

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Bergholt, Drago; Canova, Fabio; Furlanetto, Francesco; Maffei-Faccioli, Nicolò; Ulvedal, Pål

Working Paper What drives the recent surge in inflation? The historical decomposition roller coaster

Working Paper, No. 7/2024

Provided in Cooperation with:

Norges Bank, Oslo

Suggested Citation: Bergholt, Drago; Canova, Fabio; Furlanetto, Francesco; Maffei-Faccioli, Nicolò; Ulvedal, Pål (2024) : What drives the recent surge in inflation? The historical decomposition roller coaster, Working Paper, No. 7/2024, ISBN 978-82-8379-316-1, Norges Bank, Oslo, https://hdl.handle.net/11250/3172780

This Version is available at: https://hdl.handle.net/10419/310427

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



NC ND https://creativecommons.org/licenses/by-nc-nd/4.0/



WWW.ECONSTOR.EU



7 | 2024

Working Paper

What drives the recent surge in inflation? The historical decomposition roller coaster

Norges Bank Research

Authors: Drago Bergholt Fabio Canova Francesco Furlanetto Nicolò Maffei-Faccioli Pål Ulvedal

Keywords:

Bayesian vector autoregression, deterministic component, single unit root prior, inflation dynamics Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles på epost: <u>servicesenter@norges-bank.no</u>

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail: servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)

ISBN 978-82-8379-316-1 (online)

WHAT DRIVES THE RECENT SURGE IN INFLATION? THE HISTORICAL DECOMPOSITION ROLLER COASTER*

Drago Bergholt[†], Fabio Canova[‡], Francesco Furlanetto[§], Nicolò Maffei-Faccioli[¶] and Pål Ulvedal[∥]

This version: April 8, 2024

Abstract: What drives the recent inflation surge? To answer this question, one must decompose inflation fluctuations into the contribution of structural shocks. We document how whimsical such a historical shock decomposition can be in standard vector autoregressive (VAR) models. We show that the deterministic component of the VAR tends to be imprecisely estimated, making the shock contributions poorly identified under general conditions. Our preferred approach to solve this problem—the single-unit-root prior—can massively shrink the uncertainty around the estimated deterministic component. Once this uncertainty is taken care of, demand shocks unambiguously appear as the main drivers of the inflation surge in the United States, the euro Area, and in four small open economies.

Keywords: Bayesian vector autoregression, deterministic component, single unit root prior, inflation dynamics.

JEL Classification: C11, C32, E32.

^{*}This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We are grateful to Juan Antolin-Diaz for detailed comments on a draft of the paper. We also thank Dante Amengual, Todd Clark, Olivier Coibion, Juan Dolado, Luca Fosso, Domenico Giannone, Michele Lenza, Francesca Loria, Lorenzo Mori, Giorgio Primiceri, Juan Rubio-Ramirez, and seminar participants at Bank of Canada, NTNU, Milan Time Series Seminar, University of Namur, University of Cyprus, Tallinn Tec, the 2022 CEF conference in Dallas, the 2023 SNDE conference in Orlando, the 2023 Barcelona Summer Forum, the 2023 Norges Bank Mini-Workshop in Monetary Economics, the 2023 Dolomiti Macro Meetings in San Candido and the 2023 New Developments in Business Cycle Analysis conference at Danmarks Nationalbank for useful comments.

[‡]BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Corresponding author. E-mail: fabio.canova@bi.no.

 [§]Norges Bank. P.O. Box 1179 Sentrum, 0107 Oslo, Norway. E-mail: francesco.furlanetto@norges-bank.no.
 [¶]Norges Bank. P.O. Box 1179 Sentrum, 0107 Oslo, Norway. E-mail: nicolo.maffei-faccioli@norges-bank.no.

^INord University Business School P.O. Box 2501, 7729 Steinkjer, Norway. E-mail: pal.b.ulvedal@nord.no.

1 INTRODUCTION

MOTIVATION After a long period of low and stable inflation, the outlook has suddenly changed in the aftermath of the COVID pandemic and most countries have seen inflation rates reaching unprecedented levels since the late 1970s. In the US, the GDP deflator started rising at the end of 2020 and peaked at 7.7 percent in 2022:Q2. In the euro area, inflation was in negative territory at the end of 2020 before rising sharply in 2021 and 2022, and peaking around 10 percent on an annual basis. In some European countries the annual inflation rate exceeded 15 percent in 2022. Initially, supply disruptions associated with the pandemic-induced sectoral reallocation of economic activity were thought to cause the rise in inflation. However, the emergence of a historically tight labor market (as measured, e.g., by the vacancy-to-unemployment rate) and the fast recovery in the employment-to-population ratio, as the result of the massive monetary and fiscal stimulus implemented in many countries, shifted attention towards demand factors. Needless to say, disentangling demand and supply factors is crucial since a proper policy response may be highly dependent on the nature of the impulses driving the inflation surge.

A natural way to separate demand and supply shocks is to rely on flexible time-series models such as Vector Autoregressions (VAR). These models are routinely employed for understanding key features of the data and for evaluating the state of the macroeconomic environment, both retrospectively and prospectively (see Crump et al. (2021), Del Negro and Schorfheide (2011), Carriero et al. (2015), Chan (2022) and Koop and Korobilis (2013), among others). Importantly, VAR models decompose observable data into a deterministic and a stochastic component. The deterministic component is regulated by the VAR model's parameters and by the initial values of the system. The stochastic component, in contrast, depends on the discounted sum of innovations to the system. Equipped with some identification restrictions, one can map these (reduced-form) innovations into various structural shocks, with the most fundamental partition being between shocks to demand and supply, respectively. The resulting structural VAR (SVAR) model allows us to decompose observed changes in macroeconomic data over time, into underlying, economic drivers. In particular, such a historical shock decomposition can effectively quantify the relative importance of demand and supply factors during the recent inflation surge.

CONTRIBUTION This paper highlights an important pitfall when computing historical shock decompositions. As a motivating example, suppose that a researcher wishes to quantify the role of demand and supply factors for inflation during the post-COVID period. She can estimate her preferred SVAR model, and then split observed inflation into its deterministic and stochastic components. Our key contribution is to show that there is considerable uncertainty in such an exercise. In particular, we establish that conditional likelihood-based estimation procedures, which fit the model to the covariance structure of the data, deliver very dispersed estimates of the deterministic component: even small perturbations of the estimated parameters of the VAR may result in large changes in the deterministic component, despite having only a minimal impact on the likelihood itself. Put differently, the deterministic component of the model is poorly identified and poorly estimated. Furthermore, because the deterministic and stochastic components sum to data, a poorly estimated deterministic component can translate into very imprecise estimates of

the shock decomposition, resulting in whimsical narratives about economic history. To the best of our knowledge, the "excess" dispersion just described has largely been ignored in the literature, which typically conditions on a point estimate of the deterministic component when historical shock decompositions are computed.

In contrast, most empirical applications of SVAR models report the uncertainty associated with estimated impulse response functions (IRFs). It is nevertheless common to pay attention to a measure of central tendency for the IRFs. For example, Fry and Pagan (2011) select the draw that is closest to the point-wise median IRF, while Inoue and Kilian (2013) choose the most likely admissible draw within the set satisfying the imposed restrictions. We show that draws exhibiting "good" impulse responses—meaning responses close to the point-wise median or to the most likely draw—may feature extreme deterministic components and thus provide unreasonable historical decompositions. In other words, there is no guarantee that "good" impulse responses are associated with "good" historical shock decompositions.

