Density forecast frequency transformation via Copulas

Conference on Real-Time Data Analysis, Methods, and Applications
Banco de España,Madrid

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

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

October 20, 2023

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

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.
- Now suppose that the researcher is asked to aggregate the density forecasts stemming from the model at a lower frequency.

Motivation

- Suppose the researcher has a forecasting model that produces *h*-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.
- Now suppose that the researcher is asked to aggregate the density forecasts stemming from the model at a lower frequency.
  - $\Rightarrow$ a situation typical in institutions, where annual average GDP growth rates are provided.

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.
- Now suppose that the researcher is asked to aggregate the density forecasts stemming from the model at a lower frequency.
  - $\Rightarrow$ a situation typical in institutions, where annual average GDP growth rates are provided.
- The researcher could just draw multiples sequences of realizations from the h estimated density forecasts and compute the distribution of annual average growth rates.

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.
- Now suppose that the researcher is asked to aggregate the density forecasts stemming from the model at a lower frequency.
  - $\Rightarrow$ a situation typical in institutions, where annual average GDP growth rates are provided.
- The researcher could just draw multiples sequences of realizations from the h estimated density forecasts and compute the distribution of annual average growth rates.
- **But** *direct* forecasting schemes imply that the individual predictions do not embed information on cross-horizon dependence...

Motivation

- Suppose the researcher has a forecasting model that produces h-step-ahead **direct density forecasts**.
- Also suppose that the estimates are carried out at quarterly frequency, leading to quarterly density forecasts.
  - $\Rightarrow$ a popular framework, when modelling and forecasting qoq GDP growth rates.
- Now suppose that the researcher is asked to aggregate the density forecasts stemming from the model at a lower frequency.
  - $\Rightarrow$ a situation typical in institutions, where annual average GDP growth rates are provided.
- The researcher could just draw multiples sequences of realizations from the h estimated density forecasts and compute the distribution of annual average growth rates.
- **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 annual average growth rates.

Contribution

- We propose to use **copulas** (Sklar, 1959) to combine the individual direct h-step-ahead predictive distributions into a joint predictive distribution.

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**

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.

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 you to compute predictive objects that are functions of several horizons (e.g. annual average growth rates).

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 you to compute predictive objects that are functions of several horizons (e.g. annual average growth rates).
  3. Implementation of the approach is 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 you to compute predictive objects that are functions of several horizons (e.g. annual average growth rates).
  3. Implementation of the approach is simple.
- The **cost**: need a pseudo-out-of-sample to compute reliable estimates of PITs' correlations.

Overview of the main results

- **Monte Carlo simulations** show that our approach improves the predictive distributions relative to other approximation methods.

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

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

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

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
  - annual-average forecasts of US CPI inflation using year-on-year *direct* forecasts.

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
  - annual-average forecasts of US CPI inflation using year-on-year *direct* forecasts.
  - quarter-on-quarter forecasts of variables in FRED MD month-on-month *direct* forecasts.

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
    - annual-average forecasts of US CPI inflation using year-on-year *direct* forecasts.
    - quarter-on-quarter forecasts of variables in FRED MD month-on-month *direct* forecasts.
    - annual-average forecasts of US GDP growth using quarter-on-quarter *direct* forecasts

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 $Y_{t+h}, h = 1, .., H$

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.
- 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}.

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.
- 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}.
- Remember that a **copula is a multivariate CDF characterizing the dependence structure between random variables** $\Rightarrow$ any multivariate joint distribution can be expressed in terms of univariate marginals and a copula.

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.
- 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}.
- Remember that a **copula is a multivariate CDF characterizing the dependence structure between random variables** $\Rightarrow$ any multivariate joint distribution can be expressed in terms of univariate marginals and a copula.
- Hence, $Q_T(y_{T+1}, ..., y_{T+H}|R) = C_{\mathsf{Ga}}(G_{T,1}(y_{T+1}), ..., G_{T,H}(y_{T+H})|R)$.

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.
- 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}.
- Remember that a **copula is a multivariate CDF characterizing the dependence structure between random variables** $\Rightarrow$ any multivariate joint distribution can be expressed in terms of univariate marginals and a copula.
- Hence, $Q_T(y_{T+1}, ..., y_{T+H}|R) = C_{\text{Ga}}(G_{T,1}(y_{T+1}), ..., G_{T,H}(y_{T+H})|R)$.
- 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.

