BANK ĊENTRALI TA' MALTA
EUROSISTEMA
CENTRAL BANK OF MALTA

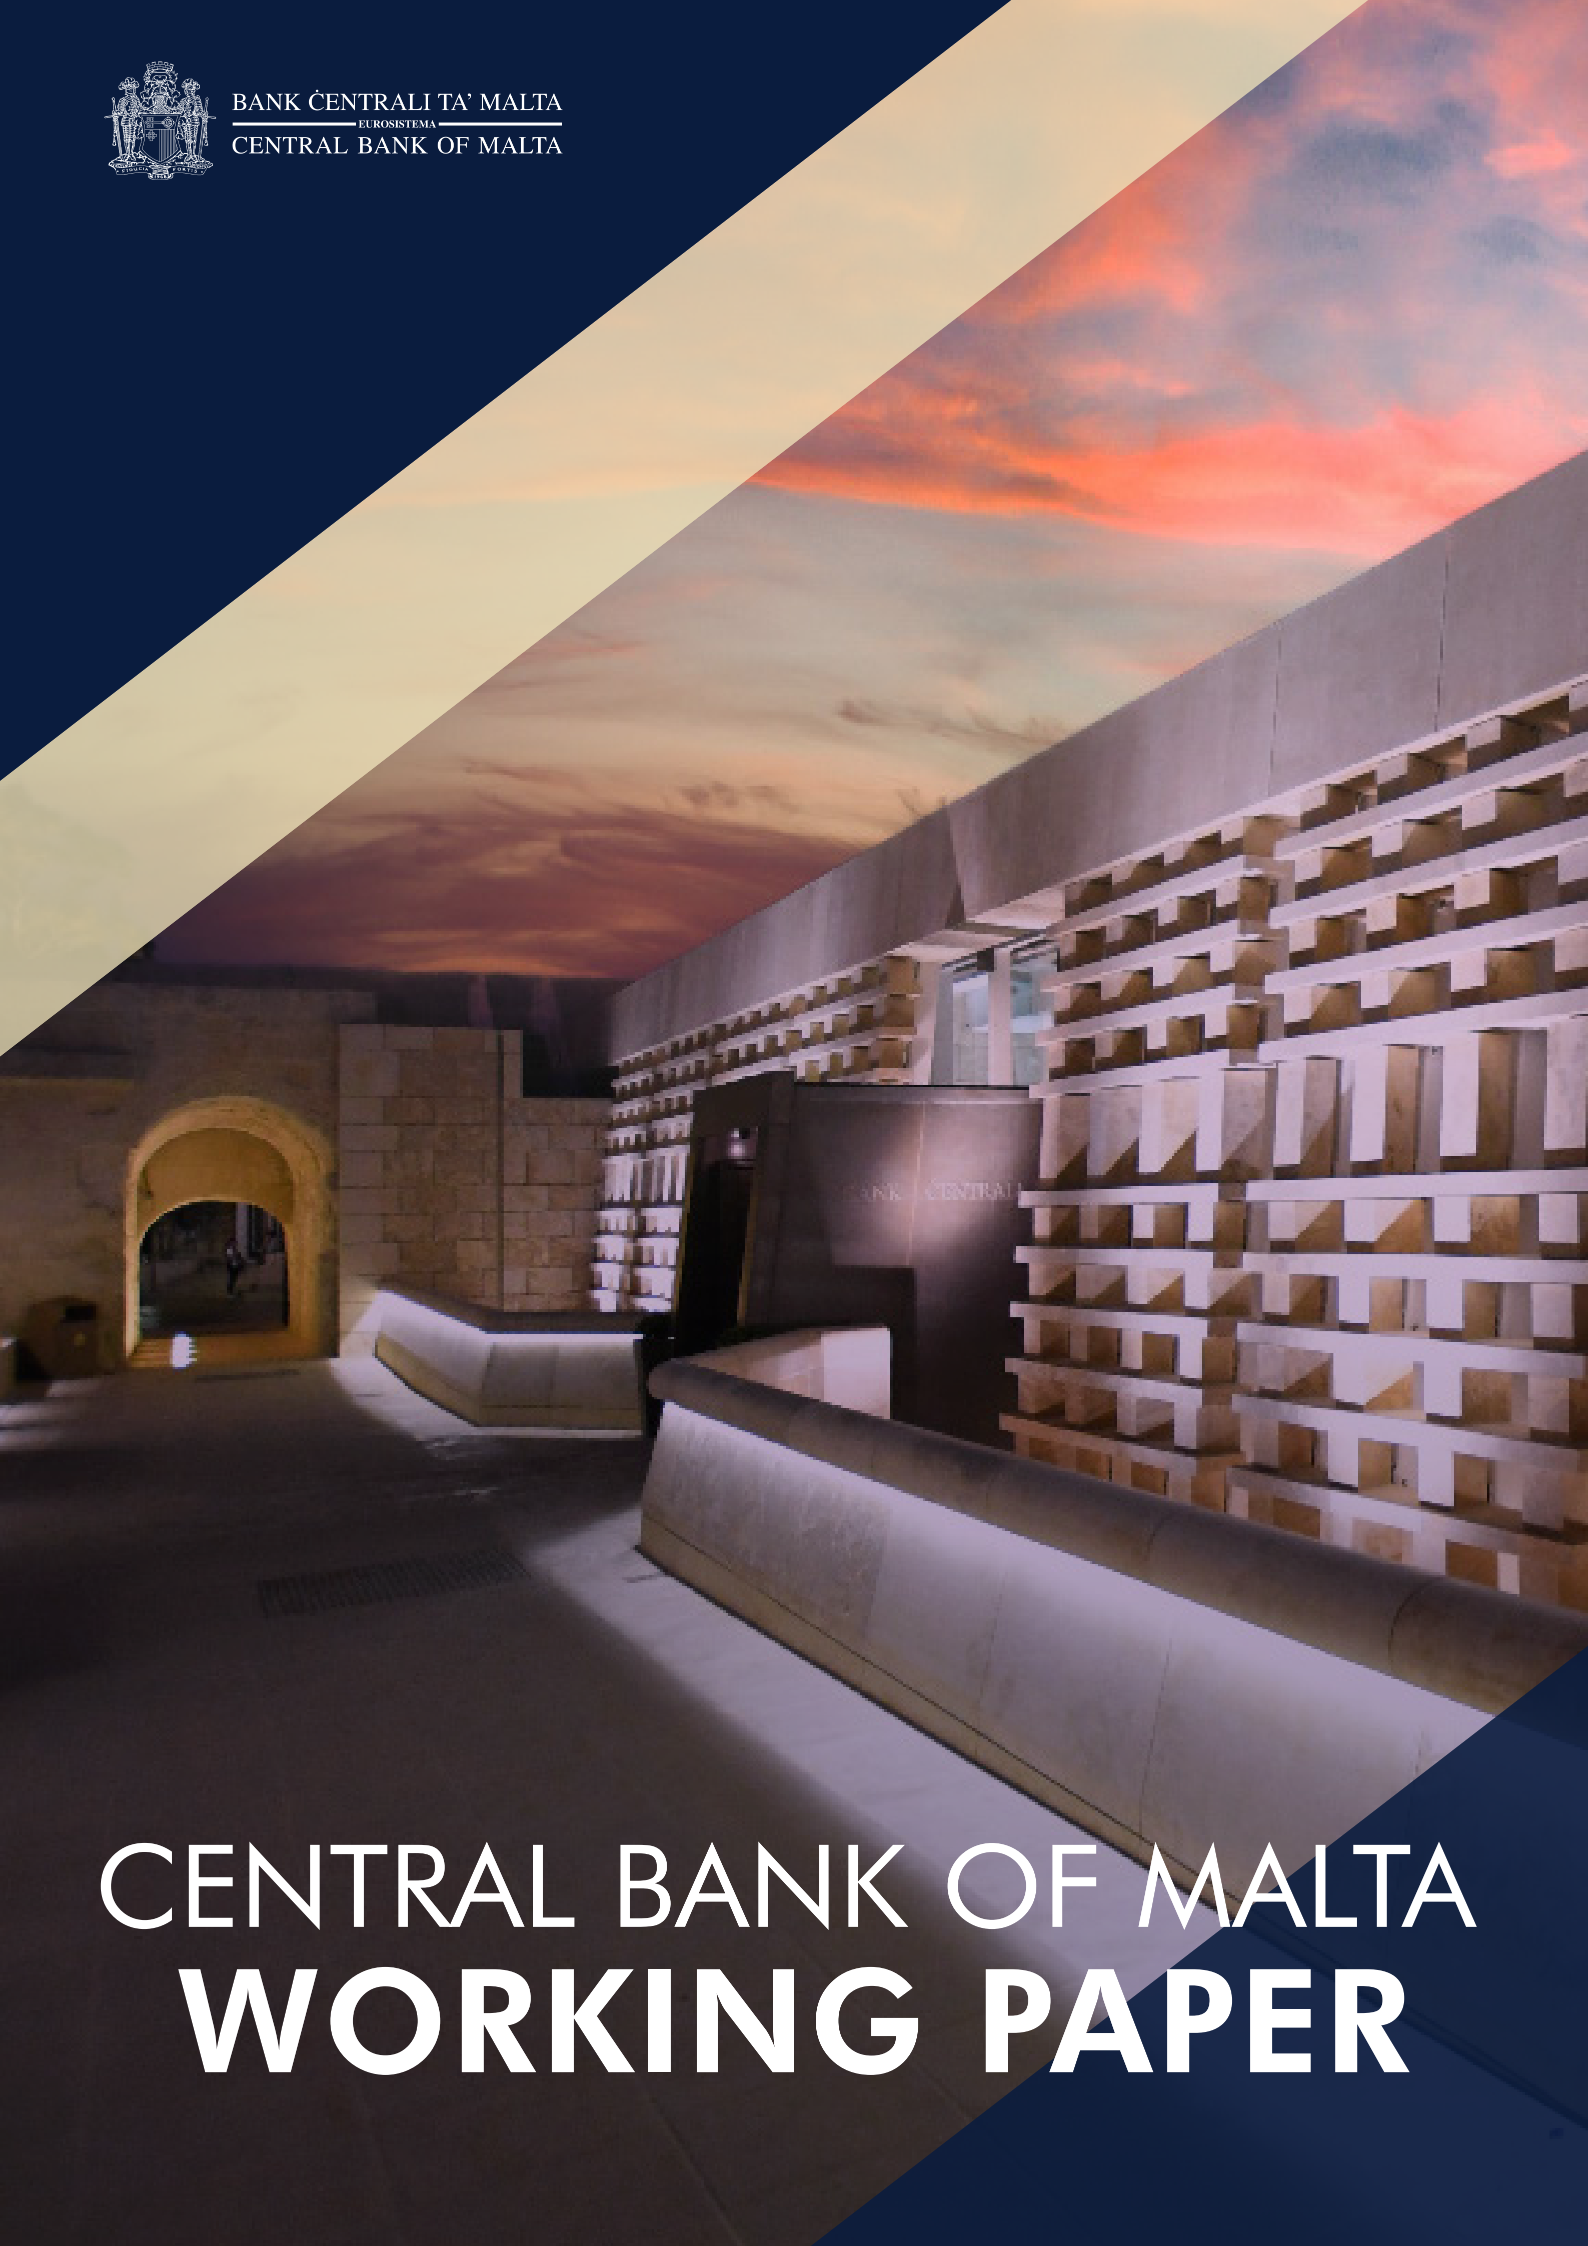

CENTRAL BANK OF MALTA
WORKING PAPER

Nowcasting the Maltese economy with a dynamic factor model*

Reuben Ellul† Germano Ruisi‡

WP/02/2022

*The model described in the Methodology section of this paper is based on the work carried out by Reuben Ellul over the last years in the effort to provide the Central Bank of Malta with a valid tool to timely forecast the economic developments taking place in the Maltese economy. Germano Ruisi extracted all the relevant results, set up the horse-race among competing forecasting models to assess the reliability of the dynamic factor model and was responsible for the writing stage of this paper. We wish to thank Davor Kunovac for kindly reviewing this paper and we also would like to thank Deputy Governor Alexander Demarco, Brian Micallef and Willam Gatt Fenech for their helpful suggestions. Any remaining errors are our own.

†Principal Economist, Economic Analysis Department, Central Bank of Malta until 2021.

‡Senior Research Economist, Economic Research Department, Central Bank of Malta.

Abstract

This paper describes a dynamic factor model for the Maltese economy. The model mainly serves as a tool to timely provide the Central Bank of Malta with nowcasts as well as short-term forecasts of the growth rate of the real gross domestic product, which in turn are used as an input in the forecasting process. Such forecasts reflect and incorporate the flow of information that periodically becomes available. Furthermore, the model can handle mixed frequencies that are likely to exist in large datasets used to summarise the Maltese economy and, as an additional advantage, it is able to deal with any path of missing data. This last feature is of crucial importance as data releases that are used to update the model do not take place in a synchronous way. The forecasting power of the dynamic factor model is compared with those of several other models available at the Central Bank of Malta. Overall, the results point towards a higher forecast accuracy of the dynamic factor model at very short horizons while, at longer ones, bayesian vector autoregressions appear to be more reliable.

JEL Classification: C53, E37

Keywords: dynamic factor models, missing data, nowcasting, forecasting, back-dating, forecasting horse-race.

Contents

1 Introduction **4**

2 Methodology **6**

2.1 The model 6

2.2 Data 8

3 Main results **12**

3.1 The smoothed series 12

3.2 Backdating 13

3.3 Nowcasting and forecasting 14

3.4 Forecasting performance 18

3.4.1 Competing Models 18

3.4.2 Forecasting the real GDP growth 21

4 Conclusion **26**

Appendix A Bayesian VARs Data **29**

Appendix B Forecasting GDP with the large BVARX GLP **30**

1 Introduction

This paper aims at developing a dynamic factor model (DFM) for the Maltese economy to nowcast and produce short-term forecasts of the real gross domestic product (GDP) growth rate and other main macroeconomic variables. The methodology, similar to those adopted by several other central banks and institutions around the world, exploits the recent advances in computational and statistical methods that have led to the development of automated real-time solutions that help solve the problem of achieving stable forecasts. For the Maltese case, the model is shown to return reasonable estimates of output growth.

The DFM exploits the information contained in medium and large sized datasets that have become increasingly available to researchers in recent years and the Maltese case does not represent an exception. However, econometric models have to be redesigned around this new, data rich environment. Traditional time series models are generally unable to incorporate many variables. If the number of parameters that needs to be estimated is large when compared with the sample size, these models will experience a degrees of freedom problem. This, in turn, will undermine the estimation accuracy as the parameter estimates will be likely characterised by very large variances. Moreover, traditional models cannot easily handle instances where there are more variables than observations or, even more difficultly, missing observations. The latter is a particular and recurrent characteristic of Maltese datasets where, for example, data collection and compilation did not start at the same time for all the variables of interest thus resulting in certain series being longer than others from a time perspective.[1]

One strand of literature this paper is mainly related to has tackled the issue of effectively using large amounts of data to extract policy-relevant information. The latter initiated with the first works on principal component analysis (PCA) (Hotelling, 1933) where it was understood that it is possible to project each data point onto only the first few principal components extracted from a large dataset in order to obtain lower-dimensional data while preserving as much of the data's variation as possible. Subsequently, in the field of economics, methods aiming at dealing with the curse of dimensionality experienced a remarkable growth especially when the first dynamic factor models were developed. Such early developments took place in the late 70s and early 80s where it was shown that they could be estimated by means of frequency-domain methods (Geweke, 1977; Sargent et al., 1977) or via maximum likelihood techniques using time-domain state-space methods (Engle and Watson, 1983) which make use of the Kalman filter (Kalman, 1960). In subsequent years, dynamic factor models experienced further improvements aiming at, on the one hand, better handling larger and larger datasets while, on the other, producing more consistent estimates of the parameters (Connor

[1] Vector autoregressions (VARs) estimated in a bayesian fashion can handle large amounts of data as shown in Ruisi and Borg (2018) where, *inter alia*, a large VAR was built in a similar way to Bańbura et al. (2010) in order to provide short-term forecasts for the Maltese economy. However, in its latest form available at the Central Bank of Malta (CBM), the latter still cannot deal with missing data nor with the continuous and asynchronous inflow of new information as will be discussed further down in this paper.

