Density forecasts of inflation: a quantile regression forest approach

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

Real-Time Data Analyisis, Methods, and Applications
BdE- Madrid
19 October 2023

Density Forecasting

- Adequate policy response requires to assess the temporary/persistent nature of past, current (and future, to a certain extent) “facts” for inflation dynamics
- The future is uncertain: economic policy is based on an analysis of the likelihood of different events (risk analysis)

$\Rightarrow$ **Accurate density forecasts are a fundamental input for monetary policy**

Inflation dynamics

- The Eurosystem analysis of inflation dynamics is heavily skewed toward linear models (Darracq-Paries et al. 2021)
- Yet, non-linearities are often argued to play an important role for inflation dynamics
  - Long and ongoing debate about steepening/flattening of the Phillips Curve (e.g. Del Negro et al. 2020; Eser et al. 2020; Costain et al. 2022)
  - Evidence that some "puzzling" inflation dynamics may be reconciled with theory by invoking non-linearities (Linde and Trabandt, 2019; Forbes et al. 2021)
- However, so far non-linearity has not dominated the landscape of modelling in support of monetary policy

⇒ How relevant it is for inflation density forecasting to account for non-linearity? And which non-linearity?

Aim of this paper

- Design and evaluate the accuracy of a new model for euro area density inflation forecasting
- No commitment to one type of non-linearity
- 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**

What we find - comparison with state-of-the-art linear models

- The quantile regression forest (QRF) is a good forecasting model, especially at short horizons and 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

⇒ Complementarity of the approaches. Non-linearity maybe more relevant in specific episodes and for core inflation.

What we find - 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.**

Related literature

Inflation forecasting

Large literature, see Faust and Wright (2013) and my forthcoming survey with M. Banbura and J. Paredes. To be singled out: Medeiros et al. (2021), Random Forest for point US inflation forecast

Non-linearity in inflation dynamics

See, for example Akerlof et al., 1996; Costain et al., 2022; Fahr and Smets, 2010; Benigno and Ricci, 2011; Linde and Trabandt, 2019; Del Negro et al. 2020; Forbes et al., 2021; Goulet-Coulombe et al., 2022

Ensemble methods for prediction

See, for example, Athey et al., 2019; Avramov, 2002; Bai and Ng, 2009; Cremers, 2002; Faust et al., 2013; Fernandez et al., 2001; Inoue and Kilian, 2008; Jin et al., 2014; Ng, 2013; Rapach and Strauss, 2010; Sala-I-Martin et al., 2004; Varian, 2014; Wager and Athey, 2018; Wright, 2009; Giannone et al., 2021; Medeiros et al., 2021; Clark et al. 2022a; Clark et al. 2022b

Outline of the rest of the talk

- Empirical strategy
- Comparison with state-of-the-art linear models
- Comparison with institutional (BMPE) judgemental forecasts

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 $\pi_t = m(\pi_{t-h} \ldots \pi_{t-h-p}; x_{t-h} \ldots x_{t-h-k}) + \varepsilon_t$

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

Main ingredients

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

Non-linearity: Regression trees

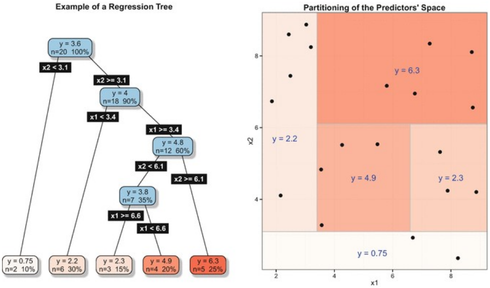

Non-linearity: Regression trees

- 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 forests

The idea of Random Forests - Breiman 2001

1. **Bootstrap** observations (and keep the "out-of bag" observations)
2. Grow **many** trees
3. In each tree, use only a (randomly chosen) **sub-set of predictors** at each node
4. **Combine** the predictions of the trees at the end

Does this make sense?

