Bank of England

A structural VAR model for the UK economy

Macro Technical Paper No. 3

July 2025

Davide Brignone and Michele Piffer

A series designed to document models, analysis and conceptual frameworks for monetary policy preparation – they are written by Bank staff to encourage feedback and foster continued model development.

Macro Technical Paper Series

Dr Bernanke's 2024 **review** of the monetary policymaking processes at the Bank of England provided a number of constructive recommendations for reform which we are taking forward.

This included improving our model maintenance and development. Macroeconomic models and frameworks play an important role in the monetary policy process. To maximise the value of macroeconomic models, they must be well documented and continuously improved as time goes on.

This Macro Technical Paper (MTP) series is part of the Bank's response to Dr Bernanke's recommendations. These MTPs are authored by Bank staff, and are intended to document models, analysis, and conceptual frameworks that underpin monetary policy preparation. The models documented in the series will typically be used to assess the current state of the economy, forecast its future, and to simulate alternative paths and policy responses.

Importantly, while each MTP will provide insights about a particular model or modelling framework that is an 'input' to policy, no single model can possibly capture all the relevant features to perform even just one of those roles adequately. Models will inevitably have to be updated and will improve over time, including as they are adapted to different constellations of macroeconomic conditions. This is a natural part of real-time monetary policy making. So, inevitably, no MTP will provide definitive answers.

The Bank seeks to encourage an active and informed debate about its modelling frameworks. Publishing and discussing the analytical work undertaken to support its monetary policy choices is central to this ambition. These MTPs will support a culture of continuous learning in monetary policy making. As time goes on, Bank staff will update and upgrade models, drawing on insights from the frontier of the academic literature. Moreover, this transparency will encourage external engagement from experts to ensure our modelling tools remain fit for purpose.

Clare Lombardelli
Deputy Governor for Monetary Policy

Bank of England

Macro Technical Paper No. 3

A structural VAR model for the UK economy

Davide Brignone[1] and Michele Piffer[2]

Abstract

We present a Structural VAR (SVAR) model for the UK economy, which is part of a wider modelling toolkit used to inform monetary policymaking. The model includes eight global and UK macroeconomic variables, and identifies a set of global and domestic shocks. We describe standard model outputs, including impulse responses, forecast-error variance decompositions, and historical decompositions. We then describe how the model can be used on a round-by-round basis to inform policy discussions by providing a structural narrative for forecast revisions.

Key words: Structural shocks, historical and forecast decomposition, macro-modelling.

JEL classification: C32, E52.

---

(1) Bank of England. Email: davide.brignone@bankofengland.co.uk
(2) Bank of England and King's Business School, King's College London.
Email: michele.piffer@bankofengland.co.uk

We thank several colleagues at the Bank of England who helped develop the model and integrate it within the policy process at the Bank. We are grateful to Ida Hjortsøe, Simon Hong, David Latto, Simon Lloyd, Andre Moreira, and Kate Reinold for helpful comments on this paper. The views expressed in this paper are those of the authors and do not necessarily represent those of the Bank of England or any of its committees.

The Bank's macro technical paper series can be found at
www.bankofengland.co.uk/macro-technical-paper/macro-technical-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 2978-3194 (online)

1 Introduction

Structural Vector Autoregression (SVAR) models are widely used by macroeconomists in academia, international institutions and central banks. This class of models is closely related to Vector Autoregression (VAR) models that, since their introduction by Sims (1980), have been extensively studied using both frequentist and Bayesian inference (for which they assume the acronym: BVARs), extended to various non-linear frameworks and applied in numerous empirical settings. While BVARs have proven to be an especially powerful tool for forecasting the economy, they miss a structural interpretation of the shocks in the model. Bayesian SVARs add a *structure* to BVARs and offer a versatile framework to analyse the transmission mechanisms of economic shocks, as well as to understand how the economy responds to different disturbances.

Being such an important tool for macroeconomic analysis, SVAR models are part of the toolkit used regularly by the Bank of England. The aim of this *Macro Technical Paper* is twofold. First, it discusses the recent specification of a SVAR model that is currently used at the Bank, in particular to disentangle *domestic* drivers of the UK business cycle. Second, it illustrates some of the most relevant applications of the model and explains how it can help monetary policymakers on a round-by-round basis. The model is designed as a tool for studying economic dynamics at business-cycle frequencies, focused on short- to medium-term horizons. It has featured in various recent speeches of the MPC members (see, e.g., Pill, 2024; Taylor, 2025; Breeden, 2025).

Main ingredients of the model. Broadly speaking, VAR models help to shed light on the relationship between variables at time t and their own past values. This delivers a set of estimated correlations that hold over the estimation sample. The identification of shocks then helps interpret those correlations as responses to structural, unexpected, shocks that hit the economic system. Due to the nature of the model, each shock is associated with a dynamic response of the variables in the system. As such, the model can help to uncover in real time which shocks have hit the economic system, and illustrate how their effects propagate in the near term. This can support policymakers in assessing the state of the economy over the business cycle.

The model discussed in this paper includes eight macroeconomic indicators, spanning three global variables and five UK variables. On the global side, the model identifies a global demand and two shocks to global supply.[1] The primary policy applications of this model are focused on the domestic shocks, for which the model identifies a UK demand, UK supply, and UK monetary policy shock. Two remaining shocks are left unidentified (one global and one domestic), to capture unexplained or additional sources of forecasting errors.

The model should be seen as a flexible platform whose variables and identified shocks can be adjusted from round-to-round depending on the set of relevant questions faced by policymakers. Identification is achieved via standard zero and sign restrictions (Uhlig, 2005; Baumeister and Hamilton, 2015; Arias et al., 2018). The data is quarterly.

Using the model in practice. This *Macro Technical Paper* describes two different uses of the model. First, we present more standard SVAR-related outputs, including impulse response functions, forecast-error variance decompositions, estimated structural shocks, and historical decompositions. These outputs are commonly used to interpret the dynamic effects of shocks on the economy, and to decompose past data.

Second, we introduce a tool that provides a *structural narrative* for forecast revisions between rounds. Forecast revisions are usually treated as reduced-form concepts. However, it can be valuable to understand the structural *drivers* behind forecast revisions, as this information can support policymakers' timely assessment of the state of the economy. To set ideas, Figure 1 shows the unconditional forecasts from the SVAR for year-on-year (y-o-y) UK GDP growth and UK inflation produced in August 2024 (solid orange line) compared to the forecast for the same variables produced in the previous quarter, May 2024 (dashed lines). As the figure shows, the forecast for UK GDP growth was revised upwards between the two quarters, while no material revisions took place for inflation. Our goal is to interpret those forecast changes through the lens of structural shocks, and assess whether these revisions have meaningful dynamic implications for key macroeconomic variables over the forecast horizon. Going back to the example shown in Figure 1, we find that forecast revisions can (at least partially) be explained by a mix of contractionary global demand and expansionary UK demand shocks that hit the economy in 2024Q2. We will describe this in more detail later on in the paper.

In summary, the first set of model applications offer a parsimonious and flexible solution for decomposing the evolution of past data into key structural shocks. It can

[1] A forthcoming *Macro Technical Paper* will explore further the *global* drivers of economic fluctuations in the UK.

Figure 1: Forecast revision

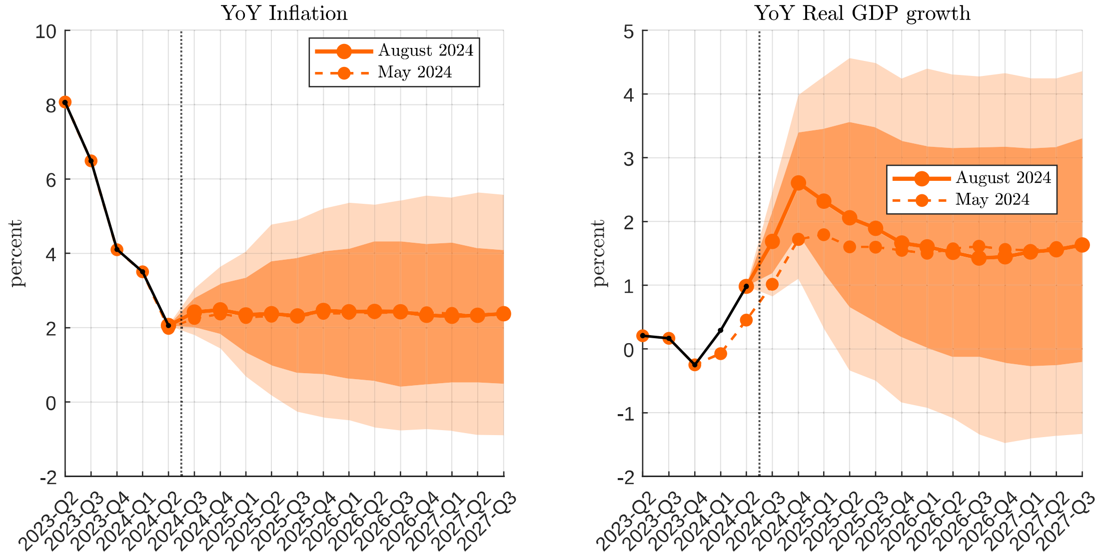

Note: The August 2024 forecast (time T) is shown as pointwise median, and 68% and 90% credible band. The May 2024 forecast (time $T-1$) is shown only via pointwise median

provide policymakers a narrative for the structural drivers of variables in the model in *absolute* terms—for instance, whether macroeconomic variables can be explained with supply or demand shocks. The second application of the model provides a structural decomposition of new data and of the forecast revisions between rounds. Therefore, it offers a narrative on the drivers of the variables and forecast in *marginal* space, along with how these revisions dynamically impact key variables over the forecast horizon. This tool can be updated on a round-by-round basis and thus offers a consistent structural narrative of the revisions across different rounds.

Relation to the literature. This paper contributes to the broader literature on the use of SVARs and BVARs in policy analysis, both within the Bank of England and across other institutions, for which providing an exhaustive list would be a challenging task. Recent examples can be found, for instance, in Cesa-Bianchi et al. (2020), Forbes et al. (2018), Brandt and Burr (2024) and Braun et al. (2025) from the Bank of England, and Bańbura et al. (2023) for the ECB.[2] Crump et al. (2025) recently focused instead on the use of a medium-size BVAR for policy scenarios for the U.S. economy. These, and other contributions, build on a modeling framework that is out-

[2] We refer readers to Ciccarelli et al. (2024) for a more complete—but still not exhaustive—list of models used at the ECB.

lined in Canova (2011) and Kilian and Lütkepohl (2017), and surveyed in Christiano (2012). Our UK-specific model contributes to the literature by employing SVAR models to decompose forecast revisions. Some contributions along the same direction were proposed by Todd (1992) and Giannone et al. (2004), although in models with fewer structural shocks, and with a more limited exploration of data revision.

Outline. The paper proceeds as follows. Section 2 outlines the model specification, data, estimation approach and shock identification. Sections 3 and 4 then briefly present standard model properties and a decomposition of past data through the lens of the identified shocks. In Section 5, we describe a new tool that helps to interpret forecast revisions and provide information relevant to the policymakers. To do this, we present 'real-time' analysis from the perspective of a forecaster in 2024Q2. Section 6 concludes.

2 Model Overview

In this section, we describe the building blocks of our model for interpreting the macroeconomic drivers of the key aggregate variables for the UK economy. We begin with the data and the model framework, before moving to the shock identification strategy and estimation.

2.1 Data

Our VAR includes a total of eight global and UK variables, listed in Table 1. The data is quarterly and our estimation sample begins in 1992Q1. Later in this paper, we carry out analysis from the perspective of a forecaster in 2024Q2. For this, we use the same data vintages that were available for the preparation of the August 2024 *Monetary Policy Report*. This gives us a sample that runs from 1992Q1 to 2024Q2 and that captures several episodes of interest for the UK economy, including the Global Financial Crisis (GFC) in 2007-2008, the Brexit referendum in 2016Q2, and the Covid-19 pandemic.

While most of the data enters the VAR in log-levels, some of the charts will be reported after transforming the data into a more informative scale. Table 1 shows the transformations applied to the data as it enters the estimated model and as it is shown in selected figures. Figure A-1 and Figure A-2 in the Appendix plot the

Table 1: Variables in the model and transformations

Variable	Enters model	Shown in some figures
Real World GDP	100·log	YoY growth
World CPI	100·log	YoY growth
Real oil price	100·log	YoY growth
Bank Rate	level	level
Exchange-rate index (ERI)	100·log	level
UK CPISA	100·log	YoY growth
UK CPI Energy	100·log	YoY growth
Real UK GDP	100·log	YoY growth

Note: World GDP and CPI are UK-trade weighted. The real price of oil is in sterling. The Bank of England exchange-rate index is a trade-weighted exchange value of sterling computed using bilateral exchange rates (a decrease in the index corresponds to a sterling depreciation). We include UK CPI Energy as it captures the combined role of oil and gas prices to the UK CPISA. This will help in the identification of the world energy shock, as recent episodes underlined the importance of accounting for the role played by gas prices in energy-price fluctuations.

variables graphically.

2.2 Model Setup

VAR models estimate correlations among a set of variables, both contemporaneously and dynamically. The starting point is a vector $\boldsymbol{y}_t$ of k variables of interest. These variables are studied as a function of the same variables at previous periods. The goal is to trace how these variables are dynamically driven by the underlying structural shocks $\boldsymbol{\epsilon}_t$.

Formally, the model is given by

$$\boldsymbol{y}_t = \sum_{l=1}^{p} \boldsymbol{\Pi}_l \boldsymbol{y}_{t-l} + \boldsymbol{c} + \sum_{c=1}^{q} \boldsymbol{\delta}_c d_{t-c} + \boldsymbol{u}_t, \tag{1}$$

where $\boldsymbol{\Pi}_l$ are $k \times k$ matrices of autoregressive model coefficients and $\boldsymbol{c}$ is a $k \times 1$ vector of constants. In addition to the autoregressive terms, we include a scalar dummy variable $\boldsymbol{\delta}_c$ to account for the Covid-19 period. This takes the value 1 if period $t-c$ includes the Covid-19 pandemic, which we define from 2020Q1 to 2021Q2, and 0 otherwise. This implies dummy variables for a total of six periods.

The $k \times 1$ vector of reduced-form innovations $\mathbf{u}_t$ are normally distributed, with $k \times k$ covariance matrix $\boldsymbol{\Sigma}$:

$$\boldsymbol{u}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}).$$

These reduced-form innovations are related to a $k \times 1$ vector of structural shocks that drive the data according to:

$$\boldsymbol{u}_t = \boldsymbol{B}\boldsymbol{\epsilon}_t, \tag{2}$$

where $\boldsymbol{B}$ is a $k \times k$ matrix and the structural shocks are assumed to be normally distributed with a covariance matrix normalised to the $k \times k$ identity matrix $\boldsymbol{I}_k$:

$$\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, \boldsymbol{I}_k).$$

The reduced-form covariance matrix $\boldsymbol{\Sigma}$ relates to the impact effect of the shocks matrix $\boldsymbol{B}$ via the expression:

$$\boldsymbol{\Sigma} = \boldsymbol{B}\boldsymbol{B}'. \tag{3}$$

Finally, the following equation relates $(\boldsymbol{\Sigma}, \boldsymbol{B})$ via the function $\chi(.)$, which is the unique Cholesky decomposition, and $\boldsymbol{Q}$, which is an orthogonal matrix:

$$\boldsymbol{B} = \chi(\boldsymbol{\Sigma})\boldsymbol{Q}. \tag{4}$$

We refer readers to Arias et al. (2018) and Canova et al. (2024) for a detailed discussion of alternative parameterisations of VARs.

As discussed in the introduction, the model estimates the contribution of individual structural shocks to the evolution of the data. In our notation, these shocks of interest are given by $\boldsymbol{\epsilon}_t$. The assumption that these shocks are statistically independent is necessary to view these shocks as structural. For instance, a shock can capture an exogenous increase in aggregate demand, which might lead the central bank to respond via changing the policy rate. Matrix $\mathbf{B}$ captures by how much each variable of the model responds to each shock. This response is contemporaneous, in that it occurs at the same horizon in which the shock hits the economic system. The combination of matrices $\mathbf{B}$ and $\boldsymbol{\Pi}_l$, $l = 1, .., p$ then helps trace out how responses evolve over time.