and Korajczyk, 1986; Forni and Reichlin, 1998; Forni et al., 2000; Stock and Watson, 2002; Bai and Ng, 2006). In addition, they found further employment in the creation of composite indices for a reliable assessment of economic conditions (Stock and Watson, 1988, 1989; Mariano and Murasawa, 2003).

The second strand of literature this paper is related to is the one focusing specifically on nowcasting. Even though the first applications were done in the field of meteorology, nowadays nowcasting models are applied to monitor the state of the economy in real-time as a proxy for official measures. Before nowcasting techniques were available, policy institutions used to rely on judgement combined with simple bridge equations (Baffigi et al., 2004) in order to obtain an early estimate of GDP. Bridge equations, by being essentially regressions relating quarterly GDP growth to one or a few monthly variables aggregated to quarterly frequency, are not able to handle the richer information that could be drawn from larger datasets and, as such, the precision of the estimates is questionable. The model here presented, instead, is designed in such a way to exploit the continuous inflow of all the available information that is relevant for forecasting purposes. The flexibility of this approach over alternative ones has contributed to the expansion of the literature. Seminal contributions in economics include Evans (2005) and Giannone et al. (2008) which provided a formal statistical framework to define a nowcasting process in detail. Key in this framework is to use a model with a state space representation with measurement equations linking observed series to a latent state process, and transition equations describing the state process dynamics. The state space representation, in turn, allows the use of the Kalman filter to obtain projections for both the state and the observed variables thus solving the recurrent issue of missing observations and/or mixed frequencies in large datasets.

By using the described techniques, therefore, the nowcasts obtainable from the DFM do not suffer from any problem related to incomplete data available at the CBM. Moreover, to keep track of the economic developments in real-time, they can be even updated on a daily basis depending on data published by the National Statistics Office (NSO), and other institutions providing data. Once the databank is updated, and the DFM estimated, the model updates its current quarter nowcast, and returns an updated forecast, discriminating between the impact of revisions to data as well as newly observed information. The model specifically presented in this paper is based on a number of DFM approaches used in the literature. It incorporates methodologies discussed in various studies, namely Bańbura et al. (2010) and Bańbura and Modugno (2012), and is closely related with dynamic-factor modelling approaches used by the European Central Bank or the Federal Reserve Bank as the one developed in Bok et al. (2018).

As far as the main results are concerned, the DFM shows a good forecasting power when compared with those of several other models currently used at the CBM. First of all, it is able to provide predictions that are at least as good as those obtainable from simple time series-based approaches, e.g., an autoregressive model of order one (AR(1)). Second, the DFM shows a higher forecast accuracy at very short horizons, i.e., one-quarter ahead, while bayesian vector autoregressions appear to be more reliable at longer ones, i.e., two or more quarters ahead.

The rest of the paper is structured as follows. Section 2 formally shows how the model is built and presents the rich dataset involved in the estimation. Section 3 shows the main results and highlights their usefulness for the Bank. Moreover, the latter section thoroughly compares the performance of the dynamic factor model vis-à-vis those of a number of competing models available at the CBM. Finally, section 4 concludes.

2 Methodology

This section aims at outlining the methodology adopted in this paper. The section proceeds by first presenting the model, how it is estimated and how it deals with missing data, and then by carefully describing the data used in the estimation.

2.1 The model

Abstracting from the estimation procedure which, by being able to deal with any pattern of missing data, represents the most computationally complex part of the model and is described in detail in Bańbura and Modugno (2012), the dynamic factor model presented in this paper has a relatively standard representation. The observation equation admits the following representation:

$$y_t = \Lambda_0 f_t + \Lambda_1 f_{t-1} + \cdots + \Lambda_L f_{t-L} + \epsilon_t \tag{1}$$

In equation (1), for each time $t = 1, \ldots, T$, $y_t = [y_{1,t}, y_{2,t}, \ldots, y_{n,t}]'$ is a n-dimensional vector of stationary time series that have been standardised in order to have a zero mean and a unit variance, while $f_t = [f_{1,t}, f_{2,t}, \ldots, f_{r,t}]'$, and similarly f_{t-l} with $l = 0, \ldots, L$, is a r-dimensional vector containing the unobserved common factors that need to be estimated. Moreover, $\epsilon_t = [\epsilon_{1,t}, \epsilon_{2,t}, \ldots, \epsilon_{n,t}]'$ is a n-dimensional vector of normally distributed idiosyncratic component that is uncorrelated with f_t at all leads and lags, i.e., $\epsilon_t \sim i.i.d.\mathcal{N}(0, R)$. Finally, Λ_l are a $n \times r$ matrices containing the factor loadings, i.e., the "slope coefficients" of each variable in y_t on each factor in f_t. The factors can be estimated from the whole, large and heterogeneous dataset or, alternatively, can be given an economic interpretation if extracted from a subset of more homogeneous series. Independently of the data they are extracted from, the factors f_t are assumed to follow a stationary VAR process of order p:

$$f_t = A_1 f_{t-1} + A_2 f_{t-2} + \cdots + A_p f_{t-p} + u_t \tag{2}$$

In the transition equation (2), the autoregressive coefficients are grouped in the $r \times r$ matrices $A_1, A_2, \ldots, A_p$. Finally, $u_t = [u_{1,t}, u_{2,t}, \ldots, u_{r,t}]'$ is a r-dimensional vector containing independent and identically distributed errors which follow a normal distribution, i.e., $u_t \sim i.i.d.\mathcal{N}(0, Q)$.

