

Tae Yeon Kim, Hee-Seok Oh and Yaeji Lim*

Forecasting High-Dimensional Non-Normal Time Series Using Averaged Quantile Regression

<https://doi.org/10.1515/jtse-2024-0014>

Received May 9, 2024; accepted August 4, 2025; published online August 19, 2025

Abstract: This paper proposes a forecasting method based on averaged quantile regression to improve predictions for non-normally distributed data. Traditional forecasting models often rely on ordinary least squares, assuming normally distributed errors, which can be restrictive in practice. By leveraging quantile regression, our approach provides a more robust alternative that captures the varying effects of predictors across different quantiles of the response variable. The key contribution of this study is to introduce averaged quantile regression (AQR) as a flexible and effective forecasting tool for high-dimensional, non-normally distributed time series. We show that AQR outperforms conventional mean-based forecasting in a factor modeling setting and remains robust across diverse heavy-tailed, skewed, and near-normal distributions. While our method can be applied broadly, we illustrate its effectiveness within the dynamic factor model framework through numerical experiments and real data analysis.

Keywords: averaged quantile regression; approximate dynamic factor model; forecasting; economic data

JEL Classification: C10; C38

1 Introduction

Effective forecasting of high-dimensional time series data is crucial in many fields, especially when the data exhibit non-normal characteristics such as heavy tails or skewness. Traditional forecasting models often rely on ordinary least squares

*Corresponding author: Yaeji Lim, Department of Applied Statistics, Chung-Ang University, Seoul, 06974, Korea, E-mail: yaeji.lim@gmail.com. <https://orcid.org/0000-0002-8698-8667>

Tae Yeon Kim and Hee-Seok Oh, Department of Statistics, Seoul National University, Seoul, 08826, Korea, E-mail: ktyppp38@snu.ac.kr (T.Y. Kim), heeseok@stats.snu.ac.kr (H.-S. Oh)

 Open Access. © 2026 the author(s), published by De Gruyter.  This work is licensed under the Creative Commons Attribution 4.0 International License.

(OLS) regression, which assumes normally distributed errors and homoscedasticity (Stock and Watson 2002; De Valk et al. 2019). However, these assumptions are frequently violated in real-world applications, leading to suboptimal predictions (Mandelbrot 1963; Koenker and Hallock 2001).

To address this limitation, we propose a forecasting method based on averaged quantile regression (AQR), which provides a more flexible and robust alternative to mean-based predictive models. Quantile regression, introduced by Koenker and Bassett (1978), allows for a more comprehensive analysis of predictor-response relationships by estimating conditional quantiles instead of relying solely on the conditional mean (Geraci and Bottai 2007). Median regression (50 % quantile regression) is particularly useful for heavy-tailed distributions, and by averaging predictive models across multiple quantiles, AQR constructs a more robust and adaptive forecasting approach (Shan and Yang 2009; Zhao and Xiao 2014).

Unlike prior studies that focus on quantile estimation within latent factor models, our approach leverages AQR specifically for prediction. While previous work has applied quantile regression to factor models (Harding et al. 2020; Chen et al. 2021), these studies primarily aim to estimate latent structures rather than improve forecasting performance. In contrast, our method integrates AQR into predictive modeling to capture the varying effects of predictors across different quantiles.

The key contribution of this study is to introduce AQR as a flexible and effective forecasting tool for high-dimensional, non-normally distributed time series. Although our approach is broadly applicable, we specifically demonstrate its advantages over conventional mean-based forecasting within a factor modeling framework. Given the high-dimensional nature of the data, factor models are particularly useful for reducing dimensionality while capturing the shared structure among the observed variables. Therefore, we propose an AQR-based forecasting approach within a factor modeling framework. Moreover, we systematically examine how variations in distributional shape influence AQR's performance by evaluating it under a diverse set of heavy-tailed, skewed, and near-normal conditions, thereby emphasizing its robustness across a wide range of data characteristics.

The remainder of the paper is organized as follows: In Section 2, we review the OLS based dynamic factor model (DFM) forecasting model. Then, proposed averaged quantile regression based forecasting method is presented in Section 3. In Section 4, we present the numerical experiments, and real data analysis is provided in Section 5. Finally, concluding remarks are presented in Section 6.