2.3 Shock Identification

As with all Gaussian SVAR models, the structural shocks $\boldsymbol{\epsilon}_t$ are not identified unless some identifying restrictions are introduced. We identify six structural shocks out of the eight in equation (2), specifically shocks to:

1. world demand;
2. world energy;

3. world supply (exogenous to world energy);

4. UK demand;

5. UK supply; and

6. UK monetary policy.

We focus on these six shocks to balance parsimony and coverage. It is important to stress that more *specific* shocks will be contained within the broader categories identified here. For instance, a U.S. monetary policy shock will be included in the world demand shock, or a domestic technology shock will implicitly be captured by the generic UK supply shock—which, at the same time, will also capture other domestic supply-side shocks. Further disentangling shocks can pose challenges for structural VARs, since these models are typically designed to identify a limited number of shocks. For this reason, SVARs should complement other models, such as DSGE (like that outlined in Albuquerque et al., 2025) and semi-structural models. Since this SVAR is partially identified—i.e., the number of identified structural shocks is smaller than the number of variables in the model—the residual unidentified shocks implicitly capture remaining volatility unexplained by our set of structural shocks.

Table 2: Identifying restrictions

	World shocks				UK shocks			
	Demand	Energy	Supply	Unid. 1	Demand	Supply	Mon. Pol.	Unid. 2
World GDP	+	+	+		0	0	0	0
World CPI	+	−	−		0	0	0	0
Oil price		−			0	0	0	0
Bank rate					+		+	
Exch. rate								
UK CPISA	+	−	−		+	−	−	
CPI Energy	+	−	0					
Real GDP	+	+	+		+	+	−	

Note: Zero and sign restrictions are introduced only on the impact effect of the shocks.

We identify the shocks using a combination of zero and sign restrictions which are summarised in Table 2. The restrictions we apply are rather standard in the literature (see, e.g. Uhlig, 2005; Baumeister and Hamilton, 2015; Arias et al., 2018). Restrictions are introduced only on the impact effect of the shocks, and no restrictions are introduced on future horizons of the impulse response, nor on the contemporaneous relationship among variables (which would affect $\mathbf{B}^{-1}$ rather than $\mathbf{B}$). The rows in Table 2 report model variables, while the columns document the identified shocks.

A future *Macro Technical Paper* will expand on the identification of world shocks. Nevertheless, in our current specification, a positive world demand shock is assumed to exert a positive impact on world and domestic GDP as well as world and domestic CPI. An expansionary (and disinflationary) world energy shock, instead, leads to an increase in world and domestic real GDP, a decrease in world and domestic CPI, and a decrease of oil prices and CPI energy. The positive world supply shock implies an increase in world and domestic real GDP, and a decrease in world and domestic CPI. To exogenise this with respect to the world energy shock, we impose that world supply shocks do not have any effect on CPI energy in the quarter of the shock.

For UK shocks, we impose a small-open economy assumption (which has been used in previous Bank analysis by Cesa-Bianchi et al., 2021) and assume that domestic shocks have zero contemporaneous impact on global variables. In turn, we then impose that an expansionary UK supply shock increases UK real GDP and decreases UK CPI, while a positive UK demand shock increases both UK real GDP, UK CPI and Bank rate. Finally, a UK monetary policy tightening shock increases Bank rate, while decreasing both UK CPI and UK real GDP.

Two shocks remain unidentified. We assume that one unidentified shocks satisfies the same zero block restrictions that separately identify UK shocks from world shocks. In doing so, we assume that one unidentified shock captures the unexplained component of global shocks, while the other unidentified shocks captures unexplained UK-specific shocks.[3]

2.4 Estimation

We estimate the model over the period from 1992Q1 to 2023Q2. Starting the estimation sample in 1992Q1 ensures that we cover the period since the introduction of inflation targeting in the UK. The estimation ends one year before the final observation in the dataset (2024Q2), as data from the most recent year are often subject to revisions. Thus, excluding the final year helps ensure that parameter estimates are not influenced by potentially significant data revisions.

To estimate the model, we use Bayesian methods. The Bayesian paradigm offers

[3]Some authors introduce the additional restrictions that the unidentified shocks do not repeat the zero/sign patterns of the identified shocks. In our specification, this is not possible with regard to the UK unidentified shock. Yet, the results are robust to a specification that removes the oil price and works with a single global unidentified shock, restricted to feature no repetition in the restrictions of the identified shocks.

a natural way of measuring uncertainty, and at the same time it allows for the introduction of non-dogmatic prior beliefs. For example, the Bayesian approach delivers a suitable tool to quantify the probability, that, say, a demand shock was positive or negative in a given period of interest. In following a Bayesian approach, we are consistent with extensive common practice in central banking.

More specifically, we follow Giannone et al. (2015) and use a Minnesota prior on the autoregressive parameters of the model along with the sum-of-coefficients priors.[4] As anticipated in the previous section, we add Covid-19 dummies for the quarters from 2020Q1 to 2021Q2 to deal with the unprecedented volatility observed for some variables during that period. To do this, we use the pandemic prior approach proposed by Cascaldi-Garcia (2022). While our baseline specification includes the Covid-19 period in the estimation sample, we have also tested the sensitivity of our modelling to ending the sample prior to Covid-19.

To pin down $\mathbf{Q}$, we draw orthogonal matrices $\mathbf{Q}$ using the method proposed by Arias et al. (2018). This ensures that the implied matrix $\mathbf{B}$ satisfies the zero restrictions. Sign restrictions are introduced via an accept/reject approach. To improve the efficiency of the algorithm, we then combine the approach by Arias et al. (2018) with the approach by Chan et al. (2025) and fully explore all possible orderings of the columns of $\mathbf{B}$, searching to see if any satisfy the sign restrictions. We found that this modification to Arias et al. (2018) delivers a significant computational gain, and significantly reduces the computational burden needed to generate accepted 10,000 draws.

3 Model Properties

In this section, we describe the standard SVAR model outputs that summarise the key properties of the model. We first present estimated impulse response functions before moving to forecast-error variance decompositions.

3.1 Impulse Responses

Figure 2 and Figure 3 report the estimated impulse responses for each identified structural shock in the model. These capture by how much, on average, each variable

[4] As shown by Bergholt et al. (2024), this prior helps to reduce the uncertainty around the deterministic component. The Minnesota prior specification centres the parameters according to a random walk, as is common for this type of data.

responds to a single structural shock, holding the other shocks at zero. The underlying shocks of interest are simulated to equal one standard deviation of the shock, which is normalised to unity. The red dots in the figure report the periods in which sign or zero restrictions were applied. The solid lines in each sub-figure report the pointwise median impulse response, with the shaded areas plotting the 68% and 90% credible sets. The estimated impulse responses highlight a number of key findings:

Global shocks. An expansionary global demand shock produces a large, positive, but relatively short-lived, impact on UK real GDP, and a significant positive response of UK CPI and Bank Rate. An expansionary world supply shock (exogenous to the world energy shock) generates a positive and more persistent impact on UK real GDP and a negative response of CPI. Moreover, the oil price variable, which is unrestricted, increases in response to the expansionary supply shock, therefore excluding the possibility that the positive world supply shock may include some spurious world energy shock. Finally, an expansionary world energy shock exerts a positive effect on both world and UK domestic activity and a negative effect on World CPI, UK CPI, CPI energy, and real Oil price. Bank Rate, which is left unconstrained, does not significantly respond to the shock. The results are quantitatively similar to Cesa-Bianchi et al. (2021).

Domestic shocks. An expansionary UK supply shock generates a persistent increase of the level of UK GDP and a fall in UK CPI and UK CPI energy. For the UK demand shock, the overall effect of the shock appears to be moderate on the domestic variables and not significant on the world variables. Finally, a UK monetary policy shock has a sizable impact on UK activity and UK CPI. In addition, it generates a negative response of CPI energy and World CPI, though not particularly significant. The results are qualitatively similar to Mountford (2005) and Braun et al. (2025).

3.2 Forecast-Error Variance Decompositions

Forecast-error variance decompositions highlight by how much, on average, the volatility of each of the variables in the dataset is driven by each shock. Figure 4 decomposes the mean of the sum of the decomposition into each identified shock (see Figure A-3 in the Appendix for the pointwise sum across shocks). The six identified shocks explain a substantial portion of the fluctuations in the variables included in our VAR: the share

Figure 2: Impulse responses to world shocks

A) World demand shock

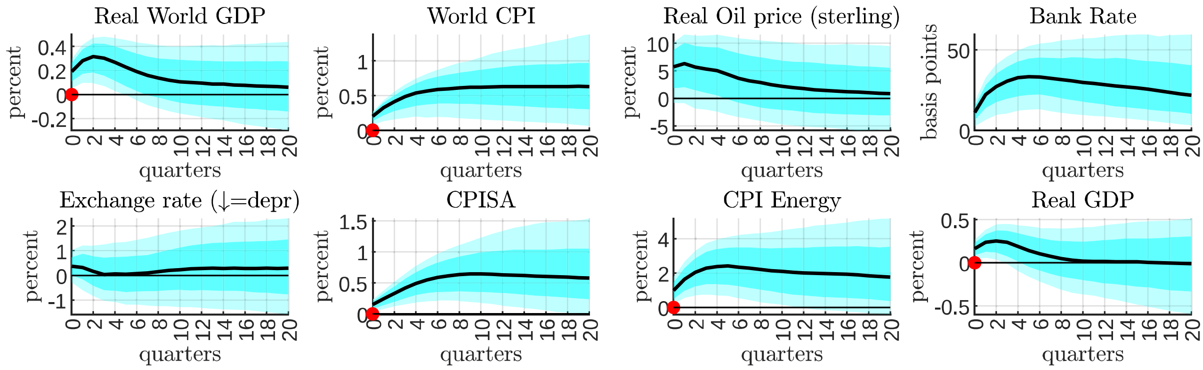

B) World supply shock

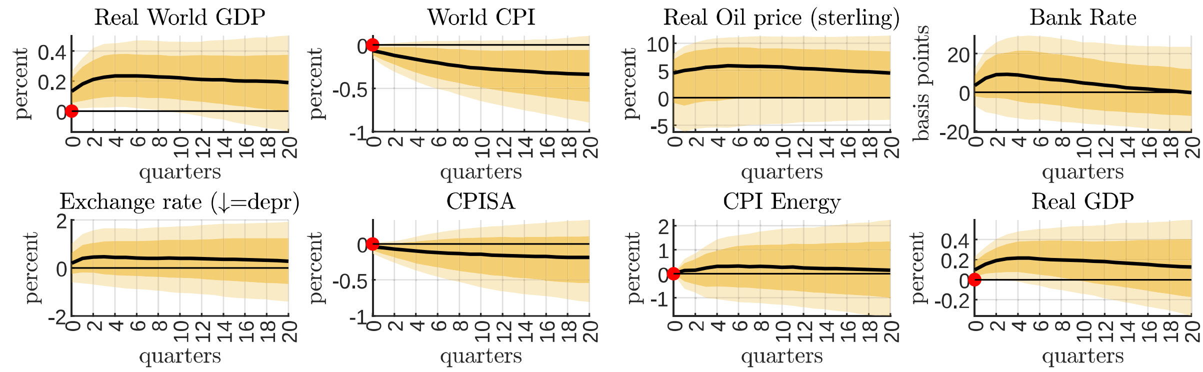

C) World energy shock

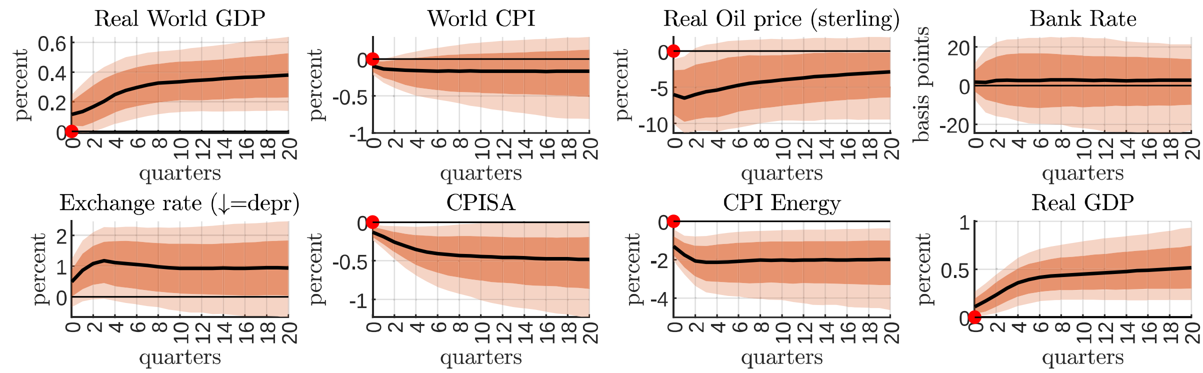

Note: Pointwise median, 68% and 90% credible set. Red dots indicate a period in which zero or sign restrictions were introduced.

Figure 3: Impulse responses to UK shocks

A) UK demand shock

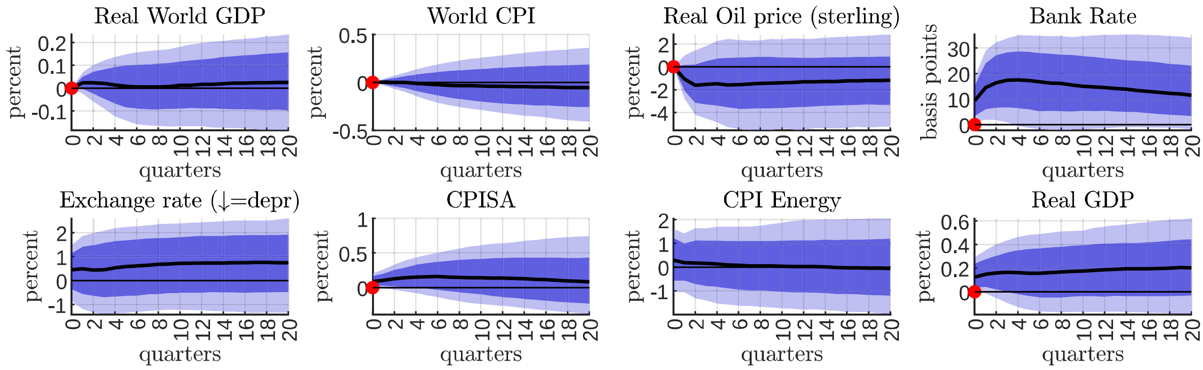

B) UK supply shock

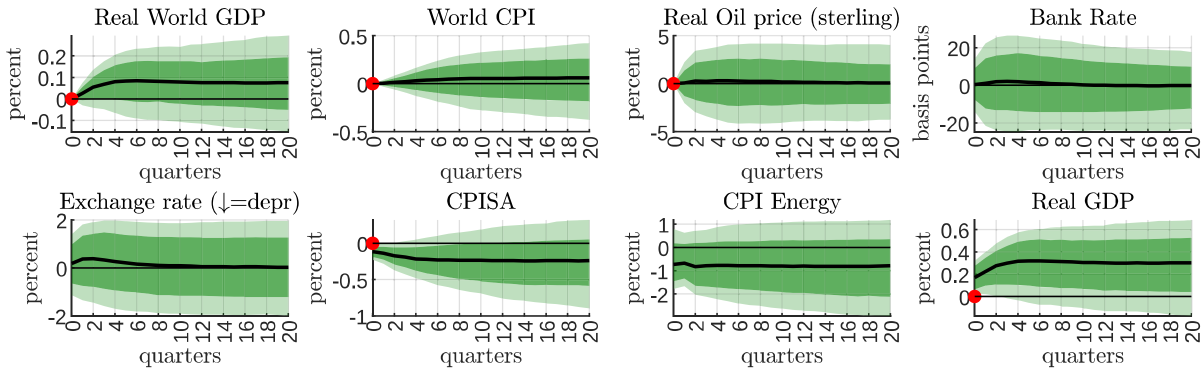

C) UK monetary policy shock

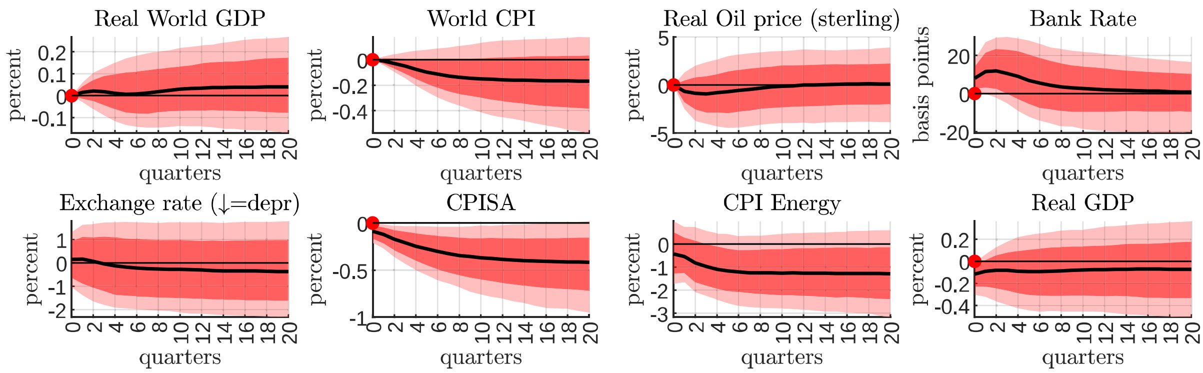

Note: Pointwise median, 68% and 90% credible set. Red dots indicate a period in which zero or sign restrictions were introduced

