Nowcasting the Finnish economy with a large Bayesian vector autoregressive model

Juha Itkonen, Bank of Finland
Petteri Juvonen, Bank of Finland and University of Jyväskylä

Abstract

Timely and accurate assessment of current macroeconomic activity is crucial for policymakers and other economic agents. Nowcasting aims to forecast the current economic situation ahead of official data releases. We develop and apply a large Bayesian vector autoregressive (BVAR) model to nowcast quarterly GDP growth rate of the Finnish economy. We study the BVAR model's out-of-sample performance at different forecasting horizons, and compare to various bridge models and a dynamic factor model.

JEL codes: C52, C53, E32, E37

1   Introduction

Policymakers and economic agents rely on timely and accurate information on the current condition of the economy. The official statistics, however, are published with a considerable delay and are subject to revisions long after their initial release. European quarterly GDP flash estimates are released with a 45 day delay and the first official statistic with a 60 day delay counting from the end of the quarter. This means that decision makers have to wait 2–5 months for the official GDP statistics of the ongoing quarter, depending on the current date within the quarter.

In economics nowcasting refers to the process of forecasting the current state and the growth rate of the economy (Banbura et al., 2013). In practice nowcasting is primarily used to forecast the previous, the present and the next quarter-on-quarter growth rate of GDP. Also other economic variables can be forecast with similar methods.

Our main contribution in this paper is to develop and apply a large Bayesian vector autoregressive (BVAR) model to nowcast quarterly GDP growth rate of the Finnish economy. A key advantage of a BVAR model is that it can be specified in levels. Stationarizing the data by taking first differences of variables, which is necessary with traditional dynamic factor models (DFMs), tends to amplify noise which is already considerable in the case of a small open economy. The long-run information contained in levels data might help to produce more consistent forecasts. Also the development of a new model was motivated by the aim to reduce the persistent upward forecasting error experienced with the DFM employed at the Bank of Finland. In this paper we specify the BVAR model and study its out-of-sample performance at different forecasting horizons, and compare it to various bridge models and the dynamic factor model.

Nowcasting models typically aim to exploit a large dataset of indicators that are published on a high frequency. For example, consumer and business confidence indicators provide "soft" information on the state of the economy.[1] Similarly, "hard" information is provided by indicators of industry and service sector turnouts, industrial production, retail and wholesale sales, and employment. In addition, financial data such as stock market indices and interest rates may be used to predict where the economy is headed.

The information set available for nowcasting has several features that complicate matters. First, the number of time series which have the potential to predict economic growth is very large, and indicators are often highly correlated with each other. Second, the information set contains data observed at different frequencies (typically at quarterly, monthly, or daily frequency). Third, different indicators

---

[1]In Finland, the Confederation on Finnish Industries collects data for businesses and the Statistics Finland for consumers. These then are combined in the European Commission's Business and Consumer Surveys.

have different publication lags creating a so called ragged edge at the end of the dataset. Soft information (e.g. confidence indicators) usually becomes available earlier than hard information (e.g. industry turnover). Fourth, indicators often have different starting dates and sometimes a fairly short history. Fifth, for most indicators, previous observations are revised when new ones are published.

Nowcasting models are designed to make use of all available information so they must be able to handle unbalanced and ever-growing information sets. As a result, nowcasting performance tends to improve as more data becomes available.

Currently popular modeling approaches include bridge models, dynamic factor models (DFM) and Bayesian vector autoregressive models (BVAR). Bridge models are simple regression models that employ a single indicator to predict GDP growth rate. A common practice is to use the average forecast of multiple bridge models. In contrast, DFMs aim extract common sources of variation, i.e. factors, from a large dataset in order to produce a more parsimonious model. BVARs employ a large vector autoregressive model with Bayesian priors to mitigate the curse of dimensionality, endemic to models with many parameters and few observations.

Nowcasting models can also be used to assess the importance of new data releases, a practice known as news analysis (Banbura et al., 2013). News is defined as the difference between the observation and model's forecast for a particular data point. DFMs and BVARs give forecasts for all missing observations in the dataset, so they can be used to analyze the news value of all data releases. What matters is not the new data release in itself but it's difference to our expectation. Only information that differs from our expectation calls for an adjustment in our view of the current state of the economy. For example, a decrease in unemployment does not necessarily mean that we should revise out GDP forecast upward, if the decrease in unemployment was anticipated.

The quarterly GDP statistics are relatively noisy in a small open economy like Finland, and this needs to be accounted for when building a picture of the current economic condition. The quarter-to-quarter growth rate, which is usually the variable of interest, is highly volatile. In a small economy, even a single business decisions by a large firm might affect GDP considerably. In Finland, one tenth of percentage point quarterly GDP represents only about 50 million Euros. Also the revisions for GDP statistics can be large and revisions occur long after the initial release.

In a small open economy with noisy data, the actual observations of quarterly GDP might deviate considerably from the short-run trend. Depending on the purpose for which the nowcasts are used, it is often necessary to build separate forecasting models to track GDP and its trend. When nowcasting is used to obtain a starting point for medium term forecasts, a model that tracks actual GDP is often more useful. On the other hand, if the purpose is to convey a wider picture of the

general state of the economy and its direction, a model for the short-run trend might be more relevant. Therefore it is important to evaluate the performance of nowcasting models also with respect to a short-run trend from which some of the noise has been filtered out.

