ИКОНОМИЧЕСКИ ИЗСЛЕДВАНИЯ
Economic Studies
Volume 31(7), 2022

Mihail Yanchev[1]

DEEP GROWTH-AT-RISK MODEL: NOWCASTING THE 2020 PANDEMIC LOCKDOWN RECESSION IN SMALL OPEN ECONOMIES[2]

Accurate forecasting of the timing and magnitude of macroeconomic recessions caused by unexpected shocks remains an area where both statistical models and judgmental forecasts tend to perform poorly. Inspired by the value-at-risk concept from financial risk management, a growing body of research has been focused on developing a framework to model and quantify macroeconomic risks and estimate the likelihood of adverse macroeconomic outcomes, which has become known as growth-at-risk assessment. The current study proposes an improvement to an established two-step procedure for empirical evaluation of the future growth distribution, which involves directly modelling the parameters of the conditional distribution in one step within an artificial neural network. The proposed procedure is tested on macroeconomic data from four small European open economies covering the coronavirus pandemic lockdown period and the recession related to it. The model achieves a better performance across the four countries compared to the established two-step procedure.
Keywords: Forecasting; Macroeconomic Risks; Artificial Neural Networks; Density Forecasts; Recessions
JEL: C53; E17; E27; E32

1. Introduction

Recessions are not rare events, according to An et al. (2018). The authors analyzed data on 153 recession episodes across 63 countries between 1992 and 2014, and found that countries, on average, are in a recession 12% of the time. However, recession events and their timing and magnitude remain hard to predict for both experts and statistical models (Lewis, Pain, 2014). On the other hand, more impactful events like the great recession that occurred between 2007 and 2009 and the recent recession caused by the coronavirus pandemic lockdown are an even greater challenge for forecasters and decision-makers as they represent realizations of low probability risks (Makridakis et al., 2009; Chen, 2019; Antipova, 2020). While the great recession was caused by a build-up of systemic risk, which in retrospect turned out to be visible in the data (Altunbaz et al., 2017), the coronavirus pandemic

[1] *Mihail Yanchev, Ph.D. Candidate, Sofia University "St. Kliment Ohridski", Faculty of Economics and Business Administration, phone: +359886435629, e-mail: yanchev.mihail@gmail.com.*
[2] This paper should be cited as: *Yanchev, M. (2022). Deep Growth-at-Risk Model: Nowcasting the 2020 Pandemic Lockdown Recession in Small Open Economies. – Economic Studies (Ikonomicheski Izsledvania), 31(7), pp. 20-41.*

lockdown was caused by an unusual and unexpected shock. Therefore, this latest crisis can be considered one of the biggest challenges for the forecasting profession in recent decades.

In answer to such challenges, the International Monetary Fund (IMF), among other institutions, has been using a framework for quantifying macroeconomic risks to growth, which has become known as growth-at-risk (Prasad et al., 2019). Since models designed to forecast a central feature of the distribution of interest, like the mean or the median, are unable to capture asymmetries between upside and downside risks, the assessment of the uncertainty surrounding point forecasts becomes necessary (Clemens, 2004). One way to address this necessity, which is supported by a growing body of research recently and is at the core of the IMF growth-at-risk framework, is to model empirically the future growth distribution on the basis of current macroeconomic and macro-financial conditions. While different models have been used to achieve this task, including Bayesian VAR models (Carriero, 2020), stochastic volatility models (Iseringhausen, 2021) and GARCH models (Brownlees, Souza, 2021), this paper focuses on methods based on quantile regression.

In an influential paper, Adrian et al. (2019) use a two-step procedure of constructing conditional quantiles using a quantile regression model and consequently fit a probability distribution to the estimated quantiles. The authors studied the conditional US growth distribution with an emphasis on financial conditions. They identified several stylized facts about the conditional distribution of growth for the USA, among which a strong negative correlation between the conditional mean and variance and a significant relationship between current financial conditions and future shifts in the lower tail of the conditional distribution. The same conclusion was confirmed by De Santis and Van der Veken (2020), who performed a similar exercise, including data from the beginning of 2020 and a separate dataset covering the Spanish flu pandemic period across a number of countries. Figuerez and Jarociński (2020) confirm the same stylized facts identified by Adrian et al. (2019) for the Euro Area.

The current paper proposes an improvement upon the semi-parametric two-step procedure used by Adrian et al. (2019) and De Santis and Van der Veken (2020) by proposing a one-step model, which is based on artificial neural networks and directly outputs the parameters of the conditional growth distribution. The model still depends internally on the estimation of conditional quantiles and for this purpose, it is based on two separate loss functions, which are being dynamically weighted. The improvements proposed here lie in four separate areas:

1. a simultaneous generation of quantiles, as proposed by Rodriguez and Pereira (2020), in order to alleviate the quantile crossing problem;
2. the introduction of quantile crossing loss to the tilted loss function, which further prevents quantile crossing as proposed by Bondell et al. (2010);
3. using artificial neural network architecture based on long short-term memory (LSTM) layers (Hochreiter, Schmidhuber, 1997) to model non-linear relationships between the predictors and the target variable and better capture the recession related to the pandemic lockdown compared to a linear model;
4. combining the two steps of the procedure into a single model, which is being optimized by minimizing two loss functions simultaneously – the tilted absolute loss function used

for estimating the conditional quantiles and a least squares loss for evaluating the final conditional distribution parameters.

This combination of improvements is called deep growth-at-risk model or DeepGaR in short for the purposes of this paper. Initially, the focus of the study was on the macroeconomic developments in Bulgaria, but after preliminary results were generated, it was decided to test the proposed approach on three other small open European economies, relatively similar in terms of the size and structure of the economy. Therefore, the proposed procedure was tested on data for Bulgaria, Estonia, Lithuania and Romania, covering the coronavirus pandemic lockdown period and the recession related to it, and achieved a better out-of-sample performance across four of them compared to the established two-step procedure.

The paper is structured as follows. The next section covers the methodology, including an overview of the established procedure, as well as the proposed improvements. The third section covers the data used in this study, while the fourth section summarizes the performance test results. The last section contains a discussion of the results and the conclusions of the study.

2. Methodology

A central idea underlying the quantile regression model, as defined by Koenker and Bassett (1978), is that a task of sorting can be turned into an optimization problem. Just as finding a sample mean can be done by minimizing the sum of squared errors, finding the median can result from minimizing the sum of absolute errors. They further elaborate to show that an asymmetrical loss function which gives different penalties to positive and negative residuals, can yield any quantile for a given sample. Solving for the following equation (1) yields the τ-th quantile as its solution:

$$min_{\xi \in \mathbb{R}} \sum_{i=0}^{n} \rho_\tau(y_i - \xi) \tag{1}$$

Where $0 < \tau < 1$ and $\rho_\tau(\cdot)$is the titled absolute value function, which can be seen on figure 1, for a sample of size n. In this equation, if τ is set to equal 0.5, the equation will yield the median. Therefore, if the scalar ξ in equation (1) is replaced with a parametric function $\xi(x_i, \beta)$and τ is set to equal 0.5, one could obtain the estimate of the conditional median function.

$$min_{\xi \in \mathbb{R}} \sum_{i=0}^{n} \rho_\tau(y_i - \xi(x_i, \beta)) \tag{2}$$

Setting τ to different values will lead to the estimation of different conditional quantile functions.

Figure 1

Tilted absolute value function

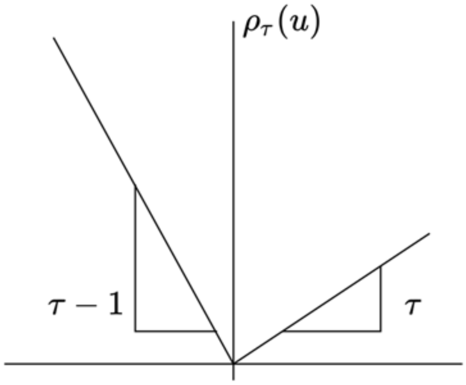

Source: Author.

