

Structural forecast analysis*

Davide Brignone[†] Michele Piffer[‡]

January 7, 2026

Abstract

This paper shows how the structural representation of a vector autoregressive model can support forecast analysis. We offer a unified framework that formalizes how the structural form of the model can help form a narrative for two key statistics in real-time VAR forecasting: the forecast errors relative to the outturn of the data, and the consequent revisions of the forecast. To illustrate the method developed, we conduct a stylised real-time exercise on the UK, focusing on the inflation surge that followed the pandemic. We show that the inflation forecast produced by a four-variable VAR was revised upwards not only due to contractionary supply-side shocks, but also due to a mix of expansionary demand-side shocks, and a revision in the past shocks.

JEL classification: C32, E52.

Keywords: VAR modeling, forecasting, structural shocks, decomposition.

*We would like to thank Domenico Giannone, Christoph Grosse Steffen, Lutz Kilian, Simon Lloyd, Helmut Lutkepohl, Simon Lloyd, Alessandro Notarpietro, Andrea Renzetti and colleagues at the Bank of England for helpful comments and suggestions. We also thank the participants of the 5th Sailing the Macro Workshop, and seminar participants at Catholic University of Milan, Padova University, Bank of Italy and Bank of Spain. The views expressed in this paper do not represent the views of the Bank of England nor of its committees.

[†]Bank of England, London, United Kingdom. e-mail: davide.brignone@bankofengland.co.uk

[‡]Bank of England and King's Business School, King's College London, UK. e-mail: m.b.piffer@gmail.com.

1 Introduction

Vector autoregressive (VAR) models are widely used both in the academic community and in policy institutions. Since their introduction by Sims (1980), an extensive literature has expanded the VAR framework in multiple directions, including developing estimation procedures for inference, deriving identification techniques for causal analysis, and studying the relation between VAR models and other macroeconomic models (for instance, DSGE models).

This paper formalizes how the structural representation of a VAR model can support forecast analysis. Consider a forecast simulated using data until time T . When data for period $T+1$ become available, one uncovers the forecast error made in the time- T predictions of the model. In addition, new data could lead to a new and different forecast, which, in turn, may imply forecast revisions. As already acknowledged in the literature, when the forecast of interest is unconditional or reduced-form conditional (Waggoner and Zha, 1999), the econometrician can rely solely on the reduced-form specification of the VAR model. The use of a structural VAR becomes necessary only when the objective is to produce a structural conditional forecast (Baumeister and Kilian, 2014, Bańbura et al., 2015, Antolin-Diaz et al., 2021). Yet, we argue that irrespective of the forecast of interest, the structural representation of the model can still be valuable to form a narrative of what drives forecast errors and forecast revisions.

To appreciate the importance of structural narratives for forecast analysis, consider forecasting at policy institutions. Significant resources are invested not only in generating accurate forecasts, but also in forming a narrative that illustrates the possible economic channels consistent with the forecast. The narrative of a forecast is no less important than the forecast itself. A narrative is also required to interpret forecast errors and the forecast revisions. If the researcher only uses reduced form information, then one can only *report* forecast errors and forecast revisions, i.e. document to what extent the forecast made in the previous periods was incorrect, and how the latest forecast differs from the previous one. By contrast, we show how structural analysis can aid forecasting by decomposing forecast errors and revisions into five components: (a) the effects of new shocks that hit the economy at time $T+1$ as estimated from the new data realizations, (b) revisions in how long the effects from past estimated shocks are predicted to last over the forecast horizon, (c) (for conditional forecasts) changes in the simulated shocks generated to support the conditioning path over the forecast

horizon, (d) changes in the estimated unconditional mean that the model converges to in the medium term, and (e) changes in any other deterministic component of the model.

We first illustrate the methodology using a bivariate simulation. We work with pseudo data on variables that, for convenience, we refer to as output growth and inflation, driven by demand and supply shocks. The illustration starts from a period of high growth and high inflation, when the forecast at T predicts a slow reverse to the unconditional mean for both variables. At time $T+1$ the new data reveals two facts. First, that the forecast for output growth is to be revised upwards, having output growth materialized above the previous forecast. Second, that the outturn for inflation was in line with the forecast (i.e. the forecast error for inflation is zero), but the revised forecast now features an *undershooting* relative to the unconditional mean. The reduced form approach to forecasting can only document these statistics, which, however, call for an economic interpretation. Structural analysis can help by uncovering that (a) the positive forecast error on output growth is the joint response to expansionary demand and supply shocks, (b) the zero forecast error on inflation is due to the matching effects on inflation of the two shocks, (c) the undershooting of inflation in the new forecast comes from the delayed response of the recent deflationary supply shock.

We then apply the methodology to the data and describe how our framework can provide a narrative to the forecast paths and forecast revisions in applied work. We build a stylized four-variable SVAR model for the UK economy, and identify four shocks: a demand shock, a supply shock, an energy shock, and a monetary policy shock. Using this simple illustrative model, we then perform a real-time evaluation exercise, focusing on the period following the Covid-19 pandemic characterized by elevated volatility and the inflation surge. For each quarter, we use the actual vintage of the data, and we produce an unconditional forecast for each variable along with the full decomposition. This illustration suggests that the initial surge and positive revision of the inflation forecast from the four-variable VAR model in 2022Q2 were accounted for not only by a combination of inflationary supply and energy shocks, but also by expansionary demand and monetary policy shocks. In addition, we document that part of the forecast errors and revisions over 2022 and 2023 are explained by a change in the importance attributed to past shocks.

The literature on VAR modeling is very large, and key references include Koop and Korobilis (2010) and Kilian and Lütkepohl (2017). One reason behind the great

success of VAR models is that they provide a flexible framework for two separate types of analysis: forecasting and structural analysis. Yet, so far, the two strands of the literature on forecasting and structural modeling have largely developed separately. This may lead one to believe that there is only limited scope for these two sides of the model to work together.¹ Our paper challenges this view. Some of the elements in the analysis of this paper have appeared in isolation in previous work in the literature. Early traces of decompositions of forecast errors and revisions can be found in Todd (1992), who provides a purely narrative discussion and an algorithm. The relation between forecast revision, impulse responses and structural shocks is also discussed in Giannone et al. (2004). Brazdik et al. (2014) discuss a DSGE model decomposing forecast errors and forecast revisions in terms of the changes in the conditioning path of the forecasts. Giannone and Primiceri (2024) explore forecast errors as indicators of prevailing contemporaneous structural demand and supply shocks. Compared to this literature, we develop a single, comprehensive framework that provides a narrative of real-time forecasts.^{2,3}

The decomposition of forecast errors and forecast revisions into deterministic components and structural shocks does not strictly require working with VAR models and can be extended to different specifications of VAR models (e.g. trend-cycle VARs as in Del Negro et al., 2017 and Ascari and Fosso, 2024), Dynamic Stochastic General Equilibrium models and structural factor models. We work with VARs for their high tractability. We view this framework as illustrative of the broader potential of deriving the connection between reduced form forecast analysis and structural representations.⁴

The rest of the paper proceeds as follows: Section 2 illustrates the methodology. Section 3 shows a bivariate illustration. Section 4 shows an application to demand

¹One direct point of contact between the two strands of the literature is with regard to conditional forecasting, which can be simulated from structural rather than reduced form shocks in order to build scenarios (references in the text, as well as Chan et al., 2025 and Crump et al., 2025). Another example is the recent literature on optimal policy adjustments, which develops policy evaluation techniques that combine structural impulse responses and reduced form forecasts (Barnichon and Mesters, 2023, Caravello et al., 2024). We are not aware of additional work at the intersection of reduced form and structural VAR models.

²The decomposition of forecast revisions is also discussed in a DSGE model by the New York Federal Reserve Bank documented in <https://frbny-dsge.github.io/DSGE.jl/stable/>. Their analysis remains reduced form, focusing on the role played separately by latest data release, revision of past data, and changes in the parameter estimates.

³Our paper elaborates over earlier work circulated in Brignone and Piffer (2025).

⁴The decomposition of the forecast revisions into its structural drivers offers a new dimension along which identifying restrictions can be introduced, in the spirit of narrative sign restrictions (Antolín-Díaz and Rubio-Ramírez, 2018, Giacomini et al., 2022). We leave this for future research.

and supply shocks in real time. Conclusions follow.

2 Methodology

In this section we summarize the SVAR model used for the analysis and describe the decomposition of the forecast errors and the forecast revisions.

2.1 The model

The reduced form model is given by

$$\mathbf{y}_t = \sum_{l=1}^p \Pi_l \mathbf{y}_{t-l} + \mathbf{c} + \mathbf{u}_t, \quad (1a)$$

$$\mathbf{u}_t \sim N(\mathbf{0}, \Sigma), \quad (1b)$$

while the structural form of the model also adds

$$\mathbf{u}_t = B\boldsymbol{\epsilon}_t, \quad (2a)$$

$$\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, I), \quad (2b)$$

where it holds that $\Sigma = BB'$ and $B = \chi(\Sigma)Q$. $\mathbf{y}_t$ is a $k \times 1$ vector of endogenous variables. $\boldsymbol{\epsilon}_t$ is a $k \times 1$ vector of structural shocks driving the data, and are assumed Normally distributed with diagonal covariance matrix normalized to the identity matrix. The reduced form innovations $\mathbf{u}_t$ are a linear function of the structural shocks via equation (2a). The reduced form covariance matrix Σ is functionally constrained to the $k \times k$ impact matrix of the shocks B , $\chi(\Sigma)$ is the Cholesky decomposition of Σ (although any other unique decomposition of Σ is also admissible), and Q is a $k \times k$ orthogonal matrix. Π_l represents the reduced form autoregressive parameters of the model at horizon l , $l = 1, \dots, p$, while $\mathbf{c}$ is a constant term. We refer to Arias et al. (2018) for a discussion of alternative parametrizations.

