NBS Working Paper
3/2023

# Bias-Correction in Time Series Quantile Regression Models

Marián Vávra

NATIONAL
BANK OF
SLOVAKIA
EUROSYSTEM

www.nbs.sk

research@nbs.sk

This publication is available on the NBS website
www.nbs.sk/en/publications-issued-by-the-nbs/research-publications



ISSN 2585-9269 (online)

# Bias-Correction in Time Series Quantile Regression Models*

Marián Vávra†

April 17, 2023

**Abstract**

This paper examines the small sample properties of a linear programming estimator in time series quantile regression models. Under certain regularity conditions, the estimator produces consistent and asymptotically normally distributed estimates of model parameters. However, despite these desirable asymptotic properties, we find that the estimator performs rather poorly in small samples. We suggest the use of a subsampling method to correct for a bias and discuss a simple rule of thumb for setting a block size. Our simulation results show that the subsampling method can effectively reduce the bias at very low computational costs and without significantly increasing the root mean squared error of the estimated parameters. The importance of bias correction for economic policy is highlighted in a growth-at-risk application.

*Keywords:* Quantile autoregressive model; bias; subsampling; growth-at-risk.
*JEL-Codes:* C15, C22.

---

*The author is grateful to Paulo M. M. Rodrigues, Ján Klacso, and Pavel Gertler for their helpful comments on an earlier version of the paper.
†National Bank of Slovakia. E-mail: marian.vavra@nbs.sk.

# 1. Introduction

Although the vast majority of regression applications in economics are concerned with modelling a conditional mean of economic variables, there is a rapidly growing interest in modelling other aspects of a conditional distribution of these variables. A quantile regression, which allows for a flexible analysis of the conditional quantile distribution of the variables under consideration, has become a workhorse method in applied economics. The examples include: autoregressive modelling (Koenker and Xiao (2006); Tsong and Lee (2011); Mumtaz and Surico (2015); Montes-Rojas (2019)), unit root testing (Koenker and Xiao (2004)), testing for causality (Troster (2018); Song and Taamouti (2021)), or forecasting (Manzan (2015); Korobilis (2017)), macroeconomic tail risk modelling (Adrian, Boyarchenko, and Giannone (2019); Adams, Adrian, Boyarchenko, and Giannone (2021); Chavleishvili and Manganelli (2019)), just to name a few.

It is well known that under certain regularity conditions, quantile regression estimators with time series data produce consistent and asymptotically normally distributed estimates of model parameters (see Koenker and Xiao (2006); Gregory, Lahiri, and Nordman (2018); Galvao, Montes-Rojas, and Park (2009)). However, many consistent estimators used in time series analysis produce severely biased estimates in small samples[1], something which has a negative impact on making statistical inference and forecasting from these models. Despite rapidly growing applications of quantile regression models in economics, very little is known about the behavior of quantile regression estimators with time series data in small samples usually encountered in applied macroeconomics.

The aim of this paper is twofold. First, we examine the small sample performance of a popular linear programming estimator in quantile autoregressive models by means of Monte Carlo simulations. Second, in contrast to similar studies that are limited to documenting a small sample bias in various time series models, we suggest the use of a subsampling method to correct for a bias and discuss a simple rule of thumb for

[1] See MacKinnon and Smith (1998) for linear autoregressive models, Psaradakis and Sola (1998) for Markov switching autoregressive models, Kapetanios (2000) for threshold autoregressive models, and Deb (1996) for conditional volatility models.

setting a block size.

The paper is organized as follows. Section 2 describes a quantile autoregressive regression model and discusses the asymptotic properties of the estimated parameters. Section 3 examines the small sample properties of the linear programming estimator in quantile autoregressions by means of Monte Carlo experiments. Section 4 discusses a subsampling method for a bias correction of the estimated quantile regression parameters. Section 5 discusses a growth-at-risk application to output growth. Section 6 summarizes and concludes.

# 2. Quantile Autoregressive Model

Given a real-valued time series $Y_1, \ldots, Y_n$, a $p$-th order quantile autoregressive model (QAR) can be written as follows

$$Y_t = \phi_0(U_t) + \phi_1(U_t)Y_{t-1} + \cdots + \phi_p(U_t)Y_{t-p}, \tag{1}$$

where $\{U_t\}$ are i.i.d. standard uniform random variables. The first term $\phi_0(U_t)$ can also be written as $\phi_0(U_t) = \phi_0 + \mathbb{F}^{-1}(U_t) = \phi_0 + \epsilon_t$, where $\mathbb{F}$ is a continuous distribution function of model errors. It is worth remarking that the QAR model in (1) can be interpreted as a specific type of functional-coefficient autoregressive (FCAR) model and nests an autoregressive conditional heteroskedasticity (ARCH) model in terms of the second-order properties (see Xiao (2012) for details). A sufficient condition for the QAR model in (1) to admit a stationary solution is that roots of the polynomial equation $c(z) = 1 - c_1 z - \cdots - c_p z^p$, where $c_i = \max_U(|\phi_i(U)|)$, for $i = 1, \ldots, p$, lie outside the unit disk.[2] The conditional quantile function of the model in (1) is given by

$$\mathbb{Q}(\tau|\mathcal{F}_{t-1}) = \phi_0(\tau) + \phi_1(\tau)Y_{t-1} + \cdots + \phi_p(\tau)Y_{t-p}, \tag{2}$$

where $\tau \in (0, 1)$ is the quantile parameter and $\mathcal{F}_{t-1}$ denotes the sigma-field that contains information up to and including time $t - 1$. The quantile function is sometimes denoted as $\mathbb{Q}_{t-1}(\tau)$ in the literature (we will use both terms interchangeably). The con-

[2] A weaker, yet less operational, stationarity condition is discussed in Koenker and Xiao (2006, p. 981).

ditional quantile function $\mathbb{Q}(\tau|\cdot)$ is supposed to be a monotonically increasing function in the quantile parameter $\tau$ in order to ensure a non-crossing property of the conditional quantiles and thus consistency of the estimated parameters (see Chernozhukov, Fernández-Val, and Galichon (2010) for further details).

For any quantile parameter $\tau \in (0,1)$, the estimated quantile regression parameters are a solution to the following minimizing problem:

$$\hat{\boldsymbol{\phi}}(\tau) = \underset{\boldsymbol{\phi}\in\mathbb{R}^{p+1}}{\text{argmin}} \sum_{t=1}^{n} \rho_\tau(Y_t - \boldsymbol{X}'_t\boldsymbol{\phi}), \tag{3}$$

where $\rho_\tau(z) = z\left[\tau - \mathbb{I}(z<0)\right]$ denotes the check function with $\mathbb{I}(\cdot)$ being a standard indicator function, $\boldsymbol{\phi} = (\phi_0, \phi_1, \ldots, \phi_p)'$ is a $(p+1)\times 1$ vector of model parameters and $\boldsymbol{X}_t = (1, Y_{t-1}, \ldots, Y_{t-p})'$ is a $(p+1)\times 1$ vector of lags of the dependent variable. It has been recognized for a long time that the optimization problem in (3) can be formulated as a linear programming problem and solved efficiently using either a simplex method or an interior point method. We opt for the latter method, which is more convenient, especially for large-scale problems and thus widely used in practice. The interested reader is referred to Chen and Wei (2005, pp. 401–404) for a detailed description of linear programming estimators in quantile regressions. It can be shown that under certain regularity conditions, quantile regression estimators produce consistent and asymptotically normally distributed estimates of model parameters (see Theorem 1 and conditions C1$-$C5 in Gregory, Lahiri, and Nordman (2018, pp. 1144-1145)). However, despite these desirable asymptotic properties, very little is known about the behavior of these estimators in small samples usually encountered in applied macroeconomics.

# 3. Simulation Study

In this section we present and discuss the results of a simulation study examining the small sample properties of the linear programming (LP) estimator in quantile autoregressive models under various data-generating mechanisms.

