

WORKING paper



Density Forecast Transformations

Matteo Mogliani¹ and Florens Odendahl²

December 2025, WP #1027

ABSTRACT

The popular choice of using a *direct* forecasting scheme implies that the individual predictions ignore information on their cross-horizon dependence. However, this dependence is needed if the forecaster has to construct, based on *direct* density forecasts, predictive objects that are functions of several horizons (*e.g.* when constructing annual-average growth rates from quarter-on-quarter growth rates). To address this issue, we propose to use copulas to combine the individual *b*-step-ahead predictive distributions into a joint predictive distribution. Our method is particularly appealing to practitioners for whom changing the *direct* forecasting specification is too costly. In a Monte Carlo study, we demonstrate that our approach leads to a better approximation of the true density than an approach which ignores the potential dependence. We show the superior performance of our method in several empirical examples, where we construct (i) quarterly forecasts using month-on-month *direct* forecasts, (ii) annual-average forecasts using monthly year-on-year *direct* forecasts, and (iii) annual-average forecasts using quarter-on-quarter *direct* forecasts.

Keywords: Joint Predictive Distribution, Frequency Transformation, Path Forecasts, Cross-horizon Dependence.

JEL classification: C53, C32, E37

¹ Banque de France, 31 rue Croix des Petits Champs, 75049 Paris, France. *Corresponding Author*, matteo.mogliani@banque-france.fr.

² Banco de España and CEMFI, Calle de Alcalá 48, 28014 Madrid, Spain. florens.odendahl@bde.es.

NON-TECHNICAL SUMMARY

Forecasting has become increasingly probabilistic, with central banks and research institutions routinely relying on density forecasts that characterize the full distribution of possible outcomes for key economic variables. In practice, however, multi-horizon forecasts are typically produced for short-term changes in the target variable (such as month-on-month or quarter-on-quarter growth rates), whereas users often require predictions for aggregated, calendar-consistent, horizons, including annual average or year-on-year growth rates. When the input densities are combined by a “dependence-inattentive” forecaster, the resulting predictive distributions fail to account for the serial-dependence that naturally arises across successive periods. Consequently, uncertainty is systematically underestimated, particularly in the tails of the distribution, precisely where policymakers’ interests frequently lie.

To address this problem, this paper develops a general methodology for transforming individual marginal predictive densities into coherent aggregated forecasts while explicitly capturing the serial-dependence across horizons. The approach relies on copulas, i.e. multivariate cumulative distribution functions that reconcile the marginal densities with a dependence structure, where the correlation parameters are estimated from historical forecast errors using the probability-integral transforms of realized outcomes. Incorporating this dependence allows the “dependence-attentive” forecaster to simulate the full joint distribution and, in turn, to construct predictive densities for any aggregated target variable, such as annual average growth rates derived from quarter-on-quarter growth forecasts.

The proposed methodology is intentionally simple, model-agnostic, and practical. It can be implemented using existing forecasts without altering the underlying econometric models. As a result, forecasters can enhance the accuracy and coherence of their risk assessments across horizons while preserving their established forecasting frameworks.

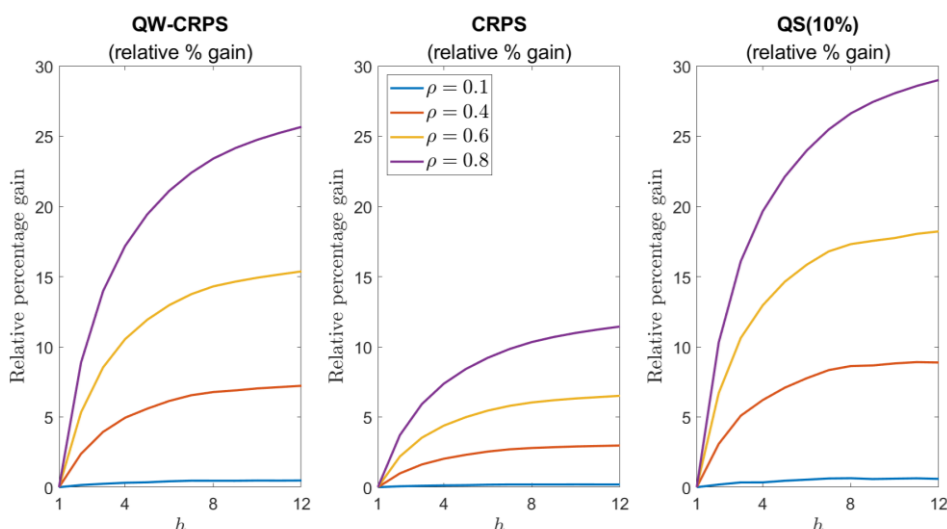
The paper evaluates this new approach using both Monte Carlo experiments and empirical applications. The simulation results demonstrate that incorporating cross-horizon dependence yields better forecasting, with more accurate predictive densities, especially when the underlying economic variables exhibit persistence or when the aggregation spans multiple periods. The improvements are particularly pronounced for scoring rules that place greater weight on tail accuracy, measures that are central in risk management.

Three empirical exercises are presented in the figure below in order to further illustrate the method’s practical value. First, using a large-scale set of bivariate regressions with US data from the FRED-MD database, density forecasts for month-on-month growth rates are transformed into quarter-on-quarter growth rates. Second, monthly density forecasts of year-on-year US inflation are converted into annual-average predictive distributions. Third, annual-average densities of US GDP growth are obtained from quarterly predictive distributions. In all cases, the copula-based approach consistently outperforms benchmarks that assume independence across horizons, as measured by the difference between the “dependence attentive” forecaster (who implements our proposed methodology) and the “dependence inattentive” (who does not).

The findings underscore three key messages. First, ignoring serial dependence across forecast horizons leads to a systematic underestimation of risk and uncertainty. Second, the proposed copula-based transformation offers a practical and cost-effective means of correcting this deficiency without the need to redesign existing forecasting models. Third, the largest gains arise for tail-sensitive and long-horizon applications, which are particularly relevant for policy analysis under uncertainty.

Overall, the study provides a robust and easy-to-implement framework for transforming density forecasts into new target distributions. By more accurately capturing joint uncertainty across time, the approach enhances both the reliability and the policy usefulness of probabilistic forecasts.

Scores for density forecasts: “dependence-attentive” vs “dependence-inattentive”



Note: Forecasting performance is measured using different metrics. QW-CRPS denotes the quantile-weighted version of the Continuous Ranked Probability Score, which places greater emphasis on the tails of the distribution. CRPS refers to the standard Continuous Ranked Probability Score. QS(10%) denotes the quantile score evaluated at the 10th percentile. All three metrics, CRPS, QW-CRPS, and QS(10%), are reported as relative percentage gains of the “dependence-attentive” forecaster (who implements the proposed methodology) over the “dependence-inattentive” forecaster (who does not).

Transformations de Densités Prédicatives

RÉSUMÉ

Le recours fréquent à un schéma de prévision directe implique que les prévisions individuelles ignorent la dépendance existante entre les horizons. Or, cette dépendance doit être prise en compte lorsque le prévisionniste souhaite construire des objets prédictifs qui sont des fonctions de plusieurs horizons (par exemple, lors du calcul d’une moyenne annuelle à partir de taux de croissance trimestriels). Pour résoudre à cette problématique, nous proposons d'utiliser des copules afin de combiner les distributions prédictives marginales à h pas d’avance en une distribution prédictive conjointe. Notre méthode est particulièrement intéressante pour les praticiens pour lesquels modifier la spécification de prévision directe serait trop coûteux. Dans une étude Monte Carlo, nous montrons que notre approche fournit une approximation plus précise de la densité véritable qu’une méthode qui ignore la dépendance entre horizons. Nous démontrons également la supériorité de notre approche au travers de plusieurs applications empiriques, dans lesquelles nous construisons (i) des prévisions trimestrielles à partir de prévisions mensuelles, (ii) des prévisions annuelles à partir de prévisions mensuelles en glissement annuel, et (iii) des prévisions annuelles à partir de prévisions trimestrielles.

Mots-clés : distribution prédictive conjointe, transformation de fréquence, prévisions de trajectoire, dépendance inter-horizon.

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ils sont disponibles sur publications.banque-france.fr

1 Introduction

Forecasting models are often specified to produce *direct* h -step-ahead forecasts, which implies that the predictions ignore information on their cross-horizon dependence. As a consequence, individual h -step-ahead predictive distributions cannot easily be transformed into predictive objects that depend on several horizons. For example, the literature on macroeconomic risk often uses quantile regression models that, in brief, produce *direct* density forecasts of quarter-on-quarter (qoq) real GDP growth (Adrian et al., 2019; Ferrara et al., 2022). However, it is not straightforward in this framework to construct, for instance, density forecasts for annual-average growth rates from the estimated quarter-on-quarter density forecasts because of their potential serial dependence.

To address this issue, we propose using Gaussian copulas to combine the information of the marginal *direct* h -step-ahead predictive densities into a joint distribution, which reflects the serial dependence between the marginals. This enables the practitioner to draw from the multi-horizon distribution and, therefore, to construct predictive distributions that are functions of several horizons, which we label target-frequency predictive densities.¹ The resulting multi-horizon predictive objects provide forecasts that are coherent with the underlying marginal predictive densities, in the sense that the moments of the multi-horizon objects are a function of the marginals and the copula parameters.

While separate forecasting specifications could be used to make predictions for the different predictive objects, there are several reasons why a single specification can be preferable. First, the scenario of having a single forecasting specification is particularly common among professional forecasters and institutions, such as central banks, where changing the forecasting process is costly and yet transformations of the existing forecasts to other target-frequencies are often required. Second, using a single forecasting specification leads to predictive objects with moments that are coherent across the target-frequencies. For instance, using our approach, the mean of the annual-average forecast is consistent with the mean of the quarter-on-quarter predictions. Third, some density forecasts are derived from surveys that only report one frequency but not the required target-frequency. In general, our approach treats the underlying density forecasts as the primitives such that it does not require a specific approach for their construction and, therefore, helps to broaden the usability of already available individual predictive densities that are based on a *direct* forecasting scheme.

As an alternative to our approach, the researcher could use a simple approach that assumes independence between the different marginal predictive densities, i.e., no correlation between the direct h -step-ahead predictions at different horizons. However, this approach neglects the serial dependence typically present in macroeconomic indicators, which is important for the tails of the multi-horizon distribution. For example, accounting for the positive serial correlation in quarter-on-quarter GDP growth forecasts leads to annual-average predictive distributions with considerably larger variance. This is related to the fact that the variance of the sum is not equal to the sum of the variances for random variables that are correlated. In other words, an approach that constructs the annual-average growth while taking into account the path of quarterly GDP growth

¹Note that for models that produce *iterative* h -step-ahead predictive densities, for instance Vector Autoregressions, forecasts conditional on specific paths or the annual-average frequency transformation are in general readily available since the iterative approach allows to draw conditionally on all the other forecast horizons.

over several quarters and, in combination with a positive correlation, would reflect the feature that large positive (negative) growth is typically followed by positive (negative) growth. This makes our approach particularly appealing for macroeconomic and financial risk applications.

In this paper, we show analytically that the proposed copula approach for density forecast transformations may outperform substantially a benchmark approach that neglects the cross-horizon dependence of the marginal densities. We also show how the resulting gain in the forecasting performance depends on the persistence of the underlying process. Further, we show in several Monte Carlo studies that our approach provides a better approximation to the true underlying annual-average density forecasts under different DGPs. This result holds under misspecified forecasting models, i.e., when the true multivariate distribution is not Gaussian, and for small training samples for the copula parameter estimation.

For the application of our approach, the researcher only needs to compute the correlation between the empirical PITs of the individual h -step-ahead predictive distributions for different horizons in a training sample. In particular, the forecaster needs to (i) compute the sequence of realized PITs for the marginal predictive densities at each forecast horizon $h = 1, \dots, H$, from a pseudo out-of-sample exercise over a training sample, and to (ii) combine the marginal distributions into a joint distribution via a multivariate Gaussian copula, where the maximum-likelihood estimator of the copula parameters is equal to the rank correlation of the realized PITs.

We demonstrate the usefulness of our methodology in three empirical applications, where we compare the forecasting performance of the proposed copula approach with a benchmark approach that ignores the serial dependence between the multi-horizon marginal distributions. The first application is a large-scale forecasting exercise based on monthly data from FRED-MD (McCracken and Ng, 2016) entering bivariate ARDL models regressing a large number of pairs of randomly selected variables. We first compute density forecasts for month-on-month values from these regressions and we then use these predictive densities to compute quarter-on-quarter density forecasts. Results show that the copula approach outperforms the benchmark for the majority of randomly specified bivariate models. The second empirical application aims at emulating a situation in which a forecaster dispose of predictive densities for year-on-year U.S. monthly inflation, but she needs to summarize the picture of the expected inflation environment by converting the target frequency of the predictive densities from monthly year-on-year rate to annual-average rate. Importantly, the year-on-year and annual-average predictions need to be coherent, i.e., they should be based on the same predictors and model type, and the central tendency of the forecasts across the two frequencies should be very similar. The results show that our copula approach provides significantly better density forecasts of annual-average inflation, in particular at the tails of the distributions, than the benchmark. Finally, in the third empirical application, we use the predictive densities of quarter-on-quarter U.S. real GDP growth from Adrian et al. (2019), which are based on *direct* forecasts, and transform them into annual-average forecasts. We find again that the annual-average forecasts based on our copula approach leads to more accurate predictive densities than the benchmark approach.

Our paper contributes to the literature on density forecasts and economic risk prediction. While Patton (2006) introduced conditional copulas to economics with a focus on modeling the cross-sectional dependence of predictive objects (see Fan and Patton, 2014, for review), we show

how copulas can be used for density forecast transformations. In particular, compared to existing contributions, the proposed copula-based approach is designed to combine the marginal forecast densities, unimodal (Adrian et al., 2019; Ferrara et al., 2022) or multimodal (Mitchell et al., 2024), into new predictive objects of several horizons. Differently from Smith and Vahey (2016), who use copulas to model both the time series and cross-sectional correlation of U.S. macroeconomic variables, we propose to use copulas for modeling the time series correlation and transforming marginal density forecasts into multi-horizon objects. Finally, compared with recent works also proposing to model the joint distribution of predictive objects (e.g., Clark et al., 2020, for the joint distribution of point forecast errors obtained from surveys via multi-variate stochastic volatility models; Grothe et al., 2023, for joint forecasts of hourly electricity prices from point forecasts; Ganics et al., 2024, for the construction of fixed-horizon density forecasts out of fixed-event survey density forecasts), we propose to use Gaussian copulas to combine marginal predictive densities of macroeconomic indicators to obtain predictive objects that are transformed to a new target-frequency.

The paper is structured as follows. Section 2 provides an analytical example of our forecasting environment. Section 3 describes the methodological framework. Section 4 presents Monte Carlo results. In Section 5 we provide a robustness analysis and discuss alternative simpler approaches. Section 6 presents the results from the three empirical exercises. Finally, Section 7 concludes.

