

Beware of large shocks!

A non-parametric structural inflation model

Elena Bobeica¹ Florian Huber² Sarah Holton¹ Catalina Martínez Hernández¹

¹European Central Bank ²University of Salzburg

Monetary policy and structural tectonic shifts, OeNB SUERF Annual Economic Conference,
23 May 2025

The views expressed herein are solely those of the authors and do not necessarily reflect the views of the European Central Bank or the Eurosystem.

Motivation

- ▶ Linear models (even large ones) are not able to explain the full extent of the post-pandemic inflation surge. A notable share tends to be left unexplained (Bańbura et al., 2023)
 - ▶ Unusually large underpredictions of inflation by professional forecasters in the post-pandemic recovery
- ▶ Gap in the literature: an inflation model that is agnostic on the type of non-linearity that governs the macroeconomic relationships and allows for structural inference.

This paper in a nutshell

- ▶ We build a nonparametric structural inflation model
 - ▶ Flexible on the shape possible non-linearities
 - ▶ Allows for identification of shocks
- ▶ We focus on energy shocks
 - ▶ We find that larger shocks have a disproportionately larger impact on inflation
 - ▶ Non-linearity more visible at the beginning of the pricing chain
 - ▶ Other shocks also elicit non-linear reactions
- ▶ Non-linear model better fits the data compared to a linear version in an out-of-sample forecasting exercise

Inflation and synthetic energy commodity price indicator

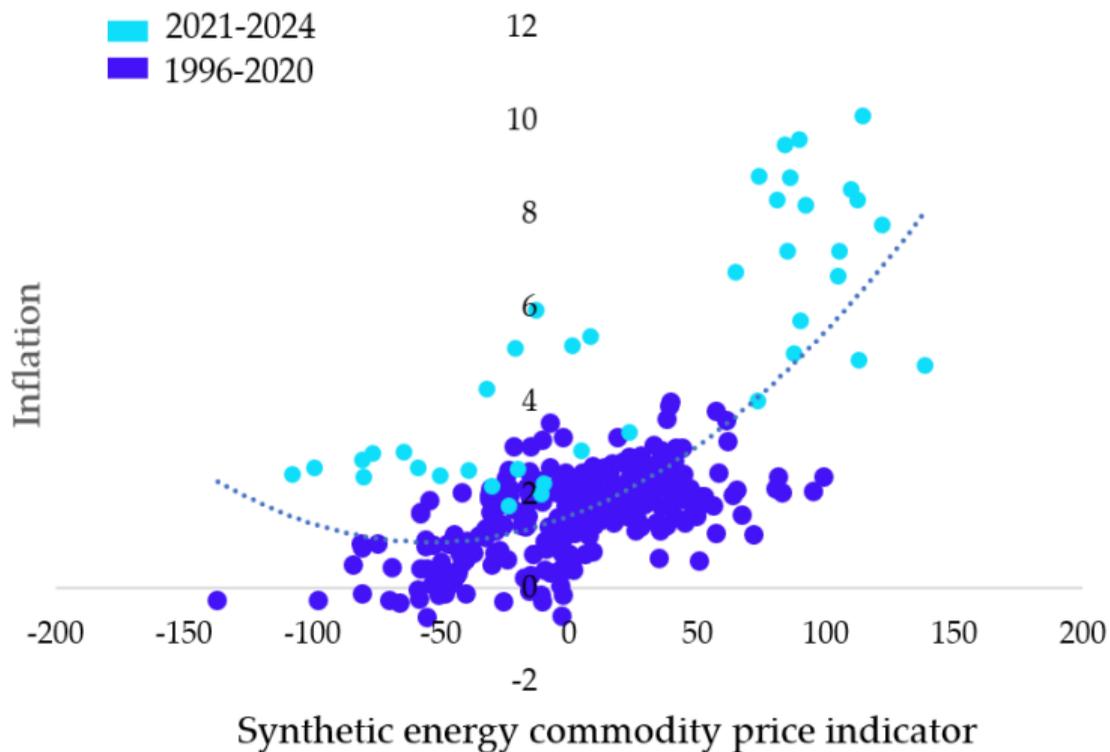


Figure: Inflation and energy commodity prices

Literature on non-linear reaction of inflation to large shocks - Theory

- ▶ **Menu costs** – firms have a range of inaction in response to shocks; small shocks do not trigger price reactions (Ball and Mankiw (1995), Harding et al. (2023))
- ▶ **Prices become more flexible** – micro studies support state-dependent pricing; the inflation spike was associated with a large increase in the frequency of price revisions (Cavallo et al. (2023), Hall (2023))
- ▶ **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, making it harder for firms to increase prices (Head et al. (2010))

Literature on non-linear reaction of inflation to large shocks - Empirics

- ▶ **Larger shocks have stronger impacts** - see Ascari and Haber (2022) for monetary policy shocks.
- ▶ **Stronger reactions in a high inflation regime** - see De Santis and Tornese (2023) for the case of energy shocks, Bobeica et al. (2020) for the case of labour cost shocks), Ascari and Haber (2022) for monetary policy shocks.
- ▶ **Stronger reactions after the pandemic** – inflation IRFs after the pandemic change shape and magnitude (Adolfson et al. (2024), Szafranek et al. (2024), Allayioti et al. (2024), Zlobins (2025)).

