NBS Working Paper
7/2024

A Growth-at-Risk Model in Slovakia

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)

A Growth-at-Risk Model in Slovakia*

Marián Vávra†

October 10, 2024

Abstract

Monitoring financial conditions can provide central banks with valuable information about risks to future GDP growth and other macroeconomic variables. In this paper, we follow the recent literature on growth-at-risk and use a linear quantile regression model to exploit the information content of the financial conditions index for tail-risk forecasting of output growth in Slovakia.

Keywords : Quantile regression model; financial conditions index; growth-at-risk.

JEL-Codes : C15, C22.

*The author is grateful to Patrik Kupkovič (NBS), Jean Guillaume Sahuc (BDF), Matteo Mogliani (BDF), and other participants of the ECB Expert Group on Macro-at-Risk (EGMAR) for interesting discussions, constructive comments, and challenging suggestions.

†National Bank of Slovakia. E-mail: marian.vavra@nbs.sk.

1. Introduction

In the aftermath of the global financial crisis in 2008, the focus of central banks on monitoring financial conditions and modeling macroeconomic tail risks has increased significantly, leading to the development of new tools (e.g. growth-at-risk) for macroprudential policy. The growth-at-risk concept offers, at least in theory, several attractive features to central bankers by (i) bringing information about the entire distribution of future output growth, which encompasses both downside and upside macroeconomic risks and goes beyond more traditional point forecasts; (ii) quantifying the impact of various risk factors on future output growth and its entire distribution, which helps to guide macroprudential policy measures; (iii) facilitating communication of the sources of macroeconomic risks to professionals and the public.

To account for possible non-linearities in the macro-financial linkages within a relatively simple modeling framework, a quantile regression method has become widely used for estimating and forecasting macroeconomic tail-risks in practice (see, e.g., Giglio, Kelly, and Pruitt (2016); De Nicolò and Lucchetta (2017); Adrian, Boyarchenko, and Giannone (2019); Figueres and Jarociński (2020), Adams, Adrian, Boyarchenko, and Giannone (2021); Kiley (2022); Ferrara, Mogliani, and Sahuc (2022) to name a few recent applications).[1] In this paper we follow the mainstream literature and employ linear quantile regression models to exploit the information content of the recently constructed financial conditions index for tail-risk forecasting of output growth in Slovakia. Attention is given to several issues related to the practical implementation of the procedure such as data adjustment for COVID-19 outliers and numerical optimization of the parameters of a skewed-t distribution.

The remainder of the paper is organized as follows. Section 2 provides a brief econometric background to linear quantile regression models, including asymptotic properties of the estimated parameters and basic diagnostic tests. The quantile regression methodology is then applied to Slovak data and the results are discussed in Section 3.

[1] See also Chahad and Mogliani (2024) for a comprehensive literature review.

Macroeconomic tail-risks and density forecasts for Slovak real GDP growth are reported and discussed in Section 4. And finally, Section 5 summarizes and concludes.

2. Econometric Framework

2.1. Quantile Regression

Given a real-valued stationary time sequence $\{(y_1, \boldsymbol{x}_1), \ldots, (y_n, \boldsymbol{x}_n)\}$, a linear (predictive) quantile regression model (QR) can be formally written as follows

$$y_{t+h} = \beta_0(u_{t+h}) + \beta_1(u_{t+h})x_{1,t} + \cdots + \beta_p(u_{t+h})x_{p,t} = \boldsymbol{x}_t'\boldsymbol{\beta}(u_{t+h}), \quad t = 1, \ldots, n, \tag{1}$$

where y_t denotes a dependent variable at time t (e.g. output growth), u_{t+h} is a standard uniform random variable, $\boldsymbol{\beta}(u_{t+h}) \in \mathbb{R}^{p+1}$ is a $(p+1 \times 1)$ vector of unknown parameters, $\boldsymbol{x}_t = (1, x_{1,t}, \ldots, x_{p,t})'$ is a $(p+1 \times 1)$ vector of explanatory variables (e.g. financial conditions index, lagged dependent variable, etc.), and an integer $h \in \{1, \ldots, H\}$ denotes a forecast horizon. The model in (1) can be interpreted as a specific type of a functional-coefficient model and nests an autoregressive conditional heteroskedasticity model in terms of the second-order properties (see Xiao (2012) for details). This makes the quantile regression a flexible yet simple method that is widely used in applied economics.

The corresponding conditional quantile function of the QR model in (1) is given by

$$Q_{y_{t+h}}(\tau|\boldsymbol{x}_t) = \boldsymbol{x}_t'\boldsymbol{\beta}(\tau), \tag{2}$$

where $\tau \in (0, 1)$ is the quantile level. The conditional quantile function in (2) is supposed to be a monotonically increasing function in the quantile parameter τ to ensure a non-crossing property of the conditional quantiles and thus the consistency of the estimated parameters (see Chernozhukov, Fernández-Val, and Galichon (2010)). For any quantile level $\tau \in (0, 1)$, the estimated quantile regression parameters are a solution to the following minimizing problem:

$$\tilde{\boldsymbol{\beta}}(\tau) = \underset{\boldsymbol{\beta}(\tau)\in\mathbb{R}^{p+1}}{\text{argmin}} \sum_{t=1}^{n} \rho_\tau(y_{t+h} - \boldsymbol{x}_t'\boldsymbol{\beta}(\tau)), \tag{3}$$

where $\rho_\tau(w) = w\,[\tau - \mathbb{I}(w < 0)]$ denotes the check (tick) function with $\mathbb{I}(\cdot)$ being a standard indicator function. It has long been recognized 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 is widely used in practice (see Chen and Wei (2005) for technical details). 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)).