We explain how excess volatility in deterministic component relates to the issue of overfitting discussed by, e.g., Sims (1996), Sims (2000) and Giannone et al. (2019). The two problems are related, yet distinct. In particular, the problem of imprecise shock decompositions is pervasive even when overfitting is not an issue and it appears in SVARs with different number of variables, estimated with samples of different size, and identified with different types of restrictions. The issue appears in a largely homogeneous sample, such as our baseline (1983-2022) sample for the United States and it is exacerbated if we use a longer sample (1949-2022), which is subject to time series instabilities. In fact, the excess dispersion in shock decompositions that we highlight arises during estimation of the reduced-form parameters and is therefore independent, e.g., from the identification scheme used to transform reduced-form residuals into structural shocks, as well as the approach (frequentist or Bayesian) used for inference.

After explaining why an estimated shock decomposition can be imprecise, using a VAR with US data as illustrative examples, we turn our attention to potential solutions. Our preferred approach considers a specific prior distribution, the single-unit-root prior, also known as the "dummy initial observations prior", introduced by Sims (1993). Such a prior effectively shrinks the dispersion of the deterministic components, and allows the econometrician to center it around some reasonable value, e.g. the unconditional mean of the data. The tightness of the prior can be calibrated, if one wishes to do so. We treat it instead as a hyperparameter and estimate its posterior distribution, using the approach of Giannone et al. (2015). The modal value we obtain is rather small, implying tight prior restrictions. As a result, a SVAR estimated with the single-unit-root prior features almost no uncertainty around the deterministic component. Hence, draws featuring similar impulse response functions will also feature similar historical decompositions. While the use of the single unit root prior in empirical applications is not new, see e.g. Antolin-Diaz et al. (2021) and Giannone et al. (2019), its ability to shrink the uncertainty around the estimated deterministic component, therefore improving historical decomposition inference has not been discussed in the literature. Notably, alternative prior assumptions often used in empirical work to reduce parameter uncertainty (e.g., a Normal-Inverse Wishart prior or a Minnesota-like prior) are unable to solve the problem we are concerned with.

RESULTS We estimate a conventional two-variable SVAR with sign restrictions to quantify the drivers of the recent inflation surge. When a diffuse prior is employed, there is substantial uncertainty around estimates of the deterministic component of inflation. In addition, the three draws which are closest to the point-wise median deliver totally different historical decompositions despite exhibiting almost identical impulse responses: one draw implies a prevalent role for supply shocks, one a prevalent role to demand shocks and one a balanced role for both shocks. When the single-unit-root prior is employed instead, draws leading to similar impulse responses are also associated with similar historical decompositions of inflation. Demand shocks are prevalent over the sample and, in particular, over the last few years: in the US they account for more than 50 percent of inflation fluctuations in 2021 and almost 80 percent in 2022. We also provide some international evidence by estimating a SVAR with a single-unit-root prior and the same identification restrictions using data for the euro area and four small open economies (Norway, Sweden, Canada and Australia). Broadly speaking, the dominant role of demand factors in the post-COVID recovery is confirmed. For example, demand shocks account for more than 50 percent of inflation fluctuations in 2022 in the euro area. However, while supply shocks are almost irrelevant in the US since mid 2022, they still matter in the euro area until the mid-2023 as the region is substantially more exposed than the US to the Ukraine war, an important supply shock. In the other countries, the outcomes are surprisingly similar to the US: supply shocks drive inflation in 2020 but already in 2021 demand forces become prevalent, and almost dominant in some cases.

For researchers who are reluctant to use a prior for inferential purposes, we offer two alternative pragmatic solutions make historical decompositions more robust. The first is demeaning the data prior to estimation. Such an approach in part addresses the problem since poor estimates of the VAR constant contribute to making the excess volatility problem important. The second alternative is to construct the distribution of the historical decomposition using draws from the (asymptotic) distribution of the parameters and taking the pointwise median contribution of each shock, as in Bergholt et al. (2023). Such an approach does not reduce the excess volatility of the deterministic component but, at least, takes it into account when computing a summary historical decomposition measure.

RELATED LITERATURE Our work contributes to three separate strands of the literature. First, we complement existing work warning about the mechanical use of SVARs for inference and for stylized fact collection, see e.g. Canova and Ferroni (2022) for a recent contribution. We point out that a careful evaluation of the uncertainty surrounding the deterministic component is crucial to compute credible historical decompositions.

Second, while other papers have studied the behavior of the deterministic component of the data and their pathologies (see e.g. Sims (1996) and Sims (2000)), we believe that our point is novel. One problem discussed in the literature is that the deterministic component tends to explain an implausibly large share of the low frequency variation of the data (the so-called overfitting problem), yielding inaccurate out-of-sample forecasts. Giannone et al. (2019) discuss the issue in detail and propose a prior specification (the prior for the long run), based on long-run theoretical predictions of macroeconomic models, that addresses the problem. In that context, the deterministic component is a measure of central tendency across all draws and its uncertainty is ignored. We highlight that the excess volatility problem, that makes inference about historical decompositions whimsical, is related but separate from the overfitting problem. This means that solving the overfitting problem does not ensure a solution to the excess volatility problem, as the latter may appear even when the former is not an issue.

We also contribute to the recent debate on the drivers of the inflation boom. Bernanke and Blanchard (2024) find that most of the inflation surge resulted from shocks to prices given wages in a simple dynamic model. Beaudry et al. (2024) stress the role of broadbased supply shocks in combination with deviations from rational expectations. Shapiro (2022) provides a decomposition of demand-driven and supply-driven inflation based on sectoral data. Eickmeier and Hofmann (2022) estimate a factor model using a large number of inflation and real activity series and find that the recent inflation dynamics in the US are driven mainly by demand and, to a lesser extent, also by supply factors. Ascari et al. (2023) estimate a Bayesian SVAR on euro area data and find a crucial role for demand factors since the fall 2020. Rubbo (2023) disentangles cross-sectional demand and supply disruptions from aggregate stimulus policies and finds that three quarters of the increase in CPI since 2021 is driven by aggregate demand. Di Giovanni et al. (2023) focus on the impact of fiscal policy on current inflation in a multi-sector model with a network structure while Gagliardone and Gertler (2023) emphasize the role of oil shocks and accommodative monetary policy in a New Keynesian model. Ball et al. (2022) find that the very high labor market tightness in 2021–2022 can explain much of the rise in monthly core inflation, especially during 2022. Our finding that demand factors are important is consistent with the conclusions of most of these papers. We contribute by stressing the importance of obtaining a robust interpretation of the historical decomposition of inflation and provide international evidence on the phenomena.

OUTLINE The paper proceeds as follows: Section 2 illustrates the problem. Section 3 proposes our solution. Section 4 presents estimates of the contribution of demand and supply disturbances for the US, for the euro area and for four other countries. Section 5 discusses the connection between excess volatility and overfitting. Section 6 illustrates two alternative pragmatic solutions to the whimsical historical decomposition problem. Finally, Section 7 concludes.