## 3.1. Experimental Design

We assess the performance of the LP estimator under different patterns of dependence by considering artificial data generated according to the first-order quantile autoregressive model

$$Y_t = \phi_0(U_t) + \phi_1(U_t)Y_{t-1},$$

with the following parameter configurations:

**M1**: $\phi_0(U_t) = 1.0 + \Phi^{-1}(U_t) = 1.0 + \epsilon_t, \quad \phi_1(U_t) = 0.5 - \frac{U_t}{4},$

**M2**: $\phi_0(U_t) = 1.0 + \Phi^{-1}(U_t) = 1.0 + \epsilon_t, \quad \phi_1(U_t) = 0.25 + \frac{U_t}{4},$

**M3**: $\phi_0(U_t) = 1.0 + \Phi^{-1}(U_t) = 1.0 + \epsilon_t, \quad \phi_1(U_t) = 0.8 - \frac{U_t}{2},$

**M4**: $\phi_0(U_t) = 1.0 + \Phi^{-1}(U_t) = 1.0 + \epsilon_t, \quad \phi_1(U_t) = 0.3 + \frac{U_t}{2},$

where $\{U_t\}$ are i.i.d. standard uniform random variables and $\Phi(\cdot)$ is a standard normal cumulative distribution function. All DGPs satisfy quantile monotonicity and stationarity conditions, but differ in generated persistence and heterogeneity: M1 (M2) exhibits relatively low, linearly decreasing (increasing), persistence across quantiles ($0.25 \leq \phi_1(\cdot) \leq 0.50$), whereas M3 (M4) exhibits strong, linearly decreasing (increasing), persistence with a higher degree of heterogeneity across quantiles ($0.30 \leq \phi_1(\cdot) \leq 0.80$).[3]

For each design point, $N = 100,000$ independent realizations of $\{Y_t\}$ of length $100 + n$, with $n \in \{100, 200, 500\}$, are generated. The first 100 data points of each realization are then discarded in order to eliminate start-up effects and the remaining $n$ data points are used to compute the quantile regression estimates and related quantities. The properties of the LP estimator are evaluated in 9 equally spaced quantile points $\tau \in \{0.1, \ldots, 0.9\}$.[4]

[3]Note that models M1 and M3 are calibrated according to a QAR(1) model estimated using the (annualized) quarterly real GDP growth rates for United States (M1) and Euro Area (M3) in the period spanning from 1985 Q1 to 2019 Q4 (see Adrian, Boyarchenko, and Giannone (2019, pp. 1272–1275) for a recent application). Models M2 and M4 are just modifications of models M1 and M3.

[4]Extreme quantiles (i.e. $\tau \to 0$ and $\tau \to 1$) are not considered here since they require special treatment (see Chernozhukov, Fernández-Val, and Kaji (2017) for details).

## 3.2. SIMULATION RESULTS

The simulation results in graphical form are reported in Figures 1–4, where bias functions of the estimated parameters are depicted over a pre-specified range of the quantile parameter $\tau$.[5] The simulation results reveal that the LP method produces severely biased estimates of the QAR model parameters regardless of the data persistence. In particular, the intercept $\phi_0(\tau)$ is systematically upward biased for all values of the quantile parameter $\tau$ considered, whereas the persistence parameter $\phi_1(\tau)$ is downward biased. The maximum bias amounts to up $11\%$ of the QAR model parameters, depending on the sample size n and the quantile parameter. Interestingly, the shape of the bias functions of $\phi_0(\tau)$ depends on the correlation between $\phi_0(\tau)$ and $\phi_1(\tau)$ parameters: in the case of positive correlation (models M2 and M4), the bias function of $\phi_0(\tau)$ is "flat" across the quantile parameter $\tau$, whereas in the case of negative correlation (models M1 and M3), the bias function of $\phi_0(\tau)$ is decreasing across the quantile parameter $\tau$. The shape of the bias function of $\phi_1(\tau)$ is proportional to the level of persistence: the higher the persistence the higher the bias (in absolute terms). Although the behavior of the estimated QAR model parameters improves quickly with the sample size $n$, 500 observations are actually needed to obtain (almost) unbiased parameters (see the results for model M3 depicted in Figure 3 where some remaining bias is still observed even for the sample of $500$ observations).

[5] For a true model parameter $\phi \in \boldsymbol{\phi}$, the bias function is calculated as $Bias(\phi) = N^{-1}\sum_{i=1}^{N}\hat{\phi}_i - \phi$, where $\hat{\phi}_i$ denotes the estimated parameter in the $i$-th Monte Carlo repetition.

