Beware of large shocks!
A non-parametric structural inflation model*

Elena Bobeica
European Central Bank, Germany

Florian Huber
University of Salzburg, Austria

Sarah Holton[1]
European Central Bank, Germany

Catalina Martínez Hernández
European Central Bank, Germany

Abstract

We propose an empirical structural inflation model that allows for non-linear transmission of shocks. Most existing methods assume a linear contemporaneous effect of shocks on the economy. Using state-of-the-art machine learning techniques, we develop a non-linear Bayesian time series model that combines VARs with regression trees and allows for non-linear effects of shocks for all impulse response horizons. Shock identification is achieved via sign, zero, and magnitude restrictions through a non-linear factor model. Using euro area data and focusing on energy shocks, we find evidence supporting non-linearities with respect to the size of the shock. Inflation reacts stronger to large shocks, while small shocks trigger no significant reactions. Such non-linearity is present along the pricing chain, more visible upstream and gets attenuated moving downstream. For policy makers, the finding that large shocks transmit differently implies that they may require a differentiated response.

JEL: E31, C32, C38, Q43

KEYWORDS: Inflation, non-linear model, machine learning, energy, euro area.

*We are grateful to Christiane Baumeister, Jesús Fernández-Villaverde, as well as the participants of the ECB DG-Economics seminar and the DG ECFIN Workshop Post-COVID labour markets: Understanding recent developments for their comments and suggestions.

[1]The views expressed are those of the authors and do not necessarily reflect those of the ECB. Please address correspondence to:

1 Introduction

Economists tend to agree that economic shocks have a non-linear impact on inflation and the effects might be subject to time-variation. Starting with the seminal idea of Phillips (1958), wage inflation is linked to unemployment in a non-linear way. In a tight labour market, wages are more flexible, as firms bid up wages rapidly, while in a downturn, wages would fall more slowly. Yet, mainstream inflation models tend to be linear. When non-linearities are investigated, modellers frequently impose ex-ante a certain type of non-linearity in a rigid way. This increases the risk of model mis-specification, eventually translating into larger forecast errors of inflation such as the ones observed during the post-pandemic recovery.

The extent to which inflation surprised economists during the post-pandemic recovery made it clear that some elements were missing from traditional models. There was an interest to employ larger models that are able to account for more (and new) sources of inflationary pressures (such as global supply side bottlenecks or more types of energy price shocks). Yet, even when employing a large linear model for inflation (e.g., Bańbura, Bobeica, and Martínez Hernández (2023)), the post-pandemic surge cannot be fully reconciled. We argue that an important element that economists missed in this period is that large shocks can amplify the pass-through to inflation whereas small shocks should not trigger substantial reactions of inflation. Investigating this non-linearity is the main contribution of this paper. But to do so, we also need to develop a new econometric framework that allows us to answer such questions with minimal assumptions on the form of non-linearities.

There has been much work on estimating and analysing non-linear relationships using parametric models. Recently, Benigno and Eggertsson (2023) and Gitti (2024) show that inflation reacts differently to labour market tightness when it exceeds a certain threshold. To capture this effect, they construct a dummy variable interacted with labour market slack. More broadly, empirical macroeconomic studies make various assumptions regarding possible non-linear relationships: parameters change abruptly (e.g., Hamilton and Lin (1996), Krolzig (2013), Hamilton and Perez-Quiros (1996)), slowly varying (Cogley and Sargent (2001), Cogley (2005), Primiceri (2005), Del Negro and Primiceri (2013)), evolve smoothly or after a threshold (Hubrich and Teräsvirta (2013), Galvao and Marcellino (2014), Bruns and Piffer (2023)), a combination of the abrupt and slow variation (Prüser (2021)), or use quantile regressions (Adrian, Boyarchenko, and Giannone (2019), Ghysels, Iania, and Striaukas (2018), Lopez-Salido and Loria (2022)).

In contrast to this literature, machine learning methods allow for remaining agnostic on the precise form of the regression relationship. While machine learning models have proved to be promising for forecasting (Huber and Rossini (2022), Clark, et al. (2023a), Huber, et al. (2023), Lenza, Moutachaker, and Paredes (2023), Hauzenberger, et al. (2024), Goulet Coulombe (2024), among many others), so far their use for structural analysis has been limited. Notable examples are Hauzenberger, et al. (2025), using a recursive identification strategy, and Mumtaz and Piffer (2022), for non-linear local projections. However, all these studies have a serious shortcoming: shocks are allowed to transmit non-linearly to the macroeconomy only with a lag while impact reactions are restricted to be linear.

The model we propose aims to address this shortcoming. We develop a flexible non-parametric framework that builds on the model proposed in Clark, et al. (2023b). This model combines Bayesian Additive Regression Trees (BART, see Chipman, George, and McCulloch (2010)) with vector autoregressions (VARs). As mentioned above, one shortcoming of this specification is that economic shocks are allowed to influence a set of macroeconomic aggregates only linearly on impact, possibly masking important non-linear dynamics that arise immediately after the shock hit the model economy. Our model assumes that the reduced-form errors are driven by a few latent factors. These factors enter the model in a non-linear way through an additive specification that is the sum of a linear function and an unknown, non-linear component. This enables non-linear impact effects of the factors. The factors are interpreted as a set of fundamental economic shocks and, following Korobilis (2022), economic identification is achieved through sign, zero, and magnitude restrictions on the linear component. The resulting empirical framework therefore addresses an important gap in the literature by providing a structural and fully non-linear approach that can also be used to compute competitive forecasts.

In our empirical work, we use a moderately-sized euro area dataset and identify four structural shocks. We first provide a brief overview on how inflation reacts to differently-sized shocks and then zoom into energy price shocks. This is predicated by the fact that they can be an important source of inflation volatility and the years since the outbreak of the COVID-19 pandemic have seen a sequence of large energy shocks unprecedented for the euro area. Also, a large part of what we missed from the post-pandemic inflation surge is attributable to energy shocks, as shown for the case of the euro area by Chahad, et al. (2023).

Our results suggest that inflation reacts non-linearly with respect to the size of the energy shocks. In particular, small shocks trigger no significant reactions of prices whereas larger shocks have disproportionally stronger effects. We, moreover, find that these non-linearities increase in a smooth fashion with shock sizes, providing evidence that models that introduce a few regime shifts are not necessarily consistent with actual inflation dynamics. Taking a more granular view suggests that the non-linear behaviour of prices is present along the pricing chain, more visible upstream when looking at the reaction of energy commodity prices and producer prices and gets attenuated moving downstream towards consumer prices. Also, the non-linearity is more visible for positive shocks, but also present for negative ones.

To highlight that our approach can also be used to produce accurate density forecasts, we employ our model to predict the conditional distribution of year-on-year inflation. This exercise reveals that the flexible non-parametric model yields predictions that are on par with the ones of a linear Bayesian VAR for one-month-ahead but improvements become sizeable for higher-order forecasts, reaching more than 10 percentage points vis-á-vis the linear model.

The remainder of the paper is organised as follows. The next section provides empirical evidence, discusses some theoretical channels for non-linear effects and highlights the related literature while Section 3 introduces our structural BART model with a non-linear factor structure. In Section 4 we give details about our data base, the identification scheme, and the key results of the paper. We further validate our structural model in terms of forecasting performance in Section 5. The last section concludes the paper.

2 Empirical evidence, theoretical channels and related literature

Empirical evidence suggests that inflation reacts more strongly to large shocks, particularly in the post-pandemic period. To provide some empirical motivation, Figure 1 illustrates the relationship between euro area headline inflation and the inflation rate of a synthetic energy commodity price index, which combines oil and gas prices based on the euro area's energy import shares. Given the unprecedented composition of energy shocks in recent years, we employ this synthetic energy commodity price index to better capture energy inflation dynamics.

The figure reveals two simple stylized facts. First, there has been a shift in the inflation-synthetic indicator relationship between the first part of the sample (1996 to 2020) and the second part of the sample (2021 to 2024). Second, and more importantly, this link likely changed and strengthened due to the unprecedented shocks that occurred in the post-pandemic period.

The empirical literature investigating the role of large shocks on consumer price inflation

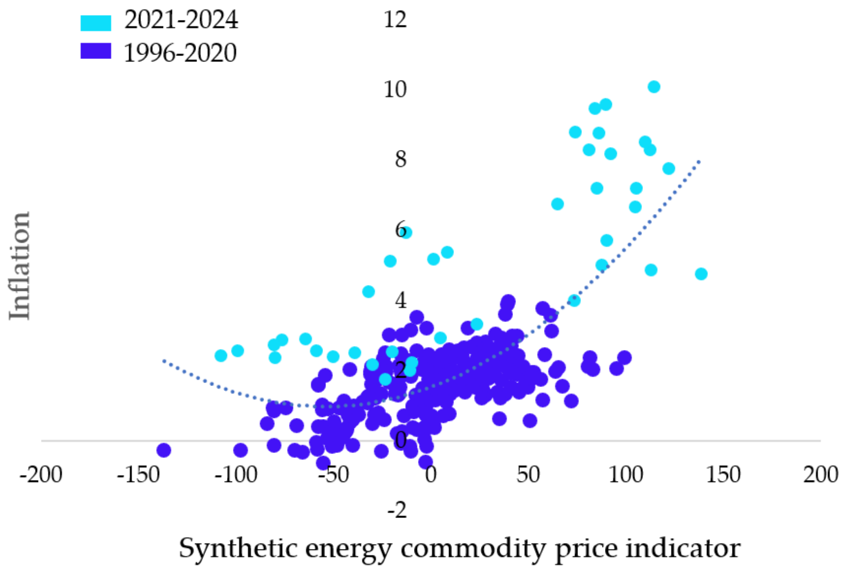

Notes: Expressed in annual growth terms; headline inflation lagged by 7 months.

Figure 1: Inflation and energy commodity prices

is rather limited. Focusing on monetary policy shocks, Ascari and Haber (2022) apply smooth local projections and find that large absolute shocks disproportionately affect prices on impact but are less persistent and have weaker real effects. In general, empirical papers tend to focus more on the role of the inflation regime for the transmission of shocks. The literature documents that inflationary shocks have a stronger impact in high-inflation regimes. De Santis and Tornese (2023) find that energy shocks propagate more intensely when inflation is elevated, while Bobeica, Ciccarelli, and Vansteenkiste (2020) show a similar pattern for labour cost shocks.

Focusing on the post-pandemic inflation surge, several studies highlight an increased sensitivity of inflation to shocks. Adolfsen, et al. (2024) examine gas market shocks and find stronger price reactions when incorporating 2020–2022 data, suggesting a structural break in the relationship between energy markets and euro area prices. Allayioti, et al. (2024) and Zlobins (2025) document a stronger reaction of inflation to monetary policy shocks when the estimation sample covers the ECB's tightening cycle, which was deployed to address the inflation surge following the pandemic. Szafranek, Szafrański, and Leszczyńska-Paczesna (2024) similarly report increased persistence of inflationary shocks since COVID-19.

Theory backs up the idea that large shocks have disproportionate effects. Ball and Mankiw (1995) argue that firms adjust prices in relation to large shocks, but not to small shocks. This is due to the presence of menu costs which make price adjustment more expensive. Insights from micro studies of price setting support the idea that large shocks are associated with a disproportionally larger impact on inflation. According to state-dependent models, the frequency of price changes and the size of these changes are endogenous, they depend on various economic states linked for instance with large shifts in demand or to the size of inflationary shocks. Available micro evidence for the post-pandemic recovery period suggests that indeed, firms practised state-dependent pricing with a large frequency of price adjustment and a rise in the overall size of price changes both in the euro area (see Gautier, et al. (2024), Dedola, et al. (2024)) and in the US (Montag and Vallenas (2023)).

In recent years, there has been significant development in state-dependent models of price setting, bringing them closer in scope to medium-scale models used in policy institutions. Cavallo, Lippi, and Miyahara (2023) adapt a New Keynesian model to the price-setting patterns observed in granular data allowing for a sizeable component of state-dependent pricing decisions. They allow for the frequency of price changes to increase sharply after a large shock and derive a punchy conclusion: 'large shocks travel faster than small shocks'. Karadi, et al. (2024) introduce a micro-founded model with state-dependant pricing, where the reaction of inflation to economic activity increases after large shocks due to an endogenous rise in the frequency of price changes. Empirical evidence supports the finding that prices are more flexible under large shocks (Karadi and Reiff (2019)). Also, state-dependent models imply that prices are more flexible when inflation is higher (Alvarez, et al. (2019)).

