



# Penalized Averaging of Quantile Forecasts from GARCH Models with Many Exogenous Predictors

Jan G. De Gooijer<sup>1</sup>

Accepted: 20 June 2022 / Published online: 13 July 2022  
© The Author(s) 2022

## Abstract

This study explores the multi-step ahead forecasting performance of a so-called hybrid conditional quantile method, which combines relevant conditional quantile forecasts from parametric and semiparametric methods. The focus is on lower (left) and upper (right) tail quantiles of the conditional distribution of the response variable. First, we evaluate and compare out-of-sample conditional quantile forecasts obtained from a hybrid method and from five non-hybrid methods, employing a large data set of exogenous predictors generated by various GARCH model specifications. Second, we compare the accuracy of these methods by calculating conditional quantile forecasts for the risk premium of the monthly S&P 500 index, using a data set of macroeconomic predictors. Monte Carlo and empirical application results indicate that the hybrid forecasting provides more accurate quantile forecasts than non-hybrid methods. The success of the hybrid method is most prominent when compared with results obtained by a simple equal-weighted combination of quantile forecasts.

**Keywords** Encompassing · Hybridization · Penalized quantile averaging · Quantile forecasting · Tick loss function

**JEL Classification** C14 · C52 · C53 · C55

---

✉ Jan G. De Gooijer  
j.g.degooijer@uva.nl

<sup>1</sup> Amsterdam School of Economics, University of Amsterdam, PO Box 15867,  
1001 NJ Amsterdam, The Netherlands

## 1 Introduction

Combining forecasts obtained by a hybrid approach has long been known to improve forecast accuracy. Because each model whose forecasts are to be combined may consider different predictors and make different assumptions about the underlying data generating process (DGP) and distributions, averaging the individual forecasts broadens the information embedded and may offset individual model biases as well; see, e.g., Zhang (2003) for additional reasons. In a one-dimensional (univariate) setting the hybridization of forecasts obtained from time series models/methods has been well researched; see, e.g., Aijaz and Agarwal (2020), Aydin and Isci Güneri (2015), Pai and Lin (2005), and Valenzuela et al. (2008) among others.

Recently, however, the enhanced availability of large databases with many time series variables as potential predictors has stimulated interest in high-dimensional forecasting. Excellent reviews of the state-of-the-art in economics and finance are given by Fan et al. (2011) and Lee (2011). Furthermore, the introductory section of the paper by Uematsu and Tanaka (2019) provides an extensive survey of the current literature on high-dimensional forecasting and variable selection. Along this direction, conditional quantile averaging procedures in conjunction with dimension reduction methods have been considered. Examples include quantile forecasting the S&P 500 equity risk premium (see, e.g., Konzen and Ziegelmann, 2016; Lima and Meng, 2017; Meligkotsidou et al., 2014; De Gooijer and Zerom, 2020) inflation forecasting (Garcia et al., 2017; De Gooijer and Zerom, 2019), quantile forecasting of macroeconomic time series (Manzan, 2015; Jiang et al., 2018), and realized volatility forecasting (Meligkotsidou et al., 2019).

Most of these studies are limited by the assumption that the high-dimensional data set comes from a linear DGP, and forecasts are obtained from finite-dimensional parametric models. One exception is the study by De Gooijer and Zerom (2020). Using a large data set of predictors, involving both macroeconomic predictors and technical indicators, these authors showed that combining quantile forecasts from parametric and semiparametric methods (called hybrid quantile averaging) can be useful in practice. Semiparametric models are infinite-dimensional. As a result, the quantile forecasts are less prone to model misspecification as may happen with forecasts obtained from parametric models. In terms of quantile forecast performance the hybrid method works well in identifying relevant predictors, and more importantly results in improved combined one-step ahead forecasts over alternative (no hybridization) conditional quantile methods.

While the empirical findings by De Gooijer and Zerom (2020) are interesting, they are sample-specific which makes it difficult to generalize the results to novel situations. In addition, the focus on one-step ahead out-of-sample prediction is somewhat restrictive. Multi-step ahead out-of-sample quantile forecasting results can provide more insight in the relative performance of the hybrid quantile averaging method over a longer time period. Unfortunately, a general theoretical comparison of quantile forecasts obtained from hybrid and non-hybrid methods is not feasible due to complicated interactions of nonlinear parameter estimation methods, sparse modelling, and correlated forecasts. Indeed, further insights regarding the

performance of the hybrid quantile averaging method can only be obtained via a Monte Carlo simulation study. In the first half of this paper, we provide such a study.

In the second half of the paper, we evaluate the out-of-sample multi-step ahead forecasting performance of the hybrid conditional quantile method and five alternative, non-hybrid, forecasting methods via an empirical application. More specifically, we report out-of-sample conditional quantile forecasts for the risk premium of the monthly S&P 500 index using a large data set of macroeconomic predictors. A simple equal-weighted combination of parametric and semiparametric conditional quantile forecasts is adopted as a benchmark. Our main finding is that that hybridization can be an effective way to improve quantile forecasts as compared to non-hybrid methods.

The rest of this paper unfolds as follows. First, for ease of reference, Sect. 2 summarizes the main features of the semiparametric and hybrid conditional quantile averaging methods. Section 3 describes six quantile forecasting methods used in the Monte Carlo experiment. Section 4 introduces the large data set of exogenous predictors via semiparametric and parametric GARCH model specifications. This section also discusses the evaluation of quantile forecasts. Section 5 contains the simulation results, while Sect. 6 presents the empirical results. Finally, Sect. 7 contains some concluding remarks.

## 2 Conditional Quantile Averaging

### 2.1 Some Notation

To simplify presentation, we only discuss the case of making  $h = 1$  step ahead quantile forecasts, where  $h$  denotes the forecast horizon. The extension to multi-horizon ( $h > 1$ ) conditional forecasting is straightforward. Let  $\{Y_t\}_{t=1}^n$  be a set of observations obtained from a strictly stationary time series process  $\{Y_t, t \in \mathbb{Z}\}$  that depends on  $q_y \geq 1$  past values of  $Y_t$ , and on a  $q_z$ -dimensional vector  $\mathbf{Z}_t = (Z_{1,t}, \dots, Z_{q_z,t})^T$  that consists of exogenous, possibly lagged, stationary time series. Let  $\mathbf{X}_t = (\mathbf{Y}_{t-1}^T, \mathbf{Z}_t^T)^T \in \mathbb{R}^q$  where  $\mathbf{Y}_{t-1} = (Y_{t-1}, \dots, Y_{t-q_y})^T$ , and  $q = q_y + q_z$ . Given the observed data set  $\{(\mathbf{X}_t, Y_t)\}_{t=1}^n$ , the one-step ahead out-of-sample  $\tau$ th ( $0 < \tau < 1$ ) conditional quantile of the unobserved random variable  $Y_{n+1}$  given  $\mathbf{X}_t = \mathbf{X}_n$  will be denoted by  $Q_{Y_{n+1}}(\tau|\mathbf{X}_n)$ .

For convenience of later analysis, we define the associated process  $\{(\mathbf{X}_t, Y_t^*)\} \in \mathbb{R}^q \times \mathbb{R}$  where the components of the predictor vector  $\mathbf{X}_t = (X_{1,t}, \dots, X_{q,t})^T$  are given by

$$X_{j,t} = \begin{cases} Z_{j,t}, & j = 1, \dots, q_z, \\ Y_{t+(j-q_z-1)}, & j = q_z + 1, \dots, q, \end{cases}$$

and

$$Y_t^* = Y_{t+q_y}, \quad t = 1, \dots, N, \quad (N \equiv n - q_y).$$

Hence, we obtain  $Q_{Y_{n+1}}(\tau|\mathbf{X}_n)$  directly via  $Q_{Y_t^*}(\tau|\mathbf{X}_{n-q_y+1}) = \inf\{y : F_{Y_t^*}(y|\mathbf{X}_{n-q_y+1}) \geq \tau\}$ , where  $F_{Y_t^*}(\cdot|\mathbf{x})$  is the conditional distribution of  $Y_t^*$  given  $\mathbf{X}_t = \mathbf{x}$ . Furthermore, for a given quantile level  $\tau$ , we define

$$\mathcal{M}_\tau = \{j : Q_{Y_t^*}(\tau|\mathbf{x}) \text{ functionally depends on } x_j\} \tag{1}$$

as the set of relevant predictors that truly influence  $Q_{Y_t^*}(\tau|\mathbf{x})$ . We assume that sparsity of the predictors exists, i.e.,  $|\mathcal{M}_\tau| \ll q$  which means that only a small subset of predictors influences  $Q_{Y_t^*}(\tau|\mathbf{x})$ .