Figure 1: Small Sample Properties of Parameters in Model M1

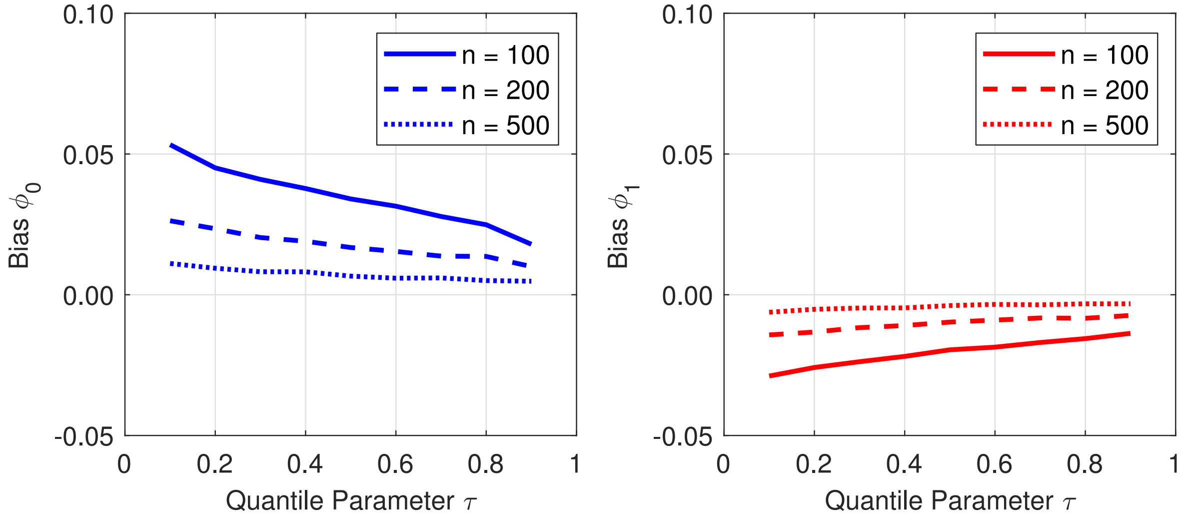


Figure 2: Small Sample Properties of Parameters in Model M2

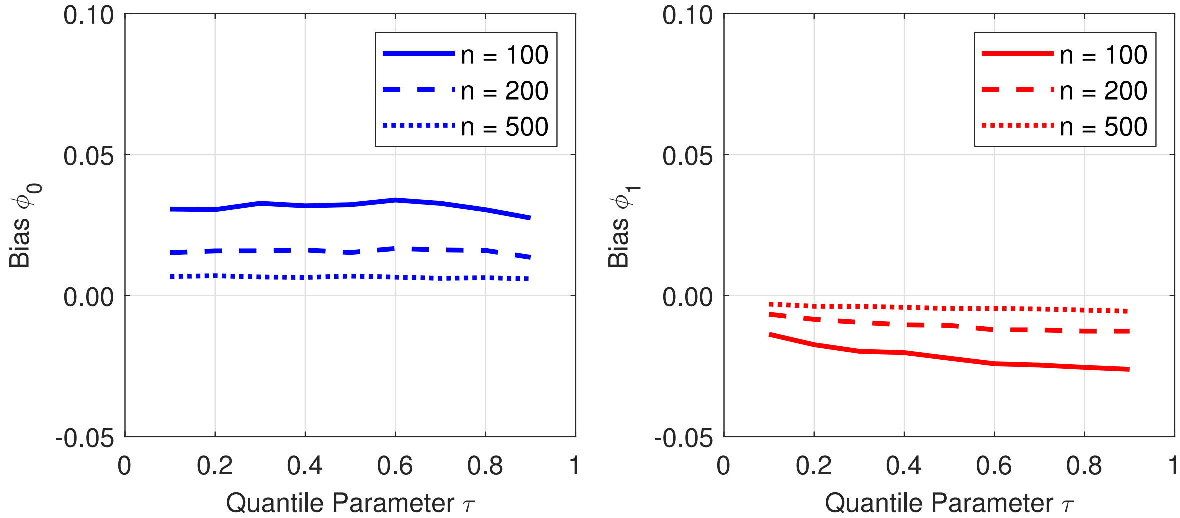

Figure 3: Small Sample Properties of Parameters in Model M3

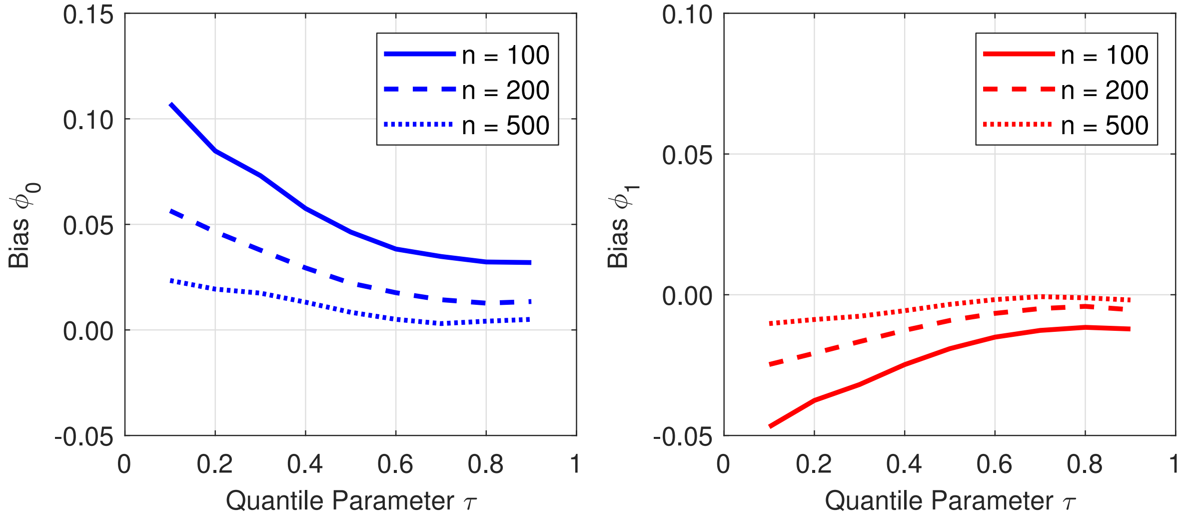


Figure 4: Small Sample Properties of Parameters in Model M4

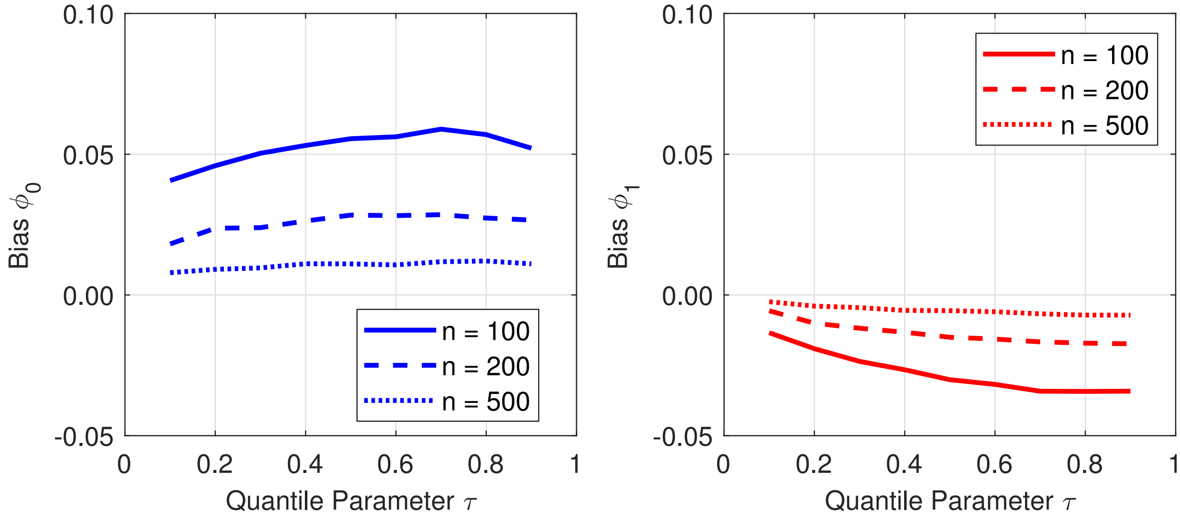

# 4. Bias Correction

We suggest using a subsampling method for adjusting a bias of the estimated quantile autoregressive model parameters.[6] Subsampling is a simple, yet powerful, resampling method that allows asymptotically valid inference under very general conditions (see Politis, Romano, and Wolf (1999)). The basic idea is to split the original series into overlapping subsamples and re-estimate the model in each of these subsamples. The approximate bias-adjusted parameters are then calculated as a linear combination of the full sample and subsample estimates, that is

$$\tilde{\phi} = \left(\frac{n}{n-\ell}\right)\hat{\phi} - \left(\frac{\ell}{n-\ell}\right)\frac{1}{n-\ell+1}\sum_{i=1}^{n-\ell+1}\hat{\phi}_i, \tag{4}$$

where $\tilde{\phi}$ denotes the bias-adjusted parameter vector, $\hat{\phi}$ is a vector of parameters estimated from the full-sample consisting of $n$ observations, $\hat{\phi}_i$ is a vector of parameters estimated from the $i$-th subsample consisting of just $\ell$ observations ($1 < \ell < n$). The interested reader is referred to Chambers (2013) for other bias-adjusting schemes.

An important issue that arises in the use of subsampling techniques in practice is the selection of a reasonable subsample size $\ell$ for a given sample size $n$, a problem akin to that of selecting the block length for blockwise bootstrap methods (see Lahiri (2003, Chap. 7)). Unfortunately, the asymptotic requirements that $\ell \to \infty$ and $\ell/n \to 0$ (as $n \to \infty$) give little guidance for the selection of an appropriate subsample size beyond the requirement that it grows at a slower rate than the sample size $n$. Several methods for selecting the subsampling size $\ell$ have been proposed in the literature (see Politis, Romano, and Wolf (1999, Chap. 9)). These methods (e.g. "minimum volatility" and "calibration") are designed namely for calculating confidence intervals and hypothesis testing about the estimated parameters and thus are not directly applicable to the prob-

[6] Although analytical bias expressions are easy to implement, while resampling methods are often computer intensive and involve some technical difficulties, most studies conducting bias-adjustment resort to resampling techniques. There are at least two good reasons in favor of resampling techniques. First, analytical bias expressions are usually based on first-order approximation techniques and thus provide only a local approximation to the true bias function (see, e.g., MacKinnon and Smith (1998, p. 210)). Second, analytical bias expressions are usually available only for simple (linear) time series models (see, e.g., Shaman and Stine (1988, p. 846)).

lem at hand - a bias correction. To overcome this difficulty, we follow Chernozhukov and Fernández-Val (2005) and employ a simple rule of thumb for selecting the subsample size $\ell$. The rule takes the following form

$$\ell = \lceil m + n^{\gamma} \rceil, \tag{5}$$

where $\gamma \in (0,1)$, $m > 0$ and $n$ denotes the sample size. Some comments on the rule are in order. The first term $m$ represents a minimum number of observations required for quantile regression estimates. We link $m$ with the tail quantiles whose estimation is data-demanding: $m = 1/\min(\mathcal{T})$, where $\mathcal{T} = \{\tau_1, 1-\tau_1, \ldots, \tau_k, 1-\tau_k\}$ is a set of all quantile parameters and their complements considered in the problem at hand (for some integer $1 \leq k < n$). In order to satisfy the condition that the block $\ell$ grows with the sample size $n$ but at a slower rate, Chernozhukov and Fernández-Val (2005, p. 264) recommended using $\gamma = 1/2$ based on some analytical results from Sakov and Bickel (2000).[7] The symbol $\lceil \cdot \rceil$ denotes the greatest integer function.

In order to keep computational costs of Monte Carlo simulations at a reasonable level, only the shortest sample size ($n = 100$) is considered when examining the performance of the subsampling-based bias adjusting procedure. Keeping in mind the sample size and the range of quantile points, the block size $\ell$ determined by the rule in (5) is $\ell = \lceil 1/0.1 + \sqrt{100} \rceil = 20$ observations.