Indeed, a rich strand of literature argues that the pass-through of shocks is stronger when inflation is high. Harding, Lindé, and Trabandt (2023) enhance a workhorse macroeconomic model with real rigidities in price and wage setting by incorporating additional strategic complementarities in firms' price-setting behaviour, which lowers the sensitivity of prices to marginal costs for a given degree of price-stickiness. As a result, cost-push shocks can propagate more than four times stronger in the non-linear model, relative to the linearized model when inflation is high. They also argue that among different types of shocks, cost-push shocks propagate stronger to inflation than demand and technology shocks when inflation is high and rising. An important policy implication of their work is that in a high-inflation regime, central banks face a more severe trade-off between inflation and output.

Head, Kumar, and Lapham (2010) suggest a different channel for why the pass-through of shocks increases with the level of inflation, linked to the search intensity of consumers. When inflation is low, consumers can better perceive changes in the price dispersion and any given shock would make them increase their search intensity. When inflation is high, price dispersion is higher and hence any given shock has only a limited impact on price dispersion and the search intensity of consumers, which makes it easier for firms to increase their prices.

Taken together, these arguments suggest a stronger transmission of shocks during the inflation surge episode. We argue that a suitable econometric model should be able to reproduce these asymmetries with minimal assumptions on the form of non-linearities.

3 Econometric framework

The main goal of this paper is to analyse how economic shocks transmit non-linearly to inflation. To do so, we avoid introducing strong parametric restrictions on the functions determining the conditional mean of the process and the contemporaneous effects from structural shocks. We introduce the features of the model, discuss structural identification, and sketch the estimation approach in the following subsections.

3.1 A non-parametric multivariate time series model

We model the non-linear interactions between a panel of M macroeconomic time series, which we store in a vector $\boldsymbol{y}_t = (y_{1t}, \ldots, y_{Mt})'$, and a lower dimensional vector of Q fundamental shocks, which are denoted by $\boldsymbol{q}_t$. Our multivariate model assumes that $\boldsymbol{y}_t$ is an unknown function of $\boldsymbol{x}_t = (\boldsymbol{y}'_{t-1}, \ldots, \boldsymbol{y}'_{t-p})'$, which is a $K(= Mp)-$dimensional vector that includes the p lags of $\boldsymbol{y}_t$ and the fundamental shocks (henceforth simply called factors) in $\boldsymbol{q}_t$. Formally, the model reads:

$$\boldsymbol{y}_t = f(\boldsymbol{x}_t) + g(\boldsymbol{q}_t) + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{q}_t \sim \mathcal{N}(\boldsymbol{0}_M, \boldsymbol{I}_M), \tag{1}$$

where $f : \mathbb{R}^K \to \mathbb{R}^M$ and $g : \mathbb{R}^Q \to \mathbb{R}^M$ are unknown (and possibly non-linear) functions that relate $\boldsymbol{x}_t$ and $\boldsymbol{q}_t$ to $\boldsymbol{y}_t$, respectively. $\boldsymbol{\varepsilon}_t$ is a set of Gaussian measurement errors with zero mean and a diagonal covariance matrix $\boldsymbol{\Omega}$.

This model is extremely flexible and nests several prominent models used in the literature. If we set $f(\boldsymbol{x}_t) = \boldsymbol{A}\boldsymbol{x}_t$ where $\boldsymbol{A}$ is an $M \times K$ matrix of coefficients and let $q(\boldsymbol{q}_t) = \boldsymbol{Q}\boldsymbol{q}_t + \boldsymbol{\varepsilon}_t$ with $\boldsymbol{Q}$ being the lower Cholesky factor of an $M \times M$ covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{Q}\boldsymbol{Q}'$ we end up with a linear VAR with factor structure (as in Korobilis (2022), Chan, Eisenstat, and Yu (2022), and Bańbura, Bobeica, and Martínez Hernández (2023)). If we set $f(\boldsymbol{x}_t) = \boldsymbol{A}\boldsymbol{x}_t + \tilde{f}(\boldsymbol{x}_t)$ we end up with the mixture BART specification used in, e.g., Clark, et al. (2023b) and Baumeister, Huber, and Marcellino (2024). In what follows, we employ this specification so that Eq. (1) becomes:

$$\boldsymbol{y}_t = \boldsymbol{A}\boldsymbol{x}_t + \tilde{f}(\boldsymbol{x}_t) + g(\boldsymbol{q}_t) + \boldsymbol{\varepsilon}_t. \tag{2}$$

The key novel feature of this model is the presence of $g(\boldsymbol{q}_t)$. This implies that the fundamental shocks are allowed to trigger an immediate and possibly non-linear effect on $\boldsymbol{y}_t$. This is in stark contrast to models such as the one proposed in Clark, et al. (2023b) which sets $g(\boldsymbol{q}_t) = \boldsymbol{\Lambda}\boldsymbol{q}_t$, with $\boldsymbol{\Lambda}$ being an $M \times Q$ matrix of factor loadings. In this case, the impact reaction of $\boldsymbol{y}_t$ to a shock to factor j, q_{jt}, is simply given by the j^{th} column in $\boldsymbol{\Lambda}$, $\partial \boldsymbol{y}_t / \partial q_{jt} = \boldsymbol{\lambda}_j$. As opposed to this specification, the model we propose is much more flexible since it allows to accommodate different forms of non-linearities. For instance, a popular, yet simple way, of handling non-linearities would be by setting $g(\boldsymbol{q}_t) = \boldsymbol{\Lambda}_2\boldsymbol{q}_t^2 + \boldsymbol{\Lambda}_1\boldsymbol{q}_t$ and thus assuming that larger shocks trigger a different effect than smaller shocks. Other, deterministic choices of g are possible and commonly used in the literature (see, e.g., Yalcin and Amemiya, 2001; Bai and Ng, 2008; Song and Dunson, 2022; Hauzenberger, Huber, and Klieber, 2023; Forni, et al., 2022). This assumption, however, is very restrictive and we wish to remain agnostic on possible non-linearities and/or interaction effects between the elements in $\boldsymbol{q}_t$. We describe how we achieve this in the next sub-section.

3.2 Non-linear structural factor models

We aim to infer the function g from the data. This is achieved through a non-parametric factor model (see, e.g., Song and Lu (2012), Li and Chen (2016), Xu, Herring, and Dunson (2023)). Subtracting $f(\boldsymbol{x}_t) = \boldsymbol{A}\boldsymbol{x}_t + \tilde{f}(\boldsymbol{x}_t)$ from $\boldsymbol{y}_t$, our model is given by:

$$\tilde{\boldsymbol{y}}_t = g(\boldsymbol{q}_t) + \boldsymbol{\varepsilon}_t,$$

with $\tilde{\boldsymbol{y}}_t = \boldsymbol{y}_t - f(\boldsymbol{x}_t)$. Without further restrictions, this specification is not identified. There are, at least, four identification issues. First, as noted in Yalcin and Amemiya (2001), if we consider a one-to-one non-linear transformation $\tilde{\boldsymbol{q}}_t = h(\boldsymbol{q}_t)$ so that $\boldsymbol{q}_t = h^{-1}(\tilde{\boldsymbol{q}}_t)$. In this case, $g(\boldsymbol{q}_t) = g^*(h(\boldsymbol{q}_t))$. The second and third relate to the sign and scale of $\boldsymbol{q}_t$, and the fourth relates to the fact that, without further information, the different factors have no economic meaning.

If interest is solely on forecast distributions, neither of these translates into major issues. However, given our interest on structural inference and economic interpretation of the factors, we need to introduce further restrictions. We have already introduced the restriction that the factors are standard normally distributed. This identifies the scale of the factors but does not fix the sign nor gives the factors an economic interpretation.

To attach an economic meaning to each of the elements in $\boldsymbol{q}_t$ and to fix the sign, we decompose $g(\boldsymbol{q}_t)$ in a linear and non-linear component:

$$g(\boldsymbol{q}_t) = \boldsymbol{\Lambda}\boldsymbol{q}_t + \tilde{g}(\boldsymbol{q}_t). \tag{3}$$

Hence, $g(\boldsymbol{q}_t)$ consists of two ingredients. First, we include a linear part that implies a linear relationship between $\tilde{\boldsymbol{y}}_t$ and $\boldsymbol{q}_t$. This part serves to capture the main effects of variations in $\boldsymbol{q}_t$ on the responses and we will place sign, zero, and magnitude restrictions on $\boldsymbol{\Lambda}$. The second ingredient is yet another unknown and possible non-linear function $\tilde{g}(\boldsymbol{q}_t)$. This part serves to control for higher-order effects such as interactions between shocks or other forms of non-linearities. A similar decomposition in the context of Bayesian non-parametric instrumental variables estimation is used in McCulloch, et al. (2021). The inclusion of the linear part will form the basis of our structural identification strategy.

Our model also implies that if $\boldsymbol{y}_t$ is mostly driven by the fundamental shocks $\boldsymbol{q}_t$, the measurement error variance matrix $\boldsymbol{\Omega}$ will be close to zero. In this case, learning f and g becomes difficult since the likelihood becomes flat (see Stella and Stock (2013) for a similar issue in the context of a multivariate unobserved components model). As a solution, we introduce the restriction that around five percent of the total variation in $\boldsymbol{y}_t$ is neither explained through $f(\boldsymbol{x}_t)$ nor $g(\boldsymbol{q}_t)$. This is achieved by fixing $\boldsymbol{\Omega} = 0.05 \times \text{diag}(\hat{\sigma}^2_{y,1}, \dots, \hat{\sigma}^2_{y,M})$ with $\hat{\sigma}^2_{y,j}$ denoting the empirical variance of y_{jt}.

3.3 Structural identification using truncated priors

We adopt a Bayesian approach to estimate our model. This calls for setting up suitable priors on all the unknowns of the model. In our case, these are the functions $\tilde{f}$ and $\tilde{g}$, the coefficients in $\boldsymbol{A}$, and the linear factor loadings $\boldsymbol{\Lambda}$. We will discuss the precise priors on $\tilde{f}$ and $\tilde{g}$ in subsection 3.4. In this section, our focus is on the prior on $\boldsymbol{\Lambda}$ and how we use it to achieve structural identification of the model and to determine the signs of the factors.

To attach an economic meaning to the different factors in $\boldsymbol{q}_t$, we follow Korobilis (2022) and specify sign, zero, and magnitude restrictions on the elements of $\boldsymbol{\Lambda}$. This is not done by exploiting the rotation invariance of $\boldsymbol{\Lambda}$ and $\boldsymbol{q}_t$ but through truncated priors with truncation sets defined to capture possible signs of the responses consistent with economic theory (see

Baumeister and Hamilton, 2015, 2018).

Let λ_{ij} denote the $(i, j)^{th}$ element of $\mathbf{\Lambda}$. Moreover, we let $\boldsymbol{S}$ denote an $M \times Q$ matrix that determines the sign restrictions, with $s_{ij} \in \{-, 0, +, c, *\}$. Values of $-, 0, +$ imply negative, no, or positive linear effects of q_{jt} on variable i, while c indicates that we restrict the corresponding loading $\lambda_{ij} = c$. An asterisk points towards an unrestricted effect.

The resulting restrictions are then introduced by truncating the prior on λ_{ij}. The prior takes the form:

$$p(\lambda_{ij}|s_{ij}) = \begin{cases} \mathcal{N}_{0,\infty}(\lambda_{ij}|0, \phi_{ij}) & \text{if } s_{ij} = + \\ \mathcal{N}_{-\infty,0}(\lambda_{ij}|0, \phi_{ij}) & \text{if } s_{ij} = - \\ \delta_0(\lambda_{ij}) & \text{if } s_{ij} = 0 \\ \delta_c(\lambda_{ij}) & \text{if } s_{ij} = c \\ \mathcal{N}(\lambda_{ij}|0, \phi_{ij}) & \text{if } s_{ij} = *, \end{cases}$$

where δ_c denotes the Dirac delta function with pole at c, $\mathcal{N}_{l,u}$ is the truncated Gaussian distribution restricted to the open set (l, u), l and u denote a lower and upper bound, respectively. The prior hyperparameters ϕ_{ij} can be fixed or estimated. In this paper, we set them equal to $\phi_{ij} = 10^2$ for all i, j.

The set of restrictions helps us determine the precise ordering of the factors in $\boldsymbol{q}_t$. For instance, if the restrictions on $\boldsymbol{\lambda}_1$, the first column of $\mathbf{\Lambda}$, are consistent with our theoretical intuition about how an energy shock impacts the economy, we can set up suitable sign restrictions for selected series, allowing us to interpret q_{1t} as the energy shock. Similarly, if the restrictions on $\boldsymbol{\lambda}_2$ are consistent with the signs one would expect from a demand shock, we can label q_{2t} as the demand shock.