We specify the BVAR model following closely Giannone, Lenza, and Primiceri (2015). The curse of dimensionality is dealt with Bayesian shrinkage. We use informative priors that push the estimates of parameter rich VAR to more parsimonious processes, namely to unit root processes as originally proposed by Litterman (1979), Doan et al. (1984) and Sims (1993), i.e. so called Minnesota type priors. Since the unit root process is a good approximation for many macroeconomic variables, this brings little estimation error while greatly reducing estimation uncertainty. Although Minnesota type priors dates back to 80's and more sophisticated approaches for Bayesian shrinkage have been developed, they have regained popularity recently. Giannone, Lenza, and Primiceri (2015) and Bańbura, Giannone, and Lenza (2015) show that the forecast accuracy of large VAR with Minnesota type Bayesian shrinkage is comparable to forecast accuracy of factor models in US and Euro area data.

Our BVAR model contains 48 variables of which 9 are observed at a quarterly frequency and 39 at a monthly frequency. The data consists of variables which we have considered to be potential predictors for the state of economy. A list of all variables is provided in Table 2.

Two distinct methods for handling mixed frequency data in a VAR model have been proposed in the literature. Schorfheide and Song (2015) show how to specify a VAR on monthly frequency, where quarterly series are treated as having missing monthly observations that can be estimated with a Kalman smoother. McCracken et al. (2015) provide an alternative approach where a VAR on quarterly frequency is specified so that, for monthly series, the three monthly observations within a quarter are treated as separate variables.

We introduce a third approach for handling mixed frequency data. We specify two separate VARs for the monthly and the quarterly variables. First, we use the monthly VAR (specified for all monthly series) to fill in quarters with missing monthly observations (usually the last quarter of the data). Second, we aggregate the monthly series to quarterly frequency while treating forecasted monthly variables as noisy signals. The precision of the noisy signal is obtained from the forecast error variance-covariance matrix of a Kalman filter/smoother. Finally, we combine the quarterly series and time aggregated monthly series, and use the quarterly VAR for forecasting and news analysis.

The paper is organized as follows. Then next section specifies the large Bayesian VAR model for the Finnish economy. Section 3 discusses and compares the forecasting performance of BVAR model against benchmark models. Section 4 illus-

trates how new data releases can be analyzed with the BVAR model and Section 5 concludes.

2 Large Bayesian vector autoregressive model

We specify separate VARs for monthly and quarterly variables to handle the mixed frequency problem. Both VAR models can be represented as:

$$y_t = c + A_1 y_{t-1} + ... + A_p y_{t-p} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma), \tag{1}$$

where y_t is a n-dimensional vector of observed variables at time t, c is vector of constants, $A_1, \ldots, A_p$ are is the coefficient matrices, and ε_t is a vector of residuals distributed normally with a variance-covariance matrix Σ.

Using a large number of variables in a VAR model typically leads to good in-sample-fit but to a poor out-of-sample performance. To mitigate the overfitting problem and to improve the forecasting performance it is useful to implement priors that bring the parameter rich VAR towards a more parsimonious model.

We assume a conjugate normal-Wishart prior for the coefficients (A's) and for the variance-covariance matrix of the residuals Σ. When a relatively uninformative prior is assumed, the estimates are centred on ordinary least squares estimates. Due the properties of conjugate normal-Wishart prior, the distributions of both the coefficients and covariance matrix are obtained through direct sampling.

As in the original Minnesota prior by Litterman (1979), we assume that each variable in the VAR follows an independent unit root process. Following closely the notation in Bańbura et al. (2015) we set prior means and variance-covariance matrix for VAR coefficients as:

$$E\left[(A_s)_{ij} | \Sigma, \lambda, \Psi\right] = \begin{cases} 1, & \text{if } i = j \text{ and } s = 1 \\ 0, & \text{otherwise} \end{cases}$$

$$cov\left[(A_s)_{ij}(A_r)_{hm} | \Sigma, \lambda, \Psi\right] = \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\Psi_{jj}}, & \text{if } m = j \text{ and } s = r \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

where $s, r \in \{1, \ldots, p\}$, $i, j \in \{1, \ldots, n\}$, and λ and Ψ are hyperparameters of the prior distributions.

In each equation of the VAR the coefficient for the first lag of the same variable is assumed to equal unity, whereas all other coefficients are assumed to equal zero. The more distant the lag, the smaller is the prior variance and hence the prior is tighter. This is determined by the s^2 term in the denominator. The overall tightness is determined by the λ-parameter. The term $\frac{\Sigma_{ij}}{\Psi_{jj}}$ accounts for the relative

scales of variables, and $E[\Sigma] = \Psi$ is the prior mean for the variance-covariance matrix of the residuals. For the intercepts of the VAR, c, an uninformative prior is assumed.

In addition, two other priors are implemented. Note that the prior implemented in (2) supports only first order unit root processes. To put prior weight also on unit root processes of higher order, we implement the "sum of coefficients" -prior as suggested originally in Doan et al. (1984). This prior is implemented by adding artificial dummy observations to data the set. First, define $x \equiv [1, y_{t-1}, y_{t-2}, ...y_{t-p}]$ and $B \equiv vec([c, A_1, A_2, ...A_p]')$. The VAR model in (1) can be rewritten as

$$y_t = X_t B + \epsilon_t, \tag{3}$$

where $X_t \equiv I_n \otimes x_t$.

To implement 'sum of coefficients" -prior, let $\bar{y}_0$ be an $n \times 1$ vector that contains the averages of first p observations for each variable. Now, the dummy observations can be defined as

$$y_\mu = \text{diag}\left(\frac{\bar{y}_0}{\mu}\right) \tag{4}$$

$$x_\mu = \left[\underset{n \times 1}{0}, \bar{y}_0, \ldots, \bar{y}_0\right]. \tag{5}$$

To understand the idea of implementing prior beliefs through dummy observations consider following. Equation (4) contains a system of n equations. When stacked to data set, (4) represents one observation. If in an equation i in (4) the sum of coefficients of lags of variable i sum exactly to 1, the size of the error term in that equation is exactly zero. If the coefficients deviate from this an estimation error is resulted to that equation for this observation. μ determines how much weight the prior gets with respect to sample information. If μ is large even a big deviations from the sum-of-coefficients prior do not generate large errors since in practice errors of dummy observations are divided by μ. In contrast, if μ approaches to zero even small deviations from the prior result large errors and hence estimates are pushed close to sum-of-coefficient prior.