The simulation results in graphical form are reported in Figures 5–8, where the bias functions are depicted for both estimated (labelled "Est") and bias adjusted (labelled "Adj") model parameters. It can be concluded from the simulation results that the subsampling method employed is capable of eliminating most of the bias from quantile regression parameters. As expected, the subsampling procedure is slightly less successful for highly-persistent stochastic processes (see model M3). Although a bias reduction is an important and attractive feature of the subsampling method, there are also other distributional properties of estimators worth considering (e.g. the root mean square error).[8] Our simulation results indicate that the subsampling method can effec-

[7]The authors confirmed that other values of $\gamma$ may also lead to reliable, powerful, and computationally attractive inference.

[8]For a true model parameter $\phi \in \boldsymbol{\phi}$, the root mean square error function is calculated as $RMSE(\phi) =$

tively reduce the bias without substantially increasing the root mean squared error of the estimated parameters.

As a robustness check, we investigate the sensitivity of the subsampling results on changes of the parameter $\gamma$ in the block size rule (5). Apart from our benchmark configuration of the rule with $\gamma = 1/2$, we also consider the rule with $\gamma = 1/4$ as in Chernozhukov and Fernández-Val (2005). This leads to the block size $\ell = \lceil 1/0.1 + \sqrt[4]{100} \rceil = 14$ observations. The simulation results for model M4 are reported in Figure 9. Extremely low sensitivity of the bias and root mean squared functions on changes in the the $\gamma$ parameter are observed. This fact can be partly explained by the constant term $m$ in the block-size rule which guarantees a minimum number of observations for estimation regardless of $\gamma$.

$\sqrt{N^{-1} \sum_{i=1}^{N} (\hat{\phi}_i - \phi)^2}$, where $\hat{\phi}_i$ denotes the estimated parameter in the $i$-th Monte Carlo repetition.