Structural impulse responses are recovered by simulating recursively from equations (1a)-(2a) after setting $\boldsymbol{\epsilon}_t = \mathbf{e}_j \bar{\epsilon}$, with $\mathbf{e}_j$ a $k \times 1$ vector of zeros except for entry j , which is set to 1. This procedure generates impulse responses to a single structural *scalar*-shock of size equal to $\bar{\epsilon}$. For the analysis of this paper, it is helpful to generalize this concept to a structural *vector*-shocks $\bar{\boldsymbol{\epsilon}}$, where $\bar{\boldsymbol{\epsilon}}$ can now take nonzero value in more than one entry. For simplicity, we refer to this impulse response as a *composite impulse*

response associated with the generic impulse vector $\bar{\epsilon}$, and write it as $\phi(h, \bar{\epsilon})$, with h the number of periods after the shock occurs. Formally, iterate model (1a) backwards to rewrite the data at time t as a function of the data in periods $(t-1-\tau, \dots, t-p-\tau)$, with $\tau \geq 0$. This gives

$$\mathbf{y}_t = \sum_{l=0}^{p-1} \Phi_{l+1, \tau} \mathbf{y}_{T-\tau-l} + \Phi_{0, \tau-1} \mathbf{c} + \sum_{l=0}^{\tau-1} \Phi_{1, l} B \epsilon_{T-l}, \quad (3)$$

where the formulas for Φ are available in Section A of the Online Appendix as well as in Kilian and Lütkepohl (2017). Composite impulse responses for periods $h = 0, 1, \dots, \tau$ are given by $\phi(h, \bar{\epsilon}) = \Phi_{1, h} B \bar{\epsilon}$, with $\Phi_{1, 0} = I$. By construction, if $\bar{\epsilon} = \mathbf{e}_j \bar{\epsilon}$, only shock j is subject to an impulse, and composite impulse responses coincide with conventional impulse responses.

2.2 Interpreting the forecast error and forecast revision

We are interested in how to use composite impulse responses to interpret forecast errors and forecast revisions. For this purpose, define $\mathbf{y}_{T+h}^{(T)}$ as the h -steps period ahead forecast made at time T , with $h = 1, \dots, H$ the forecast horizon. The $k \times H$ array of forecasts $Y^{(T)} = [\mathbf{y}_{T+1}^{(T)}, \dots, \mathbf{y}_{T+h}^{(T)}, \dots, \mathbf{y}_{T+H}^{(T)}]$ is made when the data $[\mathbf{y}_1, \dots, \mathbf{y}_T]$ is available. At time $T+1$, the data realization $\mathbf{y}_{T+1}$ becomes available, and $Y^{(T+1)} = [\mathbf{y}_{T+2}^{(T+1)}, \dots, \mathbf{y}_{T+h}^{(T+1)}, \dots, \mathbf{y}_{T+H}^{(T+1)}]$ is generated using data $[\mathbf{y}_1, \dots, \mathbf{y}_T, \mathbf{y}_{T+1}]$. Note that we hold the end of the forecast horizon at $T+H$ (rather than extending it to $T+H+1$) for simplicity. Last, when no data revisions occur between forecasts, the observables $[\mathbf{y}_1, \dots, \mathbf{y}_T]$ are the same for both forecasts. Instead, some of the observables change if data revisions take place.

We are interested in using composite impulse responses to interpret the forecast error

$$\mathbf{v}_{T+1} = \mathbf{y}_{T+1} - \mathbf{y}_{T+1}^{(T)}, \quad (4)$$

and the forecast revision

$$\begin{pmatrix} \mathbf{y}_{T+2}^{(T+1)} - \mathbf{y}_{T+2}^{(T)} \\ \vdots \\ \mathbf{y}_{T+H}^{(T+1)} - \mathbf{y}_{T+H}^{(T)} \end{pmatrix}. \quad (5)$$

The forecast error reports the difference between the data realization at time $T+1$ and the forecast made for that period at time T . The forecast revision, instead, is the

change in the full profile of the forecast over the rest of the forecast horizon.⁵

Define $U^{(T)} = [\mathbf{u}_{T+1}^{(T)}, \dots, \mathbf{u}_{T+h}^{(T)}, \dots, \mathbf{u}_{T+H}^{(T)}]$ the reduced form innovations simulated to generate the forecast $Y^{(T)}$, and define $E^{(T)} = [\boldsymbol{\epsilon}_{T+1}^{(T)}, \dots, \boldsymbol{\epsilon}_{T+h}^{(T)}, \dots, \boldsymbol{\epsilon}_{T+H}^{(T)}]$ the underlying simulated structural shocks, with $\boldsymbol{\epsilon}_{t+h}^{(T)} = B^{-1}\mathbf{u}_{t+h}^{(T)}$. Three cases summarize the alternative approaches to forecasting with VAR models: (a) if $Y^{(T)}$ is an unconditional forecast, the researcher draws $U^{(T)}$ from the distribution (2a), sometimes directly setting $U^{(T)}$ equal to zero; (b) if $Y^{(T)}$ is a conditional forecast simulated from *reduced* form shocks, the researcher draws $U^{(T)}$ from equation (2a) subject to linear restrictions that ensure the conditioning path of interest (Waggoner and Zha, 1999); (c) if $Y^{(T)}$ is a conditional forecast simulated from *structural* shocks, the researcher draws $E^{(T)}$ from equation (2b) subject to linear restrictions that ensure the conditioning path of interest (Baumeister and Kilian, 2014, Bańbura et al., 2015, Antolin-Diaz et al., 2021, Chan et al., 2025 and Crump et al., 2025). Our method works irrespectively of the type of the simulated forecast as long as both $U^{(T)}$ and $E^{(T)}$ are available. For simplicity, the illustrations and applications shown in this paper only use unconditional forecasts, setting all entries of $(U^{(T)}, E^{(T)})$ to zero, but the method is derived in a more general setting.

2.3 A simplified setting

This section helps set ideas by working under selected simplifying assumptions: (a) the model includes no constant and only one lag of the endogenous variables, (b) the true parameter values of the model are known, and hence also the realizations of the shocks up to when the forecast is made, and (c) no data revision occurs between periods. Equation (1a) hence simplifies to

$$\mathbf{y}_t = \Pi \mathbf{y}_{t-1} + \mathbf{u}_t, \quad (6)$$

where (Π, B, Σ) are now treated as known parameters, and $(\mathbf{u}_t, \boldsymbol{\epsilon}_t)$ are known, with $\boldsymbol{\epsilon}_t = B^{-1}\mathbf{u}_t$.

⁵Our method can be extended to forecast errors and forecast revisions at time $T+1$ relative to forecasts made in periods *earlier* than T .

$Y^{(T)}$ can be computed as

$$\begin{pmatrix} \mathbf{y}_{T+1}^{(T)} \\ \mathbf{y}_{T+2}^{(T)} \\ \mathbf{y}_{T+3}^{(T)} \\ \vdots \\ \mathbf{y}_{T+H}^{(T)} \end{pmatrix} = \begin{pmatrix} \Pi \\ \Pi^2 \\ \Pi^3 \\ \vdots \\ \Pi^H \end{pmatrix} \mathbf{y}_T + \begin{pmatrix} I & 0 & 0 & \dots & 0 \\ \Pi & I & 0 & \dots & 0 \\ \Pi^2 & \Pi & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \Pi^{H-1} & \Pi^{H-2} & \Pi^{H-3} & \dots & I \end{pmatrix} \begin{pmatrix} \mathbf{u}_{T+1}^{(T)} \\ \mathbf{u}_{T+2}^{(T)} \\ \mathbf{u}_{T+3}^{(T)} \\ \vdots \\ \mathbf{u}_{T+H}^{(T)} \end{pmatrix}, \quad (7)$$

$$= \begin{pmatrix} \Pi \\ \Pi^2 \\ \Pi^3 \\ \vdots \\ \Pi^H \end{pmatrix} \mathbf{y}_T + \begin{pmatrix} I \\ \Pi \\ \Pi^2 \\ \vdots \\ \Pi^{H-1} \end{pmatrix} \mathbf{u}_{T+1}^{(T)} + \begin{pmatrix} 0 & 0 & \dots & 0 \\ I & 0 & \dots & 0 \\ \Pi & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \Pi^{H-2} & \Pi^{H-3} & \dots & I \end{pmatrix} \begin{pmatrix} \mathbf{u}_{T+2}^{(T)} \\ \mathbf{u}_{T+3}^{(T)} \\ \vdots \\ \mathbf{u}_{T+H}^{(T)} \end{pmatrix}, \quad (8)$$

with $\Pi^d = \prod_{i=0}^d \Pi$ and $\Pi^0 = I$. $Y^{(T)}$ could be a conditional or an unconditional forecast in a way reflected by the selection of $U^{(T)}$.

It is instructive to notice that at time $T+1$ the data realization $\mathbf{y}_{T+1}$ under model (6) differs from the forecast $\mathbf{y}_{T+1}^{(T)}$ according to equation

$$\mathbf{y}_{T+1} = \Pi \mathbf{y}_T + \mathbf{u}_{T+1}, \quad (9)$$

$$= \underbrace{\Pi \mathbf{y}_T + \mathbf{u}_{T+1}^{(T)}}_{\mathbf{y}_{T+1}^{(T)}} + \underbrace{(\mathbf{u}_{T+1} - \mathbf{u}_{T+1}^{(T)})}_{\mathbf{v}_{T+1}}, \quad (10)$$

$$\mathbf{v}_{T+1} = \mathbf{u}_{T+1} - \mathbf{u}_{T+1}^{(T)}, \quad (11)$$

$$= B(\boldsymbol{\epsilon}_{T+1} - \boldsymbol{\epsilon}_{T+1}^{(T)}). \quad (12)$$

Put differently, the forecast error $\mathbf{v}_{T+1}$ (which, by definition, equals the difference between the data realization $\mathbf{y}_{T+1}$ and the forecast $\mathbf{y}_{T+1}^{(T)}$) coincides with the difference between the realization of the actual innovation generating the data ($\mathbf{u}_{T+1}$) and the draws used at time T to simulate the forecast ($\mathbf{u}_{T+1}^{(T)}$). Since the innovations $\mathbf{u}_t$ are ultimately driven by structural shocks $\boldsymbol{\epsilon}_t$, the forecast error is driven by the difference between the actual realizations of the structural shocks behind the data at time $T+1$ ($\boldsymbol{\epsilon}_{T+1}$) and the values $\boldsymbol{\epsilon}_{T+1}^{(T)}$ of the structural shocks consistent with the reduced form innovations $\mathbf{u}_{T+1}^{(T)}$ used to simulate the forecast. If equation (6) is the true model, it is the inability to correctly predict $\boldsymbol{\epsilon}_{T+1}$ that drives the forecast error made at time