### 2.2 Semiparametric (SP) Conditional Quantile Averaging

For  $q$  large (many predictors), the semiparametric (SP) quantile prediction averaging method consists of the following three steps.

1. Using gradient boosting, obtain  $\hat{\theta}_{0,j}(\tau|x_j)$  as a nonparametric estimate of a marginal conditional quantile, defined as

$$\theta_j(\tau|x_j) = \inf\{y : F_{Y_t^*}(y|X_{j,t} = x_j) \geq \tau\}, \quad (j = 1, \dots, q).$$

Next, approximate  $Q_{Y_t^*}(\tau|\mathbf{x})$  by a linear combination of  $\hat{\theta}_{0,j}(\tau|x_j)$ 's. That is

$$\tilde{Q}_{Y_t^*}(\tau|\mathbf{x}) = \gamma_0(\tau) + \sum_{j=1}^q \gamma_{0,j}(\tau) \hat{\theta}_{0,j}(\tau|x_j), \tag{2}$$

where  $\gamma_{0,j}(\tau)$  are weights depending on the quantile level  $\tau$ .

2. Using  $\hat{\theta}_{0,j}(\tau|x_{j,t})$  as regressors and applying penalized quantile regression, obtain the set  $\widehat{\mathcal{M}}_\tau = \{j : \hat{\gamma}_{0,j}(\tau) \neq 0\}$ . This is an estimate of the set of relevant one-step ahead predictors  $\widetilde{\mathcal{M}}_\tau = \{j : \gamma_j(\tau) \neq 0\}$  identified by the averaging model in (2). Then, compute an estimate  $\hat{\boldsymbol{\gamma}}(\tau) = (\hat{\gamma}_1(\tau), \dots, \hat{\gamma}_q(\tau))^T$  of the  $q \times 1$  parameter vector of optimal weights  $\boldsymbol{\gamma}(\tau) = (\gamma_1(\tau), \dots, \gamma_q(\tau))^T$  by considering the following weighted  $L_1$ -penalized quantile estimator

$$\hat{\boldsymbol{\gamma}}(\tau) = \arg \min_{\boldsymbol{\gamma}(\tau)} \left\{ \sum_{t=1}^N (\rho_\tau(Y_t^* - \gamma_0(\tau) - \hat{\boldsymbol{\theta}}_{0,t}^T \boldsymbol{\gamma}(\tau)) + \lambda \sum_{j=1}^q w_j |\gamma_j(\tau)| \right\}, \tag{3}$$

where  $\rho_\tau(z) = \{\tau - \mathbb{1}(z < 0)\}z$  is the quantile ‘tick’ loss function,  $\mathbb{1}(\cdot)$  the indicator function,  $\lambda > 0$  is a tuning (or penalization) parameter,  $w_j = |\hat{\gamma}_j^*|^{-c}$ ,  $c > 0$ , and  $\hat{\gamma}_j^*$  is an initial parameter estimate. Observe that for  $w_j = 1, \forall j$ , (3) becomes the usual LASSO (L) penalty. For the adaptive LASSO (aL) penalty function used in this paper, the typical choice is given by  $w_j = (\hat{\gamma}_j^{(L)} + 1/N)^{-1}$ .

3. Finally, using all the ingredients given above, compute the one-step ahead penalized averaged (PA) and aL-based quantile forecast of  $Y_{n+1}$ , i.e.,

$$\widehat{Q}_{Y_t^*}^{(PA-aL)}(\tau|\mathbf{X}_{n-q_y+1}) = \widehat{\gamma}_0(\tau) + \sum_{u \in \widetilde{\mathcal{M}}_{\lambda, \tau}^{(aL)}} \widehat{\gamma}_u^{(aL)}(\tau) \widehat{\theta}_u^{(aL)}(\tau|X_{u, n-q_y+1}), \tag{4}$$

where  $\widetilde{\mathcal{M}}_{\lambda, \tau}^{(aL)}$  is an estimate of the set of nonzero quantile predictors  $\widetilde{\mathcal{M}}_{\lambda, \tau}^{(aL)} = \{u : \theta_u(\tau|X_{u, n-q_y+1}) \neq 0\}$  associated with  $\widehat{Q}_{Y_t^*}^{(aL)}(\tau|\mathbf{X}_{n-q_y+1})$ , and where  $\widehat{\gamma}_u^{(aL)}(\tau)$  is the  $u$ th aL-based quantile regression estimate, or weight.

We adopt a prediction-based criterion for the selection of  $\lambda$ . For a given  $\lambda$ , let  $\widehat{\gamma}_\lambda(\tau) = (\widehat{\gamma}_{1, \lambda}(\tau), \dots, \widehat{\gamma}_{q, \lambda}(\tau))^T$  be the penalized estimates and  $|\widehat{\mathcal{M}}_{\lambda, \tau}|$  be the corresponding number of non-zero estimates. We choose  $\lambda$  by minimizing the following high-dimensional criterion

$$QBIC_\tau(\lambda) = \log \left( \sum_{t=1}^N \rho_\tau(Y_t^* - \widehat{\gamma}_{0, \lambda}(\tau) - \widehat{\theta}'_{0, t} \widehat{\gamma}_\lambda(\tau)) \right) + |\widehat{\mathcal{M}}_{\lambda, \tau}| \frac{\log N}{2N} C_N, \tag{5}$$

where  $C_N = \log(q) \log(\log(N)) / \log(N)$ .

### 2.3 Hybrid (H) Conditional Quantile Averaging

In addition to the forecasting approach of Section 2.2, there are also situations where a marginal conditional quantile forecast of  $Y_{t+1}$  can be computed from a fully specified parametric DGP. Combining such a parametric quantile with the marginal semiparametric conditional quantile  $\widetilde{Q}_{Y_t^*}(\tau|\mathbf{x})$  in (2) is the first step in a so-called hybrid (H) quantile averaging method. To formalize, we assume that a parametric model is available with the corresponding estimated marginal quantile functions  $\{\widehat{\theta}_1(\tau|x_{1,t}), \dots, \widehat{\theta}_{q^*}(\tau|x_{q^*,t})\}$  ( $q^* \geq 1$ ). Then, we define a hybrid extension of  $\widetilde{Q}_{Y_t^*}(\tau|\mathbf{x})$  by

$$\widetilde{Q}_{Y_t^*}^{(H)}(\tau|\mathbf{x}) = \underbrace{\gamma_0(\tau) + \sum_{j=1}^q \gamma_{0,j}(\tau) \widehat{\theta}_{0,j}(\tau|x_j)}_{\text{semiparametric}} + \underbrace{\sum_{j=1}^{q^*} \gamma_j(\tau) \widehat{\theta}_j(\tau|x_j)}_{\text{parametric}}, \tag{6}$$

where  $\{\gamma_{0,j}(\tau)\}_{j=1}^q$  and  $\{\gamma_j(\tau)\}_{j=1}^{q^*}$  are sets of weights depending on the quantile level  $\tau$ , which we summarize by the set  $\{\gamma_u^{(H)}(\tau)\}_{u=1}^{q+q^*}$  with  $q + q^* \equiv q_y + q_z + q^*$ . The design matrix  $\mathbf{X}_t$  consists of the set of potential predictors:

$$\{\widehat{\theta}_{0,1}(\tau|X_{1,t}), \dots, \widehat{\theta}_{0,q}(\tau|X_{q,t}), \widehat{\theta}_1(\tau|X_{1,t}), \dots, \widehat{\theta}_{q^*}(\tau|X_{q^*,t})\}; \tag{7}$$

its  $u$ th component will be denoted by  $\widehat{\theta}_u^{(H)}(\tau|X_{u,t})$  ( $u = 1, \dots, (q + q^*)$ ).