2 Review: OLS Based DFM Forecasting Model

In this section, we review the traditional DFM forecasting framework, which is commonly used for high-dimensional time series data. We consider the following DFM:

$$\mathbf{X}_t = \mathbf{\Lambda} \mathbf{F}_t + \boldsymbol{\epsilon}_t, \quad (1)$$

where $\mathbf{X}_t$ is an observed N -dimensional stochastic process, $\mathbf{F}_t$ is an unobservable latent factor, $\mathbf{\Lambda}$ is the $N \times r$ factor loading matrix, and $\boldsymbol{\epsilon}_t$ is an idiosyncratic component. The model assumes that the observable variables can be decomposed into a common component, which represents the shared variance across variables, and an idiosyncratic component, which accounts for the unique variance specific to each observed variable.

The forecasting is carried out by inputting the estimated common factors into the factor-augmented regression model as follows:

$$y_{t+1} = \beta_0 + \hat{\mathbf{F}}_t' \boldsymbol{\beta} + \xi_{t+1}, \quad (2)$$

where $\hat{\mathbf{F}}_t$ is estimated factor obtained from the method of Doz et al. (2011), $\boldsymbol{\beta} = (\beta_1, \dots, \beta_r)'$ is a regression coefficient vector, and ξ_{t+1} is i.i.d. error term. As in Stock and Watson (2002), we observe $\{y_t, \mathbf{X}_t\}_{t=1}^T$ and the goal is forecasting y_{T+1} using the DFM and regression model. If we consider a normally distributed error model, the following OLS method provides best linear unbiased estimator:

$$(\hat{\beta}_0, \hat{\boldsymbol{\beta}}) = \arg \min_{\beta_0, \boldsymbol{\beta}} \sum_{t=1}^T \left(y_t - \beta_0 - \hat{\mathbf{F}}_t' \boldsymbol{\beta} \right)^2.$$

Then, the OLS-based prediction value at time $T + 1$ is

$$\hat{y}_{T+1}^{OLS} = \hat{\beta}_0 + \hat{\mathbf{F}}_T' \hat{\boldsymbol{\beta}}.$$

3 Averaged Quantile Regression Based DFM Forecasting Model

In this paper, we extend the forecasting model in Section 2 and consider the weighted average of quantile regression estimators for the forecasting model. The rationale for this approach can be described as follows. For $Y \sim F_Y(y)$, it is well-known that $F_Y^{-1}(U)$ has the same cumulative distribution function (CDF) as Y when $U \sim \text{Unif}(0, 1)$. Therefore the expectation of Y , whenever it exists, is the same as the expectation of $F_Y^{-1}(U)$,

$$E(Y) = E(F_Y^{-1}(U)) = \int_0^1 F_Y^{-1}(u) du.$$

The relationship between the conditional quantile function and the mean is still true in the context of quantile regression. The conditional quantile function of Y given $\mathbf{X}$ is defined as.

$$Q_{Y|\mathbf{X}}(\tau) = F_{Y|\mathbf{X}}^{-1}(\tau) = \beta_{0\tau} + \mathbf{X}'\beta_\tau$$

where $F_{Y|\mathbf{X}}(y) = P(Y \leq y|\mathbf{X})$. Then, we have

$$E(Y|\mathbf{X}) = \int_0^1 Q_{Y|\mathbf{X}}(u) du,$$

and the integration can be approximated by the Riemann sums as

$$E(Y|\mathbf{X}) = \int_0^1 Q_{Y|\mathbf{X}}(u) du \approx \sum_i \Delta\tau_i Q_{Y|\mathbf{X}}(\tau_i).$$

This implies that we can estimate conditional expectation of Y given $\mathbf{X}$ by the weighted average of conditional quantile functions. Therefore, we first compute conditional quantile functions under the model (2), and their weighted average is our final estimator. That is, we consider τ th quantile regression model as follow:

$$(\hat{\beta}_{0\tau}, \hat{\beta}_\tau) = \arg \min_{\beta_0, \beta} \sum_{t=1}^T \rho_\tau(y_t - \beta_0 - \hat{\mathbf{F}}_t' \beta),$$

where $\rho_\tau(x) = x(\tau - I_{(x < 0)})$ is a τ -th loss function. Then, we combine quantile regression estimators under various quantiles (Lima and Meng 2017):

$$\hat{y}_{t+1} = \sum_{k=1}^g w_{\tau_k} (\hat{\beta}_{0\tau_k} + \hat{\mathbf{F}}_t' \hat{\beta}_{\tau_k}), \quad (3)$$

where $w_\tau \in (0, 1)$ is the weight value correspond to each quantile estimator satisfying $\sum_\tau w_\tau = 1$, and the $\{\tau_1, \tau_2, \dots, \tau_g\}$ is the discrete grid of quantiles.