Literature on non-parametric models - focus on forecasting

- ▶ A fast-growing literature evaluates the use of machine learning techniques for macroeconomic forecasting with success
 - ▶ Random forests work particularly well (see Goulet Coulombe, 2020, arXiv; Medeiros et al., 2021, JBES; Clark et al., 2023, IER; Lenza et al., 2023, Clark et al., 2024, AoAS)
 - ▶ Recent papers use neural networks for modeling inflation (Goulet Coulombe, 2021, arXiv; Hauzenberger et al., 2022, IJoF)
 - ▶ Other papers use approximation techniques to estimate complicated models of inflation (see, e.g., Korobilis, 2019, JBES; Hauzenberger et al. 2021, JBES)
- ▶ These models are capable of capturing outliers in macro series flexibly and thus produce superior density forecasts (see Huber et al. 2023, JoE)

The model

- ▶ We assume that the M -dimensional data vector \mathbf{y}_t is a function of $\mathbf{x}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$ and a set of Q fundamental shocks \mathbf{q}_t :

$$\mathbf{y}_t = f(\mathbf{x}_t) + g(\mathbf{q}_t) + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}),$$

The model

- ▶ We assume that the M -dimensional data vector \mathbf{y}_t is a function of $\mathbf{x}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$ and a set of Q fundamental shocks \mathbf{q}_t :

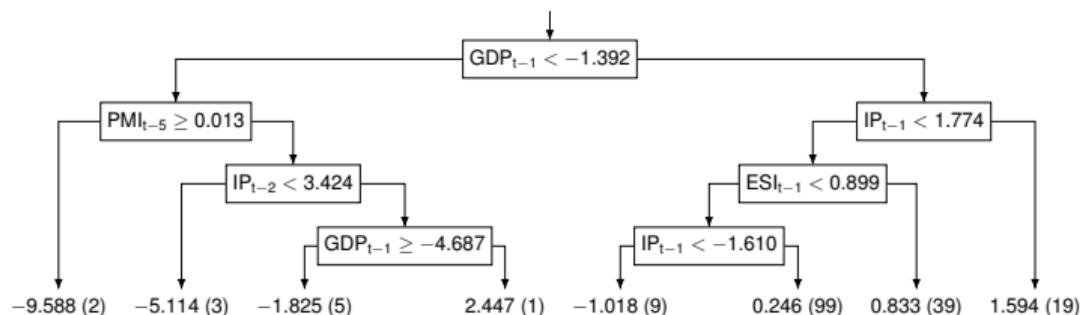
$$\mathbf{y}_t = \underbrace{f(\mathbf{x}_t)}_{\mathbf{A}\mathbf{x}_t + \tilde{f}(\mathbf{x}_t)} + \underbrace{g(\mathbf{q}_t)}_{\Lambda\mathbf{q}_t + \tilde{g}(\mathbf{q}_t)} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}),$$

where

- ▶ $f : \mathbb{R}^K \rightarrow \mathbb{R}^M$ and $g : \mathbb{R}^Q \rightarrow \mathbb{R}^M$ are unknown (and possibly non-linear) functions
- ▶ \mathbf{A} is a $M \times K$ dimensional coefficient matrix and Λ is a matrix of factor loadings
- ▶ \tilde{f} and \tilde{g} are unknown functions
- ▶ $\boldsymbol{\eta}_t$ are measurement errors with diagonal covariance matrix $\boldsymbol{\Omega}$
- ▶ $f(\mathbf{x}_t) = \mathbf{A}\mathbf{x}_t + \tilde{f}(\mathbf{x}_t)$ is a mixture Bayesian Additive Regression Tree (BART) specification used in, e.g., Clark et al. (2023)
- ▶ The static factors \mathbf{q}_t enter the model in terms of a linear ($\Lambda\mathbf{q}_t$) and a non-linear ($\tilde{g}(\mathbf{q}_t)$) component

Approximation of non-linear functions: Bayesian Additive Regression Trees

- ▶ For estimating \tilde{f} and \tilde{g} we opt for Bayesian Additive Regression Trees (Chipman, George & McCulloch, 2010, AoAS)
- ▶ Example: suppose we wish to explain German GDP growth using lagged GDP, industrial production (IP), the Economic Sentiment Indicator (ESI) and the Purchasing Manager Index (PMI)
- ▶ An estimated tree would look like this:



Data

- ▶ Focus on HICP headline inflation
- ▶ Additional endogenous variables: Synthetic commodity price index, consumer confidence indicator (CCI), GSCPI of Benigno et al (2022), selling price expectations in manufacturing, PMI supplier delivery times, PPI total, PPI energy, EUR/USD exchange rate, Euribor 1Y
- ▶ Sample covers January 1996-November 2024, transformed to year-on-year rates when needed
- ▶ Model with 2 lags of data