Similar to (3), let  $\{\widehat{\gamma}_u^{(H-aL)}(\tau)\}_{u=1}^{q+q^*}$  denote the set of hybrid quantile regression estimates of  $\{\gamma_u^{(H)}(\tau)\}_{u=1}^{q+q^*}$ . Here, we select the tuning parameter  $\lambda$  by minimizing the

hybrid version of (5), and for the weights  $w_u$  we use the aL regularization technique. Let  $\widehat{\mathcal{M}}_\tau^{(H-aL)}$  denote an estimate of the resulting set of selected nonzero quantile predictors  $\widetilde{\mathcal{M}}_\tau^{(H)} = \{u : \theta_u(\tau) \neq 0\}$  associated with  $\widetilde{Q}_{Y_t^*}^{(H)}(\tau | \mathbf{X}_{n-q_y+1})$ . Then, similar to (4) the one-step ahead hybrid and aL-based quantile forecast of  $Y_{n+1}$  is given by

$$\widehat{Q}_{Y_t^*}^{(H-aL)}(\tau | \mathbf{X}_{n-q_y+1}) = \widehat{\gamma}_0(\tau) + \sum_{u \in \widehat{\mathcal{M}}_\tau^{(H-aL)}} \widehat{\gamma}_u^{(H-aL)}(\tau) \widehat{\theta}_u^{(H)}(\tau | X_{u,n-q_y+1}). \quad (8)$$

### 3 Forecasting Procedures

In Sect. 5, we report recursive prediction results for the following six forecasting methods.

- (1) Equal Weighting (EW):  $\widehat{Q}_{\tau,t}^{(EW)} = (1/q) \sum_{j=1}^q \widehat{\theta}_{j,t}(\tau | x_{j,t})$ , where  $\widehat{\theta}_{j,t}(\tau | x_{j,t})$  is the conditional quantile estimates of  $\theta_j(\tau | x_j)$ , and  $q = q_y + q_z$ .<sup>1</sup> So, the final EW estimator does not discard irrelevant predictors.
- (2) Penalized Quantile Averaging with L-based penalty (PA-L): Forecasts are based on weighted averaging similar to the EW method, but the weights are data driven. Thus, PA-L only considers the second component of (6) with  $q^*$  marginal conditional quantiles.
- (3) Penalized Averaging of nonparametric quantiles with aL-based penalty (PA-aL): The conditional quantile forecast is defined by (4), and the set of relevant predictors consists of  $q$  marginal conditional quantiles. Thus, PA-aL only considers the first component of (6).
- (4) Penalized Linear Quantile Regression (P-Lin-QR) with L-based penalty: This is the well-known quantile estimation method of Koenker (2005). It does not allow for nonlinearities in the predictors. In this case the forecast results are based on a set of  $q$  marginal conditional quantiles.
- (5) Additive QR (Ad-QR): This method is a generalization of method (3) where each additively entered predictor has a non-parametric (possibly nonlinear) effect. Unlike methods (2)–(4), we do not conduct predictor selection for Ad-QR. Instead, we use the predictor selection results of method (3) and implement a low dimensional additive model.
- (6) Hybrid with aL-based penalty (H-aL): The conditional quantile forecast is defined by (8). In this case the forecasting results are based on a set with  $(q + q^*)$  marginal conditional quantiles. In Sect. 4, we extend the set of potential predictors with another set of marginal conditional quantiles representing a parametric approximation for the GARCH-type model under study.

Note, methods (1) – (3) and (6) are logically connected in the sense that they are based on the concept of averaging forecasts, assigning weights (equal or varying) to

<sup>1</sup> For conditional mean forecasts, the simple average of multiple forecasts is often used as an effective way of improving forecast accuracy of a certain target variable.

all predictors. Forecasting methods (4) and (5) are popular approaches used in the quantile regression literature. Unlike methods (4) and (5), penalized averaging does not assume that  $Q_{Y_{n+h}}(\tau|\mathbf{X}_n)$  has an additive structure.

### 4 Simulation Design

In Sect. 5, we evaluate/compare the multi-step out-of-sample forecasting performance of the quantile prediction methods using a recursive forecasting strategy with an increasing window size for parameter re-estimation. More concretely, the  $h = 1$  conditional quantile forecast is based on a method with parameter estimates using  $\{Y_t\}_{t=1}^{300}$  observations, the second one on a parameter vector using  $\{Y_t\}_{t=1}^{301}$  observations, and so on. The last one is based on a parameter vector using  $\{Y_t\}_{t=1}^{399}$  observations. Thus, in total there are  $N^* = 100$  one-step ahead out-of-sample conditional quantile forecasts, and where at each step of the recursive forecasting strategy the available in-sample size  $N$  is updated with one observations, starting from  $N = 300$  till  $N = 399$ . No lagged values of  $Y_t$  are included in the set of predictors, i.e.,  $q_y = 0$ . This implies that at each step of the recursive strategy  $n \equiv N$  and  $q \equiv q_z$ . For  $h > 1$ , the recursive scheme starts at observation number  $N - h + 1$  and ends at  $N - h + N^*$ .

#### 4.1 GARCH Specifications

**Semiparametric:** For the SP design, we consider an absolute value GARCH (avGARCH) model (Taylor 1986; Schwert 1990) of order  $(1, q)$  with exogenous variables  $\{Z_{j,t}\}_{j=1}^q$ , i.e.

$$Y_{t+1} = \sigma_{t+1}F_v^{-1}(V_{t+1}), \tag{9}$$

$$\sigma_{t+1} = \alpha_0 + \sum_{j=1}^q g_j(Z_{j,t}) + \alpha_1|Y_t| + \beta\sigma_t, \tag{10}$$

where  $g_j(\cdot)$  is an unknown real-valued function. In (9), we assume that  $\{V_{t+1}\}$  are i.i. d. random variables from  $U(0, 1)$  and  $F_v^{-1}(\cdot)$  is the inverse of the cumulative distribution function (CDF) of the Student  $t_v$  distribution with  $v = 4$  degrees of freedom. So, the model has heavier tails than a traditional GARCH model with Gaussian innovations.

Assuming that the root of  $1 - \beta z = 0$  lies outside the unit circle,  $\sigma_{t+1}$  can be written in an infinite-dimensional ARCH( $\infty$ ) representation. Under some regularity conditions (Xiao and Koenker, 2009) the coefficients of this model decrease geometrically. Let  $p$  ( $p \ll n - h$ ) denote a truncation parameter which depends on the number of in-sample observations. Then, the finite-order approximation of the ARCH( $\infty$ ) model is given by

$$\sigma_{t+1} \approx a_0 + \sum_{j=1}^q f_j(Z_{j,t}) + \sum_{u=1}^p a_u |Y_{t+1-u}|, \tag{11}$$

where  $a_0 = \alpha_0 / (1 - \beta)$ ,  $f_j(\cdot) = g_j(\cdot) / (1 - \beta)$  ( $j = 1, \dots, q$ ), and  $a_u = \alpha_1 \beta^{u-1}$  ( $u = 1, \dots, p$ ). To ensure the identifiability of  $\sigma_{t+1}$ , we normalize (11) so that  $a_0 = 1$ .

The truncation parameter  $p$  is small relative to the sample size  $n$ , but large enough to avoid bias of the parameter estimates. Xiao and Koenker (2009) use  $p = 3n^{1/4}$  in their study of conditional quantile estimators for GARCH models, which in our case gives  $p = 12$  or  $13$  with sample sizes  $n = 300, \dots, 399$ . However, being conservative, we set  $p = 15$  and  $q = 28$  throughout the simulation study. Thus, the design matrix  $\mathbf{X}_t$  for the SP approach of Sect. 2.2 consists of  $d = p + q = 43$  predictors. Moreover, we consider the following set of parameter values:  $\alpha_0 = 0.1$ ,  $\alpha_1 = 0.3$  and  $\beta = 0.5$ .

