. (2021). We therefore suggest this less restrictive set of assumptions. It turns out that the estimated coefficients for the two parameters are close to zero even without these assumptions.

¹³Consistent with the discussion in the previous section, we treat the short-term oil supply news shock as the contemporaneous policy shock $\nu_{0,t}$ necessary for the construction of the Sims-Zha approximation.

Figure 2: What if EA monetary policy shocks do not affect oil prices?



Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. Horizon in months. The bronze line shows the point-wise posterior means from the Sims-Zha approximation to the policy rule counterfactual, where OPEC stabilizes the oil prices. The golden line corresponds to the counterfactual obtained using the least squares approximation. The silver line corresponds to the hybrid approximation. We only retain draws from the posterior of impulse responses for which the resulting counterfactual response does not show explosive dynamics.

local energy prices and inflation fall less. The second possible explanation is related to oil price expectations, which are embedded in the least squares approximation but neglected in the Sims-Zha approximation. In particular, because the importance of oil prices in the monetary transmission is estimated to be lower in the least squares approximation, one could also infer that expectations about OPEC stabilizing the oil price not only at time t but also in the future play an important role in the overall effect.

The hybrid approximation shown in silver allows us to shed some light on this question. In this approximation, agents know that OPEC will stabilize the oil price not only in period t , but they can also expect the oil price to be stable for the upcoming period. Deviations between the bronze and silver lines in Figure 2 can therefore be attributed to the role of one period ahead oil price expectations. Looking at the panels in Figure 2, at least the one period ahead oil price expectations do not seem to play a big role. One way to rationalize this is to consider the notion that oil prices are perceived as following a random walk. That is, once the oil price is observed to be at a certain level, it is expected to remain at that level. Therefore, once the oil price has stabilized, it does not seem to matter for agents' decisions that policy

communicates that it will continue to stabilize in the next period since agents in the data already seem to expect the oil price to remain at that level.

In the bottom right panel of Figure 2, we illustrate this insight by plotting for each period t the model-consistent one-period-ahead oil price expectations for the different approximation techniques. Looking at the bronze line representing the solution under the Sims-Zha approximation, it is clear that once the oil price is stabilized ex-post, agents expect it to remain close to that level for the next period. Therefore, the announcement that OPEC will continue to stabilize the oil price in the upcoming period, which is embedded in the hybrid approach in silver, does not significantly alter agents' decision on their allocation at time t - once they observe that the oil price is indeed stabilized at time t . This explains why the Sims-Zha approximation and the hybrid approximation closely match.

For this reason, we tentatively attribute a larger share of the differences between the Sims-Zha approximation and the least squares approximation to the fact that the least squares approximation does not fully enforce OPEC's counterfactual policy rule for the spot price of oil. Regardless of the approximation technique and thereby the assumptions on the underlying model structure, we conclude that all approximations to the policy rule counterfactual agree that global energy prices, as measured by global oil prices, matter for the transmission of euro area monetary policy to inflation and inflation expectations. Furthermore, the least squares approximation and the Sims-Zha approximation appear to provide lower and upper bounds, respectively, on the importance of this channel.

Finally, it is important to note that, regardless of the approximation technique, this analysis should be seen as a lower bound on the importance of energy prices in the transmission of monetary policy in the euro area. Recall that the euro-zone HICP energy index is a weighted average of all local energy prices (in euro), not only the oil price. Therefore, movements in the oil price, despite being the dominant component of this basket, correspond to only a fraction of total euro area consumer energy prices. Thus, our fully empirical counterfactual analysis does not exactly correspond to a case where the ECB's decisions do not affect all components of the euro area energy basket, which would allow us to fully quantify the role of energy prices in monetary transmission. We expect that if the ECB's decisions would not affect local energy prices altogether, the short- and medium-term effects of a change in the ECB's policy stance on euro area inflation would be further muted.

5 Optimal monetary policy and energy prices

The previous sections revealed that the endogenous response of energy prices plays an important role in the transmission of exogenous monetary policy shocks. In this section, we go beyond exogenous monetary policy shocks and examine how energy prices matter for the optimal conduct of endogenous monetary policy in response to an exogenous increase in energy

prices.¹⁴ More specifically, we first derive the optimal response of the ECB to an increase in energy prices and then examine how the optimal policy strategy depends on the ECB’s ability to affect global energy prices. The analysis builds on recent developments in sufficient statistics approaches to optimal policy evaluation (Barnichon and Mesters (2023), McKay and Wolf (2023)), which allows us to address this question in an empirical framework without having to rely on a theoretical model.

5.1 Computing optimal policy counterfactuals

The approach of MW to estimating policy rule counterfactuals, discussed in Section 4.1, readily extends to computing impulse responses under the optimal policy. In particular, in line with Barnichon and Mesters (2023), MW define the optimal policy response as the response that implements an allocation that allows 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, we follow the definition of McKay and Wolf (2022a) and Barnichon and Mesters (2023) for the following reasons. First, the primary mandate “is arguably the relevant objective function for real-world central banks” (McKay and Wolf (2022a, p.3)), and second, it allows us to refrain from a precise parametrization of relative objectives in the loss function and therefore provides a simple benchmark against which we can evaluate the role played by energy prices in the mandate-optimal conduct of monetary policy.

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

$$\mathcal{L} = \frac{1}{2} \sum_{i=1}^{n_x} \lambda_i \mathbf{x}'_i \mathbf{W} \mathbf{x}'_i = \frac{1}{2} \mathbf{x}' (\Lambda \otimes \mathbf{W}) \mathbf{x} \quad (14)$$

where the \mathbf{x}_i contains the time path of the endogenous variable i , λ_i describes the policy weights attached to that variable with $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{n_x})$. The matrix \mathbf{W} summarizes the effects of time discounting in the policymaker’s preferences. It is potentially parameterized using a single discount factor β . The textbook solution to the optimal policy problem would amount to minimizing Equation (14) subject to the constraints embedded in Equation (5) in order to choose the optimal path of the policy instrument and thus the optimal allocation (see Galí (2015, Ch.5)).

MW show that the optimal policy problem can be analogously stated in impulse-response space. This implies minimizing the loss function subject to Equation (7) instead of Equation

¹⁴While the question of the role of the endogenous monetary policy response in the transmission of energy price shocks was raised as early as 1997 by Bernanke et al. (1997), it has not been rigorously addressed in an empirical model. For example, Bernanke et al. (1997) modestly argue that their VAR-based method and its application to this question constitute only “some modest [...] first steps toward sorting out the effects of systematic monetary policy on the the economy [...]” (Bernanke et al. (1997, p.92))

(5).¹⁵ In particular, the approach utilizes the observation that the implementable space of allocations for the endogenous variables \mathbf{x} and for the policy instrument \mathbf{z} is fully characterized by the impulse responses $\Theta_{\nu, \mathcal{A}}$ to the sequence of policy (news) shocks ν :

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{z} \end{bmatrix} = \Theta_{\nu, \mathcal{A}} \times \nu. \quad (15)$$

As shown in Appendix H.1, minimizing the loss function in Equation (14) subject to the constraint that the allocation must be in the implementable space described in Equation (15), yields an optimality condition that can be embedded in the policy block of Equation (6) by setting

$$\begin{aligned} \mathcal{A}_{\mathbf{x}}^* &= (\lambda_1 \Theta'_{x_1, \nu, \mathcal{A}} W, \lambda_2 \Theta'_{x_2, \nu, \mathcal{A}} W, \dots, \lambda_{n_x} \Theta'_{x_{n_x}, \nu, \mathcal{A}} W), \\ \mathcal{A}_{\mathbf{z}}^* &= \mathbf{0}, \end{aligned} \quad (16)$$

with $\Theta_{x_i, \nu, \mathcal{A}}$ as the matrix of impulse responses of variable i to all shocks in ν under the baseline policy rule characterized by \mathcal{A} . Although numerically equivalent to previous approaches in the literature (see, e.g., Svensson (1997)), the implied optimal policy rule is fully characterized by impulse responses to policy (news) shocks, all of which can theoretically be estimated from the data. Therefore, the optimal policy counterfactual is a special case of the approach to estimating empirical counterfactual impulse responses of Section 4.1, with the main difference being that the counterfactual policy rule is given by Equation (16).

5.2 Optimal euro area monetary policy response to an oil supply shock

The framework in Section 5.1 requires three key ingredients. First, we need to take a stance on the ECB's relevant loss function. Since we are primarily interested in how the response of energy prices shapes the optimal conduct of monetary policy, we follow McKay and Wolf (2022a) and Barnichon and Mesters (2023) by making the conservative choice of deriving the loss function from the primary mandate of the central bank. This not only allows us to largely refrain from taking a stance on a parametrization of the loss function but also yields a quantification of the output loss required to optimally fulfill the mandate in light of an oil price shock. In particular, the ECB's primary mandate is to maintain price stability, which it defines as an inflation target of 2% over the medium term.¹⁶ 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

¹⁵While MW provide a formal proof of this result and establish the equivalence of the resulting optimal policy rule to more classical results from the optimal policy literature as in Svensson (1997) and Giannoni and Woodford (2002), we provide a more heuristic explanation below and in more detail in Appendix H.1.

¹⁶In Figure H.5 of Appendix H.4 we explore an alternative loss function, where we give equal weights to deviations of inflation and GDP from steady state, which is arguably an extreme assumption for the ECB. The results remain qualitatively the same and are quantitatively similar.

is good evidence based on 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:

$$\mathcal{L} = \lambda_{\pi} \boldsymbol{\pi}' W \boldsymbol{\pi}, \quad (17)$$

with $\lambda_{\pi} = 1$. The weighting matrix is defined as $W = (\text{diag}(\beta^{24}, \dots, \beta^2, \beta, 1))$.¹⁷ Additionally, $\boldsymbol{\pi} = \mathcal{D} \mathbf{P}^{\text{HICP}}$ represents the transformed impulse responses of the (log) level of the HICP, denoted as \mathbf{P}^{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%.¹⁸

A second key component of this analysis is the choice and identification of the shock whose propagation we compare under the empirically identified baseline policy rule and the optimal policy rule. Given our focus on the interplay of monetary policy and energy prices, we choose to identify a contemporaneous oil supply shock, which we refer to interchangeably as “oil price shock” for short. In line with the previous section, we do so using high-frequency changes in 1-month oil price futures around the relevant OPEC meetings identified by Känzig (2021).

Third, we need to identify several euro area monetary policy shocks in order to approximate the optimal policy rule. Using a strategy similar to that used for the OPEC-related oil supply news shocks in section 4.2, we identify shocks to different dimensions of euro area monetary policy using high-frequency changes in interest rates at different maturities. In particular, we retain high-frequency changes in the 3-month OIS rate as our preferred proxy for conventional/contemporaneous monetary policy shocks, while we use changes in the 2-year OIS rate as a proxy for forward guidance shocks, which contain news about the future level of the policy instrument.¹⁹ Following the literature that uses the one-year yield as an indicator for the stance of short-run monetary policy, we add the 5-year Bund yield as a measure of the longer-run monetary policy stance into the baseline SVAR model.²⁰ We disentangle the

¹⁷This 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.

¹⁸We 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 15. 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.

¹⁹We use the terms ‘conventional’ and ‘contemporaneous’ monetary policy interchangeably.

²⁰At the start of the euro area in 1999, press releases announcing the Governing Council’s decisions were

three structural shocks of interest by making the following assumptions. First, we assume that all three proxies are uncorrelated with the remaining structural shocks. Second, we assume that the proxies for the conventional and forward guidance shocks are unaffected by the oil supply shock. Third, we assume that the conventional monetary policy shock has a stronger influence on high-frequency changes in the 3-month interest rate than the forward guidance shock. Finally, we assume that the forward guidance shock has a stronger effect on high-frequency changes in the 2-year interest rate futures than the conventional monetary policy shock.²¹ The impulse response functions to the identified shocks are shown in Figures H.1 and H.2 in Appendix H. They are in line with the standard results in the literature.

Given that the two identified policy shocks only constitute a subset of all policy shocks ($\boldsymbol{\nu}^{ident} \subset \boldsymbol{\nu}$), the set of hypothetical, feasible allocations under which the central bank operates is no longer described by Equation (15) but is rather given by

$$\mathbf{y} = \Theta_{\nu, \mathcal{A}} \times \boldsymbol{\nu}^{ident} \quad (18)$$

with $\mathbf{y} = (\mathbf{x}', \mathbf{z}')'$. In line with the intuition sketched above, the monetary policy authority selects the allocation of \mathbf{y} and the corresponding policy rule (see Equation (16)) that minimizes the loss function in Equation (14) - within that empirically identified space.

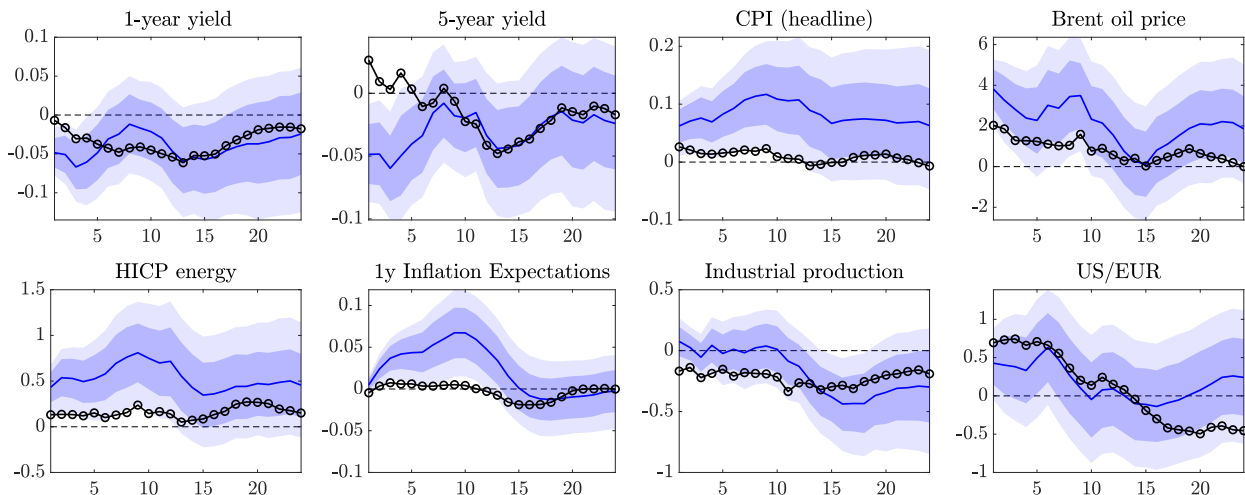
The impulse responses to the identified oil supply shock under the optimal policy response are shown in Figure 3 as the black dotted line. In this figure, we contrast the optimal response to an oil supply shock with the estimated empirical response (blue lines). The figure shows stark differences between these two responses indicating that the estimated empirical response did not optimally achieve the mandate. In particular, the estimated impulse response functions show that an oil supply shock leads to a large and long-lasting increase in the Brent oil price, and consequently to higher consumer energy prices, inflation, and inflation expectations. Crucially, this increase in energy costs and inflation rates is not a temporary phenomenon but will persist over the medium term. In addition, there is a delayed but significant economic contraction. Surprisingly, monetary policy does not seem to be counteracting the rise in inflation. However, this is consistent with the conventional wisdom that “in the past, central banks have typically looked through energy shocks” (Schnabel (2022)). Specifically, in response to oil supply shocks, the euro area’s monetary policy not only tolerates the implied increase in energy prices and inflation, but even slightly lowers interest rates, possibly as a measure to mitigate the economic downturn.²²

not regularly followed by a press conference. Since November 2001, however, each policy decision has been regularly accompanied by a press conference. This change in the ECB’s operational framework is an important consideration when identifying forward guidance shocks using high-frequency changes in OIS rates as a proxy, as Altavilla et al. (2019) finds that press releases do not contain forward guidance surprises, while press conferences do. Therefore, we estimate the model here using data starting in 2002.