4 Numerical Experiments

Model 1

To evaluate the performance of the proposed method, we consider the following model (Doz et al. 2011),

$$x_{it} = \sum_{j=1}^r \Lambda_{ij} f_{jt} + \epsilon_{it}, \quad i = 1, 2, \dots, N, \quad t = 1, 2, \dots, T,$$

and the factors follow an AR(1) model as $\mathbf{A}(L)\mathbf{F}_t = \mathbf{u}_t$, where $\mathbf{A}(L) = (1 - \phi L)\mathbf{I}_r$ and $\{\mathbf{u}_t\} \stackrel{\text{iid}}{\sim} N(0, (1 - \phi^2)\mathbf{I}_r)$. The idiosyncratic components ϵ_{it} follow an AR(1) model as

$$\epsilon_{it} = d\epsilon_{i(t-1)} + v_{it}, \quad i = 1, 2, \dots, N, \quad t = 1, 2, \dots, T,$$

where $\{\mathbf{v}_t\} \stackrel{\text{iid}}{\sim} N(0, \boldsymbol{\Psi})$. Here $\boldsymbol{\Psi}$ is an $N \times N$ matrix with $\psi_{ij} = \sqrt{\alpha_i \alpha_j} \kappa^{|i-j|} (1 - d^2)$, where $\alpha_i = \frac{\beta_i}{1 - \beta_i} \sum_{j=1}^r \Lambda_{ij}^2$, $\{\Lambda_{ij}\} \stackrel{\text{iid}}{\sim} N(0, 1)$, and the coefficient β_i is generated as $\{\beta_i\} \stackrel{\text{iid}}{\sim} \text{Unif}(u, 1 - u)$. The number of factor is set as $r = 2$, and we consider various sizes for models: $N = 10, 25, 50, 100$, and $T = 50, 100$. For other parameters, we use $\phi = 0.9, d = 0.5, \kappa = 0.5$, and $u = 0.1$ to satisfy the assumptions of the approximate DFM (Doz et al. 2011). Furthermore, we generate the scalar response variable to be forecasted as (Stock and Watson 2002),

$$y_{t+1} = \mathbf{1}'\mathbf{F}_t + \xi_{t+1}, \quad t = 1, \dots, T, \quad (4)$$

where $\mathbf{1}$ denotes a vector of 1's and $\{\xi_{t+1}\}$ are i.i.d. errors. We consider four different distributions for $\{\xi_{t+1}\}$: (i) $N(0, 1^2)$, (ii) t -distribution with degrees of freedom 2, (iii) and a mixture normal distribution, a mixture of $N(0, 1^2)$ and $N(5, 10^2)$ in a ratio of 9 to 1, and (iv) Cauchy(0,1) distribution.

For comparison, we consider the following methods.

- AQR: Our proposed quantile-based forecasting method that averages over three quantiles $\tau = 0.25, 0.50, 0.75$ with respective weights $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.
- QR: A single-quantile-based forecasting method that uses only the median quantile $\tau = 0.5$.
- OLS: OLS based forecasting model described in Section 2. We assume that the number of dynamic factors r is known as 2.

We perform 2000 Monte Carlo simulations and calculate the mean squared forecast error (MSFE) of the estimated value $\hat{y}_{T+1}$ over the true value y_{T+1} following the procedure in Stock and Watson (2002), and Table 1 presents the average MSFEs of the models. Under heavier-tailed and skewed distributions, such as mixture normal and Cauchy, the quantile-based forecasting methods, AQR and QR, yield lower MSFEs than OLS. This improvement stems from the ability of quantile regression to mitigate the influence of extreme observations. In contrast, OLS often exhibits higher MSFEs due to its sensitivity to outliers. When comparing QR and AQR across different sample sizes, our findings reveal that at $T = 50$, these two estimators produce similar outcomes, with QR occasionally exhibiting slightly better performance. However, once the training sample size increases to $T = 100$, AQR consistently provides more robust forecasts. This improvement can be attributed to the

Table 1: The average (and std. dev. in parentheses) of MSFEs by noise scenario (N , T), plus skewness and kurtosis. Boldface denotes the lowest MSFE—the best-performing model—in each row.