Identification

- ▶ Identification restrictions on the VAR part:
 - ▶ Tight priors on the linear VAR part to account for most of the persistence in the data \Rightarrow implies that $\tilde{f}(\mathbf{x}_t)$ soaks up the remainder term
 - ▶ To avoid over-fitting we fix the explained share of variance in the data to 95%
- ▶ Economic identification:
 - ▶ Economic identification of shocks via sign and zero restrictions on the elements of Λ as in Korobilis (2022), assuming truncated priors
 - ▶ Sign restrictions ensure that each element of \mathbf{q}_t has an economic meaning (see next slide)
 - ▶ We normalize each structural shock by linking them to relevant indicators:(e.g. we restrict the elements of Λ such that a one unit energy shock creates a one standard deviation increase in the synthetic indicator)

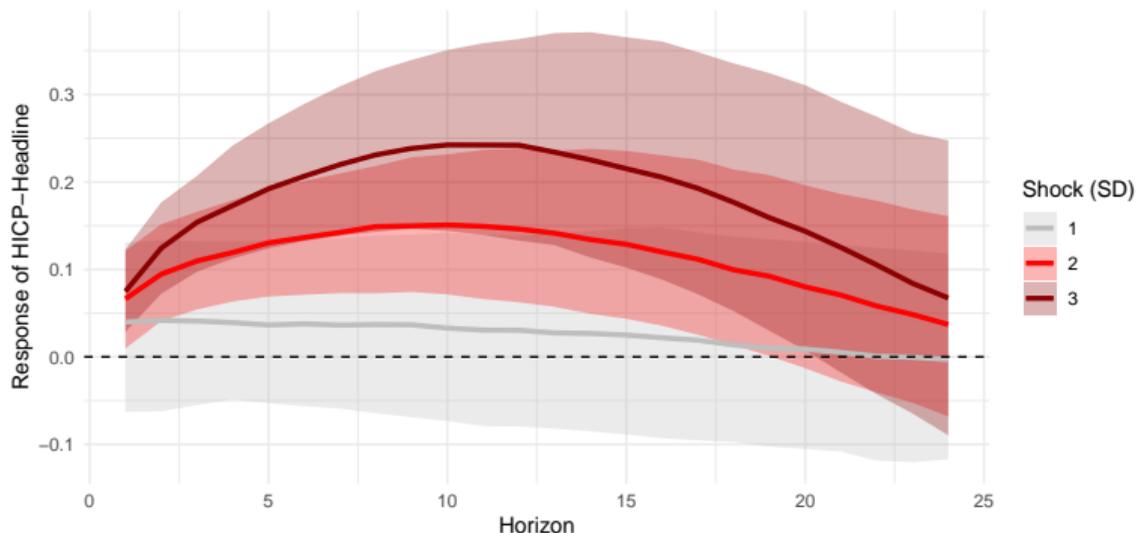
Identification scheme

Table: Structural shock identification

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

Notes: Restrictions are contemporaneous, σ_y reflects a normalisation restriction.

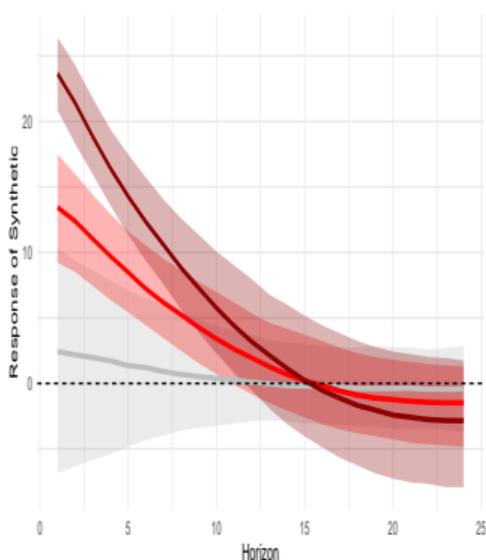
Reaction of headline inflation to energy price shocks of different sizes



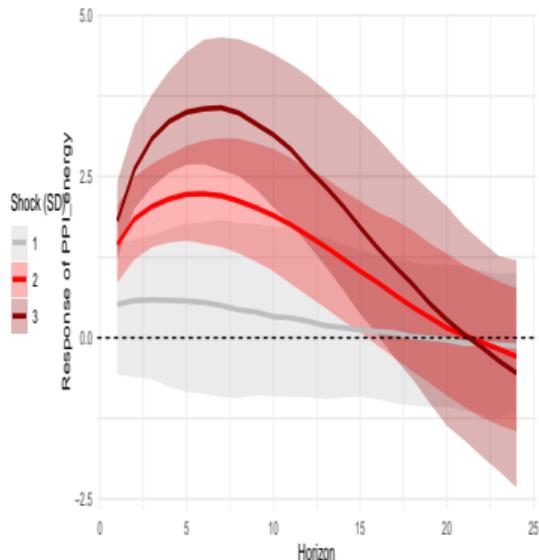
Thick lines are median estimates and the shaded areas are the 68% credible interval. IRFs are normalized for comparability.