Applying this idea to modelling the conditional quantiles of GDP growth (similarly to Adrian et al., 2019), we would model the relation between the conditional quantile of $y_{t+h}$and a vector of predictors X and optionally their lags, for a given time period t and a forecasting horizon h. In order to estimate the quantile regression of y_{t+h}on X, the regression coefficients β_τ for a given τ is chosen to minimize the weighted absolute value of errors:

$$\hat{\beta}_\tau = \text{argmin}_{\beta_\tau \in \mathbb{R}^k} \sum_{t=1}^{T-h} (\tau \cdot 1_{(y_{t+h} \geq X\beta_\tau)} |y_{t+h} - X\beta_\tau| + (1-\tau) \cdot 1_{(y_{t+h} < X\beta_\tau)} |y_{t+h} - X\beta_\tau|) \qquad (3)$$

where $\mathbf{1}(\cdot)$ is the indicator function, which subsets negative and positive errors, and T is the total length of the time series. The output value from the model is the quantile of y_{t+h} conditional on the model input X:

$$\hat{Q}_{y_{t+h}|X}(\tau|\text{X}) = X\beta_\tau \qquad (4)$$

Rodriguez and Pereira (2020) propose a multi-output deep learning approach for modelling several conditional quantiles jointly to address the problem of crossing quantiles, which often occurs when quantiles are estimated separately. The authors' suggestion is to aggregate the loss function for the separate quantiles and evaluate it for all conditional quantiles at once at every step of the optimization process. Bondell et al. (2010) propose a further addition to the tilted absolute loss function, which deals with the notorious problem of the tendency of separately estimated conditional quantiles to cross. This additional penalty for crossing can be expressed in the following way:

$$\sum_{j=1}^{J-1} max(0, X\hat{\beta}_{\tau_j} - X\hat{\beta}_{\tau_{j+1}}) \qquad (5)$$

where J is the number of quantiles sorted by the increasing value of τ. This term can be added to the loss function described in equation (3). Both of these suggestions have been incorporated into the newly proposed procedure.

Quantiles of the conditional distribution of GDP growth can be expressed either as linear or non-linear functions of the observed predictors. This framework allows one to study the skewed and fat-tailed distribution of GDP growth documented by multiple authors in the past, like Fagiolo et al. (2008), Adrian et al. (2019) and De Santis and Van der Veken (2020). Also, the predictive power of different independent variables possibly exhibits heterogeneity across different quantiles of GDP growth, as was suggested by Giglio et al. (2016), and Adrian et al. (2019), among others.

After constructing the quantiles, one could fit a probability distribution function to them in order to generate a density forecast. Adrian et al. (2016) propose using a skewed t-distribution for this purpose. In order to estimate the four parameters related to the skewed t-distribution (following Wurtz et al., 2006), the problem can be formulated as a least squares optimization problem, using the estimated conditional quantiles[3] and the inverse cumulative probability function:

$$\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{v}_{t+h}, \hat{\alpha}_{t+h}\} = \text{argmin}_{\mu,\sigma,\nu,\alpha} \sum_{j=1}^{J} (\hat{Q}_{y_{t+h}|\text{X}}(\tau_j|X) - F^{-1}(\tau_j; \mu, \sigma, \nu, \alpha))^2 \quad (6)$$

where $\hat{\mu}_{t+h} \in \mathbb{R}$ (mean or location shift), $\hat{\sigma}_{t+h} \in \mathbb{R}^+$ (standard deviation or scaling parameter), $\hat{v}_{t+h} \in \mathbb{R}$ (skewness parameter), and $\hat{\alpha}_{t+h} \in \mathbb{R}^+$ (kurtosis or tailweight parameter). $F^{-1}$is the inverse cumulative distribution function and $\hat{Q}_{y_{t+h}|X}(\tau_j|X)$ is the estimated quantile of y_{t+h} for a given τ and conditional on X. This method can be used to estimate a density based on the conditional quantiles, as well as the unconditional or observed quantiles of the actual GDP growth.

The established procedure used by Adrian et al. (2019) and partly described in the previous paragraphs is illustrated on the top of the diagram. Its first step consists of a linear regression model with a loss function similar to equation 3, which is used to generate conditional quantiles. The second step uses the conditional quantiles as an input and performs a least squares optimization between the input and the inverse CDF of the distribution of choice (in this case, the skewed t-distribution) as in equation 6. Adrian et al. (2019) follow two alternative approaches to show that the results from the two-step procedure are robust across methods. They employ fully parametric and fully non-parametric approaches to the same task and find very similar characteristics of the resulting conditional distributions. However, the authors use a fully parametric specification, which explicitly indicates the relationship between the conditional mean and the conditional variance and skewness.[4] In this sense, the

[3] The .05, .25, .75 and .95 quantiles are used for the estimation of the conditional distribution.

[4] In preliminary experiments for this study it was attempted to model the conditional distribution features without such explicit parametrization, but unfortunately such models fail to converge towards a loss minimization and often exhibited exploding gradients.

two-step procedure is less parametric, less hard-coded and perhaps superior in terms of its flexibility.

A departure from this procedure can be seen in the lower segment of Figure 2 – a combined approach that unites the two steps. This method would directly estimate the parameters of the distribution similarly to probabilistic regressions (for a deep learning application, see Salinas et al., 2020), but retaining the internal consistency of the original approach. The proposed procedure relies on an intermediate estimation of conditional quantiles in order to estimate the parameters of the conditional distribution. For this purpose, the model uses two separate loss functions – an intermediate tilted absolute loss function (as in equation 3 in combination with equation 5) and a final least squares function (as in equation 6).

Figure 2

Comparison between the two approaches for estimating the conditional distribution

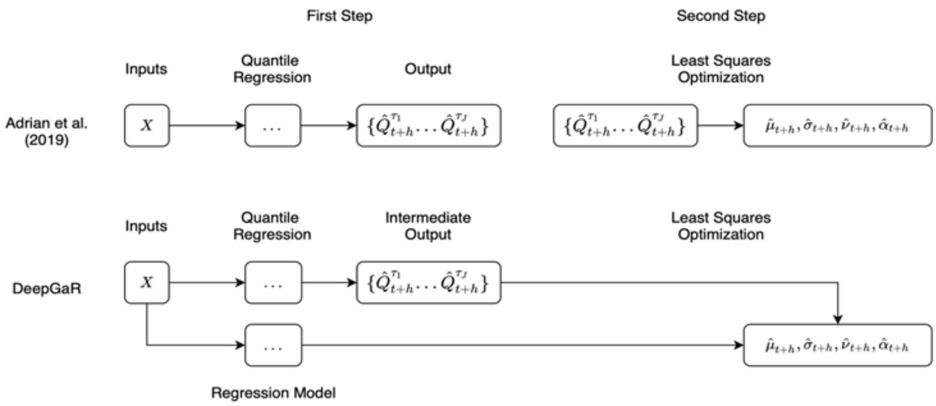

Source: Author.

Figure 3 presents an overview of the architecture used for the artificial neural network (ANN) model. Since the network is built to model time-series – a temporal sequence problem – it is more precise to call it a recurrent neural network. The first dimension of the various layers is the batch size used for the model training[5], but here it is kept equal to 1, which means that the training is performed on the whole dataset at once. Therefore, this dimension is omitted from the diagram. The length of the input time-series is denoted by *t*, while *m* denotes the number of predictor variables. *J* denotes the number of quantiles, which in this exercise equals to 4 and *h* denotes the time-horizon of the forecast in quarters, which is set to 1.