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 $Y_{t+h}, h = 1, .., H$
- The subscript h denotes the forecast horizon and the subscript t denotes the forecast origin.
- 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}.
- Remember that a **copula is a multivariate CDF characterizing the dependence structure between random variables** $\Rightarrow$ any multivariate joint distribution can be expressed in terms of univariate marginals and a copula.
- 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.
- In the paper we have an analytical example that illustrates how the copula approach captures the cross-horizon dependence.

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 $\{\{G_{t,h}\}_{h=1}^{H}\}_{t=1}^{T-H}$.

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 $\widehat{R}$.

Estimation algorithm

Algorithm 1: Joint Predictive Distribution

1. Compute the realized PITs, $\{\{\mathrm{PIT}_{t,h}\}_{h=1}^{H}\}_{t=1}^{T-H}$, of the predictive CDFs $\{\{G_{t,h}\}_{h=1}^{H}\}_{t=1}^{T-H}$.
2. Compute the rank correlations of $\mathrm{PIT}_{t,h}$ 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 $\widehat{Q}_T(y_{T+1}, .., y_{T+H}|\widehat{R})$.

Monte Carlo simulations - design of the experiments

- Simulate quarterly growth rates using an AR(1).

$$Y_t = \tau + \rho Y_{t-1} + e_t$$

where e_t may follow 3 different distributions:

$$e_t \sim \mathcal{N}(0, \sigma^2) \qquad e_t \sim \text{Skew-}\mathcal{N}(\mu, \xi, \delta) \qquad e_t \sim \text{Skew-}t(\mu, \xi, \delta, \nu)$$

with $\delta = -0.83$, $\nu = 8$, and μ and ξ calibrated such that mean $= 0$ and variance $= \sigma^2 = 0.25$.

Monte Carlo simulations - design of the experiments

- Simulate quarterly growth rates using an AR(1).

$$Y_t = \tau + \rho Y_{t-1} + e_t$$

where e_t may follow 3 different distributions:

$$e_t \sim \mathcal{N}(0, \sigma^2) \qquad e_t \sim \text{Skew-}\mathcal{N}(\mu, \xi, \delta) \qquad e_t \sim \text{Skew-}t(\mu, \xi, \delta, \nu)$$

with $\delta = -0.83$, $\nu = 8$, and μ and ξ calibrated such that mean $= 0$ and variance $= \sigma^2 = 0.25$.

- 2 types of forecasting models, each producing direct h-step-ahead forecasts:

$$Y_{t+h} = \tau_h + \gamma_h Y_t + u_{t+h} \qquad \text{linear regression when } e_t \text{ Normal}$$

$$Y_{t+h}(q) = \tau_h(q) + \gamma_h(q) Y_t + u_{t+h}(q) \qquad \text{quantile regression otherwise}$$

Monte Carlo simulations - design of the experiments

- Simulate quarterly growth rates using an AR(1).

$$Y_t = \tau + \rho Y_{t-1} + e_t$$

where e_t may follow 3 different distributions:

$$e_t \sim \mathcal{N}(0, \sigma^2) \qquad e_t \sim \text{Skew-}\mathcal{N}(\mu, \xi, \delta) \qquad e_t \sim \text{Skew-}t(\mu, \xi, \delta, \nu)$$

with $\delta = -0.83$, $\nu = 8$, and μ and ξ calibrated such that mean $= 0$ and variance $= \sigma^2 = 0.25$.

- 2 types of forecasting models, each producing direct h-step-ahead forecasts:

$$Y_{t+h} = \tau_h + \gamma_h Y_t + u_{t+h} \qquad \text{linear regression when } e_t \text{ Normal}$$

$$Y_{t+h}(q) = \tau_h(q) + \gamma_h(q) Y_t + u_{t+h}(q) \qquad \text{quantile regression otherwise}$$

- 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 of the experiments

- **We produce annual-average forecasts every four quarters for horizons of one, two, and three years ahead.**

Monte Carlo simulations - design of the experiments

- **We produce annual-average forecasts every four quarters for horizons of one, two, and three years ahead**.
- The benchmark is the true annual average distribution.