2 Motivating example

Consider the following simple mean-zero autoregressive model:

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

with $|\rho| < 1$ and $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. The optimal h -steap ahead prediction (under both iterated and direct forecasting approach) is given by

$$Y_{t+h|t} = \rho^h Y_t$$

It follows that the forecast error $e_{t+h|t} = \sum_{j=0}^{h-1} \rho^j \varepsilon_{t+h-j}$ has second moment

$$\mathbb{V}(e_{t+h|t}) = \sigma_\varepsilon^2 \left(\frac{1 - \rho^{2h}}{1 - \rho^2} \right)$$

and auto-covariance and auto-correlation functions:

$$\begin{aligned} \text{Cov}(e_{t+h|t}, e_{t+h-k|t}) &= \sigma_\varepsilon^2 \rho^k \left(\frac{1 - \rho^{2(h-k)}}{1 - \rho^2} \right) \\ \text{Corr}(e_{t+h|t}, e_{t+h-k|t}) &= \rho^k \sqrt{\frac{1 - \rho^{2(h-k)}}{1 - \rho^{2h}}} \end{aligned}$$

for $h > k > 0$. $Corr$ denotes the Pearson correlation. Thus, the process has the following conditional predictive distribution:

$$Y_{t+h|t}|\rho, \sigma_\varepsilon \sim \mathcal{N}\left(\rho^h Y_t, \sigma_\varepsilon^2 \frac{1-\rho^{2h}}{1-\rho^2}\right)$$

with the pdf of $Y_{t+h|t}$ denoted by $\phi_{Y_{t+h|t}}(y_{t+h}|\rho, \sigma_\varepsilon)$. Now consider a linear transformation of the forecast sequence $\{Y_{t+j|t}\}_{j=1}^h$, such as

$$Z_{t+h|t} = w_1 Y_{t+1|t} + w_2 Y_{t+2|t} + \dots + w_h Y_{t+h|t}, \quad (1)$$

with $\mathbf{w} = (w_1, \dots, w_h)$ a vector of weights. This transformation is often useful in macroeconomic applications when the original forecasts need to be converted into a different target periodic measure of the same variable. For instance, if Y_t is a month-on-month growth rate sampled at monthly frequency and $Y_{t+h|t}$ is its h -step ahead forecast, then for $h = 12$ and $\{w_j\}_{j=1}^h = 1$, the transformed forecast

$$Z_{t+12|t} = Y_{t+1|t} + Y_{t+2|t} + \dots + Y_{t+12|t} \quad (2)$$

is (approximately) the 12-months ahead forecast of the year-on-year growth rate.

Following the example in (1), the conditional predictive distribution of the “dependence-attentive” transformed forecast is:

$$Z_{t+h|t}|\rho, \sigma_\varepsilon, \mathbf{w} \sim \mathcal{N}\left(\sum_{j=1}^h w_j \rho^j Y_t, \sigma_\varepsilon^2 \sum_{j=1}^h w_j^2 \left(\frac{1-\rho^{2j}}{1-\rho^2}\right) + 2\sigma_\varepsilon^2 \sum_{j=2}^h \sum_{k=1}^{j-1} w_j w_{j-k} \rho^k \left(\frac{1-\rho^{2(j-k)}}{1-\rho^2}\right)\right), \quad (3)$$

which is the sum of random variables from the joint multivariate Normal forecast distribution

$$(Y_{t+1|t}, \dots, Y_{t+h|t})'|\rho, \sigma_\varepsilon, \mathbf{w} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (4)$$

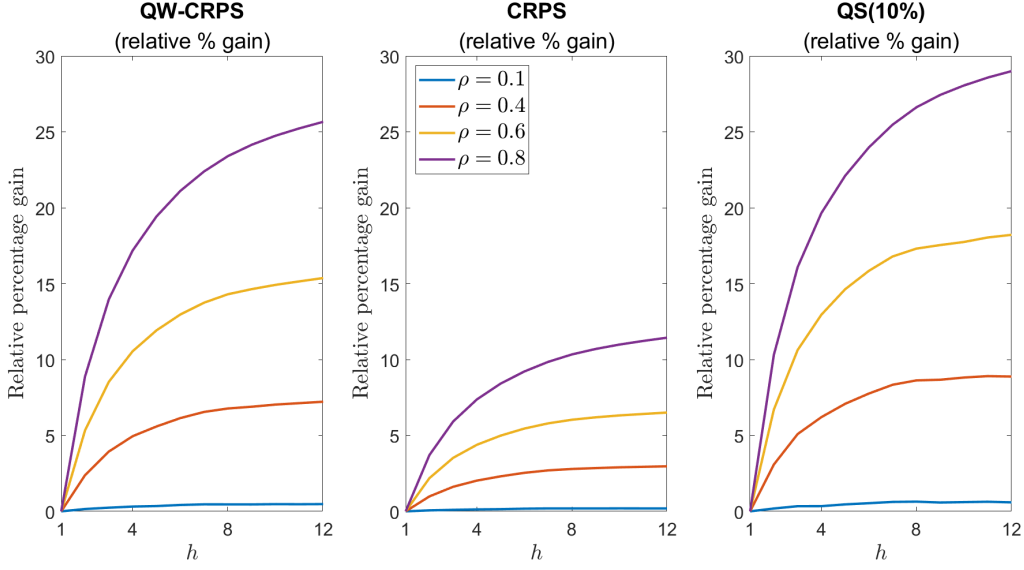
where $\boldsymbol{\mu} = (w_1 \rho Y_t, w_2 \rho^2 Y_t, w_3 \rho^3 Y_t, \dots, w_h \rho^h Y_t)'$ and

$$\boldsymbol{\Sigma} = \sigma_\varepsilon^2 \begin{pmatrix} w_1^2 & w_1 w_2 \rho & w_1 w_3 \rho^2 & \dots & w_1 w_h \rho^{h-1} \\ w_2 w_1 \rho & w_2^2 \left(\frac{1-\rho^4}{1-\rho^2}\right) & w_2 w_3 \rho \left(\frac{1-\rho^4}{1-\rho^2}\right) & \dots & w_2 w_h \rho^{h-2} \left(\frac{1-\rho^4}{1-\rho^2}\right) \\ w_3 w_1 \rho^2 & w_3 w_2 \rho \left(\frac{1-\rho^4}{1-\rho^2}\right) & w_3^2 \left(\frac{1-\rho^6}{1-\rho^2}\right) & \dots & w_3 w_h \rho^{h-3} \left(\frac{1-\rho^6}{1-\rho^2}\right) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ w_h w_1 \rho^{h-1} & w_h w_2 \rho^{h-2} \left(\frac{1-\rho^4}{1-\rho^2}\right) & w_h w_3 \rho^{h-3} \left(\frac{1-\rho^6}{1-\rho^2}\right) & \dots & w_h^2 \left(\frac{1-\rho^{2h}}{1-\rho^2}\right) \end{pmatrix}$$

The pdf of $Z_{t+h|t}$ is denoted by $\phi_{Z_{t+h|t}}(z_{t+h}|\rho, \sigma_\varepsilon, \mathbf{w})$. The predictive distribution in (3) can be compared to the predictive distribution of the forecaster who ignores the cross-horizon dependence of the forecasts. We label this forecaster the “dependence-inattentive” forecaster because no attention is paid to the potential correlation structure of the forecast errors across horizons. This forecaster draws from the following conditional forecast distribution:

$$\tilde{Z}_{t+h|t}|\rho, \sigma_\varepsilon, \mathbf{w} \sim \mathcal{N}\left(\sum_{j=1}^h w_j \rho^j Y_t, \sigma_\varepsilon^2 \sum_{j=1}^h w_j^2 \left(\frac{1-\rho^{2j}}{1-\rho^2}\right)\right), \quad (5)$$

Figure 1: Scores for density forecasts: "dependence-attentive" vs "dependence-inattentive"



Note: QW-CRPS denotes the quantile weighted versions of the continuous ranked probability score, with emphasis on the tails. CRPS denotes the continuous ranked probability score. QS(10%) denotes the quantile score at the 10% quantile. CRPS, QW-CRPS, and QS(10%) are expressed in relative % gain of the "dependence-attentive" forecaster with respect to the "dependence-inattentive" forecaster.

where the pdf of $\tilde{Z}_{t+h|t}$ is denoted by $\tilde{\phi}_{\tilde{Z}_{t+h|t}}(\tilde{z}_{t+h}|\rho, \sigma_\varepsilon, \mathbf{w})$. Note that for $\{w_j\}_{j=1}^h = 1$, the conditional predictive distribution of the "dependence-attentive" forecaster in (3) simplifies to:

$$Z_{t+h|t}|\rho, \sigma_\varepsilon, \mathbf{w} \sim \mathcal{N}\left(\sum_{j=1}^h \rho^j Y_t, \sigma_\varepsilon^2 \sum_{j=1}^h \left(\frac{1-\rho^j}{1-\rho}\right)^2\right). \quad (6)$$

The panels of Figure 1 show the average relative accuracy of the "dependence-attentive" density forecast in (6) compared to the "dependence-inattentive" density forecast in (5) with $\{w_j\}_{j=1}^h = 1$, evaluated through proper scoring rules, such as the quantile-weighted continuous ranked probability score (QW-CRPS; Gneiting and Ranjan, 2011) which emphasizes the tails, the continuous ranked probability score (CRPS; Gneiting et al., 2007), and the 10% quantile score (QS(10%); Giacomini and Komunjer, 2005). All these metrics show robust gains for the "dependence-attentive" forecast, which increases monotonically with ρ and h . For instance, with $\rho = 0.6$ and $h = 12$ (the year-on-year growth rate transformation for month-on-month predictions), the gain would stand about 7% according to the CRPS, while the QW-CRPS and the QS(10%) point to a gain of about 15-18%. Not surprisingly, the latter suggest that substantial accuracy gains can be obtained when evaluating the tails of the predictive densities, rather than their central region. This is due to the fact that the two densities differ solely in their variance, while the other moments are identical.²

In this stylized example, the researcher knows the dependence structure across forecasting

²It is nevertheless worth noting that these results depend on the sign of ρ . With $\rho < 0$, the monotonicity feature is in part lost, in particular for large (negative) autoregressive coefficients and small h . However, the "dependence-attentive" forecaster outperforms the "dependence-inattentive" forecaster even under this parameterization, in particular for large forecast horizons. These results are reported in Figure A.1 in the Appendix.

horizons and, therefore, the joint distribution of the $Y_{t+h|t}$. However, in practice when using direct forecasting schemes, this dependence structure is unknown and not estimated alongside the predictions. The next section, therefore, introduces a methodology to estimate the dependence structure for a given set of direct density forecasts.

3 Constructing multivariate densities with copulas

This section first introduces the necessary notation and then explains the methodology to transform the marginal density forecasts into the desired target multi-horizon object.

3.1 Notation

$\{Y_{t+j|t}\}_{j=1}^h$ denotes the higher frequency prediction that the researcher would like to transform into $Z_{t+h|t}$. The variable $Z_{t+h|t}$ is a function of $Y_{t+1|t}, \dots, Y_{t+h|t}$ and potentially of Y_t and its lags. The time subscript t and the forecast horizon subscript h denote the time units of $Y_{t+h|t}$. For instance, if $Y_{t+h|t}$ denotes monthly predictions then $Z_{t+h|t}$ denotes a transformation of the monthly forecasts using at most predictions up to month $t+h$. Recalling the example of the previous section, if Y_t is a month-on-month growth rate sampled at monthly frequency, then for $h = 12$ and $\{w_j\}_{j=1}^h = 1$, the transformed forecast

$$Z_{t+12|t} = Y_{t+1|t} + Y_{t+2|t} + \dots + Y_{t+12|t} \quad (7)$$

is approximately the 12-month-ahead forecast of the monthly year-on-year growth rate. In this case, we shall denote $Z_{t+h|t}$ a *periodic-transformation* of the forecast sequence $\{Y_{t+j|t}\}_{j=1}^h$.

Similarly, a linear transformation can map the forecasts generated at the original sampling frequency of the data into a forecast sequence sampled at a desired lower target frequency. For instance, let Y_t be a month-on-month growth rate sampled at monthly frequency. Then, for $h = 6$, the transformed forecast

$$Z_{t+6|t} = \frac{1}{3}Y_{t+2|t} + \frac{2}{3}Y_{t+3|t} + Y_{t+4|t} + \frac{2}{3}Y_{t+5|t} + \frac{1}{3}Y_{t+6|t} \quad (8)$$

is (approximately) the two-quarter-ahead forecast of the quarterly quarter-on-quarter growth rate (see [Mariano and Murasawa, 2003](#)).³ In this case, we shall denote $Z_{t+h|t}$ a *frequency-transformation* of the forecast sequence $\{Y_{t+j|t}\}_{j=1}^h$.

Throughout the paper, we use a boldface notation for vectors, matrices, and functions that take vectors as input.

3.2 Methodology

From the previous section, it is clear that constructing the predictive density of transformed forecasts requires drawing from the joint predictive distribution across the forecasting horizon. However, this can be often impractical in empirical applications, such as those relying on direct multi-step forecasting.

³This formula can be generalised to other frequency transformations, such as from month-on-month or quarter-on-quarter growth rates to annual-average growth rates, by changing the sequence of weights \mathbf{w} .

To address this issue, in this paper we propose to resort to (Gaussian) copulas (Sklar, 1959), which allow to model the marginals and the dependence separately while ensuring that the researcher obtains a valid multivariate distribution; see Nelsen (2006) for an introduction to copulas. Since we consider a forecasting environment, we work with the conditional copulas defined in Patton (2006) but drop the conditioning sets from the notation for simplicity.

Continuing with the example in Section 2, the predictive density $\phi_{Z_{t+h|t}}(\cdot|\rho, \sigma_\varepsilon, \mathbf{w})$, i.e. the density of $Z_{t+h|t}$, can be constructed by drawing from the joint distribution of $(Y_{t+1|t}, \dots, Y_{t+h|t})'|\rho, \sigma_\varepsilon$ described in eq. (4). Denoting by $\Phi_{Y_{t+h|t}}(y_{t+h})$ the CDF of the Gaussian predictive distribution of $Y_{t+h|t}$ evaluated at y_{t+h} , we note that:

$$\text{Corr}\left(\Phi_{Y_{t+h|t}}(y_{t+h}), \Phi_{Y_{t+h-k|t}}(y_{t+h-k})\right) = \frac{6}{\pi} \arcsin\left(\frac{1}{2}\rho^k \sqrt{\frac{1-\rho^{2(h-k)}}{1-\rho^{2h}}}\right) \quad (9)$$

which is the Pearson correlation coefficients of the probability integral transforms (PITs). The Gaussian copula is given by:

$$C_R = \Phi_R\left(\Phi^{-1}(U_1), \dots, \Phi^{-1}(U_h)\right)$$

where Φ^{-1} denotes the inverse CDF of a standard Normal, Φ_R the joint CDF of a standard multivariate Normal with covariance matrix R , and $U_i, i = 1, \dots, h$, are the PITs of the univariate predictive densities over the h forecasting horizons. Note that R is hence the correlation matrix, whose elements are defined in (9). Further note that the predictive distribution of the “dependence-inattentive” forecaster in (5) is also equivalent to the joint distribution of the forecasts constructed through a Gaussian copula, but with $R = \mathbf{I}_h$. Since C_R is the multivariate distribution of the random variables $\Phi^{-1}(U_i), i = 1, \dots, h$, given the copula, it is easy to resample the $Y_{t+1|t}, \dots, Y_{t+h|t}$ from their joint distribution:

$$(Y_{t+1|t}, \dots, Y_{t+h|t}) = \left(\Phi^{-1}(U_1), \dots, \Phi^{-1}(U_h)\right)$$

and then compute the desired (periodic or frequency) transformed density forecast from the sampled joint forecasts.

