Bulletin of Applied Economics, 2025, 12(2), 63-76
https://doi.org/10.47260/bae/1224

Greek GDP Forecasting Using Bayesian Multivariate Models

Zacharias Bragoudakis[1*] and Ioannis Krompas[2]

Abstract

Building on a proper selection of macroeconomic variables for constructing a Gross Domestic Product (GDP) forecasting multivariate model (Kazanas, 2017), this paper evaluates whether alternative Bayesian model specifications can provide greater forecasting accuracy compared to a standard Vector Error Correction model (VECM). To that end, two Bayesian Vector Autoregression models (BVARs) are estimated, a BVAR using Litterman's prior (1979) and a BVAR with time-varying parameters (TVP-VAR). The BVAR is found to have statistically significant forecasting gains against the benchmark and the TVP-VAR. Furthermore, the BVAR requires only minimal modifications to account for the effect of pandemic observations on its coefficients, only for longer-term forecasts.

JEL classification: C11, C51, C52, C53.

Keywords: Bayesian VARs, Forecasting, GDP, VECM.

[1] Bank of Greece & Department of Economics, National and Kapodistrian University of Athens.
*Corresponding author.
[2] NBG Economic Research, Ph. Cand. Panteion University.

Article Info: Received: February 4, 2025. Revised: June 14, 2025. Published online: July 1, 2025.

1. Introduction

Macro-econometrics according to Stock and Watson (2001), serves a quadruple purpose: Data description, forecasting, structural inference, and policy analysis. To that end, several types of models, from single-equation to large models with hundreds of equations have been used, like Klein's LINK model in 1980 (Klein, 1976) and more recently, Dynamic Stochastic General Equilibrium (DSGE) models (Christiano et al., 2018).

Lucas and other new classical economists were especially critical of the use of large-scale macro-econometric models to evaluate policy impacts when they were purportedly sensitive to policy changes (Lucas, 1976). Given that the optimal decision rules vary systematically with changes in the structure of series relevant to the decision maker, it follows that any policy change will systematically alter the structure of econometric models.

Sims (1980) framework of Vector Autoregressive models (VARs) came as an answer to this critique. VARs are n-equation, n-variable linear models in which each variable is in turn explained by its own lagged values, plus past values of the remaining n - 1 variables. This simple framework provides a systematic way to capture rich dynamics in multiple time series, while at the same time, the statistical toolkit that came with VARs is easy to use and interpret. As Sims (1980) and others argued in a series of early influential papers, VARs held out the promise of providing a coherent and credible approach to data description, forecasting, structural inference, and policy analysis.

An alternative to standard OLS VARs is the Bayesian VARS (BVARs), initially proposed by Sims (1980) and Doan, Sims, and Litterman (1984), who through Bayesian shrinkage sought to further improve the forecasting performance of the multivariate econometric models available at the time. The BVARs' superiority in forecasting is well established as the literature is rich in Bayesian multivariate models that outperform either standard frequentist or DSGE models, for example, see Gupta and Kabundi (2010).

To estimate a BVAR, the formulation of priors is necessary, with the most popular one being the so-called Minnesota prior (Litterman, 1979). However, since its introduction several more advanced priors have been proposed, such as the one by Sims and Zha (1998). Furthermore, advances in Bayesian statistics and computational capabilities have enabled the use of more complex BVARs, such as the Time-Varying Parameter VARs (TVP-VARs) with the most prominent work on such models being that of Cogley and Sargent (2002;2005), Primiceri (2005), and more recently Carriero et al. (2015).

This advantage of inputting a researcher's belief or knowledge as a BVAR prior has an extra argument in favor of Bayesian specifications, when it comes to specifically forecasting Greek macroeconomic variables. This is because as the time series usually used in estimating Greek macroeconomic models' coefficients start at the year 2000, a significant portion of the sample is comprised of observations that occurred during the economic crisis. This may lead to obtaining coefficients that do not accurately reflect the data generating process of the economy over the long run, and hence it makes sense to limit the parameter space that OLS would have to "search" for coefficient estimation by imposing priors consistent with general macroeconomic stylized facts. Despite that, the application of BVARs in forecasting Greek macroeconomic activity is rather limited, with the most prominent work being that of Louzis about macroeconomic and credit variables forecasting using BVARs (2017) and Greek GDP nowcasting (2018).

Against this background and given the limited application of Bayesian methods to Greek economic data, the aim of this paper is twofold: Firstly to test whether Bayesian multivariate models provide any gains when it comes to forecasting Greek GDP and secondly to examine if the models examined require any modifications to continue provide meaningful forecasts, once observations from the COVID-19 period are included in them. To do so, we sought out to use a set of macroeconomic variables, suitable for use in a multivariate model to generate GDP forecasts, as specified in Kazanas (2017), to estimate alternative specifications of BVARs and examine the accuracy gains in GDP forecasting using as a benchmark the standard frequentist Vector Error Correction Model (VECM) estimated in Kazanas (2017). To do so, we opted to use two alternative models: A BVAR estimated using the Minnesota prior under the notion of limiting the parameter space to obtain a more parsimonious model as explained earlier, and a TVP-VAR to allow the coefficients to change throughout the sample, thus generating forecasts based on the most recent state of the business cycle. The forecasting exercise comprises of a one-quarter ahead, a two-quarter ahead and a four- quarter ahead pseudo out-of-sample forecasts, evaluated over three periods: Pre-COVID-19, during COVID and Post-COVID-19.