Dist. of ξ_t	N	T	Skewness	Kurtosis	OLS	QR	AQR
$N(0, 1)$	10	50	0.396	2.570	1.102(0.03)	1.135(0.03)	1.113(0.03)
	25		0.383	2.536	1.114(0.03)	1.146(0.03)	1.118(0.03)
	50		0.407	2.594	1.120(0.03)	1.145(0.03)	1.126(0.03)
	100		0.407	2.581	1.125(0.03)	1.159(0.03)	1.137(0.03)
	10	100	0.312	2.740	1.065(0.03)	1.080(0.03)	1.070(0.03)
	25		0.310	2.759	1.114(0.04)	1.124(0.04)	1.105(0.04)
	50		0.312	2.723	1.038(0.03)	0.954(0.03)	0.939(0.03)
	100		0.310	2.726	1.007(0.03)	1.019(0.03)	1.037(0.03)
$t(2)$	10	50	0.891	4.091	1.092(0.07)	1.088(0.07)	1.069(0.07)
	25		0.865	4.663	0.935(0.06)	0.909(0.05)	0.914(0.05)
	50		0.884	4.129	1.069(0.07)	1.050(0.07)	1.044(0.07)
	100		0.845	3.979	1.016(0.06)	1.002(0.06)	1.001(0.06)
	10	100	0.070	5.570	0.909(0.07)	0.892(0.07)	0.871(0.07)
	25		0.071	6.412	0.970(0.07)	0.957(0.07)	0.940(0.06)
	50		0.050	4.129	0.899(0.05)	0.882(0.05)	0.870(0.05)
	100		0.080	3.986	0.985(0.09)	0.974(0.08)	0.960(0.08)
Mixture normal	10	50	1.338	5.476	0.950(0.07)	0.907(0.07)	0.914(0.07)
	25		1.357	5.516	1.086(0.08)	1.059(0.08)	1.057(0.08)
	50		1.388	5.615	1.118(0.09)	1.086(0.09)	1.087(0.09)
	100		1.355	5.502	1.170(0.09)	1.133(0.09)	1.138(0.09)
	10	100	1.855	9.508	0.963(0.08)	0.945(0.08)	0.908(0.08)
	25		1.924	9.801	1.094(0.10)	1.085(0.10)	1.058(0.16)
	50		1.945	9.946	1.078(0.09)	1.068(0.09)	1.057(0.09)
	100		1.901	9.721	1.056(0.09)	1.037(0.09)	1.029(0.09)
Cauchy(0,1)	10	50	1.652	6.717	1.025(0.11)	0.974(0.11)	0.974(0.11)
	25		1.714	6.909	1.184(0.13)	1.139(0.13)	1.141(0.13)
	50		1.652	6.697	0.951(0.10)	0.890(0.10)	0.903(0.10)
	100		1.699	6.844	0.916(0.11)	0.858(0.11)	0.861(0.11)
	10	100	2.548	14.392	1.024(0.07)	0.991(0.06)	0.964(0.07)
	25		1.714	14.380	1.191(0.08)	1.119(0.08)	1.103(0.08)
	50		1.652	13.939	0.951(0.10)	0.919(0.09)	0.902(0.09)
	100		1.699	14.389	0.829(0.12)	0.798(0.12)	0.790(0.12)

fact that the additional data not only refine estimates for each individual quantile but also amplify the benefits of averaging across multiple quantiles. With more observations, outliers are identified and incorporated more accurately, and the multi-quantile framework gains greater statistical power to mitigate distortions arising from extreme values. This result is closely related to the idea of combining multiple quantile estimators, as proposed by Shan and Yang (2009), where such a

combination can yield substantial accuracy gains under global performance measures. Notably, when the data-generating process exhibits considerable skewness or kurtosis, even the median quantile may underestimate or overestimate the true central tendency if outlier effects remain unaddressed, reinforcing the advantages of an averaged quantile approach in larger samples.

