Density Forecast Transformations

14th BdE-CEMFI Research Workshop

Matteo Mogliani [1] **Florens Odendahl**[2]

[1]Banque de France
[2]Banco de España

April 25, 2024

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Banque de France, Banco de España, or the Eurosystem.

Motivation

- In the Eurosystem expert group Macro@Risk, we worked on a toolbox for predicting risks for GDP growth and inflation
- Main focus: density forecasts of quarter-on-quarter GDP growth and year-on-year inflation
- Forecast strategy: "direct forecasts" to avoid having to predict the predictors
- Direct forecasts via quantile regressions for quarterly data:

$$Y_{t+h}(q) = \tau_h(q) + \beta_h(q) Y_t + \gamma_h(q) X_t + u_{t+h}(q), \tag{1}$$

  - Grid of quantiles $q = 0.01, ..., 0.99$ to obtain a density forecast for $t + h$
  - Advantage of "direct forecasts": additional predictors X_t do not need to be predicted
  - Each forecast horizon h is a separate regression
  - Separate regression for each predictor X_t: real variables, financial variables, energy price variables etc.
  - Final density forecasts typically based on a forecast combination of best performing models

Motivation

Figure: Quarter-on-Quarter growth forecasts

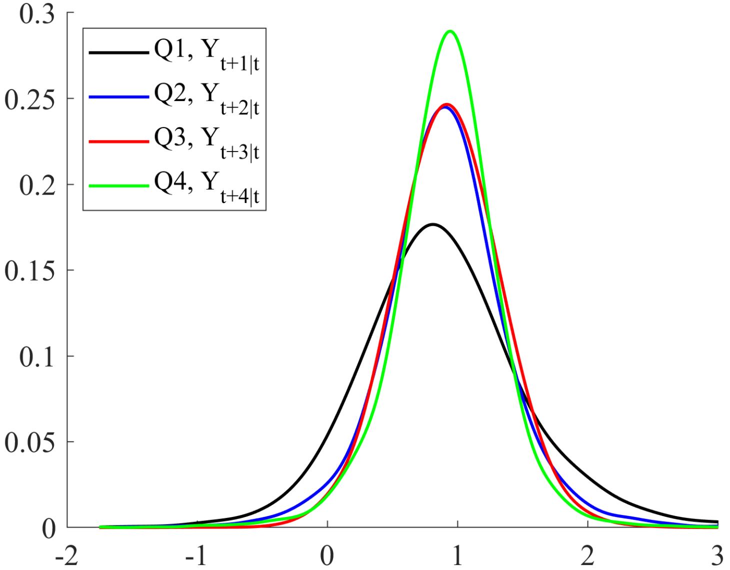

Motivation

- Then, a "new requirement was added" to the toolbox:
  - The toolbox should produce quarter-on-quarter AND annual-average density forecasts
  - The predictions should be consistent across qoq and annual-average
- How do we construct density forecasts for linear transformations of the quarter-on-quarters growth rates?
  - Density forecast of quarter-on-quarter growth rate $Y_{t+h|t}$ comes from quantile regressions...
  - ... but annual-average growth rate $Z_{t+\mathbf{4}|t} = \frac{1}{4}(Y_{t+\mathbf{1}|t} + ... + Y_{t+\mathbf{4}|t}) \sim$ **??**
  - Specific issue here due to the direct forecasting scheme: we do not know the cross-horizon dependence between $Y_{t+i|t}$ and $Y_{t+j|t}$, $i \neq j$, $i, j = 1, ..., h$
- This led to the following questions:
  - Specific question in the expert group: how do we compute annual-average density forecasts out of quarter-on-quarter density forecasts?
  - Translates to general question: how do we transform existing density forecasts coming from a direct forecasting scheme into new predictive objects that are functions of several horizons?
  - No compelling answer yet in the literature

Motivation - General set-up

The economist has a forecasting model that produces h-step-ahead **direct density forecasts**.

- (Conditional) Phillips curves for inflation
- Macro-at-Risk models such as quantile regressions
- Estimates are carried out at a given frequency/transformation of the data, and density forecast are produced accordingly (e.g. quarterly frequency/QoQ growth rates).

- The economist is then asked to transform the density forecasts (typical in institutions, where aggregated/transformed figures are usually provided.).