2 WHIMSICAL HISTORICAL DECOMPOSITIONS

This section explains why the deterministic component in an estimated VAR model may be subject to large uncertainty, and demonstrates the implications this excess volatility has for historical shock decompositions. We present evidence that the issue can be pervasive when VAR models are fitted to commonly used macroeconomic data.

2.1 AN OUTLINE OF THE PROBLEM

To fix ideas, consider the reduced-form VAR model with n variables and p lags:

$$Y_t = C + A_1 Y_{t-1} + A_2 Y_{t-2} + \dots + A_p Y_{t-p} + u_t$$
(1)

 Y_t is the $n \ge 1$ vector of the endogenous variables, $u_t \sim N(0_n, \Sigma)$ is the $n \ge 1$ vector of residuals, $A_1, ..., A_p$ are $n \ge n$ coefficient matrices associated with lagged variables, and

C is the $n \ge 1$ vector of constants. The reduced-form residuals are linear combinations of underlying economic shocks: $u_t = Fe_t$, where F is a matrix restricted by some structural identification scheme. We focus on the system (1), as our problem arises in the reduced-form model irrespective of the particular structural identification restrictions imposed. The role played by structural identification schemes is discussed later.

We find it convenient to rewrite the system in companion form. Define the $np \ge 1$ vectors $\mathbf{Y}_{\mathbf{t}} = (Y_t, Y_{t-1}, ..., Y_{t-p+1})'$, $\mathbf{u}_{\mathbf{t}} = (u_t, 0, ..., 0)'$, $\mathbf{C} = (C, 0, ..., 0)'$, and the $np \ge np$ companion matrix:

$$\mathbf{A} = \begin{bmatrix} A_1 & A_2 & \dots & \dots & A_p \\ I_n & \mathbf{0}_n & \dots & \dots & \mathbf{0}_n \\ \mathbf{0}_n & I_n & \dots & \dots & \mathbf{0}_n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_n & \mathbf{0}_n & \dots & I_n & \mathbf{0}_n \end{bmatrix}$$

With this notation the VAR model is:

$$\mathbf{Y}_{t} = \mathbf{C} + \mathbf{A}\mathbf{Y}_{t-1} + \mathbf{u}_{t}$$
(2)

For the purpose of shock decompositions we express the right-hand side of (2) as the sum of a *deterministic* and a *stochastic* component, DC_t and SC_t :

$$\mathbf{Y}_{\mathbf{t}} = DC_t + SC_t \tag{3}$$

Equation (3) follows from backward substitution, with DC_t and SC_t defined as:

$$DC_t = \left(\mathbf{I} + \mathbf{A} + \dots + \mathbf{A}^{t-1}\right)\mathbf{C} + \mathbf{A}^t\mathbf{Y_0}$$
(4)

$$SC_t = \mathbf{A}^{t-1}\mathbf{u_1} + \dots + \mathbf{A}\mathbf{u_{t-1}} + \mathbf{u_t}$$
 (5)

The deterministic component DC_t can be interpreted as the forecast of Y_t made in period 0, i.e. the counterfactual path of Y_t in the absence of any unforecastable shock realizations during the sample period. DC_t depends on the model parameters (C, A), the initial state vector Y_0 , and time t. The stochastic component SC_t , instead, is a discounted sum of all the realized shocks between period 1 and period t. Since A^t is decreasing in t when the VAR is stationary and ergodic, shocks in the distant past, as well as the initial state Y_0 , play only a negligible role for current observations as t grows.

Conceptually, a historical shock decomposition requires two steps: (a) separating DC_t from SC_t , as in (3), and (b) disentangling the contribution of individual, structural shocks e_t from $u_t = Fe_t$. Our focus is on the first of these two steps. In particular, we establish that conditional likelihood-based estimation procedures, which fit the VAR to the covariance structure of the data, provide a very imprecise estimate of DC_t : even small perturbations to the parametrization of (2) may result in large changes in DC_t despite having only a minimal impact on the likelihood itself. Given that the uncertainty in A determines the uncertainty in the (reduced form) impulse responses, the "excess" uncertainty in the deterministic components is due to the large uncertainty in C.

This is problematic for researchers who want to estimate historical shock decompositions: because the deterministic and stochastic components enter linearly in (3), the estimation uncertainty associated with DC_t must imply an equivalent estimation uncertainty associated with SC_t and the inferred in-sample paths for the residuals u_t . As a consequence, two models with similar impulse response functions may produce very different narratives regarding the realized contribution of shocks to observed fluctuations. It is in this sense that historical shock decompositions can become whimsical.

How pervasive is the issue in practice? A few remarks serve to highlight the generality of the problem. First, given that DC_t is a reduced-form object, the problem occurs regardless of whether the VAR model is estimated with frequentist or flat-prior Bayesian methods. In a Bayesian context, it leads to a wide posterior distribution for the deterministic component. A frequentist, instead, sees it in the form of large asymptotic confidence intervals around the point estimate of the deterministic component. In this paper, we take a Bayesian approach and our preferred solution to whimsical shock decompositions is inherently Bayesian in nature. The alternative solutions we propose in Section 6 are designed for those who want to remain frequentist in their inferential approach. Second, the problem is not restricted to VARs and will emerge also, for example, in other time series specifications such dynamic factor models. Third and related, to the extent that dynamic stochastic general equilibrium (DSGE) models involve estimated parameters that enter the steady state, the issue may be relevant also for these models. After all, the DSGE solution generally admits a time series representation consistent with equations (2)-(5), and it is this representation that is confronted with data. Fourth, whimsical shock decompositions occur even when $A^t Y_0$ tends to zero, suggesting that data transformations inducing stationarity do not remove the problem either. In fact, we show that the problem is present even if the VAR features variables which are mildly persistent and do not display strong low frequency movements.

2.2 AN ILLUSTRATION – A BIVARIATE VAR

To illustrate the relevance of the problem, we fit a bivariate VAR using US data on real GDP and the GDP deflator, both measured in log first differences. This is arguably the simplest possible setup to quantify structural drivers of inflation. Our baseline sample runs from 1983:Q1 to 2022:Q4. We include four lags and estimate the model with Bayesian methods. Following the literature, we employ a diffuse prior (Jeffreys, 1946) on the VAR parameters and on the covariance matrix Σ .

We use standard sign restrictions to identify two structural shocks. As in Canova and De Nicolò (2002), an aggregate demand shock is assumed to move GDP and inflation in the same direction, while an aggregate supply shock moves the two variables in opposite directions. These sign restrictions are imposed only on impact. Clearly, our demand shock bundles together various demand side disturbances including monetary and fiscal shocks (see Bianchi et al. (2023)). Similarly, the supply shock we identify is likely a mix of innovations to productivity, commodity prices, markups, and other supply-side factors. Kabaca and Tuzcuoglu (2023) disentangle an array of different structural shocks in a unified framework. As our focus is on the role of the deterministic component, a bivariate model seems sufficient.

Figure 1 reports impulse responses from the estimated structural model. We measure time in quarters on the x-axes and the estimated responses in percentage points on the y-axes. For each horizon, the black line represents the point-wise median response across 1000 draws from the estimated posterior. The shaded areas instead represent the 68% posterior credible set and the min-max identified responses (as in Baumeister and



Figure 1: Responses to identified aggregate demand and an aggregate supply shocks.