The model has multiple branches. The bottom branch starts out with a long short-term memory layer (LSTM). Each quantile is generated in a separate sub-branch and therefore, there are *J* number of sub-branches for all quantiles. There is another branch (on top) which contains an LSTM layer and leads to the final output of the model. The final output is the result of the quantile branches and the top branch being combined through the use of the least

[5] In the machine learning terminology, model training is similar to model estimation as is understood in the realm of econometrics. Different values of the batch size can be applied to optimize the training process, especially if one works with larger datasets.

squares loss function (equation 6). The total number of trainable parameters within the model is in the range of 150 and 200 thousand, depending on the number of input predictors.

Figure 3

Proposed neural network architecture

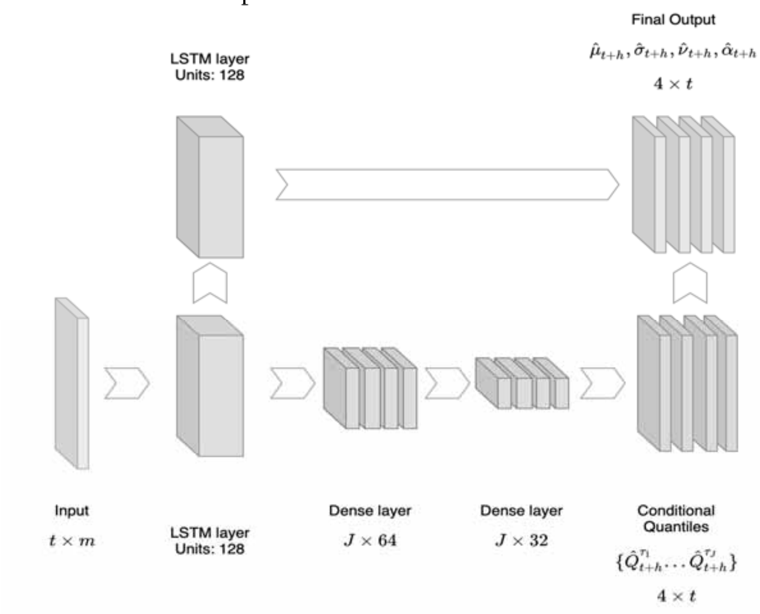

Source: Author.

The proposed architecture is reliant, to a great extent, on the LSTM layers. It is outside of the scope of this paper to explain why these layers are well-suited for time-series modelling and how they handle the task, but a good overview of the current state and implementations of these models can be found in Hewamalage, Bergmeir and Bandara (2021). Another good overview of different artificial neural network architectures in the context of economic forecasting is done by Cook and Hall (2017). The proposed model was implemented using the Tensorflow library and the ADAM optimizer was used for the model training (Kingma, Ba, 2014). The Tensorflow library uses automatic differentiation and gradient descent through time for recurrent neural networks (Rumelhart, Hinton, Williams, 1986; Williams, Zipser, 1992).

Another departure from the established procedure lies in the choice of the family of distributions, used to model the risks to growth. Instead of the skewed t-distribution an alternative distribution was used – the Sinh-Arcsinh distribution introduced by Jones and Pewsey (2009). This is a four-parameter distribution, which can account for location, scale, skewness and tailweight and is a generalization of the normal distribution. The reason for choosing it was its convenient implementation in Tensorflow and similar properties to the skewed t-distribution.

It was tested if there would be a statistically significant difference in using this family of distribution compared to the skewed t-distribution when fitting them to the GDP q-o-q data for the selected countries. The two types of distributions were fitted to the data using the unconditional quantiles and the inverse CDF functions, and then the Kolmogorov-Smirnov two-sample test was performed on randomly generated samples consisting of 1000 observations from each of the estimated distributions. The test results[6] failed to reject the null hypothesis of no differences for all four countries.

Figure 4

Comparison between the Sinh-Arcsinh distribution and skewed t-distribution

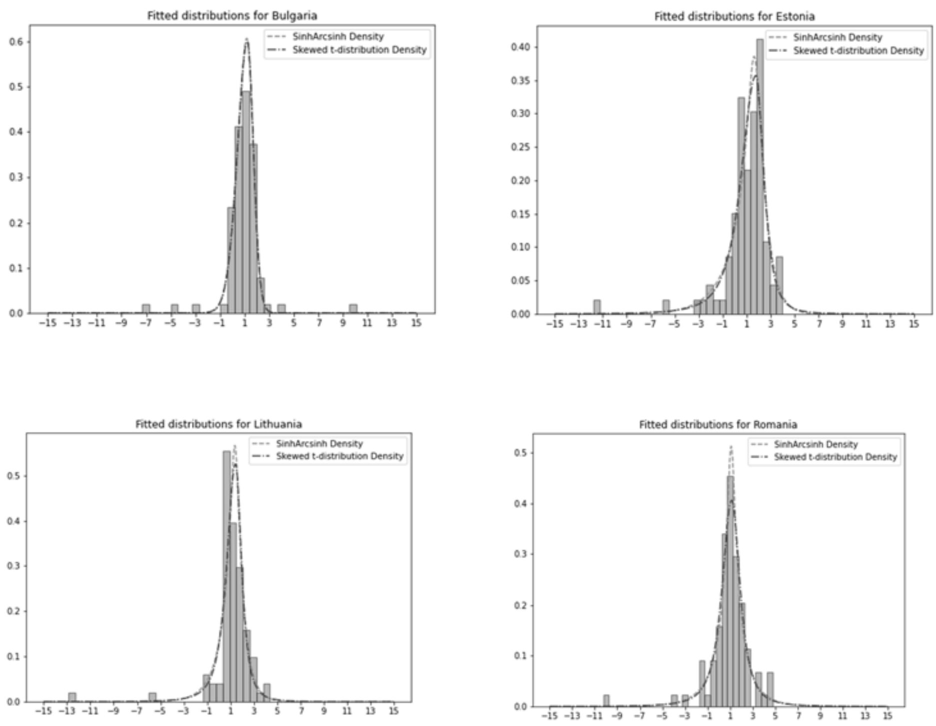

Source: Author, Eurostat.

The final output dense layers of the model architecture (marked as final output in Figure 4) uses specific parameterization in order to ensure that the scale and tailweight parameters are positive numbers, which is a prerequisite for the implementation of the Sinh-arcsinh distribution. The parametrization is the exponential linear unit plus 1 in order to ensure non-negativity.

[6] KS test results: Bulgaria (KS statistic: 0.054, p-value: 0.108), Estonia (KS statistic: 0.028, p-value: 0.828), Lithuania (KS statistic: 0.038, p-value: 0.466), Romania (KS statistic: 0.053, p-value: 0.121).

$$\begin{cases} x + 1 & x \geq 0 \\ -(e^{x} - 1) + 1 & x < 0 \end{cases} \qquad (7)$$

This parametrization is necessary and ensures convergence of the optimization algorithms, as well as the ability to generate a conditional distribution of the Sinh-Arcsinh family.

Finally, in order to combine two loss functions into the model, a type of dynamic weighting was implemented. Different versions of the weighting were tested, but the final version relied on a high weight (95%) for the tilted absolute value function for 90% of the training epochs (model training duration) and low weight (5%) for the remainder of the training. Respectively, the weight for the second least squares function, which defines the conditional distribution parameters, stays low (5%) during 90% of the duration of the training and switches to a high value (95%) for the last 10% of the duration of the training. This scheme yielded the best results, since it took more epochs for the tilted loss function to reach an optimal level of loss, compared to the least squares loss function, which usually reached optimum very quickly. The model results were robust to changes in the various parameters of the weighting.

3. Data

One of the main challenges in forecasting GDP growth is the time lag associated with the quarterly releases of the indicator. For most EU countries, the GDP flash release for a specific quarter would be published around the middle of the subsequent quarter. If one would like to use the first lag of GDP (in terms of quarterly frequency), one is already within the time frame of real-time forecasting or nowcasting. However, this time lag also presents the opportunity to use the current values of short-term indicators, which are released at a higher frequency for the purposes of real-time forecasting. Moreover, since short-term indicators are released on a monthly or daily basis, one could use intermediate values for a nowcasting exercise (e.g. use the average of the first two months of the quarter to get an earlier input to the model).