For normally distributed data, AQR provides a reasonable central estimate but may be less efficient than OLS, consistent with the known optimality of OLS under strictly Gaussian errors. Nonetheless, AQR's performance does not degrade substantially, indicating that its robustness does not entail an excessively high variance cost in more conventional settings.

Overall, the proposed AQR method yields its greatest advantages in heavy-tailed or skewed environments, where averaging across multiple quantiles leads to notably lower MSFEs compared to either QR or OLS. By effectively balancing distributional extremes, AQR enhances robustness to outliers and asymmetry. In practical forecasting scenarios where deviations from normality are common, AQR thus provides a flexible and reliable alternative that performs well across a wide range of error distributions.

Model 2

While Model 1 evaluates AQR under a fixed set of representative distributions, we now extend our analysis to systematically explore how varying tail thickness and skewness affect forecasting performance. To this end, we adopt stable-distributed errors in our factor model, whose characteristic function is given by.

$$\mathbb{E}[e^{itZ}] = \exp \{ i \delta t - \gamma^\alpha |t|^\alpha [1 - i \beta \operatorname{sgn}(t) \Phi(\alpha, t)] \},$$

where $\alpha \in (0, 2]$ governs the tail thickness, $\beta \in [-1, 1]$ controls skewness, and $\gamma > 0, \delta \in \mathbb{R}$ are scale and location parameters, respectively. Notably, when $\alpha = 2, \beta = 0$, this distribution reduces to the Gaussian case, and when $\alpha = 1, \beta = 0$, it becomes a Cauchy distribution. By choosing smaller α (e.g. $\alpha = 0.4$), we induce heavier tails, while setting $\beta = \pm 1$ introduces extreme skewness.

We adopt the same factor model setup as in Model 1, retaining the AR(1) structure for both latent factors and idiosyncratic components, and set the number of factors $r = 2$, fixing $N = 100$ and $T = 100$. We also use the same parameter choices (e.g. $\phi = 0.9, d = 0.5, \kappa = 0.5, u = 0.1$). However, we now draw each scalar error term ξ_{t+1} in (4) from a stable distribution, varying $\alpha \in \{0.4, 0.8, 1.2, 1.6, 2.0\}$ and $\beta \in \{-1, -0.5, 0, 0.5, 1\}$. This setting spans a wide range of heavy-tailed or skewed scenarios, enabling a more comprehensive examination of forecast robustness.

Table 2 presents the MSFEs obtained from OLS, QR, and our AQR across the full grid of stable-distribution parameters. When the tail-thickness index satisfies $\alpha < 2$,

Table 2: The average and standard deviation (in parentheses) of MSFEs according to various α and β . Boldface denotes the lowest MSFE—the best-performing model—in each row.

α	β	Skew	Kurt	OLS	QR	AQR
0.4	−1.0	−2.917	10.345	1.057 (0.11)	0.979 (0.11)	0.978 (0.11)
0.4	−0.5	−1.528	10.665	1.299 (0.12)	1.203 (0.12)	1.206 (0.12)
0.4	0.0	0.013	10.668	1.201 (0.12)	1.104 (0.12)	1.109 (0.12)
0.4	0.5	1.512	10.531	1.008 (0.11)	0.929 (0.11)	0.929 (0.11)
0.4	1.0	2.937	10.454	1.175 (0.12)	1.114 (0.12)	1.114 (0.12)
0.8	−1.0	−2.249	7.909	1.029 (0.09)	1.000 (0.10)	0.981 (0.09)
0.8	−0.5	−1.118	7.966	1.103 (0.09)	1.028 (0.10)	1.036 (0.10)
0.8	0.0	−0.003	8.199	1.063 (0.09)	0.980 (0.09)	0.983 (0.09)
0.8	0.5	1.114	8.118	1.018 (0.08)	0.939 (0.09)	0.946 (0.09)
0.8	1.0	2.251	7.894	0.920 (0.08)	0.867 (0.09)	0.859 (0.09)
1.2	−1.0	−1.453	5.575	1.216 (0.08)	1.210 (0.09)	1.193 (0.09)
1.2	−0.5	−0.706	5.577	1.165 (0.08)	1.116 (0.08)	1.112 (0.08)
1.2	0.0	−0.015	5.754	0.983 (0.08)	0.919 (0.07)	0.921 (0.07)
1.2	0.5	0.744	5.724	1.050 (0.07)	1.005 (0.08)	1.010 (0.08)
1.2	1.0	1.440	5.539	1.135 (0.08)	1.129 (0.08)	1.119 (0.08)
1.6	−1.0	−0.688	3.793	1.133 (0.06)	1.145 (0.06)	1.126 (0.06)
1.6	−0.5	−0.340	3.790	1.031 (0.06)	1.026 (0.05)	1.013 (0.05)
1.6	0.0	−0.009	3.900	1.161 (0.05)	1.145 (0.06)	1.137 (0.06)
1.6	0.5	0.326	3.781	1.208 (0.06)	1.205 (0.06)	1.191 (0.06)
1.6	1.0	0.679	3.744	1.104 (0.03)	1.107 (0.03)	1.087 (0.03)
2.0	−1.0	0.002	2.583	1.077 (0.03)	1.125 (0.03)	1.085 (0.03)
2.0	−0.5	0.019	2.586	1.078 (0.03)	1.120 (0.03)	1.091 (0.03)
2.0	0.0	0.007	2.575	1.040 (0.03)	1.087 (0.03)	1.053 (0.03)
2.0	0.5	0.004	2.604	1.064 (0.03)	1.110 (0.03)	1.078 (0.03)
2.0	1.0	0.004	2.601	1.083 (0.03)	1.111 (0.03)	1.094 (0.03)

