Erasmus University Rotterdam

Erasmus School of Economics

Master Thesis Quantitative Finance

Predicting Growth-at-Risk

Supervisor:

C. Zhou

Author:

Jelmar Faber (447466)

Second assessor:

B. van Os

October 14, 2021

Abstract

Growth-at-Risk (GaR) is a measure that enables policy-makers to gain insights in the future behaviour of the economy. Therefore, predicting it adequately is essential for policy-makers. Quantile regressions are the foundation for the estimation, however due to overestimating the likelihood of future recessions the estimates are contaminated. The Bayesian GaR model imposes Bayesian inference on a quantile regression to take the left tail of the future distribution of GDP growth accurately into account during recessions. This paper implements the Markov switching model to determine whether GaR predictions can be improved. The results suggest that the Bayesian GaR model is generally the superior fit when it comes to predicting GaR. This is likely due to the fact that the Bayesian GaR model takes parameter uncertainty into account when predicting the distribution of future GDP growth.

Keywords: Quantile regression, Bayesian inference, Markov switching, left tail

Contents

1 Introduction

The macro economy and its developments are of great interest to society, as the negative repercussions are vital to take into account. In order to learn the behaviour of the economy in the present and in the future we focus on macroeconomic variables measuring economic growth. In this paper, we analyse Gross Domestic Product (GDP) growth and predict its distribution in the future.

When analysing GDP growth an interesting characteristic revolves around the risk of it dropping, implying decline in economic activity. Growth-at-Risk (GaR) is a quantile from the distribution of future GDP growth that quantifies macroeconomic risk in terms of economic growth. Hence, the GaR measure enables us to learn the magnitude of the risk involved in the economy. Estimating the GaR has become more popular over the recent years, as shown in Adrian et al. (2019). They estimate GaR using quantile regressions and find that GDP growth becomes left-skewed during recessions and symmetric during expansions by modelling the full distribution of future real GDP growth. Furthermore, the volatility that is displayed by future GDP growth is approximately entirely contained in the left tail of the conditional distribution of GDP growth. This implies that recessions tend to have more influence on the distribution of GDP growth compared to expansions. These findings substantiate the relevance of further research in GaR estimation.

According to Szabo (2020) quantile regressions provide inadequate GaR predictions. Specifically, quantile regressions model the quantiles separately, which leads to overestimating the likelihood of future recessions. This suggests that other techniques may need to be developed or utilised to estimate GaR more accurately. Overall, modelling the left tail of the distribution of GDP growth or the likelihood of future recessions is highly relevant when predicting Growth-at-Risk. In this paper we aim to learn how one can deal with left-skewed GDP growth during recessions in the best manner.

In order to prevent overestimating the likelihood of future recessions multiple measures can be taken. Szabo (2020) introduces a Bayesian GaR method, which models the quantiles simultaneously opposed to quantile regressions. This approach ensures that the quantiles do not cross each other, which is a common problem when implementing quantile regresssions. By imposing this restriction the Bayesian GaR assures that information contained in the lower quantiles does not instigate GaR of upper quantiles. Therefore, the Bayesian GaR is less likely to overestimate the likelihood of recessions and provide accurate GaR predictions. Alternatively, one can choose to model recessions and expansions more explicitly when predicting GaR. A Markov switching model can describe these periods distinctly as it transitions between different regimes, according to Hamilton (1989). We can assume that the model continuously switches between a 'good' state economy and a 'bad' state economy. However, this process remains unobserved, so an Expectation-Maximization algorithm is implemented to obtain parameter estimates, as shown in Hamilton (1990). In this setting recessions are approximated in the model through the underlying states. Therefore, a Markov switching model is less likely to overestimate the likelihood of a future recession. This paper analyses Bayesian GaR and Markov switching models and examines which of the two is best suited to provide accurate predictions for future Growth-at-Risk.

Since the work of Adrian et al. (2019) interest to predict GaR has been rising in the literature, as GaR is able to quantify macroeconomic risk in terms of GDP growth. Furthermore, the distribution of future GDP growth displays important characteristics of the macro economy and is therefore vital in order to consider when conducting policy-making. Moreover, GaR proves beneficial to quantify the impact of systemic risk on future GDP growth, as stated in Prasad et al. (2019). The relation between GDP growth and macroeconomic conditions is additionally emphasised in Adrian et al. (2019). It is generally straightforward to evaluate GaR, as it is a quantile equivalent to Value-at-Risk (VaR). Therefore, one can borrow numerous techniques that have been developed to model VaR according to Brownlees and Souza (2021).

The most straightforward way to predict GaR is by implementing quantile regressions. In the context of quantile regressions, parameter estimates are obtained by minimising the sum of quantile weighted absolute value of errors, as shown in Prasad et al. (2019). Additionally, quantile regressions link downside risk of future GDP growth to exogenous variables according to Brownlees and Souza (2021). To describe the behaviour of future GDP growth as adequately as possible multiple exogenous variables are included in the models. As GDP growth tends to exhibit distinct periods differentiating between shrinking and expanding periods of growth it might prove beneficial to introduce variables that distinguish between them. The National Financial Conditions Index (NFCI) is a macro-economic variable that explains a significant amount of the variation contained in GDP growth, as shown in Adrian et al. (2019). Other variables include financial Stress (Islami and Kurz-Kim (2014)) or economic policy uncertainty (Baker et al. (2016)). Implementing quantile regressions implies modelling the quantiles separately. However, as alluded to before cross-quantile effects are not taken into account according to Szabo (2020). This causes the left tail of the distribution of future GDP growth to be overestimated. Subsequently, this has a major impact on GaR predictions, which might turn out inaccurate.

Contrary to a quantile regression, the Bayesian GaR models the quantiles simultaneously. Implying that quantile crossing is avoided, as shown in Szabo (2020). This also ensures that the likelihood of future recessions is not overestimated causing GaR predictions to be more accurate. Szabo (2020) argues that adding informative linkage and controlling stability of predictions is vital to ensure adequate GaR predictions. This is done by introducing a hierarchical prior in the Bayesian setting that guarantees steady predictions. The addition of Bayesian priors to a quantile regression setting has better performance than regular quantile regressions in terms of forecasting. This raises the question whether other models can also be useful to implement when predicting GaR, specifically in the context of overestimating the left tail of the distribution of GDP growth.

A Markov switching model, as introduced by Hamilton (1989), obtains the density of GDP growth conditioned on lagged information and unobserved states. As mentioned before, these unobserved states might differentiate between recessions and expansions causing the model to adept well in the context of left-skewed GDP growth during recessions. Intuitively, the Markov switching model should be able to estimate GaR adequately due to this characteristic. In order to predict GaR using a Markov switching model, we first estimate the Markov switching model and then use simulation to obtain GaR. Markov switching models have been implemented in various settings in the literature, however they are yet to be applied to predict GaR. Therefore, it is interesting to examine whether the Markov switching model can predict GaR more accurately than the Bayesian GaR model.

Compared to existing literature this paper estimates GaR using a Markov switching model and the Bayesian GaR model to take the left tail of the distribution of future GDP growth more accurately into account. In such a way, the paper offers a new perspective on GaR estimation by comparing two methods on the accuracy of their respective GaR predictions.

The results from this paper indicate that in absence of the Covid period the Markov switching model is relatively unable to predict GaR adequately. When the Covid period is taken into account the conclusions are not straightforward, however the Bayesian GaR model is favoured due to the results from the individual backtests. Overall, it turns out that the Markov switching model is inept in differentiating between the unobserved states in a satisfactory manner, causing inadequate GaR predictions. On the other hand, the Bayesian GaR model imposes Bayesian inference on a quantile regression setting. This results in a posterior distribution of GaR, which reflects uncertainty revolving around future GDP growth leading to accurate GaR predictions.

The remainder of the paper is structured as follows. In Section 2, we elaborate on the models and techniques implemented in this paper. In Section 3, we describe the data set and explanatory variables that we use. In Section 4, we discuss and substantiate the results. Finally, in Section 5 we conclude and discuss our findings.

2 Methodology

In this section, we elaborate on the methods implemented in this paper. Before thoroughly examining the Bayesian GaR model and Markov switching model, we first introduce some notation to model the distribution of future GDP growth.

2.1 Quantile regression

The variable of interest y_{t+h} is the GDP growth rate at time $t+h$. The data is assumed to range from $t = 1, \ldots, T$. However, as we need to evaluate GaR predictions we split the data into a test and train set. The train set ranges from $t = 1, \ldots, T-n$ with n the amount of out-of-sample test observations. For ease of notation, we note $T_0 = T - n$. We define $\boldsymbol{x}_t = (1, y_t, \ldots, y_{t-p}, z_{1t}, \ldots, z_{dt})$, where z_{jt} are exogenous variables with $j = 1, \ldots, d$. The number of lagged GDP growth variables is indicated by $p+1$ and 1 implies an intercept in the model. The variable $\boldsymbol{x}_t$ is a vector of dimension $(1 \times l)$. Adrian et al. (2019) assume that the distribution of future GDP growth $\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)$, can then be modeled as follows:

$$\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) = \boldsymbol{x}_t\boldsymbol{\beta}(\tau), \tag{1}$$

where $\boldsymbol{\beta}(\tau)$ is a vector of quantile regression coefficients with length l for the τ^{th} quantile.