  $\Rightarrow$ *periodic* transformation: the frequency does not change (e.g. QoQ to YoY).
  $\Rightarrow$ *frequency* transformation: the frequency does change (e.g. QoQ to AA).

- **But** *direct* forecasting schemes imply that the individual predictions do not embed information on cross-horizon dependence...
- ...and **this dependence is needed** if the forecaster has to construct predictive objects that are functions of several horizons, such as YoY or AA growth rates.

A simple analytical illustration

- A simple mean-zero AR(1)

$$Y_{t+1} = \rho Y_t + \varepsilon_{t+1}$$

is used to produce *h*-**step-ahead *direct* density forecasts**:

$$p(Y_{t+h|t}; \rho, \sigma_\varepsilon) \sim \mathrm{N}(\rho_h Y_t, \sigma^2_{\varepsilon,h})$$

- Note that forecast errors $e_{t+h|t}$ have auto-covariance and auto-correlation functions:

$$\begin{aligned} Cov(e_{t+h|t}, e_{t+h-k|t}) &= \sigma_{\varepsilon,hk} \\ Corr(e_{t+h|t}, e_{t+h-k|t}) &= \rho_{\varepsilon,hk} \end{aligned}$$

for $h > k > 0$.

- .. because $e_{t+1|t} = \epsilon_{t+1}$, $e_{t+2|t} = \rho\epsilon_{t+1} + \epsilon_{t+2}$,...

A simple analytical illustration

- Consider now a linear transformation of the forecast sequence $\{Y_{t+j|t}\}_{j=1}^{h}$, such as the sum over h horizons:

$$Z_{t+h|t} = Y_{t+1|t} + \cdots + Y_{t+h|t}$$

- The **"ideal" forecaster** has predictive distribution $p(Z_{t+h|t}; \rho, \sigma_\varepsilon)$:

$$p(Z_{t+h|t}; \rho, \sigma_\varepsilon) \sim \mathcal{N}\left(\sum_{j=0}^{h-1} \rho^{j+1} Y_t, \sigma_{Z,h}^2 + 2\sigma_{Z,ij}\right)$$

  - $2\sigma_{Z,ij}$ captures cross-horizon dependence

- The **"simple" forecaster**, who ignores the cross-horizon dependence of the forecasts (*i.e.* the correlation structure of the forecast errors), has predictive distribution $f(Z_{t+h|t}; \rho, \sigma_\varepsilon)$:

$$f\left(Z_{t+h|t}; \rho, \sigma_\varepsilon\right) \sim \mathcal{N}\left(\sum_{j=0}^{h-1} \rho^{j+1} Y_t, \sigma_{Z,h}^2\right).$$

- Same point forecast but **different variance**. What are the implications in terms of predictive accuracy?

A simple analytical illustration

Figure: Transformed (sum at $h = 12$) AR(1) with $\rho = 0.4$ and $\sigma^2 = 0.1$

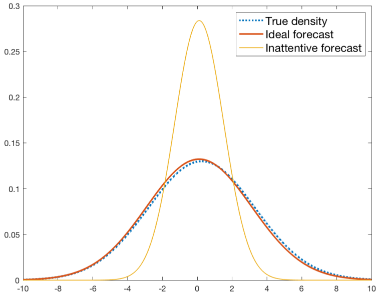

$\Rightarrow$ Simple forecaster underestimates the dispersion.

A simple analytical illustration

Figure: Scores for transformed density forecasts: cumulative sum up to $h = 12$
ideal *vs* simple