²¹For a detailed description of the implementation and estimation we refer the reader to Appendix H.2.

²²In Figure H.5 of Appendix H.4 we substantiate this notion by (i) using a loss function that assigns equal weights to deviations in inflation and GDP and (ii) documenting that the observed response of interest rates is closer to the optimal response estimated from this balanced loss function rather than the optimal response

Figure 3: Impulse responses to an oil price shock under the empirical baseline and optimal monetary policy rule



Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. The black circled lines show the least squares approximation of the responses of the endogenous variables under optimal policy with a loss function described in Equation (17).

Contrary to the observed empirical response, the optimal response does not entail that the ECB should lower short- and longer-term interest rates. Instead, the optimal strategy to achieve its mandate involves raising interest rates to counteract the initial rise in energy prices. More specifically, monetary tightening is front-loaded and somewhat more pronounced at the longer end of the yield curve. It is important to emphasize that our estimates do not support an excessive rate hike by the ECB following an oil supply shock. Instead, the analysis points to a modest immediate tightening at the longer end of the yield curve, while a "look-through" strategy should initially be adopted at the shorter end of the yield curve. The cost of this type of monetary policy, relative to the estimated impulse response function, is a front-loaded contraction in output. However, not only does this strategy optimally stabilize medium-term inflation but the initial output contraction is offset by higher output in the medium term.

One possible explanation for the observation that only a small increase in interest rates is necessary to optimally stabilize medium-term inflation and even inflation expectations, is the quick response of energy prices to a monetary contraction. As documented in the 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, the ECB does not need to tighten excessively and persistently (and possibly induce a major recession) to fulfill its mandate in the face of energy price shocks, because a large part of

from the primary mandate loss function.

the adjustment is borne by relatively flexible global energy prices. In the next section, we examine this hypothesis in more detail.

5.3 Optimal policy when the ECB cannot affect global 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 the interest rates. To do so, we conduct a thought experiment and ask: What would be the optimal monetary policy response to the exact same increase in the oil price if the ECB's decision did not affect the Brent oil price? More precisely, we are interested in the optimal allocation \mathbf{y} in the case where the empirically identified, implementable space of possible allocations is not described by Equation (18), but rather by

$$\mathbf{y} = \Theta_{\nu, \tilde{\mathcal{A}}} \times \nu^{ident} \quad (19)$$

where the subscript $\tilde{\mathcal{A}}$ indicates that the space is now characterized by counterfactual impulse responses. The special feature of these counterfactual responses is that, as in Section 4.1, the identified monetary policy shocks ν^{ident} do not impact global energy prices.

We estimate the impulse response functions of this counterfactual in the following way.²³ We first compute the counterfactual impulse responses to monetary policy shocks by applying the methodology outlined in Section 4.1 to the identified impulse responses for the contemporaneous monetary policy and forward guidance shocks. We approximate the solution by using only the least squares approximation. After all, this approximation is a conservative choice as it implies the smallest effects of energy prices in the transmission of monetary policy in Section 4.2. In the next step, given the same loss function of Equation (17), the monetary authority then chooses the optimal allocation within the empirically identified space of feasible allocations described in Equation (19) as a response to the initial oil price shock. As shown by MW this amounts to setting the implied policy rule to

$$\mathcal{A}_x^* = (\lambda_\pi \Theta'_{\pi, \nu, \tilde{\mathcal{A}}} W), \quad (20)$$

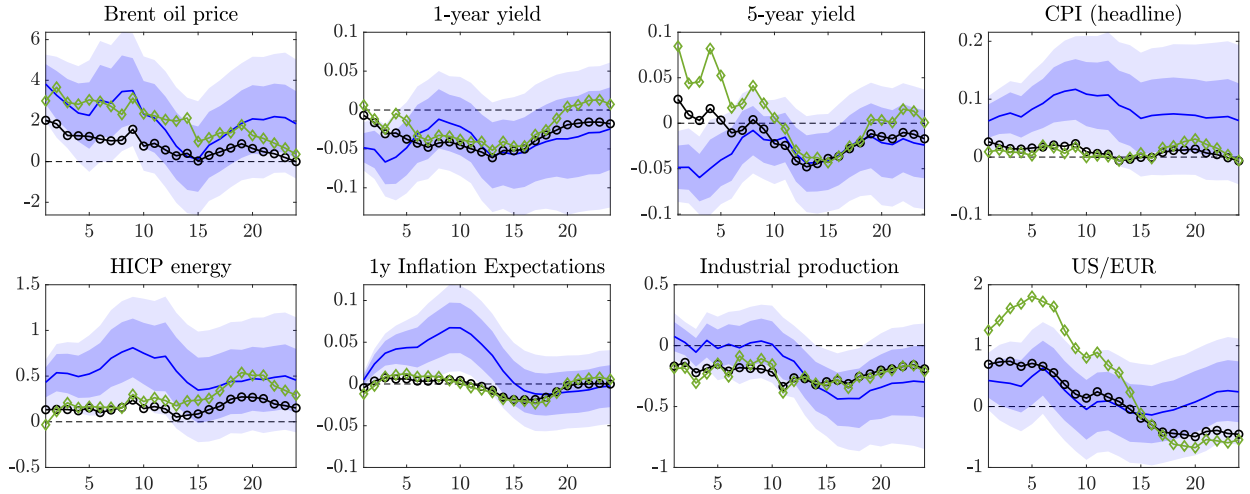
$$\mathcal{A}_z^* = \mathbf{0}, \quad (21)$$

which is now characterized by the counterfactual impulse responses of inflation to the identified monetary policy shocks $\Theta_{\pi, \nu, \tilde{\mathcal{A}}}$ instead of their estimated empirical counterparts $\Theta_{\pi, \nu, \mathcal{A}}$. The results from this exercise are depicted by the green lines in Figure 4.

First, note that in this application global energy prices are by construction unaffected by monetary policy, and thus the differences between the blue and green lines for the Brent oil price are very small and arise solely from the approximation error of the least squares

²³A detailed step-by-step summary of our approach to estimating the optimal monetary policy response to an oil price shock under the assumption that the ECB's decisions do not affect global oil prices can be found in Appendix H.3.

Figure 4: Impulse responses to an oil supply shock under optimal monetary policy when euro area monetary policy does not affect the Brent oil price



Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. The black circled lines show the least squares approximation of the responses of the endogenous variables under optimal policy with a loss function described in Equation (17). The green triangled lines show the least squares approximation of the responses of the endogenous variables under to the optimal problem computed using the impulse responses to a conventional monetary policy and forward guidance shock, which, counterfactually, do not impact the Brent oil price due to a change in the OPEC policy rule.

approximation. Second, when comparing the optimal policy given the estimated impulse response functions (black) and the optimal policy in the counterfactual scenario (green), it becomes apparent that monetary policy needs to tighten much more in order to achieve its mandate when it does not affect global energy prices. This is especially true for the 5-year yield.

As Figure 4 shows, the stronger increase in the 5-year rates allows the ECB to implement a very similar allocation as in the optimal response to the estimated impulse response functions. In fact, under the optimal counterfactual policy, inflation and inflation expectations are stabilized as well. Furthermore, our estimates show that while monetary policy needs to tighten significantly more, it does not need to engineer a significantly deeper contraction in output to bring inflation back to target. Again, this seemingly puzzling observation can be attributed to the role of energy prices. First, note that in addition to the 5-year yield, the other striking difference between the black and green lines is the exchange rate. This implies, that if the ECB cannot affect global energy prices, the optimal strategy entails a strong appreciation of the euro against the dollar. Since energy goods in the euro area are priced in euros, the strong appreciation of the euro dampens the response of local energy prices despite the inability of the ECB to fight the surge in global energy prices. In this way, the ECB

succeeds in stabilizing HICP energy inflation, thereby mitigating the rise in inflation expectations, which are shown to be particularly sensitive to changes in energy prices (Aastveit et al. (2023), Wehrhöfer (2023)). Thus, by following a strategy that relies on stabilizing local energy prices, the ECB is able to contain the inflationary pressures arising from the oil price shock. This analysis reveals that, if monetary policy cannot directly influence global energy prices (in dollars), it would find it optimal to rely on its local counterpart (in euro) to achieve its objective.

Interestingly, while monetary policy has to tighten more at the longer end of the yield curve in order to appreciate the exchange rate, this does not imply a much larger output contraction. We document in Appendix I that this outcome can be traced back to the multidimensionality of monetary policy. More specifically, we show that when monetary policy cannot directly influence global energy prices, it would be optimal to rely more heavily on forward guidance rather than contemporaneous interest rate policies to stabilize medium-term inflation. Intuitively, under the assumption that monetary policy does not affect global energy prices, forward guidance is estimated to be comparatively better suited to tackle deviations of inflation from its target than it is the case when the ECB’s decisions do affect global energy prices (see Figure I.2). Therefore this shift in strategy is fully consistent with the result of MW, who show that under the optimal monetary policy, “the policymaker will rely most heavily on the tools [...] that are best suited to offset the perturbation to its targets” (McKay and Wolf (2022a, p.9)). As forward guidance is estimated to be less recessionary, this implies that the counterfactual outcome for output is similar despite higher rates.

Therefore, we conjecture that only if the ECB’s decisions do not move energy prices at neither the global nor the local level it would be necessary to engineer a large recession to stabilize prices in the face of an energy price shock. As this counterfactual would be very difficult to formulate as an empirical policy rule counterfactual, proving this conjecture requires a fully structural model that models the interplay between monetary policy and energy prices and matches their empirically identified transmission. We leave the estimation of this counterfactual to future research, as such a model is beyond the scope of this paper.

6 Application to the most recent energy price surge

In this section, we assess the role of euro area monetary policy during the post-Covid energy-driven inflation episode considering our earlier findings on the importance of energy prices for the conduct of monetary policy. In particular, we first compute the historical decomposition of our BPSVAR model to estimate how much of the (energy price) inflation and the subsequent disinflation can be explained by the cumulative effects of exogenous monetary policy shocks. Second, we estimate how the euro area economy would have evolved during this inflationary episode if the ECB had responded optimally to the energy price shock. To this end, we apply the methods in Caravello et al. (2023), which shows how to extend the MW impulse response

counterfactuals to historical episodes.

6.1 The role of exogenous monetary policy shocks

We examine the role of monetary policy shocks in the recent rise in inflation and energy prices by first estimating the model in Section 3.2 using data through October 2023.²⁴ To explicitly model the extraordinary volatility in the data induced by the COVID-19 pandemic we incorporate the “Pandemic Prior” approach of Cascaldi-Garcia (2022) into the BPSVAR model.²⁵ We then compute the historical decomposition of the extended BPSVAR model to estimate the contributions of exogenous monetary policy shocks to the rise and fall of (energy price) inflation. To that end, we write the SVAR in structural vector moving average (SVMA) representation, which under the assumption of invertibility, allows us to decompose the state of the economy into the effects of exogenous variables, the initial conditions, the cumulative effects of the (current and past) policy shocks, and the effects of the other structural shocks.

Figure 5 presents the results for this exercise. In particular, the black solid lines depict the evolution of headline inflation and energy price inflation alongside the contribution of exogenous ECB monetary policy shocks, which are depicted by the dark blue and red bars.²⁶ Furthermore, the “Pandemic priors” approach of Cascaldi-Garcia (2022) allows us to estimate the effects of the onset and fading out of the pandemic as depicted by the yellow bars. During the year 2020 and early 2021, the fall in energy prices and HICP inflation can largely be attributed to the pandemic.

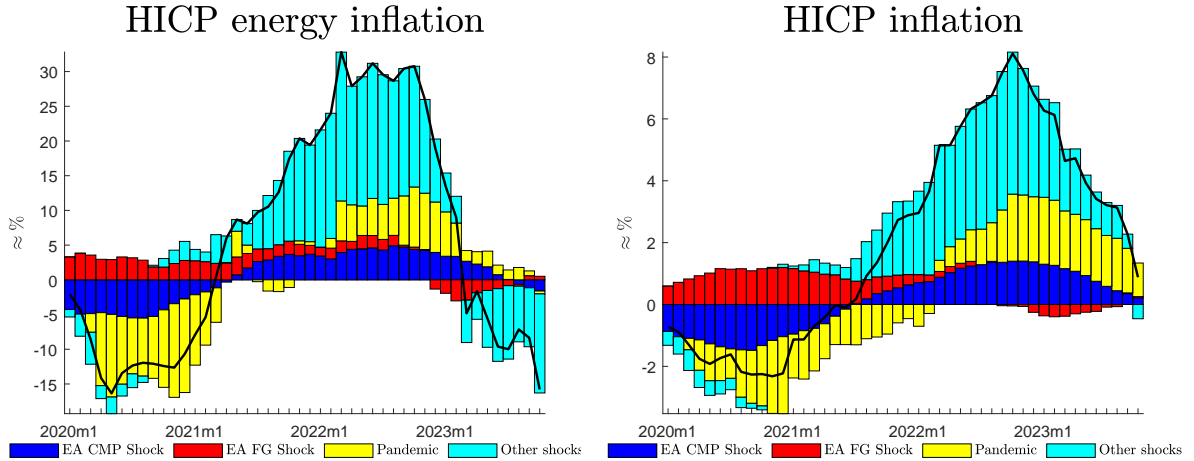
However, when inflation began to rise rapidly, monetary policy is estimated to have contributed to this increase as the effects of conventional and forward guidance policy shocks, the sum of the red and blue bars, quickly accelerate significantly into positive territory in 2021 and 2022. At their peak, the monetary policy shocks are estimated to have contributed, possibly unintentionally, to an increase in HICP inflation of around 1.5 percentage points (pp) and in energy price inflation of around 5 pp. The contribution of these expansionary shocks in a context of rapidly rising inflation then slowly fades away, with forward guidance policy tightening first, in line with ECB communication. In sum, through the lens of our model, monetary policy shocks have played a significant role in the recent rise and fall of energy- and consumer price inflation. However, this is not the full picture, as the role of endogenous monetary policy may be very different, which we turn to next.

²⁴As shown in Appendix J, our results are robust to extending the sample to include the pandemic period.

²⁵Furthermore, to preserve the stationarity of our specification we replace all price levels with year-on-year inflation rates. We also include the 5-year bund yield as we aim to identify forward guidance shocks.

²⁶For the sake of brevity we plot the evolution of these variables net of the contribution of the constant and initial conditions.

Figure 5: Exogenous monetary Policy shocks as drivers of the recent energy price inflation surge



Notes: Historical decomposition excluding the contribution of the constant and initial conditions (black line) alongside the contribution of current and past EA conventional monetary policy shocks (CMP shock, blue bars), forward guidance shocks (FG shock, red bars), effects of the 2020 COVID pandemic (yellow bars) and the *sum of the contributions* of all current and past other (non-EA MP) shocks in the system (turquoise bars). We obtained the estimates by computing the historical decomposition using the SVMA representation for each draw. We plot the resulting point-wise means for each shock and variable.

6.2 The role of endogenous monetary policy in the recent inflation surge

In this subsection, we ask how the euro area economy would have evolved in the face of the energy price surge if the endogenous component of monetary policy — the policy rule — had been optimally set to achieve the primary mandate. Traditionally, this has been addressed by, first, setting up a theoretical (DSGE) model of the economy and specifying the full menu of stochastic shocks that drive that economy; second, estimating this model using full-information methods, such as those described in Herbst and Schorfheide (2016), to back out the size and sign of those structural shocks, and third, changing the policy block and simulating the counterfactual economy using the counterfactual SVMA representation alongside the structural shocks from the second step (see Christiano et al. (2015) for an example). Instead, we adopt the approach of Caravello et al. (2023) (hereafter CMW), who extend the purely empirical approach of MW to compute impulse response counterfactuals to historical episodes.