In the following we extend this approach to a more general forecasting environment. Assume 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 cumulative distribution functions (CDF) $\{\{G_{t,h}\}_{h=1}^H\}_{t=1}^T$, for outcome variable Y_{t+h} ; the subscript h denotes the forecast horizon and the subscript t denotes the forecast origin. Further, we treat the set of predictive distributions, $\{\{g_{t,h}\}_{h=1}^H\}_{t=1}^T$, as the primitives, i.e., we do not require a specific approach for their construction. This set of predictive distributions is what the forecaster has readily available at hand, for instance, due to institutional restrictions on the forecasting model to be used.

To illustrate the application of our methodology, but without loss of generality, we will assume that the predictive density $g_{t,h}$ is a predictive density for quarter-on-quarter growth rates.⁴ In period T , the forecaster is asked to provide predictive densities for the annual-average growth

⁴Predictive distributions for month-on-month growth rates, monthly or quarterly (log-)levels or year-on-year growth rates can be handled analogously with our approach.

rates as well as for the conditional predictive density $\tilde{g}_{T,h}(y_{T+h}|y_{T+h-1}, \dots, y_{T-1})$, henceforth called path-forecast, based on $\{\{g_{t,h}\}_{h=1}^H\}_{t=1}^T$.

We propose to do this by using copula functions, developed by [Sklar \(1959\)](#). A copula can be described as a function such that for any $\mathbf{Q}(y_1, \dots, y_d)$, where \mathbf{Q} is the multivariate distribution function of the random vector (Y_1, \dots, Y_d) , there is a copula function $\mathbf{C}(\cdot|\mathbf{R})$, such that $\mathbf{Q}(y_1, \dots, y_d) = \mathbf{C}(G_{Y_1}(y_1), \dots, G_{Y_d}(y_d)|\mathbf{R})$, where G_{Y_1}, \dots, G_{Y_d} are the marginal CDFs of Y_1, \dots, Y_d , respectively, and \mathbf{R} denotes the parameter(s) that governs the dependence between $G_{Y_1}(y_1), \dots, G_{Y_d}(y_d)$. Inversely, a copula function \mathbf{C} , combined with marginal CDFs G_{Y_1}, \dots, G_{Y_d} , gives a multivariate distribution.

A popular copula family is the Gaussian copula, denoted by \mathbf{C}_{Ga} , where the dependence between the d variables is governed by the correlation matrix \mathbf{R} , with ones on the diagonal and the rank correlation of variable i and j as the respective off-diagonal element (i, j) .

Let then $\mathbf{Q}_T(y_{T+1}, \dots, y_{T+H}|\mathbf{R})$ denote the joint predictive CDF of Y_{T+1}, \dots, Y_{T+H} for forecast origin T , conditional on the correlation matrix \mathbf{R} and constructed using \mathbf{C}_{Ga} . In other words, we define $\mathbf{Q}_T(y_{T+1}, \dots, y_{T+H}|\mathbf{R}) = \mathbf{C}_{\text{Ga}}(G_{T,1}(y_{T+1}), \dots, G_{T,H}(y_{T+H})|\mathbf{R})$. Further, let $\text{PIT}_{t,h} = G_{t,h}(y_{t+h})$ the probability integral transform of the predictive densities, where y_{t+h} is the realized value. The forecaster can obtain an estimate of $\mathbf{Q}_T(y_{T+1}, \dots, y_{T+H}|\mathbf{R})$ by implementing the algorithm described below.

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 $\hat{\mathbf{R}}$.
3. Use $\hat{\mathbf{R}}$ in combination with \mathbf{C}_{Ga} to obtain the joint distribution $\hat{\mathbf{Q}}_T(y_{T+1}, \dots, y_{T+H}|\hat{\mathbf{R}})$.

The resulting multivariate distribution allows to sample the *direct* h -step-ahead predictions jointly, such that predictive objects that are functions of several horizons can be constructed. Note that the maximum likelihood estimator of \mathbf{R} is $\hat{\mathbf{R}}$, i.e. the maximum likelihood estimator of the correlation matrix under a Gaussian copula reduces to the rank correlation of the PITs.

To illustrate the use of Algorithm 1, consider the following example. The forecaster is asked in T , which is the last quarter of the year, to provide a predictive distribution of the annual-average growth for the next year. The forecaster, however, has only a set of *direct* quarter-on-quarter h -step-ahead growth rate predictions available, for $h = 1, \dots, 4$. To transform the quarter-on-quarter growth rates into annual-average predictions, the forecaster can obtain the set $\{[Y_{T,1}^{(s)}, Y_{T,2}^{(s)}, Y_{T,3}^{(s)}, Y_{T,4}^{(s)}]'\}_{s=1}^S$ of draws, for $s = 1, \dots, S$, from $\hat{\mathbf{Q}}_T^{-1}(y_{T+1}, \dots, y_{T+4}|\hat{\mathbf{R}})$. This can be easily implemented in standard statistics packages. First, do step 1 to 2 of Algorithm 1. Then, step 3 amounts to the following: compute the lower Cholesky decomposition of $\hat{\mathbf{R}}$, denoted by \mathbf{P} . Next, for each $s = 1, \dots, S$:

- (a) Draw a 4×1 vector of independent standard Normals, i.e. draw $\mathbf{X} \sim_{iid} \mathcal{N}(\mathbf{0}, \mathbf{I}_4)$ where \mathbf{I}_4 is the 4×4 identity matrix.
- (b) Compute the vector $\mathbf{U} = [U_1, U_2, U_3, U_4] = [\Phi(Z_1), \Phi(Z_2), \Phi(Z_3), \Phi(Z_4)]$, where $\Phi(\cdot)$ is the CDF of a standard Normal distribution and $[Z_1, Z_2, Z_3, Z_4]' = \mathbf{Z} = \mathbf{P}\mathbf{X}$.

(c) Evaluate $G_{T,1}(U_1), \dots, G_{T,4}(U_4)$ to get the vector of joint draws $[Y_{T,1}^{(s)}, Y_{T,2}^{(s)}, Y_{T,3}^{(s)}, Y_{T,4}^{(s)}]'$.

The set of vectors of joint draws $\{[Y_{T,1}^{(s)}, Y_{T,2}^{(s)}, Y_{T,3}^{(s)}, Y_{T,4}^{(s)}]'\}_{s=1}^S$ approximates $\hat{Q}_T(y_{T+1}, \dots, y_{T+4} | \hat{R})$ and can be used, alongside with the observations y_T, y_{T-1}, y_{T-2} , to obtain the predictive distribution of the annual-average growth rate using, for instance, an exact formula or a linear approximation formula as in eq. (8), for each $s = 1, \dots, S$. The multivariate distribution also allows to obtain draws of $Y_{t,h}$ conditional on $Y_{t,h-j}$ for $j = 1, \dots, h-1$.

Our approach aims at constructing well performing forecasts but we note that there are several potential sources of misspecification. First, the choice of the Gaussian copula might not reflect the true underlying multivariate distribution of the data. Second, if the marginal distributions are misspecified, the joint distribution will also be misspecified. Third, a potential source of misspecification comes from the conditional copulas. As shown by Patton (2006), if the conditioning set is not identical for the marginals and the copula, then the constructed conditional joint distribution might not reflect the actual conditional joint distribution. Following Patton's (2006) example, if the researcher conditions Y_1 on $W_1 = w_1$, Y_2 on $W_2 = w_2$ and the copula on W_1, W_2 , such that $\tilde{Q}_{Y_1, Y_2 | W_1, W_2}(y_1, y_2 | w_1, w_2) = C(G_{Y_1 | W_1}(y_1 | w_1), G_{Y_2 | W_2}(y_2 | w_2) | w_1, w_2, \mathbf{R})$, then \tilde{Q} will only denote the conditional joint distribution of $(Y_1, Y_2) | W_1, W_2$ if the marginal Y_1 is independent of W_2 and the marginal of Y_2 is independent of W_1 (for details see Patton, 2006). However, given the macroeconomic applications that we consider, this is unlikely to restrict the implementation of our approach for mainly two reasons. First, the marginals come from *direct* forecasting models that use the same predictors for each forecasting horizon. Second, if the predictors differ across forecasting horizons that is typically because different predictors are relevant at that horizon, i.e., the marginals are independent of the non-included predictors.

Further, our estimation algorithm assumes that (i) dependence parameters do not depend on the predictors used for the marginals and (ii) restricts the copula dependence parameters in R to be constant. Assumption (i) is common in the literature and across model classes. For instance, flexible models such as VARs with stochastic volatility do not make the time-varying variance-covariance matrix an explicit function of dependent variables. However, while we propose a simple algorithm for practitioners, both assumption (i) and (ii) can be relaxed. For instance, Patton (2006) suggests an ARMA-like specification to allow for time-variation in the copula parameters. Hafner and Manner (2012) propose a stochastic process for the copula parameters. Both specification could be extended to include exogeneous variables in the copula parameter equations.

Despite the various potential sources of misspecification, the Monte Carlo results in the next section show that our modelling approach outperforms the competitor approach and leads to forecasting performances that are often indistinguishable from the correctly specified predictive distribution.

4 Monte Carlo study

We study the performance of our suggested copula approach via Monte Carlo simulations in a scenario where the forecaster has a model that produces *direct* h -step-ahead predictive densities for qoq growth rates and then needs to transform the predictive densities into annual-average growth rates and yoy growth rates. The absolute and relative performance of the proposed approach

with respect to a simple benchmark ignoring cross-horizon dependence of the forecasts (the “dependence-inattentive” forecaster described in [Section 2](#)) are evaluated through an out-of-sample exercise.

4.1 Monte Carlo design

The underlying DGP of qoq growth rates, denoted by Y_t , takes the form of a VAR(1):

$$\begin{bmatrix} Y_t \\ X_t \end{bmatrix} = \begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} + \begin{bmatrix} \theta_1 & \theta_2 \\ 0 & \gamma \end{bmatrix} \begin{bmatrix} Y_{t-1} \\ X_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix} \quad (10)$$

where $\{\varepsilon_{j,t}\}_{t=1}^T$ are two uncorrelated sequences of independent and identically distributed (iid) structural shocks. We set $\varepsilon_{2,t} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{\varepsilon_2}^2)$, with $\sigma_{\varepsilon_2} = 0.3$, but we consider three different specifications for the error term $\varepsilon_{1,t}$: (i) a Normal distribution, (ii) a Skew-Normal distribution, or (iii) a Skew- t distribution. For cases (ii) and (iii), we adopt a location-scale-shape parameterization ([Azzalini and Capitanio, 2003](#)). All the distributions are calibrated to have mean zero and standard deviation $\sigma_{\varepsilon_1} = 0.5$, as well as negative skewness for cases (ii) and (iii) (with shape parameter $\alpha = -3$). For the Skew- t distribution, the degrees of freedom parameter is set to $\nu = 8$, which implies somewhat heavier tails than the Normal distribution. We consider these different specifications to allow for a varying degree of complexity in the DGP. As for the remaining parameters, we set $\tau_1 = 0.2$, $\tau_2 = 0$, and $\theta_2 = \gamma = 0.5$. To account for different degrees of serial correlation, and hence cross-horizon dependence in the multi-step forecasts, θ_1 takes one of the following values: $\theta_1 \in \{0.1, 0.4, 0.7\}$.

We consider two types of forecasting models, both of which are misspecified AR(1) and produce *direct* h -step-ahead forecasts. The first forecasting model is used when $\varepsilon_{1,t}$ in the DGP is drawn from a Normal distribution:

$$Y_{t+h} = \tau_h + \beta_h Y_t + u_{t+h}, \quad (11)$$

with $u_{t+h} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{u,h}^2)$.

The second forecasting model is a quantile regression specification used when $\varepsilon_{1,t}$ in the DGP is drawn from the Skew-Normal or the Skew- t distribution:

$$Y_{t+h}(q) = \tau_h(q) + \beta_h(q) Y_t + u_{t+h}(q), \quad (12)$$

where q denotes the quantile with $q \in \{0.05, 0.25, 0.5, 0.75, 0.95\}$, and $\tau_h(q)$ and $\beta_h(q)$ denote respectively the quantile specific intercept and autoregressive parameter. To obtain a full predictive distribution, we smooth the five predicted quantiles using the Skew- t of [Azzalini and Capitanio \(2003\)](#). Note that the latter introduces a second potential source of model misspecification, in addition to that implied by the specification in (12).

The parameters $\{\tau_h, \beta_h, \sigma_{u,h}^2\}$ and $\{\tau_h(q), \beta_h(q)\}_{q \in Q}$ are estimated using a rolling-window estimation scheme with sample size T_{is} , which we set to 200 (quarters). Estimated parameters are used to compute qoq predictive densities up to 12 quarters ahead, from which we get S forecast paths. We then use simple approximating formulas to construct annual-average predictive distributions from those paths. For instance, assuming that the forecast origin is the last quarter

of the year, for each forecast path $s = 1, \dots, S$ the one year-ahead ($h_a = 1$) annual-average forecast is given by:

$$Z_{t+4|t}^{(s)} = \frac{1}{4}Y_{t-2} + \frac{2}{4}Y_{t-1} + \frac{3}{4}Y_t + Y_{t+1|t}^{(s)} + \frac{3}{4}Y_{t+2|t}^{(s)} + \frac{2}{4}Y_{t+3|t}^{(s)} + \frac{1}{4}Y_{t+4|t}^{(s)} \quad (13)$$

and the four quarters-ahead year-on-year forecast:

$$Z_{t+4|t}^{(s)} = Y_{t+1|t}^{(s)} + Y_{t+2|t}^{(s)} + Y_{t+3|t}^{(s)} + Y_{t+4|t}^{(s)} \quad (14)$$

Two competing approaches are here considered: *i*) the benchmark approach, which constructs the annual-average and year-on-year forecasts by directly transforming the quarter-on-quarter forecast paths with (13) and (14), i.e., without accounting for the cross-horizon dependence; and *ii*) the copula approach, which in turn constructs the predictive distributions for the annual-average and yoy growth rates using first the methodology described in Section 3 to obtain joint draws of the quarter-on-quarter growth rates, and then expressions (13) and (14) to transform the joint draws.

The out-of-sample size used to compute the correlation of the empirical PITs for the copula approach is denoted by T_R and set to 50 (quarters).⁵ Similarly, the out-of-sample size used to evaluate the annual-average (resp. year-on-year) density forecasts is denoted by T_{oos} and also set to 50. In particular, we simulate 200 periods of quarterly data and we then produce annual-average (resp. year-on-year) forecasts every four quarters for horizons one, two, and three years ahead.⁶ Results are based on 500 Monte Carlo iterations. For each iteration we evaluate the performance of the benchmark and copula approach using both relative and absolute forecasting performance measures. For this purpose, in addition to the qoq growth rates, we also simulate the true annual-average and yoy growth rates generated by the DGP described in (10). As in Section 2, the relative performance measures include the QW-CRPS (calibrated to evaluate the predictive performance at the tails of the distribution) and the CRPS.⁷ The relative performance is evaluated both by directly comparing the two approaches and against the forecasts one could generate with knowledge of the true underlying DGP and the true parameter values. The absolute forecasting performance is measured through the test for the correct specification of the predictive distribution proposed by Rossi and Sekhposyan (2019), which is a test for uniformity of the PITs.