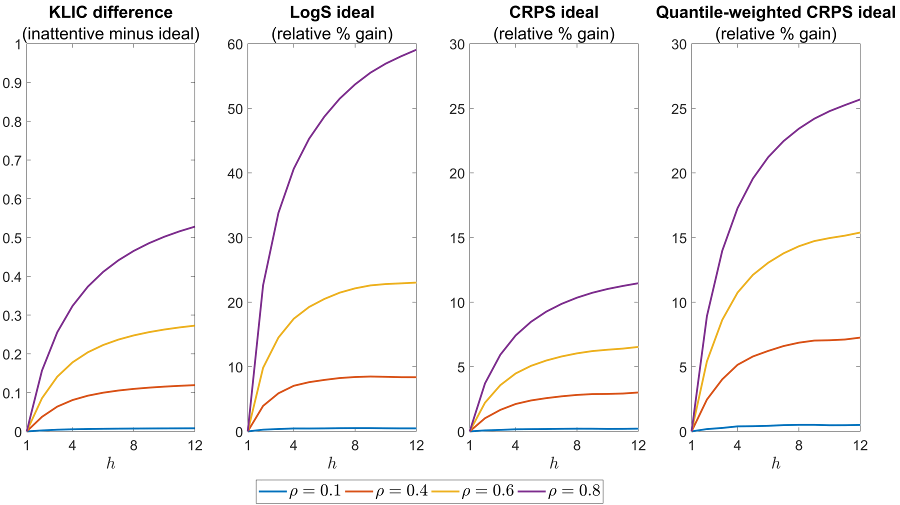

Contribution

- We propose to use **copulas** (Sklar, 1959) to combine the individual direct h-step-ahead predictive distributions into a joint predictive distribution.
- The **benefit**
  1. The joint predictive distribution takes the cross-horizon dependence into account.
  2. Allows the researcher to compute predictive objects that are functions of several horizons.
  3. Implementation of the approach is simple.
- The **cost**: need a pseudo-out-of-sample of sufficient size to compute reliable estimates of the cross-horizon dependence (PITs' correlations).

Overview of the main results

- **Monte Carlo simulations** show that our approach improves the predictive distributions relative to other approximation methods.
  - good approximations to the true underlying target-frequency density forecasts for different DGPs
  - robust to misspecified forecasting models and fairly small training samples

- In **three empirical examples**, we show that the proposed copula-approach leads to improved density forecasts in the target frequency
  - quarterly forecasts of FRED MD variables from monthly *direct* forecasts.
  - annual forecasts of US CPI inflation from year-on-year *direct* forecasts.
  - annual forecasts of US GDP growth from quarter-on-quarter *direct* forecasts

Methodology

Statistical framework

- Suppose the forecaster has a set of *direct* h-step-ahead predictive densities for T forecast origins, denoted by $\{\{g_{t,h}\}_{h=1}^{H}\}_{t=1}^{T}$ and with predictive CDF $\{\{G_{t,h}\}_{h=1}^{H}\}_{t=1}^{T}$, for outcome variables $Y_{t+h}, h = 1, .., H$
- Let then $Q_T(y_{T+1}, ..., y_{T+h}|R)$ denote the joint predictive CDF of $Y_{T+1}, ..., Y_{T+h}$ for forecast origin T, conditional on the correlation matrix R and constructed using C_{Ga}.
- A **copula is a multivariate CDF characterizing the dependence structure between random variables** $\Rightarrow$ allows to combine univariate marginals and a copula to obtain a valid multivariate distribution.
- Hence, $Q_T(y_{T+1}, ..., y_{T+H}|R) = C_{\mathsf{Ga}}(G_{T,1}(y_{T+1}), ..., G_{T,H}(y_{T+H})|R)$.
- Then, the forecaster can obtain an estimate of $Q_T(y_{T+1}, ..., y_{T+H}|R)$ using an algorithm drawing from the joint predictive distribution.

Estimation algorithm

Algorithm 1: Joint Predictive Distribution

1. Compute the realized PITs, $\{\{\text{PIT}_{t,h}\}_{h=1}^{H}\}_{t=1}^{T-H}$, of the predictive CDFs $\{\{G_{t,h}\}_{h=1}^{H}\}_{t=1}^{T-H}$.
2. Compute the rank correlations of $\text{PIT}_{t,h}$ across the different h to get an estimate of $\widehat{R}$.
3. Use $\widehat{R}$ in combination with C_{Ga} to obtain the joint distribution $\widehat{Q}_T(y_{T+1}, .., y_{T+H}|\widehat{R})$.