Figure 4: Forecast-error variance decompositions (mean)

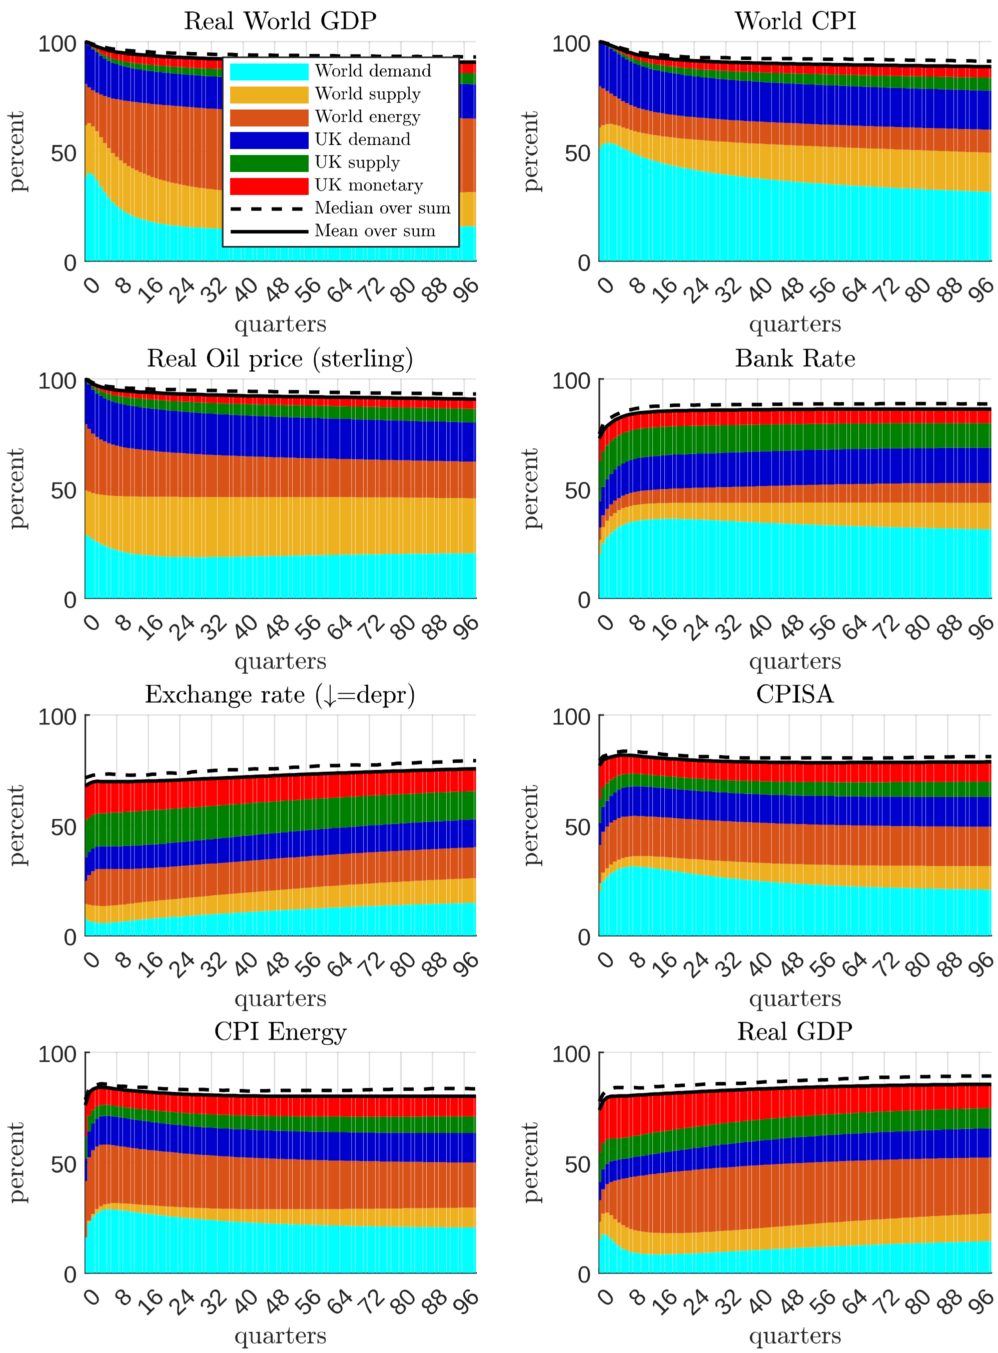

Note: Decomposition of the mean of the sum into the mean of the decompositions. The difference between the sum of the decomposition and 100 is due to the unidentified shocks.

of the explained variance is equal to around 80% for most of the domestic variables, while the corresponding figure for global variables is closer to 90%.

Global shocks are important drivers of the UK economy, with the world supply-side and demand shocks explaining around 40% of variation in real GDP variation and 50% of CPI variation one year after the shocks. The global demand shock has a greater role over the initial quarters, while the world supply-side shocks gain importance towards longer horizons.

Domestic shocks jointly explain an important share of the volatility of the two variables—around 40% for both CPI and real GDP. Of the former, the UK monetary policy shock appears to have the largest relative importance across the domestic shocks, with a smaller role left for the UK demand shock.

4 Explaining Fluctuations in Past Data

In this section, we briefly describe another standard output of the SVAR, that can be used to assess past data through the lens of the structural shocks in the model.

4.1 Estimated Shocks

While impulse responses and forecast-error variance decompositions capture average properties of the model, the actual estimates of the structural shocks $\boldsymbol{\epsilon}_t$ can begin to show how the model interprets the volatility of its endogenous variables at every period. Estimated time series for the six identified shocks are reported in Figure 5. Both panels in the figure report the pointwise median and corresponding credible sets for the six identified shocks, with panel A focusing on the pre-GFC period and panel B showing the post-GFC series. The dashed horizontal lines show two standard-deviation bounds on the shocks for reference, which complement the figure to provide a signal about the size of shocks. Together, the estimated shock series provide a narrative that is consistent with standard interpretation of key macroeconomic episodes:

Global shocks. The global demand shock series exhibits a marked drop in the second half of the 1990s around the time of the East-Asian crisis, in 2001 following 9/11 which marked a global pickup in uncertainty, and during the early stages of the GFC in 2007. The model also identifies positive global demand shocks over 2022 and 2023. The world supply shock (exogenous to world energy) series depicts strong positive

Figure 5: Estimated shocks

A) before the GFC

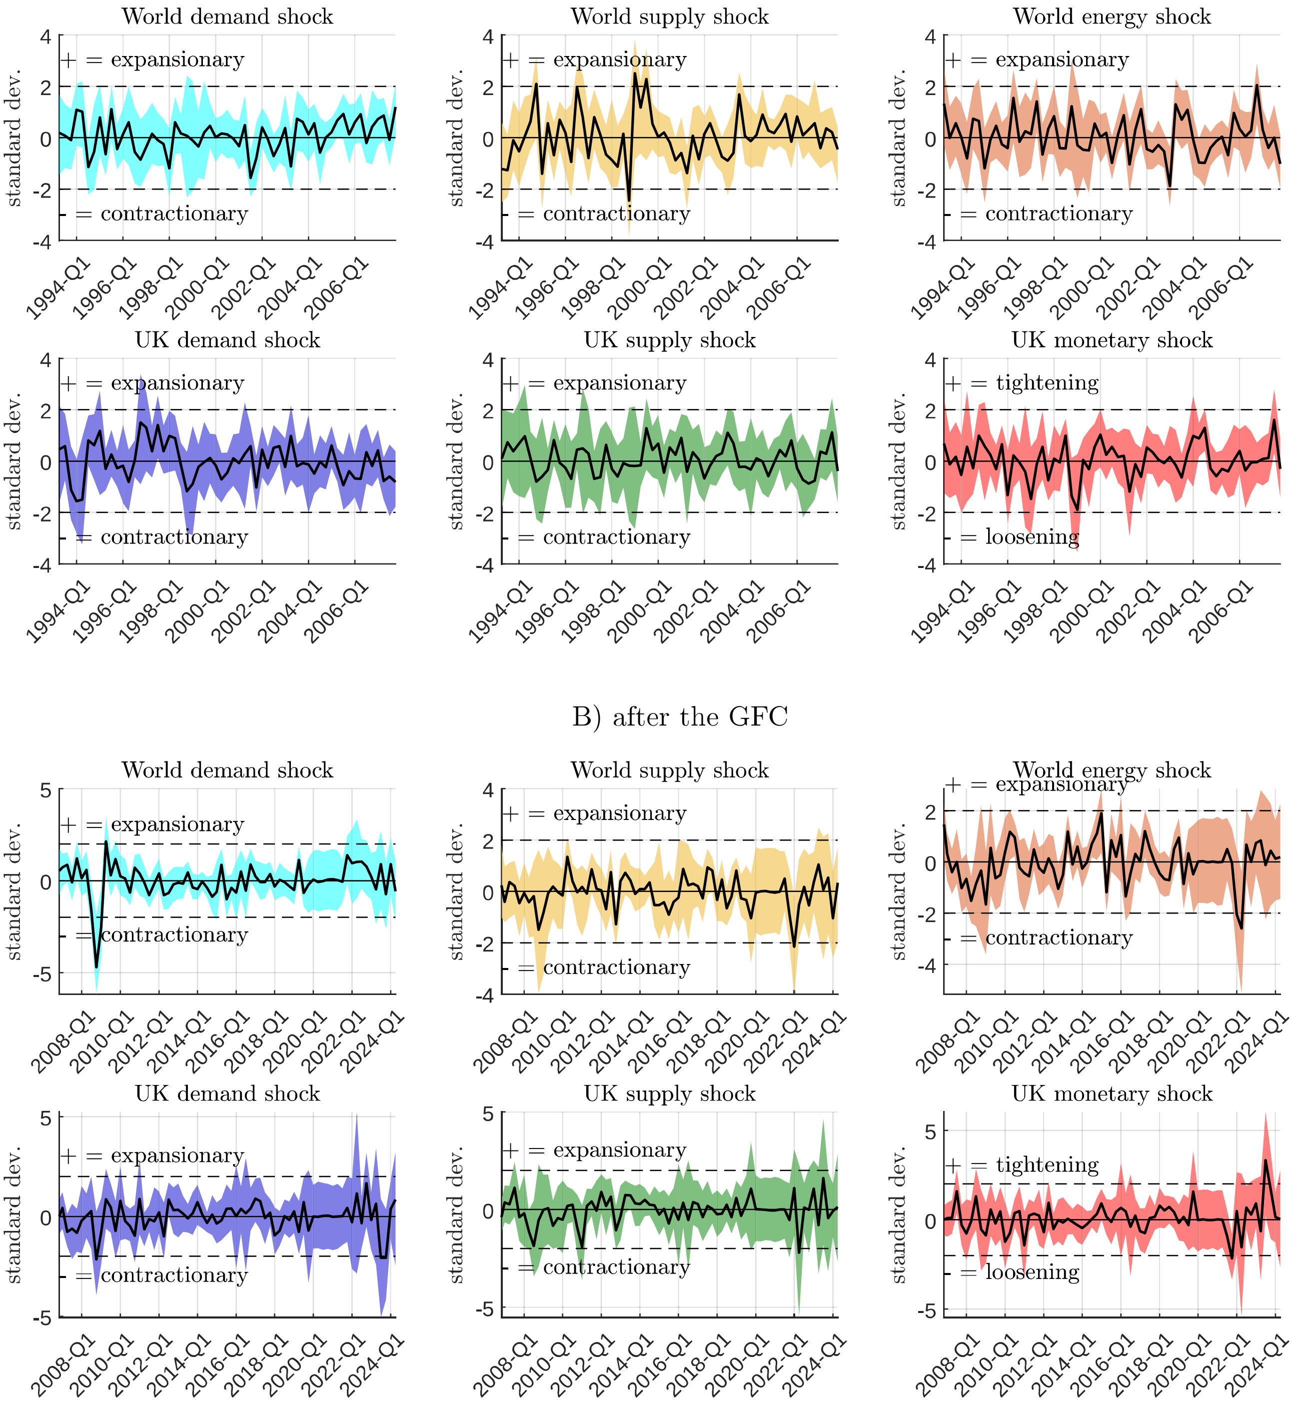

Note: Pointwise median of the estimated shocks as well as 90% credible sets. The shocks are standardised to have unit standard deviation. The figure indicates if a shock of positive sign is to be interpreted as contractionary or expansionary on UK GDP, in accordance with the signs of the impulse responses.

shocks at the end of the 1990s, which could reflect the diffusion of the internet on a world-wide scale. Conversely, 2022 is characterised by contractionary shocks, including a particularly large one in 2022Q1. Finally, the energy shock is especially negative during the Iraq invasion in 2003 and positive around 2014, which was characterised by low oil prices. The model also picks up marked contractionary shocks in 2022 following Russia's invasion of Ukraine. These results are consistent with Cesa-Bianchi and Stratford (2016), Forbes et al. (2018) and Cesa-Bianchi et al. (2021).

Domestic shocks. On the UK demand shock, the model signals a negative shock in 2008 following the GFC, and some positive shocks between 2013 and 2014. Moreover, the model also identifies marked negative UK demand shocks around 2023. The UK supply shock is overall more muted, although it is still important to highlight that the period between 2010 and 2017 is characterised by a series of negative shocks. The series of UK monetary policy shocks partially reflects the tightening and loosening cycles. For instance, the model predicts loosening shocks at the end of 1999 and around 2001, and tightening shocks during 2004 and 2007. Finally, the series also exhibits a very strong loosening shock post 2022, and a strong tightening shock during the most recent year. The results are qualitatively consistent with Mountford (2005).

It is important to remember that we have added Covid-19 dummies to our model, so the model will not use shocks to fully fit the volatility of the data over that specific period. For this reason, the shocks we identify appear relatively muted during 2020 and 2021.[5]

4.2 Historical Decompositions: Interpreting Data Outturns

We can combine the results described in the previous sections to obtain a decomposition of the time series included in the model. In other words, for each quarter, we can compute the contribution of a given shock to the series of interest, thus obtaining the historical decomposition.

Figure 6 depicts the historical decomposition of y-o-y UK CPI inflation in the left-hand panel and y-o-y growth real GDP in UK in the right-hand panel, both reported in deviations from the deterministic component associated with the initial condition and constant. The decomposition also reports the role played by the Covid-19 dummies in black bars and the portion of fluctuations due to the two unidentified shocks

[5]Nevertheless, the Bayesian Covid-19 dummy approach still allows for some volatility to be explained by the model.