$T+1$.⁶

For the forecast made at time $T+1$ until horizon $T+H$, it holds that

$$\begin{aligned}
\begin{pmatrix} \mathbf{y}_{T+2}^{(T+1)} \\ \mathbf{y}_{T+3}^{(T+1)} \\ \vdots \\ \mathbf{y}_{T+H}^{(T+1)} \end{pmatrix} &= \begin{pmatrix} \Pi^2 \\ \Pi^3 \\ \vdots \\ \Pi^H \end{pmatrix} \mathbf{y}_T + \begin{pmatrix} \Pi & I & 0 & \dots & 0 \\ \Pi^2 & \Pi & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \Pi^{H-1} & \Pi^{H-2} & \Pi^{H-3} & \dots & I \end{pmatrix} \begin{pmatrix} \mathbf{u}_{T+1}^{(T+1)} \\ \mathbf{u}_{T+2}^{(T+1)} \\ \mathbf{u}_{T+3}^{(T+1)} \\ \vdots \\ \mathbf{u}_{T+H}^{(T+1)} \end{pmatrix}, \quad (13) \\
&= \begin{pmatrix} \Pi^2 \\ \Pi^3 \\ \vdots \\ \Pi^H \end{pmatrix} \mathbf{y}_T + \begin{pmatrix} \Pi \\ \Pi^2 \\ \vdots \\ \Pi^{H-1} \end{pmatrix} \mathbf{u}_{T+1} + \begin{pmatrix} I & 0 & \dots & 0 \\ \Pi & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \Pi^{H-2} & \Pi^{H-3} & \dots & I \end{pmatrix} \begin{pmatrix} \mathbf{u}_{T+2}^{(T+1)} \\ \mathbf{u}_{T+3}^{(T+1)} \\ \vdots \\ \mathbf{u}_{T+H}^{(T+1)} \end{pmatrix}. \quad (14)
\end{aligned}$$

Note that generating the new forecast $\mathbf{y}^{(T+1)}$ requires simulating possibly new innovations $\mathbf{U}^{(T+1)}$ that might well differ from $\mathbf{U}^{(T)}$, as in conditional forecasting. Note also that the first line in the above equation features $\mathbf{u}_{T+1}$ (rather than $\mathbf{u}_{T+1}^{(T)}$), namely the innovations responsible for the data realization $\mathbf{y}_{T+1}$.

Subtracting all but the first row of (8) from equation (14) highlights the following equation pinning down the forecast revision:

$$\begin{pmatrix} \mathbf{y}_{T+2}^{(T+1)} - \mathbf{y}_{T+2}^{(T)} \\ \mathbf{y}_{T+3}^{(T+1)} - \mathbf{y}_{T+3}^{(T)} \\ \vdots \\ \mathbf{y}_{T+H}^{(T+1)} - \mathbf{y}_{T+H}^{(T)} \end{pmatrix} = \underbrace{\begin{pmatrix} \Pi \\ \Pi^2 \\ \vdots \\ \Pi^H \end{pmatrix} \underbrace{(\mathbf{u}_{T+1} - \mathbf{u}_{T+1}^{(T)})}_{\mathbf{v}_{T+1} = B(\boldsymbol{\epsilon}_{T+1} - \boldsymbol{\epsilon}_{T+1}^{(T)})}}_{\boldsymbol{\gamma}_1} + \underbrace{\begin{pmatrix} I & 0 & \dots & 0 \\ \Pi & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \Pi^{H-2} & \Pi^{H-3} & \dots & I \end{pmatrix} \begin{pmatrix} \mathbf{u}_{T+2}^{(T+1)} - \mathbf{u}_{T+2}^{(T)} \\ \mathbf{u}_{T+3}^{(T+1)} - \mathbf{u}_{T+3}^{(T)} \\ \vdots \\ \mathbf{u}_{T+H}^{(T+1)} - \mathbf{u}_{T+H}^{(T)} \end{pmatrix}}_{\boldsymbol{\gamma}_2}. \quad (15)$$

In words, two elements are responsible for the forecast revision within the simplified setting studied in this section. The first, $\boldsymbol{\gamma}_1$, is the composite effect associated with the forecast error $\mathbf{v}_{T+1}$ over the full forecast horizon. This is the composite impulse response defined in the previous section, evaluated at the difference in the structural

⁶Part of the literature refers to forecast errors as $\mathbf{u}_{T+1}$ rather than $\mathbf{v}_{T+1}$. Under the special case considered in this section, $\mathbf{u}_{T+1} = \mathbf{v}_{T+1}$ when $\mathbf{u}_{T+1}^{(T)} = \mathbf{0}$. This is no longer sufficient under the more general case considered in Section 2.4, due to the potential role played by the possible revision in the deterministic component of the model and/or the role of past shocks.

shocks driving $\mathbf{v}_{T+1}$, i.e. $\boldsymbol{\epsilon}_{T+1} - \boldsymbol{\epsilon}_{T+1}^{(T)}$ (equation 12), and delayed by one period (i.e. premultiplied by Π). The second, $\boldsymbol{\gamma}_2$, is the effects associated with the difference in the shocks $U^{(T)}$ and $U^{(T+1)}$ simulated to generate the two forecasts.

A special case simplifies things further and highlights a key idea that was also discussed in Giannone et al. (2004). Suppose that $(Y^{(T)}, Y^{(T+1)})$ are computed as an unconditional forecast that assumes zero future shocks. This corresponds to $\mathbf{u}_{T+h}^{(T)} = \mathbf{0}$ for $h = 1, \dots, H$ and $\mathbf{u}_{T+h}^{(T+1)} = \mathbf{0}$ for $h = 2, \dots, H$. Equations (12) and (15) now simplify to

$$\begin{pmatrix} \mathbf{y}_{T+1}^{(T)} - \mathbf{y}_{T+1}^{(T+1)} \\ \mathbf{y}_{T+2}^{(T+1)} - \mathbf{y}_{T+2}^{(T)} \\ \mathbf{y}_{T+3}^{(T+1)} - \mathbf{y}_{T+3}^{(T)} \\ \vdots \\ \mathbf{y}_{T+H}^{(T+1)} - \mathbf{y}_{T+H}^{(T)} \end{pmatrix} = \begin{pmatrix} I \\ \Pi \\ \Pi^2 \\ \vdots \\ \Pi^H \end{pmatrix} B \boldsymbol{\epsilon}_{T+1}. \quad (16)$$

This equation shows that the forecast error and forecast revisions (first and remaining entries of 16) coincide with the sum of the impulse responses, weighted by the shocks that hit at time $T+1$, as discussed in Giannone et al. (2004). Put differently, under the special case considered here, forecast errors and forecast revisions coincide with the composite impulse response evaluated at $\boldsymbol{\epsilon}_{T+1}$.

Equation (16) helps highlight an importance result. In general, forecast errors and forecast revisions are viewed as statistics documenting either the error made in the forecast, or the update in the full remaining forecast. Yet, policy institutions always give great importance to forming a *narrative* that helps explain the forecast errors and revisions. Equation (16) helps think of the forecast errors and revisions as the output of the structural shocks that hit the model-economy at time $T+1$. Since these shocks are structural, the forecast error and revision can now be decomposed into economically meaningful stochastic events.

2.4 Extension to a more general setting

The above section works under the assumption that the data generating process is a VAR model with no constant, whose parameters are known, and where no data revision occurs between T and $T+1$. We now generalize the analysis.

Call $[\mathbf{y}_1^{(T)}, \dots, \mathbf{y}_{T-\tau}^{(T)}, \mathbf{y}_{T-\tau+1}^{(T)}, \dots, \mathbf{y}_T^{(T)}]$ and $[\mathbf{y}_1^{(T+1)}, \dots, \mathbf{y}_{T-\tau}^{(T+1)}, \mathbf{y}_{T-\tau+1}^{(T+1)}, \dots, \mathbf{y}_{T+1}^{(T+1)}]$ the datasets available to compute the forecasts at time T and $T+1$, respectively. The notation

allows for data revision to take place between forecasts. Call $(\Pi_l^{(T)}, \mathbf{c}^{(T)}, B^{(T)})$ the parameter values used for the forecast at time T , and call $\boldsymbol{\epsilon}_t^{(T)}$, $t = 1, \dots, T$ the implied estimates of the structural shocks (for instance, $(\Pi_l^{(T)}, \mathbf{c}^{(T)}, B^{(T)}, \boldsymbol{\epsilon}_t^{(T)})$ could stand for posterior draws if working in a Bayesian setting). A similar notation holds for the forecast at time $T+1$. Equation (3) can now be rewritten as

$$\mathbf{y}_{T+h}^{(T)} = \underbrace{\sum_{l=0}^{p-1} \Phi_{l+1, h+\tau}^{(T)} \mathbf{y}_{T-\tau-l}^{(T)}}_{\mathbf{dc1}_{T+h}^{(\tau, h; T)}} + \underbrace{\Phi_{0, h+\tau-1}^{(T)} \mathbf{c}^{(T)}}_{\mathbf{dc2}_{T+h}^{(\tau, h; T)}} + \underbrace{\sum_{l=0}^{\tau-1} \Phi_{1, h+l}^{(T)} B^{(T)} \boldsymbol{\epsilon}_{T-l}^{(T)}}_{\mathbf{sc1}_{T+h}^{(\tau, h; T)}} + \underbrace{\sum_{l=1}^h \Phi_{1, h-l}^{(T)} B^{(T)} \boldsymbol{\epsilon}_{T+l}^{(T)}}_{\mathbf{sc2}_{T+h}^{(\tau, h; T)}}, \quad (17)$$

see Section A of the Online Appendix. This decomposition highlights that the forecast can be thought of as composed of four distinct parts: (a) $\mathbf{dc1}_{T+h}^{(\tau, h; T)}$ captures the role attributed to the data up to time $T-\tau$. Model stationarity implies that this part converges to zero as $\tau + h$ increases; (b) $\mathbf{dc2}_{T+h}^{(\tau, h; T)}$ can be viewed as capturing the role attributed to the unconditional mean, because as $\tau + h \rightarrow \infty$, this term converges to the unconditional mean of the model; (c) $\mathbf{sc1}_{T+h}^{(\tau, h; T)}$ captures the role played by the structural shocks that were estimated between time $T-\tau+1$ and T ; (d) $\mathbf{sc2}_{T+h}^{(\tau, h; T)}$ captures the role played by the structural shocks consistent with the innovations simulated for the forecast period between $T+1$ and $T+H$. $(\mathbf{dc1}_{T+h}^{(\tau, h; T)}, \mathbf{dc2}_{T+h}^{(\tau, h; T)})$ refer to the deterministic component of the model, while $(\mathbf{sc1}_{T+h}^{(\tau, h; T)}, \mathbf{sc2}_{T+h}^{(\tau, h; T)})$ relate to the stochastic component of the model. $\mathbf{dc1}_{T+h}^{(\tau, h; T)}$ pools together shocks up to time $T-\tau$, as the decomposition into separate structural shocks starts from $T-\tau+1$, where τ can be selected as needed. As a comparison, the historical decompositions in Kilian and Lütkepohl (2017) and Bergholt et al. (2024) set $h = 0$ and $\tau = T$.