Figure: Estimation timeline

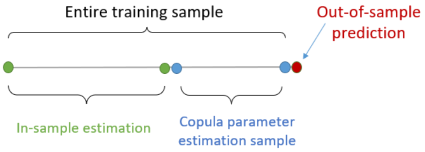

Monte Carlos

Monte Carlo simulations - design

We simulate QoQ growth rates (quarterly frequency) using a VAR(1).

$$\begin{bmatrix} y_t \\ x_t \end{bmatrix} = \begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} + \begin{bmatrix} \theta_1 & \theta_2 \\ 0 & \gamma \end{bmatrix} \begin{bmatrix} y_{t-1} \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}$$

- $\{\varepsilon_{j,t}\}_{t=1}^{T}$ two uncorrelated sequences of iid shocks.
- $\varepsilon_{2,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \sigma^2_{\varepsilon_2}\right)$, with $\sigma_{\varepsilon_2} = 0.3$.
- $\varepsilon_{1,t}$ may follow 3 different distributions:

  $$\varepsilon_{1,t} \sim \mathcal{N}(0, \sigma^2) \qquad \varepsilon_{1,t} \sim \text{Skew-}\mathcal{N}(\xi, \omega^2, \alpha) \qquad \varepsilon_{1,t} \sim \text{Skew-}t(\xi, \omega^2, \alpha, \nu)$$

  with $\alpha = -3$, $\nu = 8$, and ξ and ω^2 calibrated such that mean $= 0$ and variance $= \sigma^2 = 0.25$.
- $\tau_1 = 0.2$, $\tau_2 = 0$, and $\theta_2 = \gamma = 0.5$.
- $\theta_1 = \{0.1, 0.4, 0.7\} \Rightarrow$ account for different degrees of serial correlation, and hence cross-horizon dependence in the multi-step forecasts.

Monte Carlo simulations - design

- 2 types of forecasting models, each one a misspecified AR(1)-DMS horizon-specific regression:

$$y_{t+h} = \tau_h + \beta_h y_t + u_{t+h} \qquad \text{linear regression when } e_t \text{ Normal}$$

$$y_{t+h}(q) = \tau_h(q) + \beta_h(q) y_t + u_{t+h}(q) \qquad \text{quantile regression otherwise}$$

- We set :

$T_{is} = 200$ quarterly in-sample obs, held fixed in a rolling-window scheme

$T_{oos} = 50$ quarterly oos obs, for the computation of historical PITs

$T_{eval} = 200$ quarterly oos obs for the computation of (50) annual average

Monte Carlo simulations - design

- **AA and YoY density forecasts computed every four quarters for horizons one-, two-, and three-years-ahead.**[1]
- For every draw $s = 1, \dots, S$, transformed forecasts are based on well-known linear (approximate) formula. For instance:
  - 1-year-ahead AA from forecast origin Q4:

$$z^{(s)}_{AA,t+1|t} = \frac{1}{4}y_{t-2} + \frac{2}{4}y_{t-1} + \frac{3}{4}y_t + y^{(s)}_{t+1|t} + \frac{3}{4}y^{(s)}_{t+2|t} + \frac{2}{4}y^{(s)}_{t+3|t} + \frac{1}{4}y^{(s)}_{t+4|t}$$

  - 4-quarters-ahead YoY from forecast origin Q4:

$$z^{(s)}_{YoY,t+4|t} = y^{(s)}_{t+1|t} + y^{(s)}_{t+2|t} + y^{(s)}_{t+3|t} + y^{(s)}_{t+4|t}$$

- We compare the proposed copula approach to a "benchmark" approach which ignores cross-horizon dependence.
- We test for correct specification of the resulting transformed predictive distributions as well as for equal predictive performance **relative to the true predictive distribution**.

[1] Results for AA and YoY forecasts are quantitatively similar. We hence do not report the latter in this presentation.

Monte Carlo results: QoQ to AA transformation 1/4

Table: Relative performance for AA forecast

		Normal			Skew-Normal			Skew-t		
θ_1	Model \ h_A	1	2	3	1	2	3	1	2	3
		Log-score								
0.7	Copula − Bench.	0.36	1.11	1.22	0.32	1.01	1.14	0.40	1.10	1.22
0.4	Copula − Bench.	0.11	0.25	0.24	0.12	0.25	0.25	0.14	0.30	0.32
0.1	Copula − Bench.	0.00	-0.03	-0.03	0.00	-0.02	-0.02	0.02	0.01	0.01
		CRPS								
0.7	Copula/Bench.	0.97	0.94	0.94	0.97	0.94	0.94	0.97	0.94	0.94
0.4	Copula/Bench.	0.99	0.98	0.98	0.99	0.98	0.98	0.98	0.97	0.97
0.1	Copula/Bench.	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		Quantile-weighted CRPS								
0.7	Copula/Bench.	0.91	0.85	0.84	0.92	0.85	0.85	0.92	0.85	0.85
0.4	Copula/Bench.	0.96	0.93	0.93	0.96	0.94	0.94	0.96	0.93	0.93
0.1	Copula/Bench.	1.00	1.01	1.01	1.00	1.00	1.00	1.00	1.00	1.00