The estimation process builds on the Expectation-Maximisation (EM) algorithm whose initial version was developed in Dempster et al. (1977). The latter is able to overcome the problem of incomplete data and intractable likelihood. The essential idea of the algorithm is to write the likelihood as if the data were complete and to iterate between two steps: in the Expectation step the algorithm fills in the missing data in the likelihood, while in the Maximisation step it re-optimises this expectation in order to obtain the factors f_t as well as all the parameters in equations (1) and (2).

Given the available information set $\Omega_T \subseteq Y$, let us denote the joint-likelihood of y_t and f_t, with $t = 1, \dots, T$, by $l(Y, F|\theta)$, where $Y = [y_1, \dots, y_T]$, $F = [f_1, \dots, f_T]$ and $\theta = \{\Lambda_{l=0,\dots,L}, A_{i=0,\dots,p}, R, Q\}$.[2] The EM algorithm proceeds as follows:

1. Set $j = 0$ and initialise the algorithm by providing initial values for the factors and the parameters, i.e., $F_{j=0}$ and $\theta(0)$, and by filling in the missing data in Y;[3]

2. Set $j = j + 1$;

3. Expectation step: conditional on the data, the expectation of the log-likelihood is calculated using the estimates from the previous iteration, $\theta(j)$:

$$L(\theta, \theta(j)) = \mathbb{E}_{\theta(j)}[l(Y, F|\theta)|\Omega_T]$$

4. Maximisation step: the parameters are re-estimated by maximising the expected log-likelihood with respect to θ:

$$\theta(j+1) = \arg\max_{\theta} L(\theta, \theta(j))$$

5. Go to step 2 or otherwise stop the algorithm and get the parameter estimates θ if both the following conditions hold:

$$arg\max_{\theta} L(\theta, \theta(j+1)) - \arg\max_{\theta} L(\theta, \theta(j)) > 0$$

$$arg\max_{\theta} L(\theta, \theta(j+1)) - \arg\max_{\theta} L(\theta, \theta(j)) < threshold$$

In step 5, *threshold* represents a small number indicating that the algorithm can be stopped if the positive gain in the accuracy of the estimates is very small and below an arbitrarily chosen threshold. The latter is set to 10^{-5}.

The introduction of dynamics governing the evolution of the factors f_t allows, *inter alia*, to utilise the model for forecasting purposes. Suppose, after estimating the model and having obtained factors and all the parameters, to forecast H periods ahead. The law of motion in equation (2) can then be used to forecast the factors themselves:

$$\widehat{f_{t+h}} = \widehat{A_1} f_{t+h-1} + \widehat{A_2} f_{t+h-2} + \cdots + \widehat{A_p} f_{t+h-p} \tag{3}$$

After obtaining $\widehat{f_{t+h}}$ for $h = 1, \dots, H$ it is possible to use the observation equation in (1) to provide forecasts $\widehat{y_{t+h}}$ at the desired horizon:

[2] $\Omega_T \subseteq Y$ means that it is possible to have missing data in y_t. In case of complete data $\Omega_T = Y$.

[3] For $j = 0$ the missing data are filled in through interpolations with the available ones while the initial values of the factors are obtained by means of simple PCA. Please refer to Bańbura and Modugno (2012) for a detailed description of how to initialise the algorithm.

$$\widehat{y_{t+h}} = \widehat{\Lambda}\widehat{f_{t+h}} \tag{4}$$

As a short-term forecasting tool, H is set to one year. It is common knowledge that increasing H would make the forecast accuracy decrease. In the baseline version, the number of factors is set to four as described in subsection 2.2.

2.2 Data

This subsection describes the data used in the estimation. Specifically, table 1 lists all the variables entering the model and provides information regarding their units of measure, transformation, frequency and category. Data collection is done in such a way to have, unless there are missing data, January 2000 as the first available observation. As opposed to that, the last observation is completely series-specific and purely dependent on data availability.

The table also describes which series contribute to the extraction of the four chosen factors and, as a consequence, which factors they load onto. This information is provided in the last four rightmost columns where it is shown how the factors can be given an economic interpretation according to the variables used to extract them. To be more precise, all the 38 variables listed in table 1 contribute to the estimation of the *Global* factor which summarises the general economic developments in the Maltese economy. The remaining three are extracted from subsets of the whole dataset: *Soft* from survey data (e.g., economic sentiment indicator, services, consumer, retail and construction sentiment indicators, etc.), *Real* from data on the real economy (e.g., real GDP, industrial production, exports, etc.) and, finally, *Labour* from data regarding the labour market (e.g., unit cost of labour, unemployment level and rate, etc.).

As an example, the real gross domestic product enters the database in levels as chained million euros on a quarterly frequency, belongs to the *"National accounts"* category, is then transformed on a year-on-year growth rate and loads onto the global and the real factors. As opposed to that, the harmonised index of consumer prices (HICP) enters as an index on a monthly frequency, is then transformed into inflation rate by calculating its year-on-year growth rate and loads onto the global factor only.[4]

An important feature of the dynamic factor model is that the number of factors, as well as the information from which they are extracted, can be modified according to the forecaster's needs. As an example, which is left for future research, it might be interesting to see whether a specific set of data can be useful in estimating a factor that helps better predict inbound tourism. Moreover, for each "class" it is possible to extract more than one factor in a bid to improve the forecast accuracy at the cost, however, of an increased computational burden. These possibilities are explored section 3.

[4]Note that the transformations used are not those typically used in literature where quarter-on-quarter growth rates are matched with month-on-month ones relying on the technique described in Mariano and Murasawa (2003). Notwithstanding, it is important to emphasise that the usage of year-on-year growth rates does not undermine the precision of the estimates and, especially in the specific case of real GDP growth, this is a convenient choice as it directly returns the predicted variable with the desired transformation.

ID	Name	Frequency	Units	Transformation	Category	Global	Soft	Real	Labour
CREDIT	Credit	m	EUR, millions	YoY Percent Change	Finance	Yes	No	No	No
DEPARTOUR SA	Departing Tourists	m	Thousands of persons	YoY Percent Change	Services	Yes	No	Yes	No
ESI	Economic Sentiment Indicator	m	Index	YoY Percent Change	Surveys	Yes	Yes	No	No
GOP	Gainfully Occupied Population	m	Thousands of persons	YoY Percent Change	Labour	Yes	No	No	Yes
HICP SA	HICP	m	Index	YoY Percent Change	Prices	Yes	No	No	No
HICPXFE SA	Core HICP	m	Index	YoY Percent Change	Prices	Yes	No	No	No
INDPROD SA	Industrial Production	m	Index	YoY Percent Change	Manufacturing	Yes	No	Yes	No
MUVI	Imports UVI	m	Index	YoY Percent Change	International Trade	Yes	No	Yes	No
MVAL	Imports Customs	m	EUR, millions	YoY Percent Change	International Trade	Yes	No	Yes	No
PPI	Producer Price Index	m	Index	YoY Percent Change	Prices	Yes	No	No	No
RETRADE SA	Retail Trade	m	Index	YoY Percent Change	Real economy	Yes	No	Yes	No
RPI SA	RPI	m	Index	YoY Percent Change	Prices	Yes	No	No	No
SEATCAP SA	Seat Capacity MIA	m	Number of seats	YoY Percent Change	Services	Yes	No	Yes	No
TAX SA	Taxation	m	EUR, millions	YoY Percent Change	Real economy	Yes	No	No	No
UNEMP	Unemployment	m	Number of people	YoY Percent Change	Labour	Yes	No	No	Yes
UNRATE	Unemployment Rate	m	Percentage	YoY Difference	Labour	Yes	No	No	Yes
XUVI	Exports UVI	m	Index	YoY Percent Change	International Trade	Yes	No	Yes	No

ID	Name	Frequency	Units	Transformation	Category	Global	Soft	Real	Labour
XVAL	Exports Customs	m	EUR, millions	YoY Percent Change	International Trade	Yes	No	Yes	No
BCI Industry	Industry Confidence	m	Index	Levels	Surveys	Yes	Yes	No	No
GBPEUR	GBP - EUR	m	Exchange rate	YoY Percent Change	Finance	Yes	No	No	No
USEUR	US - EUR	m	Exchange rate	YoY Percent Change	Finance	Yes	No	No	No
BDI	Baltic Dry Index	m	US Dollars	YoY Percent Change	Prices	Yes	No	No	No
IFOBC	IFO Business Confidence Germany	m	Index	Levels	Surveys	Yes	Yes	No	No
OIL	Brent oil price	m	US Dollars	YoY Percent Change	Prices	Yes	No	No	No
DEPOSITS	Total residents deposits	m	EUR, millions	YoY Percent Change	Finance	Yes	No	No	No
MSE	Malta stock exchange index	m	Index	YoY Percent Change	Finance	Yes	No	No	No
SERVICES	Services industry sentiment	m	Index	Levels	Surveys	Yes	Yes	No	No
CONSUMER	Consumers sentiment	m	Index	Levels	Surveys	Yes	Yes	No	No
RETAIL	Retail sentiment	m	Index	Levels	Surveys	Yes	Yes	No	No
CONSTRUCTION	Construction sentiment	m	Index	Levels	Surveys	Yes	Yes	No	No
EONIA	EONIA rate	m	Percentage	First Difference	Finance	Yes	No	Yes	No
TBILL 1MONTH	1-month Tbill	m	Percentage	First Difference	Finance	Yes	No	Yes	No
MGS 10YR	10-year government bond	m	Percentage	First Difference	Finance	Yes	No	Yes	No
SEMICON BILLINGS	WSTS Semiconductors Billings (Worldwide)	m	US Dollars	YoY Percent Change	Real economy	Yes	No	No	No

ID	Name	Frequency	Units	Transformation	Category	Global	Soft	Real	Labour
GDP	GDP	q	Chained EUR, millions	YoY Percent Change	National Accounts	Yes	No	Yes	No
ULC	ULC	q	Index	YoY Percent Change	National Accounts	Yes	No	No	Yes
HPI CBM	Advertised house prices	q	Index	YoY Percent Change	Prices	Yes	No	Yes	No
MOTOR VEHICLES	Motor vehicle stock	q	Number	YoY Percent Change	Real economy	Yes	Yes	Yes	No

Table 1: Whole dataset used for the estimation of the dynamic factor model. The dataset specifies, for each variable, ID code, name, frequency, units of measure, transformation, category and factors to load onto

As an additional remark, it is important to mention that the way the factors are extracted, as outlined in the last four rightmost columns of table 1, implies that the latter are not identified and, therefore, they are not suitable for properly conducting structural analyses. This is the case as the datasets they are extracted from often overlap and, as such, the factors are likely to exhibit a certain degree of correlation which implies that they are not orthogonal to one another. As an example, variables like real GDP, industrial production as well as all the series related to international trade contribute to the extraction of both the global and the real factor.[5] Nevertheless, the lack of identification of the extracted factors does not undermine the possibility of using the model for nowcasting or forecasting purposes as only correlations between regressor and depended variable (and not causation from the first to the second) is required.