It is also worth stressing that the presence of the non-linear term $g(\boldsymbol{q}_t)$ could imply that the contemporaneous reactions to a particular shock violate the sign restrictions. However, our identification strategy hinges on the intuition that the main effects are captured by the linear factor model $\mathbf{\Lambda}$ whereas $g(\boldsymbol{q}_t)$ covers deviations from linearity.

Identification in our framework can be sharpened through the inclusion of additional restrictions, leading to a framework close to the blended identification approach developed in Carriero, Marcellino, and Tornese (2024). For instance, magnitude restrictions (as in, e.g., Corsetti, Dedola, and Leduc (2014)) could be introduced by suitable changing the truncation sets of the prior on λ_{ij}.

Narrative or historical restrictions in the spirit of Antolín-Díaz and Rubio-Ramírez (2018) can also be introduced by changing the state evolution equation on $\boldsymbol{q}_t$ as follows. Since the state equation $\boldsymbol{q}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_Q)$ can be interpreted as a prior on $\boldsymbol{q}_t$, we can modify the prior to reflect information coming from other sources. For instance, suppose we have information that a particular shock, q_{jt} was negative in a particular point in time $t = \tau$. In this case, we can assume that

$$p(q_{j\tau}) = \mathcal{N}_{-\infty,0}(q_{j,\tau}|0, 1),$$

while the prior remains unrestricted in periods $t \neq \tau$. This ensures that in time τ, the resulting shock is non-positive.

Notice, however, that this restriction does not imply a statement about the relative importance of a particular shock in a particular point in time. To do so, we can easily introduce additional restrictions. Suppose that we have information that the j^{th} shock has been the dominant driver of macroeconomic dynamics in time τ. This restriction can be incorporated by specifying:

$$p(q_{j\tau}) = \mathcal{N}(q_{j,\tau}|0, \xi_j^2),$$

where the variance ξ_j can be set so that $\xi_j \gg \xi_i$ for all $i \neq j$. This assumption, through the prior, makes large realizations of $q_{j\tau}$ more likely whereas it shrinks realizations of $q_{i\tau}$ towards zero.

We stress that the sign restrictions on $\boldsymbol{\Lambda}$ fix the ordering issue and allow us to interpret shocks structurally. If we believe that identification is weak (e.g., by having too little structural information on the sign restrictions or the necessity to identify multiple shocks within a family of shocks, such as demand or supply) one can easily incorporate these additional restrictions in our framework.

3.4 Bayesian Additive Regression Trees

Up to this point we remained silent on how we infer $\tilde{f}$ and $\tilde{g}$. We achieve this by estimating them using Bayesian Additive Regression Trees (BART, see Chipman, George, and McCulloch

(2010)). BART approximates unknown functions by summing over R trees:

$$\tilde{f}(\boldsymbol{x}_t) \approx \sum_{r=1}^{R} w(\boldsymbol{x}_t|\mathcal{T}_r^f, \boldsymbol{\mu}_r^f),$$
$$\tilde{g}(\boldsymbol{q}_t) \approx \sum_{r=1}^{R} w(\boldsymbol{q}_t|\mathcal{T}_r^g, \boldsymbol{\mu}_r^g).$$

Here, we let w denote a tree function that is parametrized by tree structures $\mathcal{T}_r^i$ and terminal node parameters $\boldsymbol{\mu}_r^i$ for $i \in \{f, g\}$. The dimension of $\boldsymbol{\mu}_r^i$ is b_r^i. These tree structures are determined by binary decision rules that ask whether x_{jt} (in the case of $\tilde{f}$) or q_{jt} (in the case of $\tilde{g}$) exceed a threshold or not. Formally, this implies that each tree is:

$$w(\boldsymbol{x}_t|\mathcal{T}_r^f, \boldsymbol{\mu}_r^f) = \sum_{s=1}^{b_r^f} \mu_{r,s}^f \mathbb{I}(\boldsymbol{x}_t \in \Omega_s^f), \quad w(\boldsymbol{q}_t|\mathcal{T}_r^g, \boldsymbol{\mu}_r^g) = \sum_{s=1}^{b_r^g} \mu_{r,s}^g \mathbb{I}(\boldsymbol{q}_t \in \Omega_s^g),$$

with $\mathcal{T}_r^i = \{\Omega_s^i\}_{s=1}^{b_r}$ and Ω_s^i, for $i \in \{f, g\}$, denoting a sequence of decision rules that run from the root of the tree down to the final node of a particular branch of the tree. This form of the tree suggests that, conditional on the tree structures, we end up with a regression model with explanatory variables being equal to a set of indicator functions. These explanatory variables are latent and effectively split the input space into disjoint subsets, with each subset being assigned a terminal node parameter $\mu_{r,s}^i$.

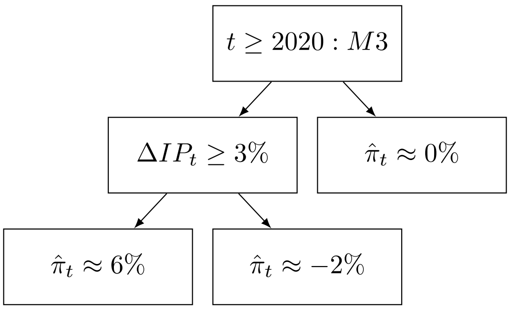

Notes: If a particular condition is fulfilled, move down the left branch of the tree.

Figure 2: A very simple regression tree for year-on-year inflation.

Trees are best understood through a simple example. Suppose that $\boldsymbol{y}_t$ includes only inflation (π_t) and $\boldsymbol{x}_t$ includes the first lag of industrial production (IP) and a time trend. A simple tree is given in Figure 2. The tree starts with a root and associated terminal node that uses time as a splitting variable. The threshold is set equal to March 2020. If time exceeds that date, we move down the left branch of the tree, ending up in a new node with decision rule $\Delta \text{IP}_t \geq 3\%$. If industrial production exceeds three percent, we end up predicting an inflation rate of around six percent. If not, we predict a deflation of around two percent. Moving back up to the root of tree, we find that for periods prior to March 2020, we end up predicting inflation to be close to zero.

This simple example shows that mean inflation can take on only three distinct values $\boldsymbol{\mu}^f = (-2\%, 0\%, 6\%)$ and hence the model might be overly simplistic. This is where BART comes in. Summing over many trees that differ in their decision rules and their depths, implies a joint model that is much more flexible and can approximate the time series behaviour of, e.g., inflation particularly well (for predictive evidence, see, e.g., Clark, et al. (2023b)). In our empirical work, we fix $R = 50$. This choice is motivated by predictive evidence for European data in Huber, et al. (2023) who document that for $R \geq 50$, predictive performance does not improve substantially. Hence, this choice strikes a balance between computational complexity and empirical performance.

3.5 Priors on the trees and remaining model parameters

We elicit standard priors on the parameters of the different trees as well as the VAR coefficients. On the tree parameters, we follow Chipman, George, and McCulloch (2010) and define a tree generating stochastic process. This process is defined recursively and consists of three steps. *First*, we start with a tree that consists simply of a root node. *Second*, with a certain probability we split this terminal node. *Third*, if the node is split a decision rule (defined through a threshold variable and a threshold) is constructed. This is done by randomly picking one of the covariates (in our case different elements in $\boldsymbol{x}_t$ for f and different shocks in $\boldsymbol{q}_t$ for g) and then, conditional on the selected covariate, randomly select a threshold over all the possible outcomes of that particular variable. This decision rule is then used to split the input dataset into a left and right node. Given this new tree, go back to step one and repeat the three steps until every node is terminal.

Step 1 of this process depends on the probability that a given node is split (i.e., it is non-terminal). Chipman, George, and McCulloch (2010) suggest that this probability equals:

$$p_{\text{split}}(\mathcal{T}_r^i, n) = \frac{\alpha}{(1 + \zeta_s)^{\beta}}, \tag{4}$$

where n refers to a particular node of a tree, ζ_s is the depth of the node within tree $\mathcal{T}_r^i$ and $\alpha \in (0, 1)$ is a base probability while $\beta \geq 0$ is a parameter that controls the penalty on growing complex trees under this prior process. Notice that if $\beta = 0$, we end up with a prior that does not penalize tree complexity and the probability that a certain node is non-terminal is equal to α. If $\beta > 1$ we end up with a prior probability of splitting nodes that decreases in the complexity (ζ_s) of the tree. We follow recommendations in Chipman, George, and McCulloch (2010) and set $\alpha = 0.95$ and $\beta = 2$.

On the terminal nodes $\boldsymbol{\mu}_r^i = (\mu_{r,1}^i, \ldots, \mu_{r,b_r^i}^i)$, we use a Gaussian prior with zero mean and standard deviation $\sigma_\mu = 0.5/(R\sqrt{\kappa})$, where κ is a hyperparameter. Chipman, George, and McCulloch (2010) suggest this choice based on rescaling $\boldsymbol{x}_t$ and $\boldsymbol{q}_t$ during estimation so that the values range from -0.5 to 0.5 and recommend fixing $\kappa = 2$. We follow their suggestions. Notice that this specification makes the prior variance a function of the numbers of trees. With a larger number of trees, the prior variance is increasingly forced to zero, implying that large values of $\boldsymbol{\mu}_r^i$ become increasingly unlikely under the prior. This has the immediate implication that each of the trees will explain only a small fraction of the variation in the responses. In machine learning terminology, this implies that each tree is a "weak learner."

On the VAR coefficients, $\boldsymbol{a} = \text{vec}(\boldsymbol{A})$, we use a Minnesota prior (Litterman (1979), Litterman (1980), Kadiyala and Karlsson (1997)). The prior mean is set so as to force the dynamics of the elements in $\boldsymbol{y}_t$ towards a persistent AR(1) process with persistence parameter 0.8 (as suggested by Carriero, Clark, and Marcellino (2015)). Within the equation for the i^{th} element in $\boldsymbol{y}_t$, the prior variance-covariance matrix is specified to induce stronger shrinkage on coefficients associated with lagged dependent variables of the endogenous variables other then variable i. This degree of shrinkage increases quadratically with the lag length. The hyperparameter controlling shrinkage on the own lags of a given variable is set equal to 0.05 while the one controlling shrinkage on other lagged variables is set equal to 0.1×0.05. These choices, while being tight, ensure that the VAR soaks up persistent dynamics in $\boldsymbol{y}_t$ whereas stationary dynamics are, a priori, explained through $f(\boldsymbol{x}_t)$ and $g(\boldsymbol{q}_t)$.

3.6 Posterior simulation

We simulate from the posterior of the parameters and latent states using a Markov chain Monte Carlo (MCMC) algorithm. Here, we briefly sketch the main steps of our sampler. Since all of the steps are standard we refer to the relevant literature.

Our sampler cycles between the following steps:

1. Sample the rows of $\boldsymbol{A}$ independently conditional on the factors from a multivariate Gaussian posterior. The moments take a standard form and further details can be found in,

e.g., Koop and Korobilis (2010).

2. Sample the latent factors $\{\boldsymbol{q}_t\}_{t=1}^T$ on a t-by-t basis using a random walk Metropolis Hastings update. The covariance matrix of the random walk proposal is set equal to the covariance of the conditional posterior of the factors under linearity. The resulting sampler achieves acceptance probabilities that are between 20 and 40% for all t.

3. Since our model is a standard BART, conditional on the factors, the tree structures and terminal nodes are sampled through the algorithm proposed in Chipman, George, and McCulloch (2010).

4. The factor loadings are simulated, row-by-row, from a truncated Gaussian posterior. The exact form of the posterior can be found in Korobilis (2022).

We repeat this algorithm 30,000 times and discard the first 15,000 draws as burn-in. To speed up computation, we use thinning to end up with 1,000 final draws on which we base our inference. The main computational bottleneck is the computation of the generalized impulse responses.

3.7 Computation of the dynamic responses to economic shocks

Our model is highly non-linear and hence standard impulse response analysis cannot be applied directly. To analyse the dynamic reactions of $\boldsymbol{y}_t$ to changes in $\boldsymbol{q}_t$, we rely on generalized impulse responses (GIRFs, see Koop, Pesaran, and Potter (1996)).

In our framework, we compute the GIRFs as follows. Let $p(\boldsymbol{y}_{t+h}|\mathcal{I}_t, q_{jt} = \delta)$ denote the predictive distribution that sets the j^{th} element of $\boldsymbol{q}_t$ equal to a constant δ. $\mathcal{I}_t$ is the available information set in time t. This predictive distribution is obtained recursively for each $h = 0, \dots, H$. The unconditional distribution is given by $p(\boldsymbol{y}_{t+h}|\mathcal{I}_t)$.