Note: for the log-score, numbers above zero indicate a superior performance of the copula approach. For the CRPS and QwCRPS, numbers smaller than one indicate a superior performance of the copula approach. The column θ_1 denotes the autoregressive parameter of Y_t in the DGP. Normal, Skew-Normal, and Skew-t indicate the distribution of the error terms in the DGP. The column label h_A denotes the annual-average horizon, *i.e.*, one-, two-, and three-years-ahead. Standard errors of the tests computed using a HAC estimator with bandwidth $= h_A - 1$.

Monte Carlo results: QoQ to AA transformation 2/4

Table: Tests of predictive performance: rejection frequency for AA forecast

		Normal			Skew-Normal			Skew-t		
θ_1	Model \ h_A	1	2	3	1	2	3	1	2	3
		Log-score								
0.7	Copula	0.07	0.07	0.07	0.07	0.06	0.07	0.11	0.04	0.07
	Benchmark	0.39	0.69	0.64	0.36	0.61	0.55	0.44	0.62	0.57
0.4	Copula	0.14	0.10	0.09	0.16	0.08	0.10	0.17	0.09	0.09
	Benchmark	0.27	0.39	0.39	0.27	0.29	0.35	0.32	0.39	0.38
0.1	Copula	0.29	0.18	0.17	0.34	0.18	0.20	0.28	0.16	0.16
	Benchmark	0.27	0.12	0.13	0.27	0.13	0.14	0.25	0.12	0.12

Note: rejection frequency of the null hypothesis of a Giacomini and White (2006) test of unconditional equal predictive ability. **The nominal size is 5%**. Standard errors of the tests computed using a HAC estimator with bandwidth $= h_A - 1$

Monte Carlo results: QoQ to AA transformation 3/4

Table: Tests of predictive performance: rejection frequency for AA forecast

		Normal			Skew-Normal			Skew-t		
θ_1	Model \ h_A	1	2	3	1	2	3	1	2	3
		Quantile weighted CRPS								
0.7	Copula	0.13	0.08	0.09	0.14	0.09	0.08	0.21	0.07	0.07
	Benchmark	0.43	0.73	0.68	0.38	0.66	0.66	0.45	0.70	0.68
0.4	Copula	0.21	0.11	0.09	0.20	0.10	0.09	0.21	0.10	0.08
	Benchmark	0.30	0.35	0.36	0.31	0.30	0.33	0.33	0.36	0.38
0.1	Copula	0.29	0.18	0.17	0.33	0.18	0.19	0.27	0.16	0.16
	Benchmark	0.26	0.12	0.11	0.26	0.14	0.13	0.25	0.13	0.12

Note: rejection frequency of the null hypothesis of a Giacomini and White (2006) test of unconditional equal predictive ability. **The nominal size is 5%**. Standard errors of the tests computed using a HAC estimator with bandwidth $= h_A - 1$

Monte Carlo results: QoQ to AA transformation 4/4

Table: Tests of predictive performance: rejection frequency for AA forecast

		Normal			Skew-Normal			Skew-t		
θ_1	Model \ h_A	1	2	3	1	2	3	1	2	3
		PIT								
0.7	Copula	0.06	0.10	0.12	0.06	0.08	0.10	0.08	0.08	0.12
	Benchmark	0.49	0.77	0.74	0.49	0.77	0.74	0.48	0.75	0.74
0.4	Copula	0.08	0.12	0.13	0.07	0.11	0.12	0.06	0.10	0.12
	Benchmark	0.24	0.40	0.38	0.25	0.37	0.37	0.25	0.37	0.36
0.1	Copula	0.07	0.14	0.15	0.09	0.16	0.18	0.07	0.11	0.12
	Benchmark	0.08	0.10	0.11	0.11	0.14	0.13	0.08	0.09	0.11