Note: The black line is the point-wise median and the shaded areas the 68 percent and min-max identified sets. The red line is the impulse response for the draw that is closest to the point-wise median; the green and blue lines are impulse responses for 2nd and 3rd draws closest to the point-wise median, respectively.

Hamilton (2015)), respectively. The point-wise median response is compared in Figure 1 with three alternatives. As the point-wise median most likely mixes different draws across horizons, applied researchers often prefer to analyze the impulse responses obtained with one particular draw instead. Following Fry and Pagan (2011), a popular choice of central tendency is the single draw, among the 1000 in our case, that is closest to the point-wise median on average. This draw is reported as the red line in Figure 1. For future purposes, we also report the impulse responses obtained from the draw second closest to the point-wise median in green, and the third closest draw in blue.

A couple of remarks are in place. First, the output growth response is short-lived while inflation responses are more persistent. Note that both shocks are allowed to have a permanent effect on the level of GDP. Second, the impulse responses from the three draws are almost indistinguishable from each other and very close to those produced with the point-wise median. Given that impulse responses are similar, one may naturally expect these draws to produce other summary statistics which are similar as well, such as the historical decomposition of shocks.

In Figure 2 we gauge whether this is the case. Panel (a) plots the historical shock decomposition of inflation over the last eight years of the sample, conditioning on each of the three draws closest to the point-wise median impulse response. The blue bars correspond to the contribution of the deterministic component, while the red and yellow bars represent the contribution of supply and demand shocks, respectively. Combined, the contributions of supply and demand sum to the stochastic component SC_t . Interestingly,



Figure 2: Historical decompositions and deterministic components of US inflation

the three draws closest to the point-wise median impulse response produce vastly different narratives about the role of shocks during the post-pandemic inflation surge: the model draw closest to the median implies that supply shocks explain more than two-thirds of the rise in inflation; the second closest draw that demand shocks account for more than two-thirds; while the third closest draw suggests that demand and supply account for about 50% each. These differences are striking given that the three draws imply almost identical impulse responses. To shed light on the underlying causes, panel (b) in Figure 2 plots the deterministic components of all 1000 draws on the left, and the 10 closest to the point-wise median on the right. Interestingly, the posterior dispersion of deterministic components is substantial even for the 10 draws closest to the point-wise median responses, and paths stabilize at very different levels.

2.3 How pervasive is the roller coaster?

To document the generality of the problem, we show that whimsical shock decompositions are produced in various contexts.



Figure 3: Estimated deterministic components of inflation (all posterior draws).

2.3.1 The role of structural identification assumptions

As a first exercise, we illustrate the irrelevance of structural identification assumptions, as the uncertainty in the deterministic component arises in the reduced-form model. To this end, let us consider two of the most popular identification strategies: the Blanchard-Quah decomposition, which imposes zero restrictions on the long-run cumulative sum of VAR coefficients (Blanchard and Quah, 1989), and the Cholesky decomposition, which imposes zero-restrictions on the impact coefficients, see (Sims, 1986). The Blanchard-Quah decomposition in particular is commonly used to separate demand from supply

under the assumption that demand shocks cannot affect the level of output.¹

Indeed, Blanchard-Quah, Cholesky and sign-restricted decompositions share the same 1000 draws from the reduced-form system, as shown in Figure 2 and in the first row of Figure 3. As such, the uncertainty in the deterministic component is equally present for all three identification schemes. Moreover, as with sign restrictions, there are notable differences in the historical decomposition even across the three closest draws. This is true both for Blanchard-Quah and Cholesky, see Figure A-2 in the Appendix. Irrespective of the identification scheme, any flat-prior VAR is in principle susceptible to the problem highlighted here.

2.3.2 The role of reduced-form priors

Given that the uncertainty we presented has to do with estimates of the VAR parameters, one may suspect that prior restrictions routinely used to reduce parameter uncertainty also may shrink the uncertainty in the deterministic component. For illustration, we consider a standard conjugate Normal-Inverse Wishart prior and a Minnesota-like prior (see Doan et al. (1984)).

In the Normal-Inverse Wishart case, the prior for the autoregressive (AR) coefficients is normal, centered at zero with a diagonal covariance matrix of 10, while the prior for the covariance matrix of the residuals is inverse Wishart with a unitary diagonal matrix as scale and 3 degrees of freedom. The Minnesota-like prior is also normally distributed for all AR coefficients and is centered around zero for all parameters, including the variables' first own lag, as VAR variables appear in growth rates. We set the hyperparameters to standard values, for example, the lag-decaying parameter is chosen to be 2. The only exception is the tightness parameter, which we treat as a random variable. The diagonal elements of the scale matrix of the inverse Wishart prior are set to the residual variance of an AR(1) process for each variables. In both cases, we use sign restrictions to identify the shocks.

As shown in the second row of Figure 3, the dispersion of the deterministic components remains large (see also Figure A-3 in the Appendix for the resulting whimsical historical decompositions). The reason for why these standard priors are unable to reduce the uncertainty around the deterministic components is that they are designed to shrink the uncertainty in AR coefficients and the covariance matrix of the disturbances, but leave the prior on the constant diffuse.

2.3.3 THE ROLE OF THE SAMPLE SIZE

In our baseline estimation, we consider a relatively homogeneous sample period starting in 1983. We re-estimate the model over the period 1949:Q1 - 2022:Q4 with diffuse priors and identify shocks with the same sign restrictions. Estimates of the deterministic component of inflation are presented in the third row of Figure 3 (the historical decompositions in Figure A-4 in the Appendix).

¹Note that neither the Blanchard-Quah decomposition, nor the Cholesky decomposition guarantee sensible signs of the price and quantity responses to shocks. Indeed, Figure A-1 in the Appendix shows that Blanchard-Quah delivers co-movement of the two variables in response to a supply shock. See Furlanetto et al. (2024) for an application that addresses this issue by combining Blanchard-Quah with sign restrictions.

The dispersion of the deterministic component is substantially larger than in the baseline model. Thus, the massive uncertainty around the deterministic component is an additional argument for not using a long and dis-homogeneous sample when estimating a SVAR model (see e.g. Furlanetto et al. (2024) for a detailed discussion of this point).

2.3.4 The role of the VAR dimension

Finally, we inspect the role played by the dimensionality of the VAR system. Canova and Ferroni (2022) have shown that small-scale VARs may be prone to deformation problems, making it desirable to estimate models of a certain size to ensure that structural objects such as impulse responses make economic sense. Here, we estimate a 5-variable VAR model and study how large the dispersion of the deterministic components is in such a medium-scale system. Importantly, since the problem is associated with the reduced-form model, we do not impose any structural identification restrictions on the system.

The following variables are used for estimation: real GDP, the GDP deflator, real gross private domestic investment, the Federal funds rate, and real wages. The latter is measured as average hourly earnings of production and non-supervisory employees, deflated with the GDP deflator. All variables except the Federal funds rate are measured in growth rates. The sample is 1983:Q1-2022:Q4, as in our baseline setup.

The deterministic components of inflation are presented in the third row of Figure 3 (see also Figure A-5 for the deterministic components of each variable in the VAR). Evidently, the problem of excessive dispersion in the deterministic component remains. If anything, the problem seems exacerbated and becomes rather dramatic for variables like inflation or the federal funds rate.