Priors in (2) and (4) support only independent unit root processes and are inconsistent with a belief that variables in the VAR share common stochastic trends. To allow for common stochastic trends, Sims' (1993) "dummy-initial-observation" prior is implemented.

$$y_\delta = \left(\frac{\bar{y}_0'}{\delta}\right) \tag{6}$$

$$x_\delta = \left[\delta^{-1}, \bar{y}_0, \ldots, \bar{y}_0\right]. \tag{7}$$

The strength of this prior increases as δ decreases.

The artificial observations, $y_\mu, y_\delta, x_\mu, x_\delta$ are stacked together with the observed data.

The strictness/informativeness of our priors is determined solely by the hyper-parameters, λ, μ and δ. Sims and Zha (1998) provide "rule of thumb" values 0.2, 1 and 1, respectively. Recently it has gained success to estimate these parameters by maximizing the marginal likelihood of the VAR model as formalized in Giannone et al. (2015) and we follow this approach.

2.1   Ragged edge and time aggregation

When producing real-time forecasts the data has a ragged edge since some variables are observed with longer delays than others. Another feature of our data is that it contains both monthly and quarterly observations. It is common to time aggregate monthly variables by taking quarterly averages or sums. However when data has ragged edge this does not work trivially. Even if time aggregation is used some approach needs to be chosen to fill the missing observations of monthly variables.

In the context of factor models Bańbura and Modugno (2014) show how to specify the model on monthly frequencies and treat the unobserved monthly observations of quarterly variables as missing observations that can be filled with the use of Kalman filter. Schorfheide and Song (2015) show how to implement this in the context of VAR. Alternative approach is given by McCracken et al. (2015) that specify the VAR on quarterly frequency. In their approach each monthly observations within a quarter are treated as different variables. The first approach requires estimation of unobserved variables and a number of iterations between parameter estimates and estimates of unobserved variables. The latter approach almost triples the number of coefficients in the VAR. In order to take into account the time ordering of monthly series, it is required to add restrictions on the VAR, which makes the estimation slightly more complicated.

Our approach relies on time aggregation. The procedure to fill the ragged edge of the data and produce nowcasts (and backcasts and forecasts) is the following. First, we specify a VAR for monthly data only and a VAR for data that consists of time aggregated monthly series and quarterly series. Monthly VAR and Kalman filter are used to fill the missing observations till the end of the last quarter that features any monthly observations. Once the missing observations of monthly series are filled, the series are time aggregate to quarterly series taking the average over the quarter. Time aggregation is done using Kalman filter because this gives us the forecast error variance-covariance matrix for those time aggregated data points that consists partly or fully of forecasts of monthly variables.

Finally forecasts of GDP and other quarterly series are obtained using quarterly VAR. Since the first official GDP statistic is released with a lag of two months

there is always more recent releases of monthly variables when doing a real time forecasting exercise. Hence forecast/nowcasts of GDP are obtained conditionally on time aggregated monthly variables. Because some data points of time aggregated series consists of forecasts it is important to take this into account. This is done by assuming that time aggregated monthly series are observed potentially with a measurement error. We obtain variance covariance matrix for measurement errors from the Kalman filter for monthly data. For time aggregated data points that consists only on observed monthly variables the measurement error is zero.

The merit of our approach is that it is much easier to implement in practice than those in Schorfheide and Song (2015) and McCracken et al. (2015). The disadvantage is that the information content of quarterly variables is not used when producing forecasts of monthly variables. However most quarterly variables are observed with a long lag so we do not consider this as a major disadvantage. Time aggregation destroys some information of monthly variables since in principle a variable might be differently related to e.g. quarterly GDP growth depending on whether it is the first, second or the third month of quarter. But taking this into account increases the number of parameters of already parameter rich model and hence the estimation uncertainty of the parameters. Each monthly observation contains also some noise and time aggregation reduces it, which can be useful in practice.

We specify the VAR in state space form following closely notation in Bańbura et al. (2015):

Measurement equation
$$Z_t = C_t S_t + v_t \tag{8}$$

Transition equation
$$S_{t+1} = G S_t + w_{t+1} \tag{9}$$

with $v_t \sim N(0, R_t)$ and $w_t \sim N(0, H_t)$. Z_t contains observables variables and S_t potentially unobserved states. We can cast the VAR in (1) a linear state space representation as $Z_t = Y_t$, $C_t = [I_n, 0_{n \times np}]$

$$S_t = \begin{pmatrix} Y_t \\ \vdots \\ Y_{t-p+1} \\ c \end{pmatrix}, G = \begin{pmatrix} A_1 & A_2 & \dots & A_p & I_n \\ I_n & 0_n & \dots & 0_n & 0_n \\ \vdots & \ddots & \dots & \vdots & \vdots \\ 0_n & \dots & I_n & 0_n & 0_n \\ 0_n & \dots & 0_n & 0_n & I_n \end{pmatrix}, H_t = \begin{pmatrix} \Sigma & \dots & 0_n \\ \vdots & \ddots & \vdots \\ 0_n & \dots & 0_n \end{pmatrix}$$

The sizes of matrices in the transition equation vary depending on which variables are actually observed. Corresponding columns and rows are deleted from C_t

and R_t matrixes as some observations are missing. In the context, of a monthly VAR R_t equals only zeros, since observed variables are observed without any noise. When Kalman filter is specified for the quarterly VAR, the corresponding elements in R_t differ from zero when Z_t contains forecasts of temporally aggregated monthly variables.