Figure 5: Properties of Bias Adjusted Parameters in Model M1

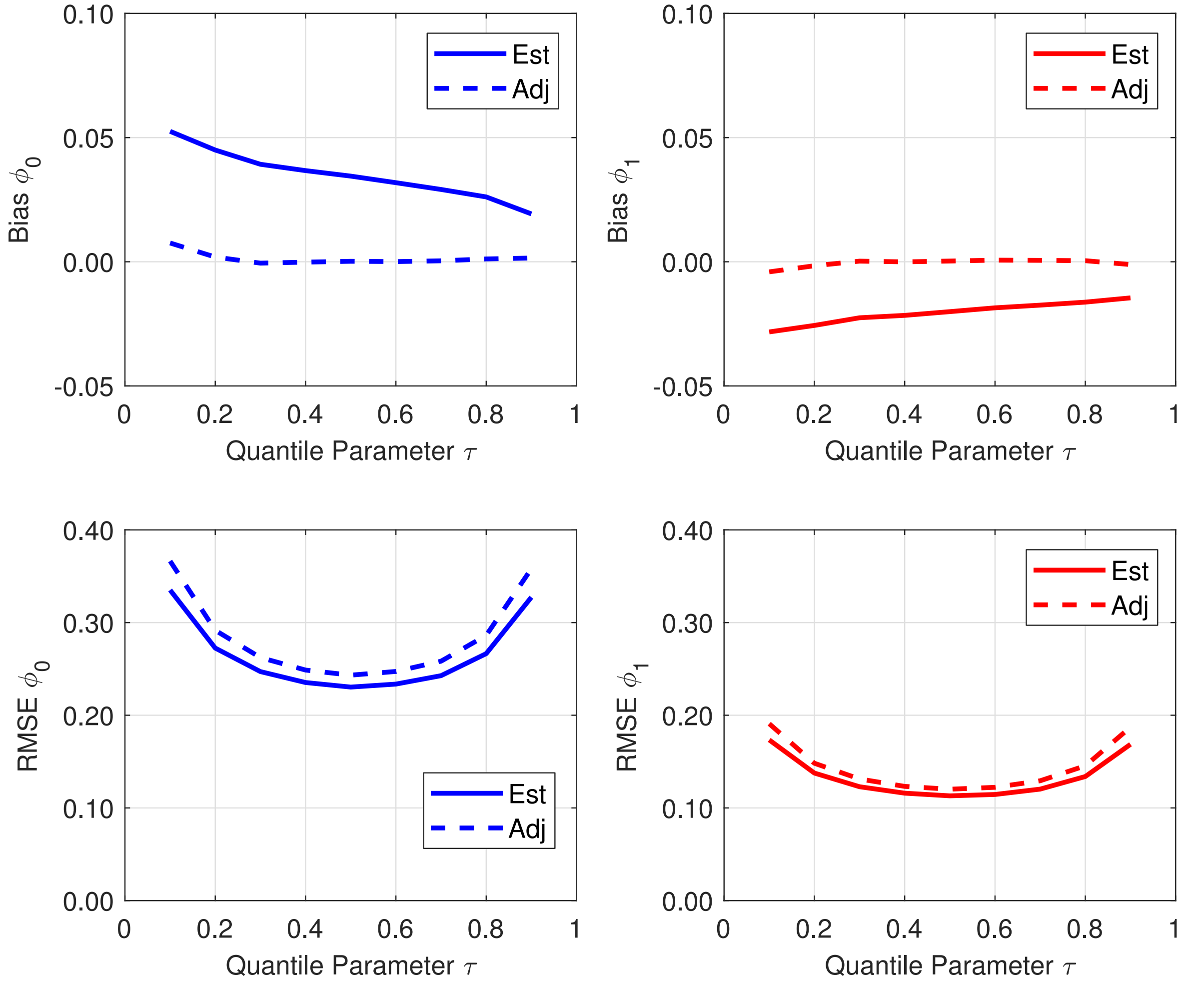

Figure 6: Properties of Bias Adjusted Parameters in Model M2

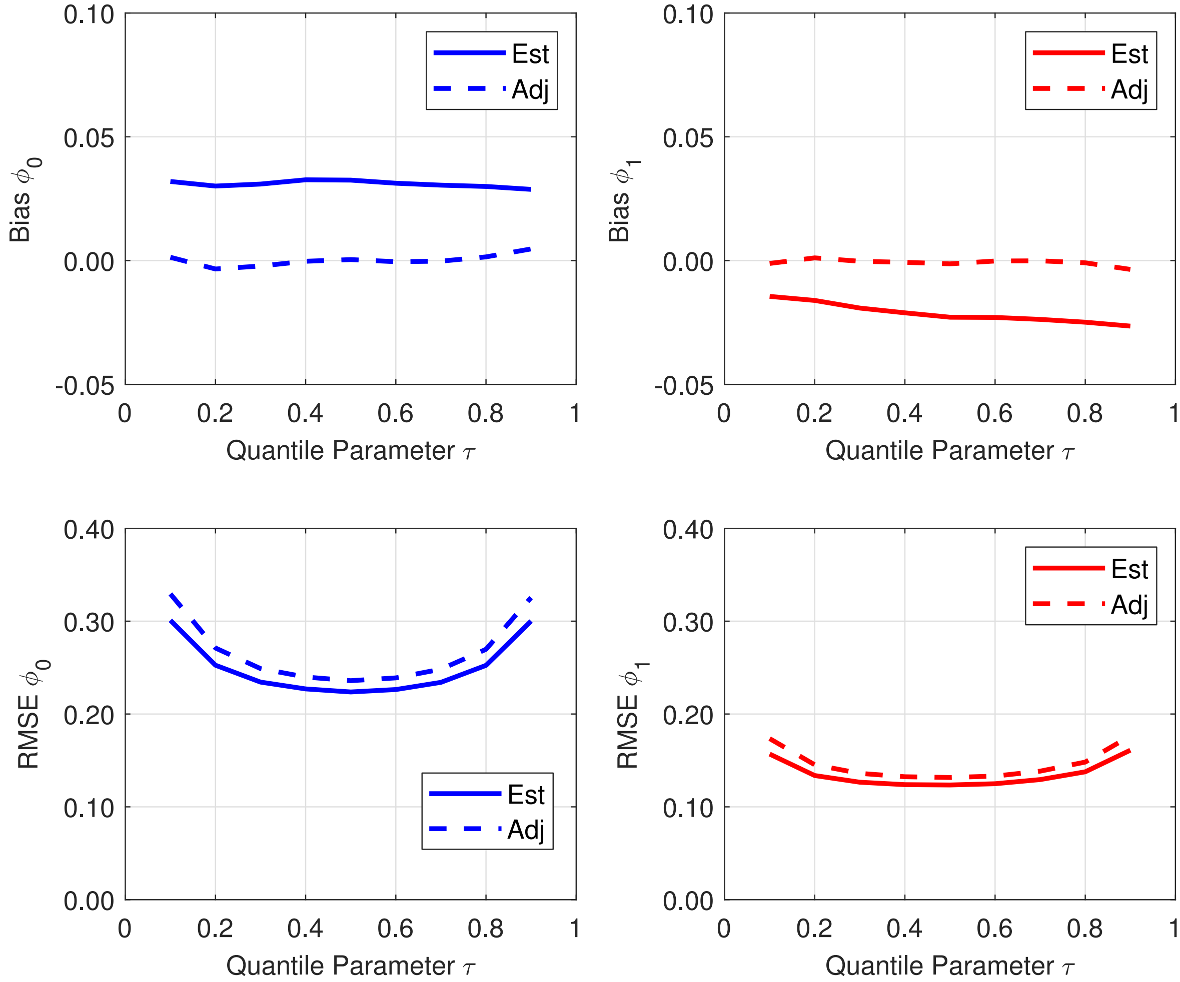

Figure 7: Properties of Bias Adjusted Parameters in Model M3

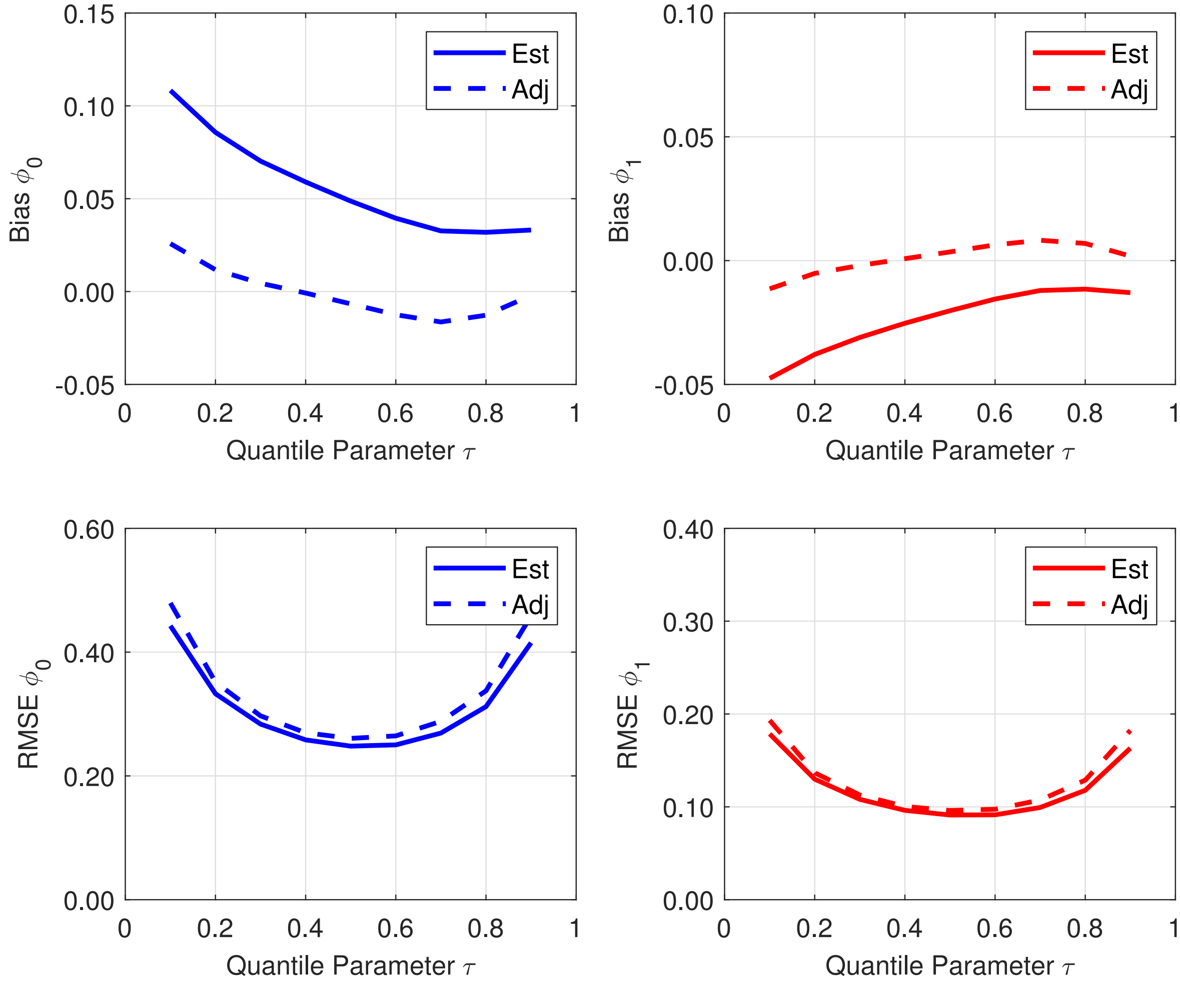

Figure 8: Properties of Bias Adjusted Parameters in Model M4

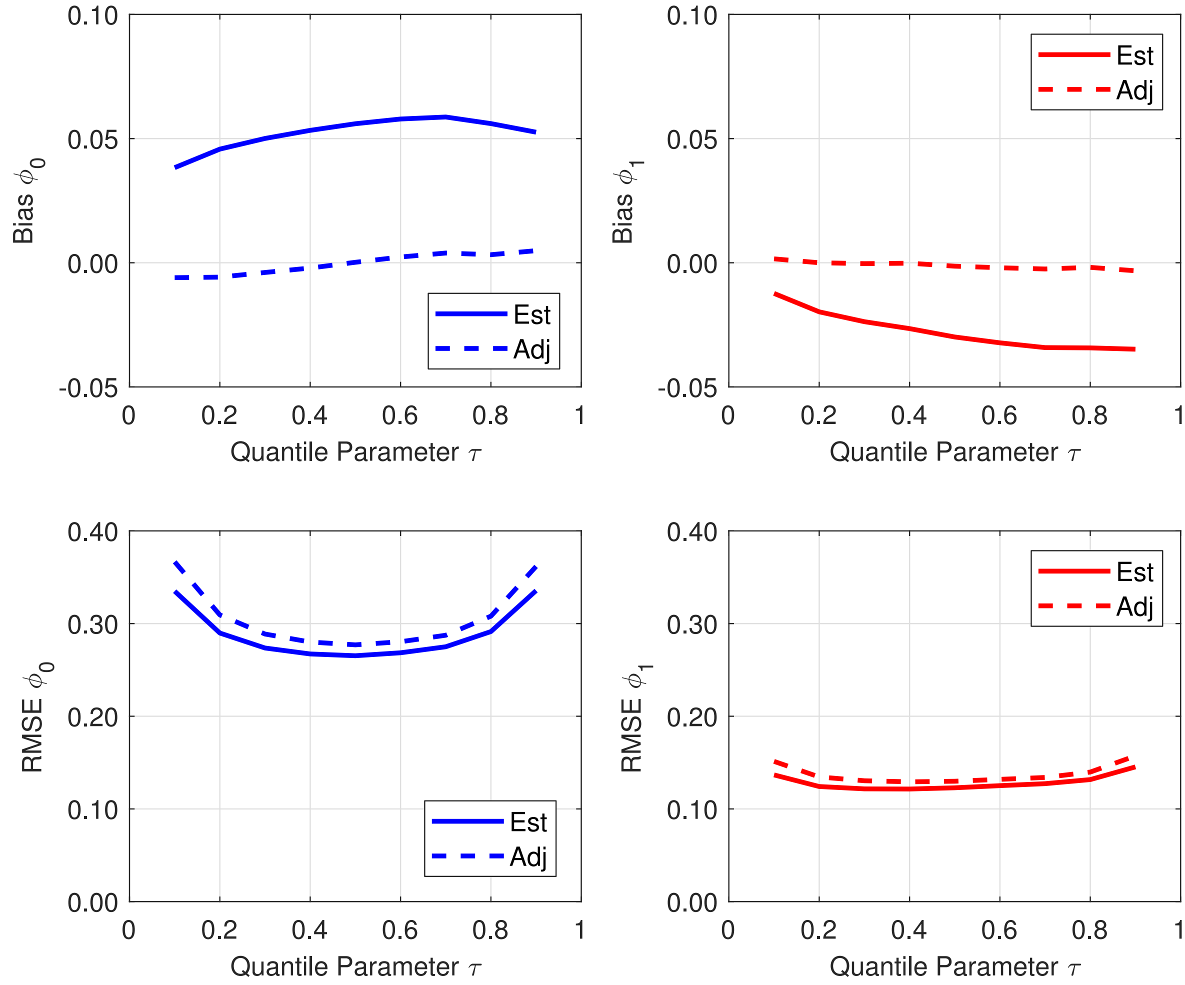

Figure 9: Sensitivity Analysis for Model M4

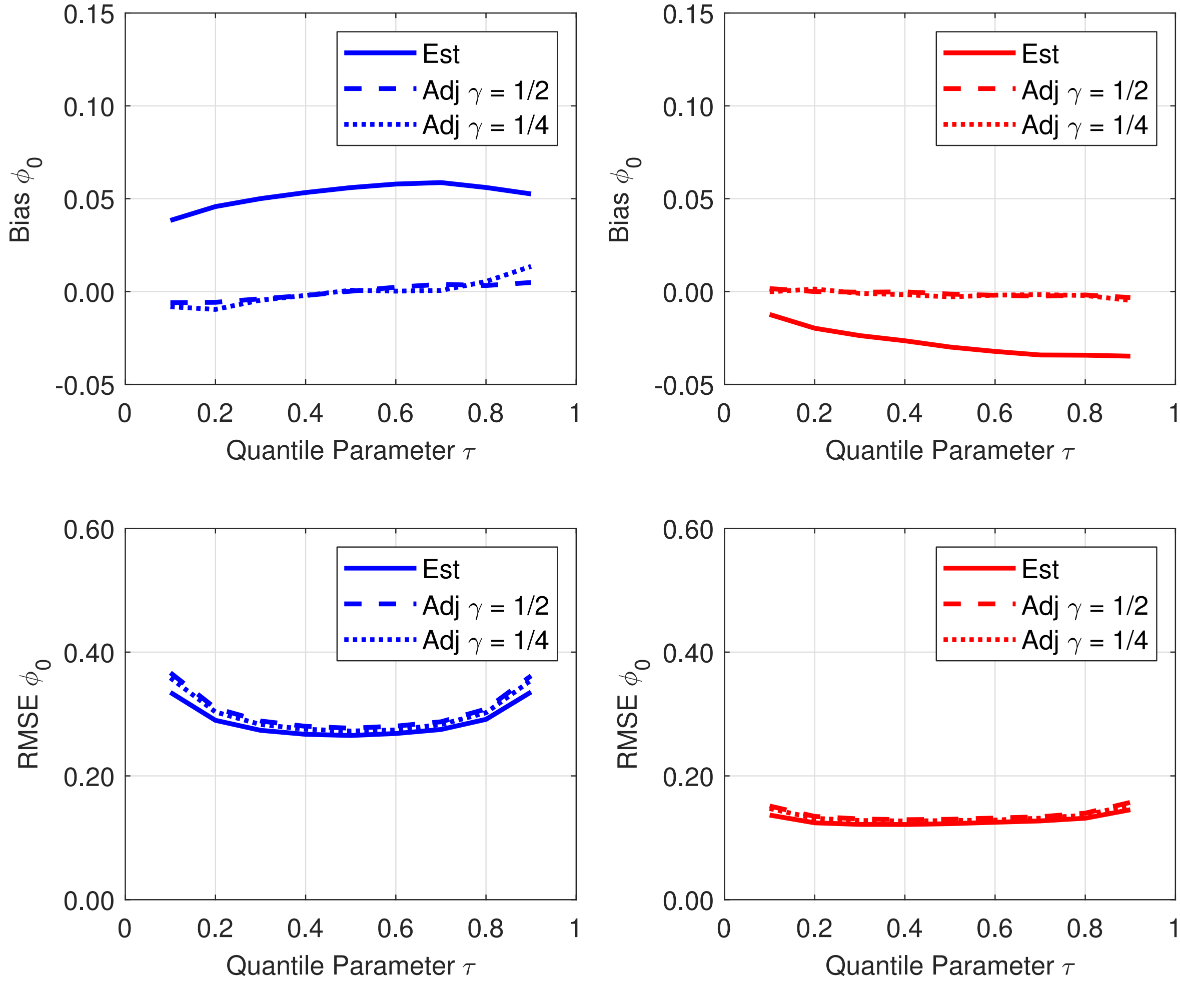

# 5. Empirical Example

Since the global economic recession in 2008, the focus of policymakers on modelling and forecasting macroeconomic tail risks has increased substantially. In order to account for possible non-linear behaviour of economic variables within a relatively simple modelling framework, a quantile regression has become a workhorse method for calculating the predictive distribution of economic variables (see, e.g., Adrian, Boyarchenko, and Giannone (2019); Chavleishvili and Manganelli (2019); Figueres and Jarociński (2020); Adrian, Boyarchenko, and Giannone (2021); Adams, Adrian, Boyarchenko, and Giannone (2021); Kiley (2022); Ferrara, Mogliani, and Sahuc (2022) to name just a few recent applications). The evolution of the predictive distribution can be used to formulate the concept of growth-at-risk (GaR) which is defined as the $\alpha$-percentile of the predictive distribution of economic variables (e.g. output growth).[9]

The density forecasting approach (including the GaR) offers a number of attractive features: (i) central bankers can benefit from information about the entire distribution of future output growth, encompassing both downside and upside macroeconomic risks, and going beyond more traditional point forecasts; (ii) it also helps central bankers to more readily communicate sources of macroeconomic risks to professionals and the public.