3 SOLVING THE PROBLEM

Having established how and why an estimated shock decomposition may be subject to undesirable uncertainty in the deterministic component, we next turn to possible solutions to the problem. In a nutshell, we propose a way to reduce the dispersion in DC_t across draws. To this end we exploit the single-unit-root prior, also known as the "dummy initial observations prior", first introduced by Sims (1993) and recently used in Antolin-Diaz et al. (2021). We show that, when such a prior is properly set, it dramatically restricts the posterior excess volatility in the estimated deterministic component. Given that uncertainty in the stochastic component remains, our prior effectively pins down the constant of the estimated VAR model.

3.1 PRIOR SHRINKAGE OF THE DETERMINISTIC COMPONENT VOLATILITY

As in Sims (1993), we add an artificial observation to the beginning of the sample, where both the current and lagged data are given by $\frac{1}{\delta} \bar{\mathbf{Y}}_0$ and the intercept is set to $\frac{1}{\delta}$. The idea is that, if lags of \mathbf{Y}_t are at a certain value $\bar{\mathbf{Y}}_0$, then also \mathbf{Y}_t tends to be close to $\bar{\mathbf{Y}}_0$. δ is a non-negative hyper parameter while $\bar{\mathbf{Y}}_0$ is a vector of size $np \ge 1$. The stochastic constraint imposed by our artificial observation on the VAR model can be written as

$$[\mathbf{I} - \mathbf{A}] \, \overline{\mathbf{Y}}_{\mathbf{0}} - \mathbf{C} = \delta \mathbf{u}_{\mathbf{0}},\tag{6}$$

where $\mathbf{u_0} = (u_0, 0, ..., 0)'$ is a vector of size $np \ge 1$ (see Miranda-Agrippino and Ricco (2019) for details). δ governs the tightness of this constraint. At one extreme, the constraint becomes uninformative when δ approaches infinity. At the other, as δ approaches zero, the constraint implies either (i) a system which contains at least one (common) unit root, with C = 0, or (ii) a stationary VAR with $C \neq 0$.

Importantly, if \mathbf{u}_0 is a random variable with zero mean and unit variance, and dispersion regulated by δ , then the constraint in (6) restricts the unconditional mean of the model, much in the same way as the steady state prior of Villani (2009).² In particular, the dummy observation restricts the unconditional mean of \mathbf{Y}_t , $(\mathbf{I} - \mathbf{A})^{-1}\mathbf{C}$, to be centered at $\mathbf{\bar{Y}}_0$ when δ approaches zero. To appreciate why this is useful, we manipulate equation (4) and substitute out \mathbf{C} , using the stochastic restriction (6) and arrive at the following expression:

$$DC_t = \left(\mathbf{A}^t (\mathbf{Y}_0 - \bar{\mathbf{Y}}_0 + (\mathbf{I} - \mathbf{A})^{-1} \delta \mathbf{u}_0) + \bar{\mathbf{Y}}_0 - (\mathbf{I} - \mathbf{A})^{-1} \delta \mathbf{u}_0\right)$$
(7)

As long as the system is stationary and ergodic, $\mathbf{A}^t(\mathbf{Y_0} - \bar{\mathbf{Y}_0} + (\mathbf{I} - \mathbf{A})^{-1}\delta \mathbf{u_0}))$ approaches zero as the sample size grows large. Thus, equation (7) implies that the deterministic component of the data a-priori fluctuates around $\bar{\mathbf{Y}_0}$, with δ regulating the tightness of this restriction. If $\delta \rightarrow 0$, the deterministic component will be fixed a-priori at $\bar{\mathbf{Y}_0}$, effectively eliminating any posterior (or asymptotic) dispersion. Notably, this happens regardless of \mathbf{A} and Σ . The stochastic component will still be uncertain—different draws of \mathbf{A} and F will still potentially affect the estimated, historical importance of different shocks for inflation. However, this type of estimation uncertainty is the same type of uncertainty as that present in the impulse responses, making the two statistics consistent.

Clearly, the choice of $\bar{\mathbf{Y}}_0$ is important for the dynamics and convergence properties of DC_t . Note that, because the single unit root prior restricts the unconditional mean of the estimated VAR, it indirectly speaks to the overfitting problem as well by potentially reducing the forecasting ability of the deterministic component for the low frequency movements in data. Based on this observation, it seems natural to set $\bar{\mathbf{Y}}_0$ to the observed sample average, so that DC_t does not deviate too much from its long run level.

We would like to highlight two important points. First, while Karlsson (2013) suggests that the single unit root prior could be used to incorporate information about the unconditional mean by centering the restriction at the expected steady state (see footnote 8, page 814), he does not use it to solve either the overfitting problem nor the excess volatility problem discussed in this paper. Antolin-Diaz et al. (2021) and Giannone et al. (2019), on the other hand, discuss how the single unit root prior may get rid of the predictability problem by restricting the steady state of the endogenous variables implied by the VAR. Still, they do not tackle the excess volatility problem we study here, nor do they show how a single unit root prior could take care of it. Second, although we set $\overline{Y_0}$ to the sample average, in the spirit of empirical Bayes approaches, one could also restrict it

²The endogenous prior suggested by Jarocinski and Marcet (2019) is not necessarily suited for our purposes, unless the specification is complemented by a hierarchical structure which restricts the unconditional mean in the appropriate way.

with an informative prior based on, say, an economic model, along the lines of Del Negro and Schorfheide (2004). In this case a hierarchical structure would be necessary, but the computations could still be manageable.

3.2 Some simulations

We perform a few simulation exercises to highlight how the stochastic restriction imposed by (6) affects the dispersion of the deterministic component. First, we compare the estimation uncertainty in DC_t with and without the single unit root prior. Second, we zoom in on the role played by δ . In both exercises we generate data from a bivariate VAR(1):

$$Y_t = C + AY_{t-1} + u_t,$$

The error term consists of two independently and normally distributed shocks with unit variance; $u_t \sim N(0, I)$. We calibrate the VAR coefficients as follows:

$$C = \begin{pmatrix} 0.4\\ 0.5 \end{pmatrix}, \quad A = \begin{pmatrix} a_{1,1} & -0.3\\ 0.3 & 0.4 \end{pmatrix}$$

We consider different degrees of persistence as well as different sample sizes. The role of persistence is assessed by varying $a_{1,1}$, the top-left element of A. For the less persistent processes we set $a_{1,1} = 0.6$; for more persistent processes we set $a_{1,1} = 0.95$. The constant term C is the same in both cases. We compare samples with 80, 150, and 500 observations respectively.

3.2.1 The implications of the single-unit-root prior

In the first exercise, the estimation step is done, first, with a diffuse prior, as in Section 2.2, and, second, with the single unit root prior. We report the deterministic component for 1000 parameter draws from the estimated posteriors to measure the magnitude of the dispersion in DC_t .

The estimated deterministic components are shown in Figure 4. We limit our attention to the first variable in the system. Panel (a) plots DC_t using a diffuse prior and panel (b) using the single-unit-root prior. Panel (a) illustrates that the more persistent is data, or the shorter is the sample, the greater will be the dispersion in the deterministic component. Nevertheless, the dispersion is substantial in all cases.

The intuition for this result is simple. When the process is more persistent or the sample size is smaller, estimates of the AR coefficients are downward biased. This bias adds to the volatility of the estimated C, making estimates of the deterministic component "excessively" volatile.