Iterating equation (17) forward to study the h -period ahead forecast made at time $T+1$ gives

$$\mathbf{y}_{T+h}^{(T+1)} = \underbrace{\sum_{l=0}^{p-1} \Phi_{l+1, h+\tau}^{(T+1)} \mathbf{y}_{T-\tau-l}^{(T+1)}}_{\mathbf{dc1}_{T+h}^{(\tau, h; T+1)}} + \underbrace{\Phi_{0, h+\tau-1}^{(T+1)} \mathbf{c}^{(T+1)}}_{\mathbf{dc2}_{T+h}^{(\tau, h; T+1)}} + \underbrace{\sum_{l=-1}^{\tau-1} \Phi_{1, h+l}^{(T+1)} B^{(T+1)} \boldsymbol{\epsilon}_{T-l}^{(T+1)}}_{\mathbf{sc1}_{T+h}^{(\tau, h; T+1)}} + \underbrace{\sum_{l=2}^h \Phi_{1, h-l}^{(T+1)} B^{(T+1)} \boldsymbol{\epsilon}_{T+l}^{(T+1)}}_{\mathbf{sc2}_{T+h}^{(\tau, h; T+1)}}. \quad (18)$$

Note that both forecasts are written as a function of the data until $T-\tau$ (rather than writing the forecast made at time $T+1$ as a function of data up to $T-\tau+1$). Note also

that the structural shocks at time $T+1$ move from the future stochastic component $\mathbf{sc2}_{T+h}^{(\tau,h;T)}$ (i.e. $\boldsymbol{\epsilon}_{T+1}^{(T)}$) to the present stochastic component $\mathbf{sc1}_{T+h}^{(\tau,h;T+1)}$ (i.e. $\boldsymbol{\epsilon}_{T+1}$), since they were simulated for the forecast made at time T , but were estimated at time $T+1$.

With this setting, the forecast revisions for horizons $h = 2, \dots, H$ can be written as

$$\begin{aligned} \mathbf{y}_{T+h}^{(T+1)} - \mathbf{y}_{T+h}^{(T)} = & (\mathbf{dc1}_{T+h}^{(\tau,h;T+1)} - \mathbf{dc1}_{T+h}^{(\tau,h;T)}) + (\mathbf{dc2}_{T+h}^{(\tau,h;T+1)} - \mathbf{dc2}_{T+h}^{(\tau,h;T)}) + \\ & + (\mathbf{sc1}_{T+h}^{(\tau,h;T+1)} - \mathbf{sc1}_{T+h}^{(\tau,h;T)}) + (\mathbf{sc2}_{T+h}^{(\tau,h;T+1)} - \mathbf{sc2}_{T+h}^{(\tau,h;T)}). \end{aligned} \quad (19)$$

This equation illustrates to what extent the forecast revision is driven by (a) an update in the role of all shocks up to $T - \tau$, (b) an update in the estimate of the unconditional mean of the model, (c) a revision in the role attributed by the two forecasts to the shocks estimated between time $T - \tau + 1$ and T , (d) an effect associated with the shocks that hit at time $T+1$ relative to the value simulated in the forecast from time T , and (e) a change in the role of future shocks over the remaining forecast horizon.

The four components from equation (19) potentially reflect a combination of changes in the parameter estimates and changes in the estimates of the shocks. The new data release for $T+1$ can lead to changes in the estimates of the parameters. This, potentially combined with the revision of the data until time T , can lead to changes in the estimates of the shocks until period T . Changes in the parameter estimates can lead to changes in the deterministic component of the model, including the unconditional mean. In addition, changes in the parameter and/or shocks can lead to changes in how the model predicts the shocks from until time T to still be unfolding over the forecast horizon.⁷

Equation (19) shows the decomposition of the forecast revision. A similar decomposition holds with respect to the forecast error. Following equation (3), the data realization $\mathbf{y}_{T+1}$ can be decomposed into the deterministic component up to time $T - \tau$ and the role of the structural shocks from $T - \tau + 1$ to $T+1$. Hence, similar to the forecast revisions, also forecast errors can be decomposed into the role attributed to the change in the deterministic component and the role of the subsequent structural shocks.

⁷We acknowledge that the revision in the deterministic component $\mathbf{dc1}$ can capture revisions in the role of the structural shocks that hit up to time $T - \tau$. For this reason, we suggest to select τ large enough to ensure that most of the revisions in the role of the shocks is attributed to $\mathbf{sc1}$, so that it can be decomposed into individual structural shocks. When identifying shocks using sign restrictions on the impulse responses, the signs might not be apparent in the role associated with the revisions of the past shock, see Section C of the Online Appendix.

3 An illustration using simulated data on a bivariate model

We use a bivariate simulation to further illustrate the decomposition of forecast errors and revisions proposed in the paper. We specify the data generating process as a bivariate SVAR model with a constant term and 12 lags. We use the model to generate data for the generic variables y_{1t} and y_{2t} , which are driven by shocks ϵ_{1t} and ϵ_{2t} . We refer to the variables as output growth and inflation, and view the structural shocks as pseudo demand and supply shocks.

To set the model parameters, we follow the approach proposed by Canova et al. (2024). We first specify the true impulse responses of output growth and inflation to the demand and supply shocks. We then set the true parameters of the model equal to the SVAR parameters consistent with the true impulse responses.⁸ A positive one-standard-deviation demand shock (blue dashed lines) increases output growth and inflation on impact by 1%. The responses then slowly revert back to zero, reaching half of the impact effect three periods after the shock. By contrast, a positive one-standard-deviation supply shock (red dotted line) increases output growth on impact by 0.5% and decreases inflation by 0.5%. Contrary to demand shocks, supply shocks generate hump-shaped responses that reach the peak effect 2 periods after the shock. The SVAR parameters consistent with these impulse responses imply model stationarity. Last, we set the true constant terms of the model such that the model-implied unconditional mean for the pseudo output growth and inflation equals 1% and 2%, respectively.

We use the model as follows. We simulate 200 periods of pseudo data, initializing the data at the unconditional mean of the model. To generate data, we randomly draw shocks from their distribution, except for the demand shock in the last five periods, which we set equal to one standard deviation. This generates a period of strong output growth and elevated inflation, which serves as starting point of the exercise. Then, starting from period $T = 200$, we simulate an unconditional forecast until horizon $T+H = 220$, assuming zero future shocks. Last, we generate data for period $T+1 = 201$ and then simulate a new unconditional forecast from the point of view of period $T+1$ over the forecast horizon $T+2 = 202$ to $T+H = 220$, still assuming zero future shocks. This framework implies forecast errors at time $T+1 = 201$, and a forecast revision from $T+2 = 202$ to $T+H = 220$. We use the structural form of the

⁸Section B of the Online Appendix discusses in details the parametrization of the model, and Figure B-1 shows the true impulse responses of the model.

model to help interpret the economic forces driving the forecast errors and revisions. Following the discussion in the previous section, we set $\tau = 5$ in equations (17)-(18)

3.1 Illustration in a simplified setting


We begin from a case where we initially assume no data revision between forecasts. This means that the data covering the period up to T is the same when forecasting at time T or at $T+1$. We also temporarily assume that the true parameters of the model are known when generating both forecasts. These are important assumptions because they imply that forecast errors and forecast revisions can only be driven by the latest shocks at period $T+1$ rather than a revision in the role of past shocks and a revision in the deterministic component of the model (we generalize the simulation in the second part of this section). Both forecasts are unconditional forecasts generated from the reduced form representation of the model, setting the future shocks to zero.

Figure 1 reports the analysis for both inflation (top panel) and output growth (bottom panel). For both panels, the left plots show the data available until time T , the unconditional mean of the model, the forecast made at time T , and its decomposition into deterministic component and stochastic components for shocks between $T-\tau+1$ and T . The middle plots show the data available both until time T and $T+1$, the unconditional mean of the model, the forecast at time T , the new forecast made at time $T+1$, and the decomposition of the new forecast into deterministic component, and the stochastic component from shocks between $T-\tau+1$ and $T+1$. The right plots show the forecast error made at time $T+1$, the forecast revision until $T+H$, and the decomposition of the forecast error and revision. In all figures, the grey vertical line indicates $T-\tau$ while the vertical black line indicates $T+1$.

The left plots of Figure 1 show that at time T , the model predicts a slow decline of output growth and inflation towards the unconditional mean, with no undershooting relative to the long term. At time $T+1$ (middle plots) the new data turns out to be in line with the forecast for inflation, but 2% above the forecast for output growth. In addition, the new forecast outlines that inflation will temporarily *undershoot* the unconditional mean, and output growth will decline much less rapidly.

