

DISCUSSION PAPER SERIES

DP15854

Nowcasting with Large Bayesian Vector Autoregressions

Jacopo Cimadomo, Domenico Giannone, Michele
Lenza, Francesca Monti and Andrej Sokol

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

Nowcasting with Large Bayesian Vector Autoregressions

Jacopo Cimadomo, Domenico Giannone, Michele Lenza, Francesca Monti and Andrej Sokol

Discussion Paper DP15854
Published 26 February 2021
Submitted 25 February 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Jacopo Cimadomo, Domenico Giannone, Michele Lenza, Francesca Monti and Andrej Sokol

Nowcasting with Large Bayesian Vector Autoregressions

Abstract

Monitoring economic conditions in real time, or nowcasting, and Big Data analytics share some challenges, sometimes called the three "Vs". Indeed, nowcasting is characterized by the use of a large number of time series (Volume), the complexity of the data covering various sectors of the economy, with different frequencies and precision and asynchronous release dates (Variety), and the need to incorporate new information continuously and in a timely manner (Velocity). In this paper, we explore three alternative routes to nowcasting with Bayesian Vector Autoregressive (BVAR) models and find that they can effectively handle the three Vs by producing, in real time, accurate probabilistic predictions of US economic activity and a meaningful narrative by means of scenario analysis.

JEL Classification: E32, E37, C01, C33, C53

Keywords: Big Data, Scenario analysis, Mixed frequency, Real time, business cycles, Nowcasting

Jacopo Cimadomo - jacopo.cimadomo@ecb.europa.eu
European Central Bank

Domenico Giannone - dgiannon2@gmail.com
Amazon.com, International Association of Applied Econometrics (IAAE) and CEPR

Michele Lenza - michele.lenza@ecb.europa.eu
European Central Bank, ECARES-ULB and CEPR

Francesca Monti - francesca.monti@kcl.ac.uk
King's Business School

Andrej Sokol - andrej.sokol@ecb.europa.eu
European Central Bank, Bank of England, Centre for Macroeconomics

Acknowledgements

Any views expressed are solely those of the authors and should not be taken to represent those of the European Central Bank, the Eurosystem, or the Bank of England. Giannone's contribution to the paper was completed prior to the author joining Amazon. This publication and its contents are not related to Amazon and do not reflect the position of the company and its subsidiaries. We are grateful to Francis Diebold, Manfred Deistler, Gabriel Perez-Quiros and Minchul Shin, as well as to participants at several conferences and seminars for insightful comments and discussions.

Nowcasting with Large Bayesian Vector Autoregressions*

Jacopo Cimadomo¹, Domenico Giannone², Michele Lenza^{1,6}, Francesca Monti^{3,5}, and
Andrej Sokol^{1,4,5}

¹European Central Bank

²Amazon

³King's Business School

⁴Bank of England

⁵Centre for Macroeconomics

⁶ECARES-ULB

February 25, 2021

Abstract

Monitoring economic conditions in real time, or nowcasting, and Big Data analytics share some challenges, sometimes called the three “Vs”. Indeed, nowcasting is characterized by the use of a large number of time series (Volume), the complexity of the data covering various sectors of the economy, with different frequencies and precision and asynchronous release dates (Variety), and the need to incorporate new information continuously and in a timely manner (Velocity). In this paper, we explore three alternative routes to nowcasting with Bayesian Vector Autoregressive (BVAR) models and find that they can effectively handle the three Vs by producing, in real time, accurate probabilistic predictions of US economic activity and a meaningful narrative by means of scenario analysis.

JEL Classification: E32, E37, C01, C33, C53.

Keywords: Big Data, Scenario Analysis, Mixed Frequency, Real Time, Business Cycles, Nowcasting.

**Disclaimer:* any views expressed are solely those of the authors and should not be taken to represent those of the European Central Bank, the Eurosystem, or the Bank of England. Giannone's contribution to the paper was completed prior to the author joining Amazon. This publication and its contents are not related to Amazon and do not reflect the position of the company and its subsidiaries. We are grateful to Francis Diebold, Manfred Deistler, Gabriel Pérez-Quíos and Minchul Shin, as well as to participants at several conferences and seminars for insightful comments and discussions.