2.2 Bayesian GaR

In order to model all quantiles simultaneously, we group the aformentioned vectors into matrices. Let $\boldsymbol{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{T_0})'$, $\boldsymbol{Y} = (y_1, \ldots, y_{T_0})$ and $\boldsymbol{B}_m = (\boldsymbol{\beta}(\tau_1), \ldots, \boldsymbol{\beta}(\tau_m))$ with $\tau_1, \ldots, \tau_m$ the m quantiles in the model, such that $\tau_1 < \cdots < \tau_m$. By introducing the prior distribution $p_m(\boldsymbol{B}_m|\boldsymbol{X})$ on the quantile regression parameters, a posterior distribution can be obtained according to Feng et al. (2015). The posterior distribution is given as follows:

$$p(\boldsymbol{B}_m|\boldsymbol{X}, \boldsymbol{Y}) \propto p_m(\boldsymbol{B}_m|\boldsymbol{X})L(\boldsymbol{Y}|\boldsymbol{X}, \boldsymbol{B}_m), \tag{2}$$

where $L(\boldsymbol{Y}|\boldsymbol{X}, \boldsymbol{B}_m)$ represents the likelihood of GDP growth given the parameters and explanatory variables. Feng et al. (2015) show that the likelihood can be obtained nonparameterically, meaning the estimation is semi-parametric. This implies that we require an approximated likelihood, which is obtained through linear interpolation.

2.2.1 Linear interpolation

The likelihood in the Bayesian representation is undefined, so we seek to approximate it through linear interpolation. Feng et al. (2015) argue that the probability density function (pdf) of a random variable can be approximated by a set of predefined quantile functions. By implementing this concept, the approximated linearly interpolated density is obtained by means of equation (3).

$$\begin{aligned}\hat{f}_t(y_{t+h}|\boldsymbol{x}_t) = &\left[\sum_{j=1}^{m-1} \mathbb{1}_{[y_{t+h}\in(\boldsymbol{x}_t\boldsymbol{\beta}(\tau_j),\ \boldsymbol{x}_t\boldsymbol{\beta}(\tau_{j+1}))]} \frac{\tau_{j+1}-\tau_j}{\boldsymbol{x}_t\boldsymbol{\beta}(\tau_{j+1})-\boldsymbol{x}_t\boldsymbol{\beta}(\tau_j)}\right] \\ &+ \mathbb{1}_{[y_{t+h}\in(-\infty,\boldsymbol{x}_t\boldsymbol{\beta}(\tau_1))]}\tau_1 g_1(y_{t+h}|\boldsymbol{x}_t) \\ &+ \mathbb{1}_{[y_{t+h}\in(\boldsymbol{x}_t\boldsymbol{\beta}(\tau_m),\infty)]}(1-\tau_m) g_2(y_{t+h}|\boldsymbol{x}_t),\end{aligned} \tag{3}$$

where g_1 is the pdf of $\mathcal{N}(\boldsymbol{x}_t\boldsymbol{\beta}(\tau_1),\sigma^2)$, g_2 is the pdf of $\mathcal{N}(\boldsymbol{x}_t\boldsymbol{\beta}(\tau_m),\sigma^2)$ and σ^2 a pre-specified parameter. By multiplying all individual densities for the respective observations, we can approximate the entire likelihood as follows:

$$L(\boldsymbol{Y}|\boldsymbol{X},\boldsymbol{B}_m) \approx \prod_{t=1}^{T_0} \hat{f}_t(y_{t+h}|\boldsymbol{x}_t)$$

2.2.2 Metropolis-Hastings algorithm

The essential information needed to obtain samples from the posterior distribution in equation (2) is known. In order to obtain samples that represent the posterior distribution, we use an alteration of the Metropolis-Hastings algorithm as shown in Szabo (2020). The parameters are simulated k times and the algorithm we implement is given in Algorithm 1.

Algorithm 1: Metropolis-Hastings

Step 1: Set $k = 0$ and initialise $\boldsymbol{B}_m^{(0)}$ by performing a quantile regression

Step 2: Approximate the likelihood $L^{(0)} = \prod_{t=1}^{T_0} \hat{f}_t^{(0)}(y_{t+h}|\boldsymbol{x}_t, \boldsymbol{B}_m^{(0)})$, or $L^{(k-1)}$ in the k^{th} simulation

Step 3: Pick a random τ_j from $\tau_1, \ldots, \tau_m$ when we are at the k^{th} iteration. Then, randomly pick $\beta_i(\tau_j)$, the i^{th} element of $\boldsymbol{\beta}(\tau_j)$ to update. In order to obtain the replacement $\beta_i^*(\tau_j)$ we compute lower and upper bounds and use these to generate $\beta_i^*(\tau_j)$ from a uniform distribution, similarly to Feng et al. (2015). This approach ensures that quantile crossing does not occur. We replace $\beta_i^{(k-1)}(\tau_j)$, the $(i,j)^{th}$ element of $\boldsymbol{B}_m^{(k-1)}$, by $\beta_i^*(\tau_j)$ to obtain $\boldsymbol{B}_m^*$.

Step 4: Compute $L^* = \prod_{t=1}^{T_0} \hat{f}_t^*(y_{t+h}|\boldsymbol{x}_t, \boldsymbol{B}_m^*)$.

Step 5: Calculate the acceptance probability

$$\alpha = min\left(1, \frac{p_m(\boldsymbol{B}_m^*|\boldsymbol{X})L^*}{p_m(\boldsymbol{B}_m^{(k-1)}|\boldsymbol{X})L^{(k-1)}}\right),$$

Let $\boldsymbol{B}_m^{(k)} = \boldsymbol{B}_m^*$ with probability α and $\boldsymbol{B}_m^{(k)} = \boldsymbol{B}_m^{(k-1)}$ with probability $1 - \alpha$.

Step 6: Repeat steps 2 to 5 until the required number of simulations is reached.

By implementing Algorithm 1 we obtain parameter samples that represent the posterior distribution of the quantile regression coefficients. We take a specific τ and select the parameter samples corresponding to this quantile. We implement these samples in equation (1) to obtain a posterior distribution of GaR corresponding to the τ^{th} quantile. In order to acquire a point estimate in this Bayesian setting we minimise an expected quadratic loss function. Solving the minimisation implies taking the mean of the posterior distribution of GaR to obtain a point estimate.

2.3 Markov switching model

The other technique one may implement to model GDP growth is the Markov switching model, as introduced by Hamilton (1989). The Markov switching model assumes a latent state that drives the dependent variable, which is GDP growth in this paper. We start by introducing some notation, in order to obtain the GaR using a Markov switching model. We assume that GDP growth can either be in one of two regimes, because the economy is generally deemed to be in an expansion or recession of the business cycle. We assume $S_t = i$ with $i = 1, 2$, which represents the underlying state at time t as done in Hamilton (1989). The model can then be formulated as follows:

$$y_t = \begin{cases} \boldsymbol{x}_t \boldsymbol{\phi}_1 + \epsilon_t & \text{if } S_t = 1 \\ \boldsymbol{x}_t \boldsymbol{\phi}_2 + \epsilon_t & \text{if } S_t = 2 \end{cases} \tag{4}$$

where $\boldsymbol{\phi}_i$ represent the coefficient vectors for the distinct states with $i = 1, 2$ and ϵ_t is the error term at time t with variance σ_i^2 for $i = 1, 2$. We assume that the error term is normally distributed, implying that the dependent variable has a Gaussian density.

2.3.1 Inference on the states

Hamilton (1989) argues that the current regime S_t follows a first-order Markov chain. The transition probabilities are then given as follows:

$$\begin{aligned} P(S_t = 1 | S_{t-1} = 1) &= p_{11} \\ P(S_t = 2 | S_{t-1} = 1) &= 1 - p_{11} \\ P(S_t = 2 | S_{t-1} = 2) &= p_{22} \\ P(S_t = 1 | S_{t-1} = 2) &= 1 - p_{22} \end{aligned}$$

Due to the latent nature of the variable S_t it is difficult to estimate the model. Hamilton (1989) introduces the Hamilton filter, which uses a prediction and updating step to estimate the parameters. We define the state variable $\boldsymbol{\xi}_t$ as a vector that contains the probabilities of being in a specific state at time t. However, as the states are unobserved we need to estimate the state variable. We refer to the estimated state variable $\hat{\boldsymbol{\xi}}_{t|t} = \begin{pmatrix} P(S_t=1|\mathcal{I}_t) \\ P(S_t=2|\mathcal{I}_t) \end{pmatrix}$ as the filtered probability and it consists of probabilities of the model being in a specific state at various points in time. The prediction step is defined as follows:

$$\hat{\boldsymbol{\xi}}_{t+1|t} = \boldsymbol{P} \hat{\boldsymbol{\xi}}_{t|t},$$

with transition probability matrix:

$$\boldsymbol{P} = \begin{pmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{pmatrix}$$

Intuitively, the prediction step is relatively easy to interpret. Given all information at time t, we implement our estimate at time t and multiply it with the transition probabilities to obtain the estimated state variable at time $t + 1$. Furthermore, the updating step is given as follows:

$$\hat{\boldsymbol{\xi}}_{t|t} = \frac{\begin{pmatrix} f(y_t|S_t=1) \\ f(y_t|S_t=2) \end{pmatrix} \odot \hat{\boldsymbol{\xi}}_{t|t-1}}{(1\ 1) \left[\begin{pmatrix} f(y_t|S_t=1) \\ f(y_t|S_t=2) \end{pmatrix} \odot \hat{\boldsymbol{\xi}}_{t|t-1} \right]},$$

where $\odot$ indicates element wise multiplication. By obtaining the log likelihood and rewriting the density, one can use maximum likelihood estimation to acquire the parameter estimates. However, the structure of the likelihood is complicated. In order to obtain estimates we resort to an Expectation-Maximization (EM) algorithm, as shown in Dempster et al. (1977).