Using (11), the  $\tau$ th conditional quantile of  $Y_{t+1}$  given  $\mathbf{X}_t = \mathbf{x}_t$  is obtained by

$$\theta^{(\text{SP})}(\tau | \mathbf{X}_t) \approx a_0(\tau) + \sum_{j=1}^q f_{j,\tau}(Z_{j,t}) + \sum_{u=1}^p a_u(\tau) |Y_{t+1-u}|, \tag{12}$$

where  $a_u(\tau) = a_u F_v^{-1}(\tau)$  ( $u = 0, 1, \dots, p$ ), and  $f_{j,\tau}(Z_{j,t}) = f_j(Z_{j,t}) F_v^{-1}(\tau)$  ( $j = 1, \dots, q$ ). For the additive functions, we set  $f_j(z_j) = 0.05z_j^2$  ( $j = 1, 2$ ) and  $f_3(z_3) = \dots = f_{28}(z_{28}) \equiv 0$ . The exogenous predictors  $Z_{j,t}$  ( $j = 1, 2$ ) are generated by the following AR(1) process

$$Z_{j,t} = 0.5 + 0.3Z_{j,t-1} + \varepsilon_t, \quad \varepsilon_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1). \tag{13}$$

The remaining 26 exogenous predictors will be considered as noncontributory to the conditional quantile. These predictors are generated to have a compound symmetry covariance structure:  $Z_{u,t} = (V_{u,t} + sW_{u,t}) / (1 + s)$  ( $u = 3, \dots, 28$ ), where  $V_{u,t}$  and  $W_{u,t}$  are i.i.d. random variables from  $U(0, 1)$ , and  $s \geq 0$ . When  $s = 0$  the predictors  $Z_{u,t}$  are independent, whereas when  $s \neq 0$  they are dependent with correlation coefficient  $r \equiv \text{Corr}(Z_{u,t}, Z_{u',t}) = s^2 / (1 + s^2)$  ( $u \neq u' = 3, \dots, 28$ ). We report quantile evaluation results for the case  $r = 0$  ( $s = 0$ ) and  $r = 0.8$  ( $s = 2$ ).

In each step of the recursive forecast strategy, the one-step ahead quantile forecast is given by

$$\widehat{Q}_{Y_t^*}^{(\text{SP})}(\tau | \mathbf{X}_{N-p+1}) = \widehat{\beta}_0(\tau) + \sum_{j \in \widetilde{\mathcal{M}}_\tau} \widehat{\beta}_j(\tau) \widehat{\theta}_j^{(\text{SP})}(\tau | \mathbf{X}_{j,N-p+1}), \tag{14}$$

where  $\widehat{\theta}_j^{(\text{SP})}(\tau | \mathbf{X}_{j,N-p+1})$  are estimates of the recursively obtained marginal conditional quantiles  $\theta_j^{(\text{SP})}(\tau | \mathbf{X}_{j,N-p+1})$ .

**Parametric:** For the parametric (P) model, we adopt the following avGARCH (1, 1) process



$$Y_{t+1} = \sigma_{j,t+1} \varepsilon_{t+1}, \quad \varepsilon_{t+1} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \tag{15}$$

$$\sigma_{j,t+1} = \xi_0 + \phi_j Z_{j,t} + \xi_1 |Y_t| + \xi_2 \sigma_{j,t}, \quad (j = 1, \dots, q), \tag{16}$$

where  $Z_{1,t}$  and  $Z_{2,t}$  are generated as in (13), and where the remaining  $Z_{j,t}$ 's are again i.i.d.  $U(0, 1)$  random variables.

Using the apARCH model specification in the R-rugarch package, we estimate the parameters in (15)–(16) for  $j = 1, \dots, q$ . The conditional quantiles are given by

$$\widehat{\theta}_j^{(P)}(\tau|Z_{j,t}) = \widehat{\sigma}_{j,t+1} \Phi^{-1}(\tau), \tag{17}$$

where  $\Phi^{-1}(\tau)$  is the  $\tau$ th quantile of the CDF of  $\varepsilon_t$ , and  $\widehat{\sigma}_{j,t+1}$  is an estimate of  $\sigma_{j,t+1}$ . These marginal parametric quantile estimates and the marginal semiparametric conditional quantiles  $\widehat{\theta}^{(SP)}(\tau|Z_{j,t})$  serve as “pilots” of the H-aL conditional quantile estimates. Note that in the case of the hybrid forecasting method, the design matrix  $\mathbf{X}_t$  consists of  $q + d = 28 + 43 = 71$  potential regressors, i.e.

$$\{\widehat{\theta}_1^{(P)}(\tau|Z_{1,t}), \dots, \widehat{\theta}_q^{(P)}(\tau|Z_{q,t}), \widehat{\theta}_1^{(SP)}(\tau|X_{1,t}), \dots, \widehat{\theta}_d^{(SP)}(\tau|X_{d,t})\}. \tag{18}$$

For instance, with  $\widehat{\theta}_u^{(H)}(\tau|\cdot)$  ( $u = 1, \dots, q + d$ ) denoting the  $u$ th component of the set (18), the recursively obtained one-step ahead H-aL quantile forecasts of  $Y_{301}, \dots, Y_{400}$  follow from (8) with  $n \equiv N = 300, \dots, 399$ , respectively.

## 4.2 Evaluating Predicted Quantiles

### 4.2.1 Expected loss

Recall that  $q_y = 0$ . Then, let  $\{e_{t,h}^{(\cdot)} = Y_{n+h} - \widehat{Q}_{Y_t}^{(\cdot)}(\tau|\mathbf{X}_{n+h})\}_{t=1}^{N^*}$  be the set of  $h$ -step ahead out-of-sample quantile prediction errors (QPEs) obtained by the recursive forecasting strategy, where the superscript  $(\cdot)$  refers to one of the six methods discussed in Sect. 3,  $Y_t^* = Y_{t+h+p-1}$ , and  $N^* = 100$  for all values of  $h$ . As a measure of performance, we calculate an average of the tick loss function  $\rho_\tau(\cdot)$  of QPE values, i.e.

$$\widehat{L}_{\tau,h}^{(\cdot)} = \frac{1}{N^*} \sum_{t=1}^{N^*} \rho_\tau(e_{t,h}^{(\cdot)}). \tag{19}$$

This average is an estimate of the expected loss  $L_{\tau,h}^{(\cdot)} = \mathbb{E}[\rho_\tau(Y_{n+h} - Q_{Y_t}^{(\cdot)}(\tau|\mathbf{X}_{n+h}))|\mathcal{F}_t]$  where  $\mathcal{F}_t$  is the sigma-algebra generated by the available information up to time  $t$ . For each recursively obtained quantile forecast,  $\widehat{L}_{\tau,h}^{(\cdot)}$  weights the difference between the observed value  $Y_{n+h}$  and the forecasted quantile  $\widehat{Q}_{Y_t}^{(\cdot)}(\tau|\mathbf{X}_{n+h})$

by  $(1 - \tau)$  when  $Y_{n+h}$  is lower than the  $\tau$ th quantile, and by  $\tau$  when  $Y_{n+h}$  exceeds the quantile. In this sense, (19) is a natural way to evaluate quantile forecasts.