2.2. Evaluation Metrics

Various evaluation metrics have been proposed in the literature to assess the accuracy of the estimated QR models. Our evaluation process includes two popular metrics (see, e.g., Giacomini and Komunjer (2005); Carriero, Clark, and Marcellino (2022)). The first metric is the average quantile score (QS) associated with the tick loss function and designed to evaluate the individual quantiles. The average quantile score for the quantile level τ and the forecast horizon h is computed as follows

$$QS_\tau(h) = \frac{1}{n-h} \sum_{t=1}^{n-h} \left(y_{t+h} - \tilde{Q}_{y_{t+h}}(\tau|\boldsymbol{x}_t)\right) \left[\tau - \mathbb{I}\left(y_{t+h} < \tilde{Q}_{y_{t+h}}(\tau|\boldsymbol{x}_t)\right)\right], \tag{4}$$

where $\tilde{Q}_{y_{t+h}}(\tau|\boldsymbol{x}_t) = \boldsymbol{x}_t'\tilde{\boldsymbol{\beta}}(\tau)$ denotes the fitted conditional quantile.

The second metric is the weighted continuous ranked probability score (CRPS) used to evaluate the entire predictive density. The metric is computed over $K = 17$ evenly spaced quantile levels as follows

$$CRPS(h) = \frac{2}{K}\sum_{j=1}^{K}\omega_j \, QS_{\tau_j}(h), \tag{5}$$

where $\tau_j \in \{0.10, 0.15, \ldots, 0.85, 0.90\}$ and $\omega_j = (1-\tau_j)^2$ is a left-tailed focused weight function that puts more weight on the left-tailed quantiles than on the right-tailed ones, something which is desirable when evaluating tail-risk models for output growth. The model with the smallest values of the diagnostic statistics is preferred.

2.3. Fitting Skewed Student Distribution

Based on the estimates of the conditional quantile function $\tilde{Q}_{y_{t+h}}(\tau|\boldsymbol{x}_t)$ over a fixed number of quantile levels, the predictive conditional density can be approximated using some flexible and smooth parametric distribution. For convenience, we use a skewed Student t distribution of Azzalini and Capitanio (2003) to compute the conditional predictive density of the dependent variable. This probability distribution allows for fat tails and asymmetry and boils down to a normal distribution in a limiting case. The unknown parameters of the skewed-t distribution are obtained by minimizing the squared distance between the estimated conditional quantiles and the conditional quantiles of the skew-t distribution, that is

$$\tilde{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}}{\text{argmin}} \sum_{j=1}^{K}\left(\tilde{Q}_{y_{t+h}}(\tau_j|\boldsymbol{x}_t) - F^{-1}_{y_{t+h}}(\tau_j|\boldsymbol{x}_t;\boldsymbol{\theta})\right)^2, \tag{6}$$

where $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4)'$ is a (4×1) vector of unknown parameters (i.e. θ_1: location, θ_2: scale, θ_3: shape, and θ_4: degrees of freedom) of the skewed-t quantile function $F^{-1}_{y_{t+h}}(\tau|\boldsymbol{x}_t;\boldsymbol{\theta})$, $\tilde{Q}_{y_{t+h}}(\tau|\boldsymbol{x}_t)$ is the estimated conditional quantile, and $\boldsymbol{\Theta}$ is a set of admissible values for the skewed-t parameters.[2] We follow arguments in Mitchell, Poon, and Zhu (2022) and employ a fine grid of $K = 17$ evenly spaced quantile points used to approximate the estimated conditional quantile function by the skewed-t quantile

[2] For example, the scale parameter θ_2 is restricted to be a positive real number $\theta_2 > 0$, the shape parameter θ_3 is restricted to be any real number within a set $[-1, 1]$, and the degrees of freedom parameter θ_4 is restricted to be an integer from a set $[5, 30]$.

function, that is $\tau_j \in \{0.10, 0.15, \ldots, 0.85, 0.90\}$. If possible, starting values required for numerical optimization of (6) are obtained directly from the estimated conditional quantile function using robust (quantile) measures of location, scale and shape (see Kim and White (2004)).[3]

Popular macroeconomic downside and upside risk measures (often called tail risk measures) can be easily calculated from the approximated skewed-t distribution. In particular, we focus on the 10% downside risk measure (i.e. $F^{-1}_{y_{t+h}}(0.1|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}})$) and the 90% upside risk measure (i.e. $F^{-1}_{y_{t+h}}(0.9|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}})$) for the dependent variable y_{t+h}. The downside risk is a particularly useful measure for output growth, whereas the upside risk is for inflation. Note that the difference between the upside risk and downside risk measures forms a prediction interval around the conditional point forecast with the probability coverage 80 $\%$. Collecting individual (marginal) prediction intervals over all h-step ahead forecasts provides a commonly used prediction band for a path (point) forecast.[4] To utilize as much information from the skewed-t distribution as possible, we also report the conditional probability of negative output growth defined as $\mathbb{P}(y_{t+h} < 0|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}}) = F_{y_{t+h}}(0|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}})$ and the conditional point (median) forecast of real GDP growth defined as $F^{-1}_{y_{t+h}}(0.5|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}})$.

3. Empirical Results

3.1. Data

In our analysis, the following two Slovak macroeconomic indicators are employed: (i) the year-on-year growth rate of real gross domestic product (RGDP), including the series adjusted for COVID-19 outliers (RGDP*)[5]; and (ii) the financial conditions index (FCI) proposed by Kupkovič and Šuster (2020). The Slovak FCI is constructed as a weighted average of ten quarterly macro-financial indicators. All economic indicators

[3] For example, the starting value for the location parameter θ_1 is set as $\tilde{Q}_{y_{t+h}}(0.5|\boldsymbol{x}_t)$.