Transmission of energy shocks along the pricing chain

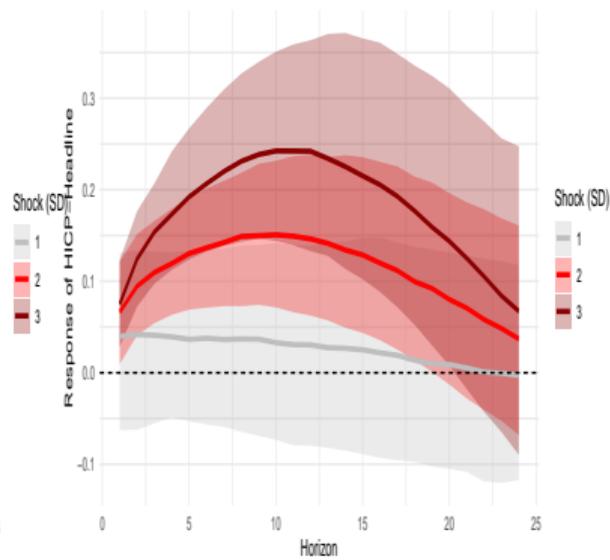
Synthetic indicator



PPI energy

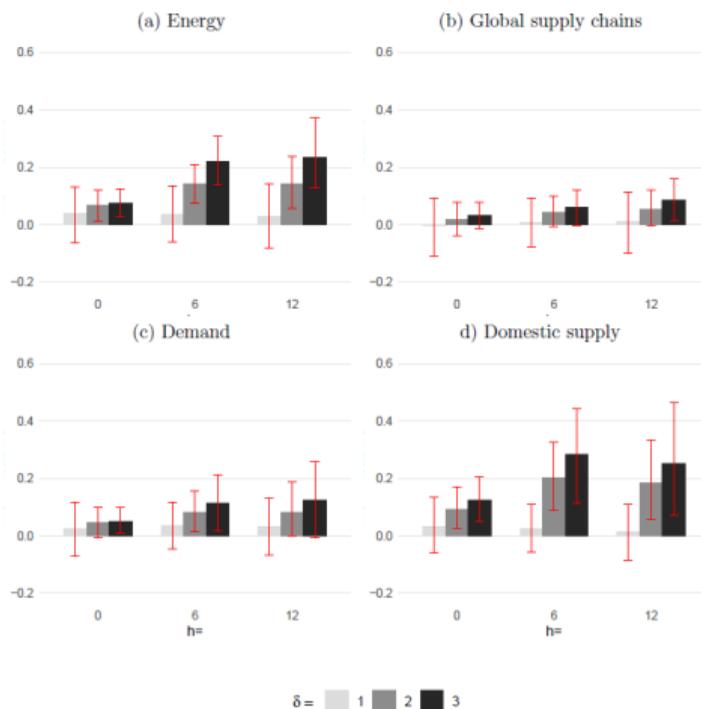


Headline inflation



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

Reaction of inflation to the identified structural shocks



Bars show the 68% credible interval. IRFs are normalized for comparability across the three shock sizes.

Validation of the model from a forecasting perspective

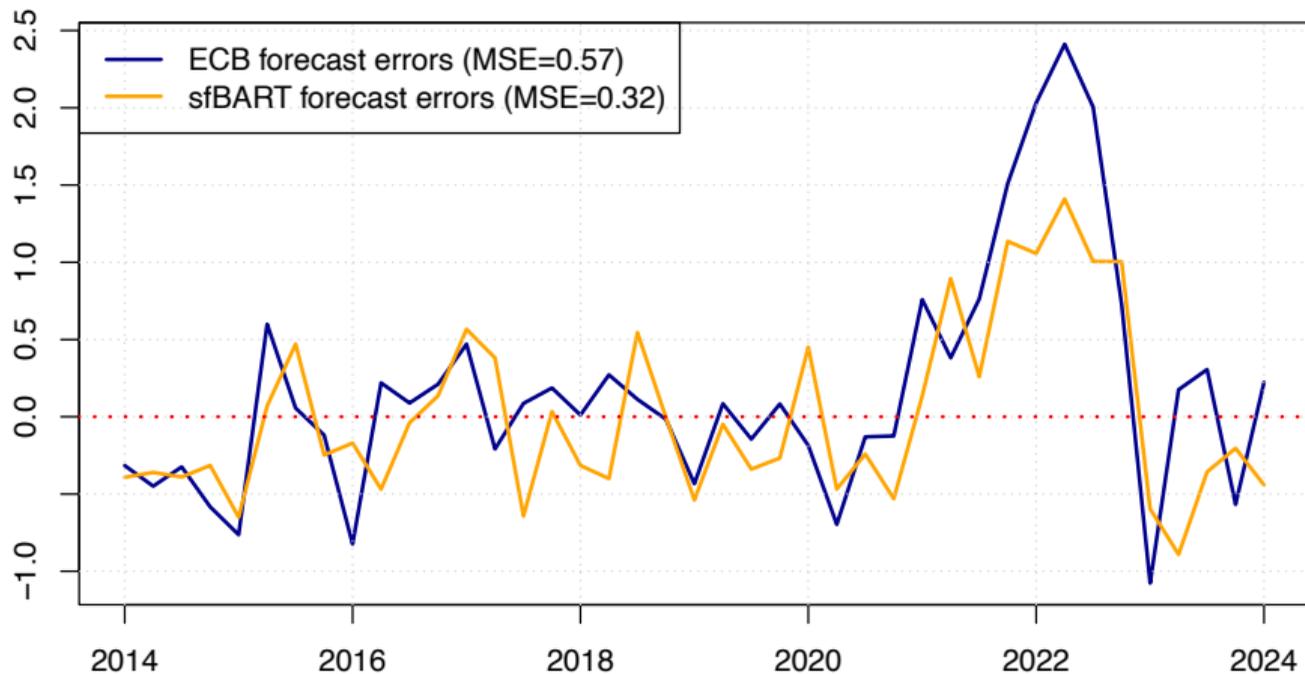
- ▶ Benchmark: Linear VAR, state of the art model with Horseshoe priors and factor stochastic volatility: $\mathbf{y}_t = \mathbf{A}\mathbf{x}_t + \mathbf{\Lambda}\mathbf{q}_t + \boldsymbol{\eta}_t$
- ▶ Metric: quantile weighted CRPS scores, can be considered as the generalization of the mean absolute forecast error that take into account the full predictive distribution.
 1. Right-tailed CRPS: puts more weight on the right tail of the forecast distribution of inflation (\approx deflation risk)
 2. Center 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 (\approx high inflation risks)