Both predictive distributions are available via simulation. The $h-$step-ahead responses are then obtained by subtracting a simulated path of the data $\{\overline{\boldsymbol{y}}_{t+h}^{(q_{jt}=\delta)} \sim p(\boldsymbol{y}_{t+h}|\mathcal{I}_t, q_{jt} = \delta)\}_{t=0}^H$ from the conditional forecast distribution as well as a simulated path $\{\overline{\boldsymbol{y}}_{t+h} \sim p(\boldsymbol{y}_{t+h}|\mathcal{I}_t)\}_{h=0}^H$. One draw from the posterior of the responses is then computed as:

$$\boldsymbol{\xi}(t+h, j) = \overline{\boldsymbol{y}}_{t+h}^{(q_{jt}=\delta)} - \overline{\boldsymbol{y}}_{t+h} \quad \text{for } h = 0, \dots, H.$$

The fact that we condition on the information set up to time t and that our model is non-linear implies that $\boldsymbol{\xi}(t+h, j)$ depends on the current state of the economy (i.e., $\boldsymbol{y}_t$ and $\boldsymbol{x}_t$).

Since we are interested in the average dynamic effects of economic shocks we need to integrate over all possible states in the model. This is achieved by computing:

$$\text{GIRF}(h, j) = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{\xi}(t + h, j),$$

leading to the GIRFs on which we focus in the following sections. These quantities are computed for each draw from the MCMC output, implying that we end up with a posterior distribution of GIRFs which we can use to carry out inference.

4 Non-linear effects of economic shocks

This section includes the main findings of the paper. Before discussing them, we give a brief description of the dataset, provide details on the specification of the model and discuss structural identification. Then, we briefly summarize how inflation reacts to all the shocks in the model before focusing on the non-linear transmission of energy shocks on inflation.

4.1 Data, model specification and identification

We estimate our model using a monthly data set, shown in Table 1. Our data ranges from 1996:01 to 2024:11. As our key inflation measure we rely on headline euro area inflation based on the total Harmonised Index of Consumer Price (HICP) index.

Table 1: Data description

Variable	Description	Source	Trans.
HICP headline	Total HICP	Eurostat	y-o-y
Synthetic	Synthetic indicator summarising oil and gas commodity prices covering 60% and 40%, respectively	ECB	y-o-y
CCI	Consumers' confidence indicator, balance of responses, euro area	European Commission	no trans
GSCPI	Global Supply Chain Pressure Index	NY Fed	no trans.
Selling price expectations	3 months ahead expectations in manufacturing	European commission	no trans.
PMI supplier delivery	Purchasing Managers' Index, manuf., supplier delivery times	Markit	no trans.
PPI total	Total Producer Price Index, domestic sales	Eurostat	y-o-y
PPI energy	Producer Price Index, domestic sales, MIG energy NACE Rev2	Eurostat	y-o-y
EUR/USD	Exchange rate EUR/US dollar	ECB	y-o-y
EURIBOR 1Y	EURIBOR rate at 1 year maturity	ECB	no trans.

We include 10 endogenous variables that apart from HICP inflation cover various aspects of inflation drivers. Since the effects from gas prices have become more prominent in the analysis of euro area inflation after the post-pandemic inflation surge, we opt for a synthetic indicator of energy prices which covers the developments of both oil and gas prices. We include two timely survey-based indicators provided by the European Commission, namely the consumers'

Table 2: Structural shock identification

Variable/Shock	Energy	Global supply chains	Domestic supply	Demand
HICP headline	+	+	+	+
Synthetic	σ_y	**0**	**0**	
Price expectations manuf.	+	+	σ_y	+
CCI	-	-	-	σ_y
PPI	+	+	+	+
PPI energy	+			
PMI supplier delivery		-		
GSCPI		σ_y	**0**	
EUR/USD				
EURIBOR 1Y				

Notes: Restrictions are assumed at the contemporaneous horizon, where $+/-$ denote a positive/negative reaction, 0 means no response, σ_y reflects a normalisation restriction, and an empty entry means no restriction.

confidence indicator (CCI) and the short-term selling price expectations in manufacturing. To account for the effects of global supply chain disruptions, we consider the Global Supply Chain Pressures Index from Benigno, et al. (2022), which comprises global dynamics, as well as the PMI supplier delivery times, targeting euro area dynamics. We further include additional drivers such as the total producer price indicator and its energy component, as well as the EUR/USD exchange rate, and the EURIBOR at the one year maturity. Most variables are transformed to a year-on-year (y-o-y) growth rate and prior to estimation we standardise the data. The model includes $p = 2$ lags of the endogenous variables.

We identify four structural shocks related to energy, global supply chains, domestic supply, and demand. Our identification approach, displayed in Table 2, is based on a set of sign and zero restrictions assumed at the contemporaneous horizon.

The three supply-side shocks – energy, global supply chains, and domestic supply – affect prices and consumer confidence in opposite directions. In our database this materialises in an increase of consumers' headline inflation and total producer prices inflation; and a decrease in consumer confidence.

Our generic energy shock bundles different aspects in the oil and natural gas markets such as supply, market-specific demand, and inventories. This shock is characterised by an increase in the synthetic indicator comprising the dynamics of oil prices and natural gas prices. Furthermore, an energy shock also rises producer prices linked to energy. For the remaining supply shocks — global supply chains and domestic supply — we assume a recursive-like identification such that the energy shock is the only supply-side shock that affects the synthetic indicator contemporaneously.

Following Bańbura, Bobeica, and Martínez Hernández (2023), we identify a global supply chain shock by assuming that it increases the Global Supply Chain Pressures Indicator (GSCPI) and decreases the PMI survey related to supplier delivery times, to sharpen the effects for the euro area.[2] Especially when considering the period that followed after the pandemic it is important to control for such supply disruptions, as they played a notable role in the inflation surge.

Moreover, we distinguish between global supply chain shocks and domestic supply shocks such that the latter does not affect the GSCPI contemporaneously. We identify domestic supply shocks as being ones that increase prices, as well as near-term price selling expectations of firms in the manufacturing sector.[3]

We identify demand shocks by assuming an increase in prices and in the consumers' confidence confidence indicator (based on the survey-based timely and monthly indicator provided by the European Commission). Such a generic demand shock in the economy can capture for instance the effects of monetary policy, fiscal spending, and global demand shocks.

To sharpen identification, we assume additional normalisation restrictions, i.e., we set $s_{ij} = c$, and $c = \sigma_y$ where c is the empirical standard deviation of the variable that is being normalised. Restrictions of this sort imply that a unit/one standard deviation shock to energy, global supply chains, domestic supply, and domestic demand contemporaneously triggers an immediate, linear reaction in the synthetic indicator, the GSCPI, the price expectations in manufacturing, and the CCI by one standard deviation, respectively. The presence of the non-linear term in Eq. (3) implies that deviations from these restrictions are possible. But these can than be exclusively attributed to non-linear dynamics in impact reactions to economic shocks and provide us with information whether the linear restrictions are at odds with the information in the data.

While the selection of the synthetic indicator and the GSCPI as normalisation variables

[2] A reading in the euro area PMI supplier delivery times below 50 indicates that delivery times have deteriorated (e.g., delays), while a reading above 50 is interpreted as an improvement in the delivery conditions. For more details see the S&P Global website.

[3] We investigated a wide range of alternative identification strategies for the four identified shocks and our results stay robust with respect to minor alterations in the restrictions used on $\mathbf{\Lambda}$. In particular, when it comes to the domestic supply shock, we also investigated robustness with respect to normalising the reactions of other economic indicators such as industrial production, the survey-based indicator of confidence in manufacturing or PMI output.

is clear, given the direct link between the variables and the shock, we further explain our rationale for the normalisation of domestic supply and domestic demand shocks. We normalise domestic supply shocks to the price expectations in manufacturing for two reasons. First, the survey is conducted to firms in the manufacturing sector and therefore it is likely to capture developments in supply to a large degree. Second, the expectations are of short-term nature and it covers timelier economic developments. In a similar fashion, we select the CCI to normalise domestic demand shocks since it is based on a timely survey conducted to consumers, whose answers likely reflect a demand component rather than a supply one.

4.2 The reaction of inflation to economic shocks

In this section, we discuss how inflation reacts to all identified economic shocks. To provide a rough overview of the effects and their asymmetries, we will focus on the reactions of inflation $h = \{0, 6, 12\}$-steps-ahead.

Figure 3 depicts the horizon-specific reactions of inflation to all identified shocks in the form of barplots. The figure suggests that non-linearities are present for all shocks and horizons considered. The finding that small shocks trigger no discernible reaction on impact carries over to demand and supply shocks. Only for global supply chains we find that neither shock size triggers an immediate reaction of inflation. This pattern carries over to higher-order responses. Only in the case of global supply chain shocks, our findings suggest significant reactions after around one year, but only if the shock is large.

When comparing asymmetries across shocks, we find somewhat less evidence of asymmetries than we do for energy shocks. Only for supply shocks we find similar a similar degree of asymmetries when moving from the small to the moderately-sized shock. Notice, however, that in this case the asymmetries when moving from a medium to a large shock is quite muted. In sum, we can say that our results point to visible non-linearities with respect to the size of the shock, putting into perspective the way one interprets results from linear models which are traditionally employed.

Why do we find more asymmetries for energy shocks as compared to the other shocks considered? One answer for this can be rooted in the statistical nature of energy shocks. Particularly during the last period of our sample, we observe a larger incidence of large energy shocks. Figure A4 in the Appendix displays the estimated energy shocks, which tend to follow the broad patterns in the developments in the synthetic energy indicator. The last period of the sample

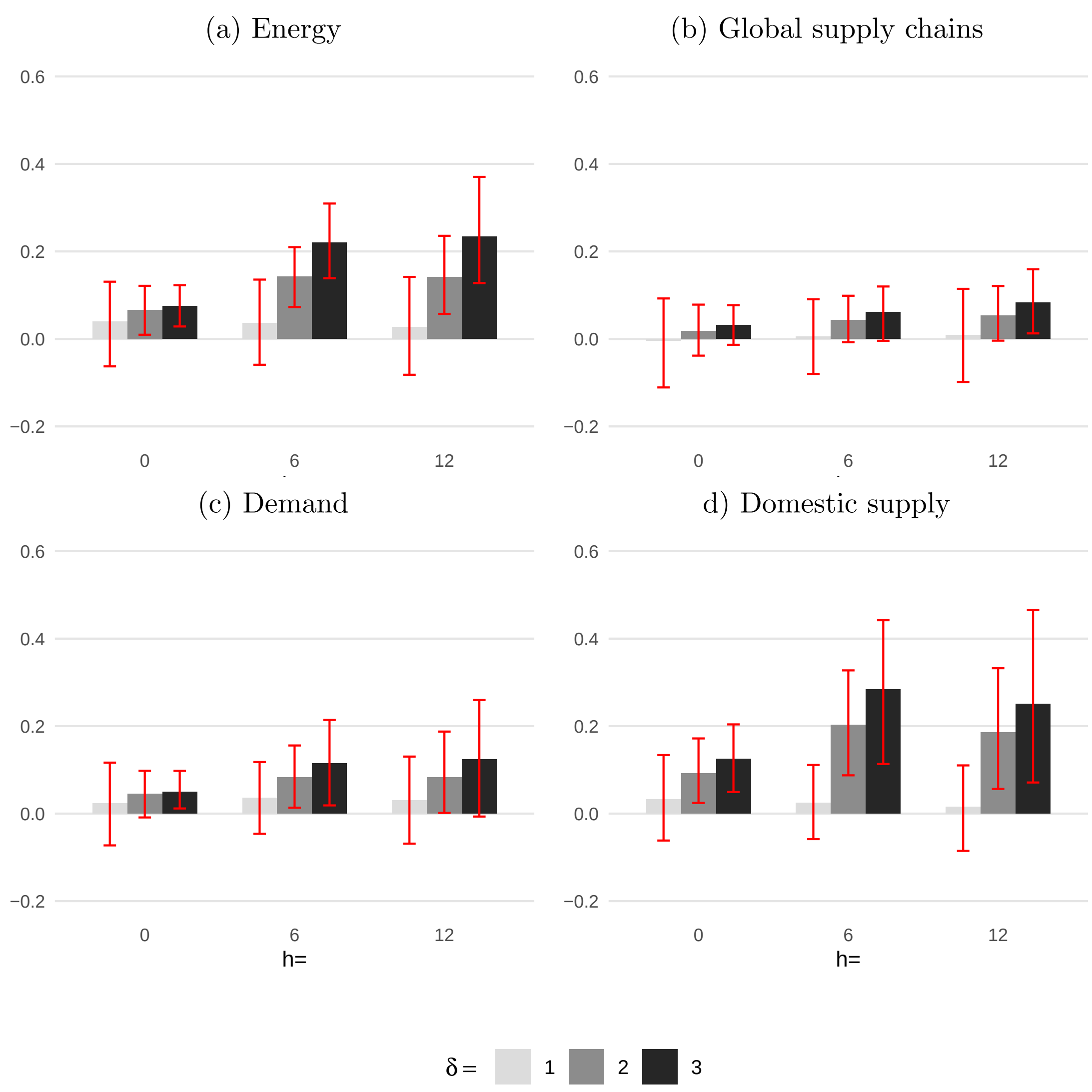

Figure 3: Reaction of inflation to the identified structural shocks

Notes: Bars show the 68% credible interval. Shock sizes and impulse response horizons are on the X-axis. IRFs are normalised for comparability across the three shock sizes.

marks a different behaviour of the estimated shock, with the series portraying both larger shocks and more volatility.