Corresponding author: Michele Lenza. Email: michele.lenza@ecb.europa.eu

1 Introduction

Vector autoregressions (VAR) gained prominence with Sims (1980) and have been a standard tool in macroeconometrics since at least the mid-1990s, due to their ability to capture complex dynamic interrelationships among macroeconomic variables in a relatively parsimonious econometric framework. This paper shows that VARs are also a powerful tool to monitor macroeconomic conditions in real time, or nowcasting, while at the same time retaining their proficiency in the tasks that they have been routinely used for, such as structural analysis, forecasting and scenario analysis.

Parsing hundreds of economic time series in order to monitor and dissect business cycle dynamics has been one of the central issues in macroeconometrics since at least Burns and Mitchell (1946). This was a very early development of the “Big Data” phenomenon, though, nowadays, the term tends to be used also in other related contexts, for example to describe the massive, often unstructured, datasets collected via the Internet. It is therefore no coincidence that one of the first appearances of the term Big Data in an academic context was during the World Congress of the Econometric Society. In 2000, with a discussion titled “‘Big Data’ Dynamic Factor Models for Macroeconomic Measurement and Forecasting”, Frank Diebold ushered the term Big Data into macroeconometrics and further stoked interest in methods apt to deal with growing amounts of data.¹

Nowcasting – defined as the prediction of the present, the very near future, and the very recent past² – mimics, in an internally consistent and automated framework, the way in which markets process data releases in real-time. It involves continuous and immediate updates of the predictions for a comprehensive measure of economic activity (usually, GDP, which is measured quarterly and released with a substantial lag) each time new data are released and typically exploits a large set of data with different release dates and reference period. Consequently, this process implies dealing with missing data, jagged edges, mixed frequencies and various other irregularities, and is very much related to the three Vs of “Volume, Velocity and Variety,” the defining properties of Big Data. In terms of *Volume*, not only did macroeconomic datasets grow exponentially; macroeconometricians have also developed, since early on, methods to deal with complex environments in which the number of parameters is large relative to the number of observations.³ *Velocity* is key in monitoring the real-time data flow, as the nowcasts are updated continuously as soon as the data are released. Finally, the data used for nowcasting refer to different sectors of the economy, are measured at different frequency, in different units,

¹See Diebold (2003).

²For a survey, see Banbura, Giannone, and Reichlin (2011); Banbura, Giannone, Modugno, and Reichlin (2013).

³We use the term “Big Data” in the sense of high-dimensional data, meaning that the number of parameters to be estimated is large relative to the number of observations. As summarized by Diebold (2012) and Bok, Caratelli, Giannone, Sbordone, and Tambalotti (2018), it was in this context that the term Big Data started to be used in the academic circles. Statisticians also refer to these Big Data problems as “large p , small n ”, as introduced by West (2002) to describe inference in factor models with many variables (p) and relatively few observations (n). The notation used in the rest of the paper, borrowed from the practice of macroeconometrics, replaces “large p , small n ” with large n , small T .

and vary substantially in terms of their sources and precision (e.g., hard data versus soft data based on qualitative information such as survey and polls). These features relate to the notion of *Variety* in Big Data.

The challenges just highlighted have been traditionally addressed by means of Dynamic Factor Models (DFMs), as these models can handle Big Data and can be naturally cast in a state-space form. This means that inference can be easily done using Kalman filtering techniques, which provide a convenient framework for handling the irregularities of the data in real time, i.e., mixed frequencies and non synchronicity of the data releases. Indeed, factor models have been, so far, the tool of choice for nowcasting starting from the contribution of [Giannone, Reichlin, and Small \(2008\)](#) and [Aruoba et al. \(2009\)](#), and are nowadays used extensively by policy institutions and market participants (for a recent survey see [Stock and Watson, 2017](#); [Luciani, 2017](#); [Bok, Caratelli, Giannone, Sbordone, and Tambalotti, 2018](#)).