3 Main results

This section shows the main pieces of output obtainable from the DFM described in section 2. Specifically, the smoothed series and the backcasting ability of the model are obtained by using a vintage of available data as of the 12^{th} of October 2021. Moreover, this section shows an example of how the forecast and/or nowcast of the 2019Q4 real GDP growth can be updated. This exercise is conducted with a number data vintages reflecting the timely inflow of information. Finally, this section presents how the DFM's performance compares with those of several other competing models available at the CBM.

3.1 The smoothed series

The first relevant piece of output of this model is the smoothed series relative to each of the variables contained in y_t. This is particularly relevant in presence of missing data or if one needs, for example, a monthly version of a series that originally has a quarterly frequency. As a matter of fact, a quarterly series can be seen as a monthly one which is observed only once every three months.[6] The model, thus, interpolates the series by exploiting all the information contained in the dataset y_t. This implies that the more detailed the dataset (especially in terms of informativeness of the variables at monthly frequency) the higher the precision with which the smoothed series is estimated. This can be relevant when trying to improve the reliability of an estimation by increasing the number of observations or if the researcher needs to increase the frequency of their data

[5] The topic of factors identification has been widely tackled and is well established in the literature. For a non exhaustive review please refer to Stock and Watson (2002), Bai and Ng (2013), Stock and Watson (2016), Williams (2020) and references therein.

[6] The need of higher frequency measures relative to variables of interest, such as economic output, has become increasingly relevant over the last years because of the COVID-19 pandemic. Notable examples are weekly indicators such as the weekly tracker of economic activity developed by the Organisation for Economic Cooperation and Development (Woloszko, 2020), the Weekly Economic Index developed by the Federal Reserve Bank of New York (Lewis et al., 2020a,b), the Weekly Activity Index developed by the Bundesbank (Eraslan and Götz, 2020) or the Weekly State-Level Economic Condition Indices for the 50 U.S. states (Baumeister et al., 2021).

in a bid to dig deeper in the nature of the data analysed.

Figure 1 shows how it is possible to obtain a monthly measure of the year-on-year real gross domestic product growth:

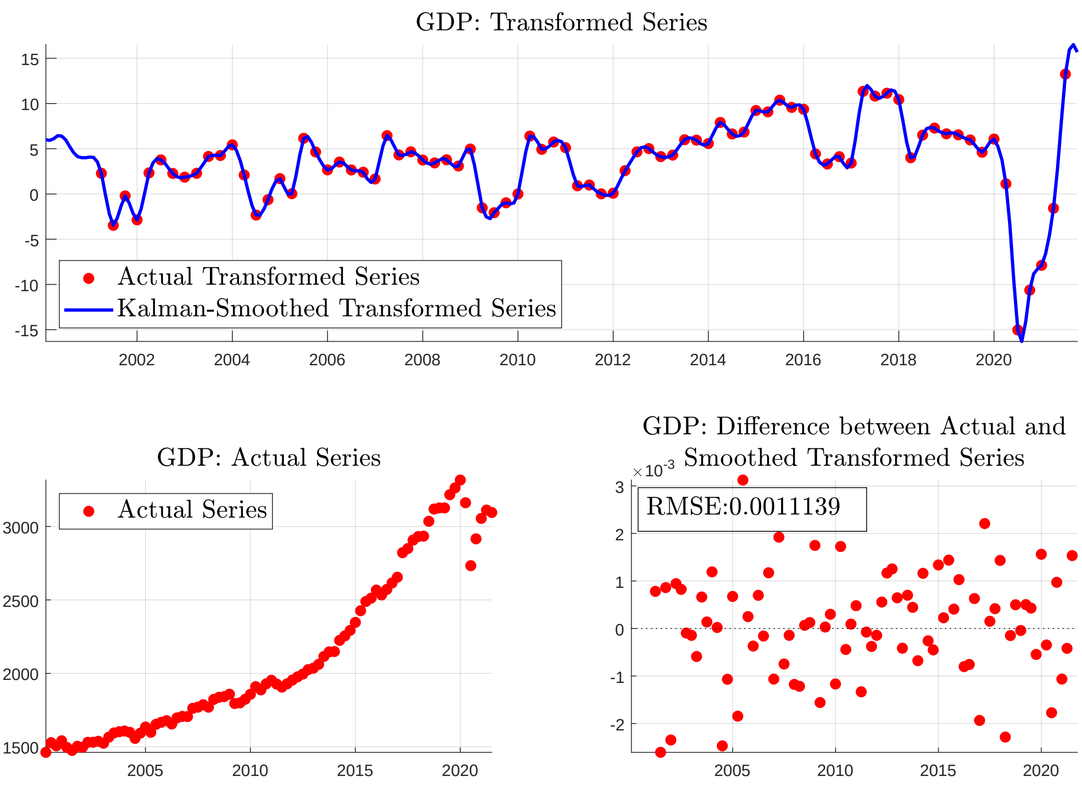

Figure 1: Real GDP - Monthly smoothed series (top panel), actual quarterly series (bottom left) and error between the two (bottom right)

The top panel shows how the smoothed series is able to closely capture the quarterly observations as well as to provide a reasonable estimate for the months in which no GDP data are available. To put it simply, as previously explained, the information contained in the whole dataset is crucial for the DFM to join the red dots. The bottom left panel shows the untransformed series while the bottom right one evidences how the interpolation error is particularly small and centred around zero. Such an error is calculated as the difference between the actual series and the interpolated one during the quarters when the data were released.

3.2 Backdating

One of the things that makes this model particularly useful is the possibility to backdate those variables in y_t that are not available from the beginning of the sample. Similarly to the monthly interpolation of the quarterly series, the backdating process is made possible by exploiting the information contained in the whole dataset. As an example, we modify the dataset by discarding the first four quarterly observations of the real GDP growth (i.e., those relative to 2001) and try to backdate them. The results are depicted in figure 2:

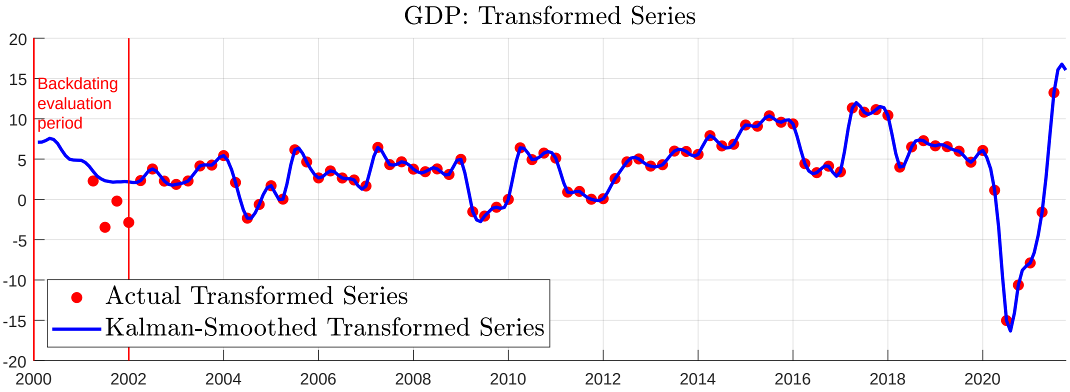

Figure 2: Real GDP growth - Actual observations and backdated series

Figure 2 demonstrates how, for the backdating evaluation period, the model is hardly able to get close to the observed series. There might be a number of reasons that explain why this happens. First, this might be due to the more limited number of variables in y_t (especially at monthly frequency) that cover the evaluation period considered as some of them were not even observed. In this case, the factors are extracted from a smaller dataset thus limiting their precision. Second, the period considered can be (through the lens of the model) particularly volatile. Third, we might be in presence of "badly behaved" data characterised by non-normality, e.g., skewness and fat tails, that the model in subsection 2.1 is not able to deal with. Fourth, the choice of the number of factors in f_t and the order p of their law of motion can also play a small role. All of this considerations should be kept in mind when in need to improve a forecasting tool kit.

Finally it is important to highlight that the backcasting ability of the DFM does not solve only the problem of variables that are compiled and/or collected starting from a later date with respect to the beginning of the dataset. More precisely, the model is able to solve any case of missing data by proving estimates of what the variable would have been like if it had been compiled and/or collected at each point in time without any interruption.

3.3 Nowcasting and forecasting

This subsection is dedicated to the most important feature of this model, i.e., the possibility to provide timely forecasts and nowcasts of the variables of interest. The case of nowcasting can be seen as the necessity of forecasting the value of a series which has not been published yet but in presence of a dataset that can comprise even more recent data releases. Put it differently, what the DFM in this paper is able to do is to predict the present, the near future and the near past of an economic indicator that still needs to be published.

Suppose to have several vintages of a dataset with the last one dating 2020M01 and to have the necessity to forecast (nowcast) the annual GDP growth rate for 2019Q4 which is not available yet. This becomes an exercise whose answer can be achieved in a similar way to what discussed, for the case of missing data, in subsections 1 and 2. The case of forecasting several steps ahead, instead, is tackled as discussed in subsection 2.1.