Monte Carlo simulations - design of the experiments

- **We produce annual-average forecasts every four quarters for horizons of one, two, and three years ahead**.
- The benchmark is the true annual average distribution.
- We compare the proposed copula approach to a “simple” approach ignoring serial dependence.

Monte Carlo simulations - design of the experiments

- **We produce annual-average forecasts every four quarters for horizons of one, two, and three years ahead**.
- The benchmark is the true annual average distribution.
- We compare the proposed copula approach to a “simple” approach ignoring serial dependence.
- Annual average distributions are based on the well-known aggregation formula of qoq growth rates:

$$AA_{Y2} = \frac{\bar{Y}_{Q2Y1}\bar{Y}_{Q3Y1}\bar{Y}_{Q4Y1}\bar{Y}_{Q1Y2}\left(1 + \bar{Y}_{Q2Y2} + \bar{Y}_{Q2Y2}\bar{Y}_{Q3Y2} + \bar{Y}_{Q2Y2}\bar{Y}_{Q3Y2}\bar{Y}_{Q4Y2}\right)}{1 + \bar{Y}_{Q2Y1}\left(1 + \bar{Y}_{Q3Y1}\bar{Y}_{Q3Y1}\bar{Y}_{Q4Y1}\right)} - 1$$

where $\bar{Y}_{QbYc}$ denotes the *gross* qoq growth rate in quarter b of year c

Monte Carlo simulations - design of the experiments

- **We produce annual-average forecasts every four quarters for horizons of one, two, and three years ahead**.
- The benchmark is the true annual average distribution.
- We compare the proposed copula approach to a "simple" approach ignoring serial dependence.
- Annual average distributions are based on the well-known aggregation formula of qoq growth rates:

$$AA_{Y2} = \frac{\bar{Y}_{Q2Y1}\bar{Y}_{Q3Y1}\bar{Y}_{Q4Y1}\bar{Y}_{Q1Y2}\left(1 + \bar{Y}_{Q2Y2} + \bar{Y}_{Q2Y2}\bar{Y}_{Q3Y2} + \bar{Y}_{Q2Y2}\bar{Y}_{Q3Y2}\bar{Y}_{Q4Y2}\right)}{1 + \bar{Y}_{Q2Y1}\left(1 + \bar{Y}_{Q3Y1}\bar{Y}_{Q3Y1}\bar{Y}_{Q4Y1}\right)} - 1$$

  where $\bar{Y}_{QbY_c}$ denotes the *gross* qoq growth rate in quarter b of year c

- We test for correct specification of the resulting annual-average predictive distributions as well as for equal predictive performance **relative to the true annual-average predictive distribution**.

Monte Carlo results: qoq to annual-average transformation 1/2

Table: Tests of predictive performance: rejection frequency for annual average forecast

		Normal			Skew Normal			Skew t		
ρ	Model	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
		Log-score								
0.8	Naïve	0.59	0.74	0.67	0.47	0.68	0.62	0.47	0.69	0.64
	Copula	0.09	0.07	0.07	0.06	0.05	0.06	0.04	0.07	0.08
0.5	Naïve	0.30	0.51	0.45	0.22	0.42	0.42	0.24	0.44	0.43
	Copula	0.05	0.06	0.09	0.04	0.10	0.09	0.03	0.07	0.09
0.1	Naïve	0.06	0.10	0.11	0.10	0.10	0.09	0.06	0.09	0.10
	Copula	0.12	0.21	0.22	0.14	0.19	0.18	0.09	0.21	0.20

Note: the 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 were computed using a HAC with a bandwidth $= h - 1$

Monte Carlo results: qoq to annual-average transformation 2/2

Table: Tests of correct specification: rejection frequency for annual average forecast

		Normal			Skew Normal			Skew t		
ρ	Model	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
		PIT								
0.8	Naïve	0.56	0.85	0.82	0.60	0.84	0.78	0.57	0.81	0.77
	Copula	0.08	0.10	0.13	0.07	0.07	0.11	0.06	0.09	0.12
0.5	Naïve	0.29	0.46	0.47	0.28	0.41	0.40	0.36	0.46	0.46
	Copula	0.11	0.15	0.18	0.05	0.09	0.10	0.08	0.11	0.13
0.1	Naïve	0.08	0.09	0.09	0.09	0.07	0.09	0.08	0.09	0.08
	Copula	0.13	0.20	0.20	0.08	0.10	0.12	0.08	0.14	0.12

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

Empirical application 1: from month-on-month to quarter-on-quarter

- Large-scale forecasting exercise based on monthly data from FRED-MD from 1959:M1 to 2019:M12.