The goal of this paper is to show that Bayesian VARs, originally proposed by [Litterman \(1979\)](#) and [Doan, Litterman, and Sims \(1984\)](#) and first used in high-dimensional environments by [Banbura, Giannone, and Reichlin \(2010\)](#), can also be used to successfully handle Big Data for real-time nowcasting and, conveniently, also for real-time policy analysis within the same framework. Indeed, VARs too can be cast in state-space form and, hence, real-time data, with all their complexities – missing data, mixed frequency and other data irregularities – can be incorporated easily in a VAR and analyzed efficiently using the Kalman filter, as shown in [Bańbura, Giannone, and Lenza \(2015\)](#). The challenge is to make inference on the model’s parameters in the presence of such data irregularities. We investigate three strategies.

A first avenue entails casting the VAR model in state-space form and modelling the low-frequency processes as latent, i.e., as if they existed at a higher frequency than the one at which they can be observed. We label this method “L-BVAR,” where L stands for “latent.” The estimates of the latent processes and the uncertainty around them are obtained by means of Kalman filtering techniques.

An alternative approach to nowcasting with VARs is to estimate the model at the lowest common data frequency, treating higher-frequency data as multiple lower-frequency variables. For example, a monthly variable would be treated as three separate quarterly variables, one for each month of the quarter. We refer to this approach as blocking or stacking, hereafter “B-BVAR”. Kalman filtering techniques can then be used to handle jagged edges resulting from asynchronous data release. This approach implies, quite conveniently, that the mixed-frequency VAR model can be estimated and analyzed as a standard VAR without latent states, therefore cutting, to some extent, the computational burden compared to the L-BVAR. Moreover, the B-BVAR in principle allows for a more flexible relationship between quarterly and monthly variables.

The third approach finds a suitable high-frequency representation of the traditional models routinely used for policy purposes, which are typically estimated on balanced quarterly datasets.

The method involves estimating the model at low frequency (quarterly) and then mapping it into a corresponding model at higher frequency (monthly). With the latter, it is then immediate to use Kalman filtering techniques to handle mixed frequencies and asynchronous data release, which imply periodically missing observations and jagged edges, respectively. The main appeal of this method is that it efficiently exploits the real-time data flow using the existing quarterly VAR models that most policy institutions have been developing and maintaining over the last decades.

We evaluate the three approaches in a real-time setting, based on US data that would have been available to an econometrician in each week from the beginning of 2005 until the end of 2019. The dataset comprises eighteen variables: some key macro variables used in most structural macroeconomic models (such as GDP, consumption, investment, labour market variables and factor prices) and other macro and financial variables (e.g., industrial production, housing starts, loans and uncertainty indices) that are monitored closely by professional and institutional forecasters and are important for their information content and the timeliness of their release.

We assess the three approaches primarily on their ability to produce accurate real-time nowcasts for US GDP. The results indicate that these tools are valid nowcasting devices: all three variants capture the information contained in the data in real time and their accuracy improves as more information becomes available over the quarter. All three BVAR approaches produce forecasts that are highly correlated with, and as accurate as, the publicly available Federal Reserve Bank of New York Staff nowcasts, which is based on a DFM. We find differences in performance across the three methods only in the first few weeks of the quarter, when no information on the current quarter is available. After that, all the mixed-frequency models are comparable and outperform a standard quarterly VAR. This result suggests that all methods are effective at distilling in real time the information contained in the continuous flow of macroeconomic releases. The implementation differences between approaches have, in practice, negligible effects on their relative nowcasting accuracy. The fact that the results are not specific to the method used indicates that the predictions reflect genuine data features.