The key insight of CMW is that, combining the assumption of invertibility of the SVAR model with the approach of MW allows the researcher to compute policy rule counterfactuals that go beyond just impulse responses.²⁷ First, they observe that the counterfactual

²⁷The assumption of invertibility is embedded in almost all prominent SVAR identification schemes (see Plagborg-Møller and Wolf (2021) for a discussion). However, in a simulation study using the Smets and Wouters (2007) model, CMW show that their approach approximates the true counterfactual very well, even

evolution of the economy can be retrieved from a counterfactual SVMA representation of the economy, where the counterfactual impulse responses to all structural shocks are computed under the counterfactual policy rule. Second, they prove that identifying all the true structural shocks and the corresponding impulse responses is not a necessary condition for estimating the counterfactual SVMA representation, as the reduced-form representation of the SVAR model proves sufficient. Third, they show how to coherently manipulate the initial conditions component of the counterfactual SVMA representation to ensure that agents' expectations in the underlying data-generating structural model incorporate the counterfactual policy rule. As in the case of the counterfactual impulse response, this method allows us to answer the question at hand, without having to specify the structural model that has generated the data.

More formally, we aim to simulate the evolution of the euro area economy under the assumption that, starting at time t^* , the ECB would have conducted the policy according to the optimal rule embedded in \mathcal{A}_x^* and \mathcal{A}_z^* in Equation (16). Thus, we are interested in estimating the counterfactual SVMA representation of the economy given by

$$\mathbf{y}_t = \underbrace{\sum_{\ell=0}^{t-t^*} \Theta_{\ell, \eta, \mathcal{A}^*} \eta_{t-\ell}}_{\text{contrib. new shocks after } t^*} + \underbrace{\mathbf{y}_{t^*}^{\mathcal{A}^*}}_{\text{contrib. initial condition at } t^*} \quad (22)$$

For each time period t , the first term in Equation (22) captures the effects of shocks $\eta_{t-\ell}$ that occur between the policy change at t^* and the period t . The propagation of these shocks is altered from $\Theta_{\ell, \eta, \mathcal{A}}$ to $\Theta_{\ell, \eta, \mathcal{A}^*}$ because euro area monetary policy no longer follows the estimated policy rule embedded in \mathcal{A} but the optimal rule that minimizes the loss function in accordance with the mandate. The second term in Equation (22) represents the contribution of the initial conditions starting at point t^* , assuming that the ECB follows the mandate-optimal policy rule. By definition, this embeds the joint effects of all the shocks that occurred before the change in the policy rule and can be thought of as the expected path of the economy given the history of shocks up to t^* .²⁸

To estimate the first term in Equation (22), we exploit the result of CMW, who show that under invertibility it suffices to apply the MW procedure sketched in Section 4.1 to each of the reduced form impulse responses corresponding to the reduced form errors of the SVAR, u_t . Their key insight is that under invertibility there is a one-to-one mapping between the reduced form residuals and the true structural shocks. Therefore, the sum of the (counterfactual) impulse responses to these true shocks must equal the sum of the (counterfactual) impulse responses to the reduced form innovations.

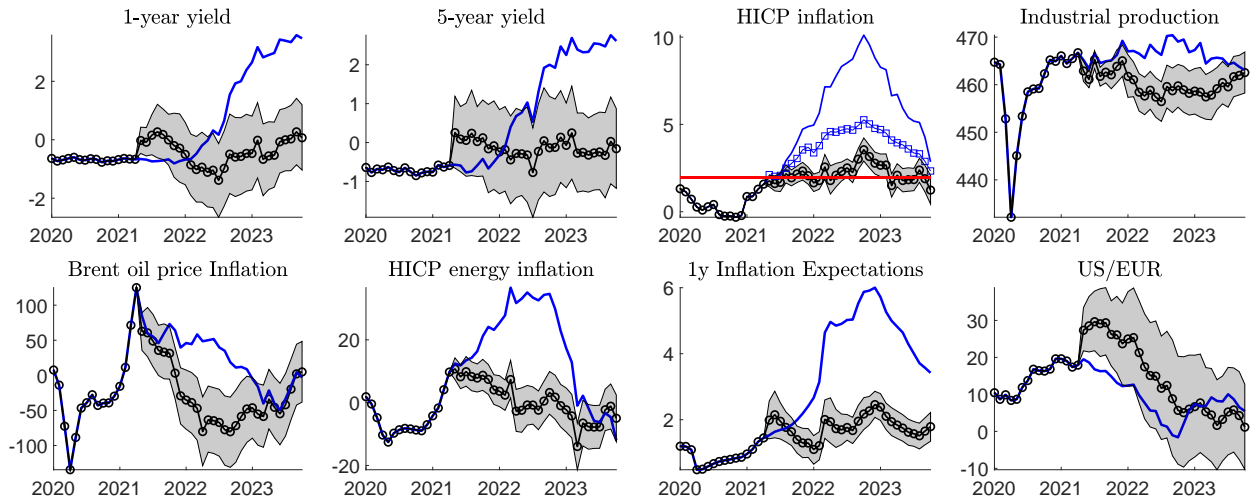
when the invertibility conditions of the SVAR model do not hold.

²⁸We leave aside the treatment of the exogenous variables because in our application the initial effects of all exogenous variables arising from the ‘‘Pandemic Priors’’ of Cascaldi-Garcia (2022) approach materialize before our choice of t^* . As such, their effect on all variables thereafter is embedded in the initial condition term $\mathbf{y}_{t^*}^{\mathcal{A}^*}$.

The second term in Equation (22) ensures that the chosen policy path for $t \geq t^*$ is such that the new policy rule is also embedded in the expectations of the agents in the economy. Under invertibility, we can recover the baseline version of these expectations $\mathbf{y}_{t^*}^A$ (i.e. the forecast formed under the baseline policy rule described by \mathcal{A}) by computing the unconditional BPSVAR forecast starting from \mathbf{y}_{t^*} . Along with the identified policy shocks this is sufficient to estimate $\mathbf{y}_{t^*}^{A^*}$, which is the expected value of the variable \mathbf{y}_t formed under the optimal policy rule given the information up to t^* . Intuitively, since impulse responses are a special form of a conditional forecast, the same proof and procedure that MW use to derive the counterfactual impulse responses also applies to constructing unconditional forecasts.

We compute the counterfactual evolution of the euro area economy under the assumption that the ECB conducts optimal monetary policy from $t^* = \text{April 2021}$.²⁹ Figure 6 plots the counterfactual evolution of the economy under the optimal policy rule (black dotted lines) alongside the actual evolution (blue solid lines). In line with Section 5.1, optimal monetary policy is defined as the achievement of the primary mandate of the ECB, i.e. medium-term inflation stability. Three interesting conclusions can be drawn from this analysis.

Figure 6: Counterfactual evolution of the endogenous variables under optimal ECB policy



Notes: The figure plots the time series of the endogenous variables (blue solid line) alongside the estimated counterfactual evolution (black dotted line) of these variables under the assumption that, from April 2021 onwards, the ECB would have conducted optimal monetary policy as described in Equation 17 and thereby minimizes the weighted squared deviations of “medium-term” inflation from the target. The red line marks the ECB’s 2% medium-term inflation target. Black shaded areas correspond to 68% point-wise credible sets.

First, the path of interest rates implied by the mandate-optimal policy rule would have

²⁹We choose April 2021 as the starting point for this exercise because it marks the period just before euro area headline inflation exceeded the ECB’s 2% target, as the euro area economy experienced a 120% year-on-year increase in the price of oil. It is also consistent with the application of CMW to the post-covid inflation period in the US.

been very different. In particular, the ECB would have had to implement an initial forceful tightening using both conventional and forward guidance policies, as shown by the counterfactual paths of the one-year and five-year interest rates. Overall, the policy response would have been faster and more front-loaded but would have avoided the interest rate levels reached in 2023. This is particularly interesting as one of the main criticisms of the ECB in this episode has arguably been the delayed response to use forward guidance and signal future rate hikes in the face of accelerating inflation.

Second, the strong increase in inflation could have been contained by the rise in interest rates. This is because energy prices and inflation expectations would have risen much less under a mandate-optimal policy. In particular, the change in the ECB’s policy stance would have mitigated some of the escalation in global oil prices and, at the same time, would have appreciated the euro, further dampening energy price inflation in the euro area.

Third, the change in the policy stance would have led to an initially sharper contraction in output, as measured by industrial production, which is estimated to quickly return to current actual levels by the end of 2023. A back-of-the-envelope calculation, taking into account the relative volatility of GDP and industrial production, suggests that, given the trough response of industrial production relative to the baseline level, the associated trough effect for GDP is estimated to be around 3.3% (see Georgiadis et al. (2024) for a similar conversion and a justification for this approximation).

The fact that it would have been optimal for the ECB to commit to a rapid tightening of monetary policy as early as 2021 deserves a closer look. To this end, we have calculated the path of HICP inflation that an agent would have expected in April 2021, given the observations up to that date, their knowledge of the underlying model, and the estimated policy rate described by \mathcal{A} . It is shown by the blue squared line in the third panel of the first row.³⁰ It becomes apparent that, even in the absence of further structural shocks, the inflation rate would have been expected to peak at around 5% in the medium term given the conditions in April 2021. Our analysis in Section 6.1 reveals that up to 3% of this increase in inflation can be explained by the fading effects of the pandemic, which initially lowered inflation substantially in 2020 but then caused it to overshoot. Through the lens of the model, these effects were predictable given the data through April 2021. To align expectations with the ECB’s commitment under optimal policy to stabilize medium-term inflation at around 2%, the ECB must commit to a substantial, early rate hike. In Appendix K we quantify this argument by showing that up to 50% of the difference between the actual and the counterfactual evolution of the economy can be attributed to the realignment of expectations. We also show that the remaining part can be attributed to the fact that the change in the policy rule also alters the propagation of incoming new shocks, making them less inflationary to begin with (see Figure K.2).

³⁰In order to improve the readability of the figure the expected paths of the other variables are shown in figure K.1 in Appendix K.

Ultimately our analysis shows that the ECB could have stabilized inflation in the face of the energy price spike. This would have come at the cost of approximately 3% lower output relative to baseline in the short run, which we estimate would then quickly return to its current actual level. Our previous analysis reveals that the ability of the ECB to directly affect global energy prices crucially shapes how monetary policy transmits to the economy and stabilizes inflation in the face of energy price shocks. If the ECB’s decisions do not translate to changes in global energy prices, it would have to tighten even more and possibly engineer an even deeper recession in order to optimally counteract a surge in inflation. We conclude this section with a word of caution. We have focused only on the primary mandate of the ECB in defining the loss function of the economy. There could, of course, be other loss functions that would alter the trade-off between price stability on the one hand and output and employment losses on the other. We leave the appropriate weighting of employment and output in the loss function to future research.

7 Conclusion

This paper challenges the prevalent view that the ECB has limited capacity to combat energy-price-driven inflation. Through a high-frequency event study and a Bayesian Proxy SVAR model, we document that energy prices are significantly affected by monetary policy and that they are more responsive than the headline consumer price index. Using Lucas critique-robust counterfactuals, as proposed by McKay and Wolf (2023), we establish that a substantial part of the impact of monetary policy on headline consumer prices is transmitted through the energy price component of the consumer price index. Moreover, our results suggest that the ECB’s monetary policy decisions affect inflation expectations primarily through their impact on energy prices. Consequently, the rapid and large response of energy prices to monetary policy can help the ECB to achieve its mandate.

We furthermore empirically analyze how monetary policy should optimally respond to an energy price shock, which we proxy by an oil supply shock. We show that, historically, the ECB has “looked through” the persistent price increases caused by this shock and has even reacted somewhat accommodatively to stabilize output. However, we show that the ECB should react more restrictively to an energy price shock in order to achieve its mandate of medium-term price stability. In particular, a small frontloaded tightening stabilizes inflation and its expectations. We document that it is precisely the ECB’s ability to influence fast-moving energy prices that allow it to successfully combat energy price spikes without excessive tightening.

Consequently, our analysis of the recent energy price-driven inflationary episode from 2021 to 2023 reveals that inflation could have been stabilized at the target through a front-loaded monetary policy tightening. We find that this policy strategy would have prevented the 8% inflation overshoot while inducing an additional 3% contraction in output, which is only

short-lived as energy prices adjust swiftly to the change in the policy stance.

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*.
- Arce, Óscar, Matteo Ciccarelli, Antoine Kornprobst, and Carlos Montes-Galdón (2024). “What caused the euro area post-pandemic inflation?” *ECB Occasional Paper* (2024/343).
- 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 Monneray, Matthew Rognlie, and Ludwig Straub (2023). “Managing an Energy Shock: Fiscal and Monetary Policy”. Tech. rep. Working Paper.
- 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*.
- 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.
- 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 (forthcoming). “Hicks in hank: Fiscal responses to an energy shock”. *Journal of Political Economy: Macroeconomics*.

- 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.
- 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.
- 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).
- Ca’Zorzi, Michele, Luca Dedola, Georgios Georgiadis, Marek Jarocinski, Livio Stracca, and Georg Strasser (2023). “Making Waves: Monetary Policy and Its Asymmetric Transmission in a Globalized World”. *International Journal of Central Banking* 19 (2), 95–144.
- 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.
- Christiano, Lawrence J, Martin S Eichenbaum, and Mathias Trabandt (2015). “Understanding the great recession”. *American Economic Journal: Macroeconomics* 7 (1), 110–167.
- Christoffel, Kai Philipp, Günter Coenen, and Anders Warne (2008). “The new area-wide model of the euro area: a micro-founded open-economy model for forecasting and policy analysis”. *ECB working paper No. 944*.
- 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.
- Galí, Jordi (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

- 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.
- Giannoni, Marc P and Michael Woodford (2002). “Optimal interest-rate rules”. *NBER Working paper* 9419.
- 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.
- Herbst, Edward P and Frank Schorfheide (2016). *Bayesian estimation of DSGE models*. Princeton University Press.
- Ider, Gökhan, Alexander Kriwoluzky, Frederik Kurcz, and Ben Schumann (2023). “The energy-price channel of (European) monetary policy”. *DIW Discussion Papers* 2033.
- 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).
- Jarociński, Marek and Peter Karadi (2020). “Deconstructing monetary policy surprises—the role of information shocks”. *American Economic Journal: Macroeconomics* 12 (2), 1–43.
- Känzig, Diego R (2021). “The macroeconomic effects of oil supply news: Evidence from OPEC announcements”. *American Economic Review* 111 (4), 1092–1125.
- Kilian, Lutz and Logan T Lewis (2011). “Does the Fed respond to oil price shocks?” *The Economic Journal* 121 (555), 1047–1072.
- Lagarde, Christin (2022a). *Committee on Economic and Monetary Affairs Monetary Dialogue with Christine Lagarde, President of the European Central Bank*. 7 February, Brussels, https://www.europarl.europa.eu/cmsdata/245081/1249437-2_EN.pdf.
- (2022b). *Transcript of the ECB Press conference 3 February*. 3 February, Frankfurt am Main, <https://www.ecb.europa.eu/press/pressconf/2022/html/ecb.is220203~ca7001dec0.en.html#qa>.
- 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.
- McKay, Alisdair and Christian Wolf (2022a). “Optimal policy rules in HANK”. *Working Paper, FRB Minneapolis*.
- (2022b). “What Can Time-Series Regressions Tell Us About Policy Counterfactuals?” *NBER Working paper* 30358.