2.3.2 Expectation-Maximization algorithm

We implement the EM algorithm due to the fact that we are dealing with a latent variable S_t. However, we can rewrite the likelihood of the system in such a way that we only require the joint likelihood of the states and y_t, according to Hamilton (1990). The rewritten expression for the likelihood is given as follows:

$$L(\boldsymbol{y}_{1:T_0}|\boldsymbol{\theta}) = \log f(\boldsymbol{y}_{1:T_0}, \boldsymbol{\xi}_{1:T_0}|\boldsymbol{\theta}) - \log f(\boldsymbol{\xi}_{1:T_0}|\boldsymbol{y}_{1:T_0};\boldsymbol{\theta})$$

where $\boldsymbol{y}_{1:T_0} = (y_1, \ldots, y_{T_0})$, $\boldsymbol{\xi}_{1:T_0} = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_{T_0})$ and $\boldsymbol{\theta} = \{p_{11}, p_{22}, \boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \sigma_1^2, \sigma_2^2\}$ contains the coefficient parameters and probability parameters that are used to maximise the likelihood. Furthermore, we define the expectation operator $\tilde{E}(.) = E_{\boldsymbol{\xi}_{1:T_0}|\boldsymbol{y}_{1:T_0}}(.)$ to show that the conditional density is of no interest in the maximisation. The location of $\boldsymbol{\theta}$ for the maximum value of the likelihood solely depends on $\log f(\boldsymbol{y}_{1:T_0}, \boldsymbol{\xi}_{1:T_0}|\boldsymbol{\theta})$, as $\tilde{E}[\frac{d}{d\boldsymbol{\theta}}\log f(\boldsymbol{\xi}_{1:T_0}|\boldsymbol{y}_{1:T_0};\boldsymbol{\theta})] = 0$. Therefore, we are only interested in the joint density of $\boldsymbol{y}_{1:T_0}, \boldsymbol{\xi}_{1:T_0}$. We define $f_i(y_t) = f(y_t|S_t = i; \boldsymbol{\theta}_i)$ as the density function in state i and $\boldsymbol{\theta}_i$ the parameter vector which is a subset of $\boldsymbol{\theta}$ containing $\boldsymbol{\phi}_i$ and σ_i^2 when $S_t = i$ for $i = 1, 2$. This enablus us to write the joint density as follows:

$$\begin{aligned}
f(y_t, s_t|s_{t-1};\boldsymbol{\theta}) = &\begin{cases} f_1(y_t)p_{11} & \text{if } s_t = 1 \text{ and } s_{t-1} = 1 \\ f_1(y_t)(1-p_{22}) & \text{if } s_t = 1 \text{ and } s_{t-1} = 2 \\ f_2(y_t)(1-p_{11}) & \text{if } s_t = 2 \text{ and } s_{t-1} = 1 \\ f_2(y_t)p_{22} & \text{if } s_t = 2 \text{ and } s_{t-1} = 2 \end{cases} \\
=&[f_1(y_t)p_{11}]^{\mathcal{I}(S_t=1)\mathcal{I}(S_{t-1}=1)} \times [f_1(y_t)(1-p_{22})]^{\mathcal{I}(S_t=1)\mathcal{I}(S_{t-1}=2)} \\
&\times [f_2(y_t)(1-p_{11})]^{\mathcal{I}(S_t=2)\mathcal{I}(S_{t-1}=1)} \times [f_2(y_t)p_{22}]^{\mathcal{I}(S_t=2)\mathcal{I}(S_{t-1}=2)}
\end{aligned} \tag{5}$$

We take the logarithm of the expression and take the expectation operator $\tilde{E}$ when initialising the algorithm. However, this results in an unknown expression according to Hamilton (1990). Fortunately, this expression closely resembles the smoothed probabilities as introduced by Kim (1994). The smoothed estimate $\hat{\boldsymbol{\xi}}_{t|T_0} = E(\boldsymbol{\xi}_t|\mathcal{I}_{T_0})$ is computed using the tower property and conditioning on the state at time $t+1$. The smoothed estimate differs compared to the filtered probability with respect to the conditioned information. Specifically, future information is incorporated regarding the smoothed estimate. The smoothed estimates can be obtained as follows:

$$\begin{aligned}\hat{\boldsymbol{\xi}}_{t|T_0} &= \hat{\boldsymbol{\xi}}_{t|t} \odot \boldsymbol{P}'(\hat{\boldsymbol{\xi}}_{t+1|T_0} \oslash \hat{\boldsymbol{\xi}}_{t+1|t}),\\ \boldsymbol{P}^*(t) &= \boldsymbol{P} \odot (\hat{\boldsymbol{\xi}}_{t|T_0}\hat{\boldsymbol{\xi}}'_{t-1|t-1}) \oslash (\hat{\boldsymbol{\xi}}_{t|t-1}[1\ 1]),\end{aligned}$$

where $\oslash$ indicates element wise division. The required variables are defined and the Expectation-step (E-step) of the EM algorithm can be implemented. In the E-step, we run the Hamilton filter forward and Kim smoother backward to obtain $\boldsymbol{P}^*(t)$ using $\hat{\boldsymbol{\xi}}_{t|T_0}$. This results in the smoothed transition probabilities given as follows:

$$p^*_{ij}(t) = P(S_t = i, S_{t-1} = j|\hat{\boldsymbol{\theta}}, \mathcal{I}_{T_0}) \quad \text{for } i, j = 1, 2,$$

where $\hat{\boldsymbol{\theta}}$ indicates the estimated parameters contained in $\boldsymbol{\theta}$. Note, in the first iteration of the algorithm the parameters $\boldsymbol{\theta}_0$ need to be initialised in order to obtain the smoothed transition probabilities. Furthermore, we define the smoothed probabilities of being in state 1 or 2 respectively as follows:

$$\begin{aligned}p^*_1(t) &= P(s_t = 1|\mathcal{I}_{T_0}) = p^*_{11}(t) + p^*_{12}(t)\\ p^*_2(t) &= P(s_t = 2|\mathcal{I}_{T_0}) = p^*_{21}(t) + p^*_{22}(t)\\ p^*_1(t-1) &= p^*_{11}(t) + p^*_{21}(t)\\ p^*_2(t-1) &= p^*_{22}(t) + p^*_{12}(t)\end{aligned}$$

After performing the E-step and obtaining the smoothed transition probabilities we continue with the Maximisation-step (M-step) of the algorithm. In order to perform the M-step, we examine the joint likelihood in more detail, which is given as follows:

$$\begin{aligned}\tilde{E}[\log f(\boldsymbol{y}_{1:T_0}, \boldsymbol{\xi}_{1:T_0}|\boldsymbol{\theta})|\boldsymbol{\theta}_0] = {} & p^*_1(0)\log[\rho_1] + p^*_2(0)\log[\rho_2]\\ & + \sum_{t=1}^{T_0}(p^*_{11}(t)\log[p_{11}f_1(y_t)] + p^*_{12}(t)\log[(1-p_{22})f_1(y_t)]\\ & + p^*_{21}(t)\log[(1-p_{11})f_2(y_t)] + p^*_{22}(t)\log[p_{22}f_2(y_t)])\end{aligned} \tag{6}$$

where the ρ_i represents the unknown distribution of $\boldsymbol{\xi}_0$, specifically $P(S_0 = i) = \rho_i$ for $i = 1, 2$. In the M-step we maximise equation (6) with respect to the parameters of interest contained in $\boldsymbol{\theta}$. This results in the following analytical expressions for the transition probabilities, according to Hamilton (1990):

$$\hat{p}_{11} = \frac{\sum_{t=1}^{T_0} p_{11}^*(t)}{\sum_{t=1}^{T_0} p_1^*(t-1)}$$
$$\hat{p}_{22} = \frac{\sum_{t=1}^{T_0} p_{22}^*(t)}{\sum_{t=1}^{T_0} p_2^*(t-1)}$$

The expressions for the remaining parameters contained in $\hat{\boldsymbol{\theta}}$ need to be derived analytically from equation (6) as well. Hamilton (1990) argues that these expressions can be obtained by solving the following equation:

$$\sum_{t=1}^{T_0} p_i^*(t) \frac{d\log[f_i(y_t)]}{d\boldsymbol{\theta}_i} = 0 \quad \text{for } i = 1, 2$$

Note, equation (6) solely holds for the first iteration of the EM algorithm. For higher order iterations we replace the initialised $\boldsymbol{\theta}_0$ by $\hat{\boldsymbol{\theta}}$. In every iteration we update the parameter estimates in $\hat{\boldsymbol{\theta}}$ and perform the EM algorithm again and keep iterating until convergence. Goodwin (1993) argues that the estimates obtained from the EM algorithm can be implemented in the following manner to forecast GDP growth:

$$\hat{y}_{t+1|t} = \hat{P}(S_{t+1} = 1|\mathcal{I}_t)(\boldsymbol{x}_{t+1}\hat{\boldsymbol{\phi}}_1) + \hat{P}(S_{t+1} = 2|\mathcal{I}_t)(\boldsymbol{x}_{t+1}\hat{\boldsymbol{\phi}}_2), \tag{7}$$

where $\hat{P}(S_{t+1} = i|\mathcal{I}_t)$ is computed using the filtered probability $\hat{\boldsymbol{\xi}}_{t|t}$ and estimated transition probabilities for $i = 1, 2$. In order to obtain the distribution of future GDP growth, we simulate the error term ϵ_t using $\hat{\sigma}_i^2$ and $\hat{P}(S_{t+1} = i|\mathcal{I}_t)$ to obtain the simulated error term in a similar manner as in equation (7). Specifically, we compute a weighted sum of the error terms for the two regimes. We repeat this procedure L times to retrieve the distribution of future GDP growth. Following this, we sort the simulated values from low to high and take the τ^{th} quantile of the distribution to obtain the GaR.