[4] Note, however, that prediction bands constructed from individual marginal prediction intervals lack, in general, the desired marginal probability coverage. As a result, multi-step ahead probability statements should be made with caution (see Wolf and Wunderli (2015) for details).

[5] Details of the outlier adjustment are provided in the Appendix.

span the period from Q1 2003 to Q2 2023 (i.e. 82 observations). The series are depicted in Figure 1.

Figure 1: Macroeconomic and Financial Data

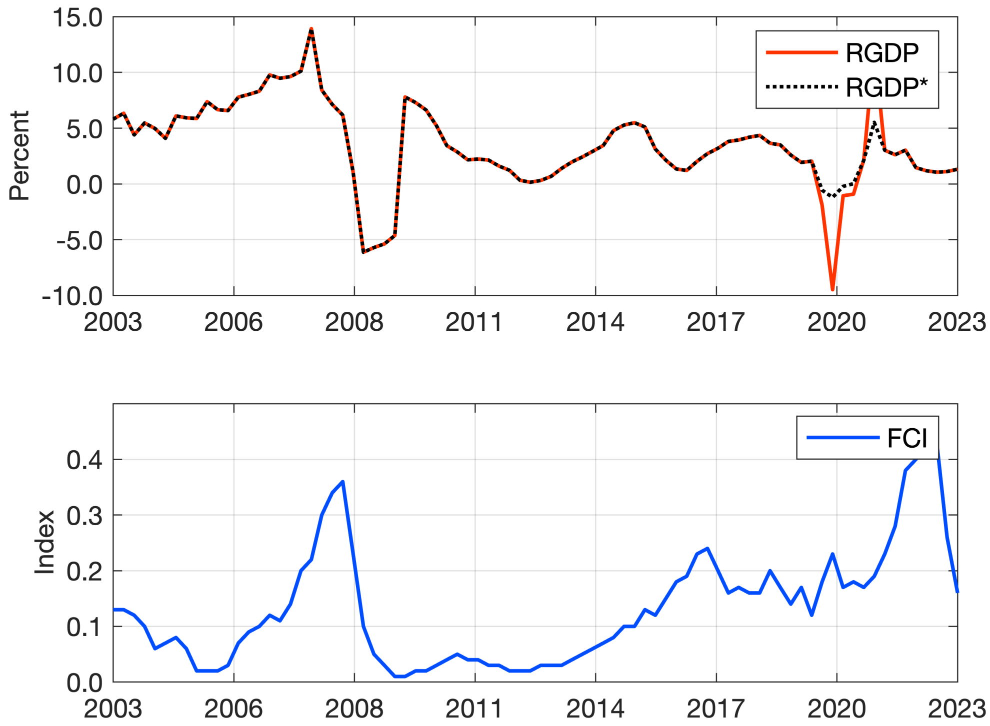

3.2. Estimation and Goodness of Fit

The benchmark QR model in (1) is estimated using the following specification of variables: (i) the outlier-adjusted real GDP is set as the dependent variable $y_{t+h} \equiv \mathrm{RGDP}^*_{t+h}$; (ii) $x_{1,t} \equiv \mathrm{RGDP}^*_t$ and $x_{2,t} \equiv \mathrm{FCI}_t$ form a set of explanatory variables. The estimated QR model parameters for the most relevant quantile levels $\tau \in \{0.1, 0.5, 0.9\}$ and the most policy-relevant forecast horizons $h \in \{1, 2, 3, 4\}$, including a pseudo goodness of fit statistic R^2, are reported in Table 1. The p-values are calculated using the moving block bootstrap with a data-driven block size (see Gregory, Lahiri, and Nordman (2018)). The blue-colored (red-colored) entries indicate the statistical significance of the estimated parameters at the nominal level 0.10 (0.05).

It can be concluded from the results that the GDP-related parameter β_1 quickly diminishes with the forecast horizon h and the effect is quite similar across all quantile points. However, the behaviour of the FCI-related parameter β_2 is a bit puzzling. As expected, the parameter increases (in absolute terms) with the forecast horizon but its values are "symmetrically" distributed around the center quantile $\tau = 0.5$. Put differently, financial conditions may be useful in explaining the scale (variance) of the predictive conditional distribution of output growth but less useful in explaining the shape (asymmetry) of the distribution. This fact is graphically demonstarted in Figure 2. The figure depicts the behaviour of real GDP growth (unadjusted for outliers) and three estimated conditional quantiles for two selected forecast horizons $h = 1$ and $h = 4$: median ($\tau = 0.5$), the lower-tail quantile ($\tau = 0.1$) and the upper-tail quantile ($\tau = 0.9$). In contrast to other studies (see, e.g., Adrian, Boyarchenko, and Giannone (2019); Figueres and Jarociński (2020)), no noticeable asymmetry in the confidence intervals is observed in the Great Recession period 2008 − 2010.

Table 1: Quantile Regression Estimates: Estimation Sample Q1 2003 − Q2 2023

	$\tau = 0.1$			$\tau = 0.5$			$\tau = 0.9$		
Horizon h	β_1	β_2	R^2	β_1	β_2	R^2	β_1	β_2	R^2
1	0.81	-7.76	0.55	0.93	-1.06	0.60	0.91	-4.28	0.45
2	0.70	-11.28	0.35	0.77	-3.33	0.34	0.37	-16.44	0.22
3	0.44	-25.12	0.25	0.65	-8.28	0.17	0.31	-17.80	0.20
4	0.12	-24.78	0.23	0.15	-9.79	0.05	0.17	-20.53	0.16

Notes: The blue-colored (red-colored) entry indicates the statistical significance of the estimated parameter at the nominal level 0.10 (0.05). The p-values are calculated using the moving block bootstrap with a data-driven block size (see Gregory, Lahiri, and Nordman (2018)).