Another reason why non-linearities are more visible for energy shocks is that these are pure upstream shocks. Unlike other shocks, the buffering efforts of firms can be mitigated because the impact of upstream shocks is amplified by the integrated network structure of the economy until it reaches more downstream sectors. Figure 3 also highlights that large shocks are estimated with larger uncertainty in the case of demand and domestic supply shocks.

Given their importance and the fact that our analysis reveals strong asymmetries with respect to the size of the shock, we now focus on how inflation reacts to energy shocks in more detail.

4.3 The non-linear transmission of energy shocks on inflation

Figure 4 illustrates the responses of euro area headline inflation to small (1 standard deviation, in gray), medium (2 standard deviations, in light red), and large (3 standard deviations, in dark red) energy shocks up to the 2-years ahead horizon. We label these shocks as being small, medium, and large based on the reaction they elicit in the synthetic energy commodity price indicator. Specifically, a one standard deviation shock triggers a small increase in the synthetic indicator, occurring in less than 10% of all increases in our sample. A two standard deviation shock elicits an average-sized increase in the synthetic indicator, with half of the increases in our sample being smaller and half larger. Finally, a three standard deviation shock causes a large increase, occurring in around 10% of all recorded increases (see Figure A2).

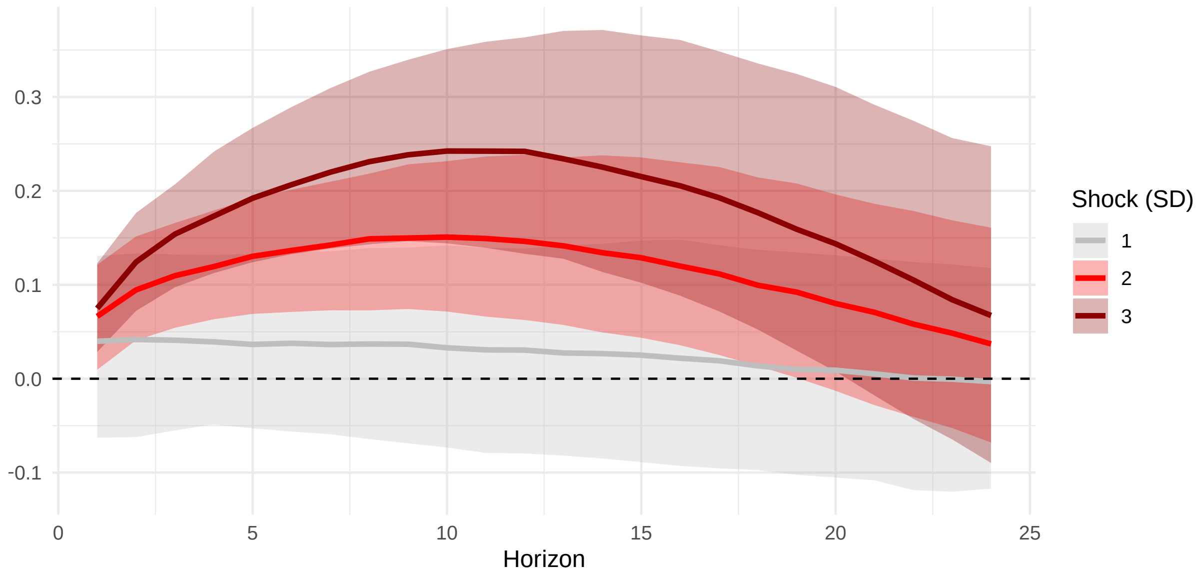

Figure 4: Reaction of inflation to small, medium, and large energy shocks

Notes: thick lines are median estimates and the shaded areas are the 68% credible interval. IRFs are normalised to a one standard deviation shock for comparability.

Figure 4 highlights the fact that larger shocks have disproportionate effects. Impulse responses are being rescaled for comparability to the reaction to a one standard deviation shock, implying that in a linear VAR all responses would evidently overlap. Small shocks do not yield significant effects on inflation, also in line with the theory laid down by Ball and Mankiw (1995) where firms have a range of inaction to small shocks as price adjustments are costly. Instead, medium-sized or larger shocks elicit significant inflation reactions, as firms pass them on to consumer prices. It is worth stressing that without adding $\tilde{g}(\boldsymbol{q}_t)$ to the model, even small shocks would trigger significant impact reactions of inflation. But these are merely artefacts that arise from truncating the prior (and hence the posterior).

In terms of magnitude (and after rescaling), a medium-sized shock would imply a 2.5 times stronger reaction of inflation to energy shocks compared to a small shock. While the most noticeable effect occurs when moving from a small to medium-sized shock, the further increase in the size of the shock adds to the non-linear effect. In this case, a large shock triggers a response which is about 1.5 times larger than the inflation response to a medium-sized one. Figure A3 in the Appendix shows that there is evidence of non-linearity also when it comes to negative shocks, albeit somewhat less pronounced.

Sources of such non-linear behaviour are provided by the theoretical literature mentioned in the introduction and broadly linked to price adjustment costs, increased frequency of price adjustments with larger shocks and decreased attention of consumers when hit by larger shocks, which allows firms to pass on costs.

Another explanation for the non-linear behaviour resides in the amplification impact via second-round effects. Variables like wages or expectations adjust to shocks with a delay and feed-through to prices amplifying the initial effects of shocks. This adjustment likely varies over time. Ampudia, Lombardi, and Renault (2024) find that wages have adjusted more prominently to inflationary shocks in the high inflation period following the pandemic. Acharya, et al. (2023) show that increased consumers' inflation expectations may have interacted with an increase in firms' pricing power to make the effects of the original supply shocks more persistent (Acharya, et al. (2023)). While, in normal times, firms are hesitant to raise prices not to lose customers, once consumers expect higher inflation, firms can more easily pass on increases in costs.

One aspect that can add to the non-linearity observed in Figure 4 is more mechanical in nature. A large part of the consumer energy price level is excise taxes for both fuel and gas applied to the price level. Together with broadly stable refining and distribution margins, this implies that a certain percentage change in the euro price of oil triggers a lower percentage change in consumer energy prices when oil prices are low compared with when they are at high levels.[4]

The discussion so far highlighted asymmetries in the reactions to shocks to three different sizes. To better understand whether there exist threshold effects of the form that inflation

[4] Another mechanical source of non-linearity could be related to the fact that higher energy prices (with little substitution possibility for energy quantities) implies higher nominal expenditures and a higher weight in the consumption basket, augmenting the impact of energy shocks on total inflation.

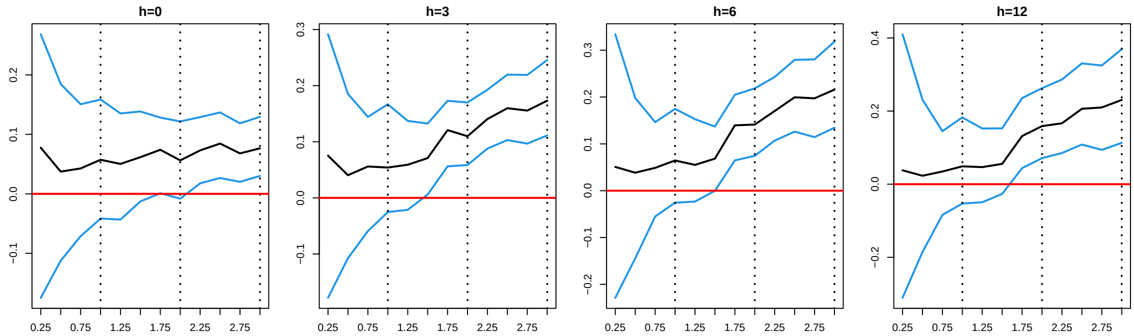

Figure 5: Reaction of inflation to energy shocks over a sizes grid

Notes: Black lines are median estimates and the blue lines denote the 16/84 percentiles of the posterior distribution. Responses are normalised to a one standard deviation shock for comparability.

reactions change markedly beyond certain shock sizes, we consider horizon-specific IRFs over a grid of shock sizes. This grid ranges from 0.25 to 3 units in 0.25 increments. The corresponding IRFs for $h \in \{0, 3, 6, 12\}$ months are provided in Figure 5. All four panels are normalised so that differences in the posterior percentiles of the response distribution suggest non-linear reactions.

The figure highlights two facts about non-linearities in inflation reactions. First, when we consider the median reaction we either have little evidence for non-linearities (this holds for the impact reaction) or a rather gradual and smooth increase in the reaction of inflation, effectively questioning the assumption of regime-switching models frequently used in the literature (see, e.g., Sims and Zha, 2006). Second, when we take into account posterior uncertainty we find that shocks below a certain size trigger no significant reactions. This result can be attributed to the non-linear factor specification. Once shocks become sizeable enough (which means that they exceed a shock size of about 1.5 to 2), posterior uncertainty declines and IRFs turn significant.

4.4 Transmission of energy shocks along the pricing chain

Next, we analyse how energy shocks transmit along the pricing chain, starting with the reaction of the synthetic energy commodity price index, then showing how producer prices for energy react and how these reactions then compare to the overall inflation reactions. The corresponding IRFs are shown in Figure 6. For comparability, the reaction of inflation is shown as well.

At a general level (and for all shock sizes except the small ones) we find that the reactions to shocks decline when moving down the pricing chain. This is expected and consistent with the declining role of energy as we move downstream. It is also worth noting that the timing of the peak reaction differs across the various prices. The synthetic energy commodity price indicator reacts instantaneously to energy shocks, while reactions of producer prices for energy peak after

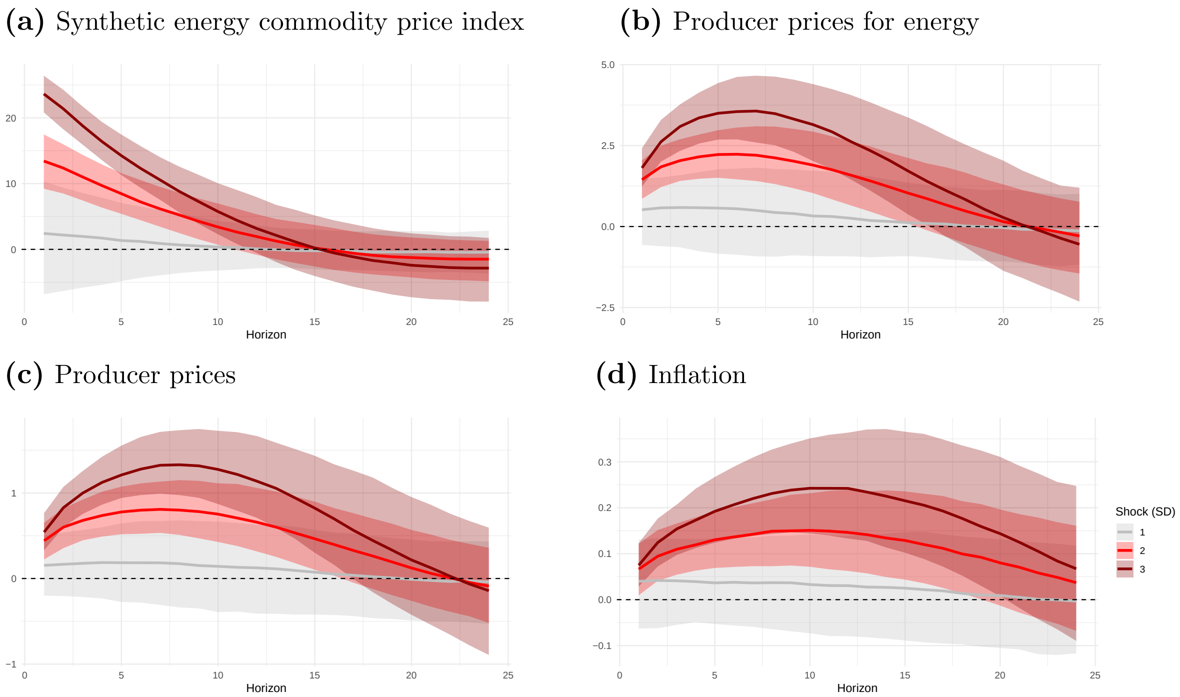

Figure 6: Transmission of energy shocks of various sizes along the pricing chain

Notes: Thick lines are median estimates and the shaded areas are the 68% credible interval. IRFs are normalised for comparability.

about 6 months and finally, the impact on consumer prices is more persistent and peaks after about one year.