4.2 Monte Carlo results

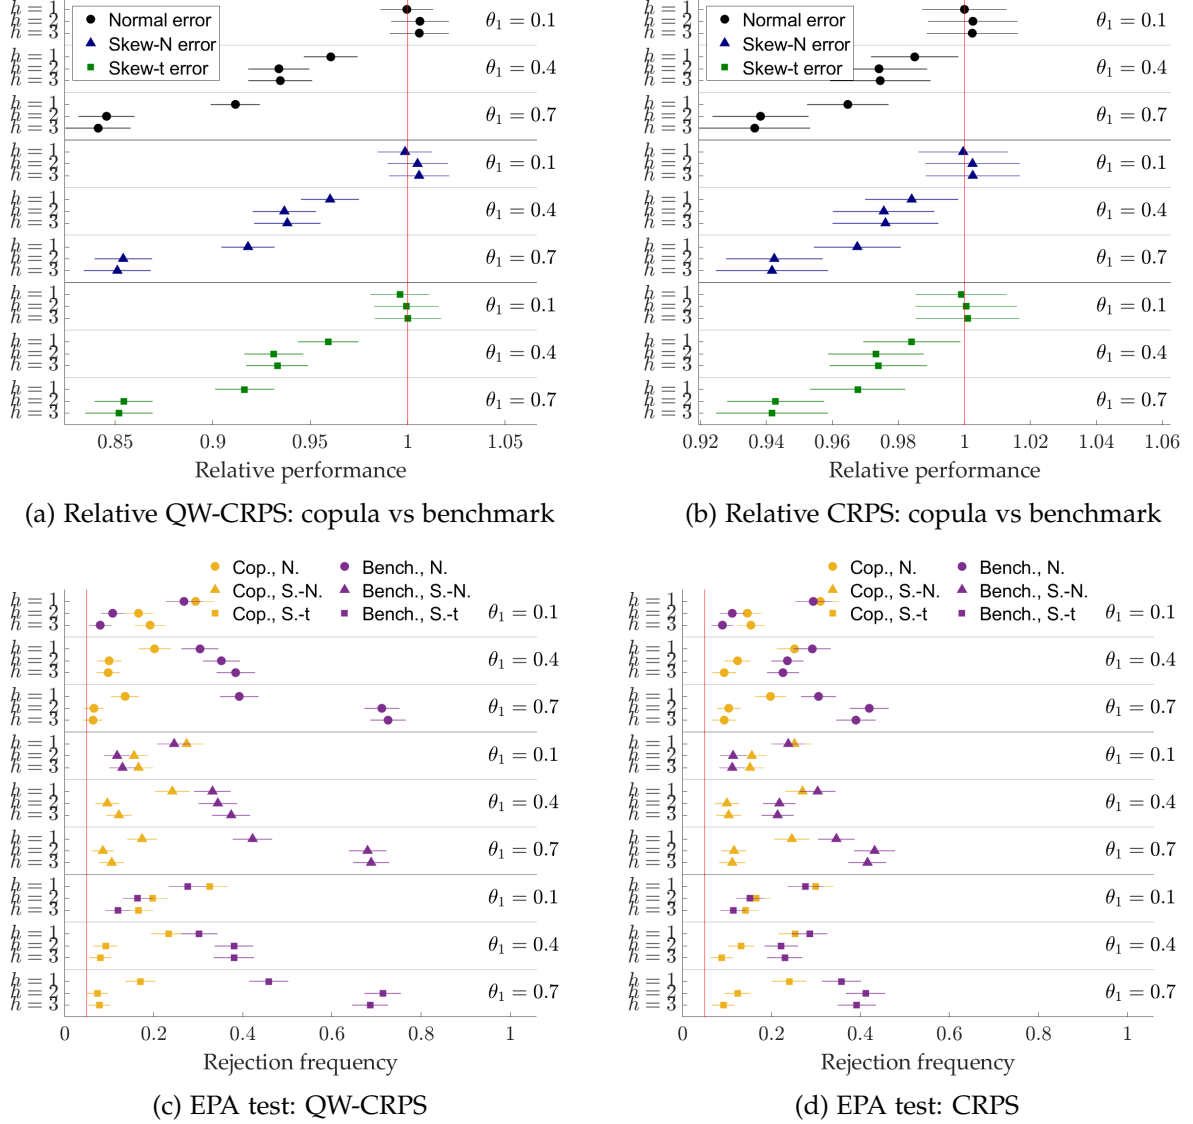
Figure 2 shows the relative forecasting performance of the benchmark and copula approach, as well as the rejection frequencies of an equal predictive ability test when comparing either of the two approaches against the optimal forecast. The markers show the average score ratios (rejection frequencies) across all Monte Carlo iterations, while the horizontal lines denote ± 2 bootstrap standard errors of the average score ratios (rejection frequencies). Panel (a) shows the ratio of the QW-CRPS of the copula approach and the benchmark approach, i.e., values < 1

⁵To reduce the computational costs in the Monte Carlo, for each simulation we compute the historical correlation of the empirical PITs once (the first forecast round) and we then use this estimate in all forecasting iterations, instead of updating it at every iteration.

⁶This strategy aims at replicating the calendar year predictions typically used in practice by professional forecasters.

⁷Results for quantile-score at the 10% lower tail, QS(10%), reported in Figure A.2 in the Appendix, are very similar to those for the QW-CRPS, and they are hence not commented for the sake of parsimony.

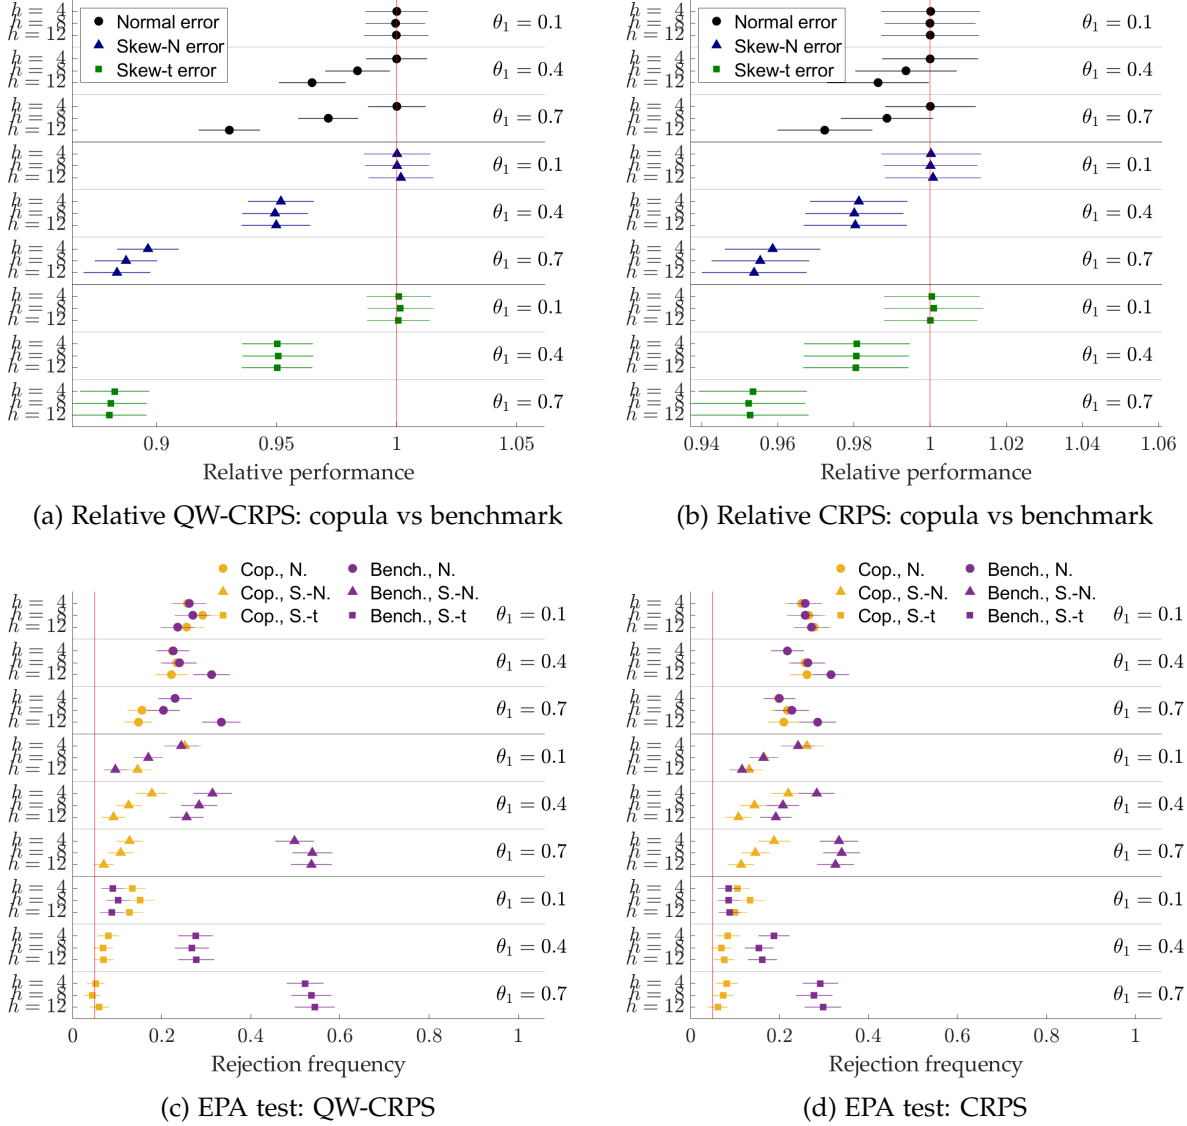
Figure 2: Monte Carlo results for annual-average forecasts: relative scores and EPA tests



Note: The θ_1 indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the annual-average horizon, i.e., one-year-, two-years-, and three-years-ahead. The x -axis in Panel (a) and (b) indicates the QW-CRPS (CRPS) of the copula relative to the benchmark approach, i.e., numbers smaller than one indicate a superior performance of the copula approach. The markers indicate the average score ratio across all Monte Carlo iterations. In Panel (c) and (d), the x -axis denotes the rejection frequency of the null hypothesis of a [Giacomini and White \(2006\)](#) test of unconditional equal predictive ability against the optimal forecasts. The markers indicate the average rejection frequency across all Monte Carlo iterations. The nominal size is 5%. The horizontal lines around the markers indicate ± 2 bootstrap standard errors of the average score ratio or rejection frequency. N., S.-N., and S.-t indicate the Normal, Skew-Normal, and Skew-t distribution of the error term $\varepsilon_{1,t}$ in the DGP. Standard errors of the tests were computed using a HAC estimator with bandwidth $= h_A - 1$.

indicate superior performance of the copula approach. Panel (b) shows the analogue ratio for the CRPS. As expected, for a serial dependence close to zero there is no improvement in the forecasting performance when using the copula approach. For medium to large values of θ_1 , the copula approach outperforms the benchmark approach by about 5% to 15% in the case of the QW-CRPS and by about 3% to 7% in the case of the CRPS. Panel (c) and (d) show the average rejection frequencies of the null hypothesis of equal predictive ability using the unconditional test

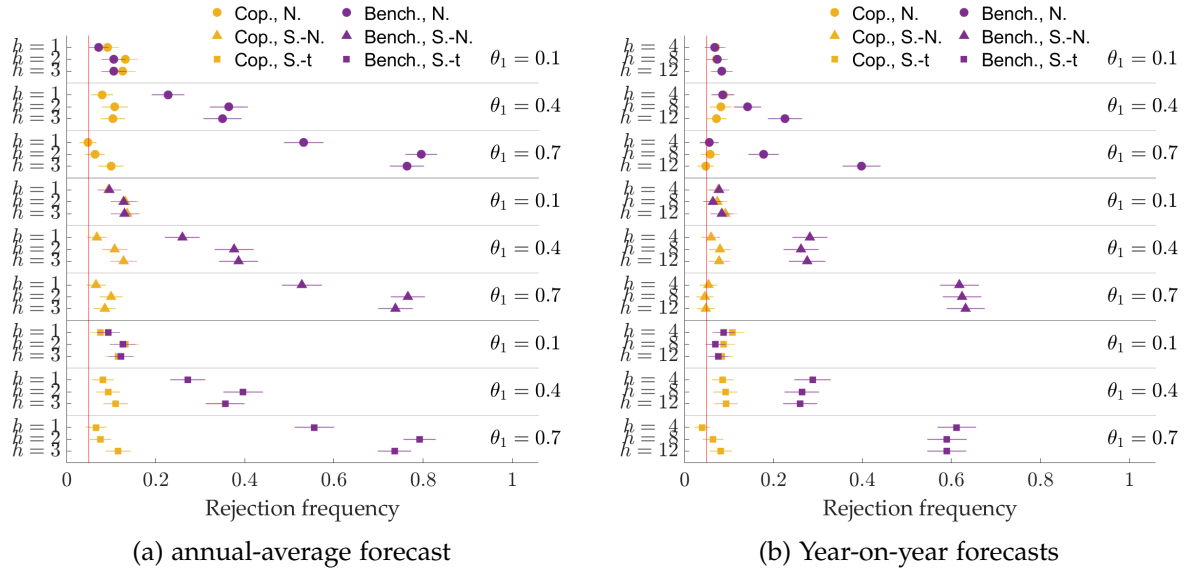
Figure 3: Monte Carlo results for year-on-year forecasts: relative scores and EPA tests



Note: The θ_1 indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the year-on-year horizon, i.e., four-quarter, eight-quarter-, and 12-quarter-ahead. The x -axis in Panel (a) and (b) indicates the QW-CRPS (CRPS) of the copula relative to the benchmark approach, i.e., numbers smaller than one indicate a superior performance of the copula approach. The markers indicate the average score ratio across all Monte Carlo iterations. In Panel (c) and (d), the x -axis denotes the rejection frequency of the null hypothesis of a [Giacomini and White \(2006\)](#) test of unconditional equal predictive ability against the optimal forecast. The markers indicate the average rejection frequency across all Monte Carlo iterations. The nominal size is 5%. The horizontal lines around the markers indicate ± 2 bootstrap standard errors of the rejection frequency. N., S.-N., and S.-t indicate the Normal, Skew-Normal, and Skew-t distribution of the error term $\varepsilon_{1,t}$ in the DGP. Standard errors of the tests were computed using a HAC estimator with bandwidth $= h_A - 1$.

of [Giacomini and White \(2006\)](#). The scoring functions are the QW-CRPS (c) and the CRPS (d) and the test compares the copula (benchmark) approach against the optimal forecast. The nominal size is 5%, i.e., rejection frequencies above 5% indicate that the null hypothesis is rejected more often than expected given the nominal size. Panel (c) and (d) of [Figure 2](#) show that for small values of θ_1 , the copula approach performs slightly worse than the benchmark approach. We attribute this result to the additional parameter estimation uncertainty induced by the estimation of the

Figure 4: Monte Carlo results of tests for correct forecast density specifications



Note: The θ_1 indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the annual-average and year-on-year horizon, respectively, in years and quarters. The x -axis denotes the rejection frequency of the Rossi and Sekhposyan (2019) test for correct specification of the predictive density. The markers indicate the average rejection frequency across all Monte Carlo iterations. The horizontal lines around the markers indicate ± 2 bootstrap standard errors of the rejection frequency. The nominal size is 5%. N., S.-N., and S.-t indicate the Normal, Skew-Normal, and Skew-t distribution of the error term $\varepsilon_{1,t}$ in the DGP.

copula parameters, which is likely to dominate the gain from taking the temporal dependence into account. In contrast, with increasing temporal dependence, i.e. increasing θ_1 , the rejection frequency associated with the copula approach tends to be close to the nominal size, whereas, for the benchmark approach, the null hypothesis of equal performance is rejected much more frequently.