The comparison with Figure 4b is striking: after about 20 periods, all 1000 DC_t draws converge to virtually identical values, irrespective of the data persistence or the sample size. Because they all converge to $\bar{\mathbf{Y}}_0$, shock decompositions will be consistent across posterior draws. The key behind this result is the discipline imposed by δ .

3.2.2 The role played by δ

The role played by δ can be assessed by comparing the deterministic component assuming $\delta = 0.2$, the "rule-of-thumb" value used in some existing literature (see, for example,



Figure 4: Estimated deterministic components of $Y_{1,t}$.

(a) Diffuse prior

Sims and Zha (1998)), and $\delta = 0.0001$, which implies that the stochastic restriction (6) holds almost exactly. In both situations, we limit our attention to the case $a_{1,1} = 0.95$ and T = 150 observations for illustration.

Results are presented in Figure 5. Clearly, the value of δ is crucial to shrink the dispersion in the estimated deterministic component: with $\delta = 0.2$ an important uncertainty remains, with the long-run values settling at vastly different estimates. This may suggest that the "rule-of-thumb" value of 0.2 can be highly problematic for researchers interested in estimating shock decompositions. $\delta = 0.0001$, instead, forces all 1000 posterior draws to share the same value after relatively few periods. As we will see, the low value of δ used here is consistent with the modal value obtained using data for many countries.



Figure 5: Estimated deterministic component of $Y_{1,t}$, different δ .

4 WHAT DRIVES THE POST-PANDEMIC INFLATION SURGE?

In this section, we revisit the drivers of the recent inflation surge across countries. To this end, we estimate the baseline, bivariate SVAR, restricted with the single-unit-root prior and estimated over the baseline sample (1983:Q1-2022:Q4). We first estimate the model on US data, and then extend the analysis to other countries.

Given the important role of δ for the dispersion in deterministic components in the simulation exercises, we choose to estimate the posterior of this hyper-parameter and to let the data determine the tightness of the stochastic constraint. We choose a Gamma density on δ with mode equal to 1.

4.1 The post-pandemic US inflation surge

Figure 6 documents the outcomes when we estimate the SVAR model on US data, applying the single-unit-root prior. Panel (a) reports selected historical shock decompositions of inflation, while the deterministic components are provided in panel (b).

As a first result, the modal value of δ returned by the algorithm of Giannone et al. (2015) is very low (0.0001), and its posterior dispersion is negligible (see Figure A-6). Put differently, US data favor a posterior model where virtually all VAR parameterizations share the same unconditional mean value.

Second, the single-unit-root prior implies that model draws featuring similar impulse responses, will now also feature similar historical shock decompositions. As an example, the three draws closest to the point-wise median response now produce a coherent narrative about post-COVID inflation drivers: as shown in Figure 6, supply factors were important in the initial phase of the inflation surge, but demand factors became the main drivers ever since 2021. They account for 56 percent of the inflation surge in 2021 and 77 percent in 2022. On average over the post-pandemic period, demand factors explain more than 50 percent of the movements in inflation.

This result is consistent with a number of recent papers in the literature supporting the



Figure 6: US: SVAR estimated using the single-unit-root prior.

(b) Deterministic components of inflation



prevalent, if not dominant, role of demand shocks in driving the inflation surge (cf. Eickmeier and Hofmann (2022), Ascari et al. (2023), Di Giovanni et al. (2023), Rubbo (2023) and Gagliardone and Gertler (2023) among others). Therefore, our distinctive contribution is not in recovering the key role of demand factors. Instead, our main message is that applied researchers should carefully deal with the uncertainty in the deterministic component prior to the computation of historical decompositions. The prevalence of demand shocks emerges clearly only when this uncertainty is massively reduced.

We would like to reiterate that the single-unit-root prior does not eliminate uncertainty around the estimated shock decomposition. However, the remaining uncertainty is of the same kind that is associated with the estimation of impulse response functions. In contrast, the uncertainty arising from the deterministic components becomes largely irrelevant, as shown in panel (b) of Figure 6.

4.2 INTERNATIONAL EVIDENCE

Although the recent surge in inflation has been a global phenomenon, it is far from obvious that the role of demand and supply factors should be similar across countries. For example, the fiscal stimulus during the COVID pandemic has been particularly pronounced in the US. In addition, European countries have been more exposed to the economic consequences of the Russian invasion of Ukraine. Given these differences, one could for example expect a larger role for demand factors in the US than in the euro area.

Figure 7: Euro area: SVAR estimated using the single-unit-root prior.



(a) Historical decomposition of inflation for top 3 draws

Therefore, it is of independent interest to estimate our SVAR model for the euro area. Because the sample is shorter, we use data on industrial production (rather than on GDP) and HICP inflation at a monthly frequency in the exercise. The sample period goes from 2001:M1 to 2023:M3. As with US data, the model features a large dispersion of deterministic components across draws (see Appendix B) when a diffuse prior is used. If anything, the problem is exacerbated. Yet, when the model is estimated with the single-unit-root prior, the uncertainty around the deterministic component is substantially reduced, as shown in Figure 7. Demand and supply factors contributed more or less equally to the recent inflation surge, with a more prevalent role for demand factors in 2022. Thus, while supply factors played a role only until mid 2022 in the US, they explain a substantial



Figure 8: Deterministic components of inflation in selected countries. Top 10 draws.

share of inflation in the euro area also in the latest part of the sample. This outcome is far from surprising, as the US are largely insulated from the consequences of the Ukraine war. Nonetheless, demand factors are prevalent also in the euro area, at least since the second quarter of 2022.

To provide further evidence on the inflation surge, we estimate the same SVAR model using data for four small open economies: Norway, Sweden, Canada and Australia. In this case, we use quarterly data on real GDP growth and year-on-year CPI inflation over the period 1993:Q1–2023:Q2 for all countries. The single unit-root-prior shrinks the uncertainty around the deterministic component in all cases, as illustrated in Figure 8. As shown in Figure 9, the role of demand factors is prevalent, if not dominant, during the recent period in all countries. All in all, despite the large heterogeneity in the policy responses to the COVID 19 pandemic, the different exposure to supply bottlenecks and to the war in Ukraine across countries, demand factors seem to drive the recent inflation surge in all the countries we considered.

5 OVERFITTING VS. EXCESS VOLATILITY

It is useful to evaluate the connection between the well-known overfitting behavior of deterministic components and the more novel excess volatility problem studied in this paper. As discussed in detail in Sims (1996), Sims (2000) and Giannone et al. (2019), flat-prior VARs tend to attribute an implausibly large share of the variation in observed time series to their deterministic components. The problem arises, in the case of stationary variables, when initial values of the variables are distant from their steady state. Complex transient dynamics from these initial conditions to the steady state are common in VAR models. Using the terminology of (Sims, 2000), this marked *temporal heterogeneity* can



Figure 9: Historical decompositions of inflation in selected countries

be exacerbated if the variables included in the VAR are persistent or display non-trivial low frequency components. In these cases, the overfitting problem persists over time and the estimated deterministic component tends to track the observed time series. Examples of these dynamics (often hump shaped) are presented in Sims (2000) and Giannone et al. (2019) and are featured in particular by labor market variables like hours worked and unemployment, which display important medium term cycles. The excess volatility discussed in this paper relates to the *uncertainty* around the estimated deterministic components and not to their *level* which was the focus in the previous literature. This is important because excess volatility can easily manifest itself even when the overfitting problem is relatively minor. As shown in Figures 2 and 3, the temporal heterogeneity discussed above is clearly present at the beginning of our sample.³ However, almost all draws stabilize after just two or three years in contrast, for example, with Sims (2000) and Giannone et al. (2019) where draws feature a clear hump-shaped behavior for many variables. Thus, our simple bivariate model estimated with a diffuse prior features a mild overfitting problem and a rather extreme excess volatility problem.