- McKay, Alisdair and Christian Wolf (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 H elene 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.
- Miranda-Pinto, Jorge, Andreas Pescatori, Ervin Prifti, and Guillermo Verduzco-Bustos (2023). “Monetary Policy Transmission through Commodity Prices”. *IMF Working Paper*, 23/215.
- 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-M oller, Mikkel and Christian K Wolf (2021). “Local projections and VARs estimate the same impulse responses”. *Econometrica* 89 (2), 955–980.
- Rosa, Carlo (2014). “The high-frequency response of energy prices to US monetary policy: Understanding the empirical evidence”. *Energy Economics* 45, 295–303.
- 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.
- Schumann, Ben (2024). “Crude Controversy - European monetary policy and energy prices”. *Four Essays in Macroeconomics*. FU Berlin. PhD thesis. Berlin: Freie Universit at Berlin. Chap. 1, 1–42.
- Sims, Christopher A and Tao Zha (2006). “Does monetary policy generate recessions?” *Macroeconomic Dynamics* 10 (2), 231–272.
- Smets, Frank and Rafael Wouters (2007). “Shocks and frictions in US business cycles: A Bayesian DSGE approach”. *American economic review* 97 (3), 586–606.
- Svensson, Lars EO (1997). “Inflation forecast targeting: Implementing and monitoring inflation targets”. *European economic review* 41 (6), 1111–1146.
- Swanson, Eric T and John C Williams (2014). “Measuring the effect of the zero lower bound on medium-and longer-term interest rates”. *American economic review* 104 (10), 3154–85.

- 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.

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 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 (1)	3-month-ahead federal funds future rate changes around FOMC announcements	Computed based on methodology of Jarociński and Karadi (2020)	Gürkaynak et al. (2005)
Fed monetary policy surprise (2)	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
US Dollar currency index (DXY)	US Dollar currency index price changes around ECB monetary policy announcements	Computed from tick data	Refinitiv Tick History database
EUR-USD exchange rate	EUR-USD changes around ECB monetary policy announcements		EA-MPD from Altavilla et al. (2019)
EURO STOXX 50 (STOXX50)	STOXX50 price changes around ECB monetary policy announcements		EA-MPD from Altavilla et al. (2019)
Dutch TTF natural gas price	Daily (closing) price changes of 1-month and 1-year Dutch TTF futures around ECB monetary policy announcements		Bloomberg

Note that, 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 in sections 3 and 4. In Section 5 we no longer do so because adding the 5-year yield to the baseline specification solves the issue.

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

Variable	Description	Notes	Source
1-year yield	Germany Government 1 year yield	End of period	Macrobond Financial AB
2-year yield	Germany Government 2 year yield	End of period	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 monetary policy proxy	3 month (monetary event window) OIS surprise	Calculated based on data and methodology by Jarociński and Karadi (2020)	Jarociński and Karadi (2020) and authors' calculations
Global oil production	Global oil production (million barrels/day)		Baumeister and Hamilton (2019)
Oil inventories	Change in global oil inventories		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	Surprise in oil futures prices around OPEC announcements	Monthly sum of daily values	Känzig (2021)

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)

	EA	EA	US	US	UK	UK
	(1)	(2)	mps_{FF4}^{pm}	mps^{\perp}	(1)	(2)
$\hat{\beta}^{100\text{bps}}$	-3.20**	-3.34**	-2.15**	-2.23***	0.37	0.36
	(1.31)	(1.54)	(1.03)	(0.83)	(0.67)	(0.68)
R^2 (%)	3.75	3.37	2.61	3.21	0.33	0.38
N	211	195	182	187	257	246
Sample	2002:1	2002:1	1996:1	1996:1	1997:6	1997:6
	2021:12	2019:12	2019:6	2019:12	2021:3	2019:12

Note: Coefficient estimates $\hat{\beta}^{100\text{bps}}$ from the Brent crude oil price event study regression equation $p_t = \alpha + \beta mps_t + \epsilon_t$, where t indexes monetary policy announcements. 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.

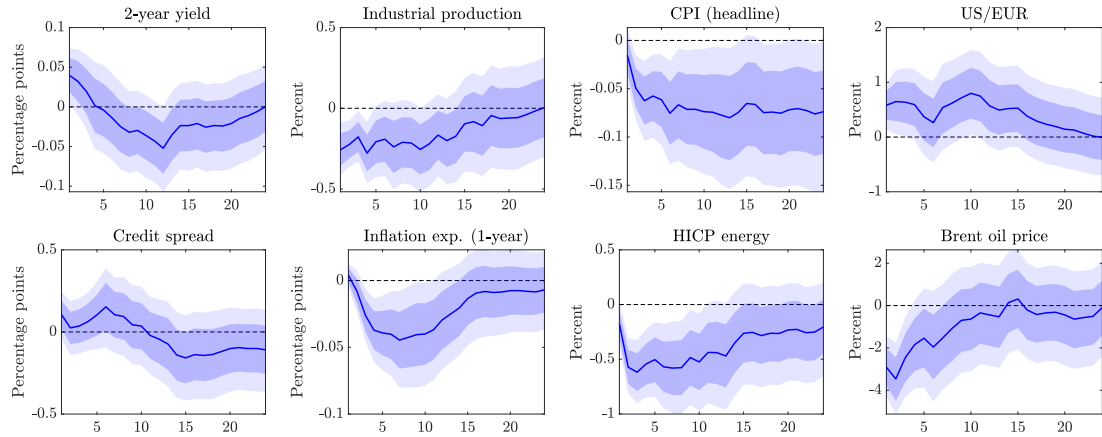
Table B.2: Coefficient estimates $\hat{\beta}^{100\text{bps}}$ from the natural gas price (Dutch TTF) event study regressions.

	1-month TTF	1-year TTF	1-month TTF	1-year TTF
$\hat{\beta}^{100\text{bps}}$	-17.42*** (4.50)	-12.32*** (3.12)	-13.85*** (3.92)	-13.41*** (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. 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.

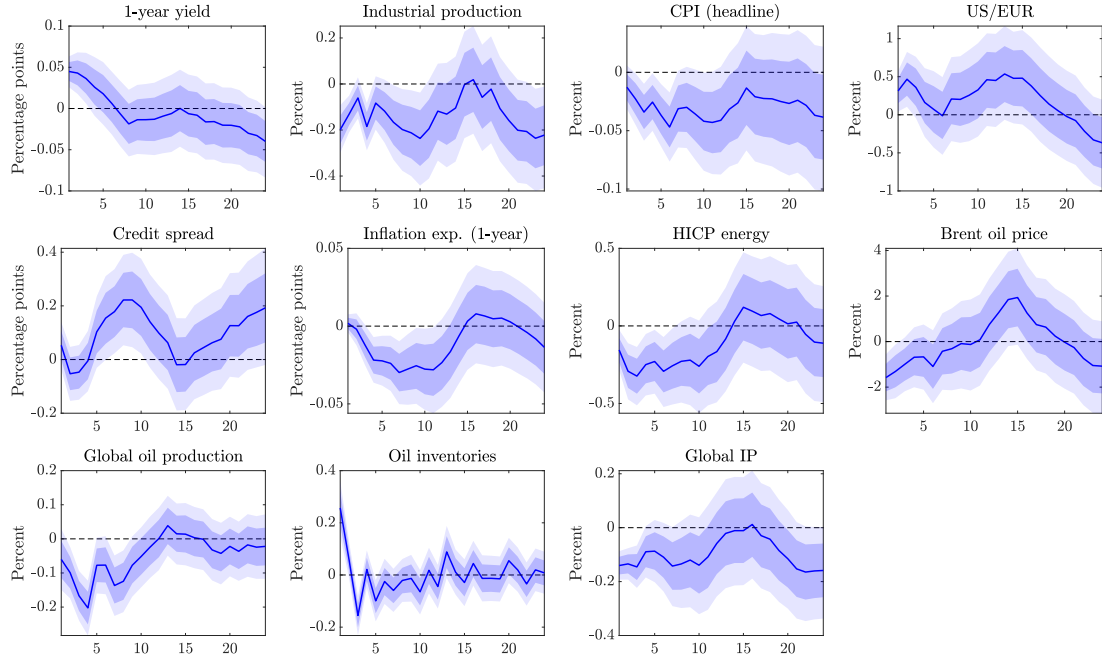
C BPSVAR robustness results

Figure C.1: Euro Area SVAR model, 2-year yield



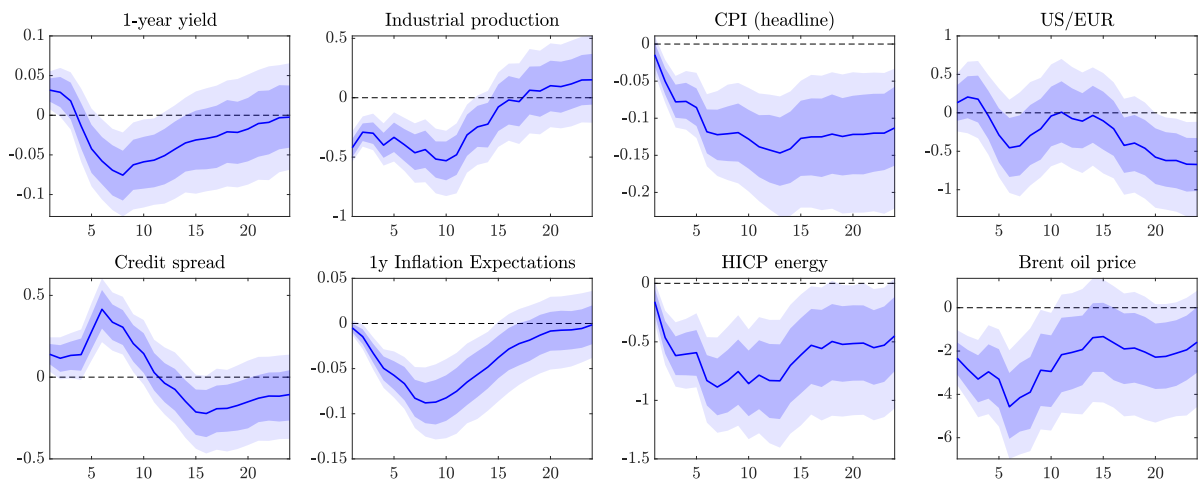
Notes: Euro Area model, 2-year Bund yield instead of 1-year Bund yield. Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Horizon in months.

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



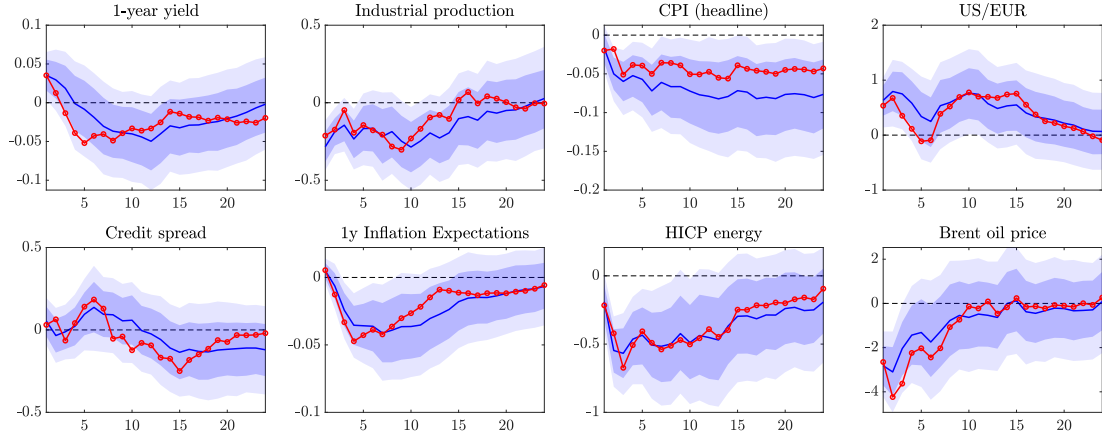
Notes: Euro Area model with world oil production, inventories, and world industrial production. Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Horizon in months.

Figure C.3: Baseline Euro Area SVAR model starting the estimation in 2002



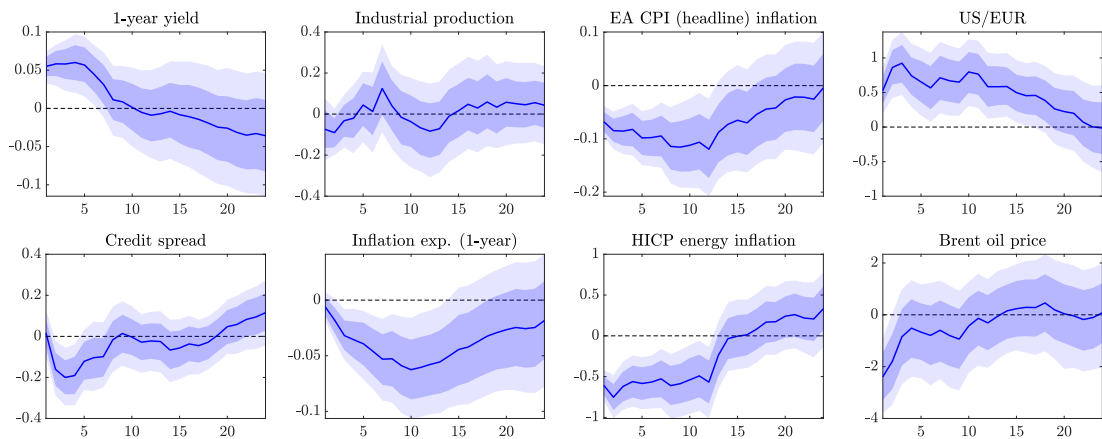
Notes: Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months.

Figure C.4: Euro Area SVAR model, external vs. internal instrument identification



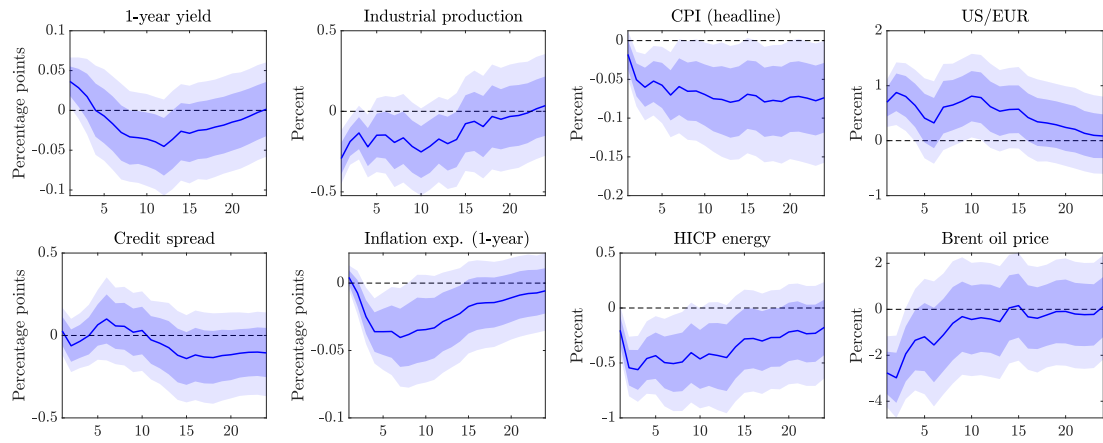
Notes: Posterior means of Impulse response functions to a one standard deviation monetary policy shock using the external-instrument BPSVAR identification alongside in blue 68% and 90% point-wise probability bands. Posterior mean impulse response functions to monetary policy shock identified using the 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 that resembles the prior distribution embedded in the BPSVAR approach of Arias et al. (2021). The impulse responses for the internal instrument approach are scaled such that they induce the same impact effect for the 1-year yield.