It is worth noting that the rejection rate for the copula approach is close to the nominal size when compared to the true predictive density. This is an important result, since for the case with Skew-Normal and Skew-t errors, the Gaussian copula is potentially misspecified. Further, for the Skew-Normal case, the marginals are potentially misspecified since we use a Skew-t distribution to smooth the quantile regression predictions. Importantly, the results suggest that these misspecifications have only a small impact on the overall forecasting performance, because the copula-based predictive distributions appear often indistinguishable from true predictive distributions in terms of forecasting performance. As regards the year-on-year forecasts, the results are reported in Figure 3 and are very similar to the annual-average results.

Figure 4 shows the rejection frequencies of the null hypothesis of a correct specification of the predictive densities, evaluated using the test of Rossi and Sekhposyan (2019). Results show that a high temporal dependence implies that the benchmark approach leads to misspecified predictive distributions. In contrast, the copula approach adjusts for the temporal dependence and shows rejection frequencies only slightly above the nominal size. Hence the results in Figure 4 suggest that the consequences of the aforementioned misspecifications on the calibration of copula-based predictive densities are fairly negligible, as the rejection rates hovers around the nominal size of the test.

5 Robustness of the copula and alternatives

This section provides some guidance on how the training sample size and the persistence of the process affect the performance of the copula approach and discusses potential alternative approaches.


5.1 Robustness of the copula approach

This subsection further explores to what extent the performance of the copula approach depends on (i) the training sample for the copula parameter R , and (ii) the persistence of the underlying process. Our main goal here is to provide a guideline to the practitioners by setting some rules-of-thumb when choosing whether or not implementing the copula approach on specific applications. Note, however, that the conclusions drawn in this section are, to some extent, dependent on the DGP considered.


For our first robustness check, we repeat the Monte Carlo experiment of [Section 4](#), keeping the design of the simulations unchanged except for the number of in-sample observations T_{is} and the number of observations T_R used to estimate the copula correlation matrix R . For this purpose, we use a grid of $T_{is} = 100, 200$, and 400 and $T_R = 25, 50, 100$, and 200 observations. Note that the values in the baseline Monte Carlo experiments discussed in [Section 4](#) are $T_{is} = 200$ and $T_R = 50$. We focus on $\theta_1 = 0.4$, the case of intermediately strong serial correlation, we compute the same scores as in [Section 4](#), and we report the relative percentage gain of the copula approach with respect to the benchmark approach (neglecting cross-horizon dependence). [Figure 5](#) reports the results for the one- and three-year-ahead annual-average predictions, and point to an overall stability of the relative performance of the copula approach across simulations. A slight improvement seems to arise from larger T_{is} and T_R , in particular when $T_R > 25$, but the results are often only marginally different and statistically indistinguishable according to (bootstrap) standard errors of the average relative scores. In other words, the Monte Carlo study suggests that *i*) a moderate training out-of-sample size should be sufficient for estimating copula parameter matrix R and obtaining relatively accurate forecasting results from the copula approach, and *ii*) the training in-sample size used to estimate the model parameters appears less relevant.

For our second robustness check, we compute the relative forecasting performance of the copula approach for different levels of persistence measured by the determinant of the estimated correlation matrix R , $\det(\hat{R})$. On the one hand, in the extreme case of completely independent marginals across all the forecast horizons, the determinant of R is equal to one. On the other hand, in the extreme case of a perfect correlation across the marginals, the determinant is equal to zero. To evaluate the effect of different degrees of path-dependence across forecast horizons, we again repeat the Monte Carlo experiment of [Section 4](#), but for each Monte Carlo iteration we now draw (with replacement) some of the VAR(1) parameters from the following uniform distributions: $(\theta_1, \theta_2, \gamma) \sim \mathcal{U}(-0.9, 0.9)$ and $\sigma_{\epsilon_2} \sim \mathcal{U}(0.3, 0.7)$. To attenuate the effect of finite samples on the results, we set the training sample at $T_{is} = T_R = 500$ and we generate $T_{oos} = 400$ observations for the out-of-sample evaluation of the transformed densities. For ease of analysis, we focus here on the Normal error case, $\epsilon_1 \sim \mathcal{N}(0, \sigma_{\epsilon_1}^2)$, and on annual-average density forecasts. Results are presented in [Figure 6](#), where the y -axis denotes the relative performance gains and the x -axis shows ranges of $\det(\hat{R})$ categorized into bins with width 0.1: smaller values on the

Figure 5: Relative performance of annual-average forecasts across ranges of T_{is} and T_R



(a) One-year-ahead annual-average

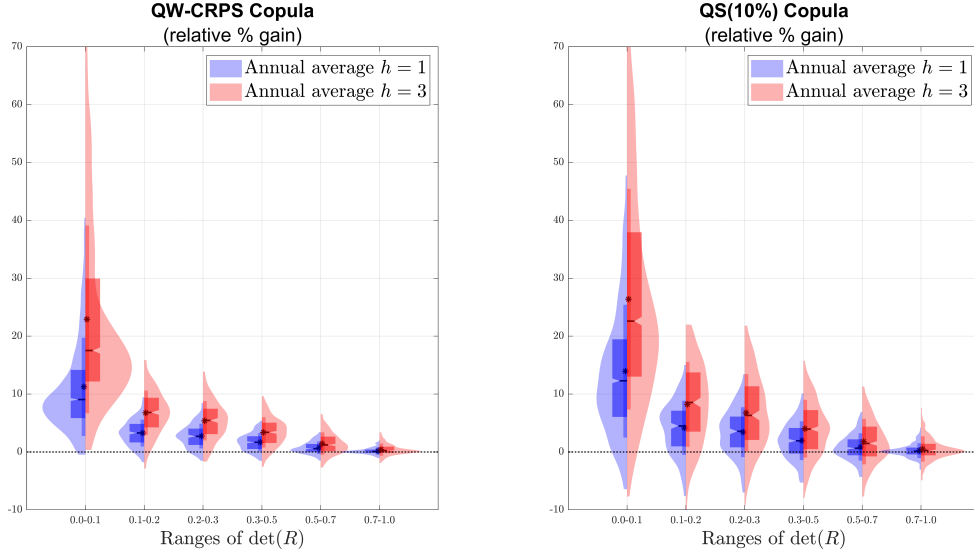


(b) Three-years-ahead annual-average

Note: Panel (a) shows results for the one-year-ahead annual-averages. Panel (b) shows results for the three-year-ahead annual-averages. T_{is} denotes the in-sample estimation size for the prediction model parameters. The rotated numbers on the x -axis denote the values of T_R , i.e. the estimation sample size for R . The units of the y -axis are the average percentage point gains in the relative forecasting performance of the copula to the benchmark approach. CRPS denotes the continuous ranked probability score. QW-CRPS denotes the quantile weighted versions of the continuous ranked probability score, with emphasis on the tails. QS(10%) denotes the quantile score at the 10% quantile. The error bars denote ± 2 Monte Carlo standard errors obtained by 999 bootstrap resamples of the average scores from 1000 Monte Carlo iterations.

x -axis indicate a higher level of persistence. The figure shows results for the QW-CRPS and the QS(10%) in so-called violin plots for the one- and three-year-ahead annual-average predictions. The curves in the figure depict the distribution of the gains over the 1000 Monte Carlo iterations, the boxes denote the inter-quartile range, the horizontal line-marker denotes the median, and the star-marker denotes the mean. For high levels of persistence, the gains are on average large and the distribution exhibits a large right-tail towards regions of big relative forecasting gains for the copula approach. For medium levels of $\det(\hat{R})$, the copula approach still shows substantial predictive gains relative to the benchmark approach. However, these gains tend to disappear with a decreasing persistence of the process.

Figure 6: Distribution of the relative performance of annual-average forecasts across ranges of $\det(\hat{R})$



Note: QW-CRPS denotes the quantile weighted versions of the CRPS. QS(10%) denotes the quantile-score at the 10% quantile. The y -axis shows the relative percentage point gain of the copula approach with respect to the benchmark. The x -axis shows different levels of the determinant of R ; larger values imply an overall smaller persistence in the process. The curves in the figure depict the distribution of the relative predictive gains over the 1,000 Monte Carlo iterations, the boxes denote the inter-quartile range, the horizontal line-marker denotes the median, and the star-marker denotes the mean.

5.2 Alternative approaches

It could be argued that the copula approach described so far may be inefficient compared to simpler alternative approaches that provide transformed forecasts in one single step. For instance, the researcher interested in the transformation in (2) could run regressions using directly the transformed variable Z_t , instead of modelling the data at the original sample. Note that while this would violate the requirement of coherence between the moments across the different frequencies, this section nonetheless entertains the idea of a direct regression of Z_t on its past history. However, it turns out that the forecasting regression based on the transformed sample, using the simplified example as in Section 2, does not lead to a significant superior forecasting performance.

Assume again that the DGP takes the form of the autoregressive model $Y_{t+1} = \rho Y_t + \varepsilon_{t+1}$ and that the target forecast follows the linear transformation $Z_{t+h|t} = Y_{t+1|t} + Y_{t+2|t} + \dots + Y_{t+h|t}$. A regression of the annual-average on its past values, or in our simplified illustrative example, Z_{t+h} on Z_t , leads to the following expression:

$$\begin{aligned}
 Z_{t+h} &= Y_{t+1} + Y_{t+2} + \dots + Y_{t+h} \\
 &= \rho^h \underbrace{\left(Y_{t-(h-1)} + Y_{t-(h-2)} + \dots + Y_t \right)}_{Z_t} + \underbrace{\left[\sum_{j=1}^h \left(\frac{1-\rho^j}{1-\rho} \right) \varepsilon_{t+h+1-j} + \sum_{j=1}^{h-1} \rho^j \left(\frac{1-\rho^{h-j}}{1-\rho} \right) \varepsilon_{t+1-j} \right]}_{u_{t+h}} \\
 &= \rho^h Z_t + u_{t+h}
 \end{aligned} \tag{15}$$

This suggests to construct a predictive density for $Z_{t+h|t}$ centered on the conditional mean

$\rho^h Z_t$ and featuring a variance based on u_{t+h} . First, note that the conditional mean is not optimal because Z_t is a linear combination of lags of Y_t , i.e., the forecaster partly conditions on "out-dated" information. This is akin to a situation where the forecaster only uses the lower frequency instead of taking a mixed-frequency approach. The information in the error terms, ε_t in (15) that already realized in t . In other words, the conditional mean forecast is not efficient. Second, note that in the regression Z_t and u_{t+h} in (15) are correlated, i.e., the OLS estimator of ρ is inconsistent.

We repeat the Monte Carlo experiment of Section 4 using the same DGP and running the alternative direct annual-average regression described in (15) using Z_t as predictor for $Z_{t+h|t}$. In this experiment we evaluate its forecasting performance relative to the copula approach. Figure 7 reports the results for $\theta_1 = 0.1, 0.4$, and 0.7 , three error distributions, and one- to three-year-ahead annual-average predictions (values <1 mean that the copula approach outperforms the alternative approach). The results interestingly point to a strong predictive gain of the copula approach compared to the alternative approach for the one-year ahead prediction, while from the two-years ahead onward the two approaches provide very similar results. Abstracting from parameter estimation errors and the parameter consistency problem in the estimation of the direct annual-average regression, the forecasting gains by the copula approach can partly be explained by better conditional mean predictions of the higher frequency regression. Consider again the autoregressive DGP $Y_t = \rho Y_t + \varepsilon_t$ and assume for simplicity and without loss of generality that the calendar year is composed of two periods and that target forecast follows the weighted linear transformation:

$$Z_{t+2|t} = \frac{1}{2}Y_t + Y_{t+1|t} + \frac{1}{2}Y_{t+2|t} \quad (\text{for one-year-ahead predictions})$$

$$Z_{t+4|t} = \frac{1}{2}Y_{t+2|t} + Y_{t+3|t} + \frac{1}{2}Y_{t+4|t} \quad (\text{for two-year-ahead predictions})$$

which is akin to an actual one-year- and two-years-ahead annual-average forecasts if the Y_t were to represent semi-annual growth rates. Note that the expression for $Z_{t+2|t}$ embeds the observation at the forecast origin Y_t , in addition to the forecasts produced at forecast origin t , while $Z_{t+4|t}$ includes only forecast terms. The mean squared forecast errors ratio of the annual-average prediction provided by the copula approach (MSFE_C) and the alternative direct annual-average approach (MSFE_{AAD}) are:


$$\frac{\text{MSFE}_C(t+2|t)}{\text{MSFE}_{AAD}(t+2|t)} = \frac{\rho^2 + 4\rho + 5}{6\rho^2 + 8\rho + 6} \quad (16)$$

$$\frac{\text{MSFE}_C(t+4|t)}{\text{MSFE}_{AAD}(t+4|t)} = \frac{\rho^6 + 4\rho^5 + 7\rho^4 + 8(\rho^3 + \rho^2 + \rho) + 6}{6\rho^6 + 8(\rho^5 + \rho^4 + \rho^3 + \rho^2 + \rho) + 6} \quad (17)$$

For $h = 1$, it can be shown that the obtained expression in (16) (a ratio of two second-order polynomials) is always lower than 1 for $|\rho| < 1$, and it ranges between 0.8 and 0.6 for $0.1 \leq \rho \leq 0.7$. Conversely, for $h = 2$, the ratio of mean square forecast errors simplifies to a ratio of two sixth-order polynomials, which is very close to 1 for $-0.6 \leq \rho \leq 0.5$ and still above 0.9 up to $\rho = 0.8$. These analytical results are consistent with the simulation results reported in Figure 7, where a slightly different DGP and additional metrics are considered.

In summary, the expected forecasting performance of the direct annual-average approach is


Figure 7: Monte Carlo results for annual-average forecasts: relative scores



(a) QW-CRPS: copula vs annual-average regression (b) CRPS: copula vs annual-average regression

Note: The θ_1 indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the annual-average horizon, i.e., one-year-, two-years-, and three-years-ahead. The x -axis in Panel (a) and (b) indicates the QW-CRPS (CRPS) of the copula relative to the annual-average regression approach, i.e., numbers smaller than one indicate a superior performance of the copula approach. The markers indicate the average score ratio across all Monte Carlo iterations. The horizontal lines around the markers indicate two standard deviations of the average score ratio frequency.