Independently of what the model is used for, nowcasting or forecasting at a higher horizon, it is able to provide evidence of those news that contributed to the forecast updates. Such news can be data revisions or new data releases that take place from one vintage to the next. The top panel of figure 3 shows the case of the example hinted above in which the annual growth rate of real GDP for 2019Q4 has been forecast and nowcast every time a new piece of information became available between the period that goes from the 28^{th} of February 2019 to the 3^{rd} of February 2020 as outlined in table 2.[7]

Date	Main data releases/additions included in the vintage
28/02/2019	Starting vintage with observations until December 2018
01/03/2019	Prices, financial, sentiment indicators
15/03/2019	GDP, prices, financial, tourism, sentiment indicators
31/03/2019	Financial
15/04/2019	Tourism, sentiment indicators
30/04/2019	Prices, financial, sentiment indicators
15/05/2019	Prices, tourism
14/06/2019	GDP, prices, financial, tourism, sentiment indicators
15/07/2019	Financial, tourism, sentiment indicators
02/09/2019	Prices, financial, tourism, sentiment indicators
17/09/2019	GDP
01/10/2019	Financial, sentiment indicators
30/10/2019	Financial, tourism, sentiment indicators
03/12/2019	Financial, tourism, sentiment indicators
03/02/2020	GDP, prices, financial, tourism, sentiment indicators

Table 2: Dates in which the 2019Q4 real GDP growth was forecast and data releases included in the new vintage

[7] For the sake of space the table lists only the most relevant additions that were done from one vintage to the other. More detailed information, as well as data vintages, is available upon request.

For ease of reading, the information additions are grouped into five classes: GDP, prices (HICP and HICP that excludes energy and food), financial (interest and exchange rates), tourism (inbound tourism and seat capacity) and sentiment indicators (economic sentiment indicator for Malta as well as several survey data relating to the construction, retail and services sectors).

Such a sequence of forecast updates is then compared with the actual value released by the Maltese NSO on the 28*th* of February 2020. It is evident how, as time goes by and the information set becomes richer, the forecast updates get closer and closer to the actual GDP growth rate for 2019Q4.

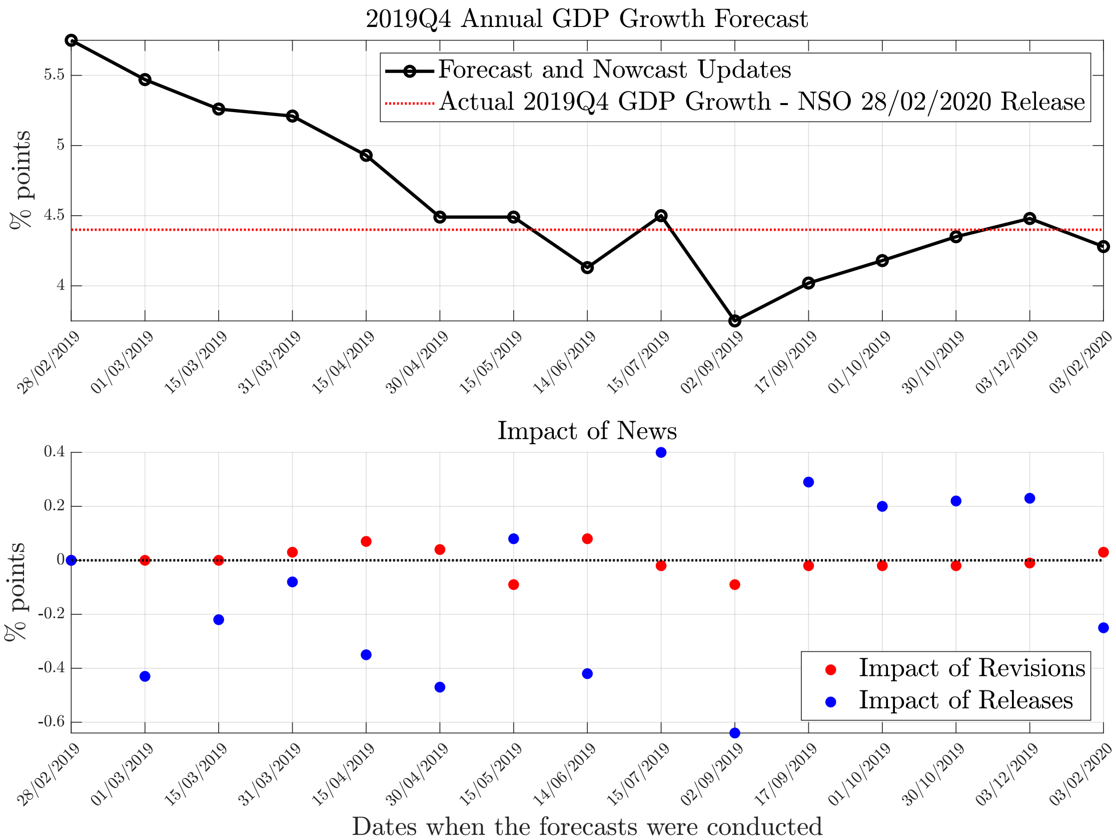

Figure 3: Forecast and nowcast updates (top panel) and news impact (bottom panel)

As an additional piece of information, the bottom panel of figure 3 shows what induces such forecast updates by decomposing the latter into the impact of data revisions and that due to new data releases. For example, the reason why the 2019Q4 forecast conducted on the 3^{rd} of February 2020 is smaller than the preceding one done on the 3^{rd} of December 2019 is a negative impact of new data releases (-0.25%) and a very small but positive one of data revisions ($+0.03\%$).

Finally, the model is able to decompose the change in forecast into the contribution of each variable in the dataset. Figure 4 shows the output relative to the forecast performed on the 3^{rd} of February 2020:

```
Nowcast Update: February 03, 2020
Nowcast for GDP (Year over Year Percent Change), 2019:Q4

  Nowcast Impact Decomposition
  Note: The displayed output is subject to rounding error

                  Dec 03 nowcast:                  4.49
      Impact from data revisions:       0.03
       Impact from data releases:      -0.25
                                     +________
                    Total impact:      -0.22
                  Feb 03 nowcast:                  4.28
```

Nowcast Detail Table

	Forecast	Actual	Weight	Impact
CREDIT	7.0633	6.8127	-0.035885	0.008994
DEPARTOUR_SA	6.4229	8.5339	0.014113	0.029793
ESI	-6.9988	-15.974	0.009259	-0.083102
GOP	5.7723	5.3725	0.064159	-0.025649
HICP_SA	1.422	1.322	-0.081255	0.0081214
HICPXFE_SA	1.3401	1.0939	-0.053252	0.013112
INDPROD_SA	2.8978	0.30731	0.0010999	-0.0028494
MUVI	-2.9702	-6.5429	-0.0027357	0.0097738
MVAL	12.189	6.4644	-0.0024084	0.013788
PPI	1.3753	1.6032	-0.0024494	-0.00055801
RETRADE_SA	3.4942	2.9875	0.017123	-0.0086757
RPI_SA	1.4072	1.1858	-0.10694	0.023676
SEATCAP_SA	6.7675	14.258	0.027556	0.20642
UNEMP	-15.788	-6.9688	-0.025342	-0.22349
UNRATE	-0.20957	-0.21859	-0.1461	0.0013181
XUVI	1.0485	1.7102	-0.0012393	-0.00081996
XVAL	7.0168	-0.086843	-0.0024134	0.017144
BCI_Industry	-6.4031	-7.8	0.015328	-0.021411
GBPEUR	-2.4299	-5.6174	-0.010375	0.033071
USEUR	-2.5614	-2.3805	0.00053683	9.7099e-05
BDI	21.306	3.414	0.00094716	-0.016947
IFOBC	95.205	96.3	0.1513	0.16565
OIL	-13.963	19.855	0.0021295	0.072013
DEPOSITS	4.228	3.8474	0.014395	-0.0054788
MSE	6.83	3.178	0.0048918	-0.017865
SERVICES	28.203	-6.5	0.0098461	-0.34169
CONSUMER	4.616	5.5	0.026877	0.023758
RETAIL	7.5227	-20.6	0.0022519	-0.063331
CONSTRUCTION	19.307	19.4	0.0044753	0.00041648
EONIA	0.0031121	-0.01	-0.11574	0.0015176
TBILL_1MONTH	-0.0064332	-0.04	-0.09197	0.0030871
MGS_10YR	-0.0085626	0.028063	-0.46986	-0.017209
SEMICON_BILLINGS	-9.5627	-8.6784	0.0053715	0.0047496
GDP	4.7419	3.3434	0.04294	-0.060049
ULC	1.8374	1.8058	-0.016936	0.00053475
MOTOR_VEHICLES	3.1756	3.1615	0.13508	-0.0019063

Figure 4: Real GDP growth - Forecast and nowcast output table

The rightmost column in figure 4 lists the impact of news relative to each of the variables. The latter is calculated as the product of the series' forecast error and its model-implied weight as described in Bańbura and Modugno (2012). In the specific case, the change in forecast from 4.49% to 4.28% is due to negative surprises outweighing positive ones. Specifically, the bulk of positive news updates came from seat capacity and the German business confidence index. Most of the negative surprises can be traced to unemployment and the services industry sentiment developed by the European Commission.

3.4 Forecasting performance

This subsection presents a forecasting horse-race between the DFM described in section 2 and a number of competing models that are currently available at the CBM to predict the growth rate of the real GDP. Specifically, the horse-race is run between a number of versions of the DFM, a simple autoregressive model of order one, and a number of vector autoregressions estimated in a bayesian fashion. The dataset used goes from January 2000 to August 2021. In order to create a number of vintages necessary to assess the performance, an entire month of observations is excluded at a time. This allows to exclude the effect that data revisions have on the forecast errors and, therefore, to assess the pure forecasting performance of each model, i.e., its ability to predict the future value of an indicator independently on the revisions that might be done on its past values (Bańbura and Modugno, 2012).[8] All the models are estimated by keeping the initial observation fixed and by adding one observation at a time as required by the forecasting exercise. The evaluation period goes from 2014Q4 to 2021Q2 thus allowing to assess the models' accuracy over 22 to 27 GDP observations depending on the forecast horizon chosen. The latter ranges from one to four steps ahead ($h = 1$ to $h = 4$) for all the models. In addition, the DFMs are also tested on the previous quarter's GDP forecast ($h = 0$).