For this analysis, the target variable of interest is the quarterly growth rate of the seasonally and calendar-adjusted chain-linked volumes of GDP. The available final release of the GDP data is used as of the writing of this article. Apart from the lags of the target variable, a list of leading indicators of financial conditions and economic activity was compiled in order to be used as candidate predictors. The choice of leading indicators was following an approach similar to Adrian et al. (2019), De Santis and Van der Veken (2020), Figuerez and Jarociński (2020) and Prasad et al. (2019). It was imperative that they are available for a longer time frame and an initial choice for a starting year of the samples was the year 2000 as this ensured a long enough training sample and the opportunity to put aside a test sample. Currently, there are a lot of interesting leading indicators which can be used for similar macroeconomic forecasting tasks, but their main disadvantage is the lack of accumulated historical data. Moreover, it was decided to include only indicators, which are available for a specific quarter by the end of the same quarter, in order to be able to use the current values of the predictors in time reference to the GDP growth values, which are released later on. Therefore, short-term indicators which are released with a significant delay were not included in the modelling data set, despite their relevance, because they have limited use in the nowcasting of GDP growth. A full list of predictors can be found in (Table 1).

Table 1

Selected Leading Indicators

Indicator	Description and Source
Country-specific sentiment indicator	Economic sentiment index collected by the Directorate-General for Economic and Financial Affairs (DG ECFIN) of the European Commission, Monthly frequency, Eurostat
Country-specific 10-year Government Bond Yield	Monthly average of daily frequency, Investing.com
Country-specific Stock Price Index Close Value (SOFIX, BET, OMXTGI, OMXVGI)	Monthly frequency, Investing.com
Dow Jones Composite Index Close Value	Monthly frequency, Yahoo Finance
Dow Jones Composite Index Volume	Monthly frequency, Yahoo Finance
DAX Index Close Value	Monthly frequency, Yahoo Finance
DAX Index Volume	Monthly frequency, Yahoo Finance
10-year US Government Bond Yield	Monthly average of daily frequency, Nasdaq Data Link
PMI Composite – Euro Area	Monthly frequency, Nasdaq Data Link
Brent Oil Price	Monthly average of daily frequency, Thomson Reuters

Source: Author.

It is important to make several clarifications regarding the country-specific financial indicators used. Estonia emitted 10-year bond yields for the first time in June 2020, and therefore the indicator was not used due to the very short length of the time series. Unfortunately, no suitable substitute was found for Estonia, as it seems that the government historically emitted only short-term government debt. For each of the other countries, the indicators were available for varying periods: October 2002 for Bulgaria, February 2003 for Lithuania, and September 2007 for Romania. Using these indicators necessitates shortening the sample and, in the case of Romania, significantly. With respect to the country-specific stock price indices, the choice was based on maximum data availability: SOFIX for Bulgaria, BET for Romania, OMXTGI for Estonia and OMXVGI for Lithuania. For each country, the indicators were available for varying periods: September 2003 for Bulgaria, May 2013 for Lithuania, and June 2002 for Romania. Due to the very short history for Lithuania, it was decided to exclude the index from the shortlist of indicators. For Estonia, the OMXTGI index has been available since before 2000.

Initially, the focus of the study was on the macroeconomic developments in Bulgaria, but after preliminary results were generated, it was decided to test the proposed approach on three other small open European economies relatively similar in terms of the size and structure of the economy. After performing a backward selection for each country individually, three leading indicators emerged as suitable candidates to be used in the forecasting models. The backward selection was based on the performance indicators discussed in the next section.

The study does not claim to use an exhaustive list of factors which influence growth. The link between the country-specific sentiment indicator and GDP could be thought to carry some degree of causality. On the other hand, the global financial indicators are exogenous in nature in reference to the countries of interest and can be considered leading indicators of financial stress. The country-specific financial variables are also leading indicators of financial stress and although they are more relevant to the specific economy, there is no justification to claim any causal relationship with respect to economic growth.

All variables were normalized with a mean 0 and standard deviation of 1 prior to use in the model. The total sample covers the period 2000Q1 to 2021Q4 and was divided into a training and a testing sample. The first 64 quarters were used for model training (2000Q1 to 2016Q2) and the last 22 quarters were used for validating the model performance (2016Q3 to 2021Q4). A rolling window approach was followed to construct the training and test samples. The final specifications are depicted in a stylized way in Table 2.

Table 2

Final Model Specifications

Country	Specifications
Bulgaria	$GDP_t = f(GDP_{t-1}, SENTIMENT_t, SOFIX_t, US\ BOND\ YIELD_t)$
Estonia	$GDP_t = f(GDP_{t-1}, SENTIMENT_t, OMXTGI_t)$
Lithuania	$GDP_t = f(GDP_{t-1}, SENTIMENT_t, LT\ BOND\ YIELD_t)$
Romania	$GDP_t = f(GDP_{t-1}, SENTIMENT_t, US\ BOND\ YIELD_t)$

Source: Author.

In the listed specifications, the left-hand side describes the target variable and the right-hand side the set of predictors. For Bulgaria and Estonia, the inclusion of the domestic stock price indices leads to optimal performance. In the case of Bulgaria, the inclusion of the US 10-year government bond yield carried additional predictive power. Similarly, for Lithuania, using the Lithuanian 10-year government bond yield results in the best model performance. For Romania, the use of the US 10-year bond yield leads to better performance compared to using domestic indicators.

With the given specifications and after all data transformations, the model training sets start at 2000Q3 for Estonia and Romania, at 2003Q4 for Bulgaria and at 2003Q2 for Lithuania. The test sets used for the estimation of the out-of-sample performance start at 2016Q4 for Estonia and Romania and 2017Q3 for Bulgaria and Lithuania. These sample lengths are a result of using 75% of the total sample size as a training sample size.

4. Results

The DeepGaR model was tested against a benchmark based on the established two-step approach, with the exception of two differences. The benchmark uses a linear quantile regression, but the conditional quantiles are generated simultaneously and use the crossing loss, similarly to the DeepGaR model. The skewed t-distribution is used for the generation of the conditional distribution of the benchmark. The performance of the two models is based on the out-of-sample performance over the test sample covering the pandemic lockdown recession occurring in 2020Q1 and/or 2020Q2. The main indicators which were used to measure and compare the performance were RMSE between actual GDP and the conditional median, the predictive score, which is the value of the PDF of the conditional distribution at the value of the realized GDP. Additional indicators were recorded to study the characteristics of the generated distributions. Detailed results can be found in Appendix tables 3 to 10.

Figure 5 presents the .05, .5 and .95 conditional quantiles generated by the DeepGaR model and the realized GDP for each of the four counties. Figure 6 is similar, but presents the conditional quantiles generated using the benchmark model.