The remainder of the paper is organized as follows: In section 2 the data used in estimating the three models (the benchmark VECM, the BVAR, and the TVP-VAR) are presented, in section 3 we concisely present the models and their estimation techniques, while in section 4 we present the forecasting exercise results and lastly in section 5 we discuss the conclusions and policy implications.

2. The data

The variable selection for the Greek GDP forecasting follows Kazanas (2017), where a VECM is constructed including data for real GDP (Y), unemployment rate (U), GDP deflator (P), 10-year government bond yield (GB), and exports as a percentage of GDP (XY). The data sample ranges from 2000:Q1 to 2023:Q2. All data are adjusted for seasonality and sourced from Eurostat's national accounts (Eurostat database code: na10), labor market survey (Eurostat database code: labour), and interest rates (Eurostat database code: irt) databases.

Descriptive statistics of the selected variables are presented in Table 1, while in Figure 1 one can see the evolution of the series over time. It is evident that the variables were performing well up to 2009 when the trend shifted due to the Greek economic crisis, the impact of which is reflected in the sharp decrease of GDP, the increase of unemployment and the bond yield volatility approximately between 2009 and 2015. Furthermore, GDP, unemployment rate, and exports display increased volatility around the second quarter of 2020, which reflects the impact of the pandemic and the corresponding lockdown. Lastly, the GDP deflator records a sharp increase towards the end of the sample (2021Q3 onwards) reflecting the post-pandemic inflation wave.

Table 1: Descriptive Statistics

	GDP (Y)	Unemployment (U)	GDP deflator (P)	Bond Yield (GB)	Exports (XY)
Unit	Euros	Percentage	2015=100	Percentage	Percentage
Mean	5.03E+10	0.156223	95.41323	0.063812	0.286479
Median	4.82E+10	0.127500	99.38350	0.049650	0.263000
Maximum	6.34E+10	0.281000	113.1220	0.254000	0.515000
Minimum	3.89E+10	0.076000	74.49300	0.007000	0.183000
Std. Dev.	6.64E+09	0.064680	9.627995	0.047900	0.085396
Skewness	0.605995	0.562050	-0.714195	2.157979	0.781572
Kurtosis	2.003059	1.855076	2.446428	8.247366	2.680626

Source: Author's calculations

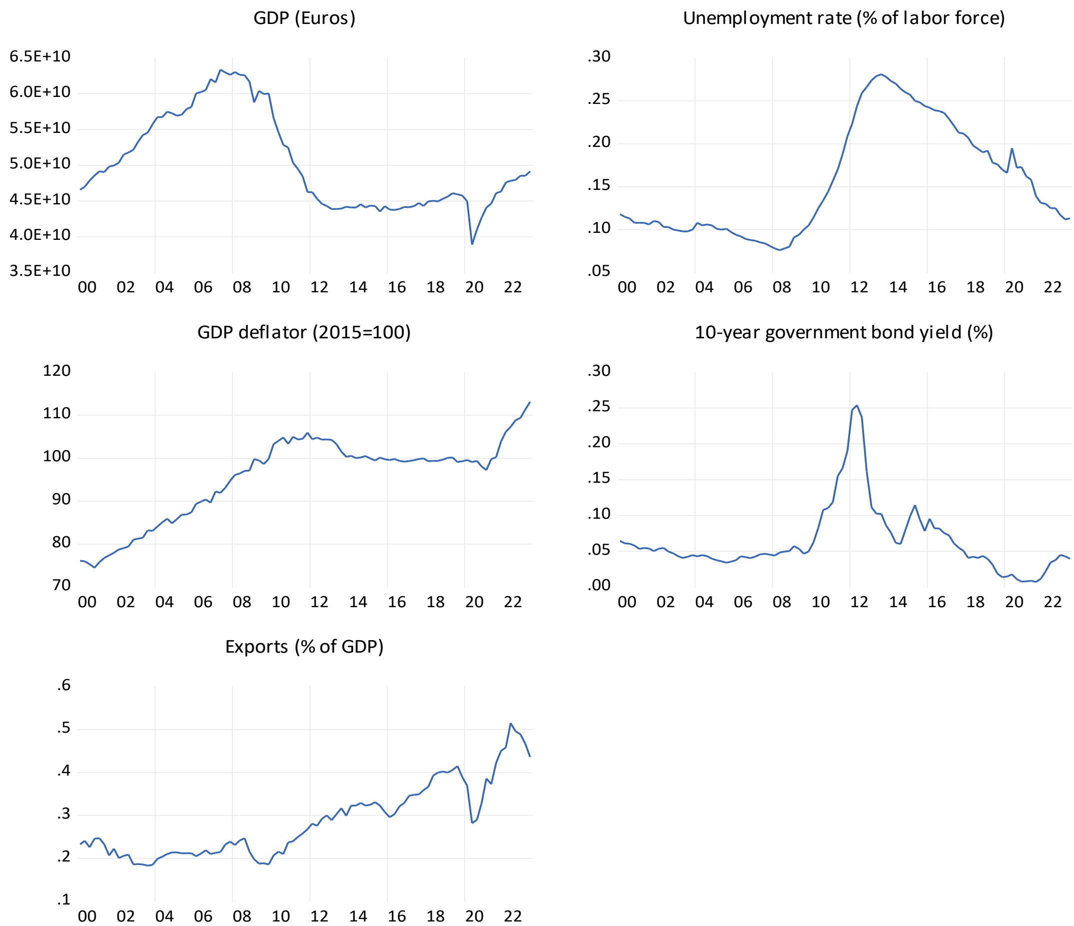

Figure 1: Sample Variables (2000Q1- 2023Q2)

Source: Eurostat

For each variable, the ADF unit root test (Dickey and Fuller, 1981) was conducted, as stationarity is a prerequisite, especially in estimating a standard VECM or a VAR model. All variables have a unit root in levels but are stationary if they are transformed. For Y, P and GB[3] the transformation consists of log differing the variables, whereas for U and XY the transformation consists of simple differencing (Henceforth, lowercase letters denote the transformed variables).