Turning to differentiated impacts, Figure 6 indicates that the non-linear behaviour of prices is present along the pricing chain, more visible upstream and gets attenuated moving downstream. As the role of energy among production inputs declines, firms are able to buffer shocks to a larger degree to smooth variation in prices. More precisely, the synthetic energy commodity price index (see panel (a)) reacts in a highly non-linear way to differently-sized energy shocks. These differences mostly stem from different impact reactions. Again, we find that small shocks trigger no significant reaction and larger shocks translate into slightly more persistent reactions for the first 5 to 7 months.

Shock size asymmetries decline when moving down the pricing chain. Substantial differences for medium-term reactions of producer prices for energy (see panel (b)) are visible. Again, these are most pronounced when moving from small to medium-sized shocks. For producer prices, we still find appreciable asymmetries but these are somewhat less pronounced than in the case of producer prices for energy and finally, differences are smaller for consumer prices.

4.5 A closer look at the pre-pandemic period

In this section, we focus on the impact that pandemic and post-pandemic periods have on the IRFs shown in Figure 4. To this end, we estimate the model over two sub-samples. The first uses data through 2019 and hence does not include any information on the pandemic and post-pandemic period. The second is based on just using data from 2020:01 onwards. This exercise tells us whether transmission mechanisms have changed during and after the pandemic.

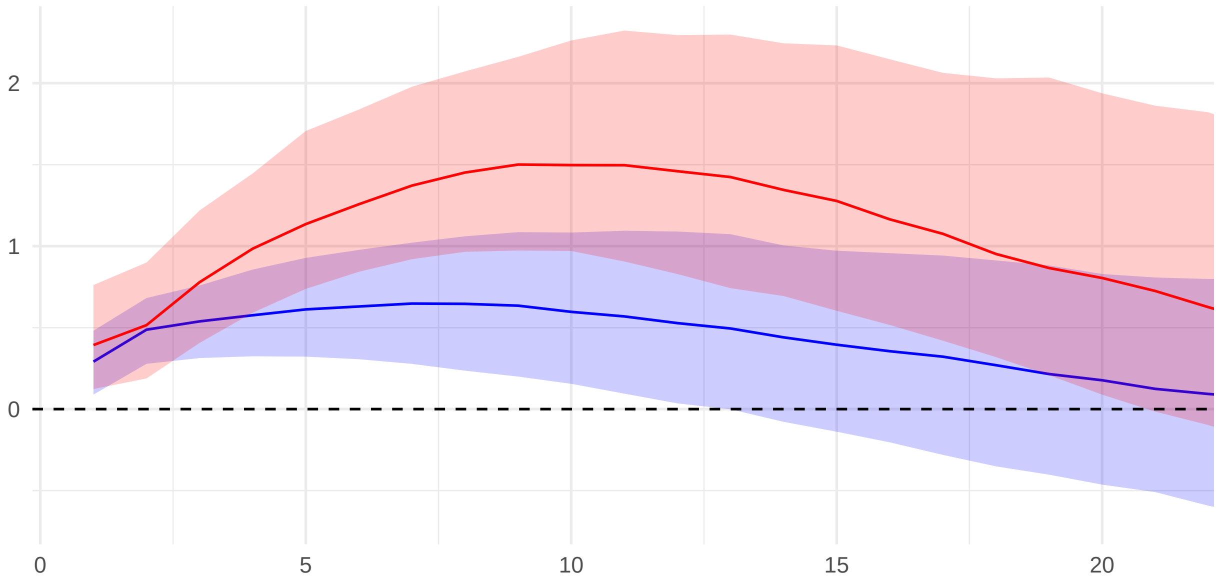

Figure 7: Reaction of inflation to large energy shocks over two sub-samples

Notes: Red areas denote the 68% credible intervals for the period from 2020:01 to 2024:11 while blue areas are the 68% credible intervals for the period from 1996:01 to 2019:12.

The resulting IRFs are depicted in Figure 7. Comparing Figure 4 and Figure 7 reveals substantial differences in the reaction of inflation before and when considering only the pandemic and post-pandemic period. Three differences across IRFs are worth highlighting. First, we can say that energy shocks have a much stronger effect on inflation when we focus on the final part of the sample for training the models.

Second, the shape of the IRFs differs appreciably when comparing both samples. The degree of persistence is significantly greater when considering the full sample. The peak response of inflation to energy shocks in the pandemic and post-pandemic sample occurs 10 months after the shock, compared to around 7 months in the pre-COVID sample.

Third, the pre-pandemic energy shock fully passes through to inflation within one year, while when adding the post-pandemic observations the energy shock only dissipates after two years. In this particular episode, this delay was amplified by the atypical role of gas, for which certain studies find a more protracted impact compared to oil (Bańbura, Bobeica, and Martínez Hernández (2023)). This holds also when it comes to direct effects. Model-based evidence for the euro area suggests that the pass-through of crude and refined oil prices to consumer prices is generally complete and quick, within 3-5 weeks. The pass-through of wholesale gas prices is slower, with about three to six months on average (Kuik, et al. (2022)).

5 Predictive evidence for our model

The IRFs provide evidence of substantial non-linearities in the pass-through of economic shocks to inflation. These responses are consistent with intuition and economic theory. However, to further test the validity of our model, we assess whether it can also generate accurate inflation predictions. This is what we do in this section. Given that linear VARs with stochastic volatility are highly competitive models when it comes to forecasting inflation (see, e.g., Banbura, Lenza, and Paredes (2024)), we benchmark the predictive properties for inflation of our proposed structural factor BART (sfBART) model to a linear BVAR with Horseshoe priors and factor stochastic volatility. This also allows us to understand the role of the non-linear features of our model on the predictive distribution of inflation.

Our forecasting experiment is recursive. We estimate the model using data through December 2008 to train the models. After doing so, we iteratively compute the predictive distribution of year-on-year inflation for $h \in \{1, 3, 6, 12\}$. Once these are obtained, we add the next observation (January 2009) to the training sample and re-do the analysis. This is done until the end of the sample is reached.

To compare the models in terms of density forecasting, we consider the quantile weighted continuous rank probability scores (qw-CRPS), where a smaller number denotes a higher forecasting accuracy. One of the merits of considering qw-CRPSs is that these measures puts more weight on certain regions of the distribution. We consider three different metrics:

1. Right-tailed CRPS: Puts more weight on the right tail of the forecast distribution of inflation (deflation risk)

2. Centered CRPS: Puts more weight on the center of the distribution

3. Left-tailed CRPS: Puts more weight on the left tail of the forecast distribution of inflation (high inflation risks)

Table 3: CRPS of BART models relative to CPRS of BVAR model

CRPS	Model	Horizon			
		1	3	6	12
Left	sfBART	0.96	0.91	0.88	0.89
Center	sfBART	0.98	0.91	0.87	0.92
Right	sfBART	0.98	0.89	0.84	0.91

We report the qw-CRPSs of sfBART relative to ones of the linear BVAR in Table 3, averaged over the hold-out period that ranges from January 2009 and November 2024. Numbers below one indicate that the our sfBART model outperforms the linear model and are highlighted by the green cells, with darker shade representing a better forecasting accuracy of the model for the respective forecasting horizon.

The table shows that our model produces forecasts that are competitive to the benchmark linVAR model at the one-month-ahead horizon. This holds for all three regions of the predictive distribution considered. When we increase the forecast horizon we find that the predictive accuracy gradually increases. For $h = 3$ we find improvements that range from around 9 percentage points (pps, Left and Center) to about 11 pps (right tail).

These improvements further increase if we turn to the 6-months-ahead predictions. In this case, we observe improvements that exceed 12 pps for all three metrics, with a maximum improvement of staggering 16 pps when it comes to predicting high inflation outcomes. Gains slightly deteriorate for 12-steps-ahead. Improvements in higher-order forecast accuracy stem from the fact that the non-linear model is less prone to mis-specification (see Clark, et al. (2023b)).

The accuracy gains for $h = 3$ and $h = 6$ in the right tail can be attributed to the high inflation period of 2021 to 2023. During this period, using a non-linear specification pays off.

As a by-product of our analysis, we compare the forecast errors of this model with those of the official ECB inflation projections in Figure 8. Our model yields similar errors in more quiet times, but improves precisely during the high inflation period. A key distinction between the ECB staff projections and our analysis is that the ECB projections are conducted as a pure real-time exercise. Although this is an important difference, we would not expect large differences in our results if we were to construct vintages of the data, since most of the variables in the model do not get revised or only have minor revisions.

Our forecasting results therefore validate the model as an adequate model that, besides

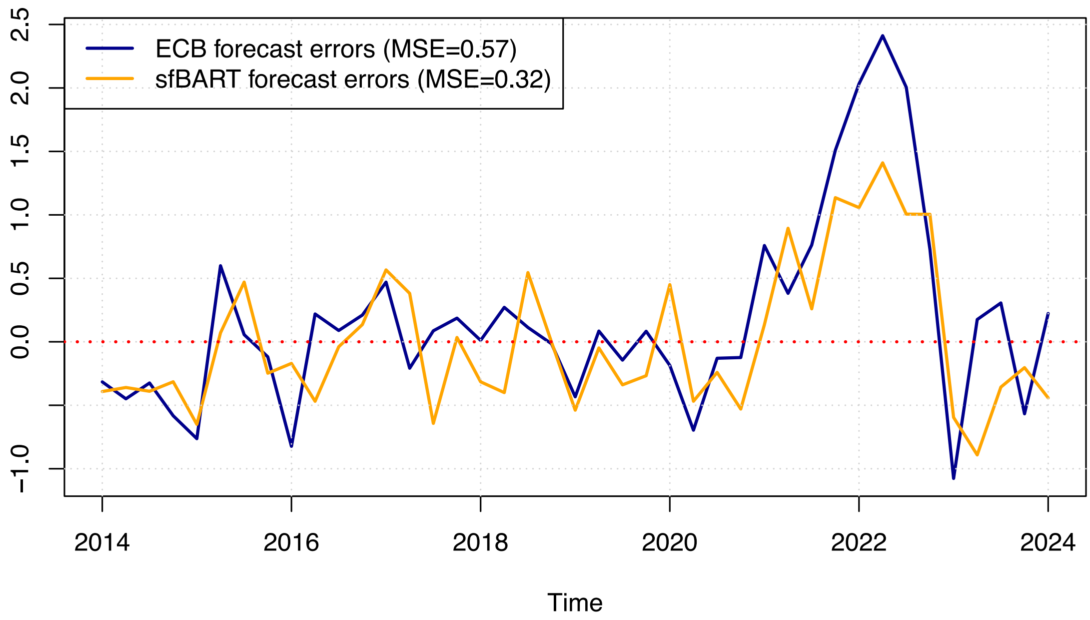

Figure 8: Forecast errors: sfBART versus ECB projections

providing a structural interpretation of inflation, it also produces accurate forecasts.

6 Conclusions

This paper addresses the non-linear transmission of structural shocks to euro area inflation. To fully answer this question, we developed a novel econometric model that accommodates non-linearities for all impulse response horizons while achieving structural identification within a unified econometric framework. Our model, leveraging state-of-the-art machine learning techniques, remains agnostic on the functional form of the relationship between lagged endogenous variables and economic shocks.

The post-pandemic inflation debate has put into the spotlight the different impact of large shocks, considered to have a disproportionately larger impact on inflation (see Cavallo, Lippi, and Miyahara (2023)). In this paper we empirically show that non-linearities are significant when inflation is hit by large shocks.

Large shocks have disproportionate effects. We find that non-linearities increase in a smooth fashion with shock sizes, suggesting that models that introduce a few regime shifts portray the behaviour of inflation in a simplified way. We also show that the non-linear behaviour of prices is present along the pricing chain, more visible upstream when looking at the reaction of energy commodity prices and producer prices and gets attenuated moving downstream towards consumer prices.