Figure 5

Out-of-sample conditional distributions generated by the DeepGaR model

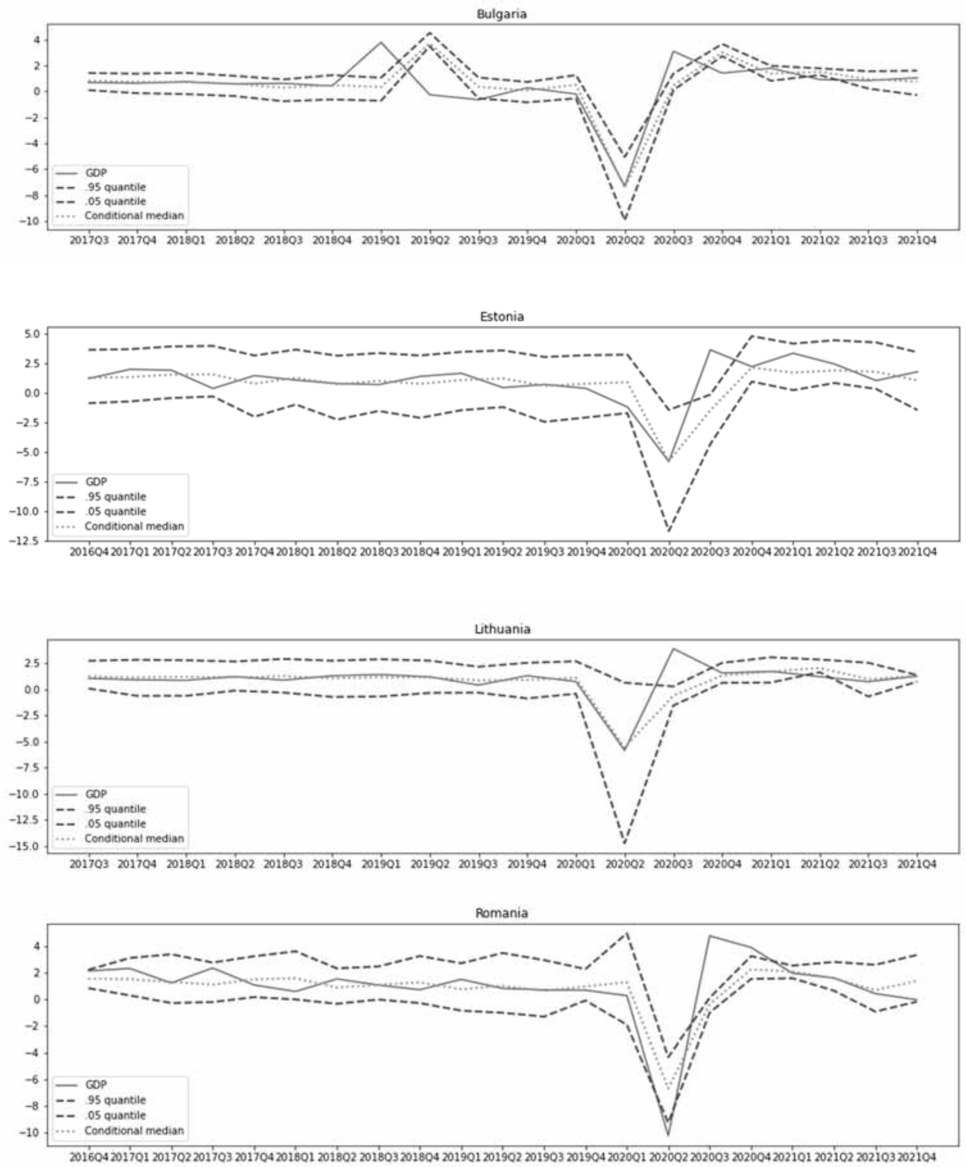

Source: Author, Eurostat.

Figure 6

Out-of-sample conditional distributions generated by the benchmark

Source: Author, Eurostat.

It is evident that the DeepGaR model achieves superior accuracy compared to the benchmark, during the negative growth period which every country experiences between 2020Q1 and 2020Q2. The model also recognizes the downside risks reliably, given the increased spread of the distribution during periods of negative growth. The benchmark fails at recognizing downside risks across all countries. However, it seems to perform satisfactory during upturn periods. Both models perform poorly when it comes to forecasting the recovery after the initial slump and at recognizing upside risks. Both models perform well during non-recession periods, but the DeepGaR model is better in modelling the recession caused by the unexpected shock of the global pandemic lockdown. By utilizing a high number of LSTM units the DeepGaR model manages to recognize a highly non-linear relationship between the predictors and the target. During part of the initial experimentation, it was observed that reducing the number of LSTM units reduces its performance during the recession periods.

The DeepGaR model consistently produces conditional distributions, which confirm some of the findings of Adrian et al. (2019). Both symmetric conditional distributions during expansions and negative skewness during periods of recession can be observed. Additionally, a negative correlation between the conditional mean and variance of the growth distribution is evident as well. These results were confirmed for all four of the countries in the sample, both in-sample (on the training set) as well as out-of-sample (on the test set).

5. Discussion

The COVID-19 pandemic and the recessions many countries experienced due to implemented lockdowns posed an unprecedented challenge to decision-makers and forecasters. Both private enterprises and government institutions had to adapt to this shock quickly and implement policies to tackle the consequences, based on limited foresight. While, it is virtually impossible to anticipate such an event like the coronavirus pandemic and its consequences ahead of time, one could forecast or nowcast its effects on the economy through leading indicators, which could help the decision-making process.

The current study demonstrates that a parsimonious model using country-specific sentiment indicators as well country-specific and global financial variables can successfully nowcast recessions caused by unexpected shocks like the coronavirus pandemic. The comparative performance of the artificial neural network DeepGaR model proves that it is a useful tool in modeling macroeconomic risks related to the 2020 coronavirus pandemic lockdown in four small open economies in Europe. Its ability to model highly non-linear relationships makes it superior to a linear benchmark in this case.

For Bulgaria and Lithuania, the DeepGaR model manages to predict very accurately the negative growth of GDP in 2020Q2, when the strongest economic effects of the lockdowns were felt. In the case of Estonia, the DeepGaR model does not accurately predict the start of the recession in 2020Q1, but manages to predict very accurately the negative growth in 2020Q2. However, it is not clear whether the growth dynamics in this quarter are not a result of seasonal adjustment. For Romania, the DeepGaR model fails to predict the full extent of the lockdown recession in 2020Q2, but still outperforms significantly the linear benchmark. Apart from its disadvantage in much lower prediction accuracy with respect to predicting the pandemic crisis, the linear benchmark achieves satisfactory performance in nowcasting growth during upturn periods.

A disadvantage shared by both the proposed DeepGaR model as well as the linear benchmark is their limited ability to predict the upturn after the initial decline in economic growth. This result is observed across all countries and at first glance, the problem is with the so-called shape of the 2020 recession, which in all countries of interested seems to have a V-shape. The models are trained on the recessions caused by the global financial crisis, which had either a U or a W-shape for the countries of interest, which might be why the models fail to anticipate a quick and strong recovery after only a quarter or two of decline in growth.

With respect to the indicators used across the four countries of interest, it is evident that both country-specific and global factors of financial stress carry predictive power with respect to economic growth and specifically in the task of predicting the pandemic lockdown recession. For Bulgaria and Estonia, it was shown that the use of domestic stock price indices' close values leads to optimal results. In the case of Bulgaria, the inclusion of the US 10-year government bond yield carried additional predictive power. For Lithuania, the inclusion of the Lithuanian 10-year government bond yield resulted in the best-performing model specification. For Romania, the US 10-year bond yield carried more predictive power with respect to predicting the pandemic recession, compared to the country-specific financial indicators. Additionally, across all four countries, it is demonstrated that a parsimonious model containing few indicators yields optimal performance.

The DeepGaR model combines a couple of recent improvements proposed by researchers working on quantile regression models, which allows it to mitigate known problems like the crossing of the quantiles. The first improvement is the simultaneous estimation of quantiles, which allows one to estimate an arbitrary number of quantiles within one estimation procedure and using a single loss function. This both speeds up the process of generating the conditional quantiles, but also is shown to alleviate the crossing problem. Indeed, the only datapoint where this problem occurs can be found in the benchmark results for Lithuania in 2021Q4. The second improvement is the explicit inclusion of a crossing loss term within the loss function, which additionally mitigates the issue. Moreover, combining the two steps of the original procedure into a single model creates one internally consistent procedure without sacrificing its flexibility. Working in the context of an artificial neural network allows one to construct a custom model with two loss functions and combine the two steps of the original estimation procedure into a one-step procedure. Combining the two steps leads to an internally consistent procedure without sacrificing the flexibility of the original approach. This new model representation is the main novel contribution of the current study.