Figure 6: Historical decompositions (deviation from deterministic component)

A) before the GFC

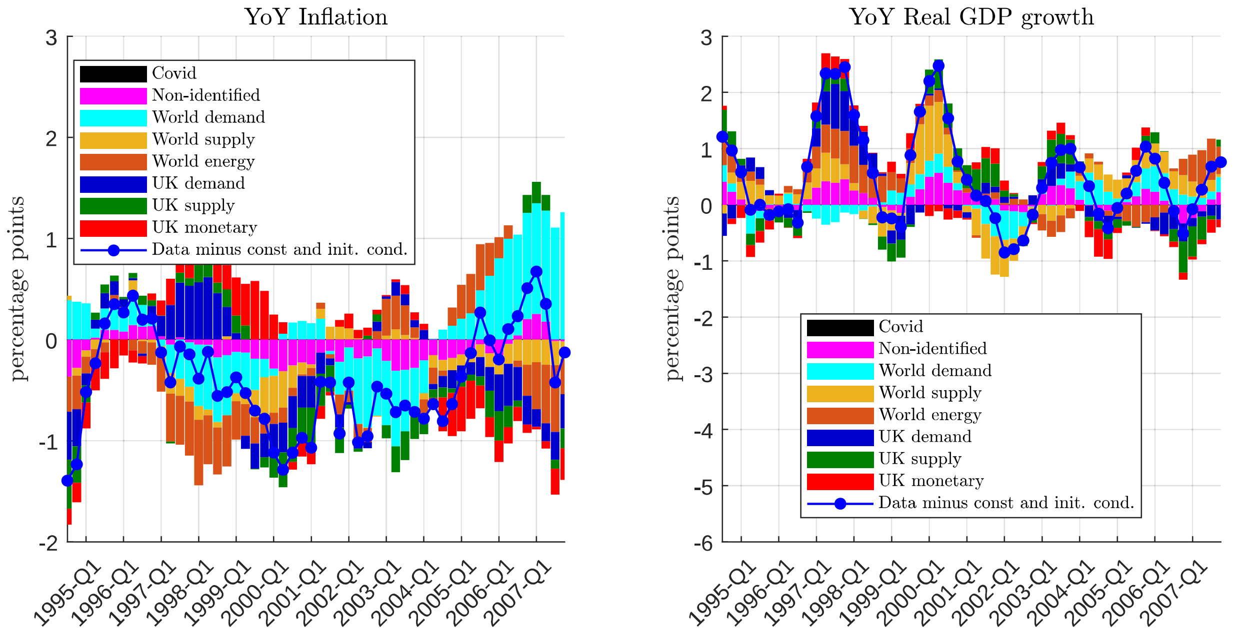

B) after the GFC

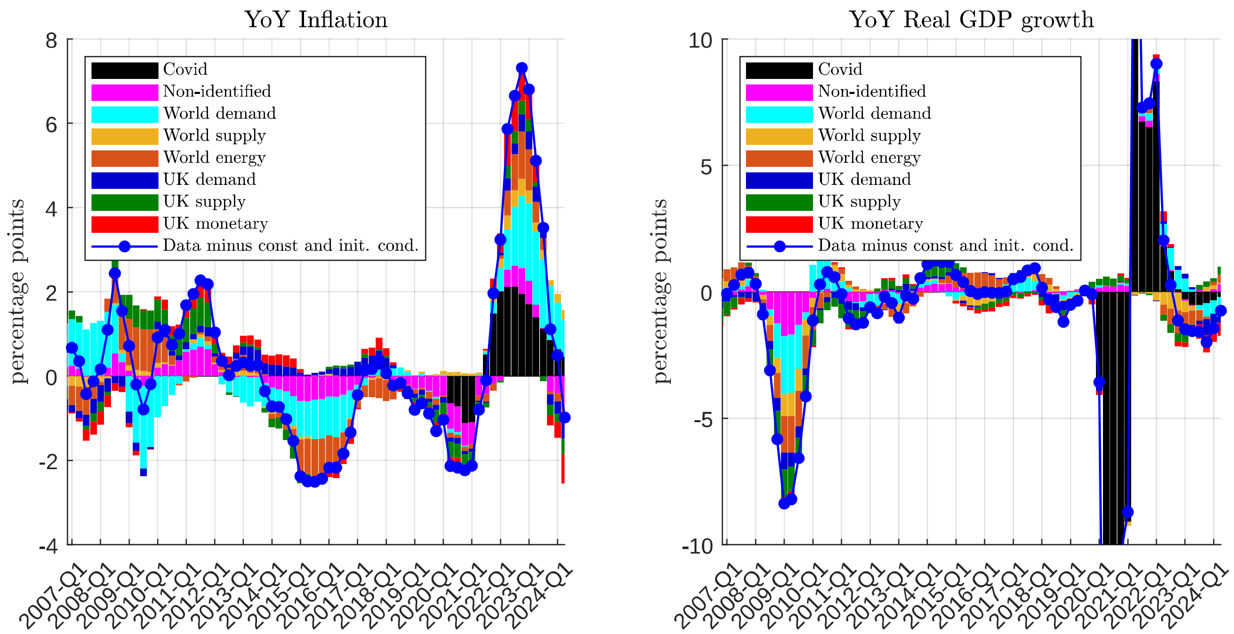

Note: The decomposition shown is associated with the pointwise mean decomposition over posterior draws. The decomposition is in deviation from the deterministic component, interpreted as the role associated with the constant term and the initial condition of the model . See Figure A-4 in the Appendix for the full decomposition, and Figure A-10 for the decomposition when ending the estimation sample before Covid-19.

in pink bars (here summed together). Although both components lack a 'structural' interpretation, they still convey valuable information. The Covid-19 dummies component helps to demonstrate the relative importance of the economic events around 2020 and 2021 and their long-lasting effects on the variables. The pink bars, on the other hand, can be seen as the net effect of all other shocks that are not captured by the identification scheme.

The historical decomposition helps to interpret historical events through the lens of structural shocks. For instance, strongly negative global demand shocks weighing on both GDP growth and CPI inflation in 2009, following the GFC. Moreover, during 2015, positive world energy shocks (orange bars) contribute positively to real GDP growth and negatively to inflation.

The model indicates that, in 2022, y-o-y UK CPI inflation was pushed up by all the shocks identified in the model. The SVAR identifies a particularly strong role for the global demand shock (aqua bars), along with some persistent effects from Covid-19 (black bars) that continued to positively contribute to inflation until the end of 2023. These effects were exacerbated by contractionary world energy (orange bars) and world supply shocks (yellow bars), and by the domestic shocks. UK monetary policy shocks turn negative in mid-2023, contributing to the decrease in inflation. Our finding about the main role played by the compounded demand shocks is consistent with many recent analyses conducted for the US and EA (see, e.g., Bergholt et al., 2024; Ascari et al., 2023; Giannone and Primiceri, 2024). Our results are also robust to an alternative specification shown in Figure A-10 in the Appendix, which depicts results obtained by stopping the estimation sample in 2019Q4 and letting that model interpret the volatility around the Covid-19 period.

Finally, the historical decomposition suggests that global demand shocks (aqua bars) contributed positively to y-o-y GDP growth over 2022, while the energy (orange bars) and world supply shocks (yellow bars) pushed down on growth. By 2023, UK monetary policy (red bars) and domestic demand (blue bars) shocks start to weigh on activity. Overall, it is important to stress that estimation uncertainty remains high around the decomposition of the contributors to both inflation and GDP around 2022.

5 Interpreting Forecast Revisions

The outputs explained so far are common in SVAR modeling. In this section, we describe a different and more novel application. Our goal is to form a narrative around the drivers of forecast revisions. This is a powerful informative tool that offers a decomposition of *marginal* changes in forecasts—compared to the absolute decomposition of data we described in Section 4.2—and implicitly provides a timely narrative for the latest shocks to hit the economy in the most recent quarter.

To fix ideas, we return to Figure 1 from the introduction. Our goal is to compare the forecast formed in May 2024 (time $T-1$) with that formed in the subseqent quarter, August 2024 (time T). For the former, we compute the unconditional forecast from the SVAR (dashed dotted line) using data for the estimation sample up to 2024Q1. For the latter, we compute the unconditional forecast using data up to 2024Q2. This is reported by the solid dotted line, along with the 68%/90% credible bands (shaded areas). As is evident from the figure, different forecasts can arise from the two different rounds. For example, between the two quarters, the forecast for output growth was revised upwards, while the forecast for inflation was left practically unchanged.

From a policy perspective, it important to establish which types of shocks *explain* these revisions, as this information may be important to interpret the state of the economy. Moreover, it is worthwhile understanding how these shocks develop over time and, more specifically, how they impact key variables of interest over the forecast horizon. Despite this, forecast revisions have typically been studied in reduced-form terms only, and often lack of a structural interpretation. Our goal is to bridge this gap and to provide a structural narrative to reduced-form forecast revisions by using our SVAR model.[6]

5.1 Decomposing Forecast Revisions: Methodology

We first discuss the foundations of our forecast-revision decomposition. Consider two unconditional forecasts constructed, in this case, with our SVAR model: one is built at time T, the other at time $T-1$. Overall, forecast revisions from $T-1$ to T can be explained by:

a) the fact that the forecast made at time $T-1$ assumed zero average shocks at time T. Yet, the realization of the data at time T allows for the estimation of the shocks that hit at time T, and hence for the assessment of the relevance of these shocks over the course of the forecast horizon;

[6] We are aware of only two attempts to use SVAR models to form a narrative over forecast revision, Todd (1992) and Giannone et al. (2004). The former takes a purely narrative approach to the model, while the latter limits the analysis to a less comprehensive set of identified shocks compared to the analysis of our paper.

b) potential data revisions in the data from $T-p+1$ to $T-1$. A SVAR forecasts by projecting ahead from the data in the latest p periods using the autoregressive nature of the model. Hence, data revisions imply that, if $p > 1$, the forecast made at time T might project ahead starting from observations that might differ from the ones used when forecasting at time $T-1$. Moreover, a revision in the data also leads to a revision in the the shocks estimated in previous rounds;

c) the fact that data revisions imply an update in the estimated coefficients of the model. Hence, the forecasts will be projecting ahead using an updated set of parameter values if it is re-estimated.

In summary, forecast revisions are generated because the economy is hit by shocks at time T that were predicted to be zero from the point of view of time $T-1$ under unconditional forecasting. These shocks generate effects over the forecast horizon, which vary depending on the nature of the shock that hits the economy. Moreover, forecast revisions are also caused by data revisions between time $T-1$ and T. This means that the new forecast builds on updated data, hence potentially using new parameter and shocks estimates, and projecting ahead from new data. We refer readers to Brignone and Piffer (2025) for further details, where we further formalize the above concepts.

5.1.1 Decomposing Forecast Revisions: An Application

To better understand the contribution of each component, we now examine the separate roles played by newly estimated shocks and data revisions in driving forecast revisions.

The role played by new shocks. Let us for now rule out the presence of any data revisions and focus only on the role played by the shocks that hit at time T. In other words, let us focus only on the (a), leaving (b) and (c) aside.

Figure 7 reports the marginal distribution of the six identified shocks estimated at time T. Put differently, the figure addresses the question 'what shocks have hit the economy at time T?', which is key to understanding how their effect will unfold over the forecast horizon. For instance, as shown in Figure 7, the model predicts with 88% posterior probability that at time T the economy was hit by a contractionary world demand shock, albeit a relatively small one (approximately one standard deviation), which generates contractionary pressures on output growth over the forecast horizon.

Figure 7: Distribution of shocks in period T

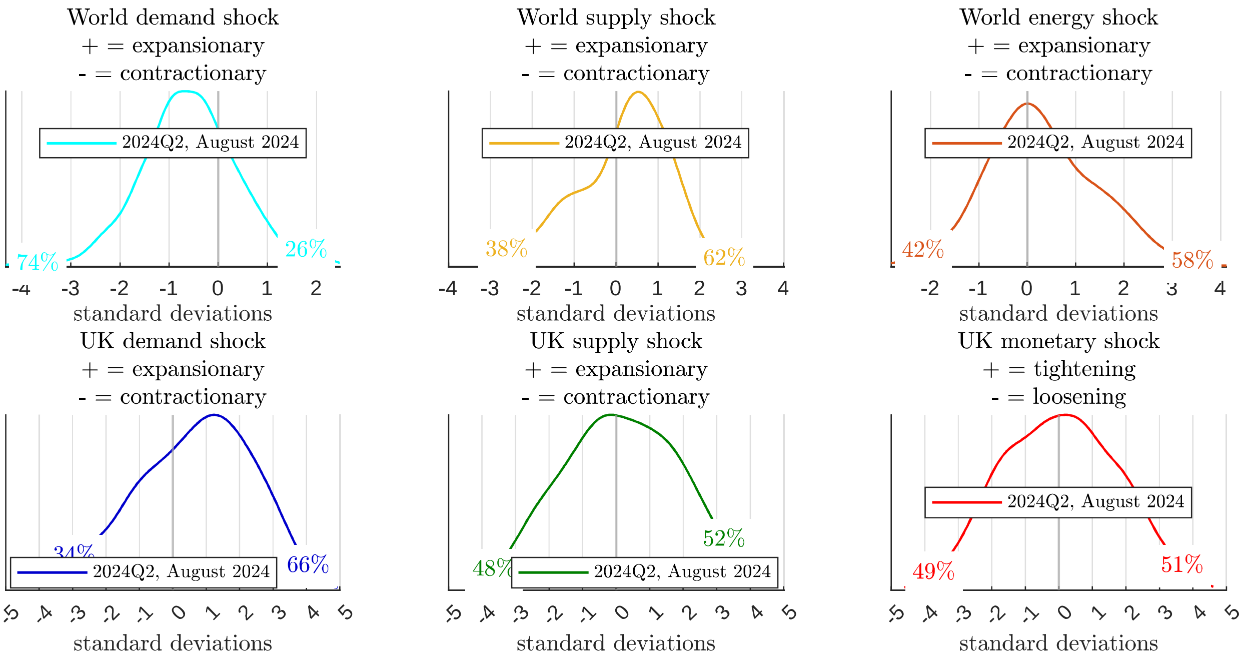

Note: The figure shows the estimate for 2024Q2 of the six structural shocks identified. The estimates are shown using a probabilistic distribution. The figure reports the probability that the estimated shock is positive or negative.

The model also estimates a small expansionary UK demand shock. There is, instead, more uncertainty around the sign of the remaining shocks.

To summarise the overall impact of the estimated shocks over the forecast, we propose the concept of '*composite impulse responses*'. This can help to inspect the role played by the latest estimated shocks—in this case, the shocks reported in Figure 7—on key variables over the forecast horizon.

The composite impulse responses associated with time-T shocks are shown in Figure 8, where we focus on the effects on inflation and output. Composite IRFs can be thought of as the weighted average of the estimated impulse responses, weighted by the estimated shocks. The top panel shows the joint effect of the shocks that hit the economy at time T, while the lower panel reports the decomposition into the structural shocks. It shows that the shocks at time T generate an overall neutral effect on inflation over all the forecast horizon. The bottom panel helps to understand why: the contractionary world demand shock, compounded with a slightly expansionary world energy shock, is counteracted by the remaining shocks, mainly a positive UK demand

Figure 8: Composite impulse responses associated with the shocks from period T

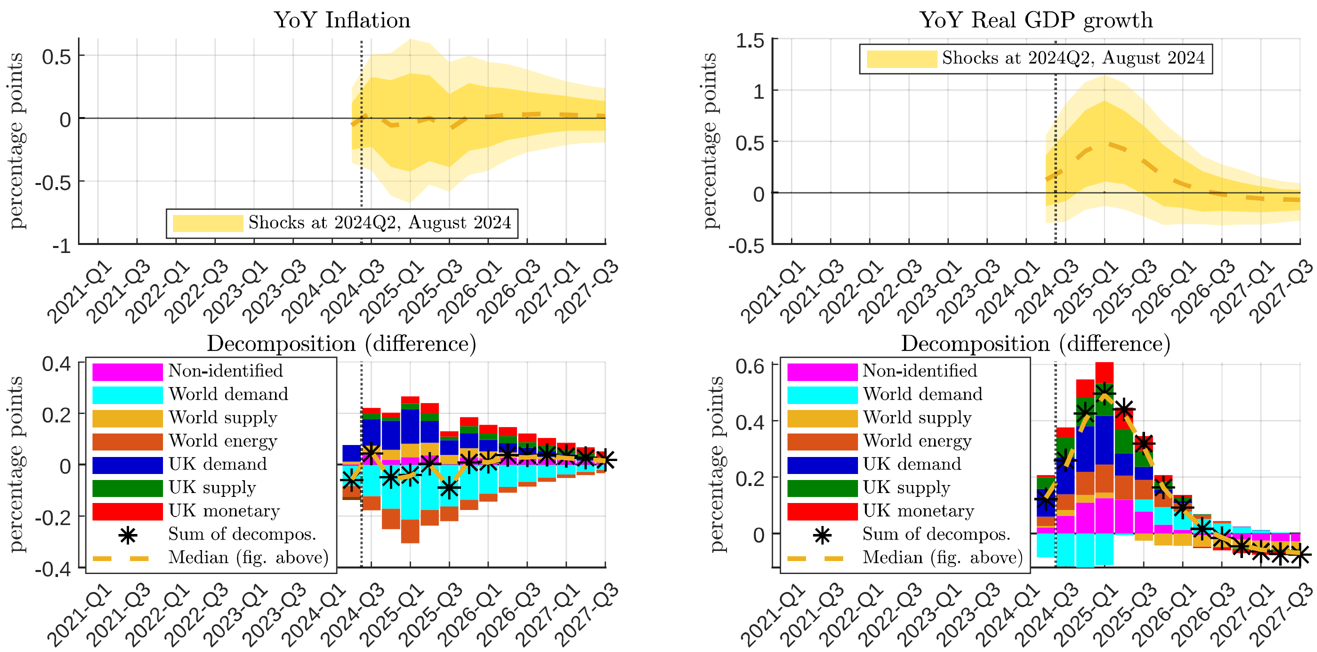

Note: The top panel shows the joint effect that the shocks estimated in 2024Q2 have on y-o-y inflation and real GDP growth. We report the pointwise median and 68%/90% credible intervals. The bottom panel shows the decomposition of the median response documented in the top panel.

shock. For real GDP, the shocks are instead jointly expansionary on output, with a peak effect to be expected in 2025Q1.