However, as shown in the previous section, quantile regression estimators produce biased estimates of model parameters in small samples, something which has a negative implications for making statistical inference and forecasting from these models. The objective of this empirical analysis is to assess the magnitude of the impact on the macroeconomic tail risk measure and the predictive distribution.

## 5.1. Data

As for the output series, the Euro Area 19 real gross domestic product (with fixed composition) is employed in our analysis. The data are seasonally adjusted, transformed

[9] The conventional values for $\alpha$ are $0.05$ or $0.10$.

into stationarity using the (annualized) quarter-on-quarter growth rates, and spanning the period from 1999 Q1 to 2019 Q4 ($n = 84$ observations).[10] The data can be downloaded from the ECB Statistical Data Warehouse. The Euro Area real GDP growth series altogether with long-run economic growth measured by a sample median (approx. $1.8\,\%$) are depicted in Figure 10.

Figure 10: Euro Area Real GDP Growth

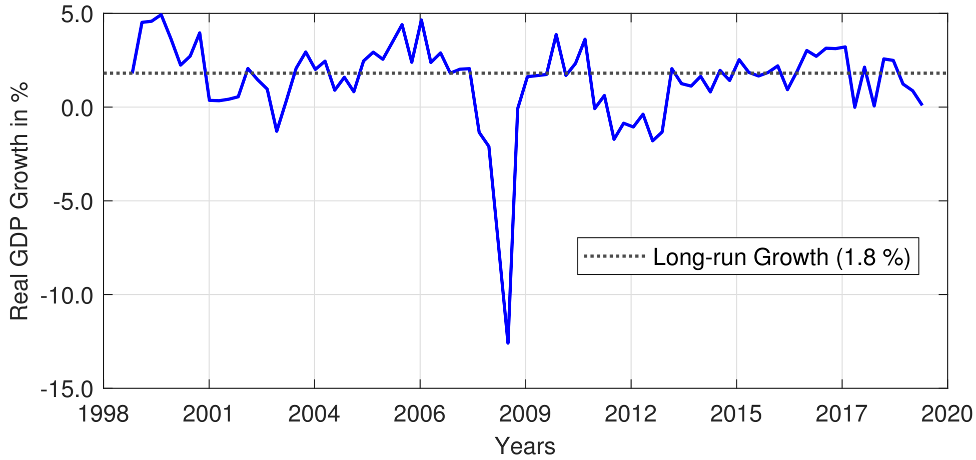


## 5.2. Empirical Results

In our empirical example, a stationary time series of the Euro Area real GDP growth rates $\{Y_t\}_{t=1}^{n}$ is supposed to follow a $p$-th order quantile autoregressive model given by

$$Y_t = \phi_0(U_t) + \phi_1(U_t)Y_{t-1} + \cdots + \phi_p(U_t)Y_{t-p}, \tag{6}$$

where $\{U_t\}$ are i.i.d. standard uniform random variables. The first term $\phi_0(U_t)$ can also be written as $\phi_0(U_t) = \phi_0 + \mathbb{F}^{-1}(U_t) = \phi_0 + \epsilon_t$, where $\mathbb{F}$ is a continuous distribution function of model errors (not necessarily Gaussian). The lag order $p$ is selected using a modified Bayesian information criterion within the range of possible values $p \in [1, 4]$ (see Galvao, Montes-Rojas, and Park (2013, p. 311)). Although the optimal lag may vary across the quantile parameter $\tau$, we follow the mainstream literature