Figure 2: Evolution of Estimated Conditional Quantiles: Horizons $h = 1$ and $h = 4$

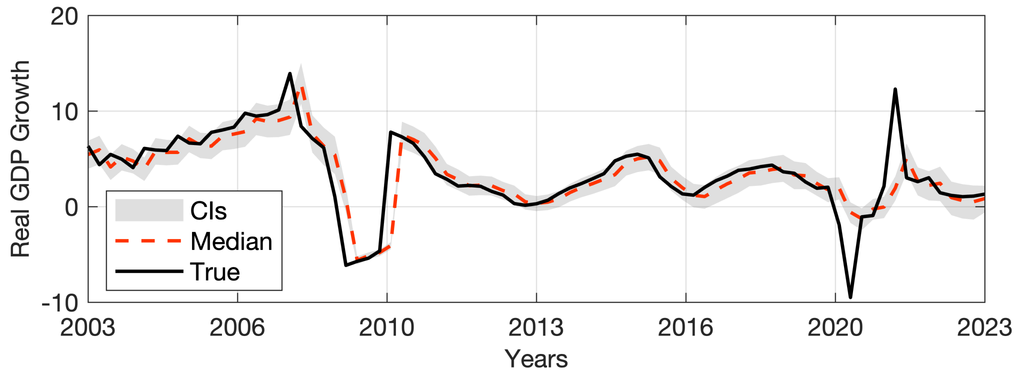

(a) $h = 1$

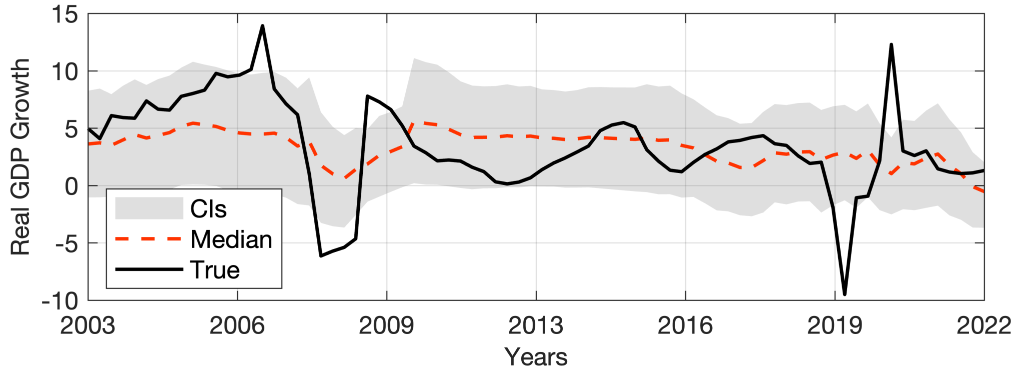

(b) $h = 4$

Notes: "True" denotes the observed real GDP growth (outlier unadjusted), "Median" denotes the median predictive quantile $\tilde{Q}_{y_{t+h}}(0.5|\boldsymbol{x}_t)$ for a given horizon h, and "CIs" represents the 80% credible intervals calculated as a difference between the upper and lower predictive quantiles $\tilde{Q}_{y_{t+h}}(0.9|\boldsymbol{x}_t) - \tilde{Q}_{y_{t+h}}(0.1|\boldsymbol{x}_t)$.

The conditional predictive quantile function $\tilde{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)$ over 17 quantile levels (red asterisks) and the fitted conditional skew-t quantile function $F^{-1}_{y_{t+h}}(\tau|\boldsymbol{x}_t;\tilde{\boldsymbol{\theta}})$ (blue line) estimated using the full sample of data (i.e. Q1 2003 – Q2 2023) are depicted in Figure 3. It is clear from the figure that the fitted skewed-t quantile function provides a very good approximation namely for the 1-step and 4-step ahead QR models, whereas relatively poor approximation for 2-step ahead and 3-step ahead QR models. These results are quite common for predictive regressions with the unbalanced predictive power of the explanatory variables over the forecast horizon (i.e. a small number of explanatory variables covering all forecast horizons).[6]

[6] A possible remedy is considering individual components of FCI in separate quantile regressions and combining the estimated predictive densities (see Busetti (2017) for details). We leave this topic for further research.

Figure 3: Conditional Predictive Quantiles and Fitted Skewed-t Quantile Function: Estimation Sample Q1 2003 – Q2 2023