- Combination reduces variance of the forecasts
- Variance reduction maximized when the predictions are not correlated
  - Bootstrap to ensure "diversity" in the trees
  - The randomization step further de-correlates the trees
- Density forecasts: rather than taking averages of the target variable in the last nodes, compute sample quantiles $\Rightarrow$ **Quantile Regression Forest** (Meinshausen, 2006)
- One issue with regression trees/random forests: they do not extrapolate. How long can it take to adapt to unprecedented developments? Is "conservativeness" good or bad?

Hyperparameters

The specification choices for Random Forests are just a few. In general, we take the default choices in the literature (but further tuning is possible)

- **Depth of trees**
  - Control overfitting
  - We experimented with several configurations (varying number of splits, number of observations in the last node) - no impact on results
- **Number of variables randomly drawn for each split**
  - We use the default value for regression trees (a third of the variables), but also experimented with other values, small differences in results
- **Number of trees**
  - 500 trees, selected by assessing "stabilization" of out-of-bag error in the first training sample

Data 1: Targets

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

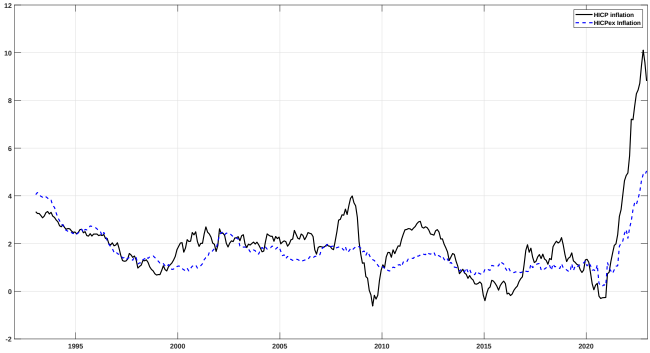

Note: Headline inflation: black solid line; Core inflation: blue dashed line.

Data 2: Predictors

- We consider 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)

Benchmark models

- Combination of 500 5-variate B-VARs (randomly drawn regressors) - VARCOMB
  - Ensemble of linear models to isolate as much as possible the linear/non-linear dimension
  - One difference in how the forecasts are produced (iterative/direct method)
- Comparison with institutional forecasts (SPF and BMPE)
  - Due to the comparison with the BMPE (which are not seasonally adjusted), we will always compute forecast accuracy statistics and produce charts for the year-on-year rates
  - SPF and BMPE are "real-time" judgemental forecasts, while we use revised data
  - SPF and BMPE quarterly frequency; we adapt to this frequency for QRF and VARCOMB, also aligning as much as possible the data availability

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

CRPS - full sample

Horizon	QRF	BVAR
Panel a: Headline Inflation		
h=3	0.29	0.28
h=6	0.50	0.49
h=9	0.74	0.67
h=12	0.93	0.88
Panel b: Core Inflation		
h=3	0.14	0.14
h=6	0.23	0.24
h=9	0.31	0.32
h=12	0.37	0.39

Note: CRPS for QRF (second column) and VARCOMB (third column)

CRPS - three years rolling window, h=6

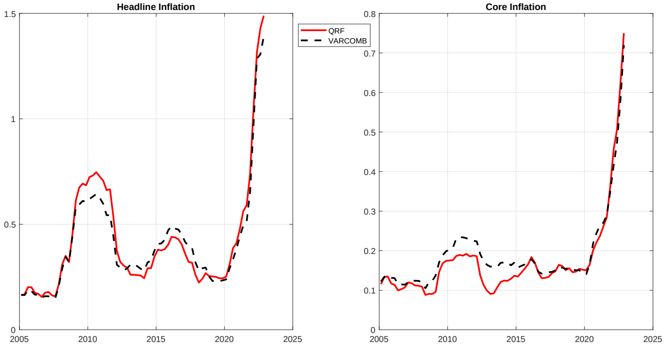

Note: Red solid line: QRF; Black dashed line: VARCOMB. The value on the vertical axis at each point refers to the average CRPS over the current quanrter and the previous eleven quarters.