To show concretely how forecast for the time aggregated variables are obtained, write transition equation with a monthly VAR(3) model.

Transition equation

$$
\begin{pmatrix} Y_{t+1}^m \\ Y_t^m \\ Y_{t-1}^m \\ c \\ Y_{t+1}^q \end{pmatrix} = \begin{pmatrix} A_1 & A_2 & A_3 & I_n & 0_n \\ I_n & 0_n & 0_n & 0_n & 0_n \\ 0_n & I_n & 0_n & 0_n & 0_n \\ 0_n & 0_n & 0_n & I_n & 0_n \\ W_n & W_n & W_n & 0_n & 0_n \end{pmatrix} \begin{pmatrix} Y_t^m \\ Y_{t-1}^m \\ Y_{t-2}^m \\ c \\ Y_t^q \end{pmatrix} + w_{t+1} \tag{10}
$$

where W_n is $n \times n$ diagonal matrix, diagonal elements equalling 1/3. We take quarterly averages since in the end we are interested on quarterly differences of flow variables for which scaling by 1/3 does not matter and for stock variables and rates quarterly average is of interest. In the equations for Y_{t+1}^q error term is alway zero, since those equations are simply definitions.

Following the Kalman recursion equations in the Appendix we obtain forecasts and forecast error covariance matrix for Y_{t+i}^q. We obtain Y_{t+i}^q for each time period (in months), but naturally we are only interested on values for periods that correspond to the first month of each quarter since this give the time aggregation for the previous quarter. Let $P_{t+i|t}^q$ be the covariance matrix for $E_t[Y_{t+i}^q]$. If on period t all monthly variables on three previous periods are observed, all elements in P_t^q equal zero. As an example, in a special case where the first two moths are fully observed and there are no observations from the third month, $P_t^q = W_n \Sigma_m W_n'$, where Σ_m is the variance-covariance matrix of the residual in the monthly VAR. Finally, when running the Kalman filter/smoother with quarterly data in order to nowcast the GDP, P_t^q determines the elements in matrix R_t that correspond to time aggregated monthly variables.

3   Assessment of the forecasting ability

Assessment of the forecasting ability of models is an important part of the development of forecasting models. Comparison of performed forecasts and actual outcomes is a natural way of assessing the forecasting accuracy of models. The forecasting error is the difference between forecast and actual outcome.

Firstly, a good forecasting model should be unbiased, which means that the model does not systematically generate higher or lower forecasts than the out-

comes. Lack of bias can be analyzed by calculating the mean forecast error (MFE) for a certain period. The MFE for an unbiased forecasting model does not significantly deviate from zero.

Lack of bias alone is not enough, as the model should also be accurate. An accurate model generates forecasts that come close to the actual outcomes. The root mean square error (RMSE) is a frequently used forecast error measure, which is obtained by computing the mean of squared forecast errors and extracting the square root of that. The closer the RMSE figure is to zero, the more accurate is the forecasting model. In relative terms, the measure gives greater weight to large forecast errors.

When realized forecast errors are assessed and new models developed, it is of vital importance to bear in mind what information was available at the time when the forecast was made. With hindsight, it is easy to build models that explain positive developments, when the variables to be forecasted are already known for a long period. However, in practice, this often gives an overly optimistic picture of the model's forecasting ability. The challenge of the economic forecaster is to assess development in a situation when future observations are not yet known.

It is always good to use pseudo out-of-sample forecasts in assessment of forecasting ability to avoid the use of hindsight (Figure 1). Thus the forecast is made only on the basis of the part of statistical data that was available prior to the period to be forecasted. In that case, the forecasting models are compared in the same situation as the one when the forecaster applied the model.

3.1 GDP revisions

The information on GDP development provided by national accounts is updated in connection with new statistical releases. The statistical information becomes more detailed as the statistical authority gets access to more information on the development of different sectors of the economy. A more accurate picture of GDP is obtained through personal tax information, for example, but this information is not final before the tax forms have been filed and the taxation decisions have been made.

Thus the first information on GDP is not final. In this respect, GDP is a moving target, which complicates the forecaster's task. Even if the forecast were to exactly hit the growth rate in the statistical release, it is very possible that the forecast will be off when the statistics get updated. From a practical point of view, this means that perfect accuracy cannot be expected from nowcasting forecasts.

The uncertainty connected with the actual statistics can be assessed by studying the difference between data in the first releases and the latest information. At times, the adjustment of GDP growth rates has been considerable. In particular, the exceptionally large drop in the first quarter of 2009 was not yet revealed in the

Figure 1: GDP Q/Q growth rate and pseudo out-of-sample forecasts at the end of the forecasted quarter

first statistical release. The official statistic has been revised for the initial growth rate of −2.7% to −6.9% in the latest release.

The adjustment of statistics can be assessed by using the same measures as in the study of forecast errors. In 2008–2015, the average statistical revision of quarterly GDP growth was 0.2%. The figure is comparable with the MFE figure presented above. Thus the growth rates in the first statistical releases have been slightly higher than the latest growth rates, i.e. the picture of GDP growth has slightly decreased as the data have been adjusted.

The deviation of statistical revisions in growth rates has been fairly wide. In half of the cases, the revision was between −0.3% and 0.5%. Correspondingly, 75% of the revisions fitted in between −0.5% and 0.8%. The standard deviation of statistical revisions was 0.99. The figure is comparable with the RMSE figure measuring the forecast error of the forecasting models.

The standard deviation of statistical revisions can also be considered as a sort of lower limit of forecast errors of forecasting models, and better accuracy should not be expected. Even if the forecasting model were always to hit exactly the figure of the first statistical release, as the statistics are updated the RMSE figure of the forecasting model in question would increase to close to one in a comparison with the latest statistical GDP release.