Interestingly, the composite impulse responses are deeply related to forecast revisions. As shown by Figure 9, the unconditional forecast made by the model at time T *coincides* with the sum of: 1) the unconditional forecast made at time $T-1$ and 2) the composite impulse responses associated with the shocks at time T (see Brignone and Piffer, 2025 for the derivations)—provided the absence of data revisions and estimation uncertainty. We show this in Figure 9. This figure plots the unconditional forecast produced at time T (orange solid line), along with a 'pseudo' $T-1$ forecast (gray dashed lines). The latter is 'pseudo' as it is generated by projecting the $T-1$ forecast ahead from the revised observations (excluding the observation at time T), practically shutting down the role of data revisions between time T and $T-1$.[7]

Ideally, the sum of the 'pseudo' $T-1$ forecast and the composite IRFs shown in

[7] Moreover, the 'pseudo' $T-1$ forecast was produced also by using the same parameter estimates used to generate the time T forecast.

Figure 9: Forecast revision: composite IRFs from time T

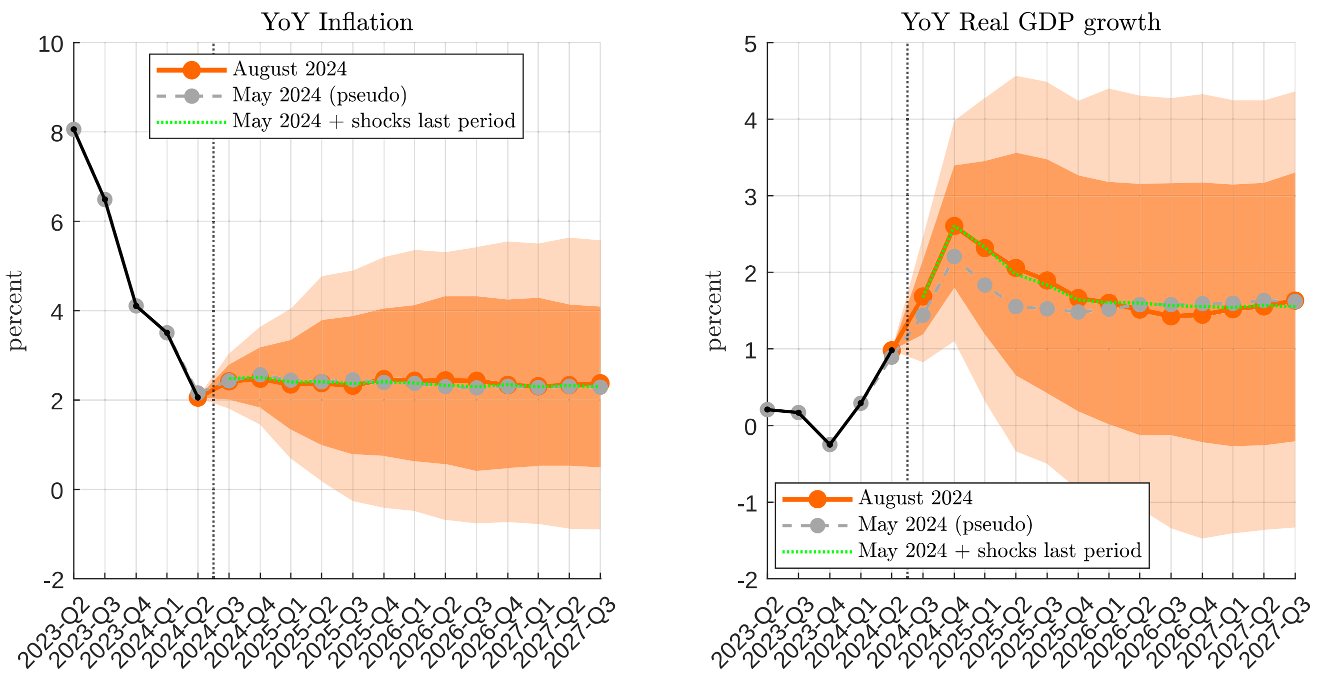

Note: The August 2024 forecast is shown as pointwise median, and 68% and 90% credible band. The pseudo May 2024 forecast, shown only via pointwise median, was computed using August 2024 parameter estimates and projecting ahead using revised data.

Figure 8 should exactly give us the new forecast produced at time T. This is what we obtain: the green line is the sum of the 'pseudo' forecast and the composite IRFs, and it is (approximately) equal to the solid orange line. This confirms that the shocks that hit the economy at time T do contribute to explaining forecast revisions, and that, thanks to the composite IRFs, we have a structural narrative of why forecast was revised.

The role played by data revisions. We now move the analysis one step forward and study the role played by data revisions, encompassing a case in which we also include (b) and (c) described earlier.

When we include data revisions, a first natural outcome is a re-assessment of the shocks previously estimated by the model. This is shown in Figure 10, which compares the time series of shocks estimated at at time T (shaded area and solid line) with the same shocks estimated at time $T-1$ (dashed line).[8] The model updates the

[8]The actual data revision is shown in Figure A-5 of the Appendix. See also Figure A-7-Figure A-7

Figure 10: Revision in estimated shocks for last year

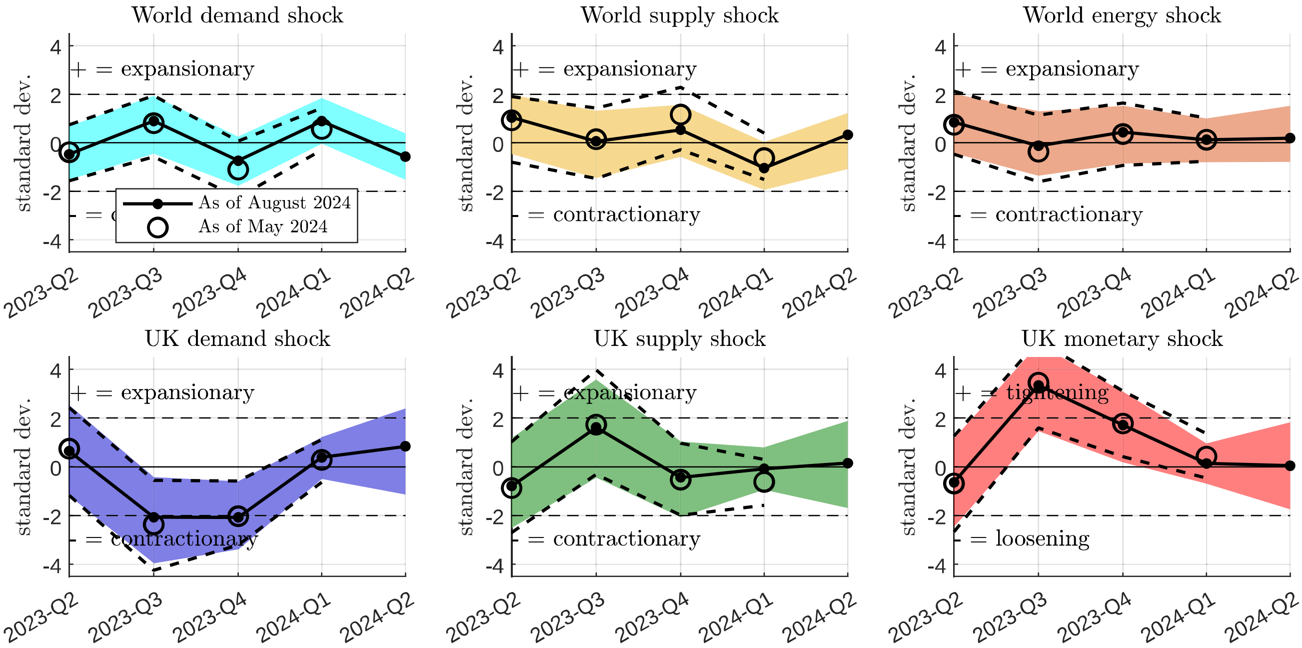

Note: Black line: Series of the shocks estimated at time T which corresponds to the August 2024 round, plotted with the 68% credible intervals. Dashed line: Series of the shocks estimated at time *T-1* which corresponds to the May 2024 round. The shocks are standardised to have unit standard deviation. The figure indicates if a shock of positive sign is to be interpreted as contractionary or expansionary on UK GDP, in accordance with the signs of the impulse responses.

estimates by finding that the world demand shock and the UK supply shock were less contractionary than initially predicted, while the world supply shock is found to be less expansionary than estimated initially.

To recap, the latest forecast made at time T projects ahead (and estimates the parameters) using revised data. These data revisions produce an updated estimate of the structural shocks previously estimated, which in turn will feed into the forecast, thus generating a revision with respect to the forecast produced at time $t-1$. Therefore, similarly to what was shown before, forecast revisions generated by data revisions can also be seen through the lens of the composite IRFs. Figure 11 reports the composite impulse responses estimated for the shocks at time $T-1$. The top panel shows the pointwise median and 68%/90% credible bands associated with the estimates made at time T (shaded areas and dashed line) and at time $T-1$ (green

for the full distribution of the revision of the shocks for selected quarters.

Figure 11: Effect of data revision on shocks from period $T-1$

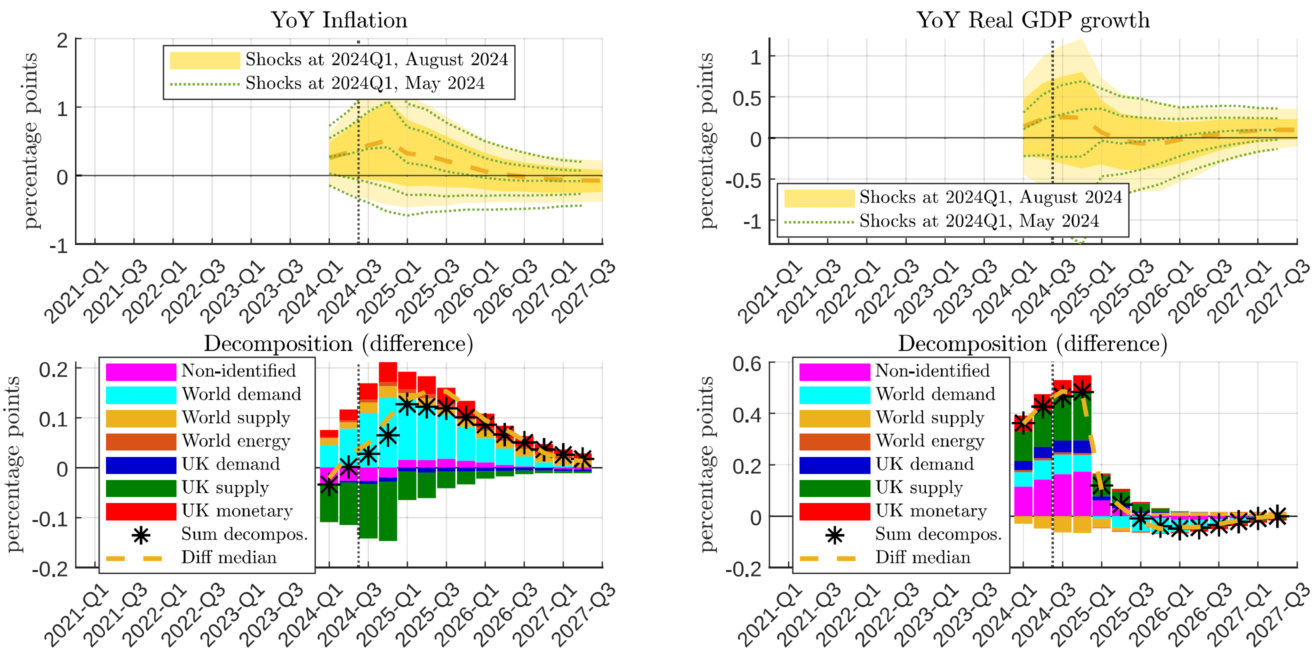

Note: The top panel shows the joint effect that the shocks in 2024Q2 have on y-o-y inflation and real GDP growth, according to estimates from either 2024Q2 (May 2024) or 2024Q3 (August 2024). We report the pointwise median and 68%/90% credible intervals. The bottom panel shows the difference in the decompositions of the median responses documented in the top panel (August 2024 minus May 2024).

dotted line). While the $T-1$ shocks had initially been estimated to jointly have a neutral/mildy contractionary effect on output, they are then found to have an overall expansionary effect on output, once re-estimated at time T. The lower panel shows instead the difference between the dashed line and the green dotted line decomposed into the different shocks, and tell us the role of the revised shocks due to data revisions in driving forecast revisions. Of course, this analysis can be extended to every previous quarter, see Figure A-8 and Figure A-9.

Summary. Finally, we bring everything together in Figure 12, which combines the ingredients of the analysis explained so far to jointly study the role of both data revisions and latest shocks in driving forecast revisions. The solid orange line refers to the unconditional forecast made at time T, and coincides with that shown in Figure 9. The dashed orange line refers to the forecast made at time $T-1$ using the data available at time $T-1$. As argued in the above paragraphs, the difference between the

Figure 12: Forecast revision: full decomposition

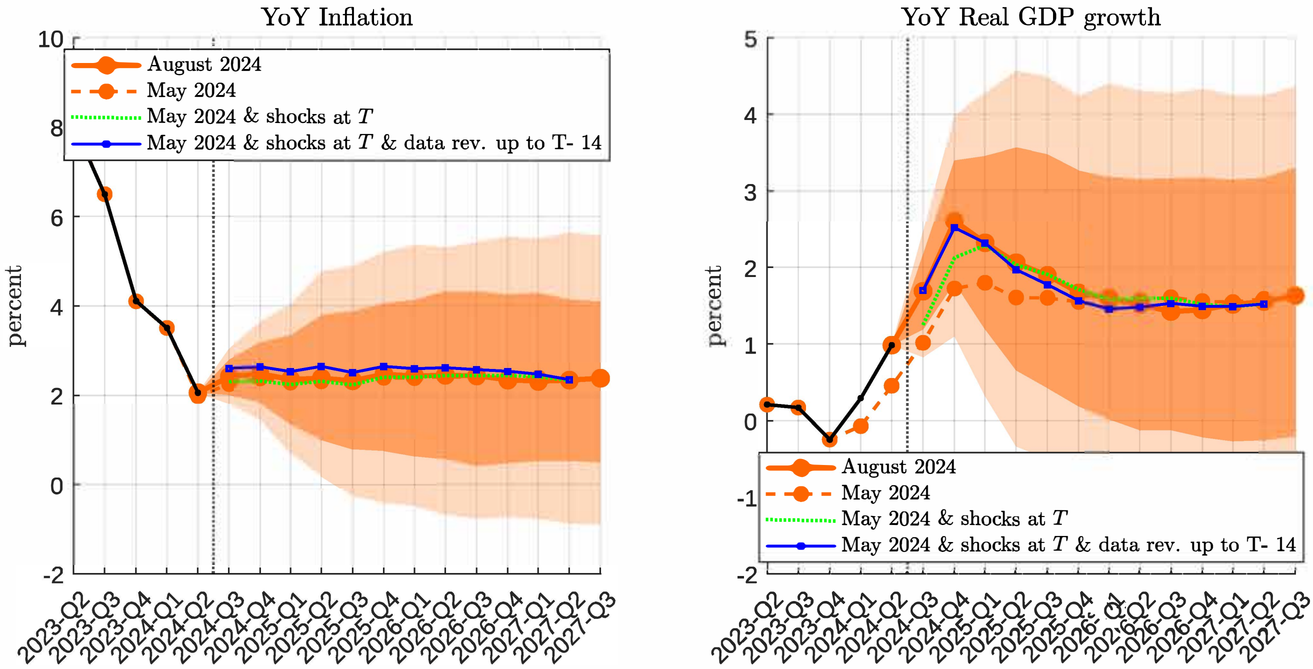

Note: The August 2024 forecast is shown as pointwise median, and 68% and 90% credible band. The May 2024 forecast is shown only via pointwise median.

dashed and solid orange lines can be attributed to: 1) the newly estimated shocks at time T and 2) data revisions. The contribution of the former is depicted by the green dotted line, which is the sum of the $t-1$ forecast and the composite IRFs of the shocks estimated at time T. This explains around half of the overall forecast revision. The role played by data revision is instead shown by the difference between the green dotted line and the blue squared line, which adds the composite impulse responses associated with up to 15 quarters before time T. As described in this section, including the composite IRFs for previous quarters accounts for the role of data revisions.

6 Conclusions

In this paper, we presented a SVAR model for the UK economy, which includes key global and UK variables and identifies demand and supply shocks on both the global and domestic side. The model is intended for use on a quarterly basis, and complements the broader modelling toolkit of the Bank used to inform policy discussions. It serves multiple purposes: cross-checking results from theoretical models, providing structural narratives to explain historical fluctuations in macroeconomic variables and to interpret forecast revisions, and delivering timely estimates of recent shocks affecting the UK economy.