A purely reduced form approach to forecasting would provide very limited support to the interpretation of the forecast errors and forecast revisions. A researcher would not be able to go beyond stating that the new forecast at $T+1$ suggests an upward revision for the forecast for output growth and a downward revision for the forecast

Figure 1: Illustration with no data revision and true parameters



Note: The left-hand side panels show the forecast at time T (solid line), while the middle panels show the new forecast at time $T+1$ (solid line) along with the one produced at time T (dashed line). Both forecasts are decomposed into the role of the different components up to time T and $T+1$ respectively: demand shock (blue bars), supply shocks (red bars) and deterministic component (grey bars). The right-hand side panels plot the forecast error and the marginal difference between the forecasts, along with the contribution of each component. Figure B-2 in the Online Appendix breaks down the stochastic components of each forecast into the contribution from the shocks of each period. These are the composite impulse responses discussed in Section 2.1.

for inflation, and that there was no forecast error for inflation and a positive forecast error for output growth. The methodology proposed in this paper offers a tool to derive a structural narrative of the forecast errors and forecast revisions. The left-hand side panels of Figure 1 show that the forecast at time T is partly driven by the strong demand shocks that have hit the system up to time T , and which are still propagating through the system. As these shocks fade away, the forecast converges to

the unconditional mean. The forecast made at time $T+1$ (middle plots) still shows a strong (yet weaker) effect from the demand shocks that have hit up to time T . It also shows a larger role associated with the shocks from time $T+1$, which were simulated to equal zero from the point of view of the forecast at time T .


In this illustration, the data realizations at $T+1$ were generated by simulating a positive one-standard-deviation demand shock and a positive (deflationary) two-standard-deviation supply shock. The joint effects of these shocks are noticeable in the middle panels. On output growth, both shocks are expansionary, and explain the strong forecast error between the two periods. Yet, for inflation, the fact that no forecast error is detected hides the fact that two opposite forces are playing out: an inflationary demand shock, and a deflationary supply shock. The undershooting of inflation predicted by the forecast made at time $T+1$ can now be rationalized as the effect of the supply shock: since supply shocks feature hump-shaped responses (see Figure B-1 in the Online Appendix), the large deflationary supply shock materializes in the medium term of the forecast, explaining the forecast revision and the undershooting of inflation.

The right-hand side plots of Figure 1 confirm that the forecast errors and forecast revisions are driven by the structural shocks that hit at time $T+1$. By contrast, no role is played by the revision in the deterministic component nor the role of the latest shocks before period $T+1$. This result is driven by the fact that no revisions apply to the data, and that the same (true) parameter values are used for both forecasts, hence the same estimates of the shocks. The forecasts made at time T assumes zero shocks at $T+1$, which hence play no role over the forecast. The forecast made at time $T+1$ infers the shocks at time $T+1$ from the data, hence these shocks will be a driving forces of the variables over the forecast horizon.

3.2 Illustration in a generalized setting

We conclude the illustration by bringing into the discussion the more realistic scenario in which the parameters are estimated, and the data is subject to revisions from one forecast to the other. This has important consequences. The combination of data revision up to time T and data release at $T+1$ leads to changes in the parameter estimates and the estimates of the shocks up to time T . This jointly implies that the forecasts can now build on different estimates of the unconditional mean of the model and of the role associated with the shocks up to time T .

Figure 2: Illustration with data revision and estimated parameters



Note: Forecasts and forecast decompositions associated with the OLS estimates. The left-hand side panels show the forecast at time T (solid line), while the middle panels show the new forecast at time $T+1$ (solid line) along with the one produced at time T (dashed line). Both forecasts are decomposed into the role of the different components up to time T and $T+1$ respectively: demand shock (blue bars), supply shocks (red bars) and deterministic component (grey bars). The right-hand side panels plot the forecast error and the marginal difference between the forecasts, along with the contribution of each component. Figure B-3 in the Online Appendix breaks down the stochastic components of each forecast into the contribution from the shocks of each period. These are the composite impulse responses discussed in Section 2.1.

We start from the original data at time T and add noise in the five periods up to T , modeling noise as the realization of independent Normal random variables with standard deviation set equal to 0.05. We then start from the original data until $T+1$ and subject it to noise, drawn in the same way as for time T . We use both datasets to estimate the reduced form parameters via Ordinary Least Squares. Last, we estimate the structural impact effect of the shocks by applying to the estimated

Cholesky decomposition of the estimated reduced form variances the true orthogonal matrix associated with the data generating process.⁹ Forecasts and decompositions are then generated using the parameter estimates and the (noisy) data used in the estimation.

The analysis with data revision is reported in Figure 2. While the narrative of the forecasts and the decompositions is similar to Figure 1, the fact that the two forecasts use different parameter and shock estimates implies that the forecasts can associate different roles to the deterministic component of the model and to the shocks in the period between $T-\tau+1$ and T . For instance, the forecast made at time $T+1$ interprets the demand shocks between $T-\tau+1$ and T as being more expansionary on output growth compared to the forecast at time T . The effects of the revised shocks on inflation and output growth are reported in the right-hand side panels by the bars highlighted in squares.

4 An application to the surge of inflation in 2022

In this section, we use our methodology in a SVAR model for the UK economy. We apply the framework explained in Section 2 to analyze the period of high inflation that followed the Covid-19 pandemic. We keep the model parsimonious and tractable, as the main intent of this section is to showcase the possible use and benefits of the methodology.

4.1 Model specification, identification, and estimation

We estimate an SVAR model of the form described in equation (1)-(2). The model includes four variables: (i) the UK policy rate captured by the Bank rate; (ii) UK Real GDP; (iii) the UK consumer price index; (iv) Real oil prices. Except for the Bank rate, all variables enter the model in log difference, in order to ensure stationarity. The frequency of the data is quarterly, and the full sample covers the period from 1992Q1 to 2025Q2.

We identify four structural shocks. We use sign restrictions to identify generic demand and supply shocks, along with a monetary policy shock and an energy shock.¹⁰

⁹We follow this approach in order to avoid entering issues related to the identification of the shocks, which we view as not strictly important for the simulation exercise from this section.


¹⁰We acknowledge that sign restrictions, pioneered by Canova and De Nicolò (2002) and Uhlig (2005), is an active field of research, for instance Baumeister and Hamilton (2015), Arias et al.

Restrictions, reported in Table D-1 of the Online Appendix, are rather standard in the literature, and are introduced only on the impact effect of the shocks.

We set $p=5$ as the number of lags of the model. We add Covid-19 dummies for the quarters from 2020Q1 to 2020Q4 to deal with the volatility over the Covid period.¹¹ We then estimate the model using Bayesian methods. We use a Minnesota prior combined with a single-unit-root prior to discipline the deterministic component, as in Bergholt et al. (2024). We follow Giannone et al. (2015) and use a hierarchical approach to the hyperparameters of the model. The estimation of the Covid dummies follows the pandemic prior approach proposed by Cascaldi-Garcia (2022).

4.2 Results from forecast analysis

Figure 3: Sequence of estimated mean forecast errors and forecast revisions



Note: Diamonds indicate pointwise mean forecast errors for each forecast produced in a specific quarter, while lines the related pointwise mean forecast revisions.

(2018), Bruns and Piffer (2023), Inoue and Kilian (2026b) and Inoue and Kilian (2026a). In this application we follow Arias et al. (2018).

¹¹The introduction of Covid dummies requires adding extra terms to the deterministic component from equations (17)-(18). Our approach is also compatible with other ways of handling the large outliers from the Covid pandemic, including period-specific stochastic volatility (Lenza and Primiceri, 2022) and fat tails (Kociecki et al., 2025).

We conduct a real-time forecast exercise in the period that goes from 2021Q4 to 2025Q2. For each quarter, which we label as time T , we estimate the model with data from 1992Q2 up to time T and we produce an unconditional forecast with zero future shocks, simulating the forecast H -steps ahead with $H=12$.¹² We then move to the next quarter and compute the forecast error given the data at time $T+1$. We re-estimate the model using data from 1992Q2 to $T+1$, simulate the new forecast made at $T+1$, and compute the forecast revisions relative to the previous quarter. We conduct the exercise using the vintages of the data available at each new quarter, so that we can also account for data revisions over time. Last, for each quarter we apply the forecast decomposition outlined in Section 2, setting $\tau = 12$. This implies that we decompose each forecast into the role of the data up to three years before the forecast and into the subsequent structural shocks.¹³

Before discussing our structural forecast decompositions, we find it helpful to document the *reduced form* results of this exercise. Figure 3 reports results for QoQ inflation (top panel) and QoQ GDP growth (lower panel). The diamond in each period T reports the pointwise mean forecast error for that period, while the line that starts from each diamond reports the subsequent pointwise mean forecast revision.¹⁴ A few findings are visible from the figure. First, forecast errors and revisions over 2022 and 2023 are sizable for both GDP and inflation, and bigger compared to subsequent periods (see also Ball et al., 2022, Koch and Noureldin, 2024 and Giannone and Primiceri, 2024). Second, the comovement in the forecast errors and forecast revisions for inflation and GDP growth can change considerably across periods. For example, 2021Q4 and 2022Q4 were marked by positive forecast errors for both inflation and GDP growth, while in 2023Q1 and 2024Q4 the correlation in forecast errors was negative. Third, in some periods, small forecast errors can be associated with large forecast revisions, as, for instance, for GDP growth in 2022Q2 and 2022Q3.


The first investigation of the drivers of forecast errors and forecast revisions can

¹²We choose $H=12$ as it coincides with the 3-years ahead forecast horizon usually analyzed by central banks.

¹³More precisely, we simulate the forecasts using data until 2021Q4 and then until 2022Q1. We decompose both as a function of the data until 2018Q4 and the subsequent shocks. Then, we simulate the forecast at 2022Q2. Decompose the forecasts made at 2022Q1 and 2022Q2 as a function of the data up to 2019Q1 and the subsequent shocks. We continue until the forecasts made at 2025Q1 and 2025Q2, which are decomposed as a function of the data up to 2022Q2 and the subsequent shocks.

¹⁴The forecast error shown for period $T+1$ is reported as the difference between the first data realization that became available at time $T+1$ and the forecast made for that period at T . The forecast revision is reported as the difference between the newly formed forecast made at time $T+1$ and the previous forecast made at T .