Additionally, the current analysis confirms that there is a negative correlation between the conditional mean and variance of the distribution of growth as well as symmetric conditional distributions during expansions and negative skewness during periods of recession for the four economies analyzed in this study, similarly to the stylized facts Adrian et al. (2019) identified for the US and which were also confirmed for the Euro Area by Figuerez and Jarociński (2020).

Future work on the subject can focus on both the applications of the model and its design. It would be interesting to study the performance of the model using data from larger economies.

Both the proposed and original procedures struggle with anticipating upside risks and rapid recoveries as the one seen after the initial recession episode in 2020 across many countries, which is certainly a disadvantage which requires further attention. Lastly, despite using a complex neural network architecture, one could use Shapley values to perform a sensitivity analysis between the inputs and output of the model. This would be a useful addition to the procedure as it would provide additional transparency into how the inputs affect each of the parameters of the generated conditional distribution.

References

Adrian, T., Boyarchenko, N., Giannone, D. (2019). Vulnerable Growth. – American Economic Review, 109, pp. 1263-1289. https://doi.org/10.1257/aer.20161923.

Altunbas, Y., Manganelli, S., Marques-Ibanez, D. (2017). Realized bank risk during the great recession. – Journal of Financial Intermediation, Vol. 32. https://doi.org/10.1016/j.jfi.2017.08.001.

An, Z., Jalles, J., Loungani, P. (2018). How Well Do Economists Forecast Recessions?. – IMF Working Paper 18/39.

Antipova, T. (2021). Coronavirus Pandemic as Black Swan Event. – In: Antipova, T. (ed.). Integrated Science in Digital Age 2020. ICIS 2020. Lecture Notes in Networks and Systems, Vol 136. Springer, Cham. https://doi.org/10.1007/978-3-030-49264-9_32.

Bondell, H. D., Reich, B. J., Wang, H. (2010). Noncrossing quantile regression curve estimation. – Biometrika, 97, pp. 825-838. https://doi.org/10.1093/biomet/asq048.

Brownlees, C., Souza, A. B. M. (2021), Backtesting global Growth-at-Risk. – Journal of Monetary Economics, Vol. 118, pp. 312-330, https://doi.org/10.1016/j.jmoneco.2020.11.003.

Carriero, A., Clark, T., Marcellino, M. G. (2020). Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions. – FRB of Cleveland Working Paper No. 20-02R.

Chen, W., Mrkaic, M., Nabar, M. (2019). The Global Economic Recovery 10 Years After the 2008 Financial Crisis. – IMF Working Paper 19/83.

Clements, M. P. (2004). Evaluating the Bank of England Density Forecasts of Inflation. – The Economic Journal, 114, pp. 844-866. https://doi.org/10.1111/j.1468-0297.2004.00246.x.

Cook, T., Smalter Hall, A. (2017). Macroeconomic Indicator Forecasting with Deep Neural Networks. RWP. https://doi.org/10.18651/RWP2017-11.

De Santis, R. A., Van der Veken, W. (2020). Forecasting macroeconomic risk in real time: great and Covid 19 recessions. – European Central Bank Working Paper Series No 2436 / July 2020.

Diebold, F. X., Gunther, T. A., Tay, A. S. (1998). Evaluating Density Forecasts with Applications to Financial Risk Management. – International Economic Review, 39, p. 863. https://doi.org/10.2307/2527342.

Fagiolo, G., Napoletano, M., Roventini, A. (2008). Are Output Growth- Rate Distributions Fat-Tailed? Some Evidence From OECD Countries. – Journal of Applied Econometrics, Vol. 23, pp. 639–669. https://doi.org/10.1002/jae.1003.

Figueres, J. M., Jarociński, M. (2020) Vulnerable growth in the euro area: Measuring the financial conditions. – Economics Letters, Vol. 191, 109126, ISSN 0165-1765, https://doi.org/10.1016/j.econlet.2020.109126.

Giglio, S., Kelly, B., Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. – Journal of Financial Economics, Elsevier, Vol. 119(3), pp. 457-471. https://doi.org/10.1016/j.jfineco.2016.01.010.

Gneiting, T. (2008). Editorial: Probabilistic forecasting. – Journal of the Royal Statistical Society Series A, Royal Statistical Society, Vol. 171(2), pp. 319-321, April. https://doi.org/10.1111/j.1467-985X.2007.00522.x.

Hewamalage, H., Bergmeir, C., Bandara, K. (2021). Recurrent Neural Networks for Time Series Forecasting: Current status and future directions. – International Journal of Forecasting, 37, N 1, pp. 388-427. https://doi.org/10.1016/j.ijforecast.2020.06.008.

Hochreiter, S., Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, Vol. 9, N 8, pp. 1735-1780, 15 Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

Iseringhausen, M. (2021). A time-varying skewness model for Growth-at-Risk. – European Stability Mechanism Working Paper, 49/2021.

Jones, M. C., Pewsey, A. (2009). Sinh-arcsinh distributions. – Biometrika, 96(4), pp. 761-780. http://www.jstor.org/stable/27798865.

Kingma, D. P., Ba, J. (2017). Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs].

Koenker, R., Bassett, G. (1978). Regression Quantiles. – Econometrica, 46, p. 33. https://doi.org/10.2307/1913643.

Lewis, C. Pain, N. (2014). Lessons from OECD forecasts during and after the financial crisis. – OECD Jounal: Economic Studies, Vol. 15. https://doi.org/10.1787/19952856.

Makridakis, S., Hogarth, R. M., Gaba, A. (2009). Forecasting and uncertainty in the economic and business world. – International Journal of Forecasting, 25, pp. 794-812. https://doi.org/10.1016/j.ijforecast.2009.05.012.

Moré, J. J. (1978). The Levenberg-Marquardt algorithm: Implementation and theory. – In: Watson, G. A. (ed.). Numerical Analysis, Lecture Notes in Mathematics. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 105–116. https://doi.org/10.1007/BFb0067700.

Prasad, A., Elekdag, S., Jeasakul, P., Lafarguette, R., Alter, A., Feng, A. X., Wang C. (2019). Growth at Risk: Concept and Application in IMF Country Surveillance. – IMF Working Paper 19/36.
Rodrigues, F., Pereira, F. (2020). Beyond Expectation: Deep Joint Mean and Quantile Regression for Spatiotemporal Problems. – IEEE Transactions on Neural Networks and Learning Systems, 31, pp. 5377-5389. https://doi.org/10.1109/TNNLS.2020.2966745.
Rumelhart, D. E., Hinton, G. E., Williams, R. J. (1986), Learning representations by back-propagating errors. – Nature, 323, pp. 533-536. https://doi.org/10.1038/323533a0.
Salinas, D., Flunkert, V., Gasthaus, J. Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. – International Journal of Forecasting, Vol. 36, N 3, pp. 1181-1191, https://doi.org/10.1016/j.ijforecast.2019.07.001.
Williams, R. J., Zipser, D. (1995). Gradient-based learning algorithms for recurrent networks and their computational complexity. – In: Chauvin, Y., Rumelhart, D. E. (eds.). Backpropagation: Theory, architectures, and applications, Lawrence Erlbaum Associates, Inc., pp. 433-486.
Wurtz, D., Chalabi, Y., Luksan, L. (2006). Parameter estimation of arma models with garch/aparch errors. An r and splus software implementation. – Journal of Statistical Software.

APPENDIX

This appendix contains the tables with the detailed results which are referenced in the results and discussion sections.