Table 2: ADF test p-values

	Y	U	P	GB	XY
Levels	0.5357	0.1718	0.9976	0.1124	0.1308
Transformed	0.0000	0.0273	0.0000	0.0000	0.0000

[3] As GB is a percentage simple differencing should be a better approach. However, this leads to models estimated having non-white noise residuals. Thus log-differencing is preferred.

3. Models presentation and estimation

The three models mentioned earlier are estimated using the abovementioned variables. The models are estimated over sixteen years (2000:Q1 to 2015:Q4), whereas the remaining sample (2016:Q1 to 2023:Q2) is to be used for evaluating their forecasting performance. The VECM and the BVAR model are estimated using EViews 10, while the TVP-VAR is estimated using the BEAR Toolbox 4.2 (Dieppe et al. 2016).

3.1 The VEC benchmark model

Based on the works of Granger (1981) and Engle and Granger (1987), Vector error correction models are essentially restricted VARs, which contain a set of variables both in differences and in levels. The differences of the variables included in the model represent the short-run interrelations of the variables, whereas the linear combination of the levels of the variables, commonly referred to as the cointegrating vector (or vectors, as more than one linear combination of a set of variables can be included), represents the long-run dynamics of the variables. Mathematically, a representative VEC model can be written as follows:

$$\Delta y_t = m + \sum_{i=1}^{p-1} B_i \Delta y_{t-i} + A y_{t-1} + \varepsilon_t \quad (1)$$

$$\varepsilon_t \sim N(0, \Sigma)$$

Where y is the vector containing the variables (in our case $y' = [\ y\ u\ p\ gb\ xy\ ]$), m is the vector containing the constants of the equations system, B_i is the matrix that contains the coefficients that describe the short-run impact of the variables' lag i , p is the lag length, A is the cointegrating vector that includes the coefficients that capture the long-run relationship between the variables, with ε_t being the vector of spherical disturbances and Σ being their variance-covariance matrix. The model can also be expanded to include exogenous variables. VECMs are very useful in modeling non-stationary time-series without having to exclude their long-run behavior, but, like their unrestricted counterparts (VARs), they suffer from the "curse of dimensionality", as the addition of a variable significantly increases the number of coefficients to be estimated.

The estimation of this model follows the Johansen procedure (Johansen, 1995). A VAR is estimated in levels (including a constant and a trend) and by incorporating lag length criteria, namely the sequential modified LR test statistic, the Final prediction error, the Akaike information criterion, the Scwartz information criterion, and the Hanna-Quinn information criterion it is found that two lags are optimal. The existence of a cointegrating relation between the variables must be confirmed in order to use a VECM specification rather than a simple VAR in differences. The max eigenvalue cointegration test is therefore used, which indicates the existence of two cointegrating vectors at the 5% level[4]. Hence a VECM is estimated, with 1 lag per variable and 2 cointegrating vectors.

[4]In Kazanas (2017) the existence of the second cointegrating vector is rejected as the hypothesis of at most 1 cointegrating vector is marginally accepted with a P-value of 0.0505, but since then a major benchmark revision of the Greek macroeconomic data has occurred causing the maximum eigenvalue cointegration test to indicate the existence of a second cointegrating vector.

Table 3: Maximum Eigenvalue Cointegration Test

Hypothesized Number of Cointegrating equations	Eigenvalue	Max-Eigenvalue statistic	5% Critical value	P-value
None *	0.567460	51.96106	38.33101	0.0008
At most 1 *	0.450144	37.08209	32.11832	0.0114
At most 2	0.296067	21.76650	25.82321	0.1571
At most 3	0.255644	18.30464	19.38704	0.0713
At most 4	0.123745	8.190059	12.51798	0.2366

* Denotes rejection of the hypothesis at the 0.05 level
**MacKinnon-Haug-Michelis (1999) p-values

Upon estimation the model's residuals are tested and found to be homoscedastic, non-autocorrelated and normally distributed.

3.2 The BVAR model

Under the Bayesian approach to econometrics, the estimated coefficients of a model are not an attempt to estimate their true value, but instead, they are perceived as a summary of the posterior distribution, which in its turn is proportional to the likelihood function times the prior distribution. Priors represent any knowledge the researcher has beforehand about the coefficients. Following this technique results in the coefficients being essentially a matrix-weighted average between the imposed priors and a regular OLS estimation (Ouliaris et al, 2016), which leads the variables to behave as if they were random walks (Del Negro and Schorfheide, 2010). Mathematically, a representative BVAR model can be written as follows:

$$\Delta y_t = m + B_1 \Delta y_{t-1} + B_2 \Delta y_{t-2} + \ldots + B_p \Delta y_{t-p} + \varepsilon_t \qquad (2)$$
$$\varepsilon_t \sim N(0, \Sigma)$$

Where (as in the previous model) *y* is the vector containing the variables (in our case *y'= [y u p gb xy]*), m is the vector containing the constants of the equations system, B_i (for *i*=1, 2,…, *p*) are the matrices that contains the coefficients of lag *i* and ε_t is the vector containing the error terms, with Σ being their variance-covariance matrix.

