Quantile Regression Forests in Central Banking: A Data-Driven Approach to Risk and Forecasting

Michele Lenza (ECB and CEPR)
Ines Moutachaker (INSEE)
Joan Paredes (ECB)

ECONDAT Knowledge session
10 April 2025

Why ML?

Why would you use ML? Predictions is one essential goal in science, but how does ML help scientists in making predictions? And what added values adds to existing predictive modelling approaches?

- transparent notion of success: low average prediction error on an unseen test dataset.
- adapts your model to the world and not the world to your model.
- handles different data structures.
- allows you to work on new questions.
- Don't undervalue machine learning because it is new.
- Further justifications for machine learning: Time and computational efficiency. . . .

Great power comes with great responsibilities

ML and other scientific goals. While machine learning has been around for years now, why do economists still have a bad gut feeling about machine learning?

Bare-bones supervised ML: practical insufficiencies or dangers

- Low empirical error on a test set is not enough
- **Domain knowledge** is overlooked
- Lack of **interpret-ability** and explanations
- Predictive performance is at odds with **causality**
- Don't undervalue machine learning because it is new.
- Lack of robustness (**replicability**)
- No **uncertainty** quantification

Aim of the project

- Design and evaluate the accuracy of a new model for euro area **density inflation forecasting**
- No commitment to one type of **non-linearity**, more general than existing models and able to handle large information set monitored in a central bank
- Assess the role of non-linearities for euro area (headline and core) inflation dynamics, by controlling for "overfitting" (out-of-sample accuracy criterion)

$\Rightarrow$ **Quantile regression forests (a variant of Random Forests) as a way to operationalize non-parametric models**

General modelling strategy

- Define our measure of prices as p_t. Assume we have data until time (i.e. month) t. $h = (3, 6, 9, 12)$ months:

$$\pi_t^h = (1200/h) \times [p_t/p_{t-h} - 1]$$

- Estimate $\pi_t = m(\pi_{t-h}...\pi_{t-h-p}; x_{t-h}...x_{t-h-k}) + \varepsilon_t$
- Project forward: $\hat{\pi}_{t+h} = m(\pi_t...\pi_{t-p}; x_t...x_{t-k})$

Main ingredients

- Direct density forecast
- m(.) quantile regression forecasts (variant of the random forest)

Targets: Our measure of prices

Figure: Headline and Core Inflation - year-on-year

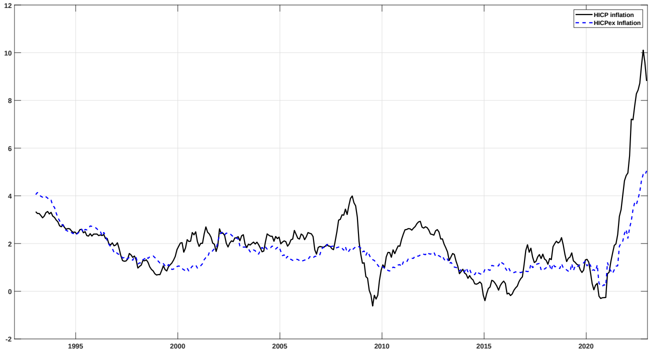

Note: Headline inflation: black solid line; Core inflation: blue dashed line.

Features: Our predictors

- We consider about sixty predictors, routinely monitored at the ECB (de Bondt et al. 2018), plus two inflation lags
- Logic for the choosing the predictors: Phillips Curve.
- Four broad groups of variables: inflation expectations, (domestic and global) cost pressures, real activity and financial variables
- No real-time database (but many variables are timely released and un-revised), stationarized, de-seasonalized (according to out-of-sample logic)

Our baseline model: Regression Tree

Regression trees allow very general relationships between predictors and the target variable

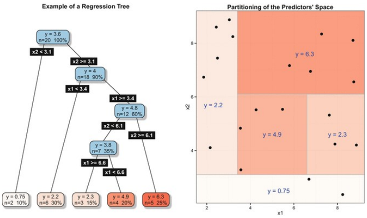

- However, regression trees are normally bad forecasting models, high variance, overfitting
  - One could "prune" them (akin to shrinkage), reducing ex ante their ability to (over-)fit
  - Normally, not the path taken in the literature

Variance reduction is rather achieved by combination of several trees: random forest