- ▶ Recursive forecasting design:
 - ▶ Use data up to 2008:12 to train the models initially
 - ▶ Forecast 1 up to 12 steps-ahead to compute predictive distributions for 2009:1 up to 2024:2
 - ▶ After obtaining these, add one month to the training sample and re-estimate the models

CRPS of BART models relative to CPRS of BVAR model 2009:1 to 2024:4

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

Forecast errors ECB forecasts vs. BART



Robustness checks conducted so far

- ▶ Economic identification - different scheme linking energy shocks to PPI energy and HICP energy inflation
- ▶ Statistical identification - robustness w.r.t. prior on A and robustness w.r.t. to prior on Lambda
- ▶ Parametrization - different number of trees

Conclusions

- ▶ New model with non-linear features and structural identification.
- ▶ Good forecasting properties compared to a linear version, especially when large shocks hit.
- ▶ Large shocks generate disproportional inflation responses.
- ▶ Policy implications: large shocks transmit differently and therefore may require differentiated response.

References I

- Adolfson, J. F., Ferrari Minesso, M., Mork, J. E., and Van Robays, I. (2024). Gas price shocks and euro area inflation. *Journal of International Money and Finance*, 149:103183.
- Allayioti, A., Górnicka, L., Holton, S., and Martínez Hernández, C. (2024). Monetary policy pass-through to consumer prices: evidence from granular price data.
- Ascari, G. and Haber, T. (2022). Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy. *The Economic Journal*, 132(641):37–57.
- Ball, L. and Mankiw, N. G. (1995). Relative-Price Changes as Aggregate Supply Shocks. *The Quarterly Journal of Economics*, 110(1):161–193.
- Bañbura, M., Bobeica, E., and Martínez Hernández, C. (2023). What drives core inflation? the role of supply shocks. *ECB Working Paper No 2875*.

References II

- Bobeica, E., Ciccarelli, M., and Vansteenkiste, I. (2020). The Link between Labor Cost Inflation and Price Inflation in the Euro Area. In Castex, G., Galí, J., and Saravia, D., editors, *Changing Inflation Dynamics, Evolving Monetary Policy*, volume 27 of *Central Banking, Analysis, and Economic Policies Book Series*, chapter 4, pages 071–148. Central Bank of Chile.
- Cavallo, A., Lippi, F., and Miyahara, K. (2023). Inflation and misallocation in New Keynesian models. Sintra paper, ECB Forum on Central Banking June 2023.
- Clark, T. E., Huber, F., Koop, G., Marcellino, M., and Pfarrhofer, M. (2023). Tail forecasting with multivariate bayesian additive regression trees. *International Economic Review*, 64(3):979–1022.
- De Santis, R. A. and Tornese, T. (2023). Energy supply shocks' nonlinearities on output and prices. Working Paper Series 2834, European Central Bank.

References III

- Hall, R. E. (2023). A major shock makes prices more flexible and may result in a burst of inflation or deflation. Working Paper 31025, National Bureau of Economic Research.
- Harding, M., Lindé, J., and Trabandt, M. (2023). Understanding post-Covid inflation dynamics. BIS Working Papers 1077, Bank for International Settlements.
- Head, A., Kumar, A., and Lapham, B. (2010). Market Power, Price Adjustment, And Inflation. *International Economic Review*, 51(1):73–98.
- Szafranek, K., Szafrański, G., and Leszczyńska-Paczesna, A. (2024). Inflation returns. Revisiting the role of external and domestic shocks with Bayesian structural VAR. *International Review of Economics & Finance*, 93(PA):789–810.
- Zlobins, A. (2025). Monetary policy transmission in the euro area: is this time different? chapter i: lags and strength. *Latvijas Banka Working Paper 01/2025*.