In a BVAR, model parameters are obtained by:

$$\hat{b} = [V^{-1} + \Sigma_e^{-1} \otimes (X'X)]^{-1} [V^{-1}\bar{b} + (\Sigma_e^{-1} \otimes X')Y] \qquad (3)$$

where $\hat{b}$ is the matrix of the estimated VAR coefficients, V is the variance matrix of the prior distribution of the model's coefficients, Σ_e is the variance-covariance matrix of the model's residuals and $\bar{b}$ is a vector containing the prior means of each variable's own first lag coefficients. *Y* contains the model endogenous variables, whereas *X* includes variable lags plus any exogenous variables that a researcher might want to include in the model.

The error variance-covariance matrix Σ_e necessary for the coefficient estimation is either estimated by fitting an AR(1) model on every variable and getting the error variances, by estimating an AR(1) and a VAR to obtain the diagonal elements of the variance-covariance matrix, or by estimating all variances-covariances as implied by a full VAR (an option not commonly used, as it can lead to a singular matrix). Under the Minnesota prior, the researcher is required to specify a set of hyperparameters in order to formulate the priors to obtain the model's coefficients: μ_1, λ_1, λ_2, and λ_3.

μ_1 is used as the prior mean of the coefficients in the matrix $\bar{b}$ and it usually takes the value of 0 (if the variables of the model are stationary) or 1 (if the variables of the model have a unit root). λ_1, λ_2 and λ_3 are used to formulate diagonal elements of the V matrix (with non-diagonal elements being set to 0). More specifically, each diagonal element of the V matrix for the j-th variable in the i-th equation at lag k is formulated as follows:

$$\left(\frac{\lambda_1}{k^{\lambda_3}}\right)^2 \text{ for } i = j, \tag{4}$$

$$\left(\frac{\lambda_1 \lambda_2 \sigma_i}{k^{\lambda_3} \sigma_j}\right)^2 \text{for } i \neq j \tag{5}$$

where σ_i, σ_j are the square roots of the corresponding elements of the Σ_e matrix.

This way λ_1 determines how binding the restrictions are. The closer to zero the value of λ_1 is, the more binding the restrictions are in the estimation of the coefficients. A value over 10 implies an uninformative prior. λ_2 determines the cross-variable effects in the equations and is set between 0 and 1. The closer the value is to 1 the more lags of variable j impact variable i (for j≠i) in the BVAR. Finally, λ_3 determines the decay rate of the own lags of a variable, excluding the first lag. As this hyper-parameter approaches zero, higher order lags decay at a slower rate.
To proceed with the estimation, we fit an AR(1) model throught each variable for variance estimation. Prior hyperparameters are chosen through optimization technique that selects the model with the highest marginal likelihood, similar to Giannone et al. (2012). This process yields as optimal hyper parameters μ_1=0.3, λ_1=0.2, λ_2=0.4 and λ_3=1.

3.3 The TVP-VAR

The time-varying parameter VAR is a model that allows model coefficients to change over time. This is particularly useful in capturing nonlinear relationships in the data as any model with time-varying parameters can successfully represent any nonlinear functional form (Swamy, 1975 and Granger, 2008). Macroeconomic variables are known to impact differently each other across the business cycle or after structural changes, hence the TVP-VAR is an interesting approach to econometric modeling. The functional form of a TVP-VAR is expressed as:

$$\Delta y_t = m_t + B_{1,t}\Delta y_{t-1} + B_{2,t}\Delta y_{t-2} + \ldots + B_{p,t}\Delta y_{t-p} + \varepsilon_t \tag{6}$$
$$\text{Where } \varepsilon_t \sim N(0, \Sigma_t)$$

With y_t being the matrix containing the variables m_t being the vector of time-varying constant parameters, $B_{i,t}$ being the matrix containing the time-varying coefficients for i=*1,2…, p.* Finally, ε_t is the vector containing the error terms, with Σ_t being their (time-varying) variance-covariance matrix. Elements $\beta_{i,t}$ of the $B_{i,t}$ matrices are assumed to follow a random walk process:

$$\beta_{i,t} = \beta_{i,t-1} + v_{i,t} \tag{7}$$

Where $v_{i,t} \sim N(0, \Omega_i)$ are normally distributed shocks with Ω being their variance -covariance matrix.

Apart from time-varying parameters of the conditional mean, TVP-VARs include stochastic volatility (hence the Σ_t). This approach makes the model heavily parametrized but is necessary to avoid bias in the coefficients across potential volatility clusters, falsely attributing variance shocks to coefficient variation (Sims, 2002). The formulation of the Σ_t matrix is based on Cogley and Sargent (2005) in the BEAR toolbox.