3.4.1 Competing Models

Dynamic factor models

For all the versions of the DFM used here, the GDP forecasts for March, June, September and December are respectively conducted:

- on the month the GDP is referred to (DFM-M0), i.e., March, June, September and December;
- one month after that the GDP is referred to (DFM-M1), i.e., April, July, October and January;

[8]In addition, retrieving real time vintages back to December 2014 for each of the variables discussed in subsection 2.2 is a nearly impossible task.

- two months after that the GDP is referred to (DFM-M2), i.e., May, August, November and February.[9]

As a matter of fact, the three above can be seen as competing models using different information sets. The main aim of this subsection is to assess the forecast accuracy of several versions of the DFM as described in section 2 by not only differing on the basis of the amount of data they contain but also by the number of factors and/or the lag order p of the transition in equation 2. To this end, the four versions are:

1. a baseline version featuring the four factors as described in table 1 and with $p = 1$ (DFM-M);
2. a version with the four factors and $p = 2$ (DFM-M p2);
3. a version with $p = 1$ but with five factors with two of which extracted from the whole dataset, i.e., $f_t = [f_{Global_1,t}, f_{Global_2,t}, f_{Soft,t}, f_{Real,t}, f_{Labour,t},]$ (DFM-M 2gf);
4. a version with five factors as in version 3 and $p = 2$ (DFM-M p2 2gf);.

Versions 2 to 3 are meant to improve upon the baseline version. Specifically, version 2 aims at better estimating the evolution of the factors by increasing the lag order of the transition equation 2. Version 3 tries to extract more information from the whole dataset by estimating an additional factor on which the variables load onto in the observation equation 1. Finally, version 4 combines versions 2 and 3.

Bayesian vector autoregressions

The Bayesian vector autoregressions considered here are all estimated in the spirit of Bańbura et al. (2010) and differ by the way the hyperparameters in the priors are set up. All of them feature a block of exogenous variables and the number of endogenous ones is either three for the *small* or 17 for the *medium* size.[10] Each BVARX has the following representation:

$$Y_t = A + \sum_{l=1}^{L} B_l Y_{t-l} + \sum_{l=1}^{L} C_l X_{t-l} + E_t \tag{5}$$

In equation 5, Y_t and Y_{t-l}, with $l = 1, \ldots, L$ are $N \times 1$ vectors containing current and past values of the endogenous variables (either three or 17), X_{t-l} are $K \times 1$ vectors of lagged values of the exogenous ones and, finally, E_t is a $N \times 1$ vector of reduced form residuals that are assumed to be normally distributed with zero mean and constant variance, i.e., $E_t \sim \mathcal{N}(0_{N\times 1}, \Sigma)$. Finally, B_l and C_l represent $N \times N$ and $N \times K$ matrices respectively containing the slope coefficients associated with the lagged values of the endogenous and the exogenous variables.

[9] The assumption that the previous quarter's GDP growth estimate is carried out after two months is in line with the observation that in Malta the GDP release takes place in the very last days of the second month after that the GDP refers to.

[10] For the sake of space all the variables entering the VARs are listed in appendix A.

All of the BVARXs are estimated by using priors that are set up through dummy observations which depend on how the following hyperparameters are chosen:[11]

- L representing the number of lags;
- λ governing the overall tightness degree of the slopes relative to the lagged endogenous variables;
- $\epsilon_{Intercept}$ for the tightness of the intercepts;
- ϵ_X governing the tightness degree of the slopes relative to the lagged exogenous variables.

The BVARXs are:

1. BVARX 2018: BVARXs set up as in Ruisi and Borg (2018):
   - Small: $L = 2$, $\lambda = 0.3$, $\epsilon_{Intercept} = 0.1$ and $\epsilon_X = 1000$
   - Medium: $L = 1$, $\lambda = 0.2$, $\epsilon_{Intercept} = 0.1$ and $\epsilon_X = 1000$
2. BVARX 2021: BVARXs whose hyperparameters are set up by minimising the 1-step ahead root mean squared forecast errors of the real GDP growth over a pre-evaluation period (2013Q1 to 2014Q4):
   - Small: $L = 2$, $\lambda = 0.4$, $\epsilon_{Intercept} = 0.1$ and $\epsilon_X = 10$
   - Medium: $L = 1$, $\lambda = 0.4$, $\epsilon_{Intercept} = 10$ and $\epsilon_X = 1000$
3. BVARX GLP: BVARXs whose hyperparameters are fine-tuned at each step as in Giannone et al. (2015) in such a way to minimise the marginal likelihood (for both sizes, i.e., small and medium).

Autoregressive model

Finally, the last competing model is a simple autoregressive model of order one for the growth rate of the real GDP:

$$GDP_t = \alpha + \beta \; GDP_{t-1} + \epsilon_t \tag{6}$$

In equation 6, α and β respectively represent the intercept and the slope coefficient while the residuals are assumed to be homoscedastic and not serially correlated, i.e., $\epsilon \sim \mathcal{N}(0, \sigma^2)$. This model is estimated with simple ordinary least squares and is the simplest one among the competing ones. As such, together with a random walk, this autoregressive model serves as a comparison for the more complex DFMs and BVARs.

[11] For a detailed discussion about how to parametrise the priors implemented with dummy observations please refer to Bańbura et al. (2010) and, for the case of Malta, to Ruisi and Borg (2018).

3.4.2 Forecasting the real GDP growth

Table 3 shows the absolute root mean square forecast error associated with the four versions of the DFMs as described previously. For each of the four, we distinguish according to the last monthly observation available, i.e., DFM-M0, M1 and M2. The forecast horizons go from zero to four.

Model	h = 0	h = 1	h = 2	h = 3	h = 4
DFM-M0	3.12	5.68	7.01	7.47	8.00
DFM-M1	2.63	4.22	6.81	7.53	7.95
DFM-M2	**2.56**	3.40	6.92	7.60	7.96
DFM-M0 p2	3.00	5.32	6.42	6.77	7.42
DFM-M1 p2	2.93	3.77	6.57	7.16	7.56
DFM-M2 p2	2.57	**3.29**	6.75	7.26	7.25
DFM-M0 2gf	2.95	5.87	6.97	7.35	7.94
DFM-M1 2gf	2.69	4.38	6.76	7.42	7.88
DFM-M2 2gf	2.67	3.56	6.74	7.39	7.70
DFM-M0 p2 2gf	2.95	5.42	6.45	**6.69**	7.19
DFM-M1 p2 2gf	2.61	4.60	**6.35**	6.88	7.33
DFM-M2 p2 2gf	2.68	3.66	6.53	7.02	**6.99**

Table 3: Absolute root mean squared forecast errors relative to the performances of the four DFMs when forecasting at several forecast horizons - Full evaluation period

Overall, three results stand out. First, the RMSFEs tend to become higher as the forecast horizon grows large. Second, for each of the four versions, the forecast accuracy tends to improve as the information included becomes richer as shown by the decreasing RMSFEs as moving from DFM-M0 to DFM-M2. Third, there is not a version that consistently outperforms all the others. Specifically, at horizon $h = 0$ the best performance belongs to the first version that contains data up to two months after the quarter the GDP release refers to (DFM-M2). At horizon $h = 1$ the best model appears to be the one with two lags in the transition equation and two additional months of data (DFM-M2 p2). For the remaining horizons the best performance is achieved by the most complete of the four versions, i.e., the one with $p = 2$ and an additional global factor. More precisely, they respectively belong to DFM-M1 p2 2gf at $h = 2$, DFM-M0 p2 2gf at $h = 3$ and DFM-M2 p2 2gf at $h = 4$.

The forecast evaluation period considered above contains also the years characterised by the slowdown and recovery induced by the current pandemic. In order to further shed light on the accuracy of the DFMs we repeat the same exercise by calculating the RMSFEs by focusing only on the period prior to the COVID-19 outbreak. Table 4 shows the results when the evaluation period is restricted to the forecasts conducted between December 2014 to February 2020.

Model	h = 0	h = 1	h = 2	h = 3	h = 4
DFM-M0	2.56	3.33	5.66	6.54	7.07
DFM-M1	**2.46**	3.12	5.62	6.62	7.1
DFM-M2	**2.46**	**3.05**	5.58	6.57	7.02
DFM-M0 p2	2.58	3.33	5.61	6.49	7.03
DFM-M1 p2	2.48	3.10	5.54	6.52	6.98
DFM-M2 p2	2.48	3.06	**5.53**	**6.50**	**6.81**
DFM-M0 2gf	2.64	3.38	5.68	6.59	7.14
DFM-M1 2gf	2.56	3.17	5.64	6.65	7.18
DFM-M2 2gf	2.55	3.14	5.61	6.62	7.11
DFM-M0 p2 2gf	2.68	3.40	5.63	6.53	7.04
DFM-M1 p2 2gf	2.59	3.18	5.57	6.57	7.06
DFM-M2 p2 2gf	2.58	3.17	5.58	6.56	6.92

Table 4: Absolute root mean squared forecast errors relative to the performances of the four DFMs when forecasting at several forecast horizons - Pre-COVID-19 outbreak evaluation period

The results in table 4, besides displaying a well expected overall improved performance as witnessed by the lower RMSFEs at any horizon, confirm what previously found with the full evaluation period. More precisely, the DFMs tend to produce more reliable real GDP predictions as their information sets get larger. Moreover, and as opposed to the previous findings, the best forecasting performance is now obtained by the two simplest models, i.e., DFM-M for very short-term projections ($h = 0, 1$) while DFM-M p2 for longer-term ones ($h = 2, 3, 4$). This suggests how when the evaluation period is not characterised by high volatility even a simpler model can produce reliable predictions.

In order to understand what the predictions look like, figure 5 depicts how the real

GDP growth rate forecasts evolve over time. All the forecasts are obtained from the DFM that features two lags in the transition equation (DFM-M p2). Being a model built for very short-term forecasts, the choice of the latter was done on the basis of the overall best performance at $h = 0$ across the three databases (M0, M1 and M2) over the full evaluation period.[12] The solid black lines represent the actual GDP series while the coloured lines marked with circles represent the forecasts done with the vintage available at each point in time. The difference between the two represents the forecast error. Notice that the filled circles represent the forecasts at $h = 0$ while the empty ones at horizons $h = 1$ to $h = 4$. The top panel shows the GDP predictions done in March, June, September and December. The middle one those in January, April, July and October and, finally, the bottom panel those in February, May, August and November.