Figure 4: Estimated structural shocks




Note: The top panel shows the marginal posterior distribution of shocks for 2022Q2 as estimated when data up to 2022Q2 became available. The bottom panel shows the same for 2022Q3. The figure also reports the share of posterior draws for which each shock is positive or negative. Sign restrictions imply that, for each shock, positive realizations of each of the shock should be interpreted as expansionary on GDP growth, see Figure D-5 in the Online Appendix.

be based on the comovement of the variables' forecast errors. Periods in which the forecast errors for inflation and GDP growth were both either positive or negative suggest a demand-side narrative, while supply-side narratives are more supported by negative correlations in the forecast errors. As an example, Figure 3 suggests that 2021Q4 and 2022Q1 were characterised by demand-side shocks. There are limitations to building narratives on the correlation of forecast errors. First, it offers a view of the prevailing shocks, but not a deeper inspection of other drivers at play. Second, it is challenging for periods in which one of the two variables depicts limited forecast errors. This the case for 2022Q2 and 2022Q3.

To better understand the drivers behind the forecast errors, one can analyse the shocks estimated by the model. Figure 4 shows the posterior marginal distribution of the estimated structural shocks that hit in 2022Q2 and 2022Q3.¹⁵ For 2022Q2,

¹⁵Figure 4 reports real-time estimates computed in 2022Q2 and 2022Q3, respectively. See Figure D-6 in the Online Appendix for the distributions estimated using subsequent three quarters of data vintages.

Figure 5: Forecast analysis for 2022Q2



Note: Pointwise mean forecasts and pointwise mean decompositions. The left-hand side panels show the forecast at time T (solid line), while the middle panels show the new forecast at time $T + 1$ (solid line) along with the one produced at time T (dashed line). Both forecasts are decomposed into the role of the different components up to time T and $T + 1$ respectively: demand (blue bars), supply (red bars), energy (orange bars), monetary policy shocks (yellow bars) and deterministic component (grey bars). The right-hand side panels plot the forecast error and the marginal difference between the forecasts, along with the contribution of each component.

the figure reports strong contractionary supply-side shocks (both supply and energy), along with marked expansionary demand-side shocks (both demand and monetary policy). This combination of shocks is consistent with upward pressure on inflation and an ambiguous effect on GDP growth. 2022Q3 shares a similar feature, as the shocks for 2022Q3 generate opposite contemporaneous forces playing out on output.


In line with the methodology discussed in Section 2, we complement the analysis of Figure 4 by jointly studying all possible forces that play out on forecast errors and forecast revisions. Figure 5 inspects the pointwise mean forecast errors and revisions for 2022Q2. In 2022Q1 (left-hand side panels) the model predicts GDP growth to reach -0.5% around 2022Q3, and inflation to reach 1.5%, before moving gradually to 0.5% towards the end of the horizon. Our decomposition suggests that the inflation forecast is elevated due to a mix of negative supply (in red) and energy (in orange) shocks, but also expansionary demand (in blue) and monetary policy (in yellow) shocks. The absolute narrative is similar for GDP growth, with negative supply and energy shocks causing GDP growth to be weak. At the same time, previously positive demand and monetary policy shocks turn negative on growth, contributing to the weak forecast. As we reach $T+1 = 2022Q2$, the data is found to have been revised, as can be seen by the difference between the dotted and the circled lines in the middle panels. Data in 2022Q2 came in almost -0.05 percentage points lower than what predicted in 2022Q1 for GDP, and almost 2.25 percentage points higher for inflation. On the latter, the new forecast produced at time $T+1$ is overall higher for inflation, but lower for GDP. In absolute space, the model interprets the elevated path for the inflation forecast similarly to time T , with a mix of inflationary supply and energy shocks, and inflationary demand and monetary policy shocks.

The right-hand side panels complement these results by plotting the decomposition of the forecast errors and of the change between the two forecasts. Only around 1 percentage points of the 2.25 forecast error for inflation was driven by the supply-side (supply and energy) shocks estimated for 2022Q2. A large part of the remaining forecast error is interpreted by the model as the effect of demand and monetary policy shocks. The revision of the role of previous shocks over the last three years before the forecast is instead found to play a more marginal role in driving the forecast errors and revisions. The interpretation for GDP growth is similar: the small forecast error for GDP growth is explained by a mix of contractionary supply-side shocks and expansionary demand-side shocks.

The analysis from Figure 5 can be carried out for every quarter. Figure 6 reports the analysis for 2022Q3. The left-hand side panels show the forecast associated with 2022Q2 and its decomposition, while the middle panels report the forecast computed in 2022Q3.¹⁶ Relative to 2022Q2, the forecast errors for inflation is negative and the

¹⁶In Figure 5, the forecast made at time 2022Q2 is decomposed as the role of the data up to 2019Q1 and the remaining components. In Figure 6 the same forecast is decomposed as the role of the data

Figure 6: Forecast analysis for 2022Q3



Note: Pointwise mean forecasts and pointwise mean decompositions. The left-hand side panels show the forecast at time T (solid line), while the middle panels show the new forecast at time $T+1$ (solid line) along with the one produced at time T (dashed line). Both forecasts are decomposed into the role of the different components up to time T and $T+1$ respectively: demand (blue bars), supply (red bars), energy (orange bars), monetary policy shocks (yellow bars) and deterministic component (grey bars). The right-hand side panels plot the forecast error and the marginal difference between the forecasts, along with the contribution of each component.

revision for inflation is downwards. As shown in the right-hand side panels, one of the factors contributing to this finding is a mix of mildly deflationary energy and supply

up to 2019Q2 and the remaining components. The difference is however immaterial, as should be expected given the selection of τ .

shocks, along with a minor restrictive monetary policy shock. However, contrary to Figure 5, a big role is also played by the revision in the role of past shocks up to time T , as the model now reviews the previous shocks to be less inflationary.¹⁷ Overall, it is interesting to notice that the model suggests an important role of demand and monetary policy shocks to explain forecast errors and revisions, a result in line with what found for the US and the Euro Area by Giannone and Primiceri (2024) (see Neri et al. (2023) for an alternative view). The increasing role played by demand-type shocks is also confirmed by the analysis for 2022Q4 (Figure D-7 in the Online Appendix), where the inflation forecast is revised up mainly due to demand shocks, and partly due to an upward revision of the unconditional mean.

Last, we computed, on average, how much forecast errors and forecast revisions are driven by revisions in the deterministic component, revisions in the role of past pooled shocks, and the role of latest pooled shocks. We found that while a large component of the forecast error is attributed to the realizations of the latest shocks, a sizeable component is still associated with revisions of the role of past shocks and partially of the deterministic component. From 2021Q4 to 2025Q5 on average, the latest shocks explain 71% and 60% of the (absolute value of the) forecast error of inflation and GDP growth, respectively. 21% and 31% are associated with the revision of the previous shocks, while the remaining 8% and 9% are due to the revision of the role of the deterministic component. The results remain qualitatively similar when inspecting the one-year-ahead forecast revisions. Figures D-10-D-11 in the Online Appendix report the results

5 Conclusions

This paper shows that the structural representation of a VAR model can offer a way to derive a narrative for forecast errors and forecast revisions in terms of structural shocks even when the forecasts of interest are unconditional reduced form forecasts. Since being able to explain the forecasts and its revisions plays a key role in forecasting – especially in policy institutions – we view the method proposed in the paper as a useful new entry to the toolkit of time series methods for macroeconomics.

The methodology proposed in the paper decomposes forecast errors and forecast revisions as a function of four components: (a) changes in what the paper refers to as the estimated deterministic component of the model, which typically captures the

¹⁷See also Figure D-9 in the Online Appendix, which reports the revisions of the shocks.

unconditional mean of the model; (b) the role associated with the estimated shocks in the periods leading up to the forecasts; (c) the role of the latest shocks that are estimated after the realization of the new data; and (b) the change in the shocks simulated to enforce conditional forecasts.

We first show our methodology by using simulated data in a bivariate VAR model. We then apply our method to the UK economy and study forecast errors from the experimental small VAR and revisions in the aftermath of the Russian invasion of Ukraine. We show that the strong upward revision in the inflation forecast derived from the simple VAR in 2022Q2 was driven not only by contractionary supply and energy shocks that hit in 2022Q2, but also by expansionary demand and monetary policy shocks in the same quarter. We also found that part of the forecast revisions in the subsequent quarters was due to a revision in the role of past shocks.

We conclude the analysis with a word of caution, remarking that careful consideration is required in the model specification to produce accurate forecasts and credibly identified shocks. We stress that the main purpose of this application is not to validate the small-scale model described in the previous section, but to show the relation between forecast errors, forecast revisions, and structural shocks, describing how the framework from this paper can be applied in a real-time forecast exercise.

References

- Antolin-Diaz, J., Petrella, I. and Rubio-Ramírez, J. F. (2021), ‘Structural scenario analysis with SVARs’, *Journal of Monetary Economics* **117**, 798–815.
- Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018), ‘Narrative sign restrictions for svars’, *American Economic Review* **108**(10), 2802–2829.
- 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. and Fosso, L. (2024), ‘The international dimension of trend inflation’, *Journal of International Economics* **148**, 103896.
- Ball, L., Leigh, D. and Mishra, P. (2022), ‘Understanding us inflation during the covid-19 era’, *Brookings Papers on Economic Activity* **2022**(2), 1–80.
- Bañbura, M., Giannone, D. and Lenza, M. (2015), ‘Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections’, *International Journal of forecasting* **31**(3), 739–756.
- Barnichon, R. and Matthes, C. (2018), ‘Functional approximation of impulse responses’, *Journal of Monetary Economics* **99**, 41–55.
- Barnichon, R. and Mesters, G. (2023), ‘A sufficient statistics approach for macro policy’, *American Economic Review* **113**(11), 2809–2845.
- Baumeister, C. and Hamilton, J. D. (2015), ‘Sign restrictions, structural vector autoregressions, and useful prior information’, *Econometrica* **83**(5), 1963–1999.
- Baumeister, C. and Kilian, L. (2014), ‘Real-time analysis of oil price risks using forecast scenarios’, *IMF Economic Review* **62**(1), 119–145.
- 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.
- Brazdik, F., Humplova, Z. and Kopriva, F. (2014), Evaluating a structural model forecast: Decomposition approach, Technical report, Research and Policy Notes 2014/02, Czech National Bank, Research and Statistics Department.