10

3.2 Benchmark models

3.2.1 Bridge models

Bridge models are simple linear models which forecast GDP quarter-to-quarter growth rate using a single monthly indicator variable that has been aggregated to the quarterly level (Baffigi et al., 2004). Although recent advances in nowcasting methodology has made more comprehensive and elaborate models available, bridge models are still widely used, and merit their place in the nowcasters tool box. First, they tend to perform well in forecasting the upcoming quarter, as we will show later. Second, bridge models are easy to implement and interpret. Third, practice has show bridge models to be relatively robust to structural breaks in the economy.

To handle missing monthly observations at the end of the indicator time series, we use an ARIMA model to complete the missing values.[2] Let x_t be the aggregated indicator at quarter t, typically calculated as the average of monthly observations (and ARIMA predictions) within the quarter. Let y_t be the GDP quarter-to-quarter growth rate. We use two lags for all models. Hence the model takes the form

$$y_t = \alpha + \beta_1 x_t + \beta_2 x_{t-1} + e_t, \quad t = 1, 2, \ldots,$$

where α, β_1, and β_2 are parameters, and e_t is the error term for the quarter t.

Several indicators perform well in a bridge model. The European Commission's Business and Consumer Surveys cover all EU countries and provide many useful indicators. Economic sentiment indicator (ESI), which is a composition of consumer and business confidence indicators, is among the most used. Similarly, the consumer confidence indicator, industrial confidence, and industrial production expectations perform well in a bridge model. Also foreign indicators, such as the ESI for the Euro area, can be used to forecast the Finnish economy. Euro area Purchasing managers' index (PMI), and other similar foreign indicators, can also be useful (Finland does no have its own PMI).

The trend indicator of output, produced by Statistics Finland based on the same data as quarterly GDP, is published monthly and is very useful for nowcasting. It is used to calculate the official quarterly GDP flash estimate. Volume indexes of industrial product also provide hard information that can be used in bridge models.

For benchmarking we use a single indicator bridge model and an average of several bridge models. The single indicator bridge model is based on the Economic sentiment indicator for Finland. The average bridge model is based on 17 bridge models whose forecasts have been weighted according to the inverse of out-of-

---

[2]We use an an automatic algorithm by Hyndman and Khandakar (2008) to produce the ARIMA forecasts for the indicator.

sample RMSEs.[3]

3.2.2 Dynamic factor models

In recent years, dynamic factor models have become one of the main tools for nowcasting at central banks and other institutions that produce economic forecasts. The advantage of the model is its capacity to use various sources of information and filter out an up-to-date picture of the economic state and direction. The Bank of Finland uses a dynamic factor model based on the approach by Giannone et al. (2008), a so-called factor-augmented vector autoregressive model (Kostiainen et al., 2013). The latest version of the Bank of Finland indicator model uses information on the economic situation from 75 different statistical series.

Dynamic factor models take advantage of the strong covariance among the indicator variables, and extract a small number of common factors that drive the co-movement of the variables.

Let y_t be a vector of n stationary variables observed at time t with mean ν. The dynamic factor model is defined by equations

$$y_t = \nu + \Lambda F_t + e_t \quad \text{and}$$

$$F_t = \Phi_1 F_{t-1} + \cdots + \Phi_2 F_{t-2} + u_t,$$

where F_t is a r-dimensional vector of common factors, Λ is a $n \times r$-matrix of factor loadings, idiosyncratic component e_t is the n-dimensional residual, whose covariance matrix Γ is diagonal, and u_t is the r-dimensional white noise with covariance matrix Q. Stock and Watson (2011) provide a recent survey of various methods for estimating DFMs.

3.3 Out-of-sample forecast results

Results for our out-of-sample forecasting exercise are shown in Table 1. Our target variable is quarter on quarter growth rate of the seasonally adjusted GDP in percentages. We use only the latest vintage of the data but take into account the publication lags of different variables in order to construct relevant information sets. The forecasting horizon varies from one month after the end of the target

---

[3]These bridge models use the following variables: Economic sentiment indicator, Consumer confidence indicator, Consumer confidence on general economic situation over next 12 months, Consumer unemployment expectations over next 12 months, Construction confidence indicator, Construction employment expectations, Industrial confidence indicator, Industrial order books, Industrial production expectations, Services confidence indicator, Services confidence: evolution of employment in recent months, Retail trade confidence indicator, Retail trade employment expectations, Economic sentiment indicator for the euro area, IFO index ISM Manufacturing PMI, and Markit Eurozone Composite PMI.

quarter to (-1) to five months before the end of the target quarter (5). For autoregressive models that operate on quarterly GDP data, only the relevant 1 and 4 months horizons are reported (the forecast is updated only when the quarterly national account, which has a 2 month publication lag, is released).

Consistent with findings in the literature, the dynamic factor model and the BVAR are able to outperform the simpler models on the shorter horizons, but on the longer horizons the advantage is smaller. In terms of accuracy measured by root mean squared error, the BVAR slightly outperforms the DFM and other models on the short horizons. For longer forecast horizons the bridge models give relatively good accuracy. The average of several bridge models does not seem to be an improve in performance compared to the single bridge model using the Economic sentiment indicator.

Bridge models, DFM and BVAR outperform autoregressive models on the evaluation period from 2006 to 2015. When the year 2009 is removed from the evaluation period, the difference in performance between the AR model and the more advanced models becomes smaller on the 4 months forecast horizon. This indicates that autoregressive models made exceptionally bad forecast during the sharp contraction of the economy in the aftermath of the financial crisis. Measured on mean absolute error, forecasting accuracy was very similar in all models when year 2009 is excluded. For the four months ahead forecasts, the AR model performed well on MAE, but for the one month ahead forecast the BVAR won the comparison.