Under this approach the Σ_t matrix has $f_{n,m}$ non-diagonal elements which are time invariant and are assumed to follow a multivariate normal distribution.
The diagonal elements of the Σ_t matrix are of the form $\bar{s}_i e^{\lambda_{i,t}}$, with $\bar{s}_i$ being a time invariant scaling factor. On the other hand, $\lambda_{i,t}$ follows an AR(1) process:

$$\lambda_{i,t} = \gamma \lambda_{i,t-1} + u_{i,t} \tag{8}$$

Where $u_{i,t} \sim N\ (0, \varphi_i)$
With γ being a persistence factor to be determined by the researcher. Furthermore, the prior of φ_i is assumed to follow an Inverse Gamma distribution:

$$\varphi_i \sim IG(\frac{\alpha_0}{2}, \frac{\delta_0}{2})$$

With α_0, δ_0 being scaling factors that also need to be determined by the researcher. As the posterior for the $f(B_{i,t}, \Omega_i, f^{-1}, \lambda_{i,t}, \varphi_i | y_t)$ cannot be analytically solved, once the abovementioned hyperparameters have been chosen, the Gibbs sampler must be used to obtain results. For more detailed presentations of TVP-VARs, one can look up Primiceri (2005), Chan and Jeliazkov (2009), Lubik and Matthes (2015), or Dieppe et al. (2018).
The above equations imply that there is no mechanism in the model to produce future values of the coefficients of the model, as in the absence of new shocks, coefficients remain the same. It is an interesting approach, however, to attempt a forecast based on the most recent interrelations between the variables and neglect coefficient values of the past, that may not adequately represent the dynamics of the system anymore.
In our estimation of the model, we follow Primiceri (2005) in choosing the number of lags, which is set to 2. We also set $\alpha_0 = \delta_0 = 0.001$ implying a rather uninformative prior. Furthermore, we set γ= 0.95, implying strong persistence of variance shocks thus limiting the possibility of explosive behavior in the model's coefficients[5] (strong persistence of shocks is also a valid macroeconomic assumption). We set the Gibbs sampler to perform 15000 iterations, along with 10000 are burn-in iterations, with a selection of 1 draw over 10. This specification, while time consuming makes sure that the MCMC has converged (Dieppe et al. 2018), thus providing a safe estimation.

4. Forecasting Evaluation

To evaluate the forecasting performance of the models three sets of recursive pseudo out of sample forecasts are carried out: a one-quarter ahead forecast, a two-quarter ahead forecast and a four-quarter ahead forecast. As mentioned earlier, the models are estimated up to 2015:Q4 and the forecast evaluation sample expands from 2016:Q1 to 2023:Q2. Due to the fact that in the forecast evaluation sample there are periods of increased volatility, because of the COVID-19 pandemic effect, loss function results are reported across three sub-periods, so as to compare the forecast performance clearer, given that during the pandemic the forecasting performance of all models is expected to deteriorate. The three sub-periods are: The pre-COVID-19 period (2016:Q1-2019:Q4), the COVID-19 period (2020:Q1-2021:Q2) and the post-COVID-19 period (2021:Q3 to 2023:Q2). The allocation of quarters into the COVID-19 period was made based on the enforcing of the first and the lifting of the last general lockdown restrictions, which is thought to be the

[5]To further check for such behavior in the model's coefficients, upon estimation we performed stationarity test. All coefficients are found to be stationary within a 10% level of significance (most in levels and a few in first differences).

main distorting factor of economic activity. In the two and four-quarter ahead forecasting exercises, a forecast is allocated into the COVID-19 period if it includes at least one quarter of said period.
The forecasts are evaluated using the Mean Absolute Percentage Error (MAPE), the Mean Absolute Error (MAE), and the Root Mean Squared Error (RMSE) and the results are reported in table 3:

$$MAPE = \left(\frac{1}{n}\sum_{t=1}^{n}\frac{|Y_t - \widehat{Y_t}|}{Y_t}\right) * 100 = \left(\frac{1}{n}\sum_{t=1}^{n}\frac{|e_t|}{Y_t}\right) * 100 \quad (9)$$

$$MAE = \frac{1}{n}\sum_{t=1}^{n}\left|Y_t - \widehat{Y_t}\right| = \frac{1}{n}\sum_{t=1}^{n}|e_t| \quad (10)$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(Y_t - \widehat{Y_t})^2} = \sqrt{\frac{1}{n}\sum_{t=1}^{n} e_t^2} \quad (11)$$

Where $\widehat{Y_t}$ is the forecasted value from the model, Y_t being the actual value and *n* being the number for steps ahead forecasted (1,2 and 4 in our case).

As evidenced in tables 4 through 6, the BVAR outperforms both the benchmark and the TVP-VAR across subperiods, forecast horizons and loss functions[6].

Table 4: Forecast Evaluation based on MAPE

		Pre-COVID-19	COVID-19	Post-COVID-19	Total
1Q Ahead	VECM	0.81	8.41	1.47	2.51
	BVAR	**0.69**	**6.07**	**0.96**	**1.84**
	TVP-VAR	0.83	7.70	1.53	2.39
2Q Ahead	VECM	0.95	11.52	2.24	3.45
	BVAR	**0.86**	**8.17**	**1.31**	**2.48**
	TVP-VAR	1.03	10.20	2.11	3.19
4Q Ahead	VECM	1.89	15.80	4.06	5.38
	BVAR	**1.63**	**9.74**	**2.36**	**3.57**
	TVP-VAR	2.11	12.13	3.56	4.61

Source: Author's calculations. Bold values indicate best performing model.

More specifically, the BVAR performs well in the pre-Covid-19 period in with a MAPE of 0.69 for the one quarter ahead forecast, 0.86 for the two quarter ahead forecast and 1.63 for four quarter ahead forecast (vs 0.81, 0.95 and 1.89 for the benchmark VECM and 0.83, 1.03 and 2.11 for the TVP-VAR respectively).

[6] Results are found to be generally robust to BVAR specification both in terms of lag length and hyperparameters. Lag changes lead only to marginal changes in forecasting performance of the order of ±0.2% in terms of MAPE (excluding pandemic observations) across subperiods and forecasting horizons. Hyperparameter changes also lead to only marginal changes in forecasting performance, with no alternative specification consistently outperforming the initial one.