- Brignone, D. and Piffer, M. (2025), ‘A structural VAR model for the UK economy’, *Bank of England Macro Technical Paper* (3).
- Bruns, M. and Piffer, M. (2023), ‘A new posterior sampler for bayesian structural vector autoregressive models’, *Quantitative Economics* **14**(4), 1221–1250.
- Canova, F. (2011), *Methods for applied macroeconomic research*, Princeton university press.
- Canova, F. and De Nicro, G. (2002), ‘Monetary disturbances matter for business fluctuations in the g-7’, *Journal of Monetary Economics* **49**(6), 1131–1159.
- Canova, F., Kociecki, A. and Piffer, M. (2024), ‘Flexible prior beliefs on impulse responses in Bayesian vector autoregressive models’.
- Caravello, T. E., McKay, A. and Wolf, C. K. (2024), ‘Evaluating policy counterfactuals: A VAR-plus approach’.
- Cascaldi-Garcia, D. (2022), ‘Pandemic priors’.
- Chan, J., Pettenuzzo, D., Poon, A. and Zhu, D. (2025), ‘Conditional forecasts in large Bayesian VARs with multiple equality and inequality constraints’, *Journal of Economic Dynamics and Control* p. 105061.
- Crump, R. K., Eusepi, S., Giannone, D., Qian, E. and Sbordone, A. (2025), ‘A large Bayesian VAR of the US economy’, *International Journal of Central Banking* **21**(2), 351–409.
- Del Negro, M., Giannone, D., Giannoni, M. P. and Tambalotti, A. (2017), ‘Safety, liquidity, and the natural rate of interest’, *Brookings Papers on Economic Activity* **2017**(1), 235–316.
- Giacomini, R., Kitagawa, T. and Read, M. (2022), ‘Narrative restrictions and proxies’, *Journal of Business & Economic Statistics* **40**(4), 1415–1425.
- 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, Technical report, National Bureau of Economic Research.

- Giannone, D., Reichlin, L. and Sala, L. (2004), ‘Monetary policy in real time’, *NBER macroeconomics annual* **19**, 161–200.
- Inoue, A. and Kilian, L. (2026a), ‘The conventional impulse response prior in var models with sign restrictions’, *Journal of Applied Econometrics* .
- Inoue, A. and Kilian, L. (2026b), ‘When is the use of gaussian-inverse wishart-haar priors appropriate?’, *Journal of Policital Economy* .
- Kilian, L. and Lütkepohl, H. (2017), *Structural vector autoregressive analysis*, Cambridge University Press.
- Koch, C. and Noureldin, D. (2024), ‘How we missed the inflation surge: An anatomy of post-2020 inflation forecast errors’, *Journal of Forecasting* **43**(4), 852–870.
- Kociecki, A., Matthes, C. and Piffer, M. (2025), ‘A scalable framework for statistical identification in structural VARs’.
- Koop, G. and Korobilis, D. (2010), ‘Bayesian multivariate time series methods for empirical macroeconomics’, *Foundations and Trends in Econometrics* **3**(4), 267–358.
- Lenza, M. and Primiceri, G. E. (2022), ‘How to estimate a vector autoregression after march 2020’, *Journal of Applied Econometrics* **37**(4), 688–699.
- Neri, S., Buseti, F., Conflitti, C., Corsello, F., Delle Monache, D. and Tagliabracci, A. (2023), ‘Energy price shocks and inflation in the euro area’, *Bank of Italy Occasional Paper* (792).
- Sims, C. A. (1980), ‘Macroeconomics and reality’, *Econometrica* pp. 1–48.
- 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.
- Waggoner, D. F. and Zha, T. (1999), ‘Conditional forecasts in dynamic multivariate models’, *Review of Economics and Statistics* **81**(4), 639–651.

Online Appendix for “Structural forecast analysis”

Davide Brignone and Michele Piffer

A	Derivations	A-2
B	Additional material for the simulation exercise	B-3
C	A comment on changes in sign	D-8
D	Additional material for the application	D-10

A Derivations

Start from equation (1a) of the paper, which we write here for convenience,

$$\mathbf{y}_t = \sum_{l=1}^p \Pi_l \mathbf{y}_{t-l} + \mathbf{c} + \mathbf{u}_t, \quad (\text{A-1})$$

and rewrite it in companion form as

$$\tilde{\mathbf{y}}_t = \tilde{\Pi} \tilde{\mathbf{y}}_{t-1} + \tilde{\mathbf{c}} + \tilde{\mathbf{u}}_t, \quad (\text{A-2})$$

with $\tilde{\mathbf{y}}_t = (\mathbf{y}'_t, \dots, \mathbf{y}'_{t-p+1})'$, $\tilde{\mathbf{c}} = (\mathbf{c}', \mathbf{0}'_{k(p-1)})'$, $\tilde{\mathbf{u}}_T = (\mathbf{u}'_T, \mathbf{0}'_{k(p-1)})'$, $\Pi = [\Pi_1, \dots, \Pi_p]$, $\tilde{\Pi} = [\Pi; [I_{k(p-1)}, 0]]$ (Canova, 2011). Define $\tilde{\Pi}^d = \prod_{g=1}^d \tilde{\Pi}$ with $\tilde{\Pi}^0 = I$. Iterating backwards starting from the forecast made in period T for horizon $T+h$ gives

$$\tilde{\mathbf{y}}_{T+h}^{(T)} = \tilde{\Pi} \tilde{\mathbf{y}}_{T+h-1}^{(T)} + \tilde{\mathbf{c}} + \tilde{\mathbf{u}}_{T+h}, \quad (\text{A-3})$$

$$= \tilde{\Pi}^2 \tilde{\mathbf{y}}_{T+h-2}^{(T)} + [I + \tilde{\Pi}] \tilde{\mathbf{c}} + \tilde{\mathbf{u}}_{T+h} + \tilde{\Pi} \tilde{\mathbf{u}}_{T+h-1}, \quad (\text{A-4})$$

$$= \tilde{\Pi}^3 \tilde{\mathbf{y}}_{T+h-3}^{(T)} + \sum_{d=0}^2 \tilde{\Pi}^d \tilde{\mathbf{c}} + \sum_{d=0}^2 \tilde{\Pi}^d \tilde{\mathbf{u}}_{T+h-d}, \quad (\text{A-5})$$

$\vdots$

$$= \tilde{\Pi}^h \tilde{\mathbf{y}}_T + \sum_{d=0}^{h-1} \tilde{\Pi}^d \tilde{\mathbf{c}} + \sum_{d=0}^{h-1} \tilde{\Pi}^d \tilde{\mathbf{u}}_{T+h-d}, \quad (\text{A-6})$$

$$= \tilde{\Pi}^{h+1} \tilde{\mathbf{y}}_{T-1} + \sum_{d=0}^h \tilde{\Pi}^d \tilde{\mathbf{c}} + \sum_{d=0}^h \tilde{\Pi}^d \tilde{\mathbf{u}}_{T+h-d}, \quad (\text{A-7})$$

$\vdots$

$$= \tilde{\Pi}^{h+\tau} \tilde{\mathbf{y}}_{T-\tau} + \sum_{d=0}^{h+\tau-1} \tilde{\Pi}^d \tilde{\mathbf{c}} + \sum_{d=0}^{h+\tau-1} \tilde{\Pi}^d \tilde{\mathbf{u}}_{T+h-d}, \quad (\text{A-8})$$

$$= \tilde{\Pi}^{h+\tau} \tilde{\mathbf{y}}_{T-\tau} + \sum_{d=0}^{h+\tau-1} \tilde{\Pi}^d \tilde{\mathbf{c}} + \sum_{d=0}^{h+\tau-1} \tilde{\Pi}^d \tilde{B} \tilde{\epsilon}_{T+h-d}, \quad (\text{A-9})$$

where the last equation uses $\tilde{B} = \text{diag}(B, 0_{k(p-1)})$ and $\tilde{\epsilon}_t = (\epsilon'_t, \mathbf{0}'_{k(p-1)})'$. Define $\Phi_{0,h+\tau-1}$ the $k \times k$ matrix on the top-left block of $(\sum_{d=0}^{h+\tau-1} \tilde{\Pi}^d)$. Then define $\{\Phi_{l,d}\}_{l=0}^p$ the p matrices of dimension $k \times k$ forming the top row of matrix $\tilde{\Pi}^d$. Equation (A-9)

can now be rewritten as

$$\begin{aligned} \mathbf{y}_{T+h}^{(T)} &= \sum_{l=0}^{p-1} \Phi_{l+1,h+\tau} \mathbf{y}_{T-\tau-l} + \Phi_{0,h+\tau-1} \mathbf{c} + \\ &+ \sum_{d=h}^{h+\tau-1} \Phi_{1,d} B \epsilon_{T+h-d} + \sum_{d=0}^{h-1} \Phi_{1,d} B \epsilon_{T+h-d}, \end{aligned} \quad (\text{A-10})$$

$$\begin{aligned} &= \sum_{l=0}^{p-1} \Phi_{l+1,h+\tau} \mathbf{y}_{T-\tau-l} + \Phi_{0,h+\tau-1} \mathbf{c} + \\ &+ \sum_{l=0}^{\tau-1} \Phi_{1,h+l} B \epsilon_{T-l} + \sum_{l=1}^h \Phi_{1,h-l} B \epsilon_{T+l}, \end{aligned} \quad (\text{A-11})$$

which is equation (17) in the paper, before substituting $[\Phi, \mathbf{c}, \mathbf{y}, B, \epsilon]$ with $[\Phi^{(T)}, \mathbf{c}^{(T)}, \mathbf{y}^{(T)}, B^{(T)}, \epsilon^{(T)}]$ on the right hand side to indicate that the data, parameters and shocks refer to estimates made in period T . Equation (18) of the paper can be recovered by adjusting the index l in the summation terms to reflect that $\epsilon_{T+1}^{(T+1)}$ are now estimated from the data, i.e.

