Technical Background: "Outlook-at-Risk"

Federal Reserve Bank of New York, Research: Macrofinance Department

May 15, 2023

This document provides a short technical background on "Outlook-at-Risk". Specifically, we describe the technical details, including intermediate steps, of how the conditional distributions are estimated. For further details please see the references in the "Related reading" section[1] along with the answers to the frequently asked questions.[2]

1 Background on Methodology

Here, we describe the procedure that we follow to estimate the monthly conditional distributions of real GDP growth, the unemployment rate, and CPI inflation. To model the distribution, we use the two-step quantile regression methodology proposed in Adrian, Boyarchenko, and Giannone (2019). Denote by $y_{q(t)+h}$ the h quarter average – computed from quarter $q(t)+1$ until $q(t)+h$ – or the h quarter ahead value of the target variable of interest. For real GDP growth and CPI inflation, the h quarter average represents the annualized average growth rate over that period; for unemployment, the h quarter value is the unemployment rate h quarters ahead. Now, denote by $\hat{y}^{BC}_{q(t)+h|t}$ the consensus (mean) BCEI forecast for $y_{q(t)+h}$ in month t, and the associated forecast error $e^{BC}_{q(t)+h|t} \equiv y_{q(t)+h} - \hat{y}^{BC}_{q(t)+h|t}$.

1.1 Quantile Regressions

In the first step, we estimate quantile regressions (as in Koenker and Hallock (2001)) of the consensus BCEI forecast errors, $e^{BC}_{q(t)+h|t}$, on conditioning information available in month t, x_t (which includes a constant). Given a percentile $\tau \in (0,1)$, our objective is to estimate the τ-quantile, $Q_{e^{BC}_{q(t)+h|t}|x_t}$, of the conditional distribution of the errors,

[1] https://www.newyorkfed.org/research/policy/outlook-at-risk#root:overview
[2] https://www.newyorkfed.org/research/policy/outlook-at-risk#root:faq

$F_{e^{BC}_{q(t)+h|t}|x_t}$, where

$$Q_{e^{BC}_{q(t)+h|t}|x_t} \equiv \inf\{q \in \mathbb{R} \mid F_{e^{BC}_{q(t)+h|t}|x_t}(q|x_t) \geq \tau\}$$

A quantile regression returns the coefficient β_τ that minimizes the sum of quantile-weight absolute residuals, or

$$\hat{\beta}_\tau = \arg\min_{\beta_\tau \in \mathbb{R}^k} \sum_{t=1}^{T-h} (\tau \cdot \mathbb{1}_{\{e^{BC}_{q(t)+h|t} > x_t'\beta_\tau\}} |e^{BC}_{q(t)+h|t} - x_t'\beta_\tau| + (1-\tau) \cdot \mathbb{1}_{\{e^{BC}_{q(t)+h|t} < x_t'\beta_\tau\}} |e^{BC}_{q(t)+h|t} - x_t'\beta_\tau|)$$

Given an estimated coefficient, the predicted value from the quantile regression is

$$\hat{Q}_{e^{BC}_{q(t)+h|t}|x_t}(\tau|x_t) = x_t'\hat{\beta}_\tau$$

which is clearly linear in the conditioning information.

Under the assumption that the forecast errors and the forecasts are independent conditional on x_t, we can obtain the implied quantiles of $y_{q(t)+h|t}$ by adding back the consensus BCEI forecast, i.e.

$$\hat{Q}_{y_{q(t)+h|t}|x_t}(\tau|x_t) = \hat{Q}_{e^{BC}_{q(t)+h|t}|x_t}(\tau|x_t) + \hat{y}^{BC}_{q(t)+h|t}$$

1.2 Fitting Parametric Distributions

The quantile regressions provide estimates for a finite set of conditional quantiles. In order to construct full conditional distributions from these estimates, we follow Adams, Adrian, Boyarchenko, and Giannone (2021) and fit a smooth quantile function from a flexible class of probability distributions to the estimands, specifically distributions from the four-parameter skew t-family of Azzalini and Capitanio (2003).