The information that can be drawn from this exercise can be summarised as follows. First, the DFMs appear to provide reliable forecasts as the latter tend to be, overall and by visual inspection, close to the actual real GDP growth series. This ability, as seen in table 4, is more evident during periods of contained volatility of the GDP series. As opposed to that, periods of heightened volatility tend to be more challenging and the predictions appear less accurate at $h = 0$ but as h gets bigger the latter become more reliable. These considerations might highlight the necessity of a richer forecasting tool kit as different economic conditions require different predictive models. Second, as the the information set gets richer with the inclusion of more monthly observations, the coloured lines get closer to the output series. This is particularly the case of the recent turbulent period characterised by the COVID-19 pandemic that induced a sharp decline in economic activity followed by a sustained recovery.

Finally, table 5 reports the results of the DFMs vis-à-vis those of the competitors in relative terms with respect to a simple random walk.[13] As before, the exercise is conducted on the full evaluation period but also on the pre-COVID-19 one. By looking at the middle section of the table, a few things are clear. First, all the DFMs that exploit information up to the month the GDP release is referred to perform roughly as a random walk at $h = 1$. This is true also for the small BVARX that uses the setup as in Ruisi and Borg (2018). Second, apart from these few occasions, all the DFMs are able to forecast the real GDP growth rate more reliably than a random walk at any horizon. Third, for the 1-step ahead forecast the DFM-M2 p2 performs the best. Fourth, as the horizon grows large the models that tend to provide reliable predictions are the two versions of the DFM that feature two lags in the transition equation of the factors, i.e., DFM-M p2 and DFM-M p2 2gf, together with the VARs. However, the best higher-horizon forecasts are delivered by the medium size BVARX whose priors are implemented as in Giannone et al. (2015). The accuracy improvement of the latter ranges between 25% and 30% when compared with a random walk. Overall, therefore, the DFMs are better at very short-term horizons while the BVARXs should be preferred at longer ones.

[12]As a comparison, figure B.1 in appendix B shows the real GDP growth forecasts obtained with the best performing BVARX. Moreover, those produced by all the other models are available upon request.

[13]Note that, as all the other models (including the random walk) cannot forecast at $h = 0$, the results relative to the latter forecast horizon are not available.

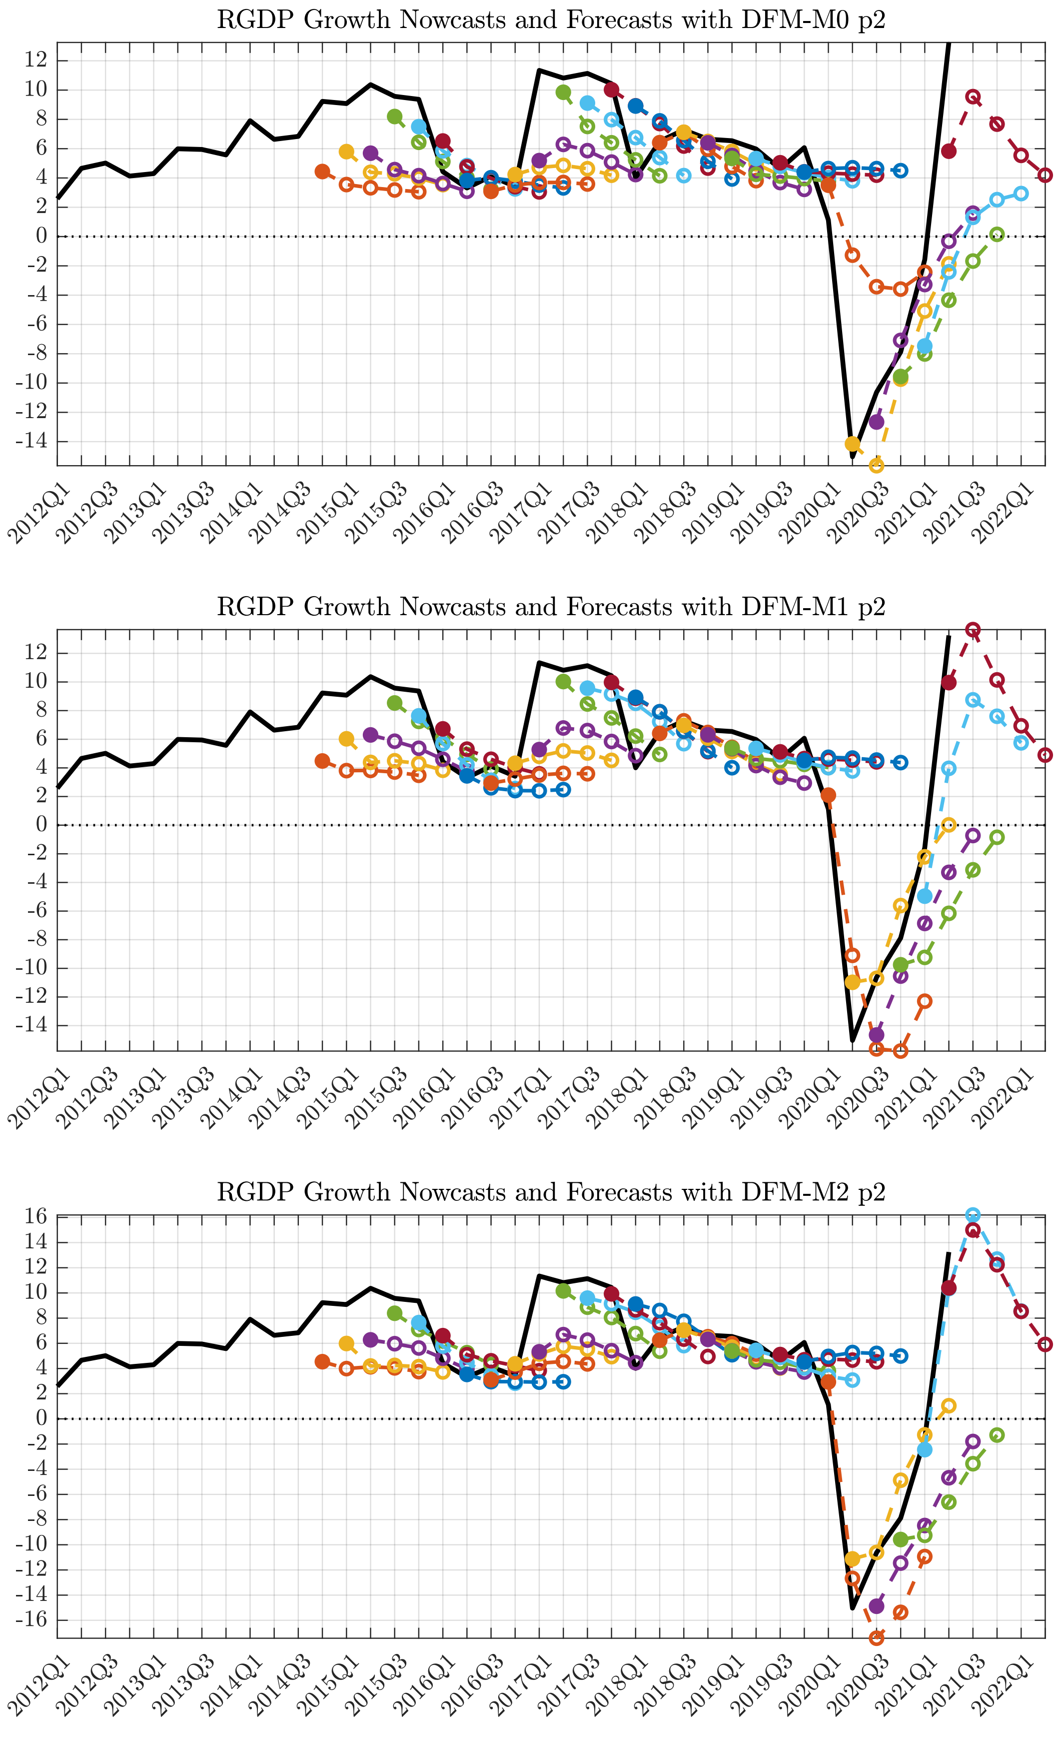

Figure 5: DFM-M p2 - Real GDP growth forecasts and nowcasts over time

Notes: the top panel shows the projections done in March, June, September and December. The middle one those in January, April, July and October. The bottom panel those in February, May, August and November.

	Full evaluation period				Priot to COVID-19 outbreak			
Model	h = 1	h = 2	h = 3	h = 4	h = 1	h = 2	h = 3	h = 4
DFM-M0	1.08	0.93	0.84	0.83	1.17	0.97	0.92	0.88
DFM-M1	0.80	0.90	0.85	0.82	1.10	0.97	0.94	0.88
DFM-M2	0.65	0.92	0.86	0.82	1.07	0.96	0.93	0.87
DFM-M0 p2	1.01	0.85	0.76	0.77	1.17	0.96	0.92	0.88
DFM-M1 p2	0.74	0.87	0.81	0.78	1.09	0.95	0.92	0.87
DFM-M2 p2	**0.63**	0.89	0.82	0.75	1.07	0.95	0.92	0.85
DFM-M0 gf2	1.12	0.92	0.83	0.82	1.19	0.97	0.93	0.89
DFM-M1 gf2	0.83	0.89	0.84	0.81	1.11	0.97	0.94	0.89
DFM-M2 gf2	0.68	0.89	0.83	0.79	1.10	0.96	0.94	0.89
DFM-M0 p2 gf2	1.03	0.85	0.75	0.74	1.19	0.97	0.92	0.88
DFM-M1 p2 gf2	0.88	0.84	0.78	0.76	1.12	0.96	0.93	0.88
DFM-M2 p2 gf2	0.70	0.86	0.79	0.72	1.11	0.96	0.93	0.86
AR(1)	0.95	0.88	0.82	0.79	0.97	0.95	0.92	0.88
Small BVARX 2018	1.00	0.95	0.91	0.89	1.04	1.04	1.03	1.02
Medium BVARX 2018	0.90	0.86	0.82	0.81	**0.92**	0.95	0.92	0.93
Small BVARX 2021	0.89	0.85	0.82	0.80	1.07	1.00	0.98	0.97
Medium BVARX 2021	0.92	0.85	0.78	0.77	0.95	0.93	**0.86**	0.87
Small BVARX GLP	0.89	0.83	0.79	0.74	0.96	0.93	0.89	0.84
Medium BVARX GLP	0.77	**0.75**	**0.72**	**0.70**	0.93	**0.88**	**0.86**	**0.83**

Table 5: Real GDP growth RMSFEs associated with the performances of all the competing models at several forecast horizons and relative to the performance of a random walk. Relative RMSFEs over the full evaluation period (middle section) and prior to the COVID-19 pandemic outbreak (right section)

When looking at the rightmost section of table 5, overall, the relative RMSFEs of all the models tend to be higher. This is not due to a worsened performance of the latter but, on the contrary, it is the result of the expected improved performance of a simple random walk when most of the turbulent years are removed from the evaluation period.[14] In such conditions, all the DFMs, the simple autoregressive model and the small BVARXs tend to have a comparable performance with that of a random walk, i.e., slightly worse for $h = 1$ and slightly better for $h = 2$. As the forecast horizons get larger their accuracy generally improves. Finally, when the current volatility is taken out, the best projections are delivered by the medium BVARXs, i.e., the 2018 version for $h = 1$ while the GLP for the remaining horizons.[15] Interestingly, the 2021 version has a comparable performance to that of the GLP for $h = 3$.