Considering forecasting bias, the mean error was clearly positive for all but the random walk, ARIMA and BVAR model. Comparing the results for the evaluation period with and without year 2009 imply that the small positive estimation bias of the BVAR was due to the recession. The BVAR makes a clear improvement in forecasting bias compared to the DFM. As can be seen from Figure 1 the forecasts of BVAR and factor model are highly correlated.

Table 1: Forecasting performance results

	Forecast horizon	Root mean squared error		Mean absolute error		Mean error	
		2006-2015	2006-2015 exl. 2009	2006-2015	2006-2015 exl. 2009	2006-2015	2006-2015 exl. 2009
Random Walk	4	1.93	1.08	1.16	0.84	0.03	0.10
	1	1.75	1.23	1.20	0.95	0.00	0.04
AR	4	1.55	0.96	0.94	0.73	0.29	0.17
	1	1.50	1.00	0.97	0.76	0.19	0.13
ARIMA	4	1.74	1.23	1.12	0.88	-0.03	-0.05
	1	1.59	1.02	1.06	0.78	-0.03	0.02
Bridge model	5	1.40	0.90	0.89	0.72	0.35	0.26
	4	1.33	0.89	0.85	0.70	0.36	0.27
	3	1.32	0.89	0.85	0.70	0.35	0.27
	2	1.27	0.85	0.78	0.64	0.33	0.24
	1	1.24	0.80	0.77	0.62	0.32	0.22
	0	1.22	0.79	0.77	0.63	0.32	0.21
Average bridge model	5	1.38	0.90	0.88	0.72	0.38	0.27
	4	1.37	0.90	0.88	0.72	0.37	0.27
	3	1.36	0.90	0.87	0.71	0.37	0.27
	2	1.29	0.85	0.80	0.66	0.34	0.22
	1	1.26	0.82	0.77	0.63	0.33	0.21
	0	1.25	0.81	0.77	0.63	0.32	0.21
Dynamic factor model	5	1.53	0.93	0.95	0.74	0.43	0.31
	4	1.42	0.97	0.92	0.75	0.35	0.27
	3	1.22	0.92	0.83	0.69	0.30	0.24
	2	1.17	0.86	0.80	0.69	0.32	0.21
	1	1.02	0.75	0.74	0.63	0.22	0.13
	0	1.07	0.73	0.72	0.59	0.19	0.10
	-1	1.08	0.71	0.71	0.58	0.20	0.11
BVAR	5	1.43	0.97	0.90	0.77	0.21	0.05
	4	1.37	0.96	0.91	0.79	0.19	0.07
	3	1.34	0.92	0.88	0.76	0.19	0.06
	2	1.17	0.78	0.77	0.64	0.17	0.05
	1	1.05	0.68	0.68	0.57	0.11	0.01
	0	0.97	0.68	0.66	0.54	0.08	-0.01
	-1	0.93	0.70	0.64	0.54	0.05	-0.04

Notes: The forecast horizon measured as months before the end of the quarter being forecast. For Random walk, AR, and ARIMA models only the 1 and 4 months ahead forecasts are reported as they operate on quarterly data.

4 News and contribution analysis

In addition to obtaining nowcats the assessment of new data information is of interest when monitoring the economic conditions. Nowcasting models can also be used to assess the importance of new data releases, a practice known as news analysis (Banbura et al., 2013). The model gives forecasts for each missing observations in the data set, so we can define news as as the difference between the observation and model's forecast for that data point.

What matters is not the new data release it self but it's difference to our expectation. For example, a decrease in unemployment does not necessarily mean that we should revise out GDP forecast upward if the decrease in unemployment was anticipated. Only information that differs from our expectation is regarded as news, in the sense that it calls for an adjustment in our view of the current state of the economy. Figure 2 illustrates the news analysis in practice.

The news analysis helps to asses the importance of data releases only on current period. In order analyze the data from longer horizon we implement contribution analysis following the approach in Koopman and Harvey (2003). In the contribution analysis we express the nowcast of GDP as a weighted sum of observed data. We group the series to asses what sort of information is behind the nowcast. Figure 3 illustrates the contribution analysis.

Figure 2: News analysis

Variable			Latest observation (A)	Previous observation	Forecast (B)	News (A-B)	Changes in GDP forecasts 2016Q4	2017Q1	2017Q2
Granted building permits	M12	lnΔ M/M	-24,19	-0,59	-46,69	22,50	0,00	0,01	0,01
Volume index of newbuilding	M12	lnΔ M/M	-0,14	0,17	0,05	-0,19	0,00	0,01	0,01
Building starts	M12	lnΔ M/M	21,78	-7,92	5,09	16,70	0,07	0,02	0,07
Building completions	M12	lnΔ M/M	1,04	31,42	-10,55	11,60	0,01	0,00	0,01
World trade	M12	lnΔ M/M	0,49	2,53	1,28	-0,79	0,03	0,07	0,02
Employed, ages 15-74	M1	lnΔ M/M	-0,20	0,32	0,06	-0,27	0,00	0,04	0,03
Unemployment rate, ages 15-74	M1	%	8,96	8,38	8,31	0,64	0,00	0,06	0,07
Jobs vacant	M1	lnΔ M/M	-6,35	2,62	4,39	-10,74	0,00	0,29	0,13
Unemployment rate (Employment Service Stats)	M1	%	12,53	12,73	12,69	-0,16	0,00	0,24	0,12
Unemployed jobseekers	M1	lnΔ M/M	-1,81	-1,23	-0,42	-1,39	0,00	0,26	0,12
Consumer price index	M1	lnΔ M/M	-0,60	0,25	0,15	-0,76	0,00	0,03	0,04
Producer price index, manufacturing	M1	lnΔ M/M	0,77	0,78	0,41	0,36	0,00	0,00	0,00
Export price index	M1	lnΔ M/M	0,79	1,28	-0,07	0,85	0,00	0,01	0,01
Import price index	M1	lnΔ M/M	1,36	2,85	0,63	0,73	0,00	0,01	0,01
Consumer survey: Own economy	M2	-	9,00	8,80	8,83	0,17	0,00	0,02	0,01
Consumer survey: Finland's economy	M2	-	14,90	15,30	15,71	-0,81	0,00	0,00	0,01
Business confidence, Manufacturing	M2	-	1,50	3,60	5,36	-3,86	0,00	0,05	0,06
Business confidence, Construction	M2	-	3,30	3,80	3,61	-0,31	0,00	0,00	0,00
Business confidence, Manufacturing prod	M2	-	12,80	17,40	19,54	-6,74	0,00	0,04	0,04
USA pmi (ISM), manufacturing	M2	-	57,70	56,00	55,90	1,80	0,00	0,02	0,03
German IFO-index	M2	-	118,40	116,90	117,99	0,41	0,00	0,01	0,01
Economic sentiment indicator, Eurozone	M2	-	108,00	107,90	108,53	-0,53	0,00	0,02	0,03
Effect of revisions							0,07	0,02	0,03