From Regression Trees to Random Forest (Breiman 2001) to Quantile Regression Forest (Meinshausen 2006)

1. Grow **many trees**
2. **Bootstrap** observations (and keep the "out-of bag" observations) ⇒ to ensure "diversity" in the trees
3. In each tree, use only a (randomly chosen) **sub-set of predictors** at each node ⇒ step further de-correlates the trees
4. **Combine** the predictions of the trees at the end ⇒ reduces further variance of the forecasts (Variance reduction already maximized when the predictions are not correlated)
5. Density forecasts: rather than taking averages of the target variable in the last nodes, compute sample quantiles ⇒ **Quantile Regression Forest**

Out-of-sample accuracy

- Full sample: January 1992 - December 2022
- About twenty years of out-of-sample evaluation (first estimation sample until end 2001)
- Update by one of observation and re-estimate the model (recursive scheme)
- Forecast horizon: 3, 6, 9 and 12 months ahead; 20 years of out-of-sample evaluation
- CRPS for density forecasts. RMSE for point forecasts

Density assessment

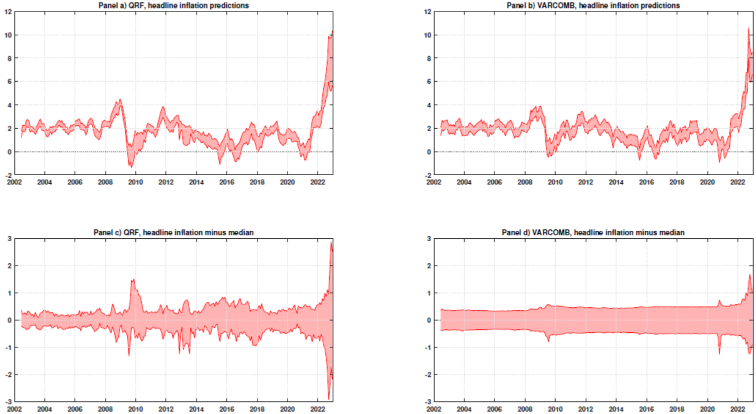

Note: Top panels: h=6 predictive density of year-on-year headline inflation, 16th to 84th quantiles; Bottom panels: 16th to 84th predictive range obtained by subtracting the median forecasts from the quantiles. Sample: June 2002 - December 2022

What we find

Comparison with state-of-the-art linear models

- The quantile regression forest (QRF) is a good forecasting model, especially for core inflation
- Overall, similar accuracy with state-of-the-art linear models on full sample. Different accuracy in sub-samples, diversity in the toolbox

$\Rightarrow$ **Complementarity of the approaches. Non-linearity maybe more relevant in specific episodes and for core inflation.**

Comparison with judgemental institutional and survey forecasts

- QRF is good in terms of relative accuracy, despite not being able to incorporate future info using judgement
- Quite strong collinearity with (judgemental) Eurosystem forecasts!

$\Rightarrow$ **Judgement may be adding mild non-linearity to the Eurosystem forecasts.**

Forecast and risks

Recent inflation developments and short-term outlook

Headline inflation, core inflation and ECB staff projections
(annual percentage changes)

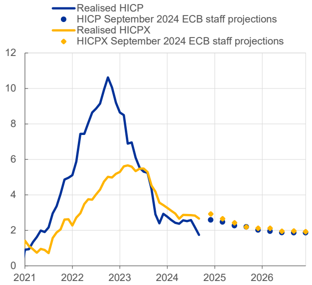

Sources: Eurostat, and September 2024 ECB staff projections.
Notes: Harmonised Index of Consumer Prices (HICP) refers to headline inflation and HICPX to HICP excluding food and energy. Realised HICP and HICPX are at a monthly frequency, and HICP and HICPX projections are at a quarterly frequency. The latest observations are for September 2024.

Short-term forecasts for HICP and HICPX
(annual percentage changes)

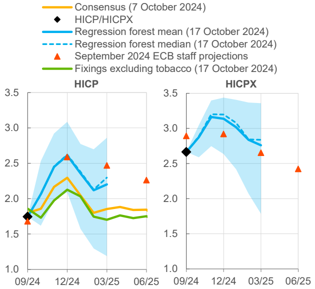

Sources: Eurostat, September 2024 ECB staff projections, Consensus Economics, Bloomberg and ECB calculations.
Notes: Quantile Regression Forest estimates are from Lenza, Moutachaker and Paredes (2023). The HICP fixings are observed market prices. The latest observations are for September 2024.