This model should not be viewed as final, but rather as a platform for further development. Its flexibility facilitates the construction of alternative specifications tailored to the evolving challenges faced by policymakers in a highly uncertain economic environment. To this end, ongoing work by Bank staff is focused on expanding the global dimension of the model to better capture the effects of a broader range of international shocks on the UK economy, and to build outputs that can better inform policymakers—including those for structural scenario analysis.

References

Albuquerque, D., Chan, J., Kanngiesser, D., Latto, D., Lloyd, S., Singh, S. and Žáček, J. (2025), 'Decompositions, forecasts and scenarios from an estimated DSGE model for the UK economy', *Bank of England Macro Technical Paper* (1).

Arias, J. E., Rubio-Ramírez, J. F. and Waggoner, D. F. (2018), 'Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications', *Econometrica* **86**(2), 685–720.

Ascari, G., Bonomolo, P., Hoeberichts, M. and Trezzi, R. (2023), 'The Euro Area great inflation surge', *Policy Brief 548, SUERF* .

Bańbura, M., Bobeica, E. and Martínez Hernández, C. (2023), *What drives core inflation? The role of supply shocks*, number 2875, ECB Working Paper.

Baumeister, C. and Hamilton, J. D. (2015), 'Sign restrictions, structural vector autoregressions, and useful prior information', *Econometrica* **83**(5), 1963–1999.

Bergholt, D., Canova, F., Furlanetto, F., Maffei-Faccioli, N. and Ulvedal, P. (2024), *What drives the recent surge in inflation? The historical decomposition roller coaster*, number 7/2024, Working Paper.

Brandt, L. and Burr, N. (2024), 'A new medium-scale proxy-SVAR for the UK economy'.

Braun, R., Miranda-Agrippino, S. and Saha, T. (2025), 'Measuring monetary policy in the UK: The UK monetary policy event-study database', *Journal of Monetary Economics* **149**, 103645.

Breeden, S. (2025), 'Reading between the lines'.

Brignone, D. and Piffer, M. (2025), 'Structural SVAR-based interpretations of reduced form forecast revisions'.

Canova, F. (2011), *Methods for applied macroeconomic research*, Princeton university press.

Canova, F., Kociecki, A. and Piffer, M. (2024), 'Flexible prior beliefs on impulse responses in Bayesian vector autoregressive models'.

Cascaldi-Garcia, D. (2022), 'Pandemic priors'.

Cesa-Bianchi, A., Dickinson, R., Kösem, S., Lloyd, S. and Manuel, E. (2021), 'No economy is an island: how foreign shocks affect UK macrofinancial stability.', *Bank of England Quarterly Bulletin* .

Cesa-Bianchi, A. and Stratford, K. (2016), 'How could a shock to growth in China affect growth in the United Kingdom?', *Bank of England Quarterly Bulletin* p. Q1.

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

Chan, J., Matthes, C. and Yu, X. (2025), 'Large structural VARs with multiple sign and ranking restrictions', *arXiv preprint arXiv:2503.20668* .

Christiano, L. J. (2012), 'Christopher A. Sims and vector autoregressions', *The Scandinavian Journal of Economics* **114**(4), 1082–1104.

Ciccarelli, M., Darracq Pariès, M. and Priftis, R. (2024), *ECB macroeconometric models for forecasting and policy analysis: development, current practices and prospective challenges*, number 344, ECB Occasional Paper.

Crump, R. K., Eusepi, S., Giannone, D., Qian, E. and Sbordonea, A. (2025), 'A large Bayesian VAR of the us economy', *International Journal of Central Banking* **21**(2), 351–409.

Forbes, K., Hjortsoe, I. and Nenova, T. (2018), 'The shocks matter: improving our estimates of exchange rate pass-through', *Journal of international economics* **114**, 255–275.

Giannone, D., Lenza, M. and Primiceri, G. E. (2015), 'Prior selection for vector autoregressions', *Review of Economics and Statistics* **97**(2), 436–451.

Giannone, D. and Primiceri, G. (2024), The Drivers of Post-Pandemic Inflation, NBER Working Papers 32859, National Bureau of Economic Research, Inc.

Giannone, D., Reichlin, L. and Sala, L. (2004), 'Monetary policy in real time', *NBER macroeconomics annual* **19**, 161–200.

Kilian, L. and Lütkepohl, H. (2017), *Structural vector autoregressive analysis*, Cambridge University Press.

Mountford, A. (2005), 'Leaning into the wind: a structural VAR investigation of UK monetary policy', *Oxford Bulletin of Economics and Statistics* **67**(5), 597–621.

Pill, H. (2024), 'Cross-checking'.

Sims, C. A. (1980), 'Macroeconomics and reality', *Econometrica: journal of the Econometric Society* pp. 1–48.

Taylor, A. (2025), 'The last half mile'.

Todd, R. M. (1992), 'Algorithms for explaining forecast revisions', *Journal of Forecasting* **11**(8), 675–685.

Uhlig, H. (2005), 'What are the effects of monetary policy on output? results from an agnostic identification procedure', *Journal of Monetary Economics* **52**(2), 381–419.

Online Appendix for
"A Structural VAR model for the UK economy"

A Additional figures

Figure A-1: Data, transformed as it enters the model

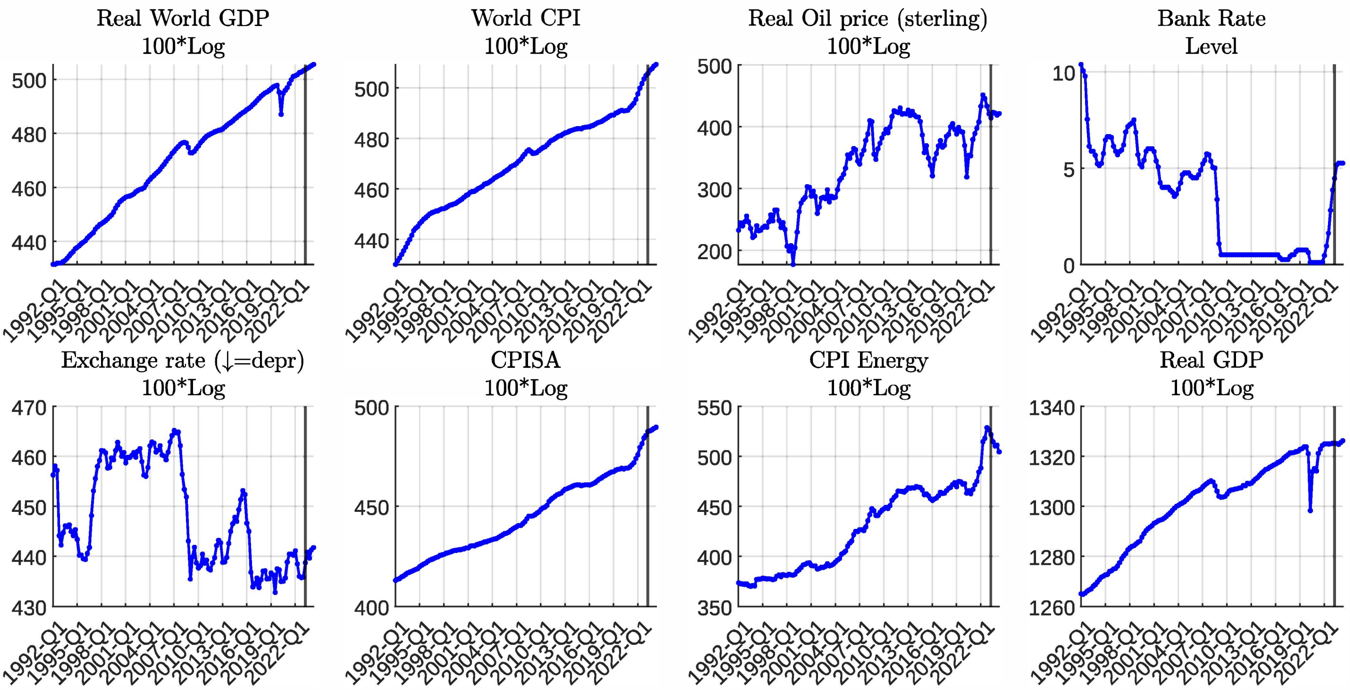

Note: Vertical line shows the time at which the estimation sample ends.

Figure A-2: Data, transformed as for selected figures

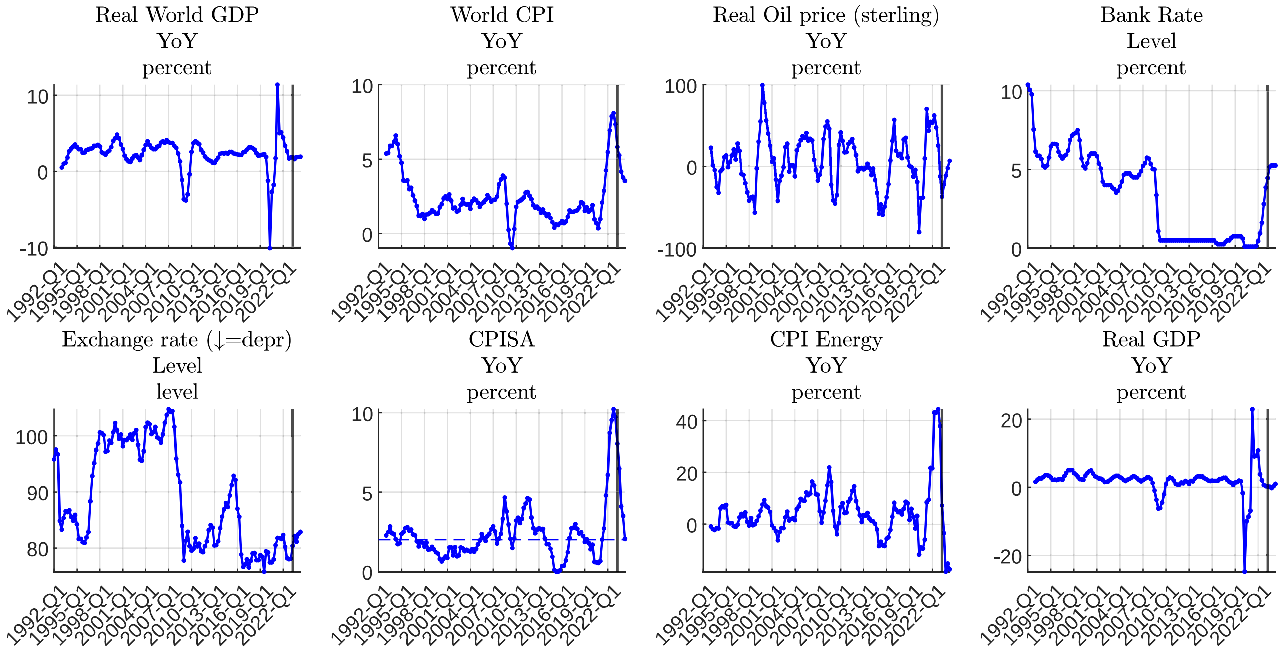

Note: Figure shows data after the transformation necessary to convert the data into a scale more informative to policymakers. The data enters the estimated model as shown in Figure A-1. Vertical line shows the time at which the estimation sample ends.

Figure A-3: Total forecast-error variance for each variable

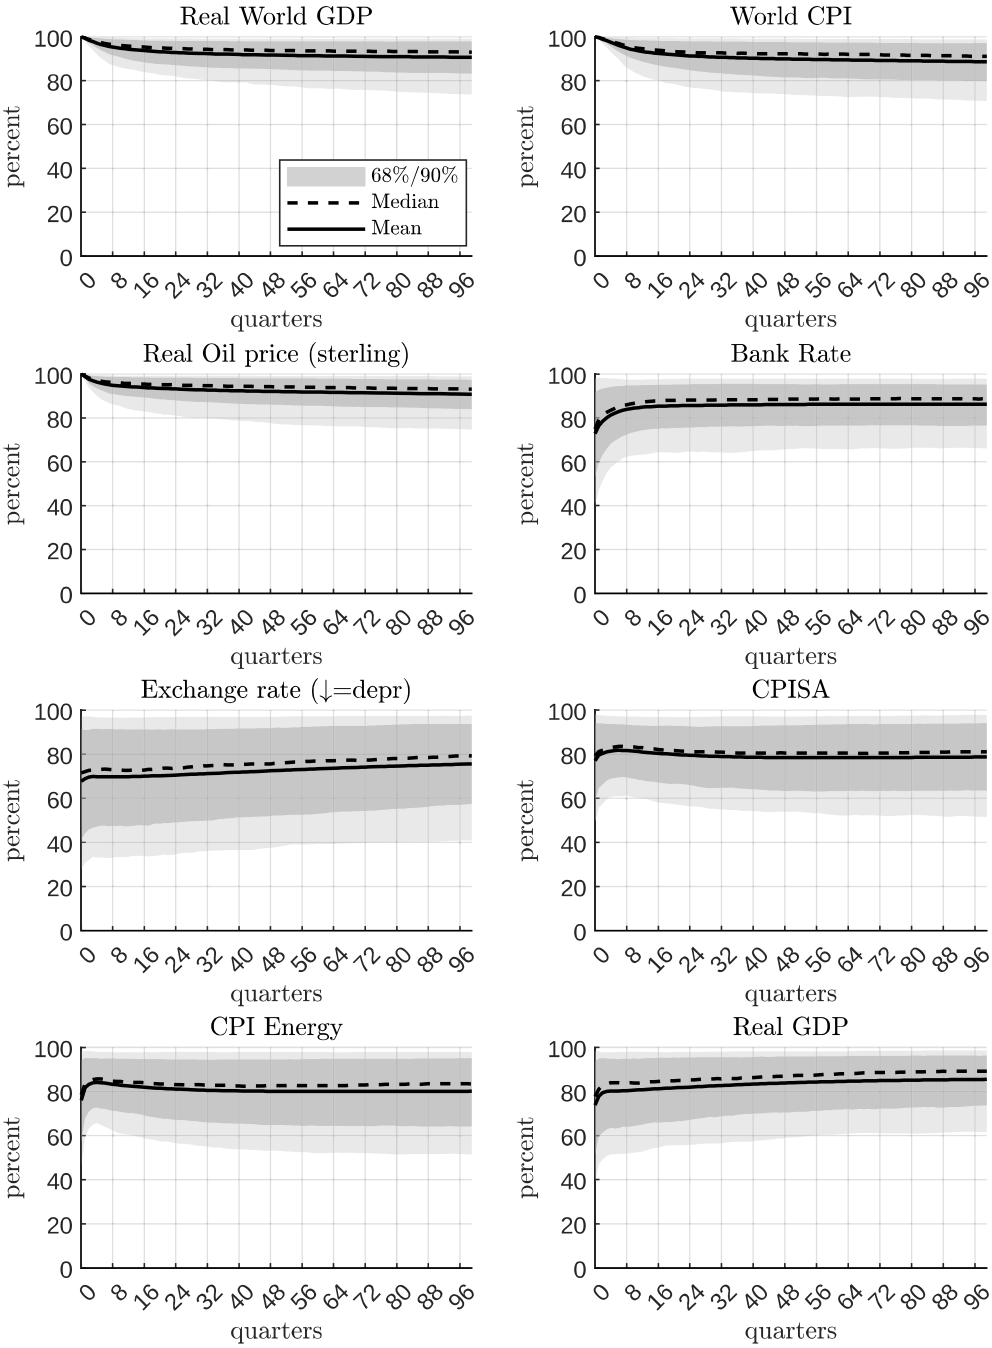

Note: Pointwise median, mean and credible sets over the sum of forecast-error variance decompositions. The difference between the sum of the decomposition and 100 is due to the unidentified shocks.