#### 4.2.2 Equal forecast accuracy

We assess the difference between two conditional quantile methods A and B via a Diebold–Mariano (DM) type test statistic. To this end, let  $e_{t,h}^{(A)}$  and  $e_{t,h}^{(B)}$  be the associated  $h$ -step ahead QPEs. For a fixed quantile level  $\tau$  and fixed forecast horizon  $h$ , the corresponding loss differential is defined as

$$D_{t,\tau,h}^{(A,B)} = \rho_\tau(e_{t,h}^{(A)}) - \rho_\tau(e_{t,h}^{(B)}). \quad (20)$$

The null hypothesis that method A produces as accurate conditional quantile forecasts as method B can be tested using the test statistic  $DM_{\tau,h}^{(A,B)} \equiv \overline{D}_{\tau,h}^{(A,B)} / \{\text{Var}(\overline{D}_{\tau,h}^{(A,B)})\}^{1/2} \sim \mathcal{N}(0, 1)$ , where  $\overline{D}_{\tau,h}^{(A,B)}$  is the average over  $t$  of  $D_{t,\tau,h}^{(A,B)}$  at forecast horizon  $h$ . Under the alternative hypothesis, we specify a one-sided test, so that rejection of the null indicates that method B is more accurate than method A. For  $h = 1$ , a consistent estimator of  $\text{Var}(\overline{D}_{\tau,h}^{(A,B)})$  is given by the sample variance of  $D_{t,\tau,h}^{(A,B)}$ . For  $h > 1$ , we use the modified DM test proposed by Harvey et al. (1989).<sup>2</sup>

#### 4.2.3 Encompassing

The topic of combining conditional quantile forecasts can also be investigated via the principle of “encompassing” (ENC). To this end, consider the set of  $h$ -step ahead quantile forecasts  $\{\widehat{Q}_{t,h}(\tau) \equiv (1, \widehat{Q}_{t,h}^{(A)}(\tau), \widehat{Q}_{t,h}^{(B)}(\tau))^T\}_{t=1}^{n-h}$ , a  $(3 \times 1)$  vector, where  $\widehat{Q}_{t,h}^{(A)}(\tau)$  and  $\widehat{Q}_{t,h}^{(B)}(\tau)$  are two alternative conditional  $h$ -step ahead quantile forecasts, such that  $\text{Corr}(\widehat{Q}_{t,h}^{(A)}(\tau), \widehat{Q}_{t,h}^{(B)}(\tau)) \neq \pm 1$ . Suppose there exists a class  $\mathcal{W}$  of weights. Let  $\widehat{Q}_{t,h}^{(c)}(\tau) = \omega_h^T \widehat{Q}_{t,h}(\tau)$  with  $\omega_h = (\omega_h^{(0)}, \omega_h^{(1)}, \omega_h^{(2)})^T \in \mathcal{W}$ , denoting a linear quantile forecast combination (c). Then, for fixed values of  $\tau$  and  $h$ , forecast  $\widehat{Q}_{t,h}^{(A)}(\tau)$  is said to encompass forecast  $\widehat{Q}_{t,h}^{(B)}(\tau)$  conditionally at time  $t$  with respect to the quantile tick loss function  $\rho_\tau(\cdot)$  if and only if  $\mathbb{E}[\rho_\tau(Y_{t+h} - \widehat{Q}_{t,h}^{(A)}(\tau)) | \mathcal{F}_t] \leq \mathbb{E}[\rho_\tau(Y_{t+h} - \widehat{Q}_{t,h}^{(c)}(\tau)) | \mathcal{F}_t]$  for all  $\omega_h \in \Omega \subseteq \mathbb{R}^3$  being a compact set. Testing this inequality for all  $\omega \in \Omega$  is, however, infeasible. Instead, let  $\tilde{\omega} = (\tilde{\omega}^{(0)}, \tilde{\omega}^{(1)}, \tilde{\omega}^{(2)})^T \in \Omega$  be the  $h$ -step ahead optimal forecast combination parameters which, conditionally at time  $t$ , minimizes the expected loss

<sup>2</sup> By allotting only a relatively small fraction of the set of observations to the forecast period, we are able to side-step issues related to the effects of in-sample parameter estimation uncertainty on the distribution of the DM test statistic.

$$\tilde{\mathcal{Q}}_h(\tau) = \arg \min_{\theta \in \Theta} \mathbb{E}[\rho_\tau(Y_{t+h} - \hat{\mathcal{Q}}_{t,h}^{(c)}(\tau)) | \mathcal{F}_t]. \tag{21}$$

We consider conducting two separate tests:  $\mathbb{H}_{1,0} : \tilde{\omega} = (0, 1, 0)^T$  versus  $\mathbb{H}_{1,a} : \tilde{\omega} \neq (0, 1, 0)^T$ , and  $\mathbb{H}_{2,0} : \tilde{\omega} = (0, 0, 1)^T$  versus  $\mathbb{H}_{2,a} : \tilde{\omega} \neq (0, 0, 1)^T$ , where for ease of notation we dropped the dependence on  $h$ . This testing framework corresponds, respectively, to testing whether  $\hat{\mathcal{Q}}_{t,h}^{(A)}(\tau)$  encompasses  $\hat{\mathcal{Q}}_{t,h}^{(B)}(\tau)$ , and whether  $\hat{\mathcal{Q}}_{t,h}^{(B)}(\tau)$  encompasses  $\hat{\mathcal{Q}}_{t,h}^{(A)}(\tau)$ , conditionally.

A root- $n$ -consistent estimator of (21) is given by

$$\begin{aligned} \tilde{\mathcal{Q}}_{n,h}^T(\tau) &\equiv (\tilde{\mathcal{Q}}_{n,h}^{(0)}(\tau), \tilde{\mathcal{Q}}_{n,h}^{(1)}(\tau), \tilde{\mathcal{Q}}_{n,h}^{(2)}(\tau)) \\ &= \arg \min_{\theta \in \Theta} \frac{1}{n-h} \sum_{t=1}^{n-h} \rho_\tau((Y_{t+h} - (\omega_0 + \omega_1 \hat{\mathcal{Q}}_t^{(1)}(\tau) + \omega_2 \hat{\mathcal{Q}}_t^{(2)}(\tau))). \end{aligned} \tag{22}$$

Under a sufficient weak dependence condition (and additional regularity conditions), it can be shown (Koenker, 2005) that within the context of conditional quantile autoregression and for fixed values of  $\tau$  and  $h$ ,  $\sqrt{n}(\tilde{\mathcal{Q}}_{n,h}(\tau) - \tilde{\mathcal{Q}}_h(\tau)) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Sigma_h)$ , with  $\Sigma_h = \tau(1-\tau)\Sigma_{1,h}^{-1}\Sigma_{0,h}\Sigma_{1,h}^{-1}$  where  $\Sigma_{0,h} = \mathbb{E}[\tilde{\mathcal{Q}}_{t,h}(\tau)\tilde{\mathcal{Q}}_{t,h}^T(\tau)|\mathcal{F}_t]$ ,  $\Sigma_{1,h} = \mathbb{E}[f_{t,h}(\tilde{\omega}^T\tilde{\mathcal{Q}}_{t,h}(\tau)|\mathcal{F}_t)\tilde{\mathcal{Q}}_{t,h}(\tau)\tilde{\mathcal{Q}}_{t,h}^T(\tau)]$ , and  $f_{t,h}(\cdot|\mathcal{F}_t)$  is the conditional density of  $Y_{t+h}$  evaluated at  $\tilde{\omega}$ .

Following Furtés and Olmo (2013), we propose the following two Wald-type test statistics for testing  $\mathbb{H}_{1,0}$  and  $\mathbb{H}_{2,0}$ :

$$\text{ENC}_{\tau,h}^{(1)} = n(\tilde{\omega}_n^T - (0, 1, 0))\hat{\Sigma}_{n,h}^{-1}(\tilde{\omega}_n - (0, 1, 0)^T), \tag{23}$$

$$\text{and ENC}_{\tau,h}^{(2)} = n(\tilde{\omega}_n^T - (0, 0, 1))\hat{\Sigma}_h^{-1}(\tilde{\omega}_n - (0, 0, 1)^T), \tag{24}$$

where  $\hat{\Sigma}_h \equiv \tau(1-\tau)\hat{\Sigma}_{1,h}^{-1}\hat{\Sigma}_{0,h}\hat{\Sigma}_{1,h}^{-1}$  is a consistent estimator of the  $3 \times 3$  covariance matrix  $\Sigma_h$  with

$$\hat{\Sigma}_{0,h} = \frac{1}{n-h} \sum_{t=1}^{n-h} \tilde{\mathcal{Q}}_{t,h}(\tau)\tilde{\mathcal{Q}}_{t,h}^T(\tau), \quad \hat{\Sigma}_{1,h} = \frac{1}{2(n-h)h_n} \sum_{t=1}^{n-h} \mathbb{1}(|Y_{t+h} - \hat{\mathcal{Q}}_{t,h}^{(c)}(\tau)| \leq h_n)\tilde{\mathcal{Q}}_{t,h}(\tau)\tilde{\mathcal{Q}}_{t,h}^T(\tau).$$

The kernel-type matrix  $\hat{\Sigma}_{1,h}$  builds upon a method proposed by Powell (1991), with  $h_n = \nu n^{-1/3}$  ( $\nu > 0$ ) a bandwidth parameter satisfying  $h_n \rightarrow 0$  and  $nh_n^2 \rightarrow \infty$  as  $n \rightarrow \infty$ .<sup>3</sup> For a fixed value of the forecast horizon  $h$ , it follows from (White, 2001, Thm. 4.3) that under  $\mathbb{H}_{i,0}$ , and as  $n \rightarrow \infty$ ,  $\text{ENC}_{\tau,h}^{(i)} \xrightarrow{d} \chi_3^2$  ( $i = 1, 2$ ). Under  $\mathbb{H}_{i,a}$ , and as  $n \rightarrow \infty$ ,  $\text{ENC}_{\tau,h}^{(i)} \rightarrow +\infty$  ( $i = 1, 2$ ).