While clearly different, overfitting and excess volatility share some similarities: they are both more pronounced in large systems and with persistent variables. In addition, both pathologies distort historical decompositions and should be taken good care in applied investigation. The single-unit-root prior was originally introduced to alleviate the overfitting problem. Our contribution is to show that it is even more successful at dealing with the excess volatility problem.

³While we use variables in growth rates rather than in levels (level variables are typically considered when discussing the issue of overfitting), one could expect the problem to arise also in our system. The AR(1) coefficients of output growth and inflation are 0.03 and 0.84, respectively, and 14.7 and 47.2 percent of the total variance of these variables is found in the low frequency spectrum (which we define as corresponding to cycles of 32 to 64 quarters). Furthermore, while the unconditional sample means are 0.41 and 0.79 respectively, the initial values of the series are 1.62 and 0.27.

6 ALTERNATIVE APPROACHES

For those researchers who are reluctant to use priors for inferential purposes, we offer two alternative pragmatic solutions that can go a long way in making the uncertainty in deterministic components less relevant for practical purposes. The first is demeaning the data prior to estimation. Such an approach is partially successful since poor estimates of the VAR constant contribute to make the problem important. The second alternative is to construct a *median* historical decomposition taking into account parameter uncertainty.

6.1 DEMEANING THE DATA

As discussed in Section 2, the deterministic component consists of two terms:

$$(I + \mathbf{A} + \mathbf{A}^2 + \dots + \mathbf{A}^{t-1})\mathbf{C} + \mathbf{A}^t\mathbf{Y_0}$$

If the VAR constant is zero, the first term disappears and, as t increases the last term goes towards zero, as long as the data is stationary and ergodic. Thus, a simple way to restrict the deterministic component to be similar across draws in finite samples is to demean the data and estimate the VAR without a constant.

We present the historical decomposition obtained in this case using US data in Figure 10. In this specification, the role of demand shocks is slightly more prevalent than with a single-unit-root prior. Note, however, that the uncertainty in the deterministic component is still larger than in Figure 6. By demeaning and estimating the VAR without a constant, the deterministic component is forced to be on the path described by A^tY_0 . Therefore, draws that imply different A's will force the estimated shocks to take different values, altering their contribution in a given historical episode.

6.2 COMPUTING THE MEDIAN HISTORICAL DECOMPOSITION

Another pragmatic approach consists in acknowledging that the uncertainty present in the deterministic components matters. To do this, one could compute a large number of historical decompositions, each associated with a draw for the parameters, and recover the point-wise median historical decomposition at each point in time. In this case, at each quarter, the data is decomposed into the median contribution of demand shocks, the median contribution of supply shocks and a residual deterministic component that absorbs the difference between data and the two median stochastic components. While such a summary measure mixes information from different draws and does not preserve the additivity property of historical decompositions, it features the non-negligible advantage of taking into account the uncertainty in the deterministic component into the inferential process while insuring robustness against non-normalities in the distribution of estimates. Bergholt et al. (2023) use such an approach when studying the causes of the reduced sensitivity of inflation to measures of economic slack in the pre-COVID period.⁴

Figure 11 computes the median historical decomposition for the four prior distributions we considered in this paper using US data. First, notice that the results are quite

⁴One could also compute the mean historical decomposition. Such a summary statistics preserves the additive property but it is not robust to the presence of outliers, which are present when a flat prior is used (see Figure B-3 in the Appendix).



Figure 10: US: Estimates without constant (demeaned data) and a diffuse prior.

(a) Historical decomposition of inflation for top 3 draws

similar across specifications. In addition, in all cases considered, demand shocks have an important role during the current inflation surge. Finally, by comparing the bottom-right panel of Figure 11 and panel a) in Figure 6 one can conclude that the point-wise median historical decomposition and the historical decomposition obtained with the single-unit-root prior provide a similar reading of the current inflation surge.

6.3 A MEASURE OF DISPERSION FOR THE HISTORICAL DECOMPOSITIONS

The evidence presented in Figure 11 may lead the reader think that, conditional on using the point-wise median historical decomposition, the choice of priors becomes irrelevant. After all, the reading of history is very similar across the panels of Figure 11. However, such a conclusion is premature. In fact, while median estimates are similar across prior specifications, the uncertainty around these median estimates can be substantially different.

Since it is hard to visually determine to what extent historical decompositions across draws are different, we propose a simple measure which could be used for the purpose. We construct it using the dispersion of the historical decompositions among the 100 draws



Figure 11: Historical decomposition of inflation, point-wise median.

(a) Diffuse prior

(b) Normal-Inverse Wishart prior

that are closest to the point-wise median IRFs (from now on referred to as the top 100 draws). Let $S_{i,j,t}$ be the contribution from the shock j to variable i in period t. Then

$$D_{i,j,t} = \max(S_{i,j,t}) - \min(S_{i,j,t})$$
(8)

$$M_{i,j} = \frac{1}{T} \sum_{t=1}^{T} D_{i,j,t}$$
(9)

 $D_{i,j,t}$ is the difference between the maximum and minimum value of $S_{i,j,t}$ among the top 100 draws for variable *i* and shock *j*, and $M_{i,j}$ is the average dispersion. For illustration, focus on the dispersion in the contribution from the demand shock to inflation for the period 2020:Q2 to 2022:Q4. Table 1 reports this statistics for the four prior specifications, and three different identification schemes. There are considerable differences for all priors. However, the dispersion becomes uniformly smaller when the single-unit-root prior is imposed, regardless of the identification scheme employed.

7 CONCLUSIONS

Most empirical applications of SVAR models report the uncertainty around impulse response functions. We show that it is also important to take into account the uncertainty

	Sign	Blanchard-Quah	Cholesky
Diffuse	1.07	0.88	2.33
Normal-Inverse Wishart	1.53	1.20	0.91
Minnesota	0.87	0.71	0.61
Single-unit-root	0.68	0.48	0.54

Table 1: Dispersion of historical decompositions, different priors and different identifications.

Note: The numbers report the measure of dispersion in the historical decompositions across the 100 draws closest to the point-wise median response, see Equation 9.

around historical decompositions. This kind of uncertainty is heavily affected by the uncertainty present in the deterministic component of the VAR. Posterior draws that exhibit similar impulse responses can feature largely different deterministic components. Draws associated with extreme deterministic components exhibit a distorted historical decomposition because shocks have to match the gap between the series of interest and the estimated deterministic component. Hence, there is no guarantee that "good" impulse response functions are related to "good" historical decompositions. We have provided one solution to the problem based on the use of the single-unit-root prior and two pragmatic alternatives based on demeaning the data and on computing a point-wise median historical decomposition. Regardless of the approach one uses to take this uncertainty into account, we find that the recent inflation surge observed in the US is mainly driven by demand factors. The prevalence of demand shocks is confirmed, and in some cases reinforced, when using data for Norway, Sweden, Canada and Australia. The contribution is generally more balanced in the euro area where supply factors are still relatively important at the end of our sample.