BACKGROUND

An Introduction to BART

- ▶ BART approximates each $f_j(\mathbf{x}_t)$ as follows:

$$f_j(\mathbf{x}_t) = \sum_{s=1}^S u_s(\mathbf{x}_t | \mathcal{T}_{js}, \mu_{js}),$$

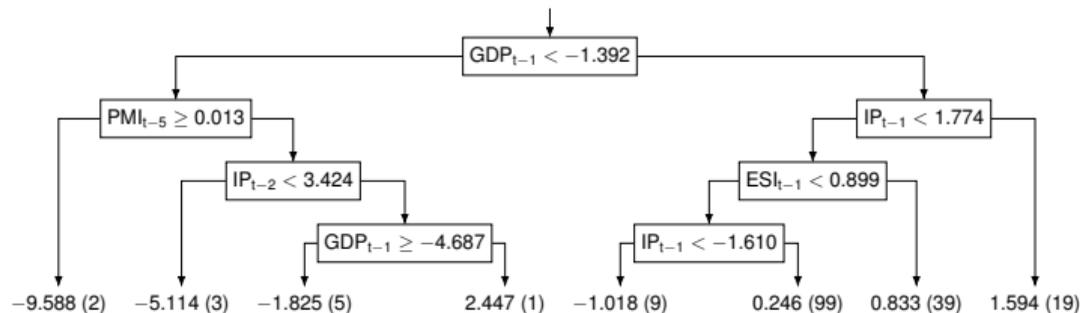
- ▶ \mathcal{T}_{js} are tree structures
- ▶ μ_{js} are tree-specific terminal nodes
- ▶ S denotes the total number of trees used.
- ▶ Dimension of μ_{js} is denoted by b_s which depends on the complexity of the tree
- ▶ Note that BART involves adding up different trees (this is the "A" in BART)
- ▶ But what is a regression tree?

Intuition of a Regression Tree

- ▶ Explain idea of BART for a single tree (suppress s subscripts)
- ▶ Conventional regression: For every value for \mathbf{X} produces a fitted value for \mathbf{y}
- ▶ BART does same thing, but in different way
- ▶ Splitting rule: Divides space of \mathbf{X} into different intervals each of which has same fitted value for \mathbf{y} (internal nodes, branches of tree)
- ▶ Fitted values are called terminal nodes (or leaves of the tree)

An estimated tree

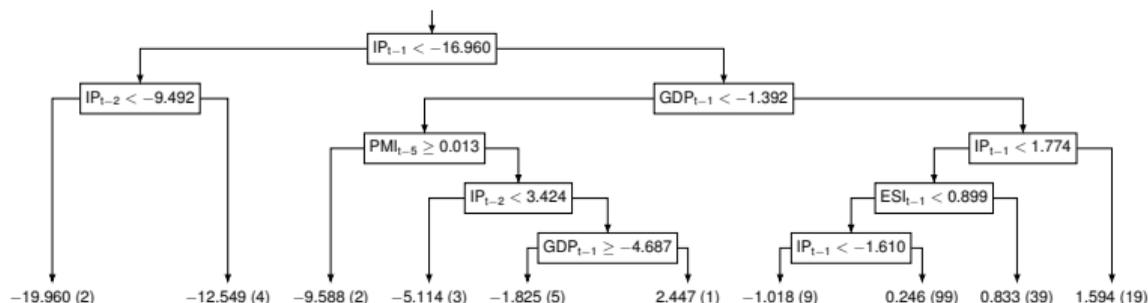
- ▶ Suppose we wish to explain German GDP growth using lagged GDP, industrial production (IP), the Economic Sentiment Indicator (ESI) and the Purchasing Manager Index (PMI)
- ▶ An estimated tree would look like this:



Estimated tree for Germany: 2005Q1 to 2019Q4. Source: Huber et al. (2023)

An estimated tree

- ▶ The previous example did not include pandemic observations
- ▶ The resulting tree looks like this:



Estimated tree for Germany: 2005Q1 to 2020Q2. Source: Huber et al. (2023)