Figure C.5: Euro Area SVAR model including the Pandemic



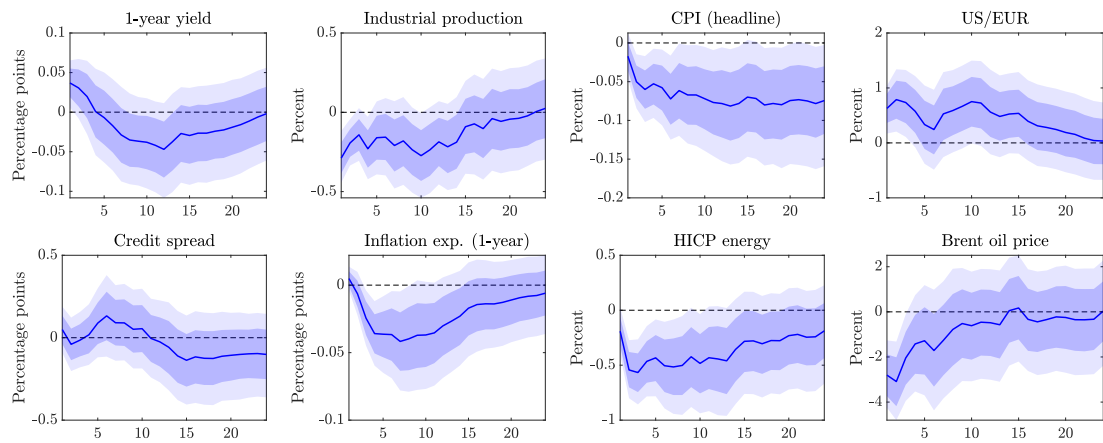
Notes: Euro Area SVAR model including the Pandemic (see text for details). Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months.

Figure C.6: Euro Area SVAR model, proxy purged of serial correlation at meeting frequency



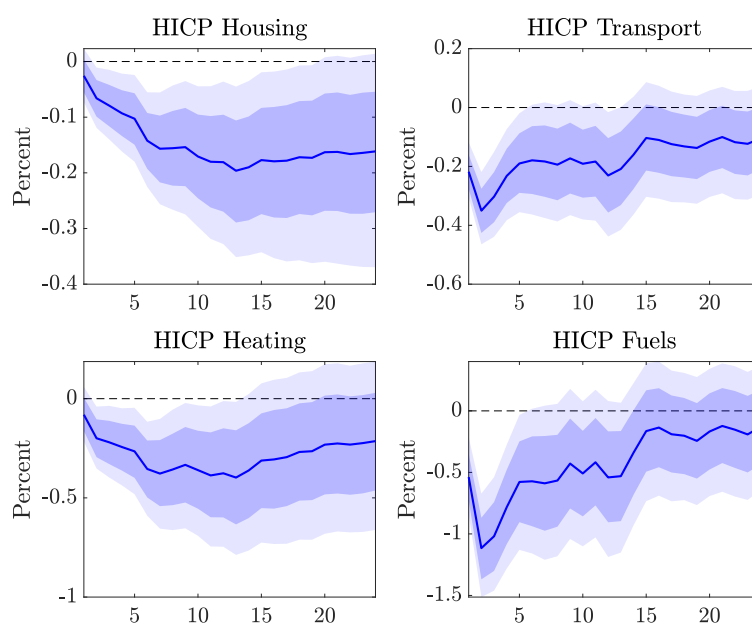
Notes: Euro Area model, where the monetary policy proxy is purged of serial correlation at the meeting frequency (see text for details). Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise probability bands. Horizon in months.

Figure C.7: Euro Area SVAR model, zero proxy relevance prior threshold



Notes: Euro Area model, with 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 C.8: Euro Area SVAR model including different subcomponents of HICP energy



Euro Area model with energy-intensive subcomponents of the HICP instead of the energy component. 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). 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. Horizon in months.

D Details on the Bayesian Proxy SVAR model

In this appendix we briefly lay out the BPSVAR model for the general case with $k \geq 1$ proxy variables and k structural shocks of interest. For convenience, we reproduce the BSPVAR model equations from section 3.1.

Following the notation of Rubio-Ramirez et al. (2010), consider without loss of generality the structural VAR model with one lag and without deterministic terms

$$\mathbf{y}'_t \mathbf{A}_0 = \mathbf{y}'_{t-1} \mathbf{A}_1 + \boldsymbol{\epsilon}'_t, \quad \boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{I}_n), \quad (\text{D.1})$$

where \mathbf{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 \mathbf{m}_t that are (i) correlated with the k structural shocks of interest $\boldsymbol{\epsilon}_t^*$ and (ii) orthogonal to the remaining structural shocks $\boldsymbol{\epsilon}_t^o$. Formally, the identifying assumptions are

$$E[\boldsymbol{\epsilon}_t^* \mathbf{m}'_t] = \underset{(k \times k)}{\mathbf{V}}, \quad (\text{D.2a})$$

$$E[\boldsymbol{\epsilon}_t^o \mathbf{m}'_t] = \underset{((n-k) \times k)}{\mathbf{0}}, \quad (\text{D.2b})$$

and represent the relevance and the exogeneity condition, respectively.

Denote by $\tilde{\mathbf{y}}'_t \equiv (\mathbf{y}'_t, \mathbf{m}'_t)$, by $\tilde{\mathbf{A}}_\ell$ the corresponding $\tilde{n} \times \tilde{n}$ coefficient matrices with $\tilde{n} = n+k$, by $\tilde{\boldsymbol{\epsilon}} \equiv (\boldsymbol{\epsilon}'_t, \mathbf{v}'_t)' \sim N(\mathbf{0}, \mathbf{I}_{n+k})$, where \mathbf{v}_t is a $k \times 1$ vector of measurement errors (see below). The augmented structural VAR model is then given by

$$\tilde{\mathbf{y}}'_t \tilde{\mathbf{A}}_0 = \tilde{\mathbf{y}}'_{t-1} \tilde{\mathbf{A}}_1 + \tilde{\boldsymbol{\epsilon}}'_t. \quad (\text{D.3})$$

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

$$\tilde{\mathbf{A}}_\ell = \begin{pmatrix} \mathbf{A}_\ell & \boldsymbol{\Gamma}_{\ell,1} \\ \mathbf{0} & \boldsymbol{\Gamma}_{\ell,2} \end{pmatrix}, \quad \ell = 0, 1. \quad (\text{D.4})$$

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

$$\tilde{\mathbf{y}}'_t = \tilde{\mathbf{y}}'_{t-1} \tilde{\mathbf{A}}_1 \tilde{\mathbf{A}}_0^{-1} + \tilde{\boldsymbol{\epsilon}}'_t \tilde{\mathbf{A}}_0^{-1}. \quad (\text{D.5})$$

Because the inverse of $\tilde{\mathbf{A}}_0$ in Equation (D.4) is given by

$$\tilde{\mathbf{A}}_0^{-1} = \begin{pmatrix} \mathbf{A}_0^{-1} & -\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1} \\ 0 & \mathbf{\Gamma}_{0,2}^{-1} \end{pmatrix}, \quad (\text{D.6})$$

the last k equations of the reduced form of the VAR model in Equation (D.5) read as

$$\mathbf{m}'_t = \tilde{\mathbf{y}}'_{t-1}\tilde{\mathbf{A}}_1 \begin{pmatrix} -\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1} \\ \mathbf{\Gamma}_{0,2}^{-1} \end{pmatrix} - \boldsymbol{\epsilon}'_t\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1} + \mathbf{v}'_t\mathbf{\Gamma}_{0,2}^{-1}, \quad (\text{D.7})$$

which shows that in the BPSVAR framework the proxy variables may be serially correlated and affected by past values of the endogenous variables and measurement error.

Ordering the structural shocks so that $\boldsymbol{\epsilon}_t = (\boldsymbol{\epsilon}_t^{o'}, \boldsymbol{\epsilon}_t^{*'})'$ yields

$$E[\boldsymbol{\epsilon}_t\mathbf{m}'_t] = -\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1} = \begin{pmatrix} \mathbf{0} \\ \mathbf{V} \end{pmatrix}. \quad (\text{D.8})$$

$\begin{matrix} ((n-k) \times k) \\ (k \times k) \end{matrix}$

The first equality is obtained using Equation (D.7) and because the structural shocks $\boldsymbol{\epsilon}_t$ are by assumption orthogonal to \mathbf{y}_{t-1} and \mathbf{v}_t . The second equality is due to the exogeneity and relevance conditions in Equations (D.2a) and (D.2b). Equation (D.8) shows that the identifying assumptions imply restrictions on the last k columns of the contemporaneous structural impact coefficients in $\tilde{\mathbf{A}}_0^{-1}$. In particular, if the exogeneity condition in Equation (D.2b) holds, the first $n - k$ rows of the upper right-hand side sub-matrix $\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1}$ of $\tilde{\mathbf{A}}_0^{-1}$ in Equation (D.6) are zero. From the reduced form in Equation (D.5) 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 (D.2a) holds, the last k rows of the upper right-hand side sub-matrix $\mathbf{A}_0^{-1}\mathbf{\Gamma}_{0,1}\mathbf{\Gamma}_{0,2}^{-1}$ of $\tilde{\mathbf{A}}_0^{-1}$ are different from zero. From the reduced form in Equation (D.5) 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 \mathbf{A}_0 and $\mathbf{\Gamma}_{0,\ell}$ such that the restrictions on $\tilde{\mathbf{A}}_0^{-1}$ implied by Equations (D.2a) and (D.2b) as well as on $\tilde{\mathbf{A}}_\ell$ in Equation (D.4) are simultaneously satisfied, and hence the estimation identifies the structural shocks $\boldsymbol{\epsilon}_t^*$.

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—as done below—using multiple proxy variables (Mertens and Ravn (2013)). 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 potentially arises when identifying multiple shocks using multiple proxies.³¹

E Discussion of the magnitude of the oil price response

Figure 1 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 3%. To shed some light on the plausibility on the magnitudes, we (i) use a representative agent version of the HANK model of Bayer et al. (forthcoming) to gauge the plausibility of our results through the lens of a state-of-the-art model and (ii) produce the impulse responses shown in Figure 1 for standardized versions of the variables to put the magnitudes into context.³²


The model of Bayer et al. (forthcoming) 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. 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. (forthcoming).

Figure E.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 goods has to fall a lot to ensure that households ultimately buy the fixed amount of energy goods supplied.

³¹We 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.

³²We thank Fabian Seyrich for sharing the code with us.


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



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 E.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 a quarter of a standard deviation and the oil prices to fall by roughly half a standard deviation. Comparing this to the responses in levels of 5 basis points and 3% it becomes apparent that the 3% fall in the oil price is by no means excessively large. The second row, which plots the IRFs in terms of their unconditional standard deviations, underscores this result. In line with the intuition that we derived above, the oil price is just an arguably much more volatile object.

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



F Details on the hybrid approximation of the structural counterfactual

As with the SZ approximation its possible to parse the sequential procedure for the “hybrid counterfactual” of MW directly into the representation in Equation (11) and thereby uncover the assumptions on the underlying structural model as described by the matrices $\tilde{\mathbf{A}}, \Theta_{\nu, \mathcal{A}}, \Theta_{\epsilon, \mathcal{A}}$, under which the resulting counterfactual path of the endogenous variables and the policy instrument actually correspond to the true counterfactual. The intuition for the underlying assumed structural model is that agents are myopic with respect to announced changes to the future level of the policy instrument. In this economy, there exists a policymaker who can credibly announce changes to the policy instruments materializing today and furthermore controls the *time t expectations* about the n_e periods ahead level of the policy instruments. In other words, the time t expectations about the future path of the policy instrument for n_e periods themselves become a policy instrument with its own policy rule and crucially with its own distinct shocks to those expectations which already enter the information set of the agents at time t . For example, in each period t there not only exists a shock to the current (0 periods ahead) level of the policy instrument ν_0 but also a shock to the *time t expectations* about the one period ahead level $\nu_{\mathcal{E}_t, 1}$. The latter already enter the information set of agents at time t , which stands in contrast to an announced change to the level of the instrument ν_1 . For the matrix notation of the hybrid counterfactual in each period there are $n_e + 1$ policy rules in which have to be embedded into Equation (11) by appropriately changing the matrix $\tilde{\mathbf{A}}$. Furthermore all n_e period ahead, model consistent, expectations of all the n_p variables that are necessary to describe the expected policy rule need to be appended to the matrices of impulse responses $\Theta_{\nu, \mathcal{A}}$ and $\Theta_{\epsilon, \mathcal{A}}$. This implies that the number of variables in the system of Equation 11 increases to $n^{\text{hybrid}} = n + (n_e \times n_p)$. The second crucial assumption is that for each period t , the policymaker has access to a set of n_e distinct policy shocks which are shocks to the n_e period ahead of expectations about the policy instrument which implies that the number of n_ν policy (news) shocks in $\boldsymbol{\nu}$ increases to $n_\nu^{\text{hybrid}} = n_h \times (n_e + 1)$. Under these assumptions its possible to show that for each distinct policy shock ν_i , i.e. for each distinct shock to the level of the policy rate or its time t expectations, and each variable j , the matrix $\Theta_{\nu_i, j, \mathcal{A}}^{\text{hybrid}}$ is lower triangular. The intuition is that, for each period t agents only observe/care about the shock of that type that materialized in that period. To be more precise, each column c of matrix $\Theta_{\nu_i, j, \mathcal{A}}^{\text{hybrid}}$, which describes the n_h responses of variable j to a news shock of type ν_i realized in c periods, is given by $\Theta_{\nu_i, j, \mathcal{A}, c} = [\mathbf{0}'_{1 \times (c-1)}, \Theta'_{\nu_i, t, j, \mathcal{A}, 0: (n_h - c)}]'$ where the vector $\Theta_{\nu_i, t, \mathcal{A}, j, 0: (n_h - c)}$ describes the impulse responses of variable j to a shock announcing a change to the $0 \leq i \leq n_h$ period ahead expectation of the policy instrument at time t for periods 0 to $n_h - i$. Thus with the estimated effects to a contemporaneous policy shock $\nu_{0, t}$ and estimated effects for n_e forward guidance shocks at hand, an econometrician can recover each of the $n_\nu^{\text{hybrid}} \times n_h$ matrices $\Theta_{\nu_i, j, \mathcal{A}}^{\text{hybrid}}$ for all n^{hybrid} variables. She can then, for the contempo-

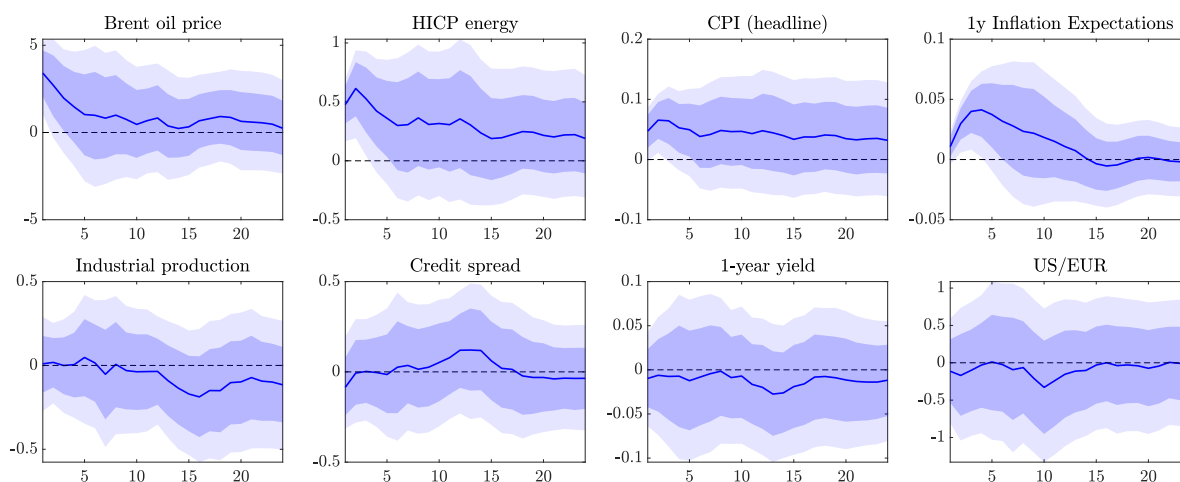
aneous shock and all forward guidance shocks, $v_0 \dots v_{n_e}$, stack the matrices across variables to form $\Theta_{\nu_i, \mathcal{A}}^{\text{hybrid}} = [\Theta_{\nu_i, 1, \mathcal{A}}^{\text{hybrid}'}, \Theta_{\nu_i, 2, \mathcal{A}}^{\text{hybrid}'}, \dots, \Theta_{\nu_i, n_{\text{hybrid}}, \mathcal{A}}^{\text{hybrid}'}]'$ and then stack all those matrices to arrive at $\Theta_{\nu, \mathcal{A}}^{\text{hybrid}} = [\Theta_{\nu_1, \mathcal{A}}^{\text{hybrid}'}, \Theta_{\nu_2, \mathcal{A}}^{\text{hybrid}'}, \dots, \Theta_{\nu_{n_e}, \mathcal{A}}^{\text{hybrid}'}]'$ which under the assumptions outlined above corresponds to $\Theta_{\nu, \mathcal{A}}^{\text{true}}$. With this structure the matrix $(\tilde{\mathcal{A}}[\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}])$ is invertible and the solution to the problem in Equation 11 is given by $\tilde{\nu}_{\text{hybrid}} = -(\tilde{\mathcal{A}}[\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}])^{-1} \tilde{\mathcal{A}}[\Theta_{A, \epsilon}]$.