Note: rejection frequency of the null hypothesis of uniformity of PITs of the Rossi and Sekhposyan (2019) test for correct calibration of the density forecasts. **The nominal size is 5%**. The test is based on the Kolmogorov-Smirnov statistic. Standard errors of the tests were computed using a HAC estimator with bandwidth $= h_A - 1$

Empirical results

Empirical application 1: from monthly to quarterly frequency 1/2

- Large-scale forecasting exercise based on monthly data from FRED-MD (McCracken and McGillicuddy, 2019).
- 101 monthly series organized in 5 groups.
- Data sample: 1974:M1 to 2019:M12
  - historical PITs sample for copula parameter estimation: 2000:M1 to 2004:M12
  - evaluation sample: 2005:M1 to 2019:M12
- We consider random bivariate systems and estimate 1000 ARDL-DMS horizon-specific monthly frequency regressions:

$$y_{t+h_M} = \alpha + \sum_{j=0}^{p-1} \beta_j y_{t-j} + \sum_{j=0}^{p-1} \gamma_j x_{t-j} + \varepsilon_{t+h}$$

where $h_M = 1, \ldots, 12$ is the monthly horizon and

$$y_{t+h_M} = \begin{cases} Y_{t+h_M} - Y_t & \text{if } Y_t \sim I(1) \\ Y_{t+h_M} - Y_t - h_M \Delta Y_t & \text{if } Y_t \sim I(2) \end{cases}$$

- Forecasts $y_{t+h_M|t}$ are used to compute $Y_{t+h_M|t}$, which are used in turn to provide forecasts of $y_{t+h_Q|t}$ at quarterly frequency, with $h_Q = 1, \ldots, 4$ the quarterly horizon.

Empirical application 1: from monthly to quarterly frequency 2/2

Table: Relative performance of copula approach for quarter-on-quarter forecasts

	$p = 4$				$p =$ BIC			
Statistics \ h_Q	1	2	3	4	1	2	3	4
	CRPS							
Mean	0.99	0.96	0.96	0.95	0.99	0.96	0.96	0.96
2S Test	0.50	0.66	0.67	0.70	0.52	0.65	0.66	0.69
1S Test	0.44	0.66	0.67	0.70	0.42	0.63	0.64	0.67
	Log-score							
Mean	0.45	2.03	2.52	2.68	0.45	2.17	2.57	2.66
2S Test	0.36	0.55	0.56	0.59	0.35	0.56	0.55	0.58
1S Test	0.34	0.55	0.56	0.59	0.31	0.55	0.53	0.57
	Quantile-weighted CRPS							
Mean	0.97	0.91	0.89	0.88	0.98	0.92	0.91	0.90
2S Test	0.60	0.75	0.73	0.73	0.62	0.71	0.68	0.72
1S Test	0.53	0.74	0.72	0.73	0.52	0.68	0.66	0.70

Note: Row "Mean" denotes the copula approach relative to the benchmark approach, i.e. numbers smaller than one (larger then zero for Log-score) indicate a better performance of the copula approach. Values in the row "Test" shows the rejection frequency (at a 5% level) of the Giacomini and White (2006) test of unconditional equal predictive ability. Standard errors of the tests were computed using a HAC estimator with bandwidth $= h_Q - 1$

Empirical application 2: Inflation-at-Risk 1/2

- Loosely follows Korobilis (2017): QR-Lasso of yoy US CPI inflation on 22 predictors.

$$\min_{\beta \in \mathbb{R}^p} \sum_{t=1}^{T} \rho_\tau(y_{t+h} - x_t'\beta) + \frac{\lambda\sqrt{u(1-u)}}{n} \sum_{j=1}^{p} \widehat{\sigma}_j^2 |\beta_j|, \tag{2}$$

- y_{t+h} denotes the monthly year-on-year US inflation rate
- $\rho_\tau(z) = (\tau - 1\{z \leq 0\})z$ denotes the tick function
- λ is a hyperparameter that determines the degree of penalization, p is the number of predictors
- $h = 1, ..., 12$ is the forecast horizon and $\widehat{\sigma}_j^2 = \sum_{t=1}^{T} x_{i,t}^2$.

Forecast environment:

- Data from 1960 to 2022, rolling window estimation scheme
- Transform the yoy density forecasts into annual average density forecasts via copula and "simple" approach
- Empirical PITs computed from 1975 to 1984. Out-of-sample from 1985 to 2022.