[10] Since the last observations from the Covid-19 period might be classified as abrupt outliers, they are purposely excluded from the sample.

and select the lag order at the median level $\tau = 0.5$, resulting in $p = 1$. The estimated parameters of the QAR(1) model are reported in Table 1 (columns labelled "Estimated"). Seven representative quantile parameters are considered in our analysis: $\tau \in \{0.10, 0.25, 0.40, 0.50, 0.60, 0.75, 0.90\}$. For comparison, the bias-adjusted parameters, based on the subsampling method described earlier, are reported in Table 1 as well (columns labelled "Adjusted"). The block size $\ell = 20$ is determined according to the rule in (5). It can be concluded from the table that the intercept $\phi_0$ is upward biased (for most of the quantile parameters), whereas the persistence parameter $\phi_1$ is slightly downward biased. These empirical results are broadly in line with our Monte Carlo findings.[11]

Table 1: Estimated and Adjusted Quantile Regression Parameters

| | Estimated | | Adjusted | |
|---|---|---|---|---|
| $\tau$ | $\phi_0(\tau)$ | $\phi_1(\tau)$ | $\phi_0(\tau)$ | $\phi_1(\tau)$ |
| 0.10 | -2.35 | 1.02 | -2.63 | 1.15 |
| 0.25 | -0.21 | 0.62 | -0.13 | 0.63 |
| 0.40 | 0.21 | 0.62 | 0.12 | 0.66 |
| 0.50 | 0.67 | 0.62 | 0.57 | 0.68 |
| 0.60 | 1.11 | 0.53 | 1.07 | 0.56 |
| 0.75 | 2.04 | 0.34 | 2.16 | 0.31 |
| 0.90 | 3.04 | 0.25 | 3.23 | 0.19 |

For better understanding of the impact of biased quantile regression parameters on macroeconomic tail risks, the $\alpha$-level GaR values are reported in the graphical form as well. Formally, the one-step ahead $\alpha$-level GaR is defined as the $\alpha$-percentile of the predictive distribution of real GDP growth, that is

$$\mathbb{P}(Y_{n+1} \leq \mathbb{Q}_n(\alpha)) = \alpha \in (0, 1), \tag{7}$$

where $Y_{n+1}$ denotes output growth in period $n+1$, $\mathbb{Q}_n(\alpha)$ the conditional $\alpha$-level quantile function conditional on $\mathcal{F}_n$. The GaR is a forecast-oriented risk measure and can be understood as the "worst-case" estimate of future output growth - output growth at time $n+1$ lower than $\mathbb{Q}_n(\alpha)$ can occur only with probability $100\,\alpha$ percent.

[11] Differences in the bias functions of the constant term $\phi_0(\cdot)$ in the empirical example and Monte Carlo simulations can be explained by differences in $\mathbb{F}$ used in (6) and $\Phi$ used in Monte Carlo simulations.

The 10%-GaR (i.e. $\alpha = 0.1$) is calculated directly from (6) as $\mathbb{Q}_n(0.1) = \phi_0(0.1) + \phi_1(0.1)\, Y_n$ using both the estimated and bias-adjusted parameters. It is obvious that the (starting) value $Y_n$ plays an important role in calculating the predictive quantiles as well. For this reason, we calculated the 10%-GaR values conditional on a set of possible starting real GDP values $Y_n \in [-10, 5]$. The resulting 10%-GaR values are depicted in Figure 11(a). The vertical and horizontal axes correspond, respectively, to starting real GDP growth values and the reference GaR values.

It can be concluded from the figure that the bias of the quantile regression parameters translates into the tail risk measure $\mathbb{Q}_n(0.1)$ in a non-trivial way, making the unadjusted tail risk measure inaccurate and even misleading. In particular, the GaR values can be significantly underestimated in times of economic downturns (up to 2 percentage points), whilst slightly overestimated in times of high growth (see Figure 11(a)). Interestingly, the estimated (unadjusted) model parameters produce unbiased tail risk measures only around long-run (median) real GDP growth (see the vertical black dotted line at the level of 1.8 in Figure 11(a)).

Since there is no easy way to obtain a smooth probability density function from the estimated quantile function, the quantile-based measures of higher-order moments (i.e. dispersion, skewness, and kurtosis), are considered for assessing the impact of the biased parameters on the whole predictive distribution of real GDP growth (see Kim and White (2004)). Formally, these quantities are defined as follows:

- Dispersion: $\zeta_2 = \mathbb{Q}_n(0.75) - \mathbb{Q}_n(0.25)$,
- Skewness: $\zeta_3 = \dfrac{\mathbb{Q}_n(0.75) + \mathbb{Q}_n(0.25) - 2\mathbb{Q}_n(0.5)}{\mathbb{Q}_n(0.75) - \mathbb{Q}_n(0.25)}$,
- Excess kurtosis: $\zeta_4 = \dfrac{\mathbb{Q}_n(0.9) - \mathbb{Q}_n(0.6) + \mathbb{Q}_n(0.4) - \mathbb{Q}_n(0.1)}{\mathbb{Q}_n(0.75) - \mathbb{Q}_n(0.25)} - 1.52$,

where each quantile $\mathbb{Q}_n(\alpha)$ satisfies (7) for any $\alpha \in (0, 1)$. The value $1.52$ serves as a norming constant for kurtosis. As in the previous case, the quantile-based higher-order moments are calculated conditional on a set of possible starting real GDP values $Y_n \in [-10, 5]$. Both estimated and adjusted one-step ahead quantile-based moment measures are depicted in Figure 11(b-d). The vertical and horizontal axes correspond, respectively, to starting real GDP growth values and the selected higher-order moments of the predictive distribution.

It can be concluded from the figure that the bias of the quantile regression parameters translates into the higher-order moments and the whole predictive distribution in a non-linear way (see the impact on skewness $\zeta_3$ and excess kurtosis $\zeta_4$ in Figure 11(c-d)). In particular, especially the left-tail of the predictive distribution, which is of the crucial importance for policymakers, can be significantly underestimated in times of economic downturns.