Shapley Values - Top Contributors

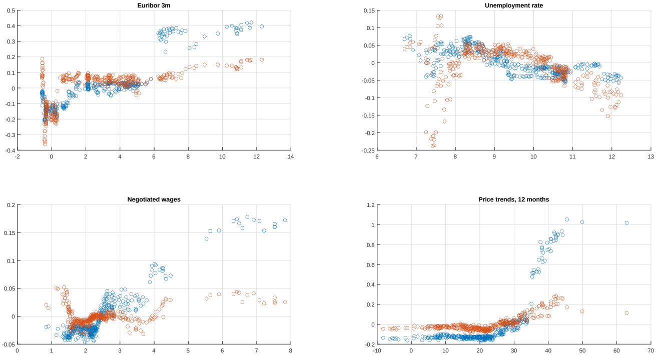

Note: Vertical axis: in-sample Shapley values for the variable indicated in the title for headline inflation (red) and core inflation (blue). Horizontal axis: value of the variable indicated in the title

Functional forms

- Most of the top contributors seem to be in a linear relationship with inflation
- Euribor: systematic monetary policy?
- Unemployment and wages: linear Phillips Curve-type correlations?
- Inflation expectations: non-linear relationship! Steeper change in slope for core inflation

Tentative conclusions on non-linearity in inflation dynamics

- Summing up the outcomes of the forecasting comparison and the analysis of the Shapley values, mild non-linearity, more evident for core
- Difference between core and headline: energy and food components
- These components, which are the most volatile of headline inflation, are driven mostly by commodity prices
- Direct effects of commodity prices have largely a linear impact on inflation

Eurosystem point forecast

- QRF median forecasts are competitive with (B)MPE up to two quarters ahead. Then the (B)MPE is more accurate
- Strong collinearity! Does it say something on judgement?

Headline Inflation, density forecasts of QRF and BMPE point forecast, h=6

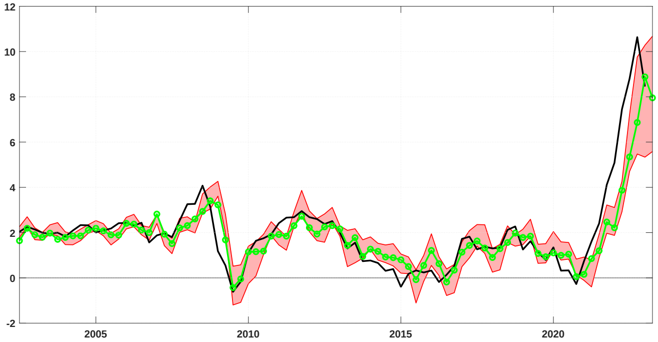

Note: Black solid line: year-on-year growth rate of HICP; Red area: 16th-84th quantiles QRF density forecasts; Green line with circles: BMPE projections.

Distance from non-linearity: gaps of BMPE and QRF versus VARCOMB

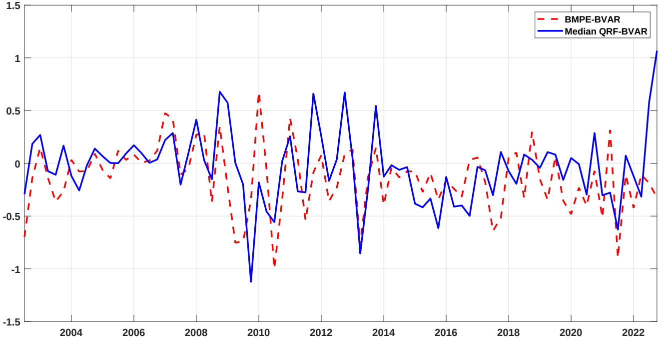

Note: Solid blue line: six months ahead (median) QRF forecast of headline inflation minus corresponding VARCOMB forecast; Dashed red line: six months ahead BMPE forecast of headline inflation minus corresponding VARCOMB forecast.

Conclusion

- The quantile regression forest (QRF) is a welcome addition to the Eurosystem toolbox to forecast inflation
- Complement rather than substitute the currently available tools
- Non-linearity in inflation dynamics: perhaps, mostly for core inflation
- QRF quite similar to BMPE/SPF, both in dynamics and accuracy - judgement partly adds "non-linearity"?