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

Lag length	Statistics	Great moderation				Full sample			
		$h=1$	$h=2$	$h=3$	$h=4$	$h=1$	$h=2$	$h=3$	$h=4$
		CRPS							
AR(4)	Median	1.00	1.01	1.10	1.21	1.00	1.01	1.12	1.23
	Test	0.08	0.25	0.62	0.73	0.08	0.27	0.66	0.74

Note: Row "Median" shows the relative CRPS of the naïve approach relative to the copula approach, i.e., numbers larger than one indicate a worse performance of the naïve approach. Values in the row "Test" shows the percentage of times that a Giacomini and West (2006) test of unconditional equal predictive ability rejects the null hypothesis at a 5% level.

Empirical application 1: from month-on-month to quarter-on-quarter

- Large-scale forecasting exercise based on monthly data from FRED-MD from 1959:M1 to 2019:M12.
- We closely follow McCracken and McGillicuddy (2019) and consider random bivariate systems, $Z_t = (Y_t, X_t)'$.

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

Lag length	Statistics	Great moderation				Full sample			
		$h=1$	$h=2$	$h=3$	$h=4$	$h=1$	$h=2$	$h=3$	$h=4$
		CRPS							
AR(4)	Median	1.00	1.01	1.10	1.21	1.00	1.01	1.12	1.23
	Test	0.08	0.25	0.62	0.73	0.08	0.27	0.66	0.74

Note: Row "Median" shows the relative CRPS of the naïve approach relative to the copula approach, i.e., numbers larger than one indicate a worse performance of the naïve approach. Values in the row "Test" shows the percentage of times that a Giacomini and West (2006) test of unconditional equal predictive ability rejects the null hypothesis at a 5% level.

Empirical application 1: from month-on-month to quarter-on-quarter

- Large-scale forecasting exercise based on monthly data from FRED-MD from 1959:M1 to 2019:M12.
- We closely follow McCracken and McGillicuddy (2019) and consider random bivariate systems, $Z_t = (Y_t, X_t)'$.
- We first compute density forecasts for month-on-month values Y_{t+h}, with $h = 1, \ldots, 12$ months, and then we use these predictive densities to compute quarter-on-quarter density forecasts through our proposed copula approach.

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

Lag length	Statistics	Great moderation				Full sample			
		$h=1$	$h=2$	$h=3$	$h=4$	$h=1$	$h=2$	$h=3$	$h=4$
		CRPS							
AR(4)	Median	1.00	1.01	1.10	1.21	1.00	1.01	1.12	1.23
	Test	0.08	0.25	0.62	0.73	0.08	0.27	0.66	0.74

Note: Row "Median" shows the relative CRPS of the naïve approach relative to the copula approach, i.e., numbers larger than one indicate a worse performance of the naïve approach. Values in the row "Test" shows the percentage of times that a Giacomini and West (2006) test of unconditional equal predictive ability rejects the null hypothesis at a 5% level.

Empirical application 2: Inflation-at-Risk

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

Figure: IaR for 2001 and 2011

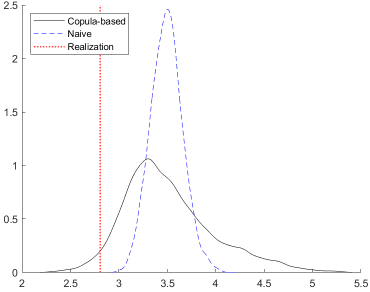

(a) 2001

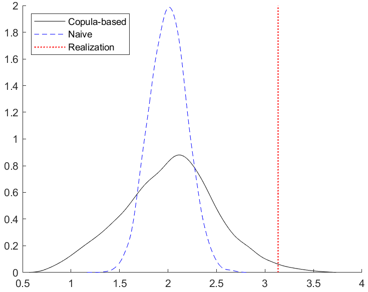

(b) 2011

Empirical application 2: Inflation-at-Risk

- Inspired by Korobilis (2017): QR-Lasso of yoy US CPI inflation on 22 predictors.
- We transform yoy density forecasts into annual average density forecasts.