Intuitively, given the assumed structure of $\Theta_{\nu, \mathcal{A}}^{\text{hybrid}}$ implied by the “hybrid counterfactual” it becomes apparent that the hybrid solution will correspond to the true solution in Equation (11) if expectations about the policy instrument more than n_e periods ahead expectations do not enter the decision problem of agents in the underlying data generating structural model. This makes the expected counterfactual path of the policy instrument for these periods irrelevant and therefore the approximation error from neglecting those expectations vanishes. This counterfactual approach could for instance be a good approximation if, in the underlying true model, agents follow some form of exponential discounting, so that the importance of the expectations decays quickly with the horizon. To summarize: Although the hybrid approach perfectly enforces the policy rule in each period and in n_e period ahead expectations, the accuracy of the approximation depends on the degree of myopia of agents in the non-policy block and/or the information set of those agents in the (unknown) structural model underlying the data. ³³

³³Despite these stark assumptions on the degree of myopia in the true underlying model, it’s possible to show as the number of explicit policy rules $n_e + 1$ approaches n_h , the resulting counterfactual from this approach converges to the true counterfactual independent of the true degree of myopia in the underlying structural model. The intuition is that, by announcing the policy rule and corresponding deviations for each period already at time t , agents at time t directly observe those rules and therefore form the correct expectations about the future path of the policy instrument. Even for a full information rational expectations HANK model, the estimated counterfactual where the counterfactual policy rule is only enforced at point time t and in $n_e = 1$ period ahead expectations, this approximation already comes very close to the true counterfactual obtained from the underlying model (see McKay and Wolf (2022a)).

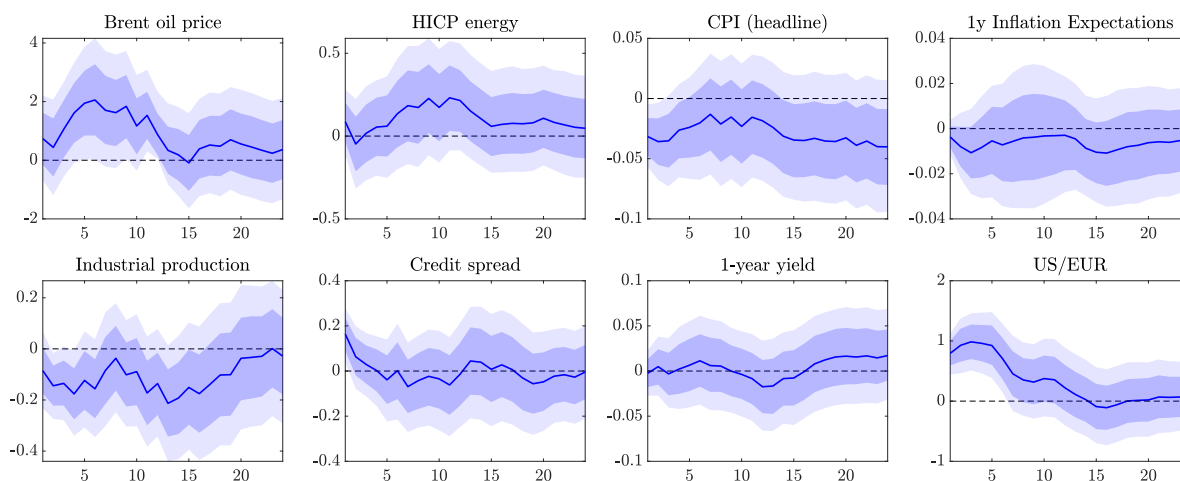
G Additional information for the OPEC policy rule counterfactual

Figure G.1: IRFs to short-run oil supply news shocks



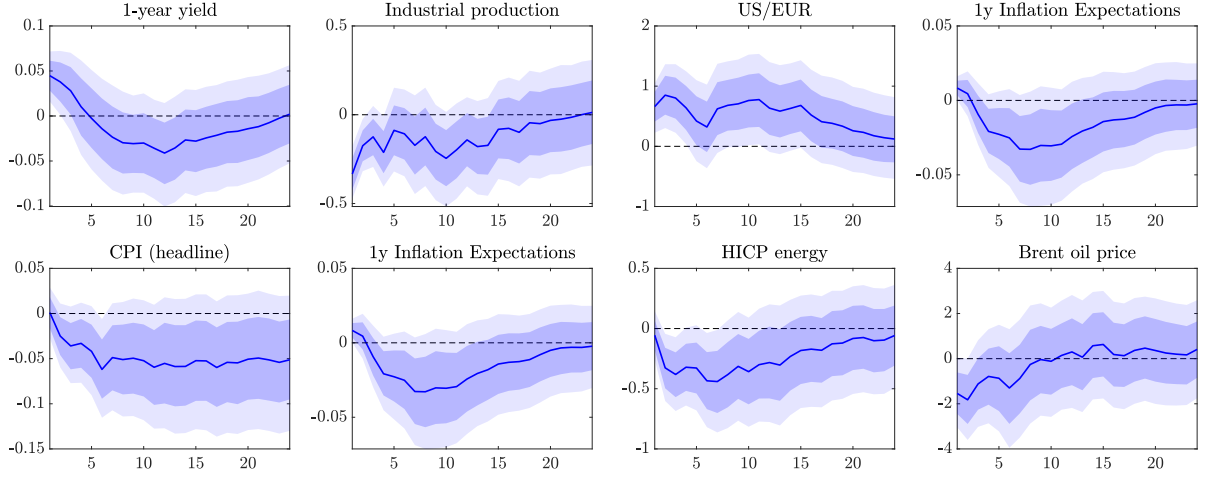
Notes: Impulse response functions to a one standard deviation short-run oil supply news shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue.

Figure G.2: IRFs to medium-run oil supply news shocks



Notes: Impulse response functions to a one standard deviation medium-run oil supply news shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue.

Figure G.3: IRFs to a euro area monetary policy shock when jointly identified alongside the oil supply shocks



Notes: Impulse response functions to a one standard deviation monetary policy shock. Point-wise posterior means along with 68% and 90% point-wise credible sets. Horizon in months.

H Further material for the optimal policy counterfactuals

H.1 Deriving the optimal policy rule

Focusing on a single variable \mathbf{x}_i , Equation (15) implies that the space of possible allocations that the policymaker can achieve for this variable is given by

$$\mathbf{x}_i = \sum_{j=1}^{n_\nu} \Theta_{x_i, \nu_j, \mathcal{A}} \times \nu_j. \quad (\text{H.1})$$

Plugging this expression into Equation (14) and taking the first-order conditions with respect to each ν_j , one arrives at the condition

$$\sum_{i=1}^{n_x} \lambda_i \Theta'_{x_i, \nu, \mathcal{A}} W \times \mathbf{x}_i = \mathbf{0}. \quad (\text{H.2})$$

For each \mathbf{x}_i the term in front of the sum describes how a change in the policy instruments ν would translate into a change in the endogenous variable \mathbf{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 (14). This rule then implies that the (weighted) sum of changes in the objective function resulting from a change in the policy instruments ν 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 $\mathcal{A}_x, \mathcal{A}_z$ of the sequence-space representation of the model in Equation (6) by noting that the optimality condition in Equation (H.2) can be written as

$$\sum_{i=1}^{n_x} \lambda_i \Theta'_{x_i, \nu, \mathcal{A}} W \mathbf{x}_i = \mathcal{A}_x^* \mathbf{x} = \mathbf{0}. \quad (\text{H.3})$$

H.2 Estimating the baseline optimal policy counterfactual

As laid out in the section 5 we construct the estimate of the optimal using the same multistep procedure as in McKay and Wolf (2023), Wolf (2023) and Caravello et al. (2023). We describe the approach in more detail below.

First, we estimate impulse responses to an identified oil price 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. We estimate the impulse responses starting the sample in 1999 as we do in Section 3.

Second, we identify the euro area conventional monetary policy and forward guidance shocks by combining the high-frequency proxies with the zero, magnitude, and sign restrictions described in the text. Again we use the same variables and transformations as in step 1. As stated in the main text for this estimation we start the sample in 2002 to take into account the liquidity concerns raised by Altavilla et al. (2019), which are particularly severe for the high-frequency changes in the longer run maturities, which we employ in this specification.

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

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

H.3 Estimating the impulse responses under counterfactual optimal policy

The procedure for the counterfactual optimal policy is very similar to the one sketched in H.2 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. We estimate the impulse responses starting the sample in 1999 as we do in section 3.

Second, we identify the euro area conventional monetary policy and forward guidance shocks by combining the high-frequency proxies with the zero, magnitude, and sign restrictions

described in the text. Again we use the same variables and transformations as in step 1. As stated in the main text for this estimation we start the sample in 2002 to take into account the liquidity concerns raised by Altavilla et al. (2019), which are particularly severe for the high-frequency changes in the longer run maturities, which we employ in this specification.

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 4.2. Again we start the sample in 1999 as we do not need to take into account the aforementioned liquidity concerns for the interest rate futures. The impulse responses to the identified shocks are plotted in Figures H.3 and H.4. The differences of the estimated with responses to those depicted in Figures G.1 and G.2 arise because we additionally included the 5-year yield.

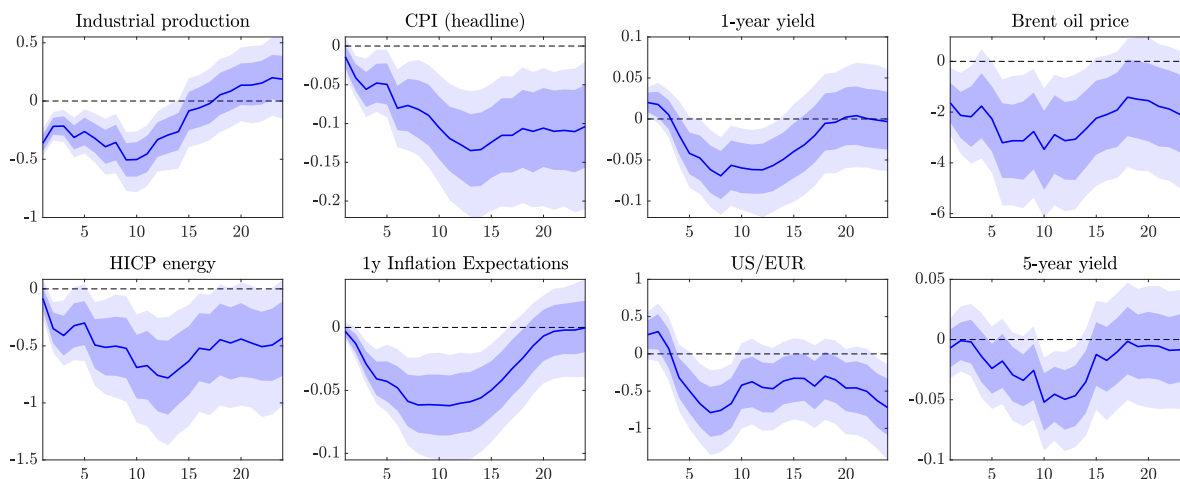
Fourth, we compute the posterior distribution of each of the counterfactual impulse responses, where the euro area monetary policy shocks from the second step do not affect the by applying the procedure of McKay and Wolf, 2023 to each draw from the posterior distribution of the second and third step.

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.

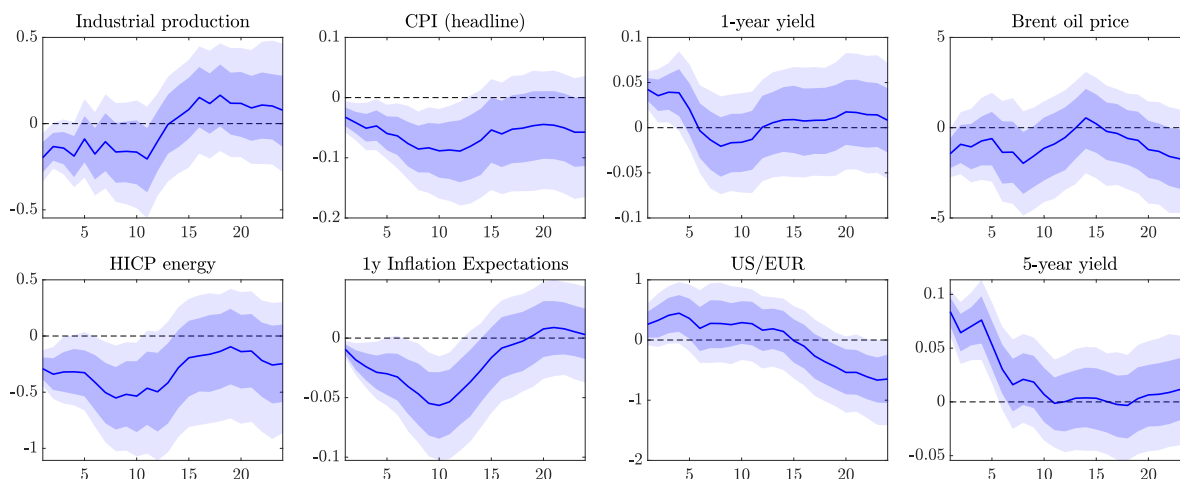
H.4 Additional figures for the optimal policy exercise

Figure H.1: IRFs to a contemporaneous EA Monetary policy shock when jointly identified alongside a forward guidance shock



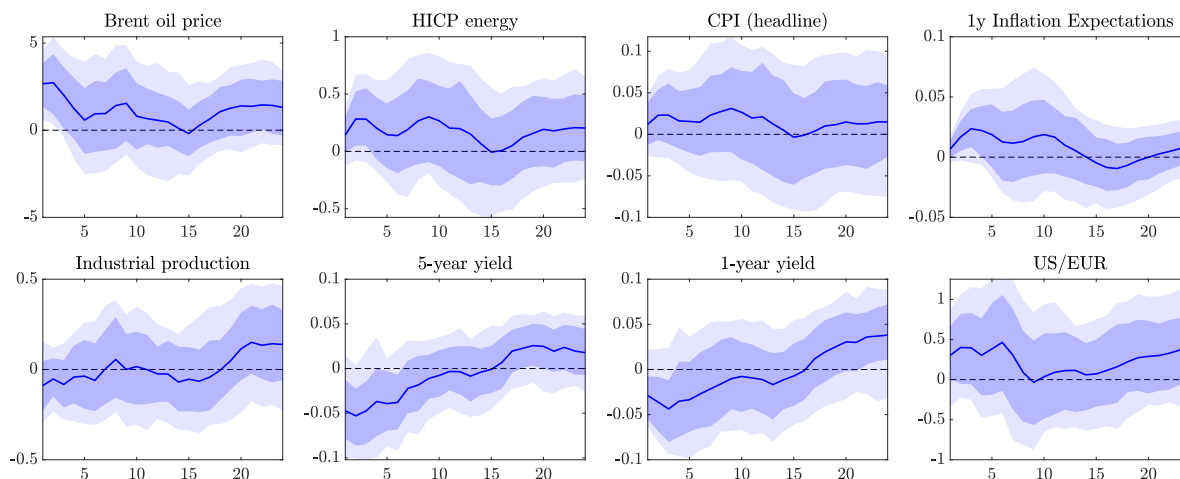
Notes: Impulse response functions to a one standard deviation conventional monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. Because we add the 5-year yield to the SVAR model, start the sample in 2002 and we jointly identify a contemporaneous monetary policy and a forward guidance shock, the estimated IRFs differ somewhat from the ones presented in 1, where we only identify a contemporaneous monetary policy shock.