The problems we highlight in this paper also appear when one wants to separate permanent from transitory components. The excess volatility problem in this case becomes even more acute because variations in the deterministic components matter for the level of the permanent component of the data. Thus, large uncertainty in the former translates into a large uncertainty in the latter. In addition, since in small samples, the long run contribution of the shocks is very imprecisely estimated, additional uncertainty may affect what is perceived as permanent and as transitory. This makes estimates of trend inflation generally very poor and inflation regimes hard to characterize. We plan to study these issues in future research.

REFERENCES

- Antolin-Diaz, J., Petrella, I., and Rubio-Ramírez, J. F. (2021). Dividend momentum and stock return predictability: A Bayesian approach. CEPR Discussion Paper 16613.
- Ascari, G., Bonomolo, P., Hoeberichts, M., and Trezzi, R. (2023). The euro area great inflation surge. *SUERF Policy Brief 548*.
- Ball, L. M., Leigh, D., and Mishra, P. (2022). Understanding US inflation during the COVID era. *Brookings Papers on Economic Activity*, pages 1–54.
- Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999.
- Beaudry, P., Hou, C., and Portier, F. (2024). The dominant role of expectations and broadbased supply shocks in driving inflation. CEPR Discussion Paper 18963.
- Bergholt, D., Furlanetto, F., and Vaccaro-Grange, E. (2023). Did monetary policy kill the Phillips curve? Some simple arithmetics. Norges Bank Working Paper 2/2023.
- Bernanke, B. and Blanchard, O. (2024). What caused the us pandemic-era inflation? *American Economic Journal: Macroeconomics, forthcoming.*
- Bianchi, F., Faccini, R., and Melosi, L. (2023). A fiscal theory of persistent inflation. *The Quarterly Journal of Economics*, page 2127–2179.
- Blanchard, O. J. and Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review*, 79(4):655–673.
- Canova, F. and De Nicolò, G. (2002). Monetary disturbances matter for business cycle fluctuations in the G7. *Journal of Monetary Economics*, 46:1131–1159.
- Canova, F. and Ferroni, F. (2022). Mind the gap! Stylized dynamic facts and structural models. *American Economic Journal: Macroeconomics*, 14(4):104–35.
- Carriero, A., Clark, T. E., and Marcellino, M. (2015). Bayesian VARs: specification choices and forecast accuracy. *Journal of Applied Econometrics*, 30(1):46–73.
- Chan, J. C. (2022). Asymmetric conjugate priors for large bayesian VARs. *Quantitative Economics*, 13(3):1145–1169.
- Crump, R. K., Eusepi, S., Giannone, D., Qian, E., and Sbordone, A. M. (2021). A large Bayesian VAR of the United States economy. Federal Reserve Bank of New York Staff Report 976.
- Del Negro, M. and Schorfheide, F. (2004). Priors from general equilibrium models for VARs. *International Economic Review*, 45:643–673.
- Del Negro, M. and Schorfheide, F. (2011). Bayesian Macroeconometrics. In *The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.

- Di Giovanni, J., Kalemli-Ozcan, S., Silva, A., and Yildirim, M. A. (2023). Quantifying the inflationary impact of fiscal stimulus under supply constraints. NBER Working Paper 30892.
- Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100.
- Eickmeier, S. and Hofmann, B. (2022). What drives inflation? Disentangling demand and supply factors. Deutsche Bundesbank Discussion Paper 46/2022.
- Fry, R. and Pagan, A. (2011). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4):938–960.
- Furlanetto, F., Lepetit, A., Robstad, Ø., Rubio Ramírez, J., and Ulvedal, P. (2024). Estimating hysteresis effects. American Economic Journal: Macroeconomics, forthcoming.
- Gagliardone, L. and Gertler, M. (2023). Oil prices, monetary policy and inflation surges. NBER Working Paper 31263.
- 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., Lenza, M., and Primiceri, G. E. (2019). Priors for the long run. *Journal of the American Statistical Association*, 114(526):565–580.
- Inoue, A. and Kilian, L. (2013). Inference on impulse response functions in structural VAR models. *Journal of Econometrics*, 177(1):1–13.
- Jarocinski, M. and Marcet, A. (2019). Priors about observables in vector autoregressions. *Journal of Econometrics*, 209(2):238–255.
- Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 186(1007):453–461.
- Kabaca, S. and Tuzcuoglu, K. (2023). Supply drivers of US inflation since the pandemic. Bank of Canada Staff Working Paper 2023-19.
- Karlsson, S. (2013). Forecasting with Bayesian autoregressions. Handbook of Economic Forecasting, (2):791–897.
- Koop, G. and Korobilis, D. (2013). Large time-varying parameter VARs. Journal of Econometrics, 177(2):185–198.
- Miranda-Agrippino, S. and Ricco, G. (2019). Bayesian vector autoregressions: Estimation. In Oxford Research Encyclopedia of Economics and Finance.
- Rubbo, E. (2023). What drives inflation? Lessons from disaggregated price data. Manuscript.

- Shapiro, A. H. (2022). Decomposing supply and demand driven inflation. Federal Reserve Bank of San Francisco Working Paper 2022-18.
- Sims, C. A. (1986). Are forecasting models usable for policy analysis? *Federal Reserve Bank of Minneapolis Quarterly Review*, 10:2–16.
- Sims, C. A. (1993). A nine-variable probabilistic macroeconomic forecasting model. In *Business cycles, indicators, and forecasting*, pages 179–212. University of Chicago press.
- Sims, C. A. (1996). Inference for multivariate time series models with trend. Manuscript.
- Sims, C. A. (2000). Using a likelihood perspective to sharpen econometric discourse: Three examples. *Journal of Econometrics*, 95(2):443–462.
- Sims, C. A. and Zha, T. (1998). Bayesian methods for dynamic multivariate models. *International Economic Review*, 39:949–968.
- Villani, M. (2009). Steady state priors for vector autoregressions. *Journal of Applied Econometrics*, 24(4):630–650.

APPENDIX

A ADDITIONAL RESULTS FOR THE US

Figure A-1: IRFs, different identification schemes



(a) Blanchard-Quah



Figure A-2: Historical decomposition of inflation, different identification schemes

(a) Blanchard-Quah decomposition



Figure A-3: Historical decompositions of US inflation, different priors (a) Normal-Inverse Wishart prior

(b) Minnesota prior





Figure A-4: Longer sample: 1949:Q1 - 2022:Q4



Figure A-5: Deterministic components, diffuse prior



Figure A-6: Posterior distribution of δ , US data

ADDITIONAL RESULTS FOR THE EURO AREA В



Figure B-1: IRFs to identified shocks

Pointwise median and three closest draws, diffuse prior



(a) Euro area

Note: The black line is the pointwise median and the shaded areas the min-max identified set. The red line is the IRF for the draw that is closest to the pointwise median; the green and blue lines are IRFS for 2nd and 3rd draws closest to the pointwise median, respectively.





(a) Historical decomposition of inflation for the 3 draws closest to the pointwise median IRFs



Figure B-3: Mean historical decompositions of inflation, euro area