indicating heavier tails than the Gaussian benchmark, the two quantile-based methods consistently outperform OLS, and their advantage grows as α becomes smaller. Conversely, under Gaussian errors with $\alpha = 2$, OLS shows better performance than the other estimators, consistent with its known efficiency under normality.

For heavy-tailed designs (e.g. $\alpha \leq 1.2$) the median-only QR and the AQR deliver almost identical MSFEs when the skewness parameter β is close to zero. As $|\beta|$ grows and the error distribution becomes more asymmetric, AQR starts to out-perform QR because the equal-weight combination of three central quantiles pools information from both tails and dampens the leverage exerted by whichever tail contains the most extreme observations. This pattern corroborates the risk-reduction arguments of Shan and Yang (2009), who show that forecast combinations

of quantile models enhance overall accuracy, and it is also consistent with the semi-parametric efficiency results for optimally combined multi-quantile estimators in Zhao and Xiao (2014).

As the error distribution moves closer to normality ($\alpha = 1.6$ or 2), the MSFE of QR rises sharply, whereas the performance of AQR remains virtually indistinguishable from that of OLS. The AQR preserves the robustness of median regression in heavy-tailed settings yet avoids the variance inflation that hampers a single-quantile forecast when the errors are nearly Gaussian. Across the entire grid of shapes considered, tail thickness dictates whether a quantile-based approach is preferred to OLS, and skewness mainly influences the size of the advantage enjoyed by the averaged or single-quantile specification.

Overall, AQR retains a high degree of robustness in the presence of tail risk and incurs only a minor loss of efficiency as the error distribution approaches normality. Accordingly, it represents a judicious baseline choice whenever the true error distribution is unknown.

5 Empirical Application to US Economic Data

In this section we apply our forecasting framework to the United States Consumer Price Index (CPI) – a widely used gauge of macro-economic conditions and a principal input into public cost-of-living adjustments. The dataset contains 55 monthly variables (see Table 3) ranging from monetary aggregates and labour-market indicators to price, interest-rate, and uncertainty measures for both the United States and the euro area (Monteforte and Moretti 2013). The panel spans January 2001–December 2022, thereby encompassing multiple business-cycle phases and major policy episodes.

To investigate how forecast accuracy depends on the surrounding economic climate, we extract four sub-periods centered on well-documented macro-financial shocks. Drawing on the chronologies of Mian and Sufi (2009) and Ivashina and Scharfstein (2010), the intervals 2005 Q3–2007 Q4 and 2007 Q3–2009 Q2 are classified as high-volatility regimes associated with the credit-boom run-up and the core of the Global Financial Crisis, respectively. In contrast, the post-2014 oil-price slump (2014 Q1–2018 Q4; Baumeister and Kilian (2016) and the calm immediately preceding the COVID-19 pandemic (2018 Q1–2019 Q4; Baker et al. (2020) serve as low-volatility benchmarks. These four windows enable a clean comparison of OLS, QR and the proposed AQR under different tail-risk environments.