10

www.ecb.europa.eu©

Note: Extracted from speech of ECB Executive Board Member Philip Lane, 22 October 2024, www.ecb.europa.eu

Contributors to the forecast

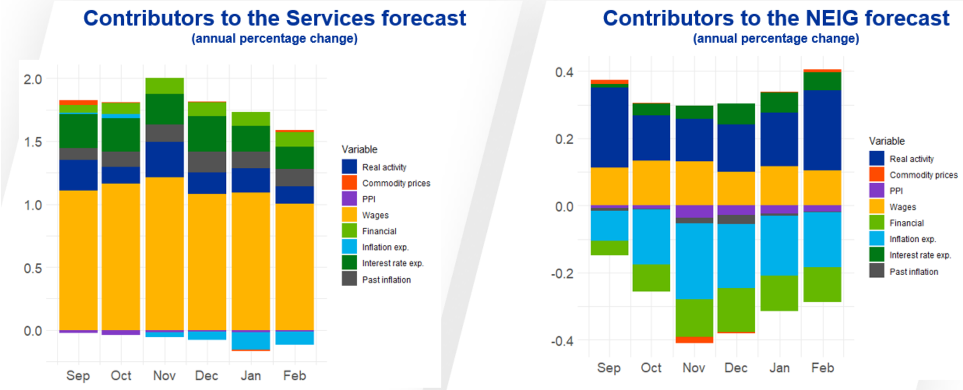

Notes: Cut-off dates: September NIPE - August 19th, 2024. MU - August 30th, 2024. QRF September 4th, 2024. The light-blue shaded area denotes the QRF 5-95 percentiles range. The contribution of each group of variables to the deviation from the QRF mean.

Dynamic decomposition

Non-energy industrial goods (NEIG) and services inflation

Dynamic QRF decomposition for the momentum of NEIG inflation

(percentage point deviations and contributions)

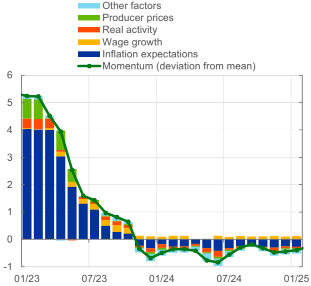

Dynamic QRF decomposition for the momentum of services inflation

(percentage point deviations and contributions)

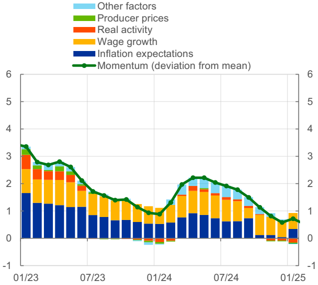

Sources: ECB calculations based on quantile regression forest (QRF) estimates from Lenza, Moutachaker and Paredes (2023) (cut-off for data is January 2025). The data sources are listed in Appendix A of their paper.
Notes: The inflation momentum is the annualised three-month-on-three-month percentage change in the seasonally adjusted price index. The latest observations are for January 2025.

8

www.ecb.europa.eu©

Note: Extracted from speech of ECB Executive Board Member Philip Lane, 12 March 2025, www.ecb.europa.eu

Identification of non-linearities

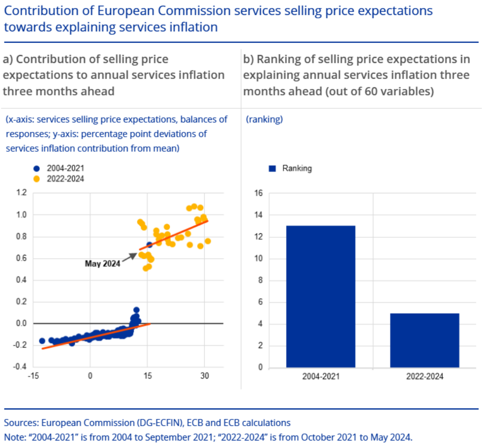

Note: Extracted from ECB Economic Bulletin, Issue 5/2024, www.ecb.europa.eu

THANK YOU!