<sup>3</sup> Following Koenker (2005), we set  $\nu = 1$ ; the results presented in Sect. 5 hardly vary for alternative values of  $\nu$ .

## 5 Monte Carlo Results

### QPE Loss Values

Table 1 presents the sample mean and the sample median values of the ratio  $\widehat{L}_{\tau,h}^{(\text{Method})} / \widehat{L}_{\tau,h}^{(\text{EW})}$  in the left tail ( $\tau = 0.05$  and  $\tau = 0.1$ ) and the right tail ( $\tau = 0.9$  and  $\tau = 0.95$ ) of the  $h$ -step ahead distribution of the QPE values ( $h = 1$  and  $h = 6$ ) for independent ( $r = 0$ ) and dependent ( $r = 0.8$ ) exogenous predictors. The main findings are summarized as follows.

1. When  $r = 0$  and  $r = 0.8$ , the estimated QPE loss ratios are virtually the same for  $h = 1$  and  $h = 6$ .
2. Relative to the EW method, the H-aL method beats all other methods in terms of the lowest QPE loss ratios, irrespective of the level of dependence  $r$  between the predictors. Compared to all other methods, the reduction in loss is less pronounced for the Ad-QR method.
3. As revealed by their lower loss ratios, there are differences between the  $h = 1$  forecasting performance of all methods in the right tail and the left tail of the distribution of the QPE values. The differences are no longer present for  $h = 6$ .
4. For all methods the QPE loss ratios are much smaller at the boundaries  $\tau = 0.05$  and  $\tau = 0.95$  than at, respectively,  $\tau = 0.1$  and  $\tau = 0.9$ . With the exception of H-aL, the differences in QPE loss ratios are small for all other quantile forecasting methods at  $\tau = 0.5$  (not shown here). In this case the set of exogenous predictors  $\{Z_{j,t}\}$  consists of 28 i.i.d.  $U(0,1)$  random variables.

### Equal Forecast Accuracy

Table 2 shows averages (based on 50 replications) of computed  $p$ -values of the  $DM_{\tau,h}^{(\text{EW,Method})}$  test statistic for  $r = 0$ , with  $h = 1$  and  $h = 6$ . Except for the PA-L method at  $\tau = 0.9$  and  $h = 6$ , all null hypotheses are rejected at a 5% nominal significance level, implying that the use of a particular nonparametric conditional quantile method yields statistically more accurate quantile forecasts than using the EW method. Again, we note that the dependencies between the exogenous predictors have hardly any impact on the reported  $p$ -values. Overall, the PA-aL method has the lowest average  $p$ -value but the differences with the P-Lin-QR, Ad-QR, and H-aL methods are minimal. The  $DM_{\tau,h}^{(\text{EW,Method})}$  test results are generally the same for  $r = 0.8$  (not shown here).

### Encompassing

Table 3 shows the average  $p$ -values of the  $ENC_{\tau,h}^{(i)}$  ( $i = 1, 2$ ) test statistics for  $h = 1$  and  $h = 6$ . Here we make pairwise comparisons between quantile forecasts obtained by methods 1)–5) relative to quantile forecasts obtained by the H-aL method. Two general results are evident from the table.

1. In the left tail of the quantile forecast distribution neither  $\mathbb{H}_{1,0}$  and  $\mathbb{H}_{2,0}$  are rejected at a 1% nominal significance level for  $h = 1$  and  $h = 6$ . This indicates that quantile forecasts of the H-aL method encompasses those of methods 1)–5).

**Table 1** Sample mean and sample median (in parentheses) of the ratios of tick loss functions  $\widehat{L}_{\tau,h}^{(Method)} / \widehat{L}_{\tau,h}^{(EW)}$  for forecast horizons  $h = 1$  and  $h = 6$ , and averaged over 50 predictions

$r$	Method	$h = 1$				$h = 6$			
		$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$
0	PA-L	0.672 (0.650)	0.832 (0.825)	0.829 (0.828)	0.652 (0.637)	0.665 (0.671)	0.837 (0.857)	0.856 (0.880)	0.695 (0.699)
	PA-aL	0.701 (0.667)	0.885 (0.869)	0.882 (0.864)	0.690 (0.650)	0.716 (0.716)	0.896 (0.884)	0.894 (0.893)	0.715 (0.727)
	P-Lin-QR	0.687 (0.658)	0.879 (0.874)	0.874 (0.869)	0.667 (0.653)	0.702 (0.713)	0.895 (0.873)	0.886 (0.882)	0.696 (0.713)
	Ad-QR	0.743 (0.700)	0.890 (0.872)	0.890 (0.882)	0.712 (0.682)	0.745 (0.758)	0.906 (0.904)	0.907 (0.914)	0.746 (0.746)
	H-aL	0.641 (0.606)	0.810 (0.803)	0.802 (0.792)	0.621 (0.602)	0.649 (0.639)	0.823 (0.827)	0.811 (0.820)	0.662 (0.672)
0.8	PA-L	0.670 (0.650)	0.832 (0.825)	0.829 (0.828)	0.652 (0.637)	0.665 (0.672)	0.837 (0.858)	0.856 (0.880)	0.695 (0.699)
	PA-aL	0.706 (0.681)	0.885 (0.869)	0.882 (0.864)	0.690 (0.650)	0.726 (0.719)	0.896 (0.888)	0.894 (0.893)	0.714 (0.722)
	P-Lin-QR	0.690 (0.656)	0.879 (0.874)	0.874 (0.869)	0.667 (0.653)	0.706 (0.703)	0.887 (0.875)	0.885 (0.894)	0.696 (0.714)
	Ad-QR	0.743 (0.710)	0.890 (0.872)	0.890 (0.882)	0.712 (0.682)	0.749 (0.745)	0.909 (0.907)	0.909 (0.916)	0.737 (0.758)
	H-aL	0.638 (0.609)	0.810 (0.803)	0.802 (0.792)	0.621 (0.602)	0.662 (0.666)	0.830 (0.831)	0.808 (0.833)	0.659 (0.668)

EW, Equal weighting; PA-L, Penalized averaging with LASSO penalty; PA-aL, Penalized averaging with adaptive LASSO penalty; P-Lin-QR, Penalized linear model with quantile regression; Ad-QR, Additive quantile regression; and H-aL, Hybrid method with adaptive LASSO penalty

**Table 2** Averages of computed  $p$ -values of the  $DM_{\tau,h}^{(EW,Method)}$  test statistic;  $r = 0$

Method	$h = 1$				$h = 6$			
	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$
PA-L	0.030	0.043	0.033	0.029	0.033	0.040	0.050	0.037
PA-aL	0.027	0.024	0.016	0.023	0.028	0.025	0.022	0.027
P-Lin-QR	0.026	0.024	0.016	0.023	0.028	0.028	0.023	0.028
Ad-QR	0.028	0.028	0.021	0.025	0.030	0.030	0.025	0.028
H-aL	0.024	0.033	0.028	0.026	0.032	0.041	0.030	0.033