In line with the simulation study, AQR is implemented with an equal-weight scheme $\mathbf{w}_\tau = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ applied to the quantile set $\tau_g = (0.25, 0.50, 0.75)$.

Table 3: Economic indicators.

Category	Index	Units
Money supply	M1, M2	Billions of dollars
	Monetary base (total)	Millions of dollars
	M1, M3 for euro area	Euros
	M2 for euro area	Euros
	Deposits (all commercial banks)	Billions of U.S. Dollars
	Federal funds effective rate	Percent
Interest rate	Market yield on U.S. Treasury securities at constant maturity	Percent
	Moody's seasoned Aaa corporate bond yield	Percent
	Treasury constant maturity difference	Percent
	Bank prime loan rate	Percent
	10-Year real interest rate	Percent
	3-Month treasury bill secondary market rate	Percent
	Interest rates for euro area	Percent
Excess demand	Industrial production (total)	Index 2017 = 100
	Federal surplus or deficit	Millions of dollars
	Unemployment rate	Percent
	Harmonised unemployment for euro area	Percent
	Labor force participation rate	Percent
	Total vehicle sales	Thousands of units
	New privately-owned housing units Started(Total)	Thousands of units
Supply cost	Eggs, grade A, large (cost per dozen) in U.S. City average	U.S. Dollars
	Crude oil prices (brent, WTI)	Dollars per barrel
	U.S. Average(All items, food and energy, used cars and trucks)	Index 1982–1984 = 100
	Producer price index (industry, all commodities)	Index Dec 1980 = 100
	Harmonized index of consumer prices for euro area	Index 2015 = 100
Currency	U.S. Dollars to euro spot exchange rate	U.S. Dollars to one euro
	Real broad effective exchange rate for euro area	Index 2020 = 100
Prediction index	Inflation expectation	Percent
	Smoothed U.S. Recession probabilities	Percent
	5-Year forward inflation expectation rate	Percent
	5-Year breakeven inflation rate	Percent
	Economic policy uncertainty index for europe	Index

Table 4: Average one-quarter-ahead MSFEs for U.S. CPI. Boldface denotes the lowest MSFE—the best-performing model—in each row.

Volatility	Period	Skew	Kurt	OLS	QR	AQR	Description
High	2005Q3–2007Q4	0.862	3.088	1.504	1.301	1.440	Pre-crisis build-up
	2007Q3–2009Q2	1.344	4.269	3.374	2.751	3.058	Global financial crisis core
Low	2014Q1–2018Q4	0.337	1.729	1.398	1.778	1.190	Low-volatility expansion
	2018Q1–2019Q4	0.290	1.849	0.931	1.175	0.847	Pre-pandemic calm

For out-of-sample assessment we adopt a rolling-window cross-validation procedure. At each evaluation quarter within a given sub-period the dynamic-factor model is re-estimated on the most recent 16 quarters, after which a one-quarter-ahead CPI forecast is generated. The window then moves forward by one quarter and the exercise is repeated until the sub-period is exhausted. The resulting series of forecast errors is summarised by the MSFE reported in Table 4. In the turbulent sub-periods, where the forecast-error distribution shows pronounced skewness and excess kurtosis, both quantile-based methods outperform OLS, with the QR achieving the lowest errors. In the tranquil intervals the hierarchy changes and the accuracy of QR falls below that of OLS, while AQR continues to match or exceed the mean-based benchmark. The evidence therefore confirms that exploiting the conditional quantile function becomes critical once heavy tails emerge, and that averaging several central quantiles protects AQR from the efficiency losses a single-quantile forecast can suffer when the distribution returns to near-Gaussian. AQR thus secures meaningful gains under elevated tail risk without relinquishing precision in placid episodes, underscoring its practical value when the underlying error shape is uncertain.

6 Conclusions

In this study, we consider an averaged quantile forecasting model within a dynamic factor framework to tackle heavy-tailed or non-normally distributed time series. By combining multiple quantiles rather than relying solely on mean-based or single-quantile approaches, the proposed method adapts more effectively to skewness and outliers, as demonstrated through our two simulation designs and real data application.