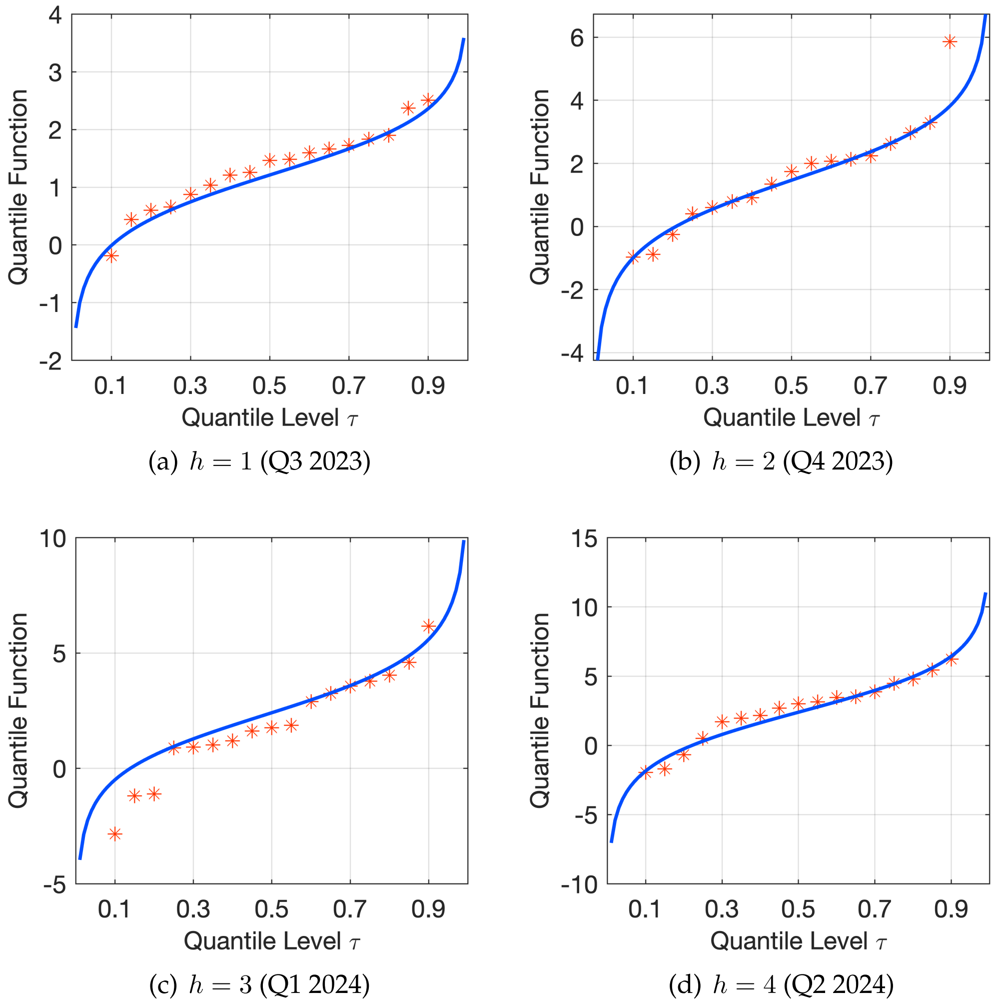

(a) $h = 1$ (Q3 2023)

(b) $h = 2$ (Q4 2023)

(c) $h = 3$ (Q1 2024)

(d) $h = 4$ (Q2 2024)

Notes: Red asterisks indicate the estimated predictive quantile points $\tilde{Q}_{y_{n+h}}(\tau|\boldsymbol{x}_n)$ whereas the blue line denotes the implied quantile function of the skewed-t distribution $F^{-1}_{y_{n+h}}(\tau|\boldsymbol{x}_n;\tilde{\boldsymbol{\theta}})$.

3.3. Comparison with Other Models

The in-sample fit of the estimated benchmark QR model (labeled "QR*" in what follows) is evaluated using the QS and $CRPS$ diagnostic statistics discussed in Section 2.[7] These results are compared with those from two additional models: (i) the QR

[7]We do not consider a standard out-of-sample forecast evaluation due to the limited number of observations available.

model where the outlier adjusted real GDP growth (RGDP*) is replaced with the original series and the set of explanatory variables is extended by two dummies capturing the beginning and the end of the COVID-19 period. This model is labeled "QRd" in what follows; (ii) a linear autoregressive model where the lag order is selected using the BIC. This model is estimated using outlier-adjusted data. This model is labeled "AR*" in what follows. All models are estimated over the full-sample spanning the period from Q1 2003 to Q2 2023. The resulting diagnostic statistics (calculated only for the comparable sample period) are shown in Figure 4 and Table 2.

Table 2: Continuous Ranked Probability Score ($CRPS$) Results: Evaluation Sample Q1 2003 – Q4 2019

Horizon h	QR*	QRd	AR*
1	0.272	0.271	0.290
2	0.436	0.435	0.477
3	0.559	0.552	0.611
4	0.620	0.615	0.682

Notes: The smaller the value of the $CRPS$ statistics the better the model in average.

Two interesting findings emerge from the results. First, the estimated QR models, although correcting data for the COVID-19 outliers in rather different ways, produce very similar in-sample fit measured either by the QS statistic (see Figure 4) or by the CRPS statistic (see Table 2). Second, both QR models clearly outperform the estimated AR model in terms of both diagnostic statistics. The differences are significant almost for all quantile points.