	GDP forecasts 2016Q4	2017Q1	2017Q2
Pre-update	0,35	1,18	0,83
Post-update	0,21	0,95	0,79

Figure 3: Contribution analysis

Contributions

5 Conclusions

In this paper we have a specified a large Bayesian VAR for the Finnish economy. According to our out-of-sample forecasting exercise BVAR's forecasting accuracy is comparable or better to forecasting accuracy of bridge models and dynamic factor model. Most importantly BVAR does not seem to suffer from as large positive mean error as other models do. Hence BVAR seems as a reliable forecasting tool. Since BVAR uses a large dataset we are able to use it for news analysis, to asses the importance of new data releases, which we have illustrated in this paper.

References

BAFFIGI, A., R. GOLINELLI, AND G. PARIGI (2004): "Bridge models to forecast the euro area GDP," *International Journal of Forecasting*, 20, 447–460.

BAŃBURA, M., D. GIANNONE, AND M. LENZA (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of Forecasting*, 31, 739–756.

BANBURA, M., D. GIANNONE, M. MODUGNO, AND L. REICHLIN (2013): "Now-Casting and the Real-Time Data Flow," in *Handbook of Economic Forecasting*, vol. 2A, 195–237.

BAŃBURA, M. AND M. MODUGNO (2014): "Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data," *Journal of Applied Econometrics*, 29, 133–160.

DOAN, T., R. LITTERMAN, AND C. SIMS (1984): "Forecasting and conditional projection using realistic prior distributions," *Econometric Reviews*, 3, 1–100.

GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): "Prior Selection for Vector Autoregressions," *Review of Economics and Statistics*, 97, 436–451.

GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): "Nowcasting: The real-time informational content of macroeconomic data," *Journal of Monetary Economics*, 55, 665–676.

HYNDMAN, R. J. AND Y. KHANDAKAR (2008): "Automatic Time Series Forecasting: The forecast Package for R," *Journal of Statistical Software*, 27, 1–22.

KOOPMAN, S. J. AND A. HARVEY (2003): "Computing observation weights for signal extraction and filtering," *Journal of Economic Dynamics and Control*, 27, 1317–1333.

KOSTIAINEN, J., S. ORJASNIEMI, AND J. RAILAVO (2013): "Lyhyen aikavälin ennustemalli Suomen kokonaistuotannolle," *BoF online*, 3.

LITTERMAN, R. (1979): "Techniques of forecasting using vector autoregressions," *Federal Reserve Bank of Minneapolis Working Papers*, 115.

MCCRACKEN, M. W., M. T. OWYANG, AND T. SEKHPOSYAN (2015): "Real-Time Forecasting with a Large, Mixed Frequency, Bayesian VAR," *Federal Reserve Bank of St. Louis Working Paper Series*, No. 2015-0.

SCHORFHEIDE, F. AND D. SONG (2015): "Real-Time Forecasting With a Mixed-Frequency VAR," *Journal of Business & Economic Statistics*, 33, 366–380.

SIMS, C. (1993): "A Nine-Variable Probabilistic Macroeconomic Forecasting Model," in *Business Cycles, Indicators and Forecasting*, National Bureau of Economic Research, 179–212.

SIMS, C. AND T. ZHA (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39, 949–968.

STOCK, J. H. AND M. W. WATSON (2011): "Dynamic factor models," *The Oxford Handbook of Economic Forecasting*.

Appendix

A   Kalman filtering/smoothing equations

$$E_t = Z_t - C_t S_t$$
$$K_t = P_{t|t-1} C_t' F_t^{-1}$$
$$S_{t|t} = S_{t|t-1} + K_t E_t$$
$$S_{t+1|t} = g + G S_{t|t}$$
$$F_t = C_t P_{t|t+1} C_t + R_t$$
$$L_t = I - K_t C_t$$
$$P_{t|t} = P_{t|t-1} L_t'$$
$$P_{t+1|t} = G P_{t|t} G' + H_t$$

Smoothing step runs backwards from $t = T$ to $t = 1$ with initializations $r_T = 0_{(np \times 1)}$ and $N_T = 0_{(np \times 1)}$.

$$r_{t-1} = C_t' F_t^{-1} E_t + L_t' r_t$$
$$S_{t|T} = S_{t|t-1} + P_{t|t-1} r_{t-1}$$
$$D_t = F_t^{-1} + G K_t' N_t G K_t$$
$$M_t = C_t' D - G' N_t G K_t$$
$$N_{t-1} = C_t F_t^{-1} C_t + G' L_t' N_t G L_t$$

Matrices D_t, N_t and M_t are needed in the calculation of the contributions, see section below.