See the note to Table 1

- Overall, the largest evidence for not rejecting  $\mathbb{H}_{i,0}$  ( $i = 1, 2$ ) seems to occur for quantile forecasts obtained by the PA-L method.
- Very different conclusions emerge for the right tail of the quantile forecast distribution. In particular, at a 5% nominal significance level both  $ENC_{\tau,h}^{(i)}$  tests plainly reject  $\mathbb{H}_{i,0}$  ( $i = 1, 2$ ) for  $h = 1$  and  $h = 6$ , indicating that quantile

**Table 3** Averages of computed  $p$ -values of the  $ENC_{\tau,h}^{(i)}$  test statistics ( $i = 1, 2$ );  $r = 0$

Test	Method	$h = 1$				$h = 6$			
		$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$
$ENC_{\tau,h}^{(1)}$	EW	0.014	0.021	0.000	0.018	0.017	0.009	0.000	0.000
	PA-L	0.081	0.075	0.000	0.000	0.127	0.071	0.000	0.000
	PA-aL	0.054	0.081	0.000	0.020	0.057	0.071	0.000	0.000
	P-Lin-QR	0.046	0.106	0.020	0.000	0.080	0.077	0.000	0.000
	Ad-QR	0.093	0.136	0.000	0.000	0.047	0.060	0.000	0.000
$ENC_{\tau,h}^{(2)}$	EW	0.071	0.071	0.000	0.018	0.043	0.124	0.000	0.000
	PA-L	0.097	0.107	0.000	0.000	0.101	0.080	0.000	0.000
	PA-aL	0.054	0.095	0.000	0.020	0.055	0.068	0.000	0.000
	P-Lin-QR	0.071	0.107	0.020	0.000	0.050	0.079	0.000	0.000
	Ad-QR	0.068	0.083	0.000	0.000	0.063	0.073	0.000	0.000

$\mathbb{H}_{1,0} : \hat{Q}_{t,h}^{(Method)}(\tau)$  encompasses  $\hat{Q}_{t,h}^{(H-aL)}(\tau)$ ;  $\mathbb{H}_{2,0} : \hat{Q}_{t,h}^{(H-aL)}(\tau)$  encompasses  $\hat{Q}_{t,h}^{(Method)}(\tau)$ . The entries 0.000 denote an average  $p$ -value  $< 10^{-3}$

forecasts obtained by the H-aL method do not encompass quantile forecasts obtained from methods 1)–5).

## 6 Empirical Application

### 6.1 Data and Forecast Procedure

In this section we make pairwise comparisons between quantile forecasts obtained by methods 2)–6) relative to quantile forecasts obtained by the EW method. The time series of interest is the risk premium of the monthly S&P 500 index, denoted by  $R_t$ . The set of potential exogenous predictors consists of  $q_z = 14$  macroeconomic variables covering the time period 1951:01–2016:12 (792 observations).<sup>4</sup> Specifically, we consider the following predictors: dividend-price ratio (DP); dividend yield (DY); earnings-price ratio (EP); dividend-payout ratio (DE); equity risk premium volatility (RVOL); book-to-market ratio (BM); net equity expansion (NTIS); treasury bill rate (TBL); long-term yield (LTY); long-term return (LTR); term spread TMS); default yield spread (DFY); default return spread (DFR); and inflation (NFL). These series have been the subject of conditional quantile prediction studies by Lima and Meng, (2017), Meligkotsidou et al. (2014, 2019), and De Gooijer and Zerom (2020) among others.

<sup>4</sup> The data set was downloaded from <http://www.hec.unil.ch/agoyal/>.

To obtain estimates of the marginal parametric conditional quantiles of the  $R_t$  series, we adopt a so-called time-varying mean model with exponential GARCH-Z volatility, i.e.,

$$\begin{cases} R_{t+1} = \beta_{0,j} + \beta_{1,j}Z_{j,t} + \varepsilon_{j,t+1}, & \varepsilon_{j,t+1} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{j,t+1}^2), \\ \log(\sigma_{j,t+1}^2) = \delta_{0,j} + \delta_{1,j}Z_{j,t} + \delta_{2,j} \log(\sigma_{j,t}^2) + \delta_{3,j} \left| \frac{\varepsilon_{j,t}}{\sigma_{j,t}} \right| + \delta_{4,j} \frac{\varepsilon_{j,t}}{\sigma_{j,t}}, & (j = 1, \dots, q_z). \end{cases} \tag{25}$$

This is a very wide model specification from the literature on the predictability of the S&P 500 risk premium; see, e.g., Cenesizoglu and Timmermann (2012). Thus, with the 14 semiparametric marginal quantiles, the complete portfolio of predictors for the hybrid quantile averaging method consists of 28 potential variables.

We use a recursive forecasting strategy similar to the one discussed in Sect. 4. For  $h = 1$ , the first forecast origin is 1965:12 and hence the first in-sample (estimation) period is from 1951:01 to 1964:12 ( $N = 168$ ). The second in-sample covers the period 1951:01–1965:01 ( $N = 169$ ). The last in-sample covers the period 1951:01–2016:11 ( $N = 791$ ). Thus, in total there are  $N^* = 624$  one-step ahead out-of-sample conditional quantile forecasts. For  $h = 6$ , the first forecast origin is 1964:07, and the last forecast origin is at 2016:06.

### 6.2 Empirical Results

Table 4, top panel, presents sample means (taken over  $N^* = 624$  values) of the ratios of the tick loss functions  $\widehat{L}_{\tau,h}^{(\text{Method})} / \widehat{L}_{\tau,h}^{(\text{EW})}$  for  $\tau = 0.05, 0.1, 0.9, \text{ and } 0.95$ , and  $h = 1$  and  $h = 6$ . With the exception of Ad-QR, the values of the estimated QPE loss ratios are all below one, irrespective of the forecast horizon  $h$ . This indicates that there are improvements in conditional quantile forecasts over out-of-sample quantile forecasts obtained by the EW method. Overall, the best forecast performance seems to occur for P-Lin-QR and H-aL. Interestingly, for  $\tau = 0.05$  and  $\tau = 0.95$  the P-Lin-QR method has slightly lower QPE loss ratios than the H-aL method irrespective of the value of  $h$ . In this respect, recall that conditional quantiles based on the P-Lin-QR method are best for linear functionals of predictors. So the choice of  $\tau$  can affect the ranking of P-Lin-QR and H-aL.

The bottom panel of Table 4 shows  $p$ -values of the  $DM_{\tau,h}^{(\text{EW}, \text{Method})}$  test statistic. Note, that the null hypothesis of equal forecast accuracy is *not* rejected for PA-aL and Ad-QR for all values of  $\tau$  and  $h$ . We also see that the  $p$ -values of  $DM_{\tau,h}^{(\text{EW}, \text{PA-L})}$  for  $h = 1, 6$ , and  $\tau = 0.9$  are close to one, indicating that there is no significant difference between the forecast performance of EW and PA-L. Overall, the top and bottom panels confirm the good performance of the H-aL method as opposed to the EW method of combining quantile forecasts.