2.4 Evaluation criteria

The Bayesian GaR and the Markov switching model enable us to derive the conditional distribution of future GDP growth. In order to determine which technique is most proficient predicting an accurate GaR, multiple evaluation criteria are required. The models are validated based on their out-of-sample performance. This is achieved by splitting the data into a train and test set. We use the observations $t = 1, \ldots, T - n$ to train the models and validate them with the remaining $t = T - n, \ldots, T$. So, we use n out-of-sample observations to evaluate the models.

As the GaR is a risk measure similar to the Value-at-Risk (VaR) we can use backtesting techniques that exist for the VaR in the context of this paper. Particularly, we examine the number of violations that correspond with the GaR using $\mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h}))$. The violations of a VaR measure are i.i.d. Bernoulli distributed random variables, according to Christoffersen and Pelletier (2004). We can aggregate them to obtain the total number of violations $Z = \sum_{t=T-n}^{T} \mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h}))$. Then, Z follows a Binomial distribution $\mathcal{B}(n, 1-\tau)$. When $n \to \infty$ the following linear transformation converges to a standard normal distribution:

$$\frac{Z - n(1-\tau)}{\sqrt{n\tau(1-\tau)}} \sim \mathcal{N}(0,1)$$

If Z is too high we reject the null hypothesis, implying that the GaR is not accurate enough and therefore the respective model is unable to provide accurate GaR estimates. In the remainder of this paper we refer to this test as the traditional backtest. In order to arrive at a robust conclusion regarding the adequacy of the models, another backtesting technique is implemented. In this setting, GaR is tested based on a scoring function to take the number of violations into account, according to Nolde and Ziegel (2017). Under the null hypothesis, the test statistic of the conditional backtest is given as follows:

$$H = \frac{1}{\sqrt{n}} \hat{\Sigma}_n^{-1/2} \sum_{t=T-n}^{T} [1 - \tau - \mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h})] \sim \mathcal{N}(0,1), \tag{8}$$

where $\hat{\Sigma}_n^{-1/2}$ is a heteroscedastic and autocorrelation consistent estimator of the asymptotic covariance matrix $\Sigma_n = cov(1/\sqrt{n} \sum_{t=T-n}^{T} [1-\tau-\mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h})])$. The metric H is asymptotically standard normal under the null hypothesis. We reject the null hypothesis if $|H|$ is too high indicating that the GaR predictions are significantly inadequate, following Nolde and Ziegel (2017).

The previously mentioned evaluation criteria offer an insight into the predictive performance for GaR of the models individually. However, to allow for a more direct comparison, we include a comparative backtest in the analysis. The null hypothesis of the comparative backtest states that the two models predict at least as well as one another, following Nolde and Ziegel (2017). To evaluate the predictions we use the following scoring function:

$$S_b(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)) = (1 - \tau - \mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h}))y_{t+h} + \mathcal{I}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t) > y_{t+h})\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)$$

where b represents the respective GaR of the Bayesian GaR (BG) or Markov switching model (MS). We define $\Delta_n \bar{S} = \frac{1}{n} \sum_{t=T-n}^{T} (S_{BG}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)) - S_{MS}(\mathbb{Q}_\tau(y_{t+h}|\boldsymbol{x}_t)))$, then the test statistic is formulated as follows:

$$W = \frac{\Delta_n \bar{S}}{\hat{\sigma}_n / \sqrt{n}} \sim \mathcal{N}(0,1), \tag{9}$$

where $\hat{\sigma}_n$ is the HAC estimator of the asymptotic variance $\sigma_n^2 = var(\sqrt{n}\Delta_n \bar{S})$. We reject the null hypothesis if $|W|$ is too high.

2.5 Implementation

After elaborating on the techniques utilised in this paper, we discuss their implementation. As alluded to before, the results are based on out-of-sample estimation. We implement a moving window and first estimate the models using $t = 1, \ldots, T - n$ observations. In the next step we expand the window to $t = 1, \ldots, T - n + 1$ and keep iterating until we reach the end of the data set. For every iteration we re-compute σ^2 in the context of the BG model as the variance of GDP growth. Furthermore, the prior that we use for the BG model is a relatively loose multivariate normal prior, similar to Szabo (2020). Specifically, we run a quantile regression using $T - n$ observations and use the respective estimates for the prior. Given a specific quantile we implement the respective quantile parameter estimates as the prior mean. Additionally, the prior covariance matrix is diagonal, with the respective quantile residual variance estimate on the diagonal. We choose to model 40 quantiles simultaneously in order to include the 5% and the 10% quantile. Moreover, to allow for straightforward evaluation we take $h = 1$ to predict the distribution of future GDP growth one quarter ahead. We set $k = 1e^5$ and define half of it as burn-in period and set thinning at 10, similar to Feng et al. (2015). In the context of the MS model we simulate $L = 1e^5$ times as well. The specific model specification we implement contains the explanatory variables discussed in section 3. Moreover, we include one lagged GDP growth variable in the model based on the autocorrelation function of the respective GDP growth series.

In order to evaluate the GaR predictions of the models, we need to adjust the context of the estimates to fit the context of the evaluation criteria. Since, these evaluation criteria are based on VaR risk measures. Therefore, we mirror the respective GDP growth series and GaR estimates around the horizontal axis and take the $1-\tau$ quantile. As we determine the quantile at the left tail of the distribution and in the context of VaR the analysis focuses on the right tail of the distribution.

3 Data

The main variable of interest in this paper is GDP growth. In order to arrive at robust conclusions regarding GaR estimation, we choose to consider multiple GDP growth rates. Specifically, we choose the US[1] and German[2] GDP growth. Both GDP growth rates are given in quarterly frequency from 1997 Q1 until 2020 Q4.

As alluded to before the NFCI is a sensible explanatory variable to include in a model that aims to describe the behaviour of GDP growth, according to Adrian et al. (2019).[3] Furthermore, the NFCI seems to explain a significant amount of the variation in GDP growth, according to Brownlees and Souza (2021). The variable constitutes as an average of numerous financial measures that gauge economic activity. Therefore, we incorporate the variable as an explanatory variable in the models. The NFCI is given in weekly frequency, so we convert it to quarterly frequency following the method presented in Adrian et al. (2019). We average weekly NFCI in the quarter and if one week starts in one quarter and ends in the following quarter, the week is attributed to the latter. Thus, we obtain NFCI in quarterly frequency.

Another variable that can increase explanatory power of the models is affiliated with financial stress. A Financial Stress Index (FSI) can be viewed as an indicator that generalises the current state of the economy, which presents a direction the economy is likely to head towards, according to Islami and Kurz-Kim (2014). The inherent construction of an FSI is likely to fit well in a model that predicts future GDP growth. So, we include the FSI in the model as an explanatory variable as well.[4] The FSI is obtained in weekly frequency, so we convert the time series to quarterly frequency in a similar manner as the NFCI.

Furthermore, policy measures are also likely to affect the state of the economy, as shown in Baker et al. (2016). Specifically, uncertainty regarding policy contributes to decreasing economic activity. The Economic Policy Uncertainty (EPU) gauges the amount of uncertainty revolving around economic policy based on newspaper coverage frequency. By quantifying the number of a specific set of words one is able to determine the amount of uncertainty in the economy according to Baker et al. (2016). This variable is therefore likely able to describe the behaviour of GDP growth well and is included in the model.[5] The EPU is given in monthly frequency, so we convert to

[1] Obtained from https://fred.stlouisfed.org/series/A191RL1Q225SBEA

[2] Obtained from https://sdw.ecb.europa.eu/browse.do?node=9683074

[3] The NFCI is obtained from https://www.chicagofed.org/research/data/index

[4] The FSI is obtained from https://fred.stlouisfed.org/series/STLFSI2

[5] The global EPU is obtained from http://www.policyuncertainty.com/global_monthly.html

quarterly frequency following Chow and Lin (1971). The method states that we average the monthly EPU that constitute as a quarter and thus obtain quarterly EPU. Ultimately, we end up with a data set that starts in 1997 Q1 and ends in 2020 Q4. In Figure 1 we plot the GDP growth series to gain more insight into the characteristics of the series.