Figure H.2: IRFs to an EA Forward Guidance Shock when jointly identified alongside a monetary policy shock



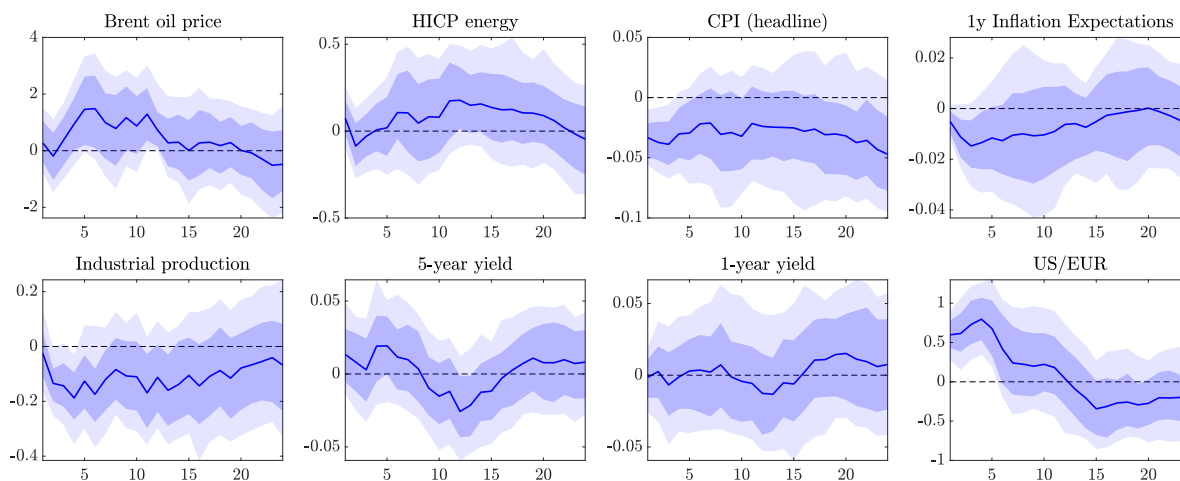
Notes: Impulse response functions to a one standard deviation forward guidance shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue.

Figure H.3: IRFs to a short-run oil supply news shock when including the 5-year yield



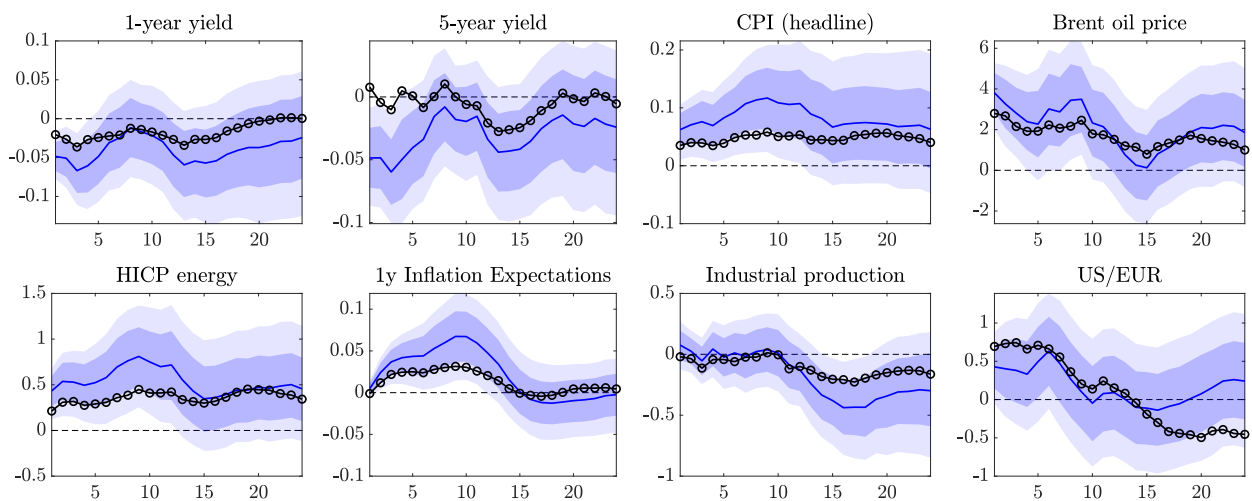
Notes: Impulse response functions to a one standard deviation short-term oil supply news shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue.

Figure H.4: IRFs to a medium-run oil supply news shock when including the 5-year yield



Notes: Impulse response functions to a one standard deviation medium-term oil supply news shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue.

Figure H.5: Impulse responses to an oil price shock under the empirical baseline and optimal (dual mandate) monetary policy rule

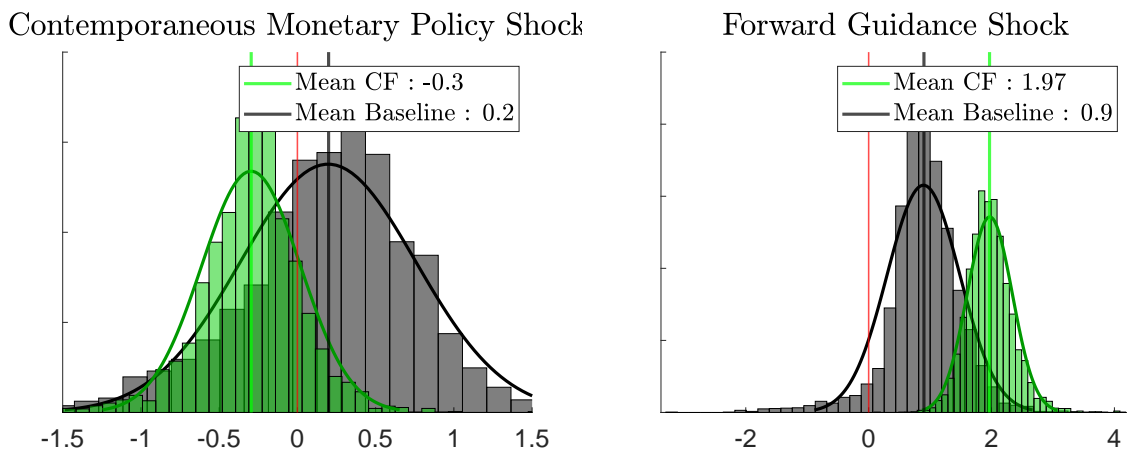


Notes: Impulse response functions to a one standard deviation monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. The black circled lines show the least squares approximation of the responses of the endogenous variables under dual mandate optimal policy with 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 ($\approx 1/3.3$).

I Implications of the global energy price response for the term structure of optimal policy

As shown in Figure 4, when the ECB’s decision do not affect global oil prices, the optimal response of the ECB to an oil price shock would entail a stronger tightening relative to the baseline optimal policy response where the euro area monetary policy shock do impact global oil prices. This is particularly true at the longer end of the yield curve, implying that the optimal strategy under this scenario would entail a significant amount of forward guidance policies. We can further corroborate this intuition by leveraging the result of McKay and Wolf (2023) who, for a large class of models such as representative and heterogeneous agents New Keynesian models, prove the equivalence between the change from the baseline to the counterfactual (optimal) monetary policy rule and a sequence of policy (news) shocks announced at date 0. While in the main body of the paper we described the differences between the baseline and counterfactuals as being caused by changes in the policy rule, we now leverage this duality to describe them in terms of the sequence of shocks necessary to mimic this policy rule. This provides additional insights into how the response of (global) energy prices changes the optimal conduct of monetary policy.

Figure I.1: Baseline and counterfactual policy shocks that characterize the optimal policy rule

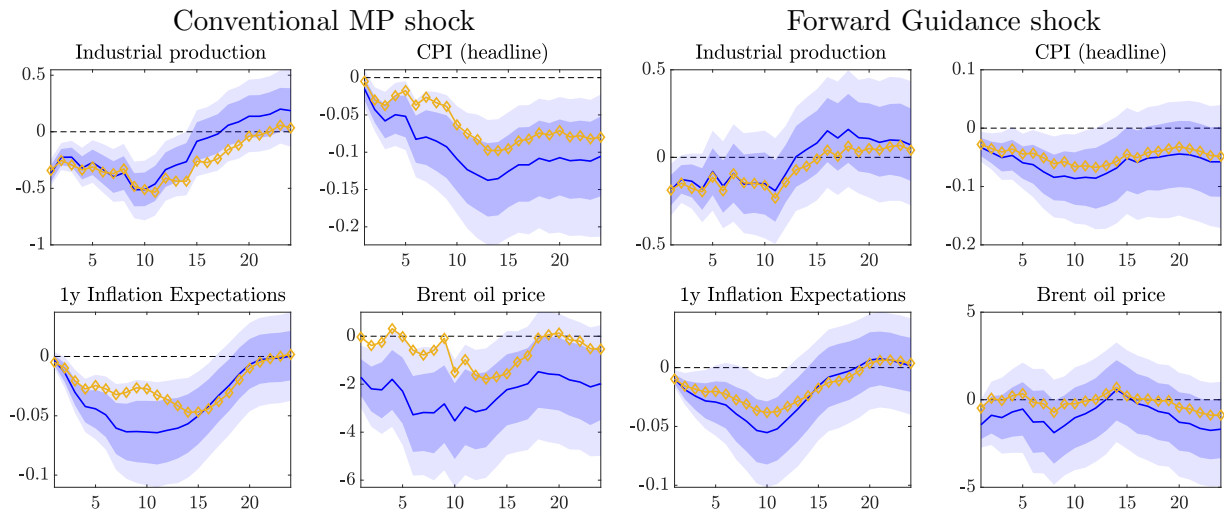


Notes: The figure plots the posterior distribution of the identified monetary policy shocks that characterized the baseline optimal policy rule (black) and counterfactual optimal policy rule (green), which is computed under the assumption that OPEC stabilizes the Brent oil price, causing euro area monetary policy to not transmit via global oil prices. Shocks are measured in standard deviations

Figure I.1 plots the posterior distribution of the estimated sequence of policy shocks that characterize the optimal policy rule under the baseline impulse responses (black) and the counterfactual impulse responses (green), where the ECB’s decisions do not impact the

global oil prices. As shown by the black bars, when the ECB’s decisions do impact oil prices, the shift from the baseline to the optimal policy rule can be characterized on average by a combination of a 0.2 standard deviation contractionary contemporaneous monetary policy shocks, which on impact raise the 1-year yield while leaving the longer run yields largely unchanged (see Figure H.1), and a 0.9 standard deviation forward guidance shock, which on impact strongly raises longer run yields and to a lesser degree the 1-year yield (see Figure H.2). This changes dramatically under the counterfactual optimal policy, where the ECB’s decisions do not impact the global oil price. In this scenario, the optimal policy rule is on average equivalent to a combination of a 0.3 standard deviation expansionary monetary policy shock alongside a whopping 2 standard deviation contractionary forward guidance shock.

Figure I.2: IRFs and counterfactual IRFs to a Conventional EA Monetary policy and a Forward Guidance shock



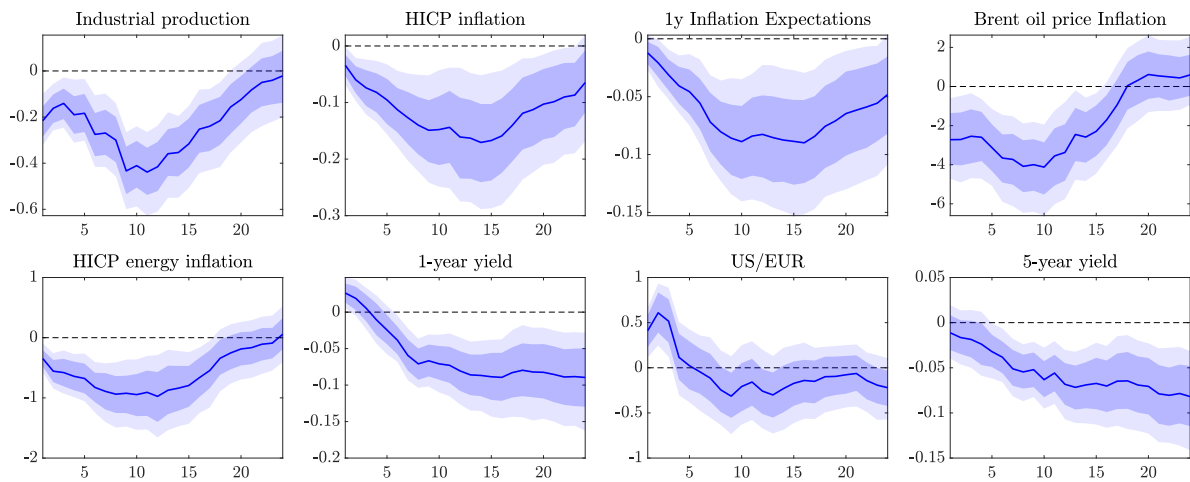
Notes: Impulse response functions to a one standard deviation conventional monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. The golden lines depict the point-wise posterior means under the assumption that OPEC, counterfactually, stabilizes the Brent oil price. Because we added the 5-year yield to the SVAR model and we jointly identify a contemporaneous monetary policy and a forward guidance shock, the estimated IRFs differ somewhat from the ones presented in Figure 1, where we only identify a generic monetary policy shock.