Priors on the tree-based parameters

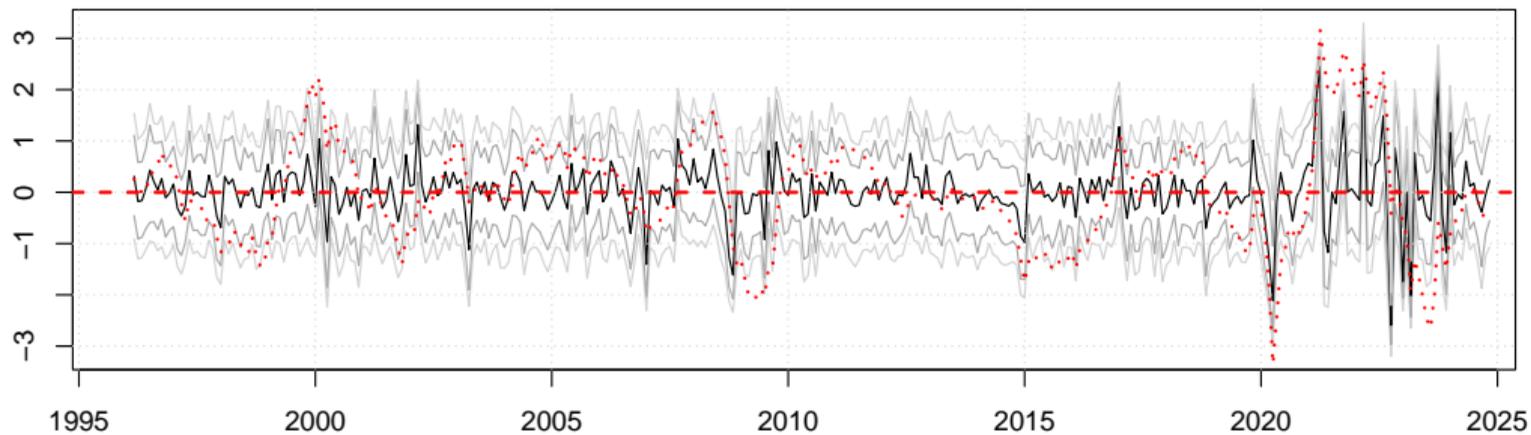
- ▶ The B in BART stands for Bayesian
- ▶ Idea: use priors that make complicated trees less likely
- ▶ Each tree should be simple and explain only a little bit of variation in the response ('weak learner')
- ▶ All this is achieved through regularization priors:
 - ▶ Probability that a node within a tree is non-terminal declines with tree complexity
 - ▶ Terminal node parameters arise from a Gaussian distribution with variance decreasing in the number of trees

Posterior simulation

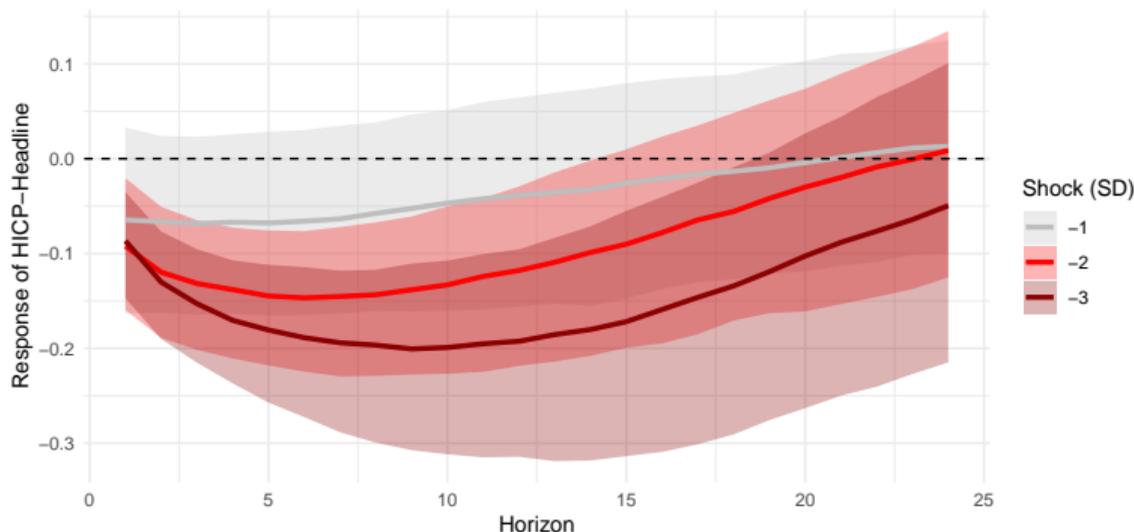
- ▶ Apart from sampling the factors and the loadings, all other steps are standard and based on the algorithm proposed in Clark et al. (2023)
- ▶ The loadings are simulated from truncated Gaussian posterior distributions as in Korobilis (2022)
- ▶ The factors are simulated using a random walk Metropolis Hastings update
 - ▶ For $t = 1, \dots, T$ we propose $\mathbf{q}_t^* \sim \mathcal{N}(\mathbf{q}_t^a, c \cdot \hat{\Sigma}_t)$

where \mathbf{q}_t^a denotes the previously accepted draw and $\hat{\Sigma}_t$ is the time t posterior covariance matrix from a linear factor model and c is a scaling parameter so that the acceptance rate is between 20 and 40 percent

Estimated energy shocks



Reaction of headline inflation to negative energy price shocks of different sizes