Figure 1: US and German GDP growth over time

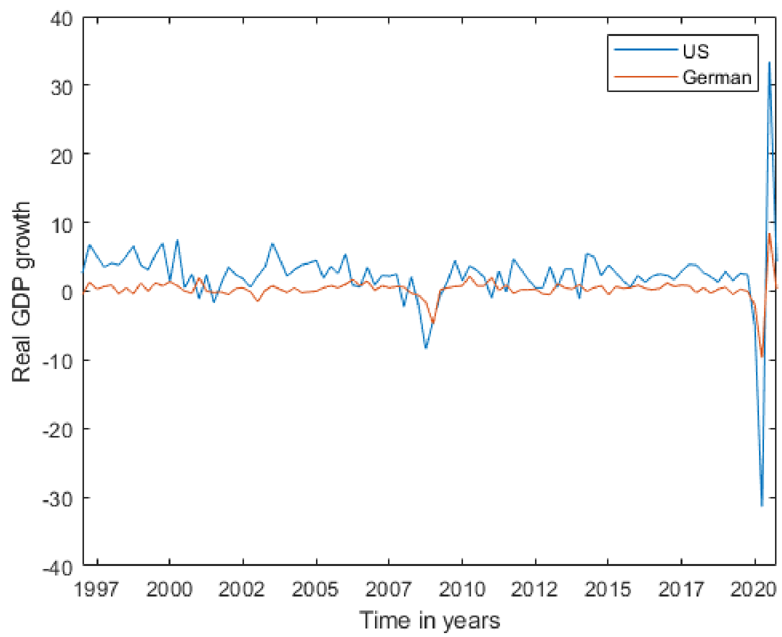

Note. Quarterly US and German GDP growth is given in the figure from 1997 Q1 to 2020 Q4. Where the blue line represents US GDP growth and the red line represents German GDP growth.

In Figure 1 the quarterly US and German GDP growth is given over time. The main difference we observe between the two series is the variation relative to their respective mean. US GDP growth seems a more volatile series compared to German GDP. Furthermore, the outliers of GDP growth over time correspond with the finanical crisis in 2008 and the corona crisis in early 2020. To continue the analysis of the series, we show the descriptive statistics of the variables in Table 1.

Table 1: Descriptive statistics

	US GDP	German GDP	NFCI	EPU	FSI
Mean	2.320	0.298	-0.360	126.601	0.041
Variance	28.191	2.507	0.255	4148.798	0.977
Skewness	-0.702	-1.529	3.391	1.565	3.700

Note. In the Table descriptive statistics of the respective time series are shown, where the mean, variance and skewness are displayed as descriptive statistics.

In Table 1 the descriptive statistics for the various variables are shown. As displayed in Figure 1 the variation in US GDP is relatively higher than the variation in German GDP, which is reflected in the variance of the respective series. Furthermore, the NFCI seems to exhibit relatively little variation, whereas the EPU and FSI display relatively much variation. The GDP growth series have negative skewness, implying that negative GDP growth has larger influence on the observations than positive GDP growth. The contrary holds for the EPU, NFCI and FSI. Furthermore, we use Augmented Dickey-Fuller tests to check the time series for possible non-stationarity. We find that the EPU series is the only non-stationary series used in this paper. The EPU series is implemented without taking first differences, as the non-stationary series does not affect parameter estimates according to Baffes (1997). Furthermore, we normalise EPU as the EM algorithm would experience great difficulty computing the estimated parameters due to the magnitude of the data.

In order to evaluate the models we use a moving window to predict out-of-sample GaR. As we are evaluating GaR estimates using backtesting techniques it is relevant to consider the length of the out-of-sample period with respect to the amount of violations that can occur. For every out-of-sample period we choose a length of 40 quarters, as GaR estimates based on the 5% quantile are expected to violate twice. This seems a reasonable amount of out-of-sample observations, as fewer out-of-sample observations would allow less than two violations and more out-of-sample observations would cause the estimation window to become too brief.

4 Results

In this section we elaborate on the results that are obtained from the techniques discussed in the methodology.

We start the analysis by performing the backtests discussed in the methodology in various settings to analyse the performance of the models. Table 2 displays the results of the various backtests when modelling US GDP growth based on the 5% quantile.

Table 2: P-values for various evaluation criteria of out-of-sample GaR predictions for US GDP growth based on the 5% quantile

	Traditional	Conditional	Comparative
Bayesian GaR	0.147	0.428	0.602
Markov switching	$1.344e^{-5}$	0.020	

Note. In this table p-values are shown for the traditional, conditional and comparative backtests using the respective Bayesian GaR and Markov switching model. The GaR predictions are based on the 5% quantile from the distribution of future US GDP growth using the full sample period. The comparative backtest shows the p-value for testing the null hypothesis that the performance of the two methods are indifferent.

We observe discrepancies between the performance of the BG and MS model. For the BG model, we do not reject the null hypotheses of the traditional and conditional backtests at a 5% significance level, implying that the Bayesian GaR model predicts GaR adequately. Contrary to the BG model, we reject the null hypotheses of both individual backtests when examining the MS model. Furthermore, we do not reject the null hypothesis of the comparative backtest at a 5% significance level. Therefore, the Bayesian GaR and Markov switching model predict GaR at least as well as one another. Although the individual backtests suggest that the MS model is unable to predict GaR adequately, the contrary holds for the BG model. These results might be due to the fact that the GaR predictions of the respective models differ substantially in absolute terms. If the MS model predicts relatively small GaR, its score is relatively small as well, which causes the violations to be less penalised. Furthermore, the out-of-sample period might be too short to evaluate GaR accurately given that these predictions are based on the 5% quantile. We continue the analysis by examining GaR predictions for the German GDP growth series. Table 3 shows p-values from the backtests when modelling German GDP growth based on the 5% quantile.

Table 3: P-values for various evaluation criteria of out-of-sample GaR predictions for German GDP growth based on the 5% quantile

	Traditional	Conditional	Comparative
Bayesian GaR	0.468	0.616	0.128
Markov switching	$2.863e^{-4}$	0.017	

Note. In this table p-values are shown for the traditional, conditional and comparative backtests using the respective Bayesian GaR and Markov switching model. The GaR predictions are based on the 5% quantile from the distribution of future German GDP growth using the full sample period. The comparative backtest shows the p-value for testing the null hypothesis that the performance of the two methods are indifferent.

From Table 3 we deduct similar results compared to Table 2. The BG model predicts GaR adequately based on the individual backtests at a 5% significance level. However, the MS model does not predict GaR adequately. In terms of the comparative backtest, we do not reject the null hypothesis, implying that both models are able to predict as well as one another at a 5% significance level. The main distinction between the results in Table 2 and Table 3 revolves around the performance between the two models. In the context of German GDP growth, the results suggest that the GaR estimates differ substantially in absolute terms, however insufficient to state that the MS model is outperformed by the BG model. Overall, GaR predictions of the models are comparable in terms of performance, though the Bayesian GaR model is favoured due to the individual backtest results. To continue our analysis we examine the differences between the GaR predictions by means of graphs.

In Figure 2 a plot of US GDP growth and predicted out-of-sample GaR is displayed based on the 5% quantile. The majority of the time the BG GaR seems to prove relatively accurate, however there are a few instances where the GaR estimate is higher than the realised US GDP growth. This holds especially in the beginning of the year 2020, which can be attributed to the corona crisis. In the first quarter of 2020 the US economy shrunk approximately 30% and in the consecutive quarter it expanded approximately 30%. After observing the enormous decline in economic activity, the model adjusts and predicts continuing economic decline in the next period. In the ensuing quarter the unusual expansion occurs and the positive and negative effects offset each other. This causes the following prediction to converge to regular values. As is the case for the predicted BG GaR, the MS GaR appears relatively adequate. However, there are multiple instances where the estimated MS GaR is higher than the realised US GDP growth. This holds for the Covid period around the year 2020 and numerous times at the start of the out-of-sample period, although US GDP growth does not exhibit major outliers at the start of the out-of-sample period. Generally, the MS GaR estimates are closer to US GDP growth than the BG GaR estimates. This is potentially due to the fact that the Bayesian GaR model takes parameter uncertainty fully into account by incorporating the posterior distribution of the parameters. In this context, the estimates reflect more uncertainty causing the BG GaR predictions to differ considerably compared to GDP growth in absolute terms. Contrary to the Markov switching model, where point estimates are used that might not fully reflect the parameter uncertainty that is inherently present in the true estimates. This might lead to the discrepancies between the GaR predictions of the models in absolute terms. Moreover, it might also explain the comparative backtest results.

Figure 2: US GDP growth and out-of-sample GaR predictions based on the 5% quantile

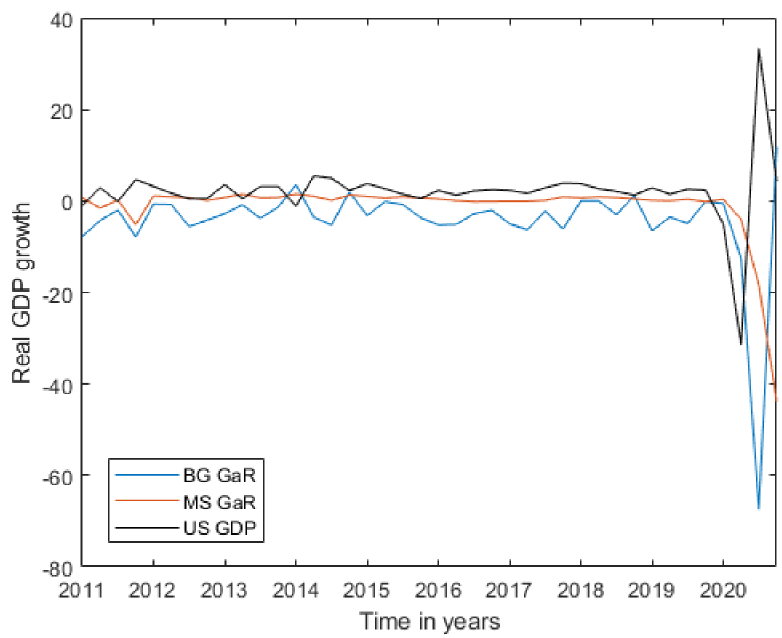

Note. This figure shows a plot of US GDP growth and their respective GaR predictions based on the 5% quantile from 2011 Q1 to 2020 Q4. Where the blue line represents the GaR predictions using the Bayesian GaR model and the red line GaR predictions using the Markov switching model and the black line represents real US GDP growth.