Figure A-4: Historical decompositions (full decomposition)

A) before the GFC

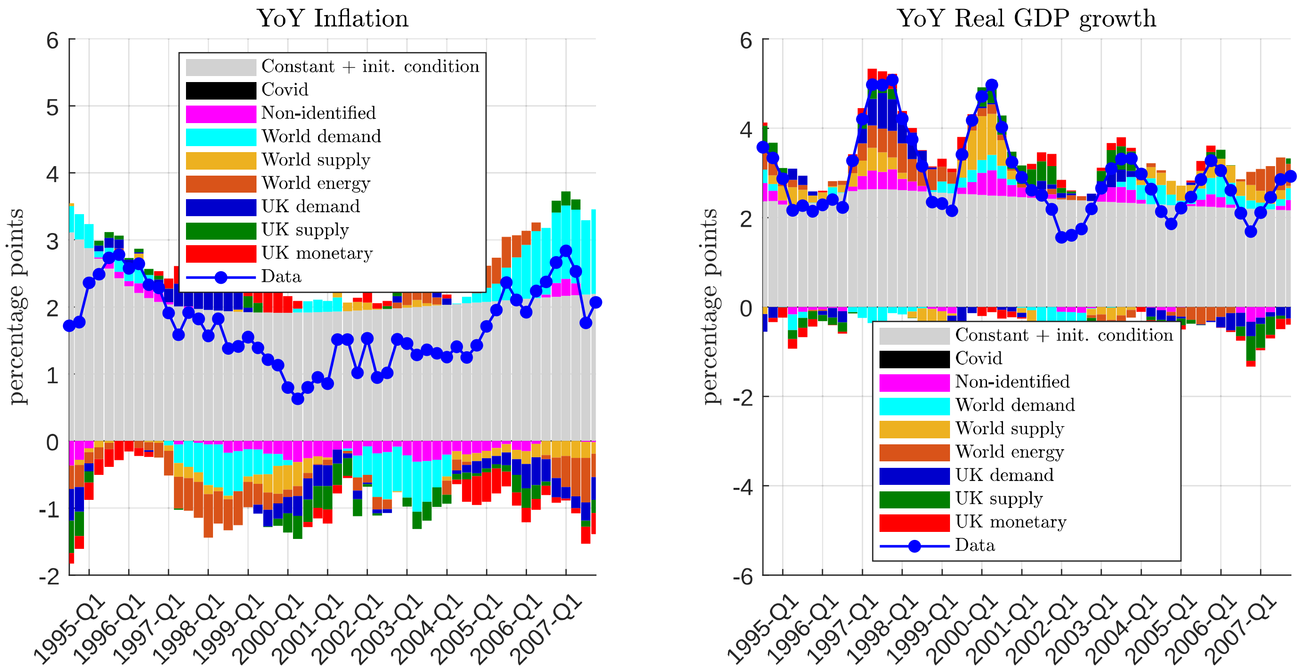

B) after the GFC

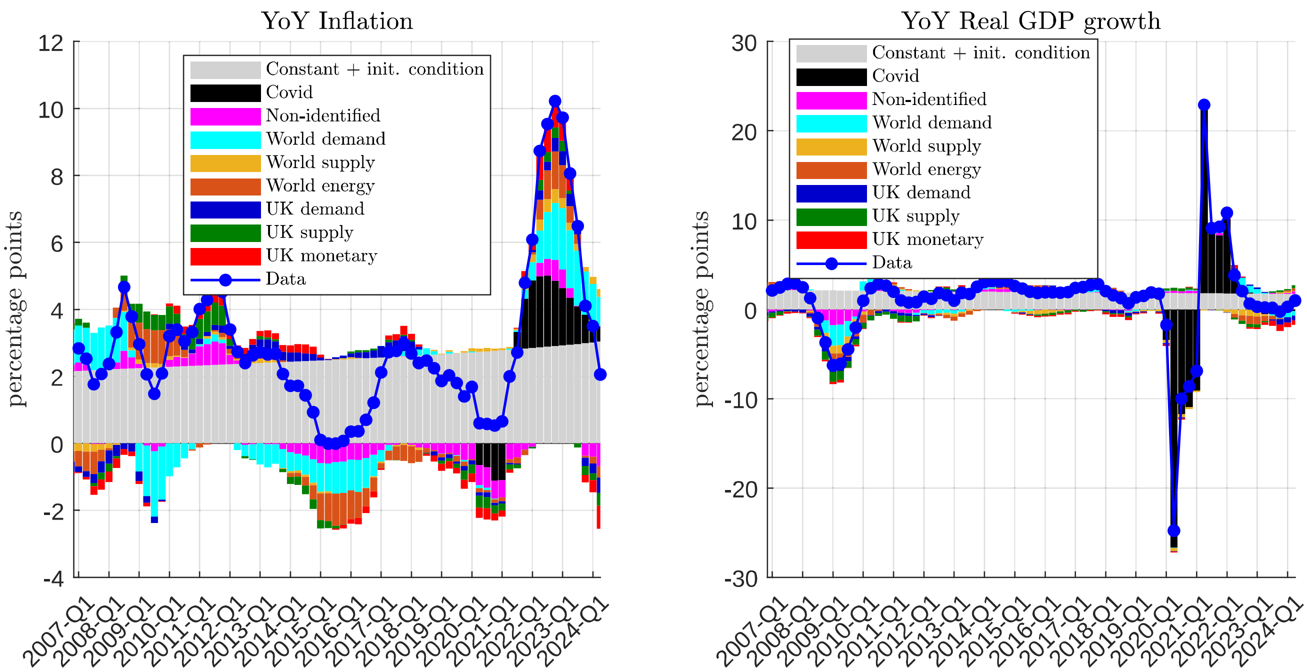

Note: The decomposition shown is associated with the pointwise mean decomposition over posterior draws.

Figure A-5: Data after transformation, revision in the latest period

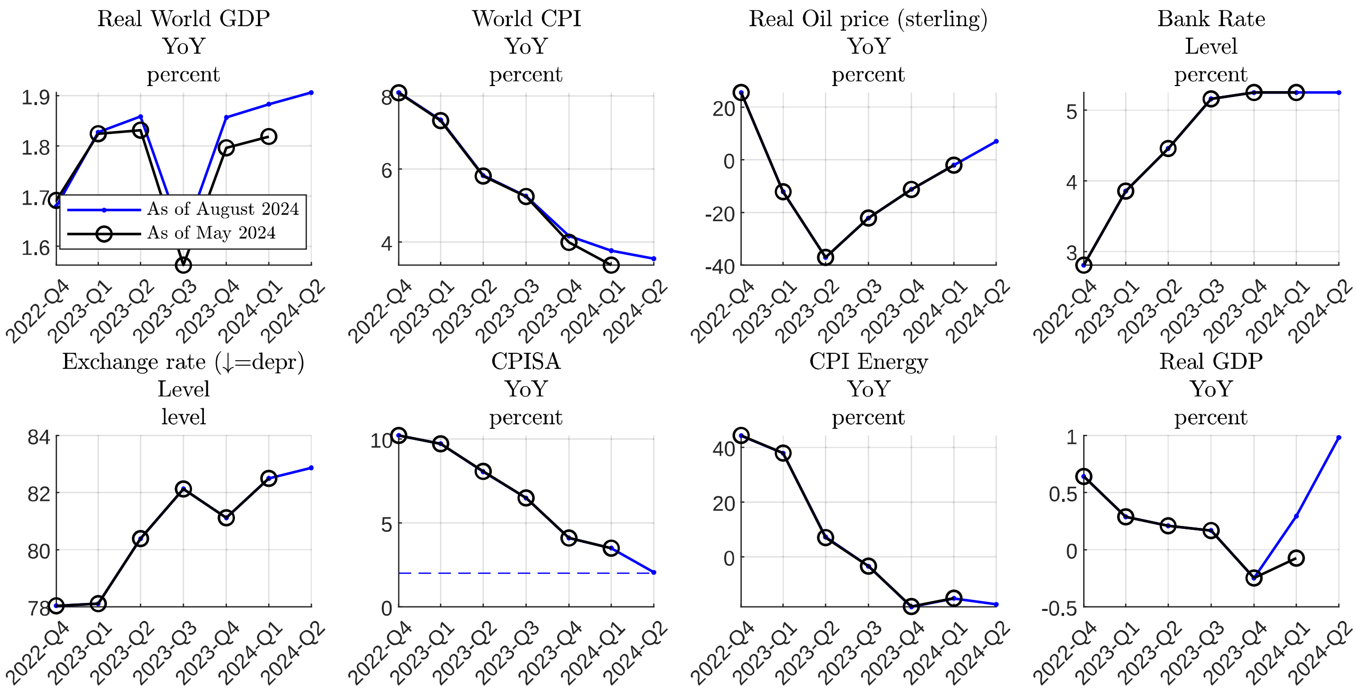

Figure A-6: Revision in estimated shocks (1/2)

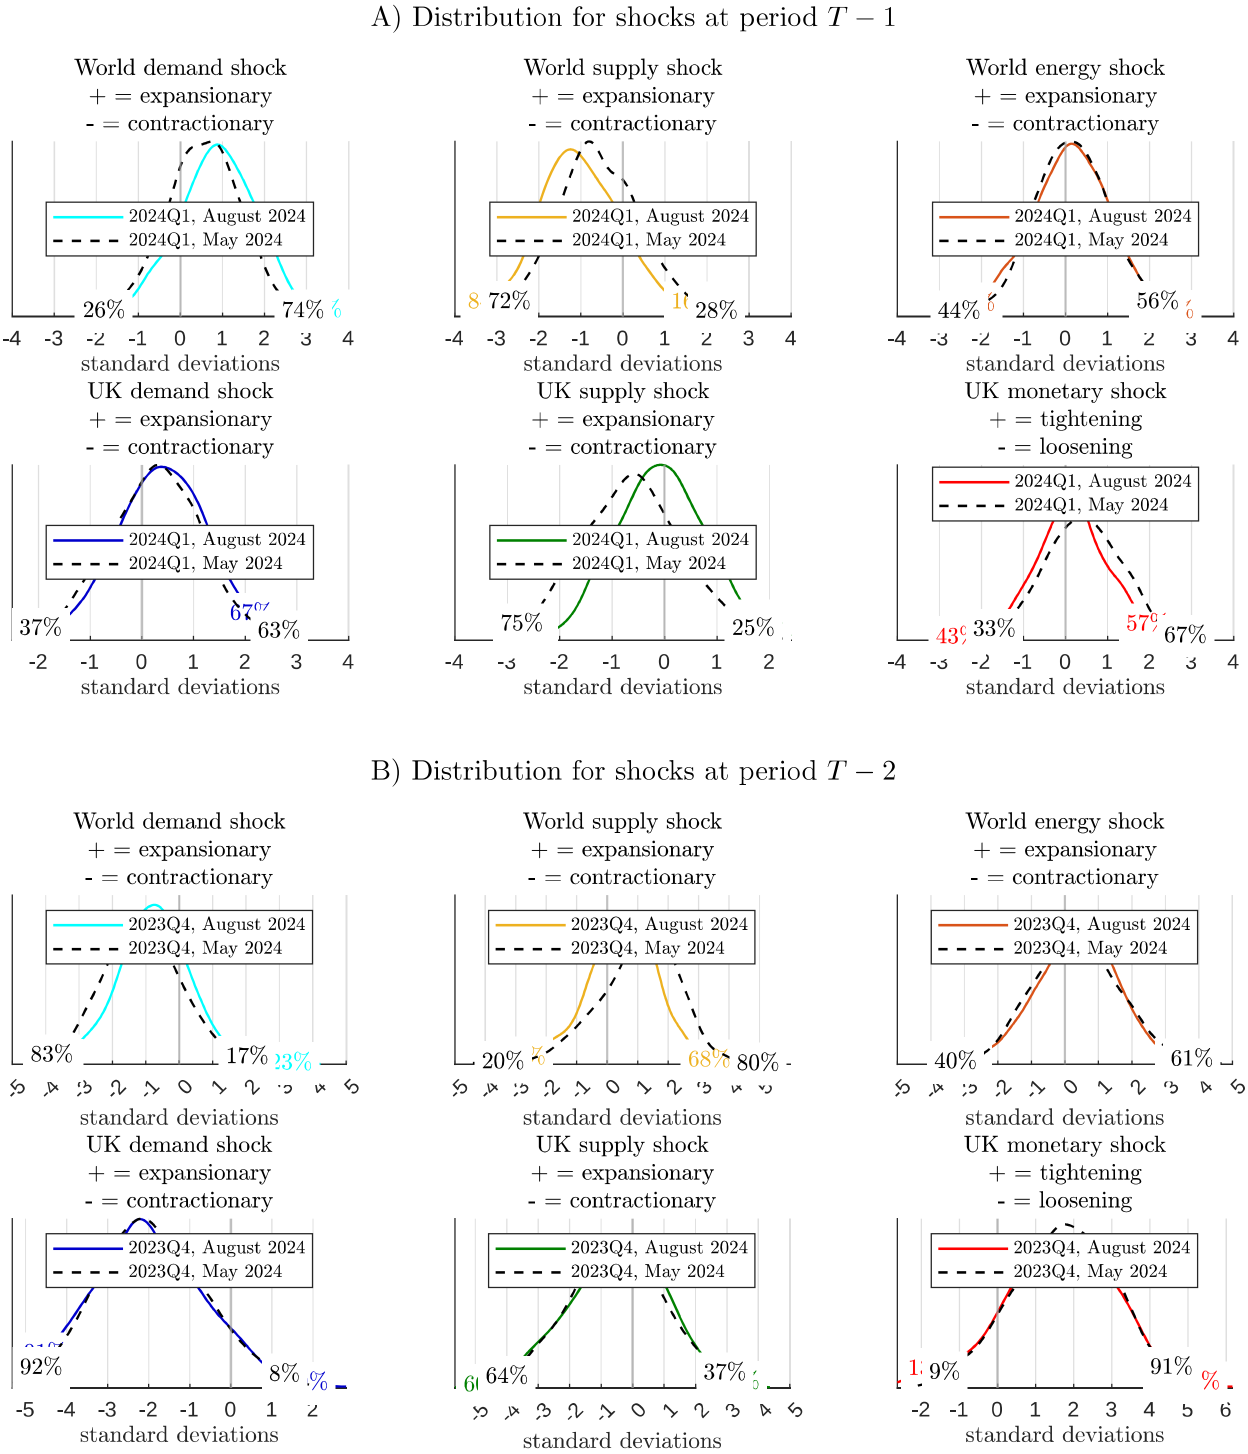

Note: Full distribution of the estimated shocks. Numbers report the probability that the shock was positive or negative.

Figure A-7: Revision in estimated shocks (2/2)

A) Distribution for shocks at period $T-3$

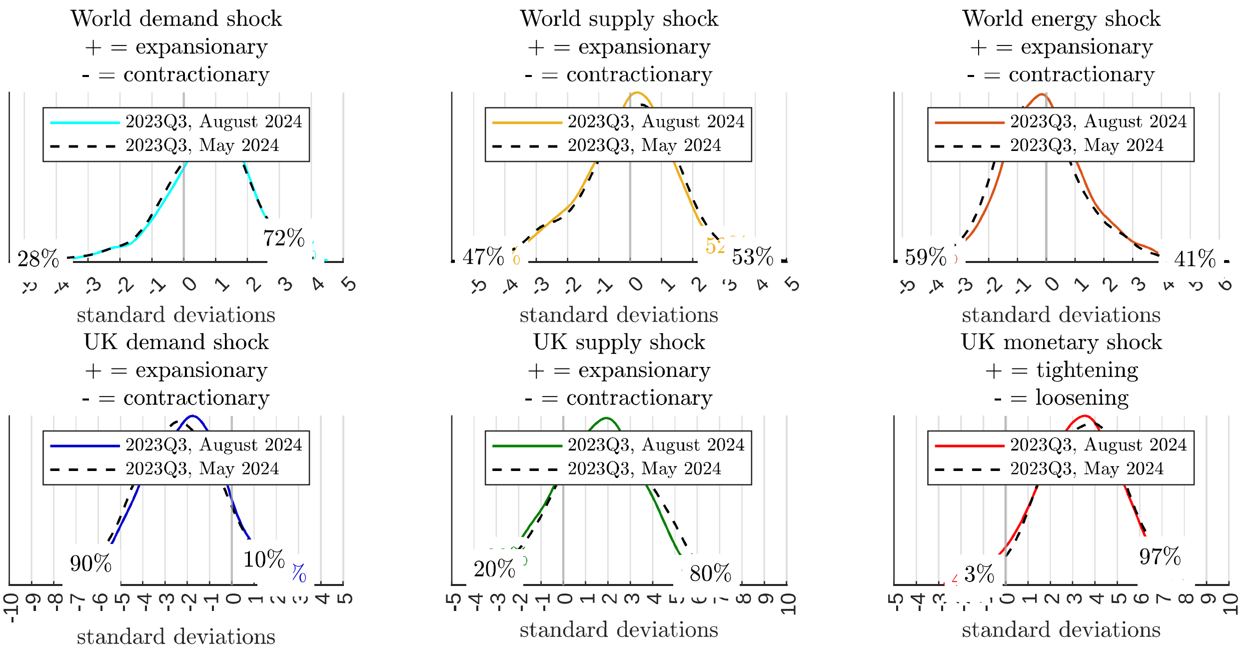

B) Distribution for shocks at period $T-4$

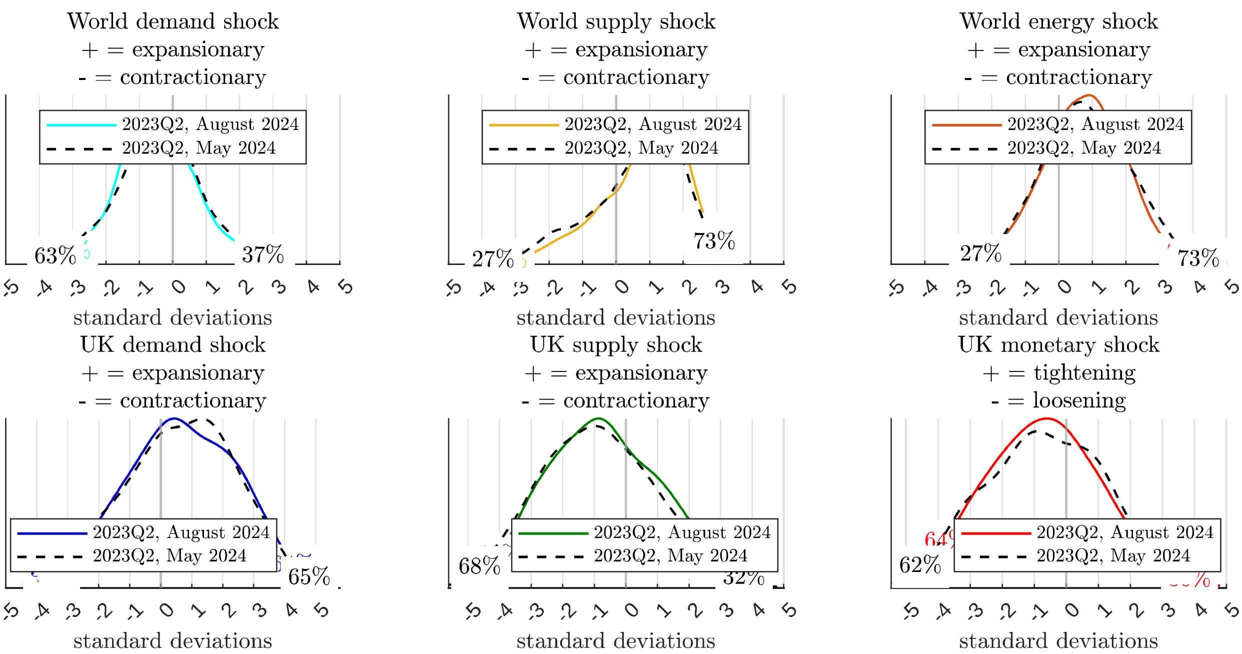

Note: Full distribution of the estimated shocks. Numbers report the probability that the shock was positive or negative.