Table 3

Performance Comparison for Bulgaria

		Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR
Time Period	GDP Actual	Median		RMSE		Predictive Score		Probability of recession	
2017Q3	0.685	0.768	0.817	0.083	0.132	0.837	1.196	0.501	0.436
2017Q4	0.628	0.691	0.727	0.063	0.098	0.781	1.133	0.556	0.551
2018Q1	0.743	0.823	0.697	0.080	0.046	0.807	1.207	0.460	0.577
2018Q2	0.572	0.681	0.584	0.108	0.012	0.751	1.211	0.560	0.705
2018Q3	0.615	0.634	0.283	0.019	0.333	0.757	0.827	0.589	0.899
2018Q4	0.416	0.677	0.484	0.260	0.068	0.661	0.893	0.563	0.739
2019Q1	3.790	0.595	0.336	3.195	3.454	0.000	0.000	0.617	0.848
2019Q2	-0.263	2.156	3.711	2.419	3.974	0.001	0.000	0.013	0.000
2019Q3	-0.642	0.275	0.354	0.917	0.997	0.215	0.121	0.802	0.853
2019Q4	0.270	0.043	0.070	0.227	0.200	0.636	1.063	0.858	0.960
2020Q1	-0.206	0.617	0.503	0.823	0.709	0.268	0.284	0.622	0.741
2020Q2	-7.324	0.302	-7.453	7.625	0.130	0.000	0.237	0.707	1.000
2020Q3	3.095	-3.024	0.460	6.119	2.635	0.000	0.003	1.000	0.798
2020Q4	1.407	1.842	3.016	0.435	1.609	0.556	0.000	0.035	0.000
2021Q1	1.786	1.264	1.366	0.522	0.421	0.527	0.415	0.161	0.042
2021Q2	0.916	1.210	1.502	0.294	0.587	0.829	0.036	0.146	0.002
2021Q3	0.837	0.905	0.941	0.067	0.104	0.851	1.295	0.399	0.296
2021Q4	1.045	0.988	0.746	0.057	0.299	0.867	0.757	0.349	0.520

Source: Author, Eurostat.

Table 4

Performance Comparison for Estonia

		Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR
Time Period	GDP Actual	Median		RMSE		Predictive Score		Probability of recession	
2016Q4	1.198	0.906	1.261	0.293	0.063	0.646	0.727	0.482	0.363
2017Q1	1.972	0.930	1.317	1.043	0.655	0.458	0.564	0.474	0.340
2017Q2	1.899	0.967	1.526	0.931	0.372	0.449	0.668	0.466	0.270
2017Q3	0.358	1.479	1.544	1.121	1.186	0.481	0.400	0.267	0.255
2017Q4	1.446	0.541	0.761	0.905	0.685	0.490	0.540	0.591	0.524
2018Q1	1.062	1.158	1.266	0.096	0.205	0.733	0.681	0.394	0.368
2018Q2	0.776	0.780	0.713	0.004	0.063	0.634	0.551	0.520	0.535
2018Q3	0.685	0.564	0.995	0.121	0.310	0.603	0.605	0.580	0.457
2018Q4	1.377	0.551	0.752	0.825	0.624	0.487	0.544	0.583	0.526
2019Q1	1.637	1.206	1.073	0.431	0.564	0.671	0.572	0.376	0.435
2019Q2	0.426	1.088	1.204	0.662	0.778	0.653	0.516	0.424	0.394
2019Q3	0.697	0.765	0.581	0.067	0.116	0.706	0.543	0.528	0.568
2019Q4	0.365	0.890	0.759	0.525	0.395	0.674	0.526	0.486	0.524
2020Q1	-1.207	0.173	0.882	1.380	2.089	0.413	0.196	0.664	0.490
2020Q2	-5.812	1.503	-5.795	7.314	0.017	0.000	0.225	0.162	1.000
2020Q3	3.619	-2.685	-1.490	6.304	5.109	0.002	0.000	0.972	0.997
2020Q4	2.197	1.386	2.123	0.811	0.074	0.443	0.938	0.351	0.040
2021Q1	3.331	1.597	1.693	1.734	1.638	0.222	0.228	0.213	0.161
2021Q2	2.419	1.242	1.884	1.177	0.535	0.376	0.591	0.391	0.055
2021Q3	1.025	1.751	1.762	0.725	0.737	0.741	0.687	0.155	0.139
2021Q4	1.766	1.108	1.044	0.658	0.723	0.622	0.537	0.408	0.443

Source: Author, Eurostat.

Table 5

Performance Comparison for Lithuania

		Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR
Time Period	GDP Actual	Median		RMSE		Predictive Score		Probability of recession	
2017Q3	1.036	1.042	1.240	0.006	0.204	0.976	1.360	0.416	0.260
2017Q4	0.897	1.083	1.135	0.186	0.238	0.983	0.931	0.397	0.373
2018Q1	0.855	1.399	1.194	0.544	0.339	0.714	0.861	0.280	0.345
2018Q2	1.205	0.915	1.139	0.291	0.066	0.818	1.292	0.475	0.330
2018Q3	0.867	1.136	1.257	0.269	0.390	0.964	0.896	0.374	0.305
2018Q4	1.284	1.004	1.054	0.280	0.230	0.859	0.956	0.434	0.409
2019Q1	1.403	1.069	1.144	0.334	0.259	0.783	0.915	0.408	0.374
2019Q2	1.181	0.954	1.139	0.227	0.042	0.743	1.150	0.462	0.352
2019Q3	0.397	0.698	0.837	0.301	0.440	1.061	0.964	0.565	0.514
2019Q4	1.290	0.870	0.901	0.420	0.389	0.813	0.873	0.495	0.480
2020Q1	0.732	0.923	1.087	0.191	0.356	0.882	0.946	0.475	0.380
2020Q2	-5.864	1.505	-5.604	7.370	0.260	0.001	0.152	0.163	0.959
2020Q3	3.882	-0.785	-0.612	4.667	4.494	0.000	0.000	1.000	0.991
2020Q4	1.555	1.157	1.299	0.398	0.256	0.611	1.242	0.372	0.126
2021Q1	1.708	1.566	1.682	0.141	0.026	1.009	1.709	0.239	0.082
2021Q2	1.202	1.340	2.040	0.137	0.838	1.132	0.068	0.259	0.002
2021Q3	0.720	0.852	0.959	0.132	0.238	0.919	0.980	0.502	0.449
2021Q4	1.254	1.087	1.273	0.167	0.020	0.294	8.498	0.467	0.072

Source: Author, Eurostat.

Table 6

Performance Comparison for Romania

		Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR	Benchmark	DeepGaR
Time Period	GDP Actual	Median		RMSE		Predictive Score		Probability of recession	
2016Q4	2.119	0.816	1.535	1.303	0.583	0.229	0.368	0.548	0.070
2017Q1	2.304	1.093	1.495	1.211	0.809	0.282	0.379	0.430	0.203
2017Q2	1.213	1.156	1.295	0.057	0.082	0.583	0.611	0.405	0.350
2017Q3	2.341	0.746	1.090	1.595	1.251	0.180	0.227	0.569	0.412
2017Q4	1.062	1.187	1.485	0.125	0.423	0.632	0.680	0.389	0.231
2018Q1	0.571	1.464	1.569	0.893	0.998	0.385	0.358	0.218	0.246
2018Q2	1.517	0.522	0.868	0.995	0.649	0.353	0.489	0.663	0.529
2018Q3	1.054	0.659	1.074	0.395	0.019	0.504	0.969	0.596	0.401
2018Q4	0.704	0.980	1.255	0.277	0.552	0.707	0.554	0.472	0.360
2019Q1	1.483	0.650	0.738	0.833	0.745	0.410	0.411	0.611	0.569
2019Q2	0.803	1.104	1.004	0.302	0.201	0.704	0.460	0.416	0.475
2019Q3	0.693	0.782	0.626	0.089	0.066	0.677	0.466	0.559	0.588
2019Q4	0.668	0.458	0.956	0.210	0.287	0.573	0.908	0.671	0.475
2020Q1	0.258	2.479	1.264	2.221	1.006	0.001	0.266	0.003	0.437
2020Q2	-10.249	-0.378	-6.694	9.871	3.554	0.000	0.022	0.948	1.000
2020Q3	4.764	-3.001	-0.386	7.765	5.150	0.000	0.000	0.999	0.999
2020Q4	3.888	1.913	2.236	1.976	1.652	0.138	0.040	0.121	0.006
2021Q1	1.946	1.727	2.074	0.219	0.128	0.492	1.960	0.196	0.004
2021Q2	1.594	1.017	1.560	0.577	0.034	0.395	1.194	0.468	0.112
2021Q3	0.408	0.757	0.672	0.349	0.263	0.638	0.594	0.562	0.594
2021Q4	-0.053	1.136	1.373	1.188	1.425	0.219	0.194	0.382	0.312

Source: Author, Eurostat.