Figure 8: Monte Carlo results of tests for correct forecast density specifications: direct annual-average regression



Note: The θ_1 indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the annual-average horizon (in years). The x -axis denotes the rejection frequency of the [Rossi and Sekhposyan \(2019\)](#) test for correct specification of the predictive density. The markers indicate the average rejection frequency across all Monte Carlo iterations; the horizontal lines around the markers indicate two standard deviations of the rejection frequency.

substantially worse than that of the copula approach, in particular for $h = 1$ and in presence of strong persistence. The intuition is that, compared with the copula approach, the direct annual-average approach likely makes an inefficient use of the available short-term information. However, for $h > 1$ this inefficiency loses importance, and the predictive performance converges across the two approaches, because the underlying DGP is mean reverting and for longer horizons the initial information advantage in the copula approach becomes irrelevant. Finally, [Figure 8](#) reports the results for the PIT test for the direct annual-average approach only. The results point overall to a correct specification of density forecasts, with rejection frequencies even closer to the nominal size than those of the copula approach reported in [Figure 4](#). The findings hence suggest that the alternative direct annual-average approach may be inefficient but correctly specified (see [Gneiting et al., 2007](#), for other examples of such predictive densities).

6 Empirical applications

This section provides three distinct empirical exercises where the copula approach is evaluated in a forecasting environment with actual macroeconomic data. Note that we do not use real-time vintages in the forecasting exercises, so all “out-of-sample” forecasts are strictly speaking pseudo out-of-sample forecasts.

6.1 Large bivariate exercise

In this section, we provide the results of a large-scale forecasting exercise based on monthly data from FRED-MD ([McCracken and Ng, 2016](#)). In the construction of the forecasting environment, we closely follow [McCracken and McGillicuddy \(2019\)](#) and consider pairs, $(y_{1,t}, y_{2,t})$, of two monthly series randomly drawn from the dataset. We first compute up to 12-months-ahead density forecasts of cumulative growth at a monthly frequency $y_{i,t+h|t}$, with $h = 1, \dots, 12$ months and $i = 1, 2$. Density forecasts are obtained through horizon-specific monthly-frequency autoregressive distributed lag (ARDL) direct-multistep regressions, estimated via OLS at each forecast origin. Then, we apply our proposed copula approach to the monthly distributions to compute density forecasts of cumulative growth at a quarterly frequency $z_{i,t+h|t}$ up to four quarters ahead. In particular, for each for $i = \{1, 2\}$, we run the following regression to obtain the monthly density forecasts:

$$y_{i,t+h} = \alpha + \sum_{j=0}^{p-1} \beta_j y_{1,t-j} + \sum_{j=0}^{p-1} \gamma_j y_{2,t-j} + \varepsilon_{t+h} \quad (18)$$

with

$$y_{i,t+h} = \begin{cases} Y_{i,t+h} - Y_{i,t} & \text{if } Y_{i,t} \sim I(1) \\ Y_{i,t+h} - Y_{i,t} - h\Delta Y_{i,t} & \text{if } Y_{i,t} \sim I(2) \end{cases} \quad (19)$$

and $Y_{i,t+h}$ denotes the levels or log-levels. Forecasts $y_{i,t+h|t}$ are then used to get $Y_{i,t+h|t}$ and the quarterly average $Z_{i,t+h|t} = \frac{1}{3} \sum_{j=1}^3 Y_{i,t+h-j+1|t}$. Finally, the quarterly frequency transformation is obtained as:

$$z_{i,t+h|t} = \begin{cases} Z_{i,t+h|t} - Z_{i,t} & \text{if } Z_{i,t} \sim I(1) \\ Z_{i,t+h|t} - Z_{i,t} - h\Delta Z_{i,t} & \text{if } Z_{i,t} \sim I(2) \end{cases}$$

For the sake of simplicity, we only consider the first month of each quarter as a forecast origin. The number of lags in (18) is either fixed at $p = 4$ or selected through the Bayesian Information Criterion (BIC) among $p \in \{1, \dots, 12\}$.

We consider two samples: a full sample starting in 1974:M1 and a reduced sample starting in 1984:M1, covering only the Great Moderation period. Different samples may help to assess the impact of potential breaks in the persistence of the series over time. The sample ends in both cases in 2019:M12. The sample is partitioned into (i) an in-sample part, of size T_{is} , for the estimation of the parameters α, β_j and γ_j , (ii) a calibration part, of size $T_R = 60 - h$ and spanning from 1995:Mh to 2004:M12, for the estimation of the copula parameters, and (iii) a forecasting part, of size $T_{oos} = 60 - h_q$ and spanning from 2005:Mh to 2019:M12, for the out-of-sample evaluation in the target frequency. Note that as in Section 4, T_{oos} denotes the number of periods in the target frequency, i.e., in quarterly observations whereas T_{is} and T_R are denoted in the monthly frequency. The estimation is carried out using a rolling-window approach, updating at each forecast origin the model parameters, the PITs, and the copula parameters.

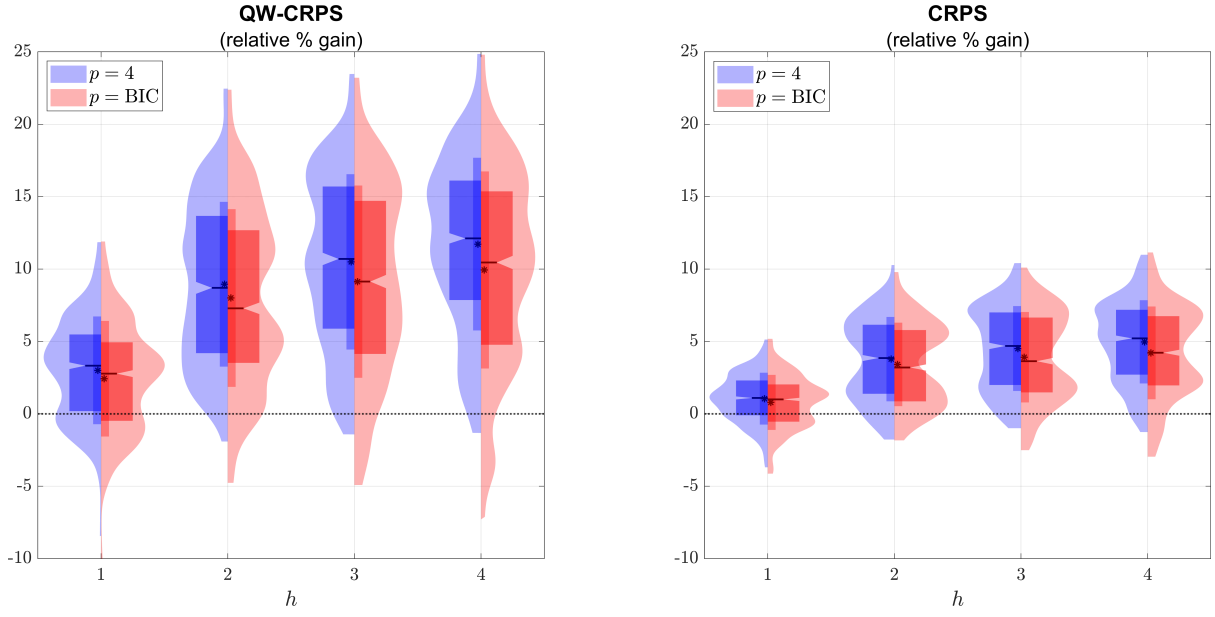
As a datasource we use the February 2023 vintage of FRED-MD. After dropping series not meeting some minimal conditions, the number of variables used in the bivariate exercise is 101, organized into 5 different groups as in Marcellino et al. (2006) and McCracken and McGillicuddy (2019): (1) income, output, sales, and capacity utilization; (2) employment and unemployment; (3) construction, inventories, and orders; (4) interest rates and asset prices; (5) nominal prices, wages, and money.⁸

Then, 650 random pairs of variables are selected from the database such that y_1 and y_2 come from distinct groups and an equal number of series pairs (y_1, y_2) comes from each of the 10 possible group pairings. For each permutation of the series, we compute density forecasts at the quarterly target frequency, based on the monthly density forecasts, using either the copula approach or the benchmark approach, and we compare their forecasting performance using their average QW-CRPS and CRPS over the out-of-sample. For each lag selection method and forecast horizon, the 650 random pairs forming the systems (y_1, y_2) and (y_2, y_1) provide a distribution of 1,300 average score ratios.

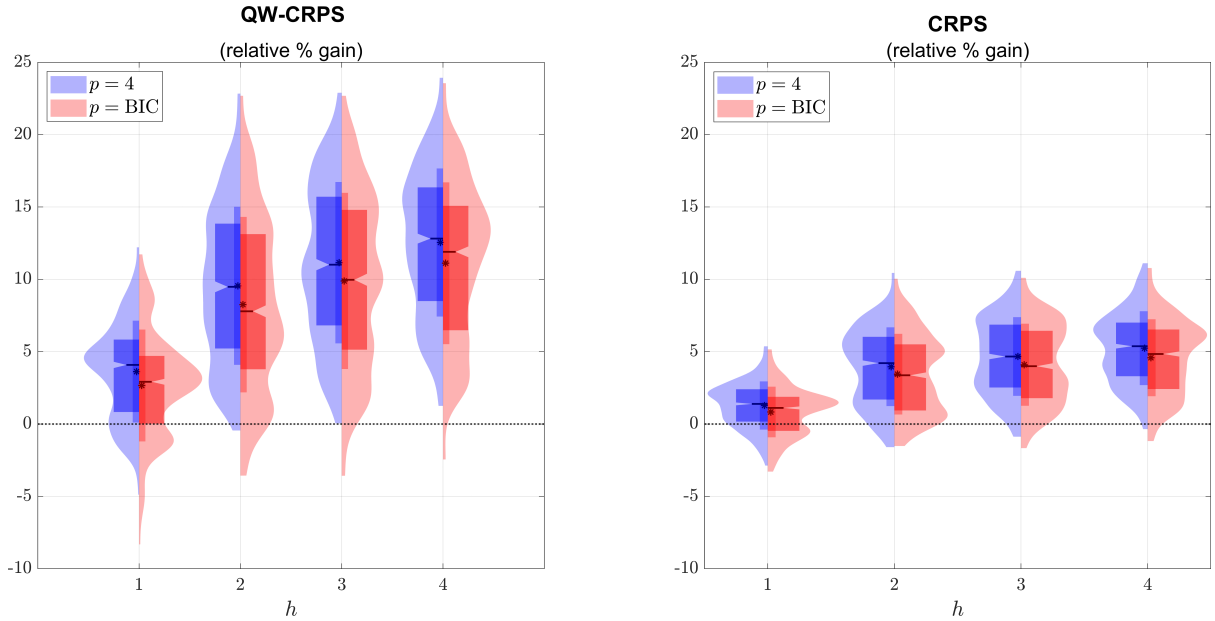
Figure 9 shows these distributions after computing the relative % gain of the copula approach with respect to the benchmark. The labels “Full sample” and “Great Moderation” indicate the samples used for the forecasting exercise. The copula approach provides, on average, a better forecasting performance for all horizons, with the predictive gain increasing with the horizon up to around 10% according to the QW-CRPS and around 5% according to the CRPS, irrespective of the sample used for the estimation and the lag length p . It is worth noting that these empirical predictive gains appear overall statistically positive, as the 95% confidence interval for the median and the inter-quartile range tend to exclude the zero value. To grasp the statistical significance of these results in a more formal way, we also compute for each pairing a test of unconditional equal predictive accuracy (Giacomini and White, 2006). The results are summarized in Figure 10, where we provide the overall rejection frequency of the null hypothesis of equal predictive accuracy and

⁸We exclude from the dataset three series starting after the 1970 (new orders for consumers goods, new orders for non-defense capital goods, trade weighted US dollar index), one series presenting missing values (consumer sentiment index), and one series switching to negative over the sample (non-borrowed reserves of depository institutions). In addition, we exclude 21 series that should be used in levels or log-levels, as we focus on series that can be expressed in first/second difference (or log-difference) in our application.

Figure 9: Forecast comparison of copula vs benchmark approach



(a) Full sample (1974:M1-2019:M12)

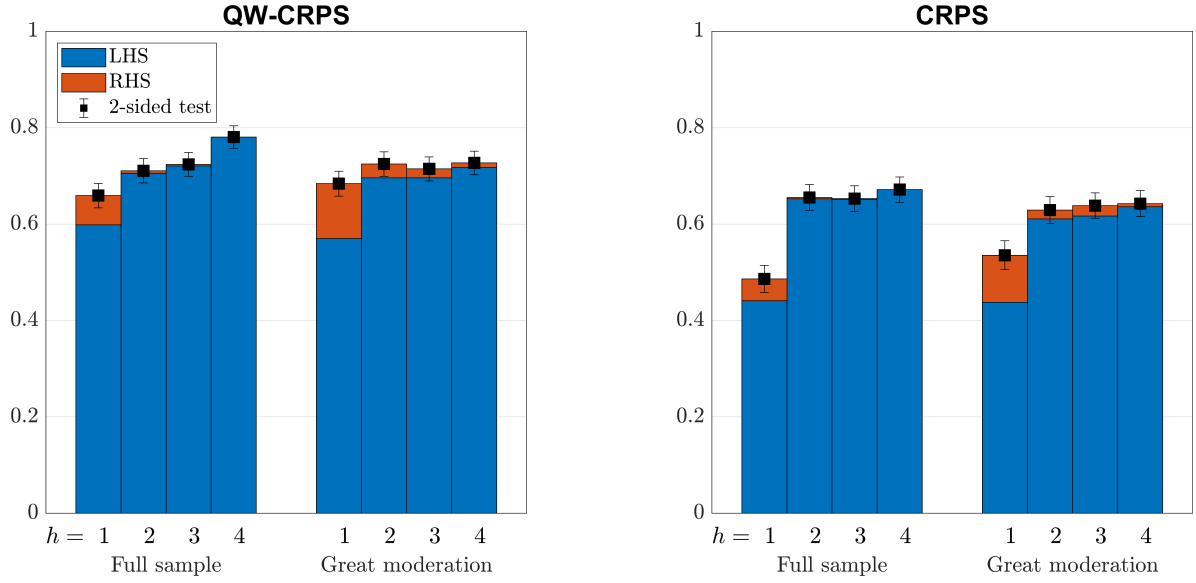


(b) Great Moderation sample (1984:M1-2019:M12)

Note: In particular, the curves in the figure depict the distribution of the relative % predictive gain in terms of the respective loss functions of the copula approach vs the benchmark approach over the 1,300 forecasting exercises. The boxes denote the inter-quartile range, the horizontal line-marker denotes the median, and the star-marker denotes the mean. Values larger than zero indicate positive gains for the copula approach. The x -axis denotes the forecast horizon in quarters.


its decomposition into the rejection frequency when the alternative is that the average loss of the copula is inferior to the average loss of the benchmark (the blue bar) and the respective opposite alternative (the red bar). The rejection frequencies appear relatively high across all samples and horizons, hovering around 0.6-0.8% for the QW-CRPS and 0.5-0.6% for the CRPS, and it can be

Figure 10: Rejection frequencies EPA test: copula vs benchmark approach



Note: The figure show the rejection frequency of the EPA tests (Giacomini and White, 2006) over the 1,300 forecasting exercises (the black square) and its decomposition into the rejection frequency when the average loss of the copula is smaller than the average loss of the benchmark (the blue bar) and when the average loss of the copula is larger than the average loss of the benchmark (the red bar). h denotes the forecast horizon in quarters. The labels “Full sample” and “Great Moderation” indicate the sample period used for the forecasting exercise. The nominal size is 5%. The error bars denote ± 2 standard errors computed through 1000 bootstrap replications.