Workhorse macroeconomic models used in policy institutions are typically linear, thereby potentially missing non-linearities when shocks are large and inflation is elevated. Nevertheless, linear models are a good approximation when inflation is stable and aggregate shocks are small.

A relevant policy implication of this paper is that optimal monetary policy should place greater emphasis on stabilising inflation in response to exceptionally large cost-push shocks. This is because the endogenous increase in the frequency of price changes leads to a stronger reaction of inflation. While central banks can afford to "look through" small supply shocks when designing monetary policy, they cannot ignore larger ones.

References

Acharya, Viral V., Matteo Crosignani, Tim Eisert, and Christian Eufinger (2023), "How Do Supply Shocks to Inflation Generalize? Evidence from the Pandemic Era in Europe," NBER Working Papers 31790, National Bureau of Economic Research, Inc.

Adolfsen, Jakob Feveile, Massimo Ferrari Minesso, Jente Esther Mork, and Ine Van Robays (2024), "Gas price shocks and euro area inflation," Working Paper Series 2905, European Central Bank.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019), "Vulnerable growth," *American Economic Review*, 109, 1263–1289.

Allayioti, Anastasia, Lucyna Górnicka, Sarah Holton, and Catalina Martínez Hernández (2024), "Monetary policy pass-through to consumer prices: evidence from granular price data."

Alvarez, Fernando, Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andrés Neumeyer (2019), "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models," *The Quarterly Journal of Economics*, 134, 451–505.

Ampudia, Miguel, Marco Jacopo Lombardi, and Théodore Renault (2024), "The wage-price pass-through across sectors: evidence from the euro area," Working Paper Series 2948, European Central Bank.

Antolín-Díaz, Juan, and Juan F. Rubio-Ramírez (2018), "Narrative sign restrictions for svars," *American Economic Review*, 108, p. 2802–29, https://doi.org/10.1257/aer.20161852.

Ascari, Guido, and Timo Haber (2022), "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy," *The Economic Journal*, 132, 37–57.

Bai, Jushan, and Serena Ng (2008), "Forecasting economic time series using targeted predictors," *Journal of Econometrics*, 146, 304–317.

Ball, Laurence, and N. Gregory Mankiw (1995), "Relative-Price Changes as Aggregate Supply Shocks," *The Quarterly Journal of Economics*, 110, 161–193.

Bańbura, Marta, Elena Bobeica, and Catalina Martínez Hernández (2023), "What drives core inflation? the role of supply shocks," *ECB Working Paper No 2875*.

Banbura, Marta, Michele Lenza, and Joan Paredes (2024), "9. forecasting inflation in the us and in the euro area," *Handbook of Research Methods and Applications in Macroeconomic Forecasting*, p. 218.

Baumeister, Christiane, and James D Hamilton (2015), "Sign restrictions, structural vector autoregressions, and useful prior information," *Econometrica*, 83, 1963–1999.

——— (2018), "Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations," *Journal of Monetary Economics*, 100, 48–65.

Baumeister, Christiane, Florian Huber, and Massimiliano Marcellino (2024), "Risky oil: It's all in the tails,"Technical report, National Bureau of Economic Research.

Benigno, Gianluca, Julian di Giovanni, Jan J. J. Groen, and Adam I. Noble (2022), "The GSCPI: A New Barometer of Global Supply Chain Pressures," Staff Reports 1017, Federal Reserve Bank of New York.

Benigno, Pierpaolo, and Gauti B Eggertsson (2023), "It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve,"Technical report, National Bureau of Economic Research.

Bobeica, Elena, Matteo Ciccarelli, and Isabel Vansteenkiste (2020), "The Link between Labor Cost Inflation and Price Inflation in the Euro Area," in *Changing Inflation Dynamics,Evolving Monetary Policy* eds. by Gonzalo Castex, Jordi Galí, and Diego Saravia, 27 of Central Banking, Analysis, and Economic Policies Book Series: Central Bank of Chile, chap. 4, 071–148.

Bruns, Martin, and Michele Piffer (2023), "Tractable bayesian estimation of smooth transition vector autoregressive models," *The Econometrics Journal*.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2015), "Bayesian vars: Specification choices and forecast accuracy," *Journal of Applied Econometrics*, 30, 46–73, https://doi.org/https://doi.org/10.1002/jae.2315.

Carriero, Andrea, Massimiliano Marcellino, and Tommaso Tornese (2024), "Blended identification in structural vars," *Journal of Monetary Economics*, 146, p. 103581, https://doi.org/https://doi.org/10.1016/j.jmoneco.2024.103581.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara (2023), "Inflation and misallocation in New Keynesian models," sintra paper, ECB Forum on Central Banking June 2023.

Chahad, Mohammed, Anna-Camilla Hofmann-Drahonsky, Adrian Page, and Marcel Tirpák (2023), "An updated assessment of short-term inflation projections by eurosystem and ecb staff,"Technical report, European Central Bank.

Chan, Joshua, Eric Eisenstat, and Xuewen Yu (2022), "Large bayesian vars with factor stochastic volatility: Identification, order invariance and structural analysis," *arXiv preprint arXiv:2207.03988*.

Chipman, Hugh A., Edward I. George, and Robert E. McCulloch (2010), "BART: Bayesian additive regression trees," *The Annals of Applied Statistics*, 4, 266 – 298, https://doi.org/10.1214/09-AOAS285.

Clark, Todd E., Florian Huber, Gary Koop, Massimiliano Marcellino, and Michael Pfarrhofer (2023a), "Tail forecasting with multivariate bayesian additive regression trees," *International Economic Review*, 64, 979–1022, https://doi.org/https://doi.org/10.1111/iere.12619.

Clark, Todd E, Florian Huber, Gary Koop, Massimiliano Marcellino, and Michael Pfarrhofer (2023b), "Tail forecasting with multivariate bayesian additive regression trees," *International Economic Review*, 64, 979–1022.

Cogley, Timothy (2005), "How fast can the new economy grow? a bayesian analysis of the evolution of trend growth," *Journal of macroeconomics*, 27, 179–207.

Cogley, Timothy, and Thomas J Sargent (2001), "Evolving post-world war ii us inflation dynamics," *NBER macroeconomics annual*, 16, 331–373.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc (2014), "The international dimension of productivity and demand shocks in the us economy," *Journal of the European Economic Association*, 12, 153–176, https://doi.org/10.1111/jeea.12070.

De Santis, Roberto A., and Tommaso Tornese (2023), "Energy supply shocks' nonlinearities on output and prices," Working Paper Series 2834, European Central Bank.

Dedola, Luca, Lukas Henkel, Christian Höynck, Chiara Osbat, and Sergio Santoro (2024), "What does new micro price evidence tell us about inflation dynamics and monetary policy transmission?," *Economic Bulletin Articles*, 3.

Del Negro, Marco, and Giorgio E. Primiceri (2013), "Time-Varying Structural Vector Autoregressions and Monetary Policy: a Corrigendum," Staff Reports 619, Federal Reserve Bank of New York.

Forni, Mario, Luca Gambetti, Nicolò Maffei-Faccioli, and Luca Sala (2022), "The nonlinear transmission of financial shocks: Some evidence."

Galvao, Ana Beatriz, and Massimiliano Marcellino (2014), "The effects of the monetary policy stance on the transmission mechanism," *Studies in Nonlinear Dynamics & Econometrics*, 18, 217–236.

Gautier, Erwan, Cristina Conflitti, Riemer P. Faber, Brian Fabo, Ludmila Fadejeva, Valentin Jouvanceau, Jean-Oliver Menz, Teresa Messner, Pavlos Petroulas, Pau Roldan-Blanco, Fabio Rumler, Sergio Santoro, Elisabeth WIeland, and Helene Zimmer (2024), "New facts on consumer price rigidity in the euro area," *American Economic Journal: Macroeconomics*, Forthcoming.

Ghysels, Eric, Leonardo Iania, and Jonas Striaukas (2018), "Quantile-based inflation risk models,"Technical report, NBB Working Paper.

Gitti, Giulia (2024), "Nonlinearities in the regional phillips curve with labor market tightness,"Technical report, mimeo.

Goulet Coulombe, Philippe (2024), "A neural phillips curve and a deep output gap," *Journal of Business & Economic Statistics*, 1–22.

Hamilton, James D, and Gang Lin (1996), "Stock market volatility and the business cycle," *Journal of applied econometrics*, 11, 573–593.

Hamilton, James D, and Gabriel Perez-Quiros (1996), "What do the leading indicators lead?" *Journal of Business*, 27–49.

Harding, Martín, Jesper Lindé, and Mathias Trabandt (2023), "Understanding post-Covid inflation dynamics," BIS Working Papers 1077, Bank for International Settlements.

Hauzenberger, Niko, Florian Huber, and Karin Klieber (2023), "Real-time inflation forecasting using non-linear dimension reduction techniques," *International Journal of Forecasting*, 39, 901–921.

Hauzenberger, Niko, Florian Huber, Karin Klieber, and Massimiliano Marcellino (2024), "Bayesian neural networks for macroeconomic analysis," *Journal of Econometrics*, p. 105843, https://doi.org/https://doi.org/10.1016/j.jeconom.2024.105843.

Hauzenberger, Niko, Florian Huber, Massimiliano Marcellino, and Nico Petz (2025), "Gaussian process vector autoregressions and macroeconomic uncertainty," *Journal of Business & Economic Statistics*, 43, 27–43.

Head, Allen, Alok Kumar, and Beverly Lapham (2010), "Market Power, Price Adjustment, And Inflation," *International Economic Review*, 51, 73–98.

Huber, Florian, Gary Koop, Luca Onorante, Michael Pfarrhofer, and Josef Schreiner (2023), "Nowcasting in a pandemic using non-parametric mixed frequency vars," *Journal of Econometrics*, 232, 52–69, https://doi.org/https://doi.org/10.1016/j.jeconom.2020.11.006.

Huber, Florian, and Luca Rossini (2022), "Inference in Bayesian additive vector autoregressive tree models," *The Annals of Applied Statistics*, 16, 104 – 123, https://doi.org/10.1214/21-AOAS1488.

Hubrich, Kirstin, and Timo Teräsvirta (2013), "Thresholds and smooth transitions in vector autoregressive models," in *VAR models in macroeconomics–New developments and applications: Essays in honor of Christopher A. Sims*: Emerald Group Publishing Limited, 273–326.

Kadiyala, K Rao, and Sune Karlsson (1997), "Numerical methods for estimation and inference in bayesian var-models," *Journal of Applied Econometrics*, 12, 99–132.

Karadi, Peter, Anton Nakov, Galo Nuno Barrau, Ernesto Pasten, and Dominik Thaler (2024), "Strike while the iron is hot: optimal monetary policy with a nonlinear Phillips curve," BIS Working Papers 1203, Bank for International Settlements.

Karadi, Peter, and Adam Reiff (2019), "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, 11, 111–146.

Koop, Gary, and Dimitris Korobilis (2010), "Bayesian multivariate time series methods for empirical macroeconomics," *Foundations and Trends (R) in Econometrics*, 3, 267–358.

Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996), "Impulse response analysis in nonlinear multivariate models," *Journal of econometrics*, 74, 119–147.

Korobilis, Dimitris (2022), "A new algorithm for structural restrictions in bayesian vector autoregressions," *European Economic Review*, 148, p. 104241.

Krolzig, Hans-Martin (2013), *Markov-switching vector autoregressions: Modelling, statistical inference, and application to business cycle analysis*, 454: Springer Science & Business Media.

Kuik, Friderike, Jakob Feveile Adolfsen, Aidan Meyler, and Eliza Lis (2022), "Energy price developments in and out of the COVID-19 pandemic – from commodity prices to consumer prices," *Economic Bulletin Articles*, 4.

Lenza, Michele, Inès Moutachaker, and Joan Paredes (2023), "Density forecasts of inflation: a quantile regression forest approach," *ECB Working Paper No. 2830*.

Li, Ping, and Songcan Chen (2016), "A review on gaussian process latent variable models," *CAAI Transactions on Intelligence Technology*, 1, 366–376.

Litterman, Robert B (1979), "Techniques of forecasting using vector autoregressions.."

——— (1980), *Bayesian procedure for forecasting with vector autoregressions*: Massachusetts Institute of Technology.

Lopez-Salido, David, and Francesca Loria (2022), "Inflation at risk," *Available at SSRN 4002673*.

McCulloch, Robert E, Rodney A Sparapani, Brent R Logan, and Purushottam W Laud (2021), "Causal inference with the instrumental variable approach and bayesian nonparametric machine learning," *arXiv preprint arXiv:2102.01199*.