Table 5 shows that the gain in forecasting performance of the H-aL method mainly comes from including the predictor variables DFR and RVOL in the parametric component. By contrast, the semiparametric component includes a variety of

**Table 4** Pairwise comparison of quantile forecasting methods for the risk premium of the monthly S&P 500 index

Method	$h = 1$				$h = 6$			
	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$	$\tau = 0.05$	$\tau = 0.1$	$\tau = 0.9$	$\tau = 0.95$
<b>QPE loss ratios</b> ( $\widehat{L}_{\tau,h}^{(Method)} / \widehat{L}_{\tau,h}^{(EW)}$ )								
PA-L	0.878	0.891	0.856	0.808	0.875	0.885	0.853	0.811
PA-aL	0.942	0.967	0.914	0.891	0.946	0.973	0.910	0.892
P-Lin-QR	<b>0.760</b>	0.824	0.818	<b>0.719</b>	<b>0.759</b>	0.821	0.815	<b>0.715</b>
Ad-QR	0.924	1.025	0.947	0.967	0.919	0.969	1.101	0.911
H-aL	0.772	<b>0.817</b>	<b>0.813</b>	0.746	0.771	<b>0.817</b>	<b>0.806</b>	0.742
<b>Equal forecast accuracy</b> ( $p$ -values of $DM_{\tau,h}^{(EW,Method)}$ )								
PA-L	0.001	0.001	0.955	0.019	0.001	0.000	0.905	0.016
PA-aL	0.922	0.987	1.000	0.999	0.992	0.993	1.000	1.000
P-Lin-QR	0.000	0.000	0.011	0.000	0.000	0.000	0.008	0.000
Ad-QR	0.377	0.501	0.997	1.000	0.309	0.137	1.000	0.988
H-aL	0.000	0.000	0.008	0.001	0.000	0.000	0.003	0.001

See the note to Table 1. Embolded entries show the lowest QPE ratios for each  $\tau$ . The entries 0.000 denote a  $p$ -value  $< 10^{-3}$

**Table 5** Frequency of predictors which are selected in at least 5% of the times by each component (Comp) of the H-aL method for forecast horizon  $h = 1$

$\tau$	Comp	Predictors										
		DP	DY	EP	RVOL	NTIS	TBL	LTR	TMS	DFY	DFR	INFL
0.05	P				6.7						6.3	
	SP		6.5				9.1		8.1			
0.1	P				7.0						7.1	
	SP		5.8				7.2	5.4	8.0			5.0
0.9	P				15.3						13.1	
	SP	6.6				11.2		10.1	8.7			
0.95	P				14.9							
	SP			8.4		10.9		11.1		6.7		

SP, Semiparametric; and P, Parametric component of H-aL method

predictor variables with significant weights on the set of potential predictors. Finally, as revealed by their selection frequencies less than 5%, there is evidence that the variables DE, BM, and LTY hardly contribute to the predictability of  $R_t$ .



## 7 Conclusion

In the first half of the paper, we have explored the multi-step ahead quantile forecasting performance of a hybrid method versus quantile forecast results obtained by five alternative, non-hybrid methods, via a Monte Carlo study. The study yields two main findings. First, relative to quantile forecasts obtained from simple averaging quantile forecasts (i.e., the EW method), the hybrid H-aL method is more accurate in terms of the lowest QPE loss ratios, irrespective of the level of dependence between the predictors. This evidence is strongest in the tails of the forecast distribution of the response variable, and most likely due to the benefits of including information from multiple predictor variables. Second, quantile forecasts obtained by the H-aL method do not encompass quantile forecasts from five competitors in the right tail of the quantile forecast distribution.

The second half of the paper provided an empirical application to forecast the risk premium of the monthly S&P 500 index. Very similar results are obtained as in the Monte Carlo simulation study. In particular, our findings strongly suggest that it is possible to use pre-selected macroeconomic predictors to produce better out-of-sample quantile forecasts than the EW combination approach which uses fixed, time-invariant, weights of conditional quantile forecasts. The H-aL method on the other hand uses stochastic rather than fixed weights. As such, we showed that variable selection and random weights jointly can achieve forecast efficiency gains in the tails of the quantile return distribution. We also provided some insight as to where the gains of the hybrid method come from.

**Acknowledgements** The author would like to acknowledge an anonymous referee for constructive suggestions and comments.

**Funding** Funding not applicable. No funding available.

## Declarations

**Competing Interests** The author declares that he has no conflict of interest.

**Compliance with ethical standards** The author declares that he has no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- 3Aijaz, I., & Agarwal, P. (2020). A study on time series forecasting using hybridization of time series models and neural networks. *Recent Advances in Computer Science and Communications*, 13, 827–832.
- Aydin, D., & Işçi Güneri, O. (2015). Time series prediction using hybridization of AR, SETAR and ARM models. *International Journal of Applied Science and Technology*, 5(6), 87–96.
- Cenesizoglu, T., & Timmermann, A. (2012). Do return prediction models add economic value? *Journal of Banking & Finance*, 36, 2974–2987.
- De Gooijer, J. G., & Zerom, D. (2019). Semiparametric quantile averaging in the presence of high-dimensional predictors. *International Journal of Forecasting*, 35, 891–909.
- De Gooijer, J.G., & Zerom, D. (2020). Penalized averaging of parametric and non-parametric quantile forecasts. *Journal of Time Series Econometrics*, 12, 1.
- Fan, J., Lv, J., & Qi, L. (2011). Sparse high dimensional models in economics. *Annual Review of Economics*, 3, 291–317.
- Fuertes, A.-M., & Olmo, J. (2013). Optimally harnessing inter-day and intra-day information for daily value-at-risk prediction. *International Journal of Forecasting*, 29, 28–42.
- Garcia, M. G. P., Medeiros, M. C., & Vasconcelos, G. F. R. (2017). Real-time inflation forecasting with high-dimensional models: The case of Brazil. *International Journal of Forecasting*, 33, 679–693.
- Giacomini, R., & Komunjer, I. (2005). Evaluation and combination of conditional quantile forecasts. *Journal of Business & Economic Statistics*, 23, 416–431.
- Harvey, D. I., Leybourne, S. J., & Newbold, P. (1989). Tests for forecasting encompassing. *Journal of Business & Economic Statistics*, 16, 254–259.
- Jiang, B., Athanasopoulos, G., Hyndman, R. J., Panagiotelis, A., & Vahid, F. (2017). Macroeconomic forecasting for Australia using a large number of predictors. *Monash Econometrics and Business Statistics Working Papers*, 2, 17.
- Koenker, R. (2005). *Quantile Regression*. MA: New York: Cambridge, University Press.
- Konzen, E., & Zieglemann, F. A. (2016). LASSO-type penalties for covariate selection and forecasting time series. *Journal of Forecasting*, 35, 592–612.
- Lee, T. H. (2011). Combining forecasts with many predictors. In M. I. Kalamazoo (Ed.), *Advances in economic forecasting* (pp. 149–172). MI: WE Upjohn Institute for Employment Research.
- Lima, L. R., & Meng, F. (2017). Out-of-sample return predictability: a quantile combination approach. *Journal of Applied Econometrics*, 32, 877–895.
- Manzan, S. (2015). Forecasting the distribution of economic variables in a data-rich environment. *Journal of Business & Economic Statistics*, 33, 144–164.
- Meligkotsidou, L., Panopoulou, E., Vrontos, I. D., & Vrontos, S. D. (2014). A quantile regression approach to equity premium prediction. *Journal of Forecasting*, 33, 558–576.
- Meligkotsidou, L., Panopoulou, E., Vrontos, I. D., & Vrontos, S. D. (2019). Quantile forecast combinations in realised volatility prediction. *Journal of the Operational Research Society*, 70, 1720–1733.
- Pai, P. F., & Lin, C. S. (2005). A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega*, 33, 497–505.
- Powell, J. L. (1991). Estimation of monotonic regression models under quantile restrictions. In W. A. Barnett, J. L. Powell, & G. E. Tauchen (Eds.), *Nonparametric and Semiparametric Methods in Econometrics and Statistics* (pp. 357–386). MA: Cambridge, MIT Press.
- Schwert, W. (1990). Stock volatility and the crash of '87. *Review of Financial Studies*, 3, 77–102.
- Taylor, S. J. (1986). *Modelling Financial Time Series*. New York: Wiley.
- Uematsu, Y., & Tanaka, S. (2019). High-dimensional macroeconomic forecasting and variable selection via penalized regression. *Econometrics Journal*, 22, 34–56.
- Valenzuela, O., Rojas, I., Rojas, F., Pomares, H., Herrera, L. J., Guillen, A., et al. (2008). Hybridization of intelligent techniques and ARIMA models for time series prediction. *Fuzzy Sets and Systems*, 159, 821–845.
- White, H. (2001). *Asymptotic Theory for Econometricians* (Rev). London: Academic Press.
- Xiao, Z., & Koenker, R. (2009). Conditional quantile estimation for generalized autoregressive conditional heteroscedasticity models. *Journal of the American Statistical Association*, 104, 1696–1712.
- Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, 50, 159–175.