Empirical application 2: Inflation-at-Risk 2/2

- According to the CRPS ratio the copula-based approach delivers a 10% better performance (statistically significant at the 1% level).

Figure: Inflation@Risk for 2001 and 2011

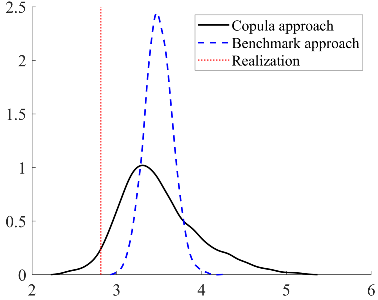

(a) 2001 (b) 2011

Empirical application 3: Growth-at-Risk 1/2

- Based on density forecasts Adrian et al. (2019): QR of qoq US GDP growth on NFCI.

Figure: G@Risk for QoQ growth for the year 2008

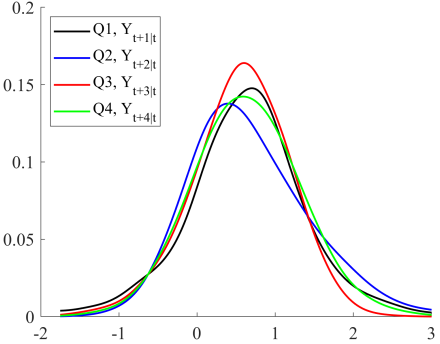

Empirical application 3: Growth-at-Risk 2/2

- We transform qoq density forecasts into annual average density forecasts.
- Empirical PITs computed from 1993 to 2001. Out-of-sample from 2002 onwards (only 12 observations).
- Forecast origin is Q4.

Figure: G@Risk for the year 2008 (forecast origin in 2007Q4)

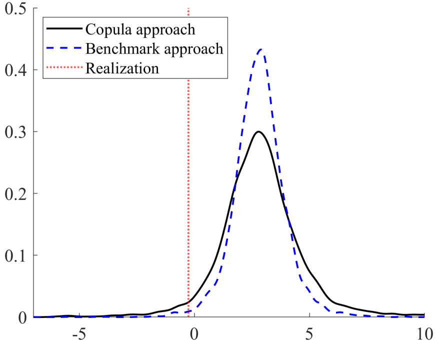

Direct annual-average regression

Direct annual-average regression 1/3

- Recurring comment: why not built annual-average density forecasts from annual-average regression.
- Let's simplify and assume there are just two quarters per year
- DGP $y_t = \rho y_{t-1} + e_t$, $e_t \sim N(0, \sigma^2)$ ...
- ... and that annual-average is defined as $Z_t = \frac{1}{2}(y_t + y_{t-1})$,$Z_{t-2} = \frac{1}{2}(y_{t-2} + y_{t-3})$,...
- MSFE of $Z_{t+2}|Z_t$ is larger than quarter-on-quarter-then-average regression by: $\frac{\rho^2\sigma^2}{4}$
- Why? Because the predictor Z_t is the average of y_t and y_{t-1}, i.e., the predictor in the annual-average regression uses the "out-dated" information y_{t-1}
- Has the flavour of a mixed-frequency regression where instead of using the higher frequency, the lower frequency is used
- The difference in the MSFE disappears with the forecasting horizon because of mean reversion of the process

- It would be better to regress Z_t on y_{t-2} instead of $\frac{1}{2}(y_{t-2} + y_{t-3})$
- The MSFE of this regression and the quarter-on-quarter-then-average is identical

Direct annual-average regression 2/3

Same Monte Carlo study as before with two additional competitor models:

Table: Relative performance for AA forecast

θ_1	Model \ h_A	Normal 1	Normal 2	Normal 3	Skew-Normal 1	Skew-Normal 2	Skew-Normal 3	Skew-t 1	Skew-t 2	Skew-t 3
		CRPS								
0.7	Copula/Benchmark	0.97	0.94	0.94	0.97	0.94	0.94	0.97	0.94	0.94
	Copula/AAonAA	0.62	0.96	0.99	0.60	0.94	0.97	0.59	0.94	0.97
	Copula/AAonQoQ	0.95	1.00	1.00	0.93	0.98	0.98	0.94	0.98	0.98
		Quantile weighted CRPS								
0.7	Copula/Benchmark	0.91	0.85	0.84	0.92	0.85	0.85	0.92	0.85	0.85
	Copula/AAonAA	0.62	0.96	1.00	0.60	0.93	0.96	0.59	0.93	0.96
	Copula/AAonQoQ	0.95	1.00	1.00	0.92	0.97	0.97	0.92	0.97	0.97