Figure 4: Quantile Scores (QS) Results: Evaluation Sample Q1 2003 − Q4 2019

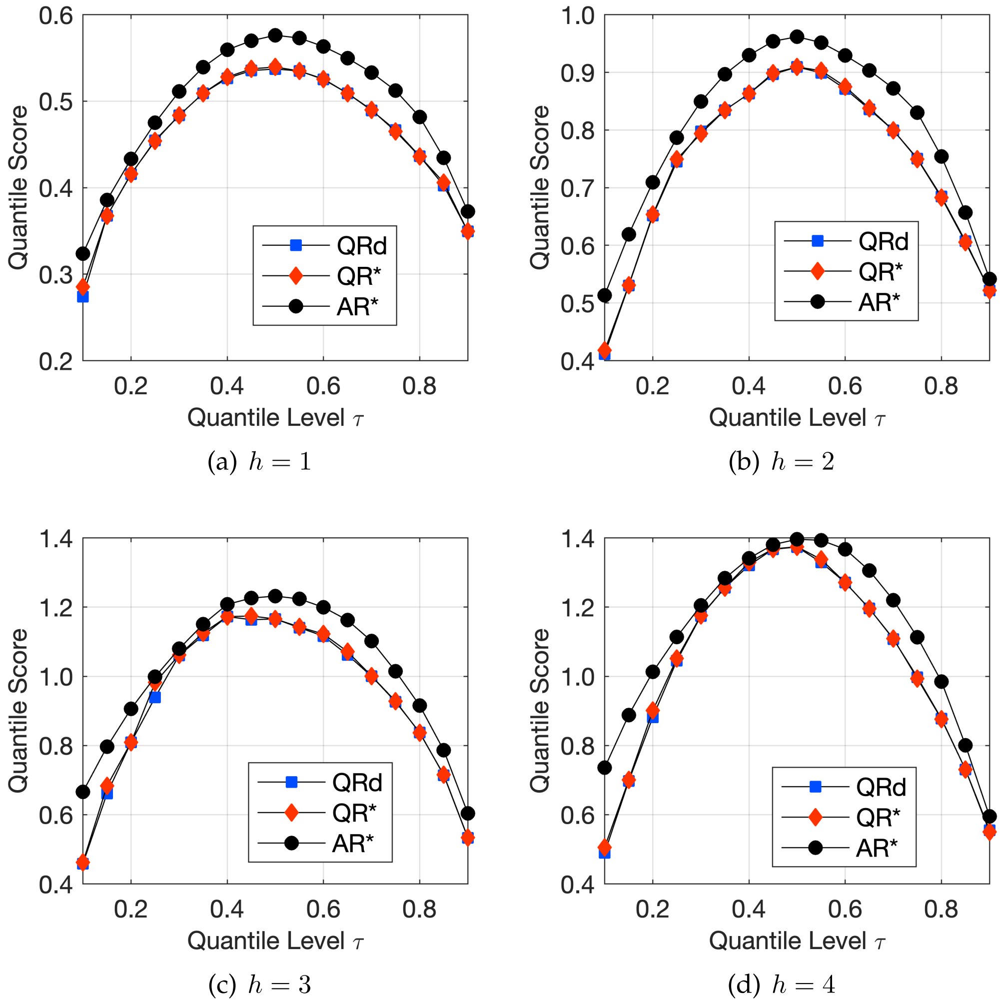

(a) $h = 1$ (b) $h = 2$

(c) $h = 3$ (d) $h = 4$

Notes: The smaller the value of the QS statistics the better the model.

4. Forecast Application

Having estimated parameters of our benchmark QR model and using the last available observations for real GDP growth and FCI (i.e. n = Q2 2023)[8], we can compute both the tail risk measures and the entire predictive density of real GDP growth for the next four quarters, that is for Q3 2023, Q4 2023, Q1 2024, and Q2 2024. The results are depicted in Figure 5. In each panel, the blue line depicts the predictive density of real GDP growth implied from the skewed-t distribution, the red diamond denotes the 10% downside risk of real GDP growth, the green square denotes the 90% upside risk of real GDP growth, the black dot denotes the median forecast of real GDP growth, and the shaded area represents the probability of negative real GDP growth.

Three interesting findings emerge from the figures below. First, no significant asymmetry in the predictive distributions of real GDP growth is observed over the forecast horizon. Put differently, possible economic risks associated with future economic growth in Slovakia are balanced around the point (median) forecast.

Second, the spread of the predictive distributions significantly increases over the forecast horizon. This fact can be easily documented by widening the marginal prediction intervals calculated as the difference between upside and downside risk measures (i.e. the difference between green squares and red diamonds). For example, we expect Slovak real GDP growth to lie between 0.0 percent and 2.4 percent with the probability of 80 percent in Q3 2023 but between -1.8 percent and 6.4 percent in Q2 2024 conditional on information (data) known up to Q2 2023.

Third, despite improving financial conditions (see Figure 1), the Slovak economy is stuck in a low-growth environment with output growth ranging between 1.0 and 2.5 percent.[9] Low economic growth is inevitably associated with a higher risk of economic

[8]It is important to remark here that the Slovak FCI is updated irregularly - Q2 2023 is that last available observation.

[9]It should be noted that the reported point (median) forecasts are model-based forecasts and do not