In Figure 3 a plot of German GDP growth and predicted out-of-sample GaR is shown based on the 5% quantile. As alluded to before, the main distinction between German GDP growth and US GDP growth is the magnitude of the outliers. Due to the less volatile nature, the performance of the models might be superior in the context of German GDP growth. For the BG model this appears to be the case as the BG GaR estimates capture the trend of GDP growth more adequately in Figure 3 compared to Figure 2. Nonetheless, the Covid period remains difficult to accurately predict, causing the GaR estimates to be inadequate. The BG GaR predictions continue to exhibit violations for this period. We observe similar characteristics for the MS GaR estimates to those displayed in previous plots. The estimates are relatively close to realised German GDP growth, which is likely caused by the point estimates that are impotent when uncertainty plays a major role in the GaR predictions. Furthermore, the response to the Covid period is stronger in the context of the Markov switching model. This might again be due to the point estimates, which remain persistent contrary to a posterior distribution of parameters. However, as the comparative backtest suggests it is difficult to determine which model is better suited to predict GaR. Therefore, we extend the analysis by examining the performance of the models in the context of the 10% quantile from the distribution of future GDP growth.

Figure 3: German GDP growth and out-of-sample GaR predictions based on the 5% quantile

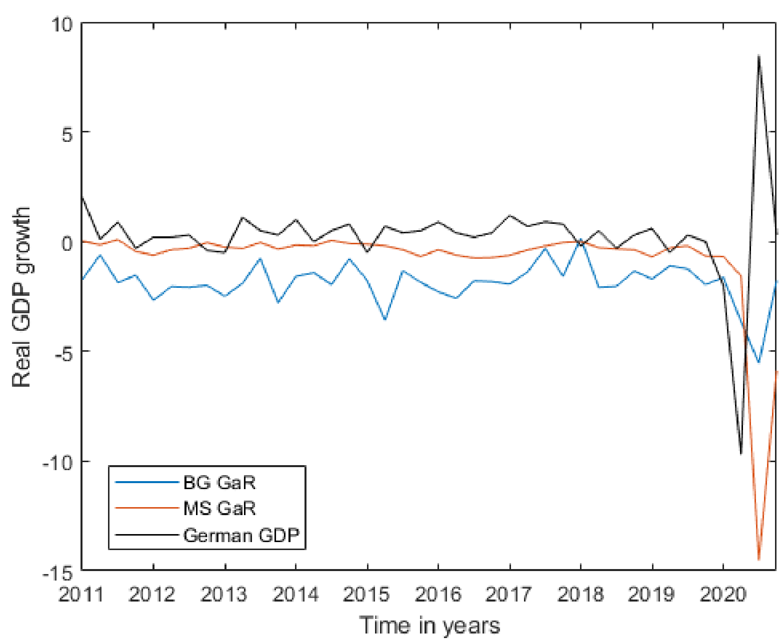

Note. This figure shows a plot of German GDP growth and their respective GaR predictions based on the 5% quantile from 2011 Q1 to 2020 Q4. Where the blue line represents the GaR predictions using the Bayesian GaR model and the red line GaR predictions using the Markov switching model and the black line represents real German GDP growth.

Next, we analyse the prediction of the 10% quantile from the distribution of future GDP growth. In Table 4 p-values are shown when modelling US GDP growth based on the 10% quantile.

Table 4: P-values for various evaluation criteria of out-of-sample GaR predictions for US GDP growth based on the 10% quantile

	Traditional	Conditional	Comparative
Bayesian GaR	0.598	0.672	0.077
Markov switching	0.008	0.112	

Note. In this table p-values are shown for the traditional, conditional and comparative backtests using the respective Bayesian GaR and Markov switching model. The GaR predictions are based on the 10% quantile from the distribution of future US GDP growth using the full sample period. The comparative backtest shows the p-value for testing the null hypothesis that the performance of the two methods are indifferent.

From Table 4 we observe similar results to those displayed in Table 2. Though, the MS model is able to predict GaR adequately based on the 10% quantile according to the conditional backtest. Furthermore, the comparative backtest implies that the models should predict at least as well as one another at a 5% significance level. However, we reject the null hypothesis at a 10% significance level, implying that the difference in adequacy between the GaR predictions is more apparent for the 10% quantile. In Table 5 p-values are shown when modelling German GDP growth based on the 10% quantile.

Table 5: P-values for various evaluation criteria of out-of-sample GaR predictions for German GDP growth based on the 10% quantile

	Traditional	Conditional	Comparative
Bayesian GaR	0.008	0.037	0.233
Markov switching	0.002	0.008	

Note. In this table p-values are shown for the traditional, conditional and comparative backtests using the respective Bayesian GaR and Markov switching model. The GaR predictions are based on the 10% quantile from the distribution of future German GDP growth using the full sample period. The comparative backtest shows the p-value for testing the null hypothesis that the performance of the two methods are indifferent.

Table 5 offers a new perspective on the analysis of the models. We observe rejections of the null hypotheses for all individual backtests, entailing that the respective GaR predictions are inadequate. The mediocre performance might be attributed to the fact that the BG model predicts all quantiles simultaneously. By considering cross-quantile effects the parameter values can incidentally be contaminated causing inadequate GaR predictions. Nonetheless, the GaR predictions of the Markov switching model do not seem to be more adequate, which is confirmed by the comparative backtest.

Overall, it remains difficult to favour one model substantially over the other due to the results of the comparative backtest. Therefore, we continue the analysis by examining characteristics of the models in more detail, in order to learn why the models display their respective performance. Figure 4 shows posterior GaR samples for predicted German GDP growth at 2020 Q4. We notice that the GaR ranges approximately between -20% and 30%. Furthermore, the majority of the mass is centered between -10% and 10%. This indicates that there exists a relatively large amount of uncertainty in the out-of-sample GaR samples. These specific samples mark the start of a relatively less volatile period after the extremely volatile Covid period in the first two quarters of 2020, as shown in Figure 1, which explains the amount of uncertainty in the samples. This causes GaR to be relatively unpredictable at this point in time. On the other hand, this figure confirms that the Bayesian GaR model takes uncertainty into account in the right and left tail when predicting GaR. The majority of the mass in the posterior distribution is located in the left tail, as negative GDP growth causes major downsides relative to positive GDP growth. Therefore, the Bayesian GaR model is likely to predict relatively low GaR, as displayed in previous plots.

Figure 4: Posterior GaR samples for German GDP growth based on the 5% quantile

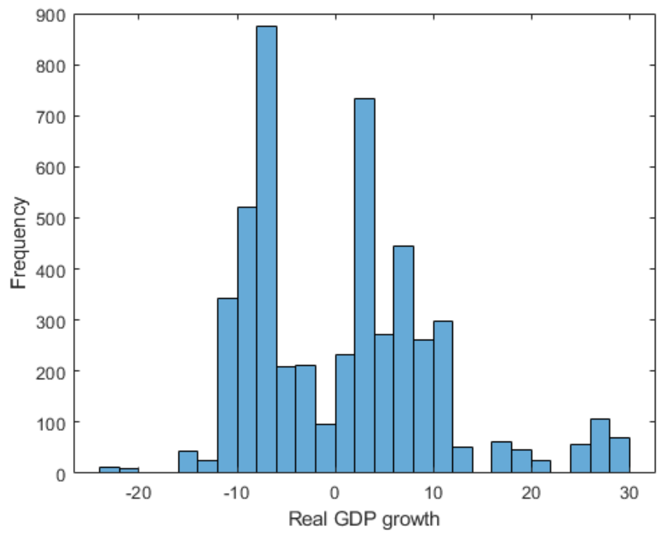

Note. This figure displays posterior GaR samples based on the 5% quantile for future German GDP growth in 2020 Q4. These samples are obtained by first implementing the Bayesian GaR model to obtain samples from the posterior distribution and then computing GaR.