Table 5: Forecast Evaluation based on MAE

		Pre-COVID-19	COVID-19	Post-COVID-19	Total
1Q Ahead	VECM	3.15E+08	3.41E+09	6.71E+08	1.03E+09
	BVAR	**3.09E+08**	**2.52E+09**	**4.44E+08**	**7.88E+08**
	TVP-VAR	3.72E+08	3.19E+09	7.26E+08	1.03E+09
2Q Ahead	VECM	4.34E+08	4.57E+09	1.02E+09	1.43E+09
	BVAR	**3.90E+08**	**3.38E+09**	**6.07E+08**	**1.06E+09**
	TVP-VAR	4.63E+08	4.29E+09	1.01E+09	1.39E+09
4Q Ahead	VECM	8.83E+08	5.78E+09	1.83E+09	2.15E+09
	BVAR	**7.48E+08**	**3.92E+09**	**1.09E+09**	**1.52E+09**
	TVP-VAR	9.04E+08	5.30E+09	1.71E+09	2.03E+09

Source: Author's calculations. Bold values indicate best performing model.

Table 6: Forecast Evaluation based on RMSE

		Pre-COVID-19	COVID-19	Post-COVID-19	Total
1Q Ahead	VECM	3.15E+08	3.41E+09	6.71E+08	1.03E+09
	BVAR	**3.09E+08**	**2.52E+09**	**4.44E+08**	**7.88E+08**
	TVP-VAR	371978647	3.19E+09	7.26E+08	1.03E+09
2Q Ahead	VECM	4.72E+08	4.87E+09	1.09E+09	1.53E+09
	BVAR	**4.20E+08**	**3.65E+09**	**6.50E+08**	**1.14E+09**
	TVP-VAR	4.92E+08	4.6E+09	1.06E+09	1.48E+09
4Q Ahead	VECM	1.09E+09	6.24E+09	1.95E+09	2.39E+09
	BVAR	**9.34E+08**	**4.26E+09**	**1.19E+09**	**1.72E+09**
	TVP-VAR	1.10E+09	5.78E+09	1.91E+09	2.29E+09

Source: Author's calculations. Bold values indicate best performing model.

As expected, forecasting performance deteriorates for the BVAR as it does for all models during the COVID-19 period (mean absolute percentage errors of 6.1 for the one quarter ahead, 8.2 for the two-quarter ahead and 9.7 for the four quarter ahead forecast), but it still performs better than the other competing models. Forecast errors are lower in the post COVID-19 period but remain elevated compared to the pre-COVID-19 period (mean absolute percentage error of 0.96 for the one quarter ahead forecast, 1.31 for the two-quarter ahead forecast and 2.36 for the four quarter ahead forecast). In Figure 2, the one step ahead forecasts for the competing models are presented.

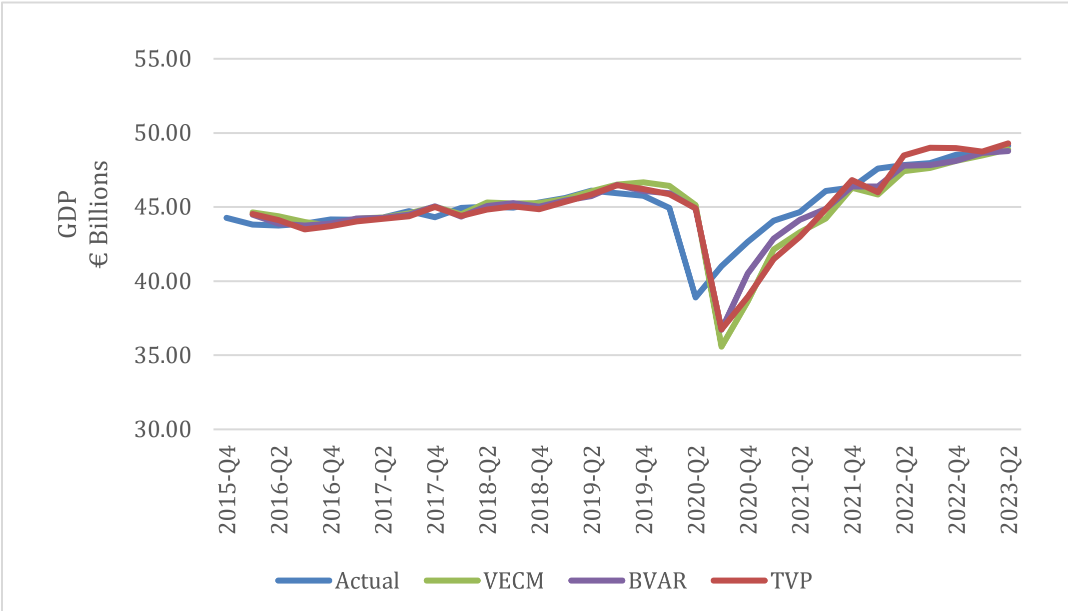

Figure 2: 1 step ahead forecast evaluation graphical representation

Source: Author's calculations

Given that we have more than two competing models, we employ the Model Confidence Set test (Hansen et al., 2011) to identify the subset of "true" models, that are accepted as statistically superior in terms of forecasting. Indeed, the BVAR is accepted as the only model to belong in the model confidence set for all three forecasting exercises and for any reasonable level of significance. The test rejects that the rest models belong in the set, for any level of significance.