Figure A-8: Effect of data revision (1/2)

A) Associated with shocks from period $T-2$

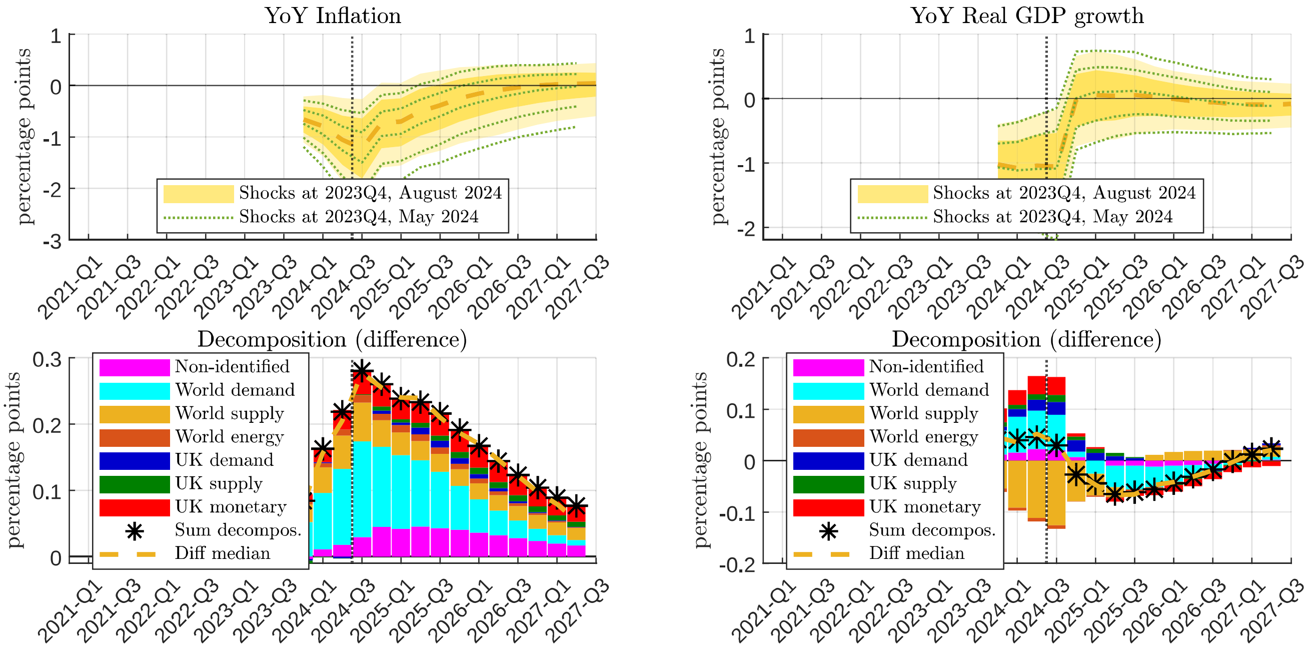

B) Associated with shocks from period $T-3$

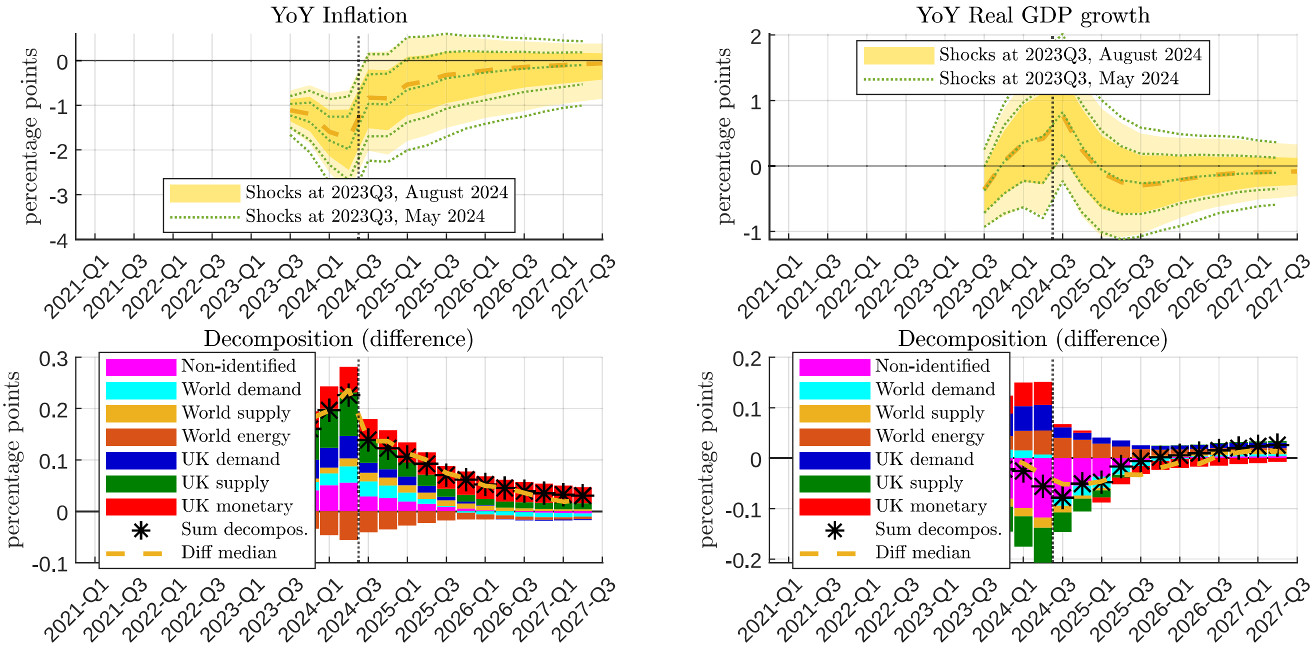

Figure A-9: Effect of data revision (2/2)

A) Associated with shocks from period $T-4$

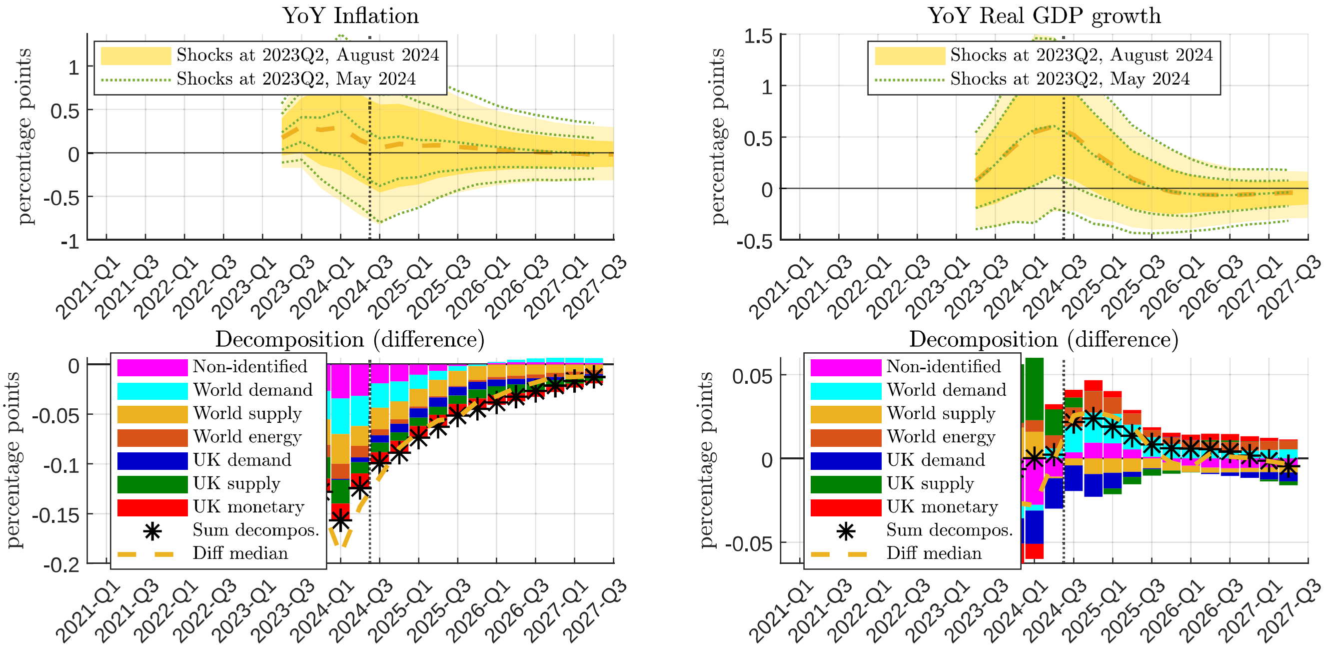

B) Associated with shocks from period $T-5$

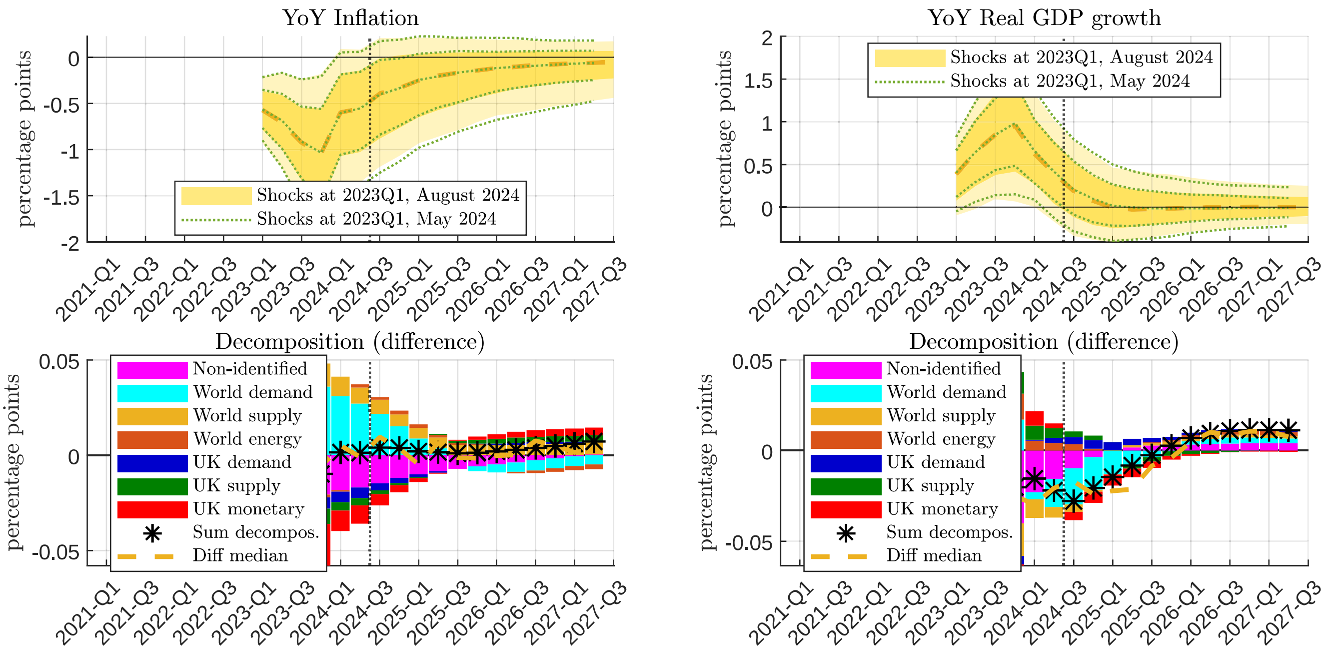

Figure A-10: Historical decompositions (estimation sample ends at Covid)

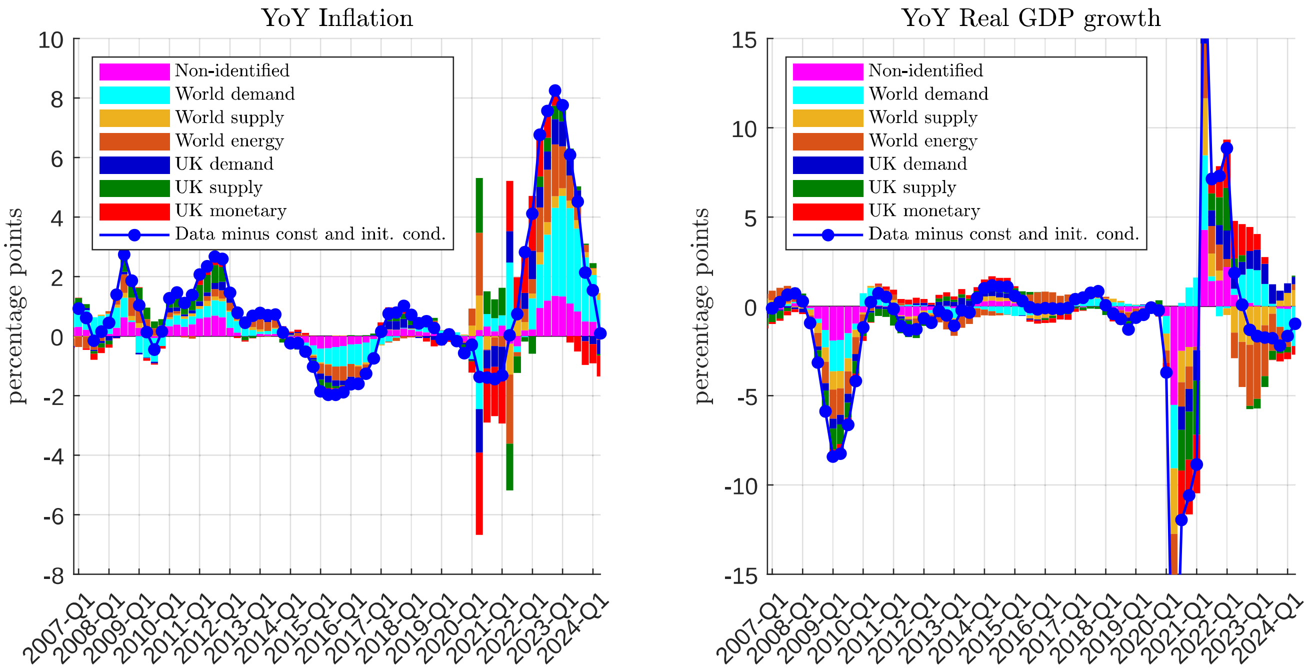

Note: The decomposition shown is associated with the pointwise mean decomposition over posterior draws. The decomposition is in deviation from the deterministic component, interpreted as the role associated with the constant term and the initial condition of the model.