Note:

Direct annual-average regression 3/3

Table: Tests of predictive performance for AA forecast

		Normal			Skew-Normal			Skew-t		
θ_1	Model \ h_A	1	2	3	1	2	3	1	2	3
		CRPS								
0.7	Copula	0.24	0.11	0.10	0.18	0.10	0.08	0.26	0.12	0.11
	Benchmark	0.34	0.44	0.40	0.29	0.36	0.37	0.36	0.40	0.38
	AAonAA	0.99	0.24	0.11	1.00	0.26	0.13	1.00	0.27	0.16
	AAonQoQ	0.37	0.12	0.10	0.42	0.16	0.13	0.43	0.20	0.16
		PIT								
0.7	Copula	0.06	0.10	0.12	0.06	0.08	0.10	0.08	0.08	0.12
	Benchmark	0.49	0.77	0.74	0.49	0.77	0.74	0.48	0.75	0.74
	AAonAA	0.03	0.06	0.08	0.06	0.06	0.08	0.03	0.06	0.08
	AAonQoQ	0.05	0.05	0.08	0.05	0.07	0.07	0.06	0.05	0.07

Note: In Panel CRPS, numbers show the rejection frequency of a equal forecast ability test with the alternative model being the optimal forecast. Panel PIT shows the rejection frequency of a test for correct specification.

- PIT results for AAreg suggest that **AAreg is an inefficient forecasts** because it uses old information but **it is correctly specified**

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are functions of several horizons (e.g. annual average growth rates).
- The approach is simple to implement and requires only enough oos observations to compute the correlation of PITs at the necessary horizons.
- In a Monte Carlo exercise, we show that our methodology outperforms the "simple" approach whenever the serial correlation across different forecasting horizons is not extremely low.
- Three empirical applications provide evidence that the copula approach can provide better density forecasts than the "simple" approach.
- The approach is already implemented in the Macro@Risk toolbox (to be published soon) and as that part of the (B)MPE

Work in progress:

- Provide some guidance on how strong the cross-horizon correlation must be for the copula approach to be preferable.

Thank you for your attention

matteo.mogliani@banque-france.fr
florens.odendahl@bde.es

Terminology

- Direct forecasts:

$$y_{t+1} = \underbrace{\alpha_1 + \beta_1 y_t}_{\text{mean prediction}} + \underbrace{u_{t+1}}_{\text{"}uncertainty\text{"}}, \dots \quad , y_{t+h} = \underbrace{\alpha_h + \beta_h y_t}_{\text{mean prediction}} + \underbrace{u_{t+h}}_{\text{"}uncertainty\text{"}} \tag{3}$$

- Iterative forecasts:

$$y_{t+1} = \underbrace{\alpha + \beta y_t}_{\text{mean prediction}} + \underbrace{e_{t+1}}_{\text{"}uncertainty\text{"}} = \hat{y}_{t+1|t} + e_{t+1} \tag{4}$$

$$y_{t+2} = \alpha + \beta \hat{y}_{t+1|t} + e_{t+2}, \dots \quad , y_{t+h} = \alpha + \beta \hat{y}_{t+h|t+h-1} + e_{t+h} \tag{5}$$

- Advantage of "direct forecasts": additional predictors x_t do not need to be predicted:

$$y_{t+h} = \underbrace{\alpha_h + \beta_h y_t + \gamma_h x_t}_{\text{mean prediction}} + \underbrace{u_{t+h}}_{\text{"}uncertainty\text{"}} \tag{6}$$

- Iterative forecasts would require:

$$y_{t+h} = \underbrace{\alpha + \beta \hat{y}_{t+h|t+h-1} + \gamma \hat{x}_{t+h|t+h-1}}_{\text{mean prediction}} + \underbrace{e_{t+h}}_{\text{"}uncertainty\text{"}} \tag{7}$$

back