Example of policy use

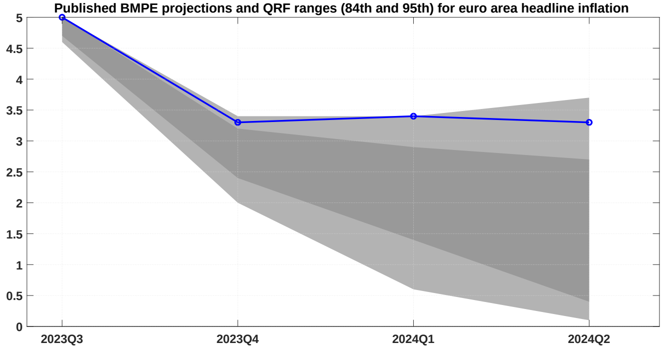

Note: Solid blue line with circles: BMPE projections, headline inflation, year-on-year rates; Grey areas: 68% coverage (dark) and 90% coverage (light).

BACKGROUND SLIDES

Calibration Headline Inflation - Rossi and Sekhposyan (2019)'s test of uniformity of PITs

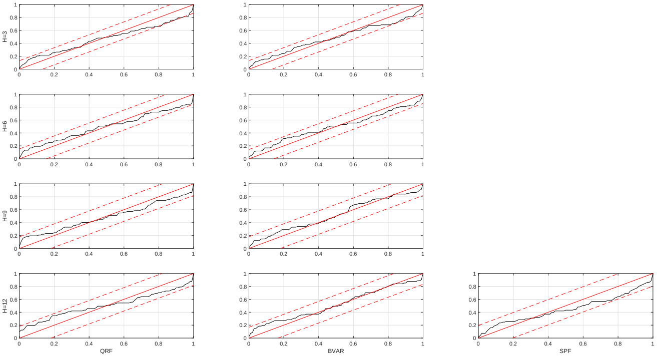

Note: Red lines: 1% critical values of the Kolmogorov-Smirnov test of PIT uniformity (dashed) and 45% degree line; Black line: Cumulative distribution function (CDF) of the PITs.

Calibration Core Inflation - Rossi and Sekhposyan (2019)'s test of uniformity of PITs

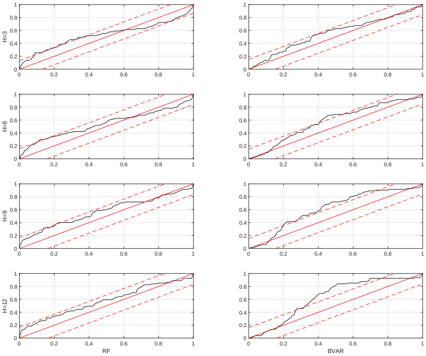

Note: Red lines: 1% critical values, Kolmogorov-Smirnov test of PIT uniformity (dashed) and 45% degree line; Black line: Cumulative distribution function (CDF) of the PITs.

SPF - density forecast of QRF and SPF for current year

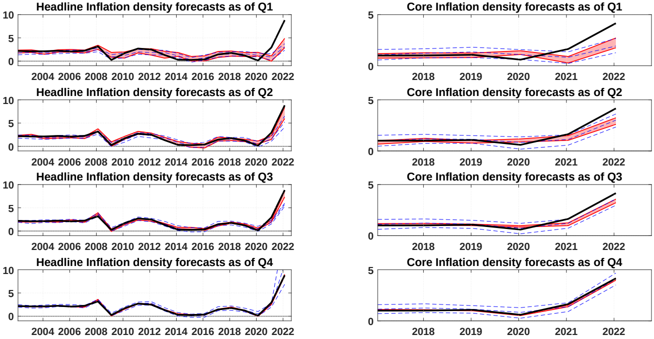

Note: Red Area: 16th to 84th quantile of the QRF, current year for headline inflation (left panels) and core inflation (right panels); Dashed Lines: 16th to 84th quantile of the SPF, current year for headline inflation (left panels) and core inflation (right panels).