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

Reaction of headline inflation to energy shocks over a sizes grid

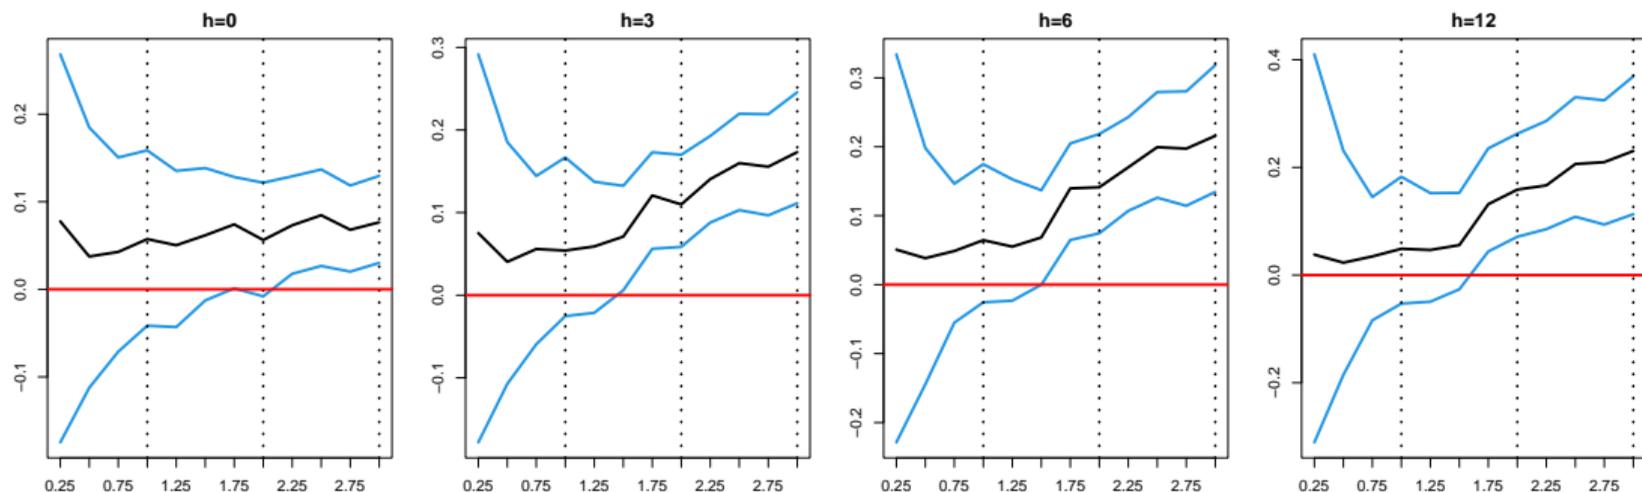
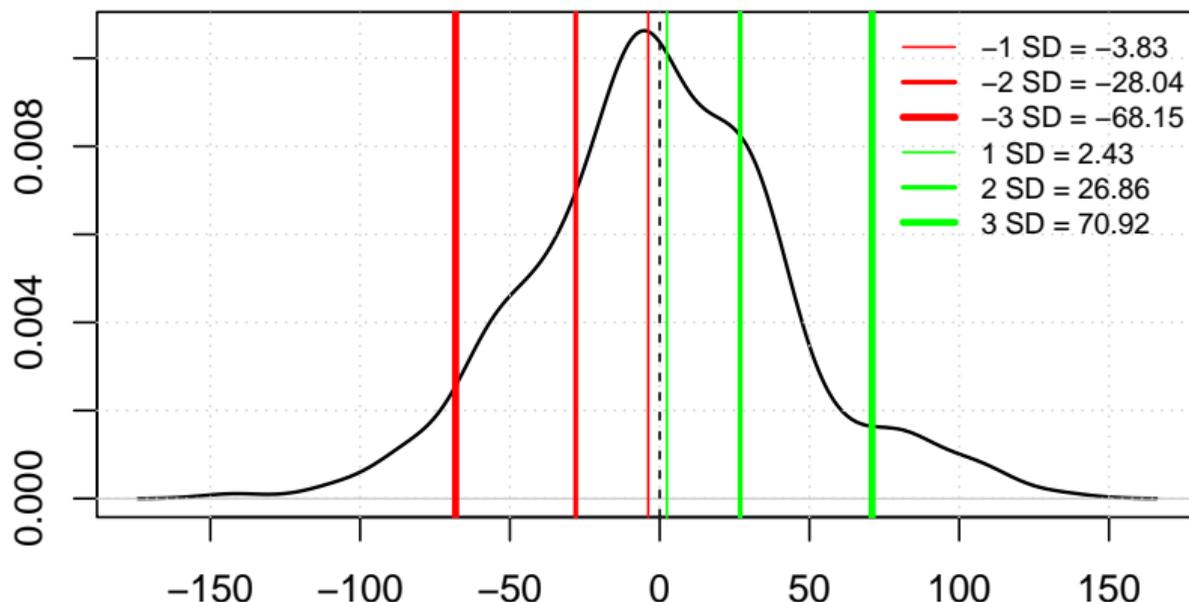


Figure: Reaction of inflation to energy shocks over a sizes grid

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.

Small, medium and large energy shocks

Empirical distribution of Synthetic vs. impact responses



Notes: Synthetic energy indicator and contemporaneous reaction to shocks of different size. Synthetic indicator expressed in annual growth terms.