4 Conclusion

This paper describes a dynamic factor model for the Maltese economy. The latter exploits the information contained in medium and large sized datasets and can handle mixed frequencies as well as any path of missing data. The GDP forecasts and nowcasts can be updated daily depending on data published by data-providing institutions. Once the databank is updated, and the DFM estimated, the model updates its current quarter nowcast, and returns an updated forecast, discriminating between the impact of revisions to data as well as newly observed information.

The model is shown to provide timely and reliable forecasts and nowcasts as well as other valuable pieces of output such as interpolated versions of low frequency series and backdating of late starting ones. Moreover, in a horse-race with a number of competing forecasting models of real GDP growth, the DFM outperforms the others at very short horizons. At longer horizons, however, bayesian VARs with a block of exogenous variables as in Ruisi and Borg (2018) and priors set up as in Bańbura et al. (2010) and Giannone et al. (2015) appear to be more reliable. These results hold both when the evaluation period contains highly volatile data and when it does not as prior to the COVID-19 outbreak.

The results contained in this paper evidence how the DFM is a valid tool and is eligible to enrich the suite of models of the CBM for its forecasting process. Moreover, the analysis conducted in this document can be considered as an important starting point for further improvements which are left for future works, e.g., state-contingent models for forecasting during boom or bust periods or, subject to data availability, models that exploit even higher-frequency series.

[14]The improved performance of the DFMs over the pre-COVID-19 evaluation period was already outlined in table 4. The absolute RMSFEs of all the competing models, as well as those of the random walk, are not shown here for the sake of space but are available upon request.

[15]The result that the medium BVARX 2018 has the best performance at $h = 1$ is not surprising as the evaluation periods mostly overlaps with that in Ruisi and Borg (2018).

References

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.

Bańbura, M. and Modugno, M. (2012). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, 29(1):133–160.

Baffigi, A., Golinelli, R., and Parigi, G. (2004). Bridge models to forecast the euro area gdp. *International Journal of Forecasting*, 20(3):447–460.

Bai, J. and Ng, S. (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica*, 74(4):1133–1150.

Bai, J. and Ng, S. (2013). Principal components estimation and identification of static factors. *Journal of Econometrics*, 176:18–29.

Baumeister, C., Leiva-León, D., and Sims, E. R. (2021). Tracking weekly state-level economic conditions. *NBER Working Papers Series*, (29003).

Bok, B., Giannone, D., Caratelli, D., Sbordone, A. M., and Tambalotti, A. (2018). Macroeconomic nowcasting and forecasting with big data. *Annual Review of Economics*, 10:615–643.

Connor, G. and Korajczyk, R. A. (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics*, 15(3):373–394.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.

Engle, R. F. and Watson, M. W. (1983). Alternative algorithms for the estimation of dynamic factor, mimic and varying coefficient regression models. *Journal of Econometrics*, 23(3):385–400.

Eraslan, S. and Götz, T. (2020). An unconventional weekly economic activity index for germany. *Deutsche Bundesbank Technical Paper.*

Evans, M. D. D. (2005). Where are we now? real-time estimates of the macroeconomy. *International Journal of Central Banking*, 1(2).

Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2000). The generalized dynamic-factor model: Identification and estimation. *The Review of Economics and Statistics*, 82(4):540–554.

Forni, M. and Reichlin, L. (1998). Let's get real: A factor analytical approach to disaggregated business cycle dynamics. *The Review of Economic Studies*, 65(3):453–473.

Geweke, J. (1977). The dynamic factor analysis of economic time series. *Latent variables in socio-economic models.*

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *The Review of Economics and Statistics*, 97(2):436–451.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676.

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45.

Lewis, D., Mertens, K., Stock, J., and Trivedi, M. (2020a). High frequency data and a weekly economic index during the pandemic. *Federal Reserve Bank of New York Staff Reports*, 954.

Lewis, D., Mertens, K., Stock, J., and Trivedi, M. (2020b). Measuring real activity using a weekly economic index. *Federal Reserve Bank of New York Staff Reports*, 920.

Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443.

Ruisi, G. and Borg, I. (2018). Forecasting using bayesian vars: A benchmark for stream. *Central Bank of Malta Policy Working Papers Series.*

Sargent, T. J., Sims, C. A., et al. (1977). Business cycle modeling without pretending to have too much a priori economic theory. *New methods in business cycle research*, 1:145–168.

Stock, J. and Watson, M. (1988). A probability model of the coincident economic indicators. *NBER Working Papers*, (2772).

Stock, J. and Watson, M. (1989). New indexes of coincident and leading economic indicators. *NBER Macroeconomics Annual 1989*, 4:351–409.

Stock, J. H. and Watson, M. W. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460):1167–1179.

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. *Handbook of Macroeconomics*, 2A:415–525.

Williams, B. (2020). Identification of the linear factor model. *Econometric Reviews*, 39(1):92–109.

Woloszko, N. (2020). Tracking activity in real time with google trends. *OECD Economics Department Working Papers*, (1634).

Appendix A Bayesian VARs Data

Tables 6 to 8 in this appendix describes the data series and transformations used to estimate the bayesian VARXs in subsection 3.4.

Series	Transformation	Source
Real GDP	YoY growth rate	NSO - National accounts
GDP deflator	YoY growth rate	NSO - National accounts
Unemployment rate	Levels	NSO - LFS

Table 6: Endogenous variables used for the estimation of the small bayesian VARXs

Series	Transformation	Source
Consumption deflator	YoY growth rate	NSO - National accounts
Capital Formation deflator	YoY growth rate	NSO - National accounts
Government consumption deflator	YoY growth rate	NSO - National accounts
Exports deflator	YoY growth rate	NSO - National accounts
Imports deflator	YoY growth rate	NSO - National accounts
House prices	YoY growth rate	CBM
Real private consumption	YoY growth rate	NSO - National accounts
Real capital formation	YoY growth rate	NSO - National accounts
Real government consumption	YoY growth rate	NSO - National accounts
Real exports	YoY growth rate	NSO - National accounts
Real Imports	YoY growth rate	NSO - National accounts
Total employment	YoY growth rate	NSO - National accounts
Harmonised index of consumer prices	YoY growth rate	NSO
Compensation per employee	YoY growth rate	NSO - National accounts

Table 7: Endogenous variables used for the estimation of the medium bayesian VARXs in addition to those listed in table 6

Series	Transformation	Source
Oil prices (in EUR)	YoY growth rate	ECB
Competitor prices (import side)	YoY growth rate	ECB
Competitor prices (export side)	YoY growth rate	ECB
Real effective exchange rate (import side)	YoY growth rate	ECB
Real effective exchange rate (export side)	YoY growth rate	ECB
Foreign demand	YoY growth rate	ECB
Monetary policy rate	Levels	ECB

Table 8: Exogenous variables, and their projections, used for the estimation of the bayesian VARXs (both small and medium)

Appendix B Forecasting GDP with the large BVARX GLP

This appendix aims at comparing the DFM's reliability with that of the best performing BVARX described in subsection 3.4. In order to do so, this appendix depicts the real GDP forecasts obtained by means of the medium BVARX GLP. The latter VAR is estimated by using the same sample size (from a time-length standpoint) and the same evaluation period as in subsection 3.4. The variables used are those listed in tables 6 to 8 in appendix A.

Figure B.1 clearly evidences how reliable the predictions are thus confirming what suggested by table 5. It is evident how this VAR is able to produce projections that are closer to the actual GDP data for horizons that go from $h = 1$ to $h = 4$ both in the per-pandemic period and after March 2020.

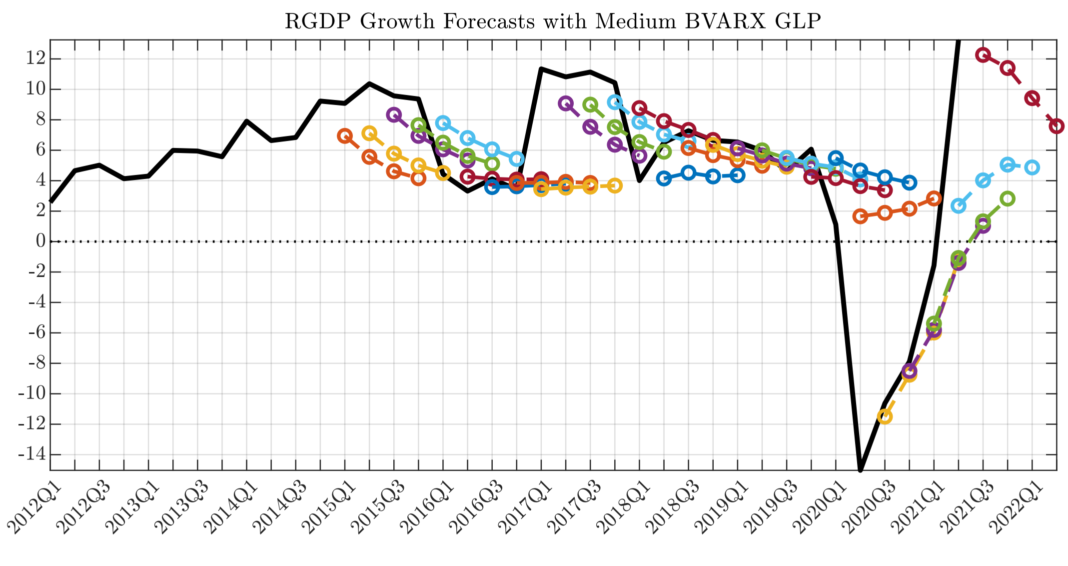

Figure B.1: Medium BVARX GLP - Real GDP growth forecasts over time