Montag, Hugh, and Daniel Villar Vallenas (2023), "Price-Setting During the Covid Era," FEDS Notes 2023-08-29, Board of Governors of the Federal Reserve System (U.S.), https://doi.org/10.17016/2380-7172.3298.

Mumtaz, Haroon, and Michele Piffer (2022), "Impulse response estimation via flexible local projections," *arXiv preprint arXiv:2204.13150*.

Phillips, Alban Williams (1958), "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, 25, 283–299.

Primiceri, Giorgio E (2005), "Time varying structural vector autoregressions and monetary policy," *The Review of Economic Studies*, 72, 821–852.

Prüser, Jan (2021), "The horseshoe prior for time-varying parameter vars and monetary policy," *Journal of Economic Dynamics and Control*, 129, p. 104188.

Sims, Christopher A, and Tao Zha (2006), "Were there regime switches in us monetary policy?" *American Economic Review*, 96, 54–81.

Song, Hanyu, and David B Dunson (2022), "Curved factor analysis with the ellipsoid-gaussian distribution," *arXiv preprint arXiv:2201.08502*.

Song, Xin-Yuan, and Zhao-Hua Lu (2012), "Semiparametric transformation models with bayesian p-splines," *Statistics and Computing*, 22, 1085–1098.

Stella, Andrea, and James H Stock (2013), "A state-dependent model for inflation forecasting," *FRB International Finance Discussion Paper*.

Szafranek, Karol, Grzegorz Szafrański, and Agnieszka Leszczyńska-Paczesna (2024), "Inflation returns. Revisiting the role of external and domestic shocks with Bayesian structural VAR," *International Review of Economics & Finance*, 93, 789–810, https://doi.org/10.1016/j.iref.2024.03.05.

Xu, Maoran, Amy H Herring, and David B Dunson (2023), "Identifiable and interpretable nonparametric factor analysis," *arXiv preprint arXiv:2311.08254*.

Yalcin, Ilker, and Yasuo Amemiya (2001), "Nonlinear factor analysis as a statistical method," *Statistical science*, 275–294.

Zlobins, Andrejs (2025), "Monetary policy transmission in the euro area: is this time different? chapter i: lags and strength," *Latvijas Banka Working Paper 01/2025*.

A Additional Results

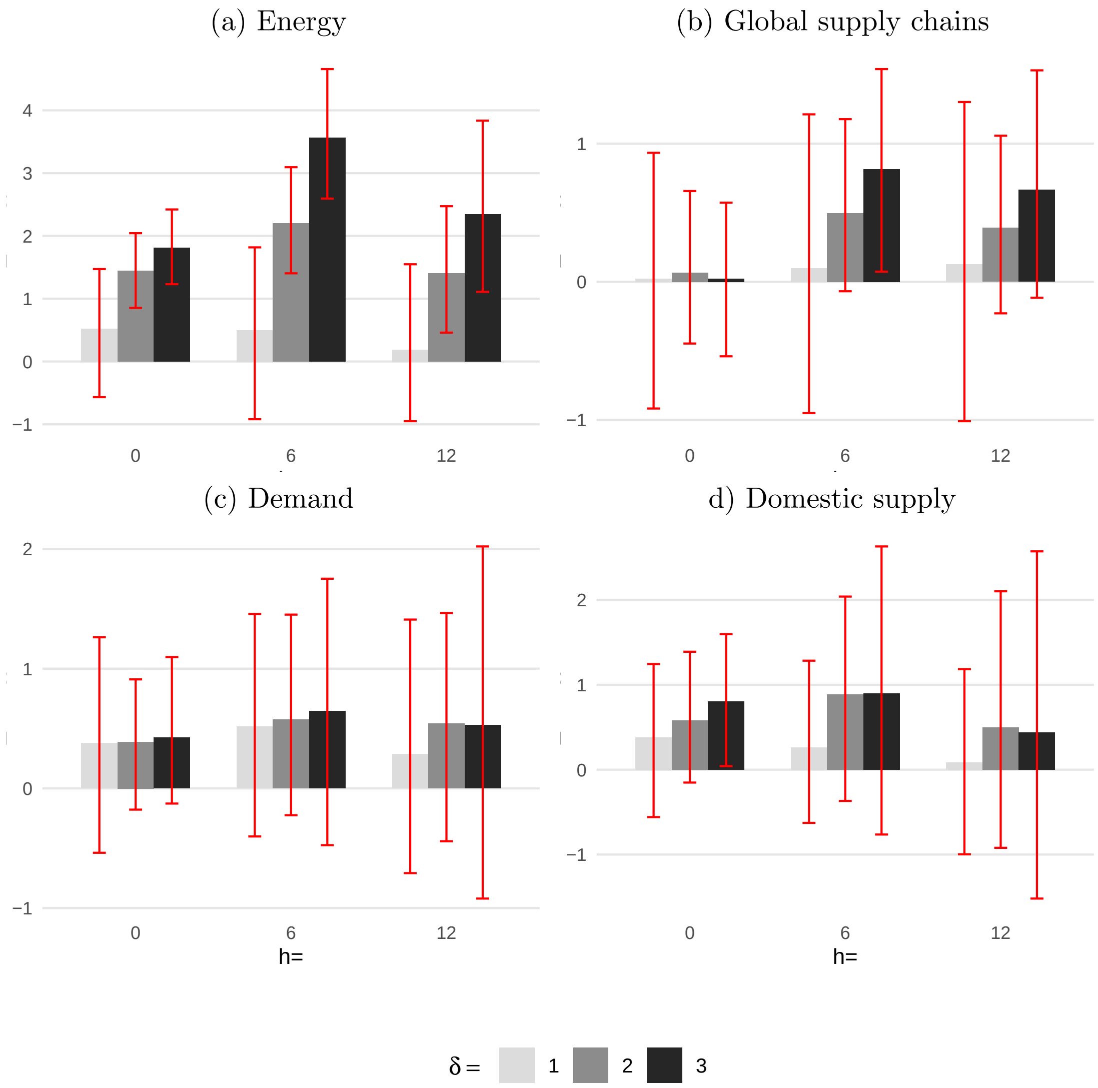

Figure A1: Reaction of PPI energy to the identified structural shocks

Notes: Bars show the 68% credible interval. Shock sizes and impulse response horizons are on the X-axis. IRFs are normalised for comparability across the three shock sizes.

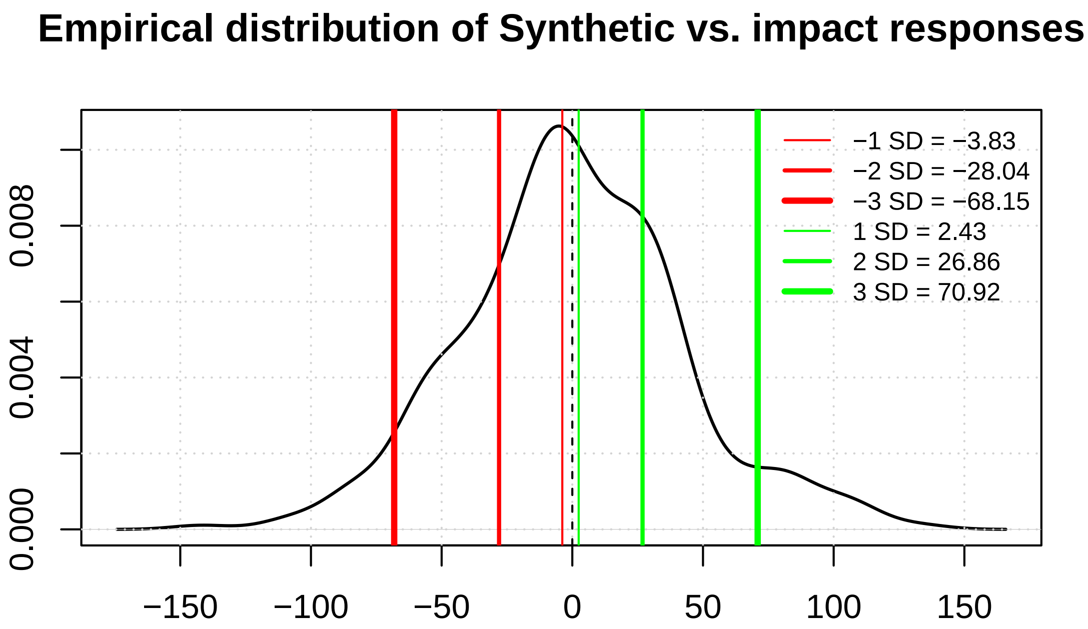

Notes: Synthetic indicator expressed in annual growth terms.

Figure A2: Synthetic energy indicator and contemporaneous reaction to shocks of different sizes

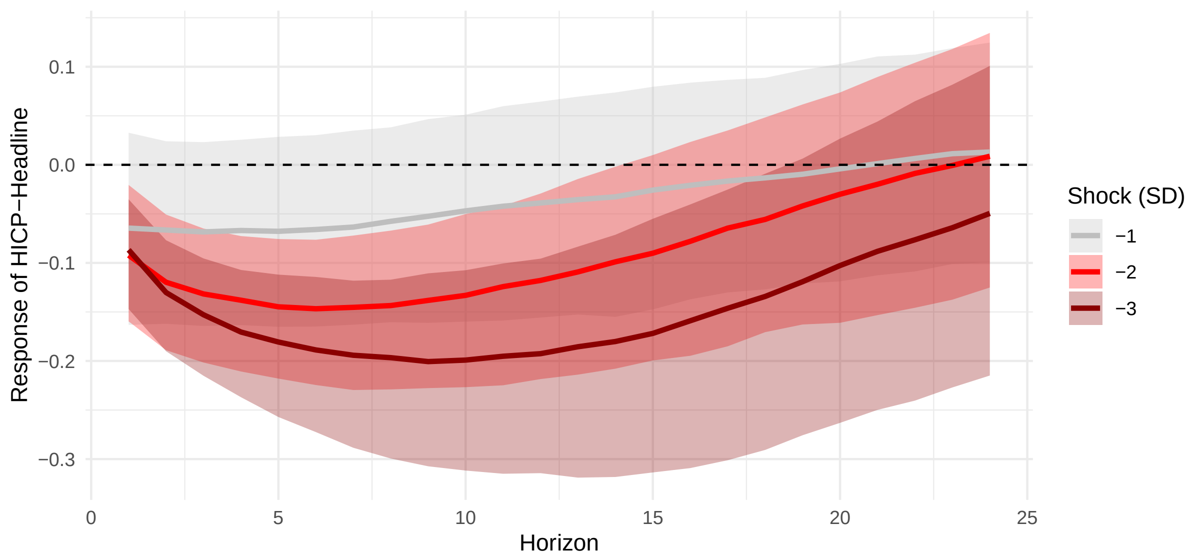

Notes: thick lines are median estimates and the shaded areas are the 68% credible interval. IRFs are normalised for comparability.

Figure A3: Reaction of inflation to negative small, medium, and large energy shocks h!]

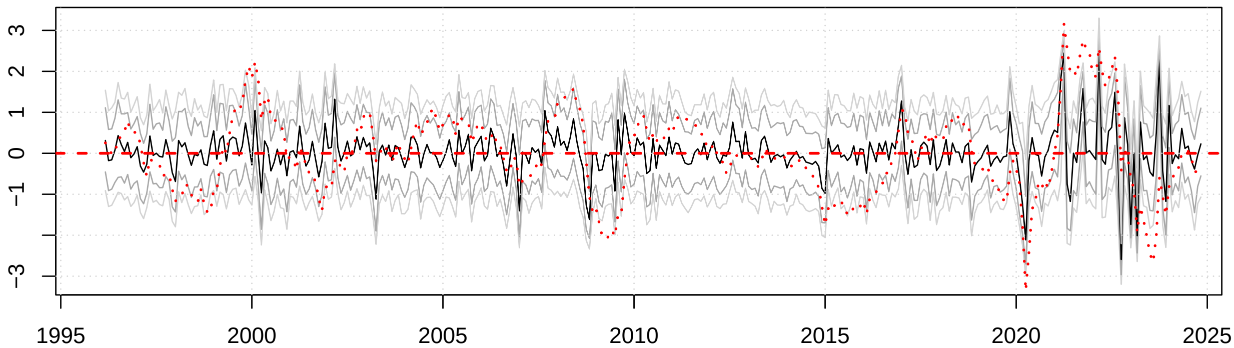

Note: Grey lines portray the 68th and the 90th credibility bands, while the black lines depict the median of the shocks' posterior distribution. The red, dotted line shows the dynamics of the synthetic indicator, the variable used to normalise the shocks.

Figure A4: Estimated energy shocks