Focusing on the post-COVID-19 period performance and on the fact that variable observations that occurred during the pandemic may have "distorted" the models' coefficients due to the exogenous effects of the time, we performed a second run of the forecasting exercise, after we modified the models to reduce the effect of the pandemic observations on the model estimations. Each model was modified in the following ways:

1. In the VECM we included a dummy variable in the short run dynamics of the model that takes the value of 1 during the pandemic period and 0 otherwise.
2. In the TVP-VAR we reduce the autoregressive prior γ to 0.75 to allow exogenous shocks to fade away faster.
3. When it comes to BVARs literature is relatively richer in regard with treatment of pandemic observations. More specifically, Lenza and Primiceri (2022) weigh pandemic observations, whereas Schorfheide and Song (2021) drop those observations altogether (even though they ultimately advise against this practice). Other methods include either allowing for fatter error distribution tails as for example in Bobeica and Hartwig (2022) or including stochastic volatility in the models as in Carriero et al. (2022). Given that these alternatives were developed with the pandemic ongoing and that we do not expect a pandemic recurrence we opted for simpler approaches. One includes simply recalibrating the BVAR's priors similar to what Sznajderska and Haug (2023) did in evaluating BVARs for the U.S. economy. The second approach follows Cascaldi-Garcia (2022) in including a dummy prior that takes the value of 1 during the pandemic and zero otherwise. We do so by including in the matrix V of eq. (3) the term $(\lambda^1\lambda^4)^2$ in the corresponding to the exogenous variabls variance elements, where λ^4 is a variance scaling parameter (usually taking large values to represent the uncertainty surrounding variables not determined in the model). Yet again, hyperparameters are optimized, resulting in the following: μ_1=0.2, λ_1=0.1, λ_2=1 and λ_3=1 for both model specifications, with λ_4=100 in the second specification.

The models are estimated up to 2021:Q2 and the exercise was repeated for 2021:Q3 to 2023:Q2. The original BVAR yet again is the best specification as no model outperforms it in the one and two quarter ahead forecasts, with the Model Confidence Set test also rejecting all other specifications. On the other hand, in the four quarter ahead forecast it is the recalibrated BVAR (without the dummy and prior adjustment) that provides both the most accurate forecasts (2.26 MAPE vs 2.36 for the original BVAR) and is also deemed the only model to be included in the Model Confidence Set by the relevant test.

5. Conclusions

Three VARs were estimated using a given set of variables aiming to examine whether Bayesian estimation could provide real GDP forecasting gains. Using two different Bayesian VAR estimation methods, namely Bayesian estimation using a Minnesota-Litterman prior and a TVP-VAR it is found that a relatively simple BVAR performs better than the VECM and a more complex model such as the TVP-VAR. Furthermore, it is shown that the model does not need any modifications to account for the effect of pandemic observations in its coefficients and that only a prior recalibration provides forecasting gains in longer term forecasts. However, a word of caution is needed, as the post-COVID-19 period sample is still small, and the analysis should most likely be repeated to check whether forecast errors eventually converge to their pre-pandemic levels.

This forecasting exercise demonstrated that even the most basic of Bayesian priors provided forecasting gains when it comes to Greek GDP forecasting, but this is only one of the available priors a researcher is available to choose from. One could extend this research to include more advanced Bayesian priors such as the Sims-Zha prior (Sims and Zha, 1998) that incorporates the existence of unit roots and cointegrating relationships in the priors (as it is found in Table 3 that cointegration relationships exist between the variables of the given set. Another way this research could be extended is by using the TVP-VAR estimated above (possibly using a larger sample if available, to account for the model's intensive parameterization), to compute the variation in the relations between macroeconomic variables, as expressed by the time-varying coefficients, and thus examine structural changes of the Greek economy over time. This model can also be used to perform impulse response analysis on specific dates, which allows examining how differently exogenous shocks would affect the Greek economy, at different points in time.

Acknowledgment

Thank D. Louzis and T. Kazanas for their constructive suggestions that helped us to improve the clarity of the paper. The paper has also benefited from the comments of the participants of the 34th Panhellenic Statistics Conference, Athens, 2022. The paper has also been presented at the IMAEF 2024 International Conference.

Declarations

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Greece and National Bank of Greece. The authors are responsible for any errors or omissions in this paper.

References

Bobeica, E. and Hartwig, B. (2023). 'The COVID-19 shock and challenges for inflation modelling'. International Journal of Forecasting, Vol 39, pp 519-539.

Carriero, A., Mumtaz, H. Theophilopoulou, A. (2015). 'Macroeconomic information, structural change, and the prediction of fiscal aggregates'. International Journal of Forecasting, Vol. 31 pp.325-348.

Carriero, A., Clark, T., Marcellino, M., & Mertens, E. (2022). 'Addressing COVID-19 outliers in BVARs with stochastic volatility. Review of Economics and Statistics, pp 1-38.

Cascaldi-Garcia, D. (2022). 'Pandemic Priors'. International Finance Discussion Papers 1352. Board Of Governors of the Federal Reserve System.

Chan, J., Jeliazkov, I. (2009). 'Efficient simulation and integrated likelihood estimation in state space models'. International Journal of Mathematical Modelling and Numerical Optimisation, Vol.1 (1-2), pp. 101-120.

Christiano, J. L., Eichenbaum, S.M. and Trabandt, M. (2018). 'On DSGE models', NBER Working Paper 24811.

Cogley, T., and Sargent, T. J. (2002). 'Evolving post-World War II US inflation dynamics'. NBER macroeconomics Annual, Vol. 16, pp. 331–388.