Figure: IaR for 2001 and 2011

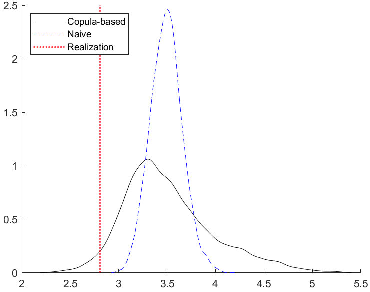

(a) 2001

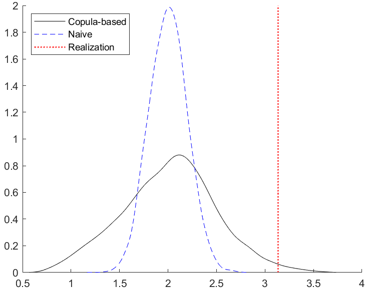

(b) 2011

Empirical application 2: Inflation-at-Risk

- Inspired by Korobilis (2017): QR-Lasso of yoy US CPI inflation on 22 predictors.
- We transform yoy density forecasts into annual average density forecasts.
- Empirical PITs computed from 1975 to 1984. Out-of-sample from 1985 onwards.

Figure: IaR for 2001 and 2011

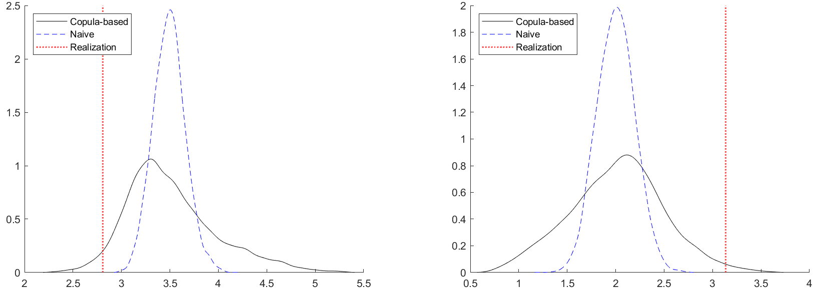

(a) 2001 (b) 2011

Empirical application 2: Inflation-at-Risk

- Inspired by Korobilis (2017): QR-Lasso of yoy US CPI inflation on 22 predictors.
- We transform yoy density forecasts into annual average density forecasts.
- Empirical PITs computed from 1975 to 1984. Out-of-sample from 1985 onwards.
- Forecast origin is YM12.

Figure: IaR for 2001 and 2011

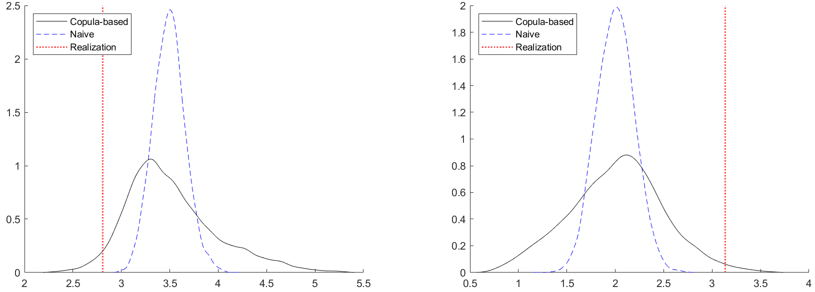

(a) 2001 (b) 2011

Empirical application 2: Inflation-at-Risk

- Inspired by Korobilis (2017): QR-Lasso of yoy US CPI inflation on 22 predictors.
- We transform yoy density forecasts into annual average density forecasts.
- Empirical PITs computed from 1975 to 1984. Out-of-sample from 1985 onwards.
- Forecast origin is YM12.
- According to the CRPS ratio the copula-based approach delivers a 10% better performance (statistically significant at the 1% level using Giacomini and White (2006)).