We examined a range of heavy-tailed and skewed distributions, systematically varying both tail thickness and skewness, and found that averaging multiple quantiles yields more robust forecasts than single-quantile or conventional mean-based methods. We also note the similarity of our approach to that of Lima and Meng

(2017), who propose averaging conditional quantiles to estimate conditional means. However, our work differs notably in focusing on high-dimensional dynamic factor models and systematically investigating robustness under heavy-tailed and skewed errors. Moreover, our weighting schemes, simulation settings, and empirical applications distinguish our contribution from theirs. This explicit comparison clarifies the specific contributions of our approach within the broader quantile regression literature.

Despite these encouraging results, challenges remain. Determining the optimal set of quantile levels and associated weights is a key area for future work. A more data-adaptive approach, perhaps by quantifying the informational value or entropy across different quantiles, could provide a principled basis for selecting both quantile grids and weighting schemes, leading to even more robust performance.

Finally, the source code used in this study is available on GitHub at https://github.com/taeyeon98/Averaged_Quantile_Regression, making our implementation openly accessible and encouraging further refinements in weighting design or integration with other factor-estimation techniques.

Research funding: This research was supported by the National Research Foundation of Korea (NRF) funded by the Korea government (2022R1F1A1074134, 2021R1A2C1091357).

References

- Baker, S. R., N. Bloom, S. J. Davis, K. Kost, M. Sammon, and T. Viratyosin. 2020. "The Unprecedented Stock Market Reaction to COVID-19." *The Review of Asset Pricing Studies* 10 (4): 742–58.
- Baumeister, C., and L. Kilian. 2016. "Lower Oil Prices and the U.S. Economy: Is This Time Different?" *Brookings Papers on Economic Activity* 2016 (2): 287–336.
- Chen, L., J. J. Dolado, and J. Gonzalo. 2021. "Quantile Factor Models." *Econometrica* 89 (2): 875–910.
- De Valk, S., D. de Mattos, and P. Ferreira. 2019. "Nowcasting: an R Package for Predicting Economic Variables Using Dynamic Factor Models." *The R Journal* 11 (1): 230–44.
- Doz, C., D. Giannone, and L. Reichlin. 2011. "A Two-step Estimator for Large Approximate Dynamic Factor Models Based on Kalman Filtering." *Journal of Econometrics* 164: 188–205.
- Geraci, M., and M. Bottai. 2007. "Quantile Regression for Longitudinal Data Using the Asymmetric Laplace Distribution." *Biostatistics* 8 (1): 140–54.
- Harding, M., C. Lamarche, and M. H. Pesaran. 2020. "Common Correlated Effects Estimation of Heterogeneous Dynamic Panel Quantile Regression Models." *Journal of Applied Econometrics* 35 (3): 294–314.
- Ivashina, V., and D. Scharfstein. 2010. "Loan Syndication and Credit Cycles." *The American Economic Review: Papers and Proceedings* 100 (2): 57–61.
- Koenker, R., and G. Bassett Jr. 1978. "Regression Quantiles." *Econometrica: Journal of the Econometric Society* 46 (1): 33–50.
- Koenker, R., and K. F. Hallock. 2001. "Quantile Regression." *The Journal of Economic Perspectives* 15 (4): 143–56.

- Lima, L. R., and F. Meng. 2017. "Out-Of-Sample Return Predictability: a Quantile Combination Approach." *Journal of Applied Econometrics* 32 (4): 877–95.
- Mandelbrot, B. 1963. "The Variation of Certain Speculative Prices." *Journal of Business* 36 (4): 394–419.
- Mian, A., and A. Sufi. 2009. "Household Leverage and the Recession of 2007 to 2009." In *Paper presented at the 10th Jacques Polak Annual Research Conference*. Washington, DC: IMF.
- Monteforte, L., and G. Moretti. 2013. "Real-Time Forecasts of Inflation: the Role of Financial Variables." *Journal of Forecasting* 32 (1): 51–61.
- Shan, K., and Y. Yang. 2009. "Combining Regression Quantile Estimators." *Statistica Sinica* 19 (3): 1171–91.
- Stock, H., and W. Watson. 2002. "Forecasting Using Principal Components from a Large Number of Predictors." *Journal of the American Statistical Association* 97 (460): 1167–79.
- Zhao, Z., and Z. Xiao. 2014. "Efficient Regressions via Optimally Combining Quantile Information." *Econometric Theory* 30 (6): 1272–314.