Cogley, T., and Sargent, T. J. (2005). 'Drifts and volatilities: monetary policies and outcomes in the post WWII US', Review of Economic Dynamics, Vol.8(2), pp. 262–302.

Del Negro, M. and Schorfheide, F. (2010). Bayesian Macroeconometrics, In: Del Negro, M. and Schorfheide, F. Handbook of Bayesian Econometrics.

Dickey, D. and Fuller, W. (1981). 'Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root', Econometrica, Vol.49, pp. 1057-1072.

Dieppe, A., van Roye, B. & Legrand, R. (2016). 'The BEAR toolbox', European Central Bank Working Paper 1934.

Dieppe, A., van Roye, B. & Legrand, R. (2018). The Bayesian Estimation, Analysis and Regression (BEAR) Toolbox Technical guide.

Doan, T., Litterman, R. and Sims, C. (1983). Forecasting and Conditional Projection Using Realistic Prior Distributions, NBER Working Paper 1202.

Engle, R. and Granger, C. (1987). 'Co-Integration and Error Correction: Representation, Estimation, and Testing', Econometrica, Vol. 55, pp. 251-276.

Giannone, D., Lenza, M. and Primiceri, G. (2012). 'Prior Selection for Vector Autoregressions' Working Papers ECARES 2012-002, Universite Libre de Bruxelles.

Giannone, D., Lenza, M. and Primiceri, G. (2015). 'Prior Selection for Vector Autoregressions', Review of Economics and Statistics, Vol. 97, pp. 436-451.

Granger, C. (1981). 'Some properties of time series data and their use in econometric model specification', Journal of Econometrics, Vol. 16, pp. 121-130.

Granger, C. (2008). 'Non-linear Models: Where Do We Go Next? – Time-Varying Parameter Models', Studies in Nonlinear Dynamics and Econometrics, Vol. 12, pp.1-10.

Hansen, P. R., Lunde, A., & Nason, J. M. (2011). 'The Model Confidence Set'. Econometrica, 79(2), 453–497.

Johansen, S. (1995). Likelihood-based Inference in Cointegrated Vector Autoregressive Models. Oxford: Oxford University Press.

Kazanas, T. (2017). A Vector Error Correction Forecasting Model of the Greek Economy, Hellenic Fiscal Council Working Paper 2.

Klein, L. R. (1976). 'Project LINK: Linking National Economic Models', Challenge, Vol. 19, pp. 25–29.

Gupta, R., Kabundi, A. (2010). 'Forecasting macroeconomic variables in a small open economy: a comparison between small- and large-scale models', Journal of Forecasting, Vol. 29, pp. 168-185.

Lenza, M., & Primiceri, G. E. (2022). How to estimate a VAR after March 2020. Journal of Applied Econometrics, 37, pp. 688–699.

Litterman, R., (1979). Techniques of forecasting with Bayesian vector autoregressions, Federal Reserve Bank of Minneapolis Working Papers 115.

Louzis, D. P. (2017). 'Macroeconomic and credit forecasts during the Greek crisis using Bayesian VARs'. Empirical Economics, Vol. 53, pp. 569-598.

Louzis, D. P. (2018). 'Greek GDP revisions and short-term forecasting', Bank of Greece Economic Bulletin, Vol. 48.

Lubik, T. and Matthes, C. (2015). 'Time-Varying Parameter Vector Autoregressions: Specification, Estimation and an Application', Federal Reserve Bank of Richmond Economic Quarterly, Vol. 101, pp. 323-352.

Lucas, R. (1976). 'Econometric policy evaluation: A critique', Carnegie-Rochester Conference Series on Public Policy, Vol. 1, pp. 19-46.

Mackinnon, J., Haug, A., and Michelis, L. (1999). 'Numerical Distribution Functions of Likelihood Ratio Tests for Cointegration', Journal of Applied Econometrics, Vol. 14, pp.563-577.

Ouliaris, S., Pagan, A. R., Restrepo J. (2016). Bayesian Vars, In: Ouliaris, S., Pagan, A. R., Restrepo J. Quantitative Macroeconomic Modelling with Structural Vector Autoregressions - An Eviews Implementation, 49-62.

Primiceri, G. (2005). 'Time Varying Structural Vector Autoregressions and Monetary Policy', The Review of Economic Studies, Vol. 72, pp. 821-852.

Schorfheide, F., & Song, D. (2021). Real-time forecasting with a (standard) mixed frequency VAR during a pandemic. Working Paper No. 29535. National Bureau of Economic Research.

Sims, C. (1980). 'Macroeconomics and Reality', Econometrica, Vol. 48, 1-48.

Sims, C. (2002). 'Comments on Cogley and Sargent's 'Evolving Post-World War II U.S. Inflation Dynamics', NBER Macroeconomics Annual 2001, Vol. 16, pp. 373-379.

Sims, C. and Zha, T. (1998). 'Bayesian Methods for Dynamic Multivariate Models', International Economic Review, Vol. 39, pp. 949-968.

Stock, J. and Watson, W. (2001). 'Vector Autoregressions', The Journal of Economic Perspectives, Vol. 15, pp. 101-115.

Swamy, P. (1975). 'Bayesian and Non-Bayesian Analysis of Switching Regressions and a Random Coefficient Regression Model', Journal of the American Statistical Association, Vol. 70, pp. 593-602.

Sznajderska, A. and Haug, A. (2023). 'Bayesian VARs of the U.S. economy before and during the pandemic'. Eurasian Economic Review Vol.13, pp.211-236.