Figure 11: Empirical distribution of $\det(\hat{R})$



(a) Full sample (1974:M1-2019:M12)

(b) Great Moderation sample (1984:M1-2019:M12)

Note: The figure shows distribution of $\det(\hat{R})$ over the 1,300 forecasting exercises. Smaller numbers of the determinant imply a larger degree of cross-horizon persistence in the forecasts.

largely attributed to the copula approach.

Figure 11 shows the empirical distribution of the determinant of the estimated copula parameter matrices \hat{R} . The determinant is a summary measure of the persistence of the forecasts

$y_{i,t+h}$ across the forecast horizons h (see [Section 5.1](#)). The results suggest that for a majority of the randomly selected bivariate systems, the persistence in $y_{i,t+h}$ is considerably large. This should explain the superior average forecasting performance of the copula approach relative to the benchmark approach depicted in [Figure 9](#). The large average cross-horizon persistence is, in part, driven by defining the monthly density forecasts as cumulative growth rates such that shocks are accumulated across the horizons h . However, predicting cumulative growth rates is not at all unusual and, in fact, the definition of the predictands in (19) is taken from [McCracken and McGillicuddy \(2019\)](#).

6.2 Inflation-at-Risk

In this section, we provide estimates of inflation at risk for year-on-year and annual-average inflation based on the U.S. Consumer Price Index. We use quantile regressions in combination with a Lasso, to select among many potential predictors, and produce monthly forecasts of the year-on-year inflation rate. Then, we construct predictive distributions of the (calendar year) annual-average inflation rates out of the year-on-year predictions. This forecasting environment aims at replicating the situation where a professional forecaster dispose of predictive densities for year-on-year inflation, but she is required to transform them into lower-frequency annual-average densities. This is a pretty common operational framework across institutions such as central banks, where the communication around the expected inflation environment is usually performed around annual-average rates, even though the monthly year-on-year rate represents the operational target.⁹ Importantly, the year-on-year and annual-average predictive densities need to be coherent, i.e., they should be based on the same predictors, model type, and the moments of the baseline frequency should be reflected in the moments of the new target frequency.

The underlying price index is the monthly and seasonally adjusted Consumer Price Index for all Urban Consumers¹⁰ (henceforth CPI) from 1960 to 2022. The baseline model predicts monthly year-on-year inflation, computed via the log-difference of t and $t - 12$, via a quantile regression model that uses a Lasso to select among a number of potential predictors ([Belloni and Chernozhukov, 2011](#)). The quantile regression coefficient vector β_τ , for the τ th quantile, is estimated by minimizing the following objective function:

$$\hat{\beta}_\tau = \underset{\tilde{\beta}_\tau \in \mathbb{R}^p}{\operatorname{argmin}} \sum_{t=t_0}^{T_{\text{is}}+t_0-1} \rho_\tau(y_{t+h} - \mathbf{x}'_t \tilde{\beta}_\tau) + \frac{\lambda \sqrt{\tau(1-\tau)}}{T_{\text{is}}} \sum_{j=1}^p \hat{\sigma}_j^2 |\tilde{\beta}_{j,\tau}|, \quad (20)$$

where y_{t+h} denotes the monthly year-on-year inflation rate, \mathbf{x}_t denotes a $p \times 1$ vector of predictors (including a constant), $\rho_\tau(z) = (\tau - \mathbb{1}\{z \leq 0\})z$, λ is a hyperparameter that determines the degree of penalization (set here as recommended by [Belloni and Chernozhukov, 2011](#)), p is the number of predictors, $h = 1, \dots, 12$ is the forecast horizon, T_{is} is the in-sample estimation, $t_0 = 1, \dots, 565$, denotes the initial observation which changes due to a rolling window estimation scheme, and $\hat{\sigma}_j^2 = \sum_{t=t_0}^{T_{\text{is}}+t_0-1} x_{j,t}^2$.

Note that because not all predictors are available over the entire range of the sample, not

⁹Note that, in fact, both the ECB and the Federal Reserve publish annual-average forecasts as part of their institutional projection exercises.

¹⁰Results are robust to using the non-seasonally adjusted CPI for all items, which has the mnemonic USACPI-ALLMINMEI on FRED.

all predictors enter the model in all forecast origins (see [Section A.2](#) for details on the data series). Therefore, the predictor vector x_t contains at most 23 exogenous predictors, similar to the predictors used in [Korobilis \(2017\)](#), as well as two lags of the endogenous variable and the constant.¹¹

We use the model in (20) to produce out-of-sample forecasts for horizons $h = 1, \dots, 12$ months ahead, with the first forecast origin being 1974:M12 and the last forecast origin being 2021:M12. At each forecast origin, the quantile regression is re-estimated over a rolling window of 15 years of monthly data, that is, $T_{is} = 180$. For example, for the forecast origin 1974:M12, the first observation used in the model estimation is 1960:M1. We let the model produce quantile predictions for $\tau \in [0.01, 0.02, \dots, 0.98, 0.99]$ and we calculate the predictive CDF based on linearly interpolating between adjacent quantiles (see [Mitchell et al., 2024](#), for a similar approach; details on the interpolation are provided in [Section A.2](#)).¹² Finally, the generated year-on-year predictive densities are then transformed into annual-average inflation densities by drawing from the joint distribution of the year-on-year forecasts using either the copula or the benchmark approach. In order to approximate the annual-average, we first compute month-on-month rates from the year-on-year forecasts, and then we transform the month-on-month predictions into annual-averages using the approximation of [Mariano and Murasawa \(2003\)](#).

For the estimation of the copula parameter matrix R , we use about 10 years of data over a rolling window starting with the observations forecast origins 1974:M12 to 1984:M12, that is, $T_R = 121$. More precisely, we start by evaluating the empirical PITs of the predictive distributions of monthly year-on-year inflation using observations 1975:M h to 1985:M h for forecast horizon $h = 1, \dots, 12$. We then estimate the copula parameter based on the empirical PITs and we use the estimated copula parameters in combination with the forecasts for 1985:M h , $h = 1, \dots, 12$, to construct an annual-average inflation predictive distribution for 1986. Next, we use the predictive distributions for year-on-year inflation from 1976:M h to 1986:M h for horizon $h = 1, \dots, 12$, and re-do the steps to construct the annual-average inflation predictive distribution for 1987. We then repeat this until we have an annual-average forecast for each year from 1986 to 2022. In addition, we construct annual-average forecasts as described above but using the benchmark approach, which is identical to setting all off-diagonal elements of R equal to zero. Therefore, the out-of-sample evaluation period for annual-average inflation forecasts spans from 1986 to 2022, leading to 37 out-of-sample (calendar year) annual-average predictive distributions, that is, $T_{oos} = 37$.


The results can be summarized as follows. Regarding the performance at the tails of the predictive distributions, the QW-CRPS ratio of the copula approach relative to the benchmark is 0.79, i.e., the copula approach provides a substantial predictive gain of about 21%. Similarly, the ratio of the QS(10%) and QS(90%) is respectively 0.72 and 0.85, i.e., the copula approach provides a relative gain of 28% and 15% for predicting the risks of low and high inflation, respectively. When evaluating the entire density forecasts, the CRPS ratio of the copula approach relative to the benchmark is 0.91, i.e., the copula approach provides a relative predictive gain of 9%. All these results are statistically significant at the 1% level according to the test for unconditional equal predictive ability of [Giacomini and White \(2006\)](#).

¹¹Results are robust to increasing the number of lags to six.

¹²Quantile crossing is dealt with by sorting the predictive quantiles.

Figure 12 illustrates the superior predictive ability of the copula approach. Panel (a) shows the predictive density of the annual-average inflation for 2001 and panel (b) shows the predictive density of the annual-average inflation for 2011. In both cases, the benchmark approach assigns ex-ante zero-probability to the realizations, whereas the copula-based density forecasts can better predict tail-risks. In particular, the latter displays fatter tails as a result of the high correlation between adjacent months' year-on-year growth rates. It is also worth noting that the predictive distributions of the copula-based approach exhibit notably stronger asymmetry.

Figure 12: annual-average predictive densities for inflation



Note: The figure shows the results for the annual-average inflation predictive density for 2001, panel (a), and for 2011, panel (b), for both the copula approach (solid line) as well as the benchmark approach (dashed line), alongside the realization (dotted line).

6.3 Growth-at-Risk

In a recent influential contribution, [Adrian et al. \(2019\)](#) use quantile regressions to produce *direct* density forecasts of quarterly quarter-on-quarter U.S. real GDP growth and show that financial conditions, captured by the National Financial Conditions Index (NFCI), are an important predictor for downside risks to GDP growth. We apply our methodology to transform their quarter-on-quarter density forecasts into annual-average densities.

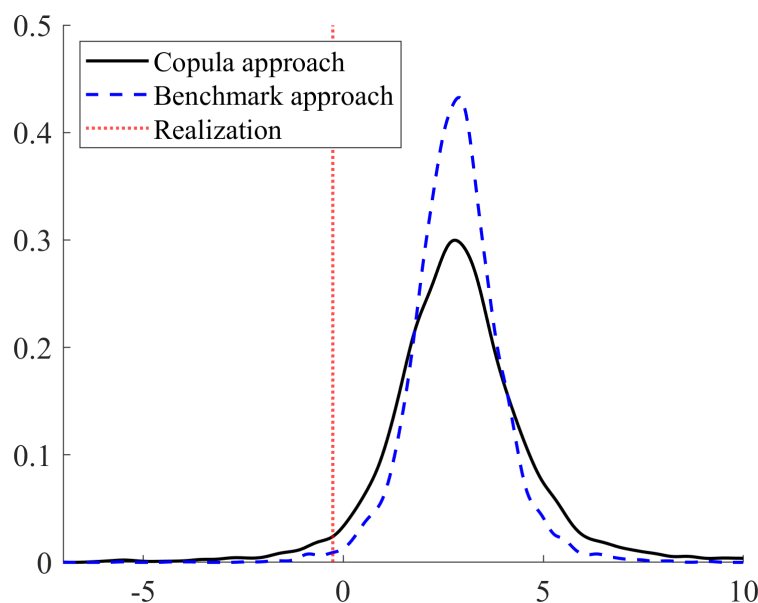
The total data sample of [Adrian et al. \(2019\)](#) ranges from 1973:Q1 to 2015:Q4 and, starting with forecast origin 1993:Q1, the authors produce (pseudo) out-of-sample forecasts using quantile regressions for up to four quarters ahead.¹³ Thus, using their exact specification leaves us with a series of out-of-sample forecasts from 1993:Q h to 2015:Q4, for forecast horizons $h = 1, \dots, 4$.

To construct annual-average predictive distributions, we proceed as follows. We first evaluate the predictive distributions for the forecast targets 1993:Q h to 2001:Q h , with $h = 1, \dots, 4$, and we obtain a series of 32 empirical PITs for each forecast horizon $h = 1, \dots, 4$. From the PITs, we calculate the copula parameter as described in [Section 3](#). We then compute the annual-average predictive distribution for 2002 from the quarter-on-quarter densities for horizon $h = 1, \dots, 4$,

¹³The predictors of their preferred specification are one lag of both the endogenous variable (the real GDP growth) and the NFCI. For details on the exact specification and variable definitions, see [Adrian et al. \(2019\)](#).

with origin 2001:Q4 (and forecasting targets 2002:Q1 to 2002:Q4). Then, we move four quarters ahead and repeat the exercise to construct the annual-average predictive distribution for 2003. In total, we repeat this algorithm until we have a series of 14 annual-average predictive distributions spanning from 2002 to 2015. The obtained annual-average predictive distributions are hence based on the original quarterly predictive distributions of [Adrian et al. \(2019\)](#) and takes into account the serial correlation across the quarterly growth rates. The benchmark annual-average predictive distributions are constructed similarly.

Figure 13: annual-average predictive densities for 2008



Note: The figure shows the results for the annual-average predictive density of 2008 for both the copula approach (solid line) as well as the benchmark approach (dashed line), alongside the realization (dotted line).

[Figure 13](#) shows the results for the annual-average predictive density of 2008 for both the copula approach (solid line) and the benchmark approach (dashed line), alongside the realized annual-average growth rate (dotted line). The copula approach leads to annual-average forecasts with larger tails due to the positive correlation between the quarterly growth rates. Indeed, the rank correlation, i.e., the elements in the Gaussian copula correlation matrix, is around 50% to 60% for adjacent quarters, depending on the forecast horizon. Importantly, the larger tails of the copula approach help to assess the downside risk to the annual-average real GDP growth for 2008, i.e., during the onset of the financial crisis. This underlines the usefulness of our approach and provides anecdotal evidence of incorrect risk assessment if the serial correlation between the original predictive distributions is not taken into account when constructing multi-horizon objects.

Turning to the full evaluation sample, the QW-CRPS and CRPS ratios point to a predictive for the copula approach with respect to the benchmark of about 2.3% and 2.5%, respectively. However, the QS(10%) ratio points to a substantial predictive gain of about 25% for the copula approach. In other words, even though the sample is too short for formal testing procedures, the empirical analysis suggests that the copula approach may provides a considerable improvement for predicting lower tail risks on activity, which are carefully and closely monitored by policy-makers.

7 Conclusion

In this work, we propose a method to combine *direct* forecasts to obtain new predictive objects that are function of several horizons. The approach is useful in a situation where the forecaster has a set of *direct* forecasts available and has to use the same set of *direct* predictive densities to construct a new predictive object; for instance, if the forecasters already has a set of *direct* quarter-on-quarter growth rates but also needs annual-average predictions. These type of situations typically arise, but are not limited to, in institutions where the forecasting process is rigid.

In a Monte Carlo exercise, we show that our methodology outperforms the benchmark approach in terms of forecasting performance whenever the serial correlation across different forecasting horizons is close to zero. In terms of the absolute forecasting performance the copula approach provides density forecasts that largely pass a correct specification test based on evaluating the uniformity of the PIT, whereas the benchmark approach fails to pass this test whenever the serial correlation is not close to zero.

In a first empirical application, we investigate in a large-scale forecasting exercise, based on monthly data from FRED-MD ([McCracken and Ng, 2016](#)), the performance of our methodology for a large number of outcome variable and predictor combinations. In this exercise, we transform month-on-month predictive densities to quarter-on-quarter density forecasts through our proposed copula approach and results show that copula approach outperforms the benchmark approach for the majority of outcome variable and predictor combinations.

In the second empirical application, we show the usefulness of the approach by transforming year-on-year predictive densities for inflation into annual-average predictive densities. The copula approach significantly outperforms the benchmark approach both in terms of the CRPS and quantile tick loss evaluation.

In the third empirical application, we transform the quarter-on-quarter *direct* forecasts of U.S. real GDP growth of [Adrian et al. \(2019\)](#) into annual-average forecasts, and provide anecdotal evidence that the copula approach provides a better forecasts of the growth at risk during the Great Recession period.