Figure 5: Predictive Densities and Tail Risk Measures of Slovak Real GDP Growth

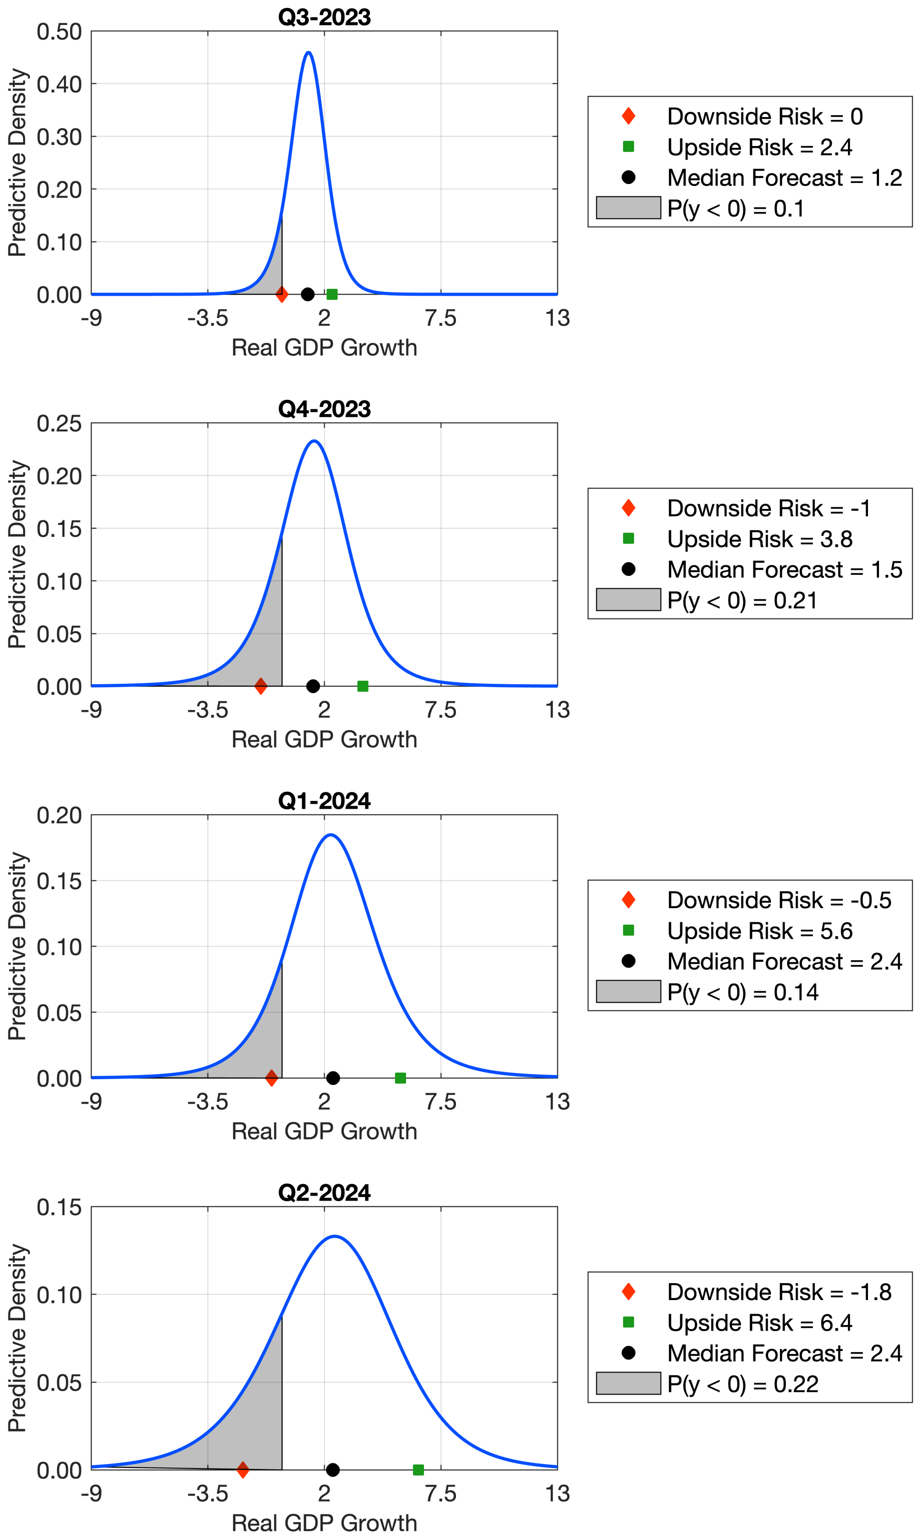

downturn. This fact is well documented by the reported probability of negative real GDP growth ranging from 10 to 22 percent over the forecast horizon.

The prediction intervals for a path forecast are depicted in Figure 6. The upside risk values, denoted by green squares, form the upper bound of the prediction interval whereas the downside risk values, denoted by the red diamonds, form the lower bound of the prediction interval. It is worth remarking here that, although the prediction intervals for individual forecasts have coverage of approximately 80 %, this is no longer true for the prediction intervals of path forecasts. As a result, all multistep-ahead probability statements need to be adjusted accordingly. In our case, the probability that the Slovak real GDP growth rate will lie within the prediction bands over the entire forecast horizon ranges from 41 % ($= 0.8^4$) to 80 %.

Figure 6: Real GDP Fan Chart

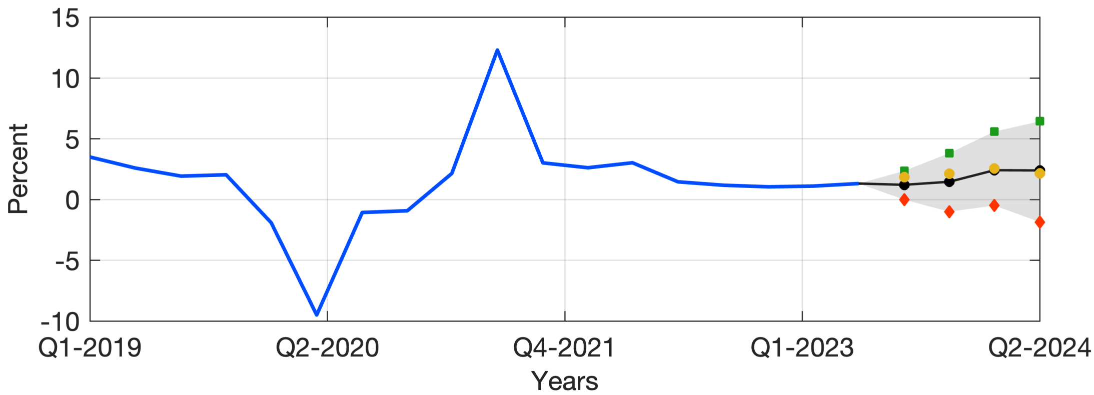

As mentioned earlier in the paper, a proper out-of-sample analysis is impossible to conduct due to a limited number observations. Nevertheless, we can at least compare the out-of-sample forecasts from our benchmark model with officially released real GDP data over the forecast horizon Q3 2023 – Q2 2024. The model point forecasts are denoted by black points and the official data by yellow points in Figure 6. It can be conclude that, keeping in mind the level of uncertainty in real GDP forecasting, the proposed QR model performs quite well.

necessarily correspond to the official forecasts published by the Bank.

5. Conclusion

In this paper, we show how to estimate a standard growth-at-risk model that can be used to assess the impact of changing financial conditions on future output growth in Slovakia. We show that the estimated linear dynamic quantile regression model captures the main stochastic features of real GDP growth sufficiently well and outperforms a simple linear autoregressive model.