In order to examine the Markov switching model in more detail, we analyse the estimated smoothed probabilities. In Figure 5 the smoothed probability series is shown when estimating the MS model for US GDP growth. In the plot we observe a relatively low number of points in time where it is likely that the model is in state 2. The smoothed probabilities for state 2 are close to 1 when real US GDP growth is extremely volatile, as the dates match with the financial crisis and corona crisis. During these periods in time a major amount of uncertainty was present which resulted in volatile GDP growth. Hence, the MS model is able to distinguish between the states based on the volatile behaviour of real GDP growth at various points in time. This is corroborated by Table 6 in Appendix A, which displays the estimated parameters of the MS model in the same setting as Figure 5. The magnitude of the state 1 estimates is relatively small compared to the state 2 estimates. In the recession regime described by state 2, future GDP growth is relatively dependent on the explanatory variables. This dependence might be due to the ability of the variables to describe the behaviour of volatile GDP growth in an adequate manner, as larger estimates in absolute terms can amplify the predictions. These predictions can vary substantially, potentially matching the volatile behaviour of GDP growth in the recession regime. Therefore, the explanatory variables exert more influence on future GDP growth in the recession regime. Additionally, we notice that there is a major amount of uncertainty present in the point estimates, as implied by their respective standard error. The table also displays the transition probabilities, which indicate that the model is likely to be situated in state 1 and not likely to transition to state 2. Incidentally, this means that the model is unlikely to switch to the recession regime, implying that the predictions are overall less volatile due to the relatively modest state 1 estimates.

Moreover, we observe a strong response from the MS model to the Covid period in terms of the GaR predictions. Figure 5 implies that US GDP growth is relatively volatile when the model is likely in state 2. During that time the large shrinkage of GDP growth occurs and the model predicts shrinkage in the following period, while GDP growth has recovered and displays positive GDP growth. In the next quarter the observed GDP growth stabilises, while the large expansion of GDP growth in the previous quarter still leads to a low GaR forecast. These events lead to the distinct GaR predictions displayed in previous plots during the Covid period. Although the MS model identifies the Covid period as a recession regime, the model is not able to predict the movements of GDP growth adequately.

Figure 5: Smoothed probability of US GDP growth over time

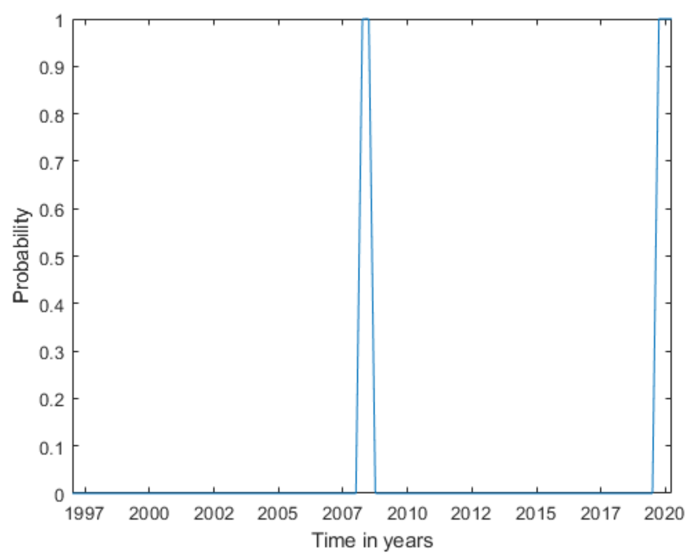

Note. This figure displays the smoothed probability of being in state 2 from 1997 Q1 to 2020 Q2.

In Appendix B, Figure 7 shows the smoothed probabilities when estimating the MS model for German GDP growth. This plot is similar to the plot shown in Figure 5, however there is a clear distinction around the year 2013, indicating that the model likely transitions to a relatively volatile regime around the year 2013.

Since the Covid period affects the results substantially, we examine the smoothed probabilities in the absence of the Covid period. Figure 8 in Appendix B displays the smoothed probabilities for US GDP growth when excluding the Covid period. As the Covid period is excluded, the transition between regimes seems less apparent. Based on the estimates state 2 implies a relatively volatile regime, however the discrepancy between the states is smaller in this context. This indicates that the MS model has difficulty predicting sudden recessions accurately and fails to take uncertainty revolving around GDP growth adequately into account. This in turn leads to inadequate GaR predictions as displayed in previous plots. Primarily, the Covid period has major influence on the results, therefore we also include analysis on a sample period in the absence of the Covid period.

Figure 6 shows predicted GaR and real US GDP growth based on the 5% quantile when omitting the Covid period. The GaR predictions using the BG model are consistently below realised US GDP growth. This confirms the impact of the Covid period on the models, along with the ability of the BG model to track the patterns of real US GDP growth relatively well. Furthermore, the MS model fails to predict GaR adequately alongside the variation contained in real US GDP growth. As mentioned above, the GaR predictions of the MS model are relatively close to the realised GDP growth series in absolute terms. Generally, the MS model is impotent in terms of predicting the variation in the GDP growth series. This is confirmed by the comparative backtest in this context, which results in a p-value of 0.039. Therefore, we reject the null hypothesis at a 5% significance level and find that the Bayesian GaR model predicts GaR significantly more adequately than the Markov switching model in the absence of the Covid period for US GDP growth.

Figure 6: US GDP growth and out-of-sample GaR predictions based on the 5% quantile excluding the Covid period

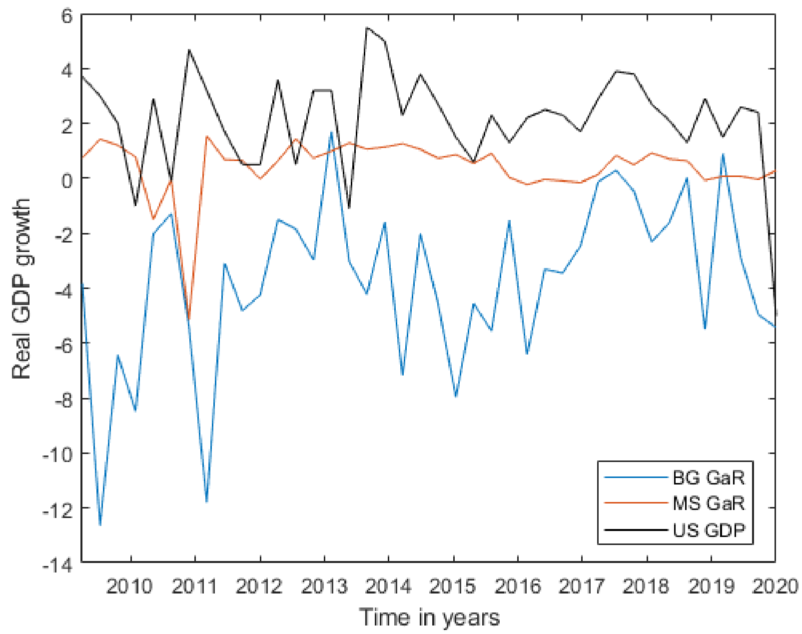

Note. This figure shows a plot of US GDP growth and their respective GaR predictions based on the 5% quantile from 2009 Q2 to 2020 Q1. Where the blue line represents the GaR predictions using the Bayesian GaR model and the red line GaR predictions using the Markov switching model and the black line represents real US GDP growth.

Figure 9 in Appendix C shows a plot of the predicted GaR and German GDP growth based on the 5% quantile without the Covid period. The patterns and characteristics in Figure 6 are present here as well. However, at the end of the out-of-sample period, the BG GaR predictions exhibit several violations. Furthermore, the MS GaR predictions are relatively close to real German GDP growth, implying that the MS model is unable to deal with small negative shocks. This is corroborated by the comparative backtest, which has a p-value of 0.017. Thus, we reject the null hypothesis at a 5% significance level and find that the Bayesian GaR model predicts GaR significantly more adequately than the Markov switching model in the absence of the Covid period for German GDP growth.

Overall, the Bayesian GaR model predicts GaR more adequately than the Markov switching model. Based on individual backtests the Bayesian GaR model is able to adequately predict GaR the majority of the time contrary to the Markov switching model. This holds in the context of German and US GDP growth and across the 5% and 10% quantile of the distribution of future GDP growth.

The rationale behind these results might revolve around the inherent characteristics the two models exhibit. The Bayesian GaR model imposes Bayesian inference on a quantile regression, resulting in a posterior distribution of GaR. This posterior distribution contains all uncertainty captured in the distribution of future GDP growth. Therefore, the resulting GaR predictions take left-skewed shocks relatively accurate into account. On the other hand, the Markov switching model imposes inference on latent states. Through estimation one obtains different parameter estimates in the distinct states the model can transition between and implement these accordingly in the predictions. Nonetheless, the model is incapable of distinguishing accurately between large recessions and small shocks in real GDP growth by means of the unobserved states. Generally, this results in inadequate GaR predictions based on the individual backtests. However, from the comparative backtests we deduct that both models should predict at least as well as each other. Furthermore, we find that the Covid period has substantial influence on the results, so further analysis in the absence of the Covid period is conducted. The Bayesian GaR model is persistent in predicting GaR adequately, while the Markov switching model falters relatively. The comparative backtest substantiates this, as the null hypotheses are rejected. This implies that the Bayesian GaR model is significantly better in predicting GaR than the Markov switching model.

5 Conclusion

To learn the behaviour of the macro economy multiple venues can be explored, this paper examines the behaviour of the economy through predicting Growth-at-Risk. The main difficulty revolves around uncertainty contained in the left tail of the distribution of future GDP growth. Negative events, such as recessions, influence the estimates excessively. A regular quantile regression is inept in dealing with such complexities. Therefore, we evaluate the Bayesian GaR model and the Markov switching model to determine which model deals with left-skewed GDP growth during recessions in the best manner.