Figure 11: Euro Area Tail Risk and Distributional Measures

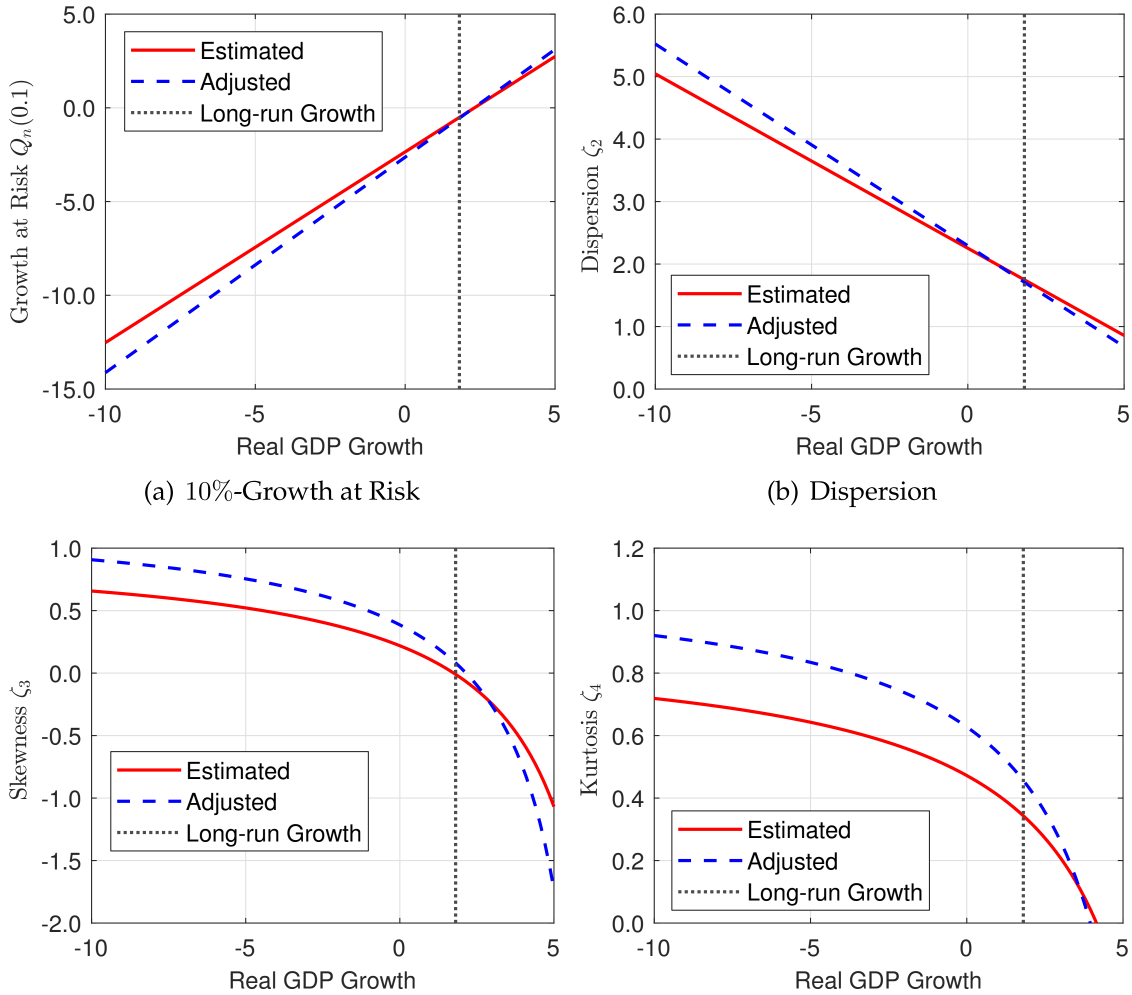


(a) 10%-Growth at Risk

(b) Dispersion

(c) Skewness

(d) Excess kurtosis

# 6. Conclusion

This paper has investigated the small sample properties of the quantile regression estimator with time series data. Our simulation results show that the estimator performs rather poorly in small samples. More than 500 observations are needed to obtain unbiased quantile regression parameters and thus reliable inference and accurate forecasts. It is also shown that subsampling can effectively reduce the bias at very low computational costs and without significantly increasing the root mean squared error of the estimated quantile regression parameters. The importance of bias correction for economic policy is highlighted in a growth-at-risk application.

Although the findings of this paper are specific to first-order quantile autoregressive models, we expect to encounter similar difficulties to arise in higher-order and/or multivariate quantile autoregressive models. In fact, the problems are likely to become more severe as the dimension of such models increases.

# References


Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021): "Forecasting macroeconomic risks," *International Journal of Forecasting*, 37, 1173–1191.

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable growth," *American Economic Review*, 109, 1263–89.

——— (2021): "Multimodality in macrofinancial dynamics," *International Economic Review*, 62, 861–886.

Chambers, M. J. (2013): "Jackknife estimation of stationary autoregressive models," *Journal of Econometrics*, 172, 142–157.

Chavleishvili, S., and S. Manganelli (2019): "Forecasting and stress testing with quantile vector autoregression," *ECB Working Paper*, 2330.

Chen, C., and Y. Wei (2005): "Computational issues for quantile regression," *Sankhyā: The Indian Journal of Statistics*, 67, 399–417.

Chernozhukov, V., and I. Fernández-Val (2005): "Subsampling inference on quantile regression processes," *Sankhyā: The Indian Journal of Statistics*, 67, 253–276.

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): "Quantile and probability curves without crossing," *Econometrica*, 78(3), 1093–1125.

Chernozhukov, V., I. Fernández-Val, and T. Kaji (2017): "Extremal quantile regression," in *Handbook of Quantile Regression*. Chapman and Hall/CRC.

Deb, P. (1996): "Finite sample properties of maximum likelihood and quasi-maximum likelihood estimators of EGARCH models," *Econometric Reviews*, 15, 51–68.

Ferrara, L., M. Mogliani, and J.-G. Sahuc (2022): "High-frequency monitoring of growth at risk," *International Journal of Forecasting*, 38, 582–595.

Figueres, J. M., and M. Jarociński (2020): "Vulnerable growth in the euro area: Measuring the financial conditions," *Economics Letters*, 191, 109126.

Galvao, A. F., G. Montes-Rojas, and S. Y. Park (2009): "Quantile autoregressive distributed lag model with an application to house price returns," *City University London, Department of Economics Discussion Paper Series*, 09/04.

——— (2013): "Quantile autoregressive distributed lag model with an application to house price returns," *Oxford Bulletin of Economics and Statistics*, 75, 307–321.

GREGORY, K. B., S. N. LAHIRI, AND D. J. NORDMAN (2018): "A smooth block bootstrap for quantile regression with time series," *The Annals of Statistics*, 46, 1138–1166.

KAPETANIOS, G. (2000): "Small sample properties of the conditional least squares estimator in SETAR models," *Economics Letters*, 69, 267–276.

KILEY, M. T. (2022): "Unemployment risk," *Journal of Money, Credit and Banking*, 54, 1407–1424.

KIM, T.-H., AND H. WHITE (2004): "On more robust estimation of skewness and kurtosis," *Finance Research Letters*, 1, 56–73.

KOENKER, R., AND Z. XIAO (2004): "Unit root quantile autoregression inference," *Journal of the American Statistical Association*, 99, 775–787.

——— (2006): "Quantile autoregression," *Journal of the American Statistical Association*, 101, 980–990.

KOROBILIS, D. (2017): "Quantile regression forecasts of inflation under model uncertainty," *International Journal of Forecasting*, 33, 11–20.

LAHIRI, S. (2003): *Resampling Methods for Dependent Data*. Springer.

MACKINNON, J., AND A. SMITH (1998): "Approximate bias correction in econometrics," *Journal of Econometrics*, 85, 205–230.

MANZAN, S. (2015): "Forecasting the distribution of economic variables in a data-rich environment," *Journal of Business and Economic Statistics*, 33, 144–164.

MONTES-ROJAS, G. (2019): "Multivariate quantile impulse response functions," *Journal of Time Series Analysis*, 40, 739–752.

MUMTAZ, H., AND P. SURICO (2015): "The transmission mechanism in good and bad times," *International Economic Review*, 56, 1237–1260.

POLITIS, D. N., J. P. ROMANO, AND M. WOLF (1999): *Subsampling*. Springer.

PSARADAKIS, Z., AND M. SOLA (1998): "Finite-sample properties of the maximum likelihood estimator in autoregressive models with Markov switching," *Journal of Econometrics*, 86, 369–386.

SAKOV, A., AND P. J. BICKEL (2000): "An Edgeworth expansion for the $m$ out of $n$ bootstrapped median," *Statistics and Probability Letters*, 49, 217–223.

SHAMAN, P., AND R. A. STINE (1988): "The bias of autoregressive coefficient estimators," *Journal of the American Statistical Association*, 83, 842–848.

SONG, X., AND A. TAAMOUTI (2021): "Measuring Granger causality in quantiles," *Journal of Business and Economic Statistics*, 39, 837–852.

TROSTER, V. (2018): "Testing for Granger causality in quantiles," *Econometric Reviews*, 37, 850–866.

TSONG, C. C., AND C. F. LEE (2011): "Asymmetric inflation dynamics: Evidence from quantile regression analysis," *Journal of Macroeconomics*, 33, 668–680.

XIAO, Z. (2012): "Time series quantile regressions," in *Handbook of Statistics*, vol. 30.