This family is quite general and allows us to capture fat tails and skewness (though not multimodality). The density function for this family is given by

$$f(y;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma} t\left(\frac{y-\mu}{\sigma};\nu\right) T\left(\alpha(\frac{y-\mu}{\sigma})\sqrt{\frac{\nu+1}{\nu+\frac{y-\mu}{\sigma}}};\nu+1\right)$$

where $t(\cdot;n)$ and $T(\cdot;n)$ denote the probability density function and cumulative distribution function, respectively, of the standard student's t-distribution with n degrees of freedom.

In the second step, we estimate the parameters of this distribution given the estimated quantiles $\hat{Q}_{y_{q(t)+h|t}|x_t}(\tau|x_t)$ for $\tau = \{.10, .25, .75, .90\}$. We do this by minimizing the squared difference between the estimated quantiles and the quantiles of the parametric distribution, i.e.

$$\{\hat{\mu}_{h,t}, \hat{\sigma}_{h,t}, \hat{\alpha}_{h,t}, \hat{\nu}_{h,t}\} = \arg\min_{\mu,\sigma,\alpha,\nu} \sum_{\tau=\{.1,.25,.75,.9\}} (\hat{Q}_{y_{q(t)+h|t}|x_t}(\tau|x_t) - F^{-1}(\tau;\mu,\sigma,\alpha,\nu))^2$$

where $F^{-1}(\tau;\mu,\sigma,\alpha,\nu)$ is the quantile function of the skew t-distribution.

1.3 Data and Implementation

To implement the aforementioned procedure we utilize available survey forecasts along with corresponding realized data for our variables of interest. Specifically, we use monthly near-term consensus forecasts (one quarter to four quarters ahead) for our economic variables of interest – real GDP growth[3], the unemployment rate, and CPI inflation – obtained from the Blue Chip Economic Indicators Survey (BCEI). Survey forecasts are available for real GNP/GDP growth starting in August 1978, for the unemployment rate starting in February 1980, and for CPI inflation starting in March

[3]Prior to 1992 the survey queried respondents on real GNP growth rather than real GDP growth.

1980.

For the realized variables, we use the third release of real output growth available from the Federal Reserve Bank of Philadelphia Real-Time Data[4] corresponding to real gross national product (GNP) growth prior to 1992 and real GDP growth thereafter. We also use the civilian unemployment rate (U-3 unemployment rate) and the consumer price index for all urban consumers, both available from the Bureau of Labor Statistics. For real GDP growth and CPI inflation we construct benchmark forecasts and corresponding forecast errors based on the average of the following four quarters' quarterly annualized growth rates. For the unemployment rate we use the quarterly average of the unemployment rate in four quarters' time. Second, we use a measure of overall U.S. financial conditions, the Composite Indicator of Systemic Stress (CISS), which is produced by the European Central Bank.[5]

To construct initial estimates of the parameters governing the conditional distributions we use the available sample ending in December 1988. Thereafter, we update the estimated conditional distributions each month and re-estimate the parameters every 3 months using an expanding window of observations. This approach ensures that, starting in January 1989, we only utilize information that would have been available at the time and avoid any "look ahead" bias.

[4] `https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/first-second-third`

[5] `https://sdw.ecb.europa.eu/quickview.do;jsessionid=45EC9C197363A48C764F2FEA7BDE2E7A?SERIES_KEY=290.CISS.D.US.Z0Z.4F.EC.SS_CI.IDX`

References

Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021): "Forecasting Macroeconomic Risks," *International Journal of Forecasting*, 37, 1173–1191.

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89.

Azzalini, A. and A. Capitanio (2003): "Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t-Distribution," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 367–389.

Koenker, R. and K. F. Hallock (2001): "Quantile Regression," *Journal of Economic Perspectives*, 15, 143–156.