Figure: IaR for 2001 and 2011

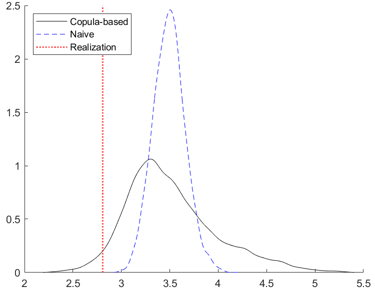

(a) 2001

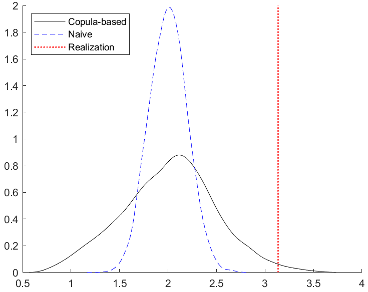

(b) 2011

Empirical application 3: Growth-at-Risk

- Inspired by Adrian et al. (2019): QR of qoq US GDP growth on NFCI.

Figure: GaR for the year 2008 (forecast origin in 2007Q4)

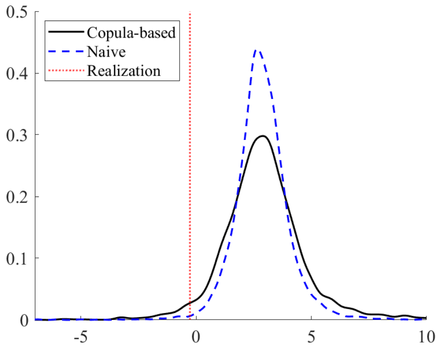

Empirical application 3: Growth-at-Risk

- Inspired by Adrian et al. (2019): QR of qoq US GDP growth on NFCI.
- We transform qoq density forecasts into annual average density forecasts.

Figure: GaR for the year 2008 (forecast origin in 2007Q4)

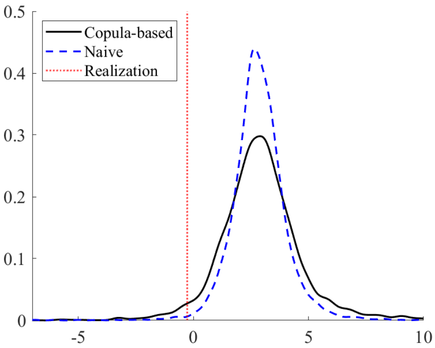

Empirical application 3: Growth-at-Risk

- Inspired by Adrian et al. (2019): QR of qoq US GDP growth on NFCI.
- We transform qoq density forecasts into annual average density forecasts.
- Empirical PITs computed from 1993 to 2001. Out-of-sample from 2002 onwards.

Figure: GaR for the year 2008 (forecast origin in 2007Q4)

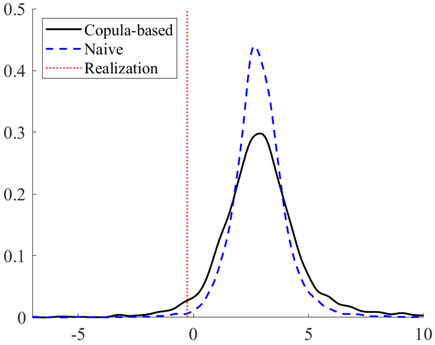

Empirical application 3: Growth-at-Risk

- Inspired by Adrian et al. (2019): QR of qoq US GDP growth on NFCI.
- We transform qoq density forecasts into annual average density forecasts.
- Empirical PITs computed from 1993 to 2001. Out-of-sample from 2002 onwards.
- Forecast origin is YQ4.

Figure: GaR for the year 2008 (forecast origin in 2007Q4)

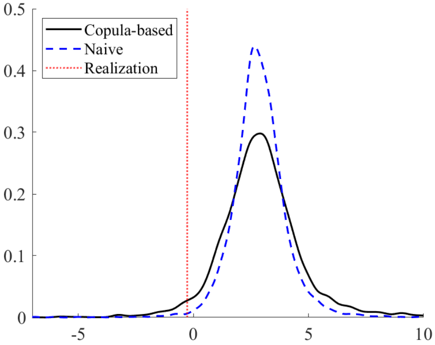

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function of several horizons (e.g. annual average growth rates).

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function 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.

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function 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.

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function 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.

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function 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.

Concluding remarks

- We provide a copula-based approach to combine direct forecasts to obtain new predictive objects that are function 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.

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