References

- Adrian, Tobias, Nina Boyarchenko and Domenico Giannone. (2019). "Vulnerable growth". *American Economic Review*, 109(4), pp. 1263–1289.
<https://doi.org/10.1257/aer.20161923>
- Azzalini, Adelchi, and Antonella 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(2), pp. 367–389.
<https://doi.org/10.1111/1467-9868.00391>
- Belloni, Alexandre, and Victor Chernozhukov. (2011). "L1-penalized quantile regression in high-dimensional sparse models". *The Annals of Statistics*, 39(1), pp. 82–130.
<https://doi.org/10.1214/10-AOS827>
- Clark, Todd E., Michael W. McCracken and Elmar Mertens. (2020). "Modeling Time-Varying Uncertainty of Multiple-Horizon Forecast Errors". *The Review of Economics and Statistics*, 102(1), pp. 17–33.
https://doi.org/10.1162/rest_a_00809
- Fan, Yanqin, and Andrew J. Patton. (2014). "Copulas in econometrics". *Annual Review of Economics*, 6, pp. 179–200.
<https://doi.org/10.1146/annurev-economics-080213-041221>
- Ferrara, Laurent, Matteo Mogliani and Jean-Guillaume Sahuc. (2022). "High-frequency monitoring of growth at risk". *International Journal of Forecasting*, 38(2), pp. 582–595.
<https://doi.org/10.1016/j.ijforecast.2021.06.010>
- Ganics, Gergely, Barbara Rossi and Tatevik Sekhposyan. (2024). "From fixed-event to fixed-horizon density forecasts: Obtaining measures of multihorizon uncertainty from survey density forecasts". *Journal of Money, Credit and Banking*, 56(7), pp. 1675–1704.
<https://doi.org/10.1111/jmcb.13105>
- Giacomini, Raffaella, and Ivana Komunjer. (2005). "Evaluation and Combination of Conditional Quantile Forecasts". *Journal of Business & Economic Statistics*, 23(4), pp. 416–431.
<https://doi.org/10.1198/073500105000000018>
- Giacomini, Raffaella, and Halbert White. (2006). "Tests of Conditional Predictive Ability". *Econometrica*, 74(6), pp. 1545–1578.
<https://doi.org/10.1111/j.1468-0262.2006.00718.x>
- Gneiting, Tilmann, Fadoua Balabdaoui and Adrian E. Raftery. (2007). "Probabilistic forecasts, calibration and sharpness". *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2), pp. 243–268.
<https://doi.org/10.1111/j.1467-9868.2007.00587.x>
- Gneiting, Tilmann, and Roopesh Ranjan. (2011). "Comparing density forecasts using threshold- and quantile-weighted scoring rules". *Journal of Business & Economic Statistics*, 29(3), pp. 411–422.
<https://doi.org/10.1198/jbes.2010.08110>

- Grothe, Oliver, Fabian Kächele and Fabian Krüger. (2023). "From point forecasts to multivariate probabilistic forecasts: The Schaake shuffle for day-ahead electricity price forecasting". *Energy Economics*, 120 (106602).
<https://doi.org/10.1016/j.eneco.2023.106602>
- Hafner, Christian M., and Hans Manner. (2012). "Dynamic stochastic copula models: Estimation, inference and applications". *Journal of Applied Econometrics*, 27, pp. 269–295.
<https://doi.org/10.1002/jae.1197>
- Korobilis, Dimitris. (2017). "Quantile regression forecasts of inflation under model uncertainty". *International Journal of Forecasting*, 33(1), pp. 11–20.
<https://doi.org/10.1016/j.ijforecast.2016.07.005>
- Marcellino, Massimiliano, James H. Stock and Mark W. Watson. (2006). "A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series". *Journal of Econometrics*, 135(1-2), pp. 499–526.
<https://doi.org/10.1016/j.jeconom.2005.07.020>
- Mariano, Roberto S., and Yasutomo Murasawa. (2003). "A new coincident index of business cycles based on monthly and quarterly series". *Journal of Applied Econometrics*, 18(4), pp. 427–443.
<https://doi.org/10.1002/jae.695>
- McCracken, Michael W., and Joseph T. McGillicuddy. (2019). "An empirical investigation of direct and iterated multistep conditional forecasts". *Journal of Applied Econometrics*, 34(2), pp. 181–204.
<https://doi.org/10.1002/jae.2668>
- McCracken, Michael W., and Serena Ng. (2016). "FRED-MD: A Monthly Database for Macroeconomic Research". *Journal of Business & Economic Statistics*, 34(4), pp. 574–589.
<https://doi.org/10.1080/07350015.2015.1086655>
- Mitchell, James, Aubrey Poon and Dan Zhu. (2024). "Constructing density forecasts from quantile regressions: Multimodality in macrofinancial dynamics". *Journal of Applied Econometrics*, 39(5), pp. 790–812.
<https://doi.org/10.1002/jae.3049>
- Nelsen, Roger B. (2006). *An Introduction to Copulas*. Springer.
<https://doi.org/10.1007/0-387-28678-0>
- Patton, Andrew J. (2006). "Modelling asymmetric exchange rate dependence". *International Economic Review*, 47, pp. 527–556.
<https://doi.org/10.1111/j.1468-2354.2006.00387.x>
- Rossi, Barbara, and Tatevik Sekhposyan. (2019). "Alternative tests for correct specification of conditional predictive densities". *Journal of Econometrics*, 208(2), pp. 638–657.
<https://doi.org/10.1016/j.jeconom.2018.07.008>
- Sklar, Abe. (1959). "Fonctions de répartition à n dimensions et leurs marges". *Publications de l'Institut de Statistique de l'Université de Paris*, 8, pp. 229–231.
<https://doi.org/10.2139/ssrn.4198458>


Smith, Michael S., and Shaun P. Vahey. (2016). "Asymmetric forecast densities for U.S. macroeconomic variables from a Gaussian copula model of cross-sectional and serial dependence". *Journal of Business & Economic Statistics*, 34(3), pp. 416–434.
<https://doi.org/10.1080/07350015.2015.1044533>

Appendix

Appendix A Additional results and data

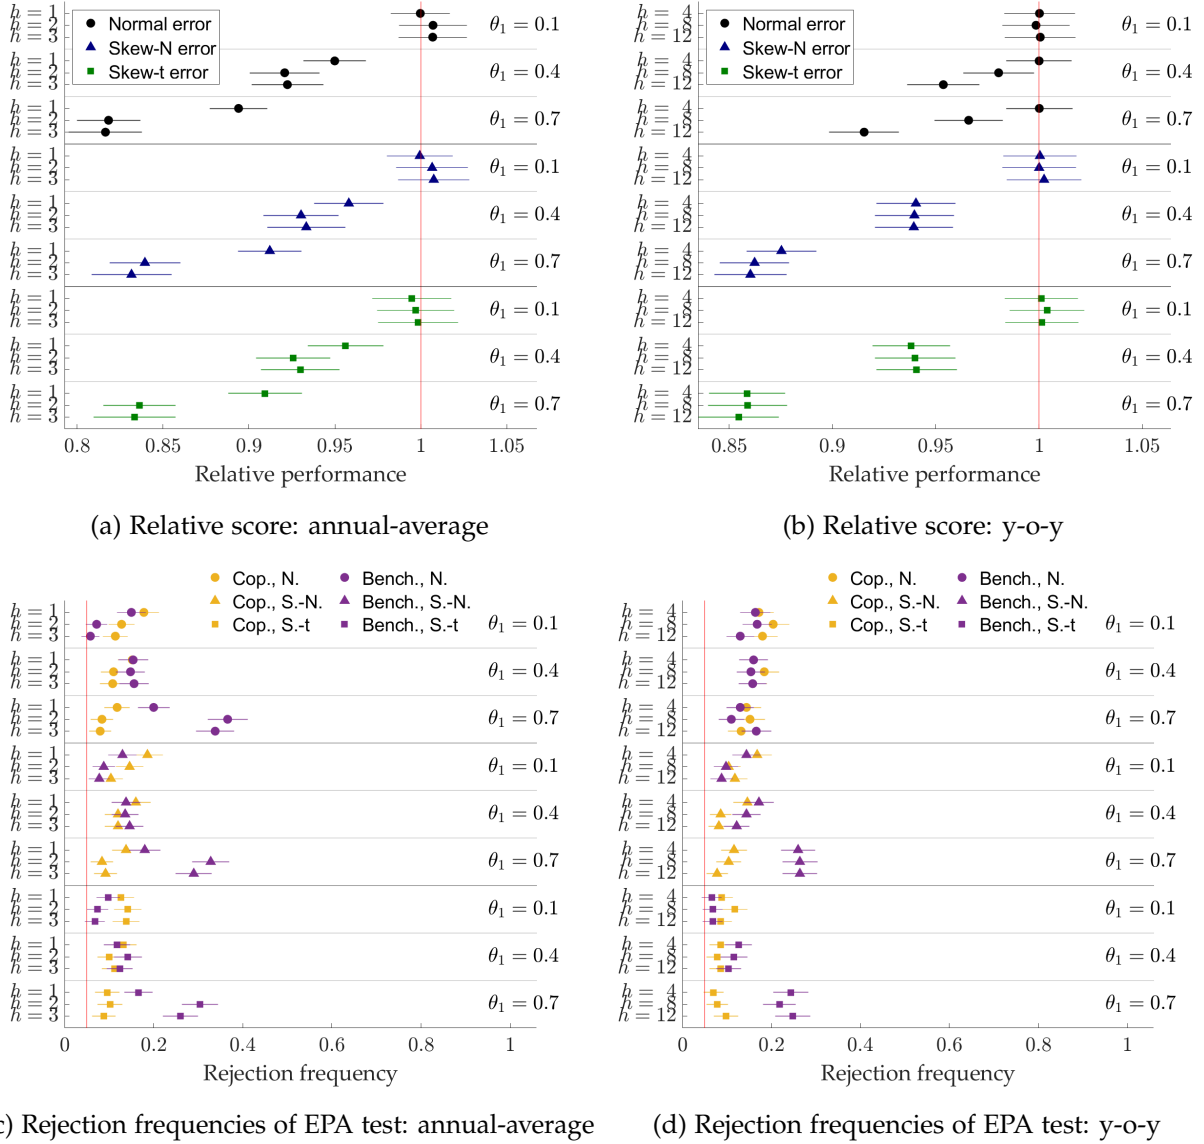
A.1 Additional simulation results

Figure A.1: Scores for density forecasts with negative autoregressive coefficients: "dependence-attentive" *vs* "dependence-inattentive" approach



Note: QW-CRPS denotes the quantile weighted versions of the continuous ranked probability score, with emphasis on the tails. CRPS denotes the continuous ranked probability score. QS(10%) denotes the quantile score at the 10% quantile. QW-CRPS, CRPS, and QS are expressed in relative % gain of the "dependence-attentive" forecaster with respect to the "dependence-inattentive" forecaster.

Figure A.2: Monte Carlo results for QS(10%): relative scores and EPA tests



Note: The QS(10%) denotes the quantile score at the 10% quantile. The θ indicates the autoregressive parameter of Y_t in the DGP. The y -axis label h denotes the annual-average horizon, i.e., one-year-, two-years-, and three-years-ahead. The x -axis in Panel (a) and (b) indicates the QS(10%) of the copula relative to the benchmark approach, i.e., numbers smaller than one indicate a superior performance of the copula approach. In Panel (c) and (d), the x -axis denotes the rejection frequency of the null hypothesis of a [Giacomini and White \(2006\)](#) test of unconditional equal predictive ability using the QS (10%). The nominal size is 5%. N., S.-N., and S.-t indicate the Normal, Skew-Normal, and Skew-t distribution of the error terms in the DGP. Standard errors of the tests were computed using a HAC estimator with bandwidth $= h_A - 1$.

A.2 Inflation at Risk data

[Table A.1](#) shows the predictors used in [Section 6.2](#). All data is seasonally adjusted were applicable. The transformation codes imply the following: 1 — no transformation; 4 — $\log(x_t)$; 5 — $100[\log(x_t) - \log(x_{t-12})]$

Table A.1: Predictors for inflation at risk forecasts

Variable	Transformation	Mnemonic
Aggregate weekly hours	5	AWHI
Commercial + industrial loans	5	BUSLOANS
Labor force participation rate	4	CIVPART
Consumer loans	5	CONSUMER
CPI All Urban Consumers	5	CPIAUCSL
Canadian dollar to U.S. exchange rate	5	EXCAUS
Japanese Yen to U.S. exchange rate	5	EXJPUS
British Pound to U.S. exchange rate	5	EXUSUK
Federal Funds Target	1	FEDFUNDS
Private Housing starts	4	HOUST
New Family houses sold	4	HSN1F
Industrial production	5	INDPRO
Fixed-rate 30-year mortgage rate	1	MORTG
Bank prime loan rate	1	MPRIME
Motor Vehicle assemblies	1	MVATOTASSS
Total non-farm employees	5	PAYEMS
Real estate loans	5	REALLN
Capacity utilization	4	TCU
Number unemployed for 15 weeks & over	4	UEMP15OV
Number unemployed for less than 5 weeks	4	UEMPLT5
University of Michigan: consumer sentiment	1	UMCSENT
Unemployment rate	1	UNRATE
WTI spot price	5	WTISPLC

Note: All data was downloaded from the database FRED of the St. Louis Federal Reserve Bank.

A.3 Inflation at Risk quantile interpolation

Let $F_{t+h|t}(y_{t+h})$ denotes the predictive cumulative distribution function in t and h horizons ahead, evaluated at y_{t+h} . Then, given the set $\{Q_{t+h|t}(\tau_i)\}_{i=1}^{99}$ of predictive quantiles, we compute $F_{t+h|t}(y_{t+h})$ as follows:

$$F_{t+h|t}(y_{t+h}) = \tau_i + \frac{\tau_{i+1} - \tau_i}{Q_{t+h|t}(\tau_{i+1}) - Q_{t+h|t}(\tau_i)}(y_{t+h} - Q_{t+h|t}(\tau_i)), \quad (21)$$

where $Q_{t+h|t}$ denotes predictive value of quantile τ_i , and τ_i and τ_{i+1} are such that $y_{t+h} \in [Q_{t+h|t}(\tau_i), Q_{t+h|t}(\tau_{i+1})]$. For values of $y_{t+h} < Q_{t+h|t}(\tau_1)$ and $y_{t+h} > Q_{t+h|t}(\tau_{99})$, we approximate the slope as $\frac{\tau_2 - \tau_1}{Q_{t+h|t}(\tau_2) - Q_{t+h|t}(\tau_1)}$ and $\frac{\tau_{99} - \tau_{98}}{Q_{t+h|t}(\tau_{99}) - Q_{t+h|t}(\tau_{98})}$ and the distance as $(Q_{t+h|t}(\tau_1) - y_{t+h})$ and $(y_{t+h} - Q_{t+h|t}(\tau_{99}))$.

Similarly, we sample from the distribution given by the conditional quantiles $Q_{t+h|t}(\tau_i)$ as follows. Let u_j denote the j -th draw from the uniform distribution and $u_j \in [\tau_i^j, \tau_{i+1}^j]$. Then,

$$y_{t+h|t}^j = Q_{t+h|t}(\tau_i^j) + \frac{Q_{t+h|t}(\tau_{i+1}^j) - Q_{t+h|t}(\tau_i^j)}{\tau_{i+1}^j - \tau_i^j} (u_j - \tau_i^j), \quad (22)$$

where $y_{t+h|t}^j$ denotes draw j of the predictive distribution for y_{t+h} , conditional on information in t . The two endpoints are treated analogously to the procedure described for equation (21).