Table 2: Data description, BVAR

	Frequency	Log	Lag
Gross domestic product	Q	x	2
Private consumption expenditure	Q	x	2
Government consumption expenditure	Q	x	2
Gross fixed capital formation, residential buildings	Q	x	2
Gross fixed capital formation, excluding residential buildings	Q	x	2
Exports of goods and services	Q	x	2
Imports of goods and services	Q	x	2
Index of wage and salary earnings	Q	x	1
Price index of dwellings	Q	x	2
Volume index of industrial output	M	x	1
Capacity utilisation rate, Manufacturing	M		1
Granted building permits	M	x	2
Turnover of retail trade, volume index	M	x	1
Turnover of wholesale trade, volume index	M	x	1
Turnover of motor vehicle trade, volume index	M	x	1
Manufacturing working on orders, Index of turnover in industry	M	x	2
Exports of goods	M	x	1
Imports of goods	M	x	1
Employed,ages 15-74	M	x	1
Unemployment rate, ages 15-74	M		1
Jobs vacant	M	x	1
Unemployed jobseekers	M	x	1
Number of bankruptcy cases instigated	M	x	1
OMX Helsinki All-Share Index	M	x	0
Consumer survey: Own economy	M		0
Consumer confidence: general economic situation	M		0
Business confidence, Manufacturing	M		0
Business confidence, Construction	M		0
Business confidence, Manufacturing: production expectations	M		0
Consumer price index	M	x	0
Building cost index	M	x	2
Turnover of construction, volume index	M	x	3
Producer price index, manufacturing	M	x	1
Export price index	M	x	1
Import price index	M	x	1
New orders in manufacturing	M	x	1
ISM's Manufacturing PMI Index, US	M		0
ifo Business Climate Index, Germany	M		0
Wages and salaries sum	M	x	1
Volume index of newbuilding	M	x	2
Economic sentiment indicator,Eurozone	M		0
Index of turnover in industry	M	x	3
Building starts	M	x	3
Building completions	M	x	3
Nights spend in all accommodation establishments, domestic visitors	M	x	2
Nights spend in all accommodation establishments, foreign visitors	M	x	2
World trade	M	x	3
Turnover of service industries	M	x	2

Note: Log refers to logarithmic transformation. Publication lag in months.

B   Calculating news

As the information set expands after new data releases the nowcast of GDP is updated. We follow Bańbura and Modugno (2014) to obtain the importance of new data releases for updating the nowcast of GDP. The new nowcast can be divided to old nowcast and a revision

$$E[GDP_t|\Omega_{v+1}] = E[GDP_t|\Omega_v] + \Big(E[GDP_t|\Omega_{v+1}] - E[GDP_t|\Omega_v]\Big) \quad (11)$$

Furthermore, the revision of GDP nowcast is due to new information in Ω_{v+1}. Let data release for indicator $y_{j,t}$ be part of Ω_{v+1} but not part of Ω_v. The new information that the release of $y_{j,t}$ contains is given as

$$I_{v,j,t} = y_{j,t} - E[y_{j,t}|\Omega_v] \quad (12)$$

By defining new information this way means the new information,the news, is orthogonal to past information. In other words, the new information that the release of $y_{j,t}$ contains is unexpected, something that was not possible to forecast based on the previous information set. The nowcast of GDP_t is only revised to the extent that the new data releases were unpredictable. In fact the revision of GDP can be written as a weighted sum of the news

$$\Big(E[GDP_t|\Omega_{v+1}] - E[GDP_t|\Omega_v]\Big) = B_{v+1}I_{v+1} = \sum_{j=1}^{n} b_{v+1,j}(y_{j,t} - E[y_{j,t}|\Omega_v]) \quad (13)$$

The diagonal weighting matrix B_{v+1} is given by

$$B_{v+1} = E[GDP_t I'_{v+1}]E[I_{v+1}I'_{v+1}]^{-1} \quad (14)$$

Let GDP_t be the first variable in S_t (following the notation of state space form in section 2.1)and y_j the j^{th} then $b_{v+1,j}$ can be obtained as

$$b_{v+1,j} = cov(GDP_t, y_{j,t}|\Omega_v)/var(y_{j,t}|\Omega_v) \quad (15)$$

The covariances and variances can be obtained from matrix $P_{t|t+1}$ which is part of the Kalman filtering recursions.

C   Calculating contributions

We calculate contributions following closely Koopman and Harvey (2003). The outcome of contribution analysis is that we can express the nowcast of GPD as a weighted sum of observed variables. The distinction to news analysis is that in the news analysis we express the nowcast revision at the some specific time point as a

weighted sum of news at that time period. In the contribution analysis we want to express the forecast as a weighted sum of all observed variables on the current and past periods. In practice, only the few latest time periods have importance.

The values of the state vector S on period t conditional on all observed data are given as

$$S_{t|T} = \sum_{j=1}^{T} w_j(S_{t|T})Z_t \tag{16}$$

where Z_t is the data and w_j is the weights. The weights for $j < t$ are obtained through backward recursions from $j = t-1$ to $t = 1$

$$w_j(S_{t|T}) = B_{t,j}^* G K_j \tag{17}$$

$$B_{t,j-1}^* = B_{t,j}^* G - w_j(S_{t|T}) C_t \tag{18}$$

with the initialization $B_{t,j}^* = I - P_{t|t-1} N_{t-1}$. The weights for $j >= t$ are obtained through forward recursions from $j = t$ to $j = T$

$$w_j(S_{t|T}) = B_{t,j}^* M_j \tag{19}$$

$$B_{t,j-1}^* = B_{t,j}^* G_j L_j \tag{20}$$

The matrices D_t, N_t and M_t are obtained from the output of the Kalman smoother.