The reason for this shift in strategy lies in the different implications of the global component of the energy price channel for respective policy shocks or what we call the “term structure of monetary policy”. As shown in Figure H.1 the contemporaneous monetary policy shock significantly decreases the Brent oil prices and this strongly affects the transmission of the contemporaneous monetary policy shock to the HICP. On the other hand, the forward guidance shock is estimated to have a strong effect on the exchange rate, output, and ultimately the HICP, while only slightly depressing the Brent oil price as depicted in Figure H.2. Therefore, as illustrated in Figure I.2, the importance of the global component of the energy

price channel for the transmission of the forward guidance shock to the ECB's objective is estimated to be smaller compared to its contemporary monetary policy shock counterpart. Note that McKay and Wolf (2022a, p.9) show that under optimal monetary policy, “the policymaker will rely most heavily on the tools [...] that are best suited to offset the perturbation to its targets”. Exactly in line with this reasoning, when the ECB's decisions do not affect global oil prices, the optimal strategy implies putting more weight on forward guidance shocks, because they are estimated to transmit to inflation less via global energy prices and more strongly via the exchange rate and therefore among other things via local energy prices. Intuitively, the slightly expansionary contemporaneous monetary policy shock is then used to smooth out any overshooting from the target that would arise if the central bank were to purely rely on forward guidance. An arguably unintended but probably highly welcome benefit of this strategy is, that because the optimal policy rule is now characterized by a combination of an expansionary contemporaneous and a contractionary forward guidance shock, output is roughly stabilized at the level that would prevail in the baseline optimal policy case despite the significant tightening at the longer end of the yield curve.

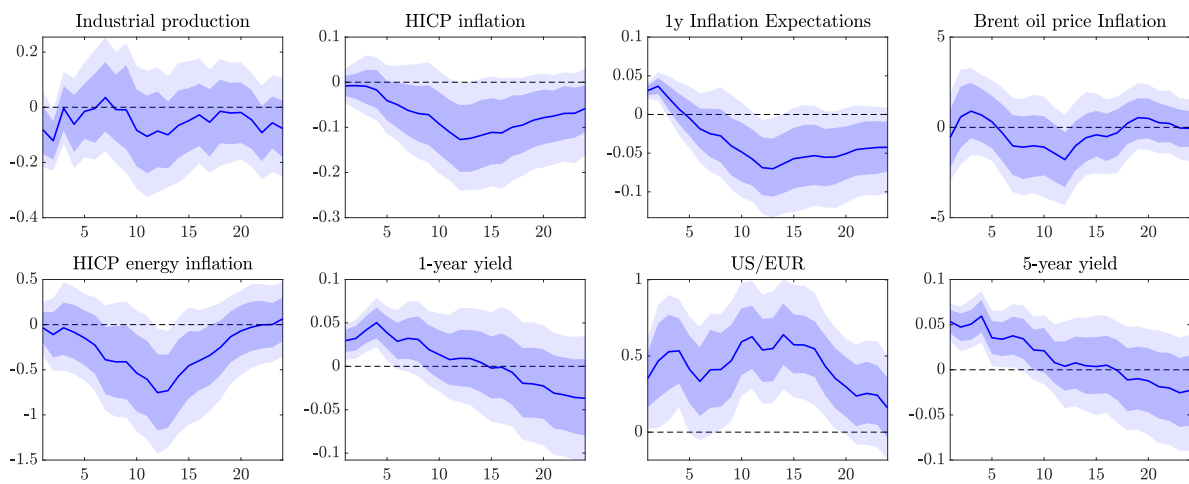
J IRFs to monetary policy shocks when including the pandemic into the estimation

Figure J.1: IRFs to a contemporaneous EA Monetary policy shock when estimating the BPSVAR using data until 2024



Notes: Impulse response functions to a one standard deviation conventional monetary policy shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. We extend the dataset for the estimation until 2024, which includes the pandemic. We explicitly model the Pandemic using the Pandemic Priors approach of Cascaldi-Garcia (2022). Furthermore, relative to the baseline, we add the 5-year yield to the SVAR model, start the sample in 2002 and we jointly identify a contemporaneous monetary policy and a forward guidance shock. Lastly we incorporate y-o-y inflation rates for prices instead of their level to ensure stationarity.

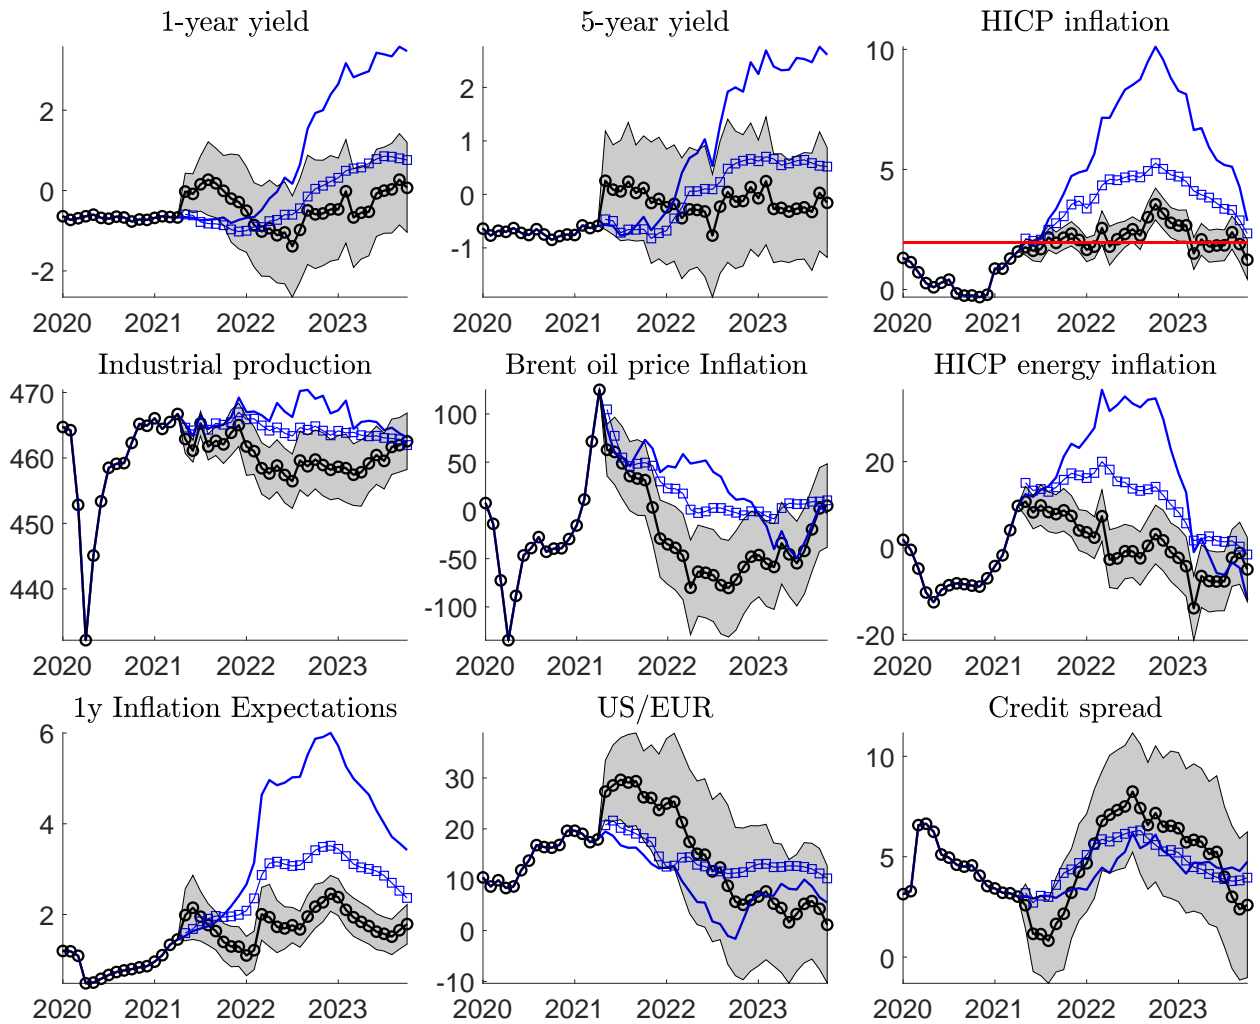
Figure J.2: IRFs to a EA forward guidance shock when estimating the BPSVAR using data until 2024



Notes: Impulse response functions to a one standard deviation forward guidance shock showing the point-wise posterior means along with 68% and 90% point-wise credible sets in blue. We extend the dataset for the estimation until 2024, which includes the pandemic. We explicitly model the Pandemic using the Pandemic Priors approach of Cascaldi-Garcia (2022). Furthermore, relative to the baseline, we add the 5-year yield to the SVAR model, start the sample in 2002 and we jointly identify a contemporaneous monetary policy and a forward guidance shock. Lastly we incorporate y-o-y inflation rates for prices instead of their level to ensure stationarity.

K Further material for the counterfactual evolution of the economy under optimal policy

Figure K.1: Counterfactual and expected evolution of the endogenous variables under optimal ECB policy



Notes: The figure plots the time series of the endogenous variables (blue solid line) alongside the estimated counterfactual evolution (black circled line) of these variables under the assumption that, from April 2021 onwards, the ECB would have conducted optimal monetary policy as described in Equation 17 and thereby minimizes the weighted squared deviations of “medium-term” inflation from the target. Furthermore, the blue solid line with diamonds represents the expected evolution (forecast) of the economy under the assumption that the ECB follows the baseline policy rule also from April 2021 onwards. We refer to this as the deterministic or expected component. The difference between the blue lines is therefore driven by surprises and represents the stochastic component.

The differences between the actual and counterfactual evolution of the economy can be decomposed into the differences in the initial conditions and the differences in the propagation

of new incoming shocks after t^* . More precisely, we first compute the respective contributions of the incoming shocks and initial conditions under the baseline policy rule characterized by \mathcal{A} . This implies that we decompose the black line in Figure K.1 along the lines of Equation (K.1)

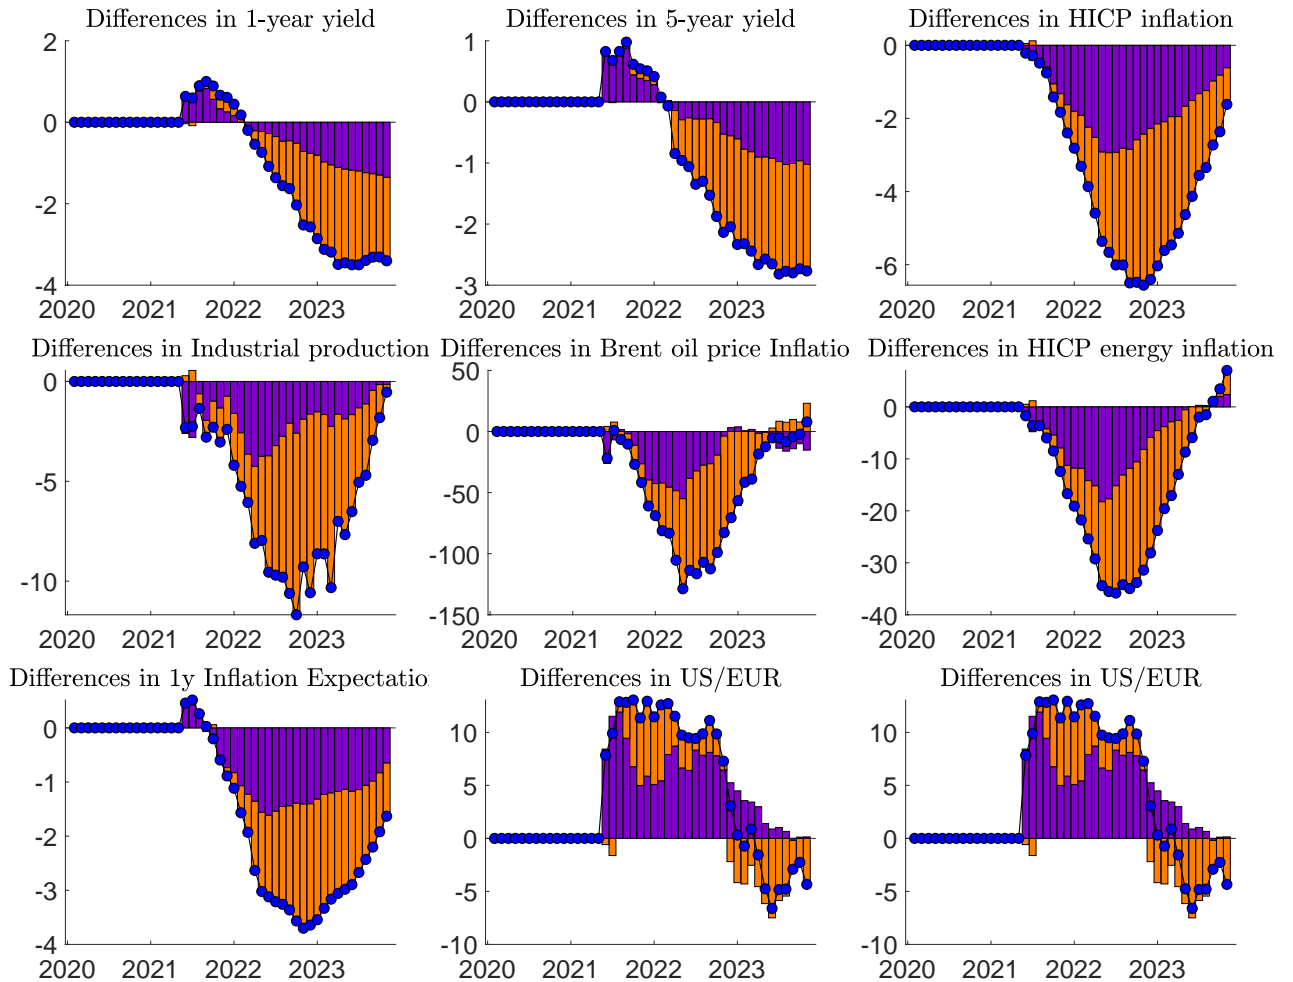
$$\mathbf{y}_t = \underbrace{\sum_{\ell=0}^{t-t^*} \Theta_{\ell, \eta, \mathcal{A}} \eta_{t-\ell}}_{\text{contrib. new shocks after } t^*} + \underbrace{\mathbf{y}_{t^*}^{\mathcal{A}}}_{\text{contrib. initial condition at } t^*}. \quad (\text{K.1})$$

We then compute the counterfactual SVMA representation, which gives the contribution of these two components under the mandate-optimal policy rule \mathcal{A}^* .

$$\mathbf{y}_t = \underbrace{\sum_{\ell=0}^{t-t^*} \Theta_{\ell, \eta, \mathcal{A}^*} \eta_{t-\ell}}_{\text{contrib. new shocks after } t^*} + \underbrace{\mathbf{y}_{t^*}^{\mathcal{A}^*}}_{\text{contrib. initial condition at } t^*}. \quad (\text{K.2})$$

It becomes apparent that the differences between the actual and counterfactual evolution of the economy can therefore be traced back to (i) the difference in the expected evolution of the economy at date t^* and (ii) the different propagation of incoming shocks after t^* , which also captures the fact that, under optimal policy, there exist no monetary policy related shocks. Those differences are shown in K.2 in purple and orange respectively. Looking at the panel reveals that the realignment of expectations with the mandate that is due to the switch in the policy rule, already explains roughly 50% of the total difference in HICP inflation between the actual and counterfactual evolution of the economy.


Figure K.2: Counterfactual and expected evolution of the endogenous variables under optimal ECB policy



Notes: The Figure decomposes the difference between the observed evolution of the economy (blue lines in Figure K.1) and the counterfactual evolution of the economy (black lines in Figure K.1). These differences arise because of the change in the policy rule from the baseline to the optimal policy rule, which we assume to happen in April 2021. They can be decomposed into differences in deterministic components (i.e. the expected path of the endogenous variables without any further shocks) which are depicted in purple and differences in the stochastic component (i.e. the way that new, unexpected shocks propagate) which are depicted in orange.

L Oil import shares

Figure L.1: Oil import shares



Notes: Figure represents the global share of crude oil imports of the euro area, China and the United States at five-year intervals from 2000 to 2020. The global share is computed as a percentage of total crude oil imports in the world using data from the International Energy Agency (IEA).