There are two potential drawbacks to our econometric approach. The first one stems from modeling macroeconomic tail risks using quarterly (financial conditions) data, which are even published with a considerable delay (almost one quarter). We believe that policymakers could benefit from a more flexible and timely approach based on a mixed-frequency model, especially in times of economic downturns (see Ferrara, Mogliani, and Sahuc (2022)). The second potential drawback stems from using the aggregate financial conditions index with pre-specified weights for individual components. We believe that using individual components in separate quantile regressions and combining the quantile functions and/or predictive densities can further improve the results (see Busetti (2017)). We leave these topics for further research.

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.

Azzalini, A., and A. Capitanio (2003): "Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65, 367–389.

Busetti, F. (2017): "Quantile aggregation of density forecasts," *Oxford Bulletin of Economics and Statistics*, 79, 495–512.

Carriero, A., T. E. Clark, and M. G. Marcellino (2022): "Specification choices in quantile regression for empirical macroeconomics," *FRB of Cleveland Working Paper*, 25.

Chahad, M., and M. Mogliani (2024): "Macro-at-risk in the Euro Area," *ECB Working Paper Series*, (draft).

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

Chernis, T., P. J. Coe, and S. P. Vahey (2023): "Reassessing the dependence between economic growth and financial conditions since 1973," *Journal of Applied Econometrics*, 38, 260–267.

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

De Nicolò, G., and M. Lucchetta (2017): "Forecasting tail risks," *Journal of Applied Econometrics*, 32, 159–170.

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.

GIACOMINI, R., AND I. KOMUNJER (2005): "Evaluation and combination of conditional quantile forecasts," *Journal of Business and Economic Statistics*, 23, 416–431.

GIGLIO, S., B. KELLY, AND S. PRUITT (2016): "Systemic risk and the macroeconomy: An empirical evaluation," *Journal of Financial Economics*, 119, 457–471.

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.

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.

KUPKOVIČ, P., AND M. ŠUSTER (2020): "Identifying the financial cycle in Slovakia," *NBS Working Paper*, 2.

MITCHELL, J., A. POON, AND D. ZHU (2022): "Constructing density forecasts from quantile regressions: Multimodality in macro-financial dynamics," *FRB of Cleveland Working Paper Series*, 12.

WOLF, M., AND D. WUNDERLI (2015): "Bootstrap joint prediction regions," *Journal of Time Series Analysis*, 36, 352–376.

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

A. Outlier Detection

A potential shortcoming of dynamic time series QR models is their lack of robustness to outliers. Since the COVID-19 pandemic caused unprecedented variation in many macroeconomic variables (including real GDP growth), there is a consensus in the literature that these outlying observations should be adjusted. Chernis, Coe, and Vahey (2023) applied a Gaussian copula method to adjust data for COVID-19 outliers. The copula approach uses ranked data which helps mitigate outlying observations but preserves the dependence between macroeconomic variables. However, the copula approach also mitigates very likely economically relevant periods such as deep recessions. To avoid this undesirable feature of the Gaussian copula method, we use this non-parametric approach to adjust GDP data but only as a "local" approximation in a specific period.

Figure 7: Real GDP Adjustment for Outliers

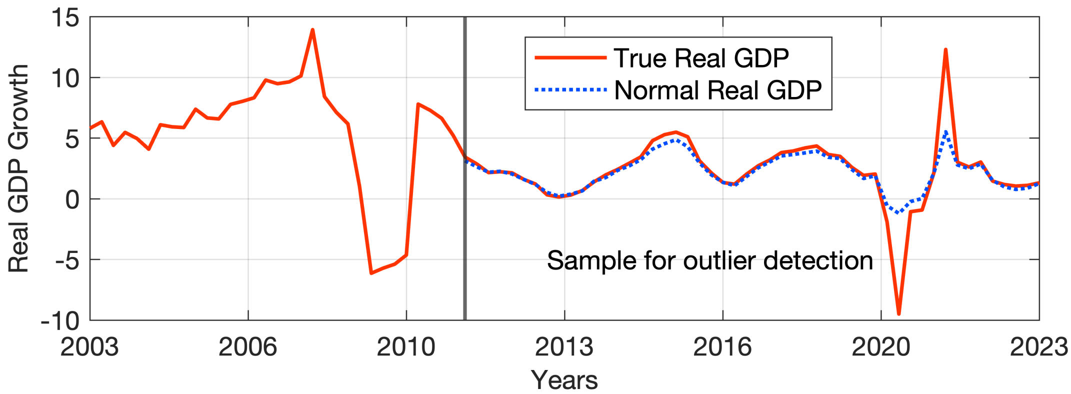

In particular, we restrict our attention to the subsample covering the period from Q1 2011 to Q2 2023 where the only outlying observations are due to the COVID-19 pandemic (see Figure 7). Denoting $R(y_t)$ as the rank of real GDP growth in period t, for $1 \leq t \leq T$, the corresponding normal scores are obtained by plugging these ranks (re-scaled to range between 0 and 1) in the quantile function of the standard normal distribution, that is

$$z_t = \Phi^{-1}\left(\frac{R(y_t)}{T+c}\right), \quad \text{for} \quad t = 1, \dots, T, \tag{7}$$

where $c = 1/2$ is Bloom's constant. Finally, the zero mean and unit variance normal scores $\{z_t\}$ are re-scaled again using a subsample median and a subsample median absolute deviation of real GDP growth to match moments of the original series. Both the official real GDP series contaminated by COVID-19 outliers (denoted "True Real GDP") and the adjusted one (denoted "Normal Real GDP") are depicted in Figure 7. It can be concluded from the figure that, apart from the specific COVID-19 period, the normal scores method approximates real GDP growth surprisingly well. The reconstructed real GDP series used in our analysis consists of the officially published real GDP data, except for the period Q1 2020 - Q2 2021 where the observations are replaced by the normal approximation.