Table 7

Conditional Distribution Parameters Generated by DeepGaR – Bulgaria

Time Period	Standard Deviation	Skewness	Tailweight	.05 Quantile	.95 Quantile	IQR
2017Q3	0.384	-0.082	1.582	0.090	1.414	1.325
2017Q4	0.425	-0.139	1.511	-0.141	1.363	1.504
2018Q1	0.468	-0.098	1.484	-0.212	1.431	1.643
2018Q2	0.426	-0.214	1.447	-0.366	1.194	1.560
2018Q3	0.443	-0.276	1.331	-0.764	0.920	1.684
2018Q4	0.507	-0.207	1.335	-0.626	1.247	1.873
2019Q1	0.479	-0.219	1.327	-0.726	1.052	1.778
2019Q2	0.260	0.340	2.173	3.480	4.519	1.039
2019Q3	0.452	-0.111	1.443	-0.538	1.063	1.601
2019Q4	0.433	-0.173	1.394	-0.847	0.722	1.569
2020Q1	0.498	-0.166	1.428	-0.546	1.250	1.796
2020Q2	1.179	-2.026	0.134	-9.944	-5.073	4.871
2020Q3	0.328	0.321	1.826	0.100	1.370	1.270
2020Q4	0.271	0.210	2.080	2.714	3.654	0.940
2021Q1	0.345	0.033	1.725	0.814	1.970	1.156
2021Q2	0.166	0.027	2.197	1.258	1.774	0.515
2021Q3	0.385	-0.075	1.575	0.216	1.546	1.330
2021Q4	0.529	-0.101	1.375	-0.290	1.598	1.888

Source: Author.

Table 8

Conditional Distribution Parameters Generated by DeepGaR – Estonia

Time Period	Standard Deviation	Skewness	Tailweight	.05 Quantile	.95 Quantile	IQR
2016Q4	0.660	0.055	1.274	-0.884	3.617	4.501
2017Q1	0.647	0.079	1.296	-0.735	3.675	4.410
2017Q2	0.638	0.104	1.312	-0.450	3.906	4.356
2017Q3	0.621	0.141	1.327	-0.318	3.949	4.267
2017Q4	0.741	-0.112	1.143	-2.027	3.142	5.169
2018Q1	0.677	0.029	1.239	-0.995	3.635	4.630
2018Q2	0.761	-0.163	1.093	-2.268	3.122	5.391
2018Q3	0.710	-0.050	1.194	-1.544	3.347	4.891
2018Q4	0.750	-0.135	1.123	-2.123	3.143	5.266
2019Q1	0.712	-0.046	1.174	-1.472	3.447	4.919
2019Q2	0.696	-0.015	1.205	-1.214	3.567	4.781
2019Q3	0.771	-0.172	1.084	-2.457	3.016	5.473
2019Q4	0.750	-0.125	1.118	-2.094	3.165	5.259
2020Q1	0.714	-0.073	1.193	-1.722	3.211	4.932
2020Q2	1.159	-1.321	0.454	-11.677	-1.453	10.224
2020Q3	0.538	-0.403	1.398	-4.339	-0.160	4.179
2020Q4	0.506	0.364	1.540	0.938	4.775	3.836
2021Q1	0.555	0.251	1.435	0.220	4.146	3.926
2021Q2	0.472	0.373	1.592	0.815	4.429	3.614
2021Q3	0.551	0.266	1.439	0.318	4.243	3.925
2021Q4	0.709	-0.036	1.186	-1.465	3.419	4.884

Source: Author.

Table 9

Conditional Distribution Parameters Generated by DeepGaR – Lithuania

Time Period	Standard Deviation	Skewness	Tailweight	.05 Quantile	.95 Quantile	IQR
2017Q3	0.405	0.101	1.542	0.068	2.721	2.653
2017Q4	0.523	-0.027	1.412	-0.643	2.818	3.461
2018Q1	0.517	-0.071	1.463	-0.639	2.770	3.409
2018Q2	0.426	0.076	1.496	-0.143	2.654	2.796
2018Q3	0.491	0.016	1.454	-0.327	2.897	3.224
2018Q4	0.524	-0.035	1.400	-0.743	2.729	3.471
2019Q1	0.537	-0.033	1.393	-0.696	2.869	3.565
2019Q2	0.469	0.030	1.436	-0.358	2.733	3.091
2019Q3	0.380	0.053	1.486	-0.335	2.153	2.489
2019Q4	0.513	-0.046	1.393	-0.882	2.527	3.409
2020Q1	0.472	0.019	1.429	-0.438	2.673	3.111
2020Q2	1.796	-0.965	0.602	-14.763	0.618	15.381
2020Q3	0.283	-0.035	1.536	-1.568	0.269	1.837
2020Q4	0.281	0.225	1.647	0.605	2.523	1.917
2021Q1	0.374	0.106	1.673	0.642	3.053	2.411
2021Q2	0.192	0.180	2.135	1.627	2.843	1.215
2021Q3	0.488	-0.027	1.398	-0.700	2.531	3.231
2021Q4	0.073	-0.387	2.794	0.719	1.353	0.633

Source: Author.

Table 10

Conditional Distribution Parameters Generated by DeepGaR – Romania

Time Period	Standard Deviation	Skewness	Tailweight	.05 Quantile	.95 Quantile	IQR
2016Q4	0.303	-0.027	1.618	0.818	2.207	1.389
2017Q1	0.580	0.147	1.341	0.274	3.097	2.823
2017Q2	0.738	0.174	1.195	-0.302	3.374	3.676
2017Q3	0.607	0.139	1.260	-0.227	2.750	2.978
2017Q4	0.628	0.151	1.296	0.147	3.227	3.080
2018Q1	0.734	0.151	1.220	-0.031	3.597	3.628
2018Q2	0.548	0.098	1.264	-0.358	2.312	2.670
2018Q3	0.515	0.118	1.335	-0.038	2.455	2.493
2018Q4	0.716	0.155	1.205	-0.307	3.241	3.547
2019Q1	0.715	0.143	1.118	-0.874	2.697	3.571
2019Q2	0.894	0.157	1.048	-1.035	3.470	4.505
2019Q3	0.844	0.147	1.012	-1.328	2.938	4.266
2019Q4	0.491	0.108	1.345	-0.109	2.258	2.367
2020Q1	1.346	0.158	0.893	-1.934	4.956	6.890
2020Q2	0.890	-1.055	0.269	-9.281	-4.339	4.942
2020Q3	0.239	-0.113	1.602	-1.020	0.094	1.114
2020Q4	0.369	0.129	1.628	1.512	3.237	1.725
2021Q1	0.214	-0.042	1.836	1.568	2.519	0.951
2021Q2	0.455	0.131	1.489	0.631	2.797	2.166
2021Q3	0.707	0.126	1.100	-0.946	2.582	3.528
2021Q4	0.714	0.137	1.213	-0.199	3.321	3.521

Source: Author.