$$\begin{aligned} \mathbf{y}_{T+h}^{(T+1)} &= \sum_{l=0}^{p-1} \Phi_{l+1,h+\tau} \mathbf{y}_{T-\tau-l} + \Phi_{0,h+\tau-1} \mathbf{c} + \\ &+ \sum_{l=-1}^{\tau-1} \Phi_{1,h+l} B \epsilon_{T-l} + \sum_{l=2}^h \Phi_{1,h-l} B \epsilon_{T+l}. \end{aligned} \quad (\text{A-12})$$

Last, equation (3) can be recovered by removing $^{(T)}$ and setting $h = 0$, $T = t$, which gives

$$\mathbf{y}_t = \sum_{l=0}^{p-1} \Phi_{l+1,\tau} \mathbf{y}_{T-\tau-l} + \Phi_{0,\tau-1} \mathbf{c} + \sum_{l=0}^{\tau-1} \Phi_{1,l} B \epsilon_{T-l}. \quad (\text{A-13})$$

B Additional material for the simulation exercise

The parameter values of the data generating process are set by first specifying the true impulse responses over 12 horizons. As in Canova et al. (2024), we use the following formulation of the Gaussian basis function for each shock j and variable i :

$$\bar{\psi}_{ij,h} = a_{ij} \cdot e^{-\left(\frac{(h-b_{ij})^2}{c_{ij}^2}\right) + \frac{b_{ij}^2}{c_{ij}^2}}. \quad (\text{A-14})$$

The function allows us to span $H + 1$ dynamic responses with only a handful of parameters: a_{ij} captures the impact effect of shock j on variable i , b_{ij} corresponds to the horizon at which the peak effect is reached, and equals 0 if no hump-shaped response is desired, c_{ij} controls for the persistence of the response. Equation (A-14) extends the specification by Barnichon and Matthes (2018).

We specify $a_{11} = a_{21} = 1$, $a_{12} = 0.5$ and $a_{22} = 0.5$. Hence, a one standard deviation positive shock to the first shock increases both variables by 1, while a one standard deviation positive shock to the second shock increases the first variable by 0.5 and decreases the second variable by 0.5. We then set $b_{11} = b_{21} = 0$ and $b_{21} = b_{22} = 2$, so that the first shock generates no hump-shaped patterns, while the second shock generates peak effects two periods after the shock. Last, we set c_{ij} so that the response to the first shock reaches 0.5 three periods after the shock, while for the second shock it leads to a peak effect 20% above the impact effect, in absolute value.

The implied impulse responses are shown in Figure B-1 of the paper. We then use the method by Canova et al. (2024) to compute the following parameters of a SVAR model with 12 lags:

$$B = \begin{pmatrix} 1 & 0.1 \\ 1 & -0.5 \end{pmatrix}, \quad (\text{A-15})$$

$$\Sigma = \begin{pmatrix} 1.25 & 0.75 \\ 0.75 & 1.25 \end{pmatrix}, \quad (\text{A-16})$$

$$Q = \begin{pmatrix} 0.8944 & 0.4472 \\ 0.4472 & -0.8944 \end{pmatrix}, \quad (\text{A-17})$$

$$\Pi_1 = \begin{pmatrix} 1.0362 & -0.1103 \\ -0.1103 & 1.0362 \end{pmatrix}, \quad \Pi_2 = \begin{pmatrix} -0.1185 & -0.0039 \\ -0.0039 & -0.1185 \end{pmatrix}, \quad (\text{A-18})$$

$$\Pi_3 = \begin{pmatrix} -0.0825 & 0.0154 \\ 0.0154 & -0.0825 \end{pmatrix}, \quad \Pi_4 = \begin{pmatrix} -0.0420 & 0.0227 \\ 0.0227 & -0.0420 \end{pmatrix}, \quad (\text{A-19})$$

$$\Pi_5 = \begin{pmatrix} -0.0115 & 0.0157 \\ 0.0157 & -0.0115 \end{pmatrix}, \quad \Pi_6 = \begin{pmatrix} 0.0043 & 0.0027 \\ 0.0027 & 0.0043 \end{pmatrix}, \quad (\text{A-20})$$

$$\Pi_7 = \begin{pmatrix} 0.0086 & -0.0061 \\ -0.0061 & 0.0086 \end{pmatrix}, \quad \Pi_8 = \begin{pmatrix} 0.0066 & -0.0072 \\ -0.0072 & 0.0066 \end{pmatrix}, \quad (\text{A-21})$$

$$\Pi_9 = \begin{pmatrix} 0.0026 & -0.0035 \\ -0.0035 & 0.0026 \end{pmatrix}, \quad \Pi_{10} = \begin{pmatrix} -0.0007 & 0.0004 \\ 0.0004 & -0.0007 \end{pmatrix}, \quad (\text{A-22})$$

$$\Pi_{11} = \begin{pmatrix} -0.0020 & 0.0022 \\ 0.0022 & -0.0020 \end{pmatrix}, \quad \Pi_{12} = \begin{pmatrix} -0.0016 & 0.0018 \\ 0.0018 & -0.0016 \end{pmatrix}. \quad (\text{A-23})$$

Last, the constant term were computed as

$$\mathbf{c} = \begin{pmatrix} 0.3409 \\ 0.4714 \end{pmatrix}. \quad (\text{A-24})$$

Figure B-1: Illustration: true impulse responses






Figure B-2: Illustration with no data revision and true parameters:
composite impulse responses



Note: The dashed blue and red lines in the left and middle plots show the individual composite impulse responses for the shocks in each period from $T - \tau + 1$ to either T (left plots) or $T + 1$ (middle plots). By contrast, Figure 1 in the paper shows the pointwise sum across composite impulse responses.

Figure B-3: Illustration with data revision and estimated parameters:
composite impulse responses



Note: Forecasts and forecast decompositions associated with the OLS estimates. The dashed blue and red lines in the left and middle plots show the individual composite impulse responses for the shocks in each period from $T-\tau+1$ to either T (left plots) or $T+1$ (middle plots). By contrast, Figure 2 in the paper shows the pointwise sum across composite impulse responses.

C A comment on changes in sign


As discussed in the paper, part of the forecast error and forecast revision can be interpreted as the change in the role that the time- T and the time- $T+1$ forecasts associate with shocks before time T . This differential effect is computed as the difference in the impulse responses estimated at time T and at time $T+1$, weighted by the size of the shock that the two forecasts estimate for the period of interest before T . The differential effect of the shocks can follow a very different sign pattern than either of the estimated impulse responses, depending on how the shape of the estimated impulse response changes from T to $T+1$.

As an illustration, suppose for simplicity that the time- T and the time- $T+1$ forecasts are being used to assess the effect that the demand shock at time $T-1$ will still exert on output and inflation over the course of the forecast horizon. Suppose that both forecasts estimate the size of the shock to equal 1, but the actual estimate of the impulse responses changes across forecasts. Consider the top row of Figure C-4, panel A). The blue dotted line illustrates a hypothetical impulse response to a demand shock estimated at time T , while the yellow line reports the estimated impulse response from the time- $T+1$ forecast. Both sets of impulse responses imply that a demand shock moves output growth and inflation in the same direction, hence the shock at $T-1$ is found to exert upward pressure on both variables. Relative to time T , the time- $T+1$ forecast estimates a stronger response of both output growth and inflation. Hence, the contribution of the revision of the shock at time $T-1$ is positive for both output growth and inflation (second row of Figure C-4, panel A).

Things are different for Figure C-4, panel B). We still assume that both forecasts estimate the size of the shock at time $T-1$ to equal 1. However, now the forecast at $T+1$ revises upwards the response of output growth, but downwards the response for inflation. While both forecasts still interpret the shock at time $T-1$ as exerting upward pressure on both output growth and inflation, the marginal role attributed to the revision in the estimate of the role of the time- $T-1$ shock is positive for output growth and negative for inflation.

Figure C-4: Illustration of how the revisions in the role of past shocks can feature different signs compared to the underlying impulse responses

A) No change in signs



B) Changes in signs




Table D-1: Identifying restrictions

	Demand	Supply	Energy	Monetary
Bank rate	+			–
Real GDP growth	+	+	+	+
Inflation	+	–	–	+
Real oil prices	+	+	–	+

Note: The rows report model variables, while the columns indicate the identified shocks. Sign restrictions are imposed only on impact.

D Additional material for the application




Figure D-5: IRFs estimated for 2022Q2

Figure D-6: Estimated structural shocks




Note: The first row shows the marginal posterior distribution of shocks for 2022Q2 estimated over different vintages, with T corresponding to 2022Q2 and $T+h$ the h subsequent quarters. The bottom row shows the same for 2022Q3.

Figure D-7: Forecast analysis for 2022Q4




Note: Pointwise mean forecasts and pointwise mean decompositions. The left-hand side panels show the forecast at time T (solid line), while the middle panels show the new forecast at time $T + 1$ (solid line) along with the one produced at time T (dashed line). Both forecasts are decomposed into the role of the different components up to time T and $T + 1$ respectively: demand (blue bars), supply (red bars), energy (orange bars), monetary policy shocks (yellow bars) and deterministic component (grey bars). The right-hand side panels plot the forecast error and the marginal difference between the forecasts, along with the contribution of each component.

Figure D-8: Series of the shocks over the last year: 2022Q2




Note: Pointwise median and 68% credible sets. Estimates using data up to 2022Q2 are shown in the dashed lines. Estimates using data up to 2022Q3 are shown as shaded band and circled line.

Figure D-9: Series of the shocks over the last year: 2022Q3




Note: Pointwise median and 68% credible sets. Estimates using data up to 2022Q3 are shown in the dashed lines. Estimates using data up to 2022Q4 are shown as shaded band and circled line.

Figure D-10: Decomposition of forecast errors



Note: Pointwise mean of the forecast error and pointwise mean of the decomposition.

Figure D-11: Decomposition of forecast revisions at one year horizon



Note: Pointwise mean of the forecast revision and pointwise mean of the decomposition.