Our analysis shows that the Markov switching model is generally unable to provide adequate GaR predictions in terms of the individual backtests. On the other hand, the Bayesian GaR model estimates GaR adequately based on the individual backtests in the majority of the settings. However, the comparative backtests suggest that both models should be able to predict at least as well as each other. We find that the Covid period has major influence on the evaluation of GaR predictions. By evaluating the models in the absence of the Covid period, we find major discrepancies between the Markov switching model and the Bayesian GaR model. Therefore, we conclude that the Markov switching model is unable to predict GaR adequately relative to the Bayesian GaR model. This is likely due to the fact that the model is unable to clearly distinguish between large recessions and small negative shocks in real GDP growth. Therefore, the Bayesian GaR model takes the left tail of the distribution of future GDP growth more accurately into account. This can be explained by the fact that the Bayesian GaR model considers all uncertainty contained in the GaR predictions when estimating the risk measure.

The Bayesian GaR model is thus better suited at predicting GaR. Therefore, it is beneficial to implement the model in practice to obtain adequate GaR predictions. When policy-makers obtain these accurate estimates, they can act on them accordingly. For example, if the model signals that a large amount of uncertainty is present in the economy, measures can be taken to restrict spending by the government and prepare for a possible decline in economic activity. Furthermore, new knowledge regarding the behaviour of GDP growth can also be utilised in portfolio management, to cover potential downside risk.

Though, we conclude that the Markov switching model is unable to predict GaR adequately there are multiple venues that might be explored further. In this paper we estimate the models under a specific model specification. Perhaps this is sub-optimal for the Markov switching model and other specifications might improve performance, such as auxiliary explanatory variables or non-linear elements. However, this is beyond the scope of this paper. Furthermore, implementing more than two states in the Markov switching model might improve the GaR predictions. On the other hand, the Markov switching model experiences great difficulty in estimating the latent state variables in this paper. Increasing the number of regimes might only increase the difficulty the model experiences estimating them. Lastly, obtaining GaR predictions using other methods could offer a new perspective. In this paper, we simulate GDP growth and obtain a distribution of future GDP growth of which we take a specific quantile. It might be possible to model future GDP growth by means of a GARCH model.

References

Tobias Adrian, Nina Boyarchenko, and Domenico Giannone. Vulnerable growth. *American Economic Review*, 109(4):1263–89, 2019.

John Baffes. Explaining stationary variables with non-stationary regressors. *Applied Economics Letters*, 4(1):69–75, 1997.

Scott R Baker, Nicholas Bloom, and Steven J Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

Christian Brownlees and Andre BM Souza. Backtesting global Growth-at-Risk. *Journal of Monetary Economics*, 118:312–330, 2021.

Gregory C Chow and An-loh Lin. Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *Review of Economics and Statistics*, 53(4):372–375, 1971.

Peter Christoffersen and Denis Pelletier. Backtesting Value-at-Risk: A duration-based approach. *Journal of Financial Econometrics*, 2(1):84–108, 2004.

Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1): 1–22, 1977.

Yang Feng, Yuguo Chen, and Xuming He. Bayesian quantile regression with approximate likelihood. *Bernoulli*, 21(2):832–850, 2015.

Thomas H Goodwin. Business-cycle analysis with a Markov-switching model. *Journal of Business & Economic Statistics*, 11(3):331–339, 1993.

James D Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989.

James D Hamilton. Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45(1-2):39–70, 1990.

Mevlud Islami and Jeong-Ryeol Kurz-Kim. A single composite financial stress indicator and its real impact in the euro area. *International Journal of Finance & Economics*, 19(3):204–211, 2014.

Chang-Jin Kim. Dynamic linear models with Markov-switching. *Journal of Econometrics*, 60(1-2): 1–22, 1994.

Natalia Nolde and Johanna F Ziegel. Elicitability and backtesting: Perspectives for banking regulation. *Annals of Applied Statistics*, 11(4):1833–1874, 2017.

Ananthakrishnan Prasad, Selim Elekdag, Phakawa Jeasakul, Romain Lafarguette, Adrian Alter, Alan Xiaochen Feng, and Changchun Wang. *Growth at Risk: Concept and application in imf country surveillance*. International Monetary Fund, 2019.

Milan Szabo. Growth-at-Risk: Bayesian approach. Technical report, 2020.

Appendices

A Parameter estimates using the Markov switching model

Table 6: Point estimates using the Markov switching model for US GDP growth

	State 1 estimates	State 2 estimates
Constant	2.419(10.919)	30.220(121.926)
Lagged GDP	0.079(0.956)	-1.804(10.680)
NFCI	-1.493(16.405)	54.302(183.194)
EPU	-62.810(789.014)	-552.888(8.811e^3)
FSI	-0.002(8.386)	-42.190(93.646)
p_{11}	0.977	-
p_{22}	-	0.742
σ_1^2	3.3733	-
σ_2^2	-	2.500e^{-4}

Note. This table displays estimates using the Markov switching model for US GDP growth with standard errors of the estimates given in parentheses. Furthermore, the estimates for the transition probabilities are given combined with the respective estimated variance of the error term.

B Smoothed probabilities in various settings

Figure 7: Smoothed probability of German GDP growth over time

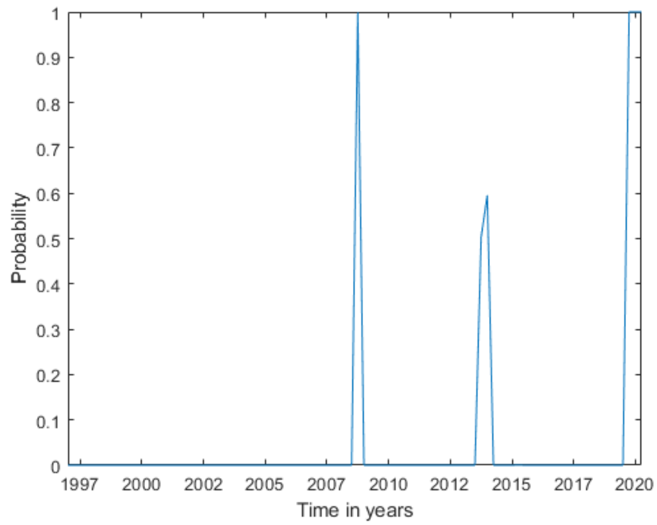

Note. This figure displays the smoothed probability of being in state 2 from 1997 Q1 to 2020 Q2

Figure 8: Smoothed probability of US GDP growth over time excluding the Covid period

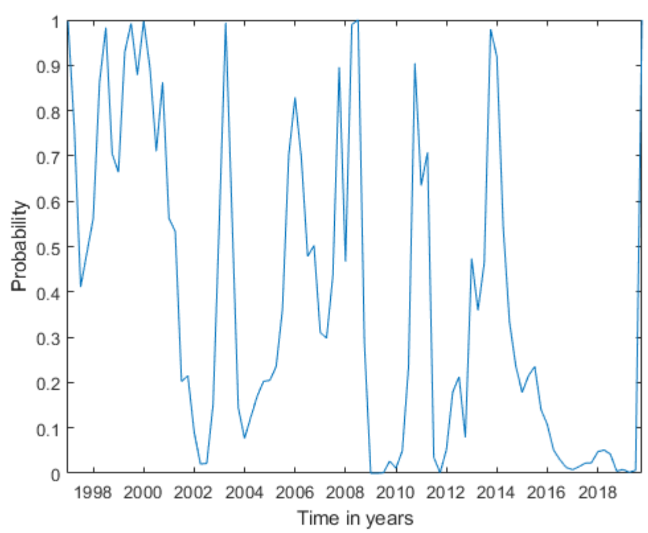

Note. This figure displays the smoothed probability of being in state 2 from 1997 Q1 to 2019 Q4

C GaR predictions excluding the Covid period

Figure 9: German GDP growth and out-of-sample GaR predictions based on the 5% quantile excluding the Covid period

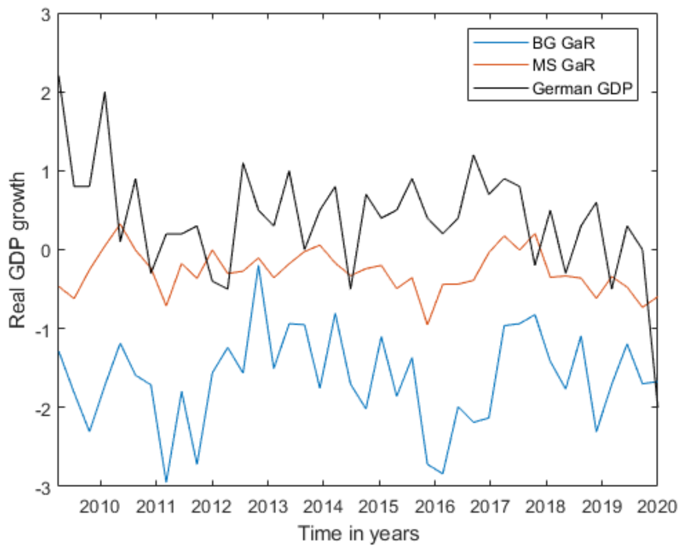

Note. This figure shows a plot of German GDP growth and their respective GaR predictions based on the 5% quantile 2009 Q2 to 2020 Q1. Where the blue line represents the GaR predictions using the Bayesian GaR model and the red line GaR predictions using the Markov switching model and the black line represents real German GDP growth.