The real-time nowcasting accuracy of the models should already attest their usefulness for policy analysis. However, we propose three additional policy exercises to gain further insight into the ability of such models to capture the complex dynamic interactions among macroeconomic variables. First, we focus on the period just before and entering the current Covid-19 crisis, and we report a real-time nowcast for the first quarter of 2020, and a joint forecast of GDP growth and PCE inflation for 2020, based on the information available until the end of 2020Q1. Second, we report a counterfactual exercise aimed at tracking the 2008Q4 Fed Funds “shadow rate”, using data which became progressively available in real time in that year. Finally, we show that our mixed-frequency models can also be used to track the transmission mechanism of shocks hitting the US economy, exemplified by generalized impulse response functions in response to an exogenous shift in GDP. The similarity in the transmission of shocks estimated with our three methods, which differ in the degree of temporal aggregation imposed on the data,

suggests that the bias arising from the fact that economic agents might not be taking decisions at the same frequency at which the data is sampled (e.g. Sims (1971), Hansen, Sargent et al. (1981), Christiano and Eichenbaum (1986)) is negligible, in practice. Hence, the importance of using mixed-frequency data resides mainly in their *timeliness*.

The remainder of the paper is organised as follows. Section 2 describes and compares the three mixed-frequency BVAR approaches, Section 3 discusses the dataset and nowcasting results, and Section 4 presents the three policy applications. Finally, Section 5 concludes.

2 Methodology

As discussed, factor models have been, so far, the tool of choice for nowcasting. These models were introduced in macroeconomics by Geweke (1977), Sargent and Sims (1977) and Engle and Watson (1981), and then extended to high-dimensional time series by Stock and Watson (1999), Forni, Hallin, Lippi, and Reichlin (2000), and Doz, Giannone, and Reichlin (2012).⁴ DFMs exploit the pervasiveness of business cycle fluctuations to provide a representation of macroeconomic dynamics that is, at the same time, accurate and parsimonious. The aim of this paper is to show that large Bayesian VARs are a very successful alternative to DFMs.

VAR models are standard components of the macroeconomist’s toolkit since the pioneering work of Sims in the early 80s and have been extensively used for forecasting and policy analysis (for a systematic review, see Stock and Watson, 2001; Karlsson, 2013; Kilian and Lütkepohl, 2018), but their use for real-time monitoring of economic conditions has not yet been explored. The use of BVARs for nowcasting is new because it was recognized only recently that Bayesian shrinkage is a powerful alternative for controlling the high estimation uncertainty due to the proliferation of parameters in a high-dimensional setting (De Mol et al., 2008). BVARs offer several advantages compared to DFMs. First, factor models generally assume away the dynamic heterogeneity present in the data, i.e., they posit that shocks affect all variables in a factor model at the same time, without leads or lags.⁵ BVARs have a more general and flexible structure, and capture more accurately the salient features of the data.⁶ Second, factor models generally require the data to be made stationary,⁷ while VARs can be easily estimated also on non-stationary data (Sims et al., 1990). Third, in factor models there are many modelling choices to be made, notably the number of lags, the number of factors and the block structure. Usually, the uncertainty coming from these choices is not taken into account. By adopting the hierarchical

⁴For recent surveys, see Stock and Watson (2016); Doz and Fuleky (2019).

⁵Notable exceptions are the works of D’Agostino, Giannone, Lenza, and Modugno (2016) and Antolin-Diaz, Drechsel, and Petrella (2017), which allow for some degree of dynamic heterogeneity.

⁶Formally, large BVARs encompass DFMs, in the sense that if the data being analyzed actually have a factor structure, then the Bayesian VAR captures it, as shown in De Mol et al. (2008); Banbura et al. (2010), and the bias introduced by the imposition of priors disappears asymptotically as the number of variables increases. Recent applications of large BVARs include Altavilla et al. (2016); Ellahie and Ricco (2017); Giannone et al. (2019b); Angelini et al. (2019); Domit et al. (2019); Del Negro et al. (2020); Miranda-Agrippino and Rey (2020). For a survey of the literature see Koop (2017); Miranda-Agrippino and Ricco (2018).

⁷See Barigozzi et al. (2016) for recent advances in the estimation of non-stationary dynamic factor models for large datasets.

approach of [Giannone, Lenza, and Primiceri \(2015\)](#), we produce probabilistic forecasts that reflect all sources of uncertainty, including that coming from the setting of hyperparameters underlying the prior distributions.

In the remainder of this section, we describe a baseline quarterly model, and then discuss three approaches to deal with the real-time data flow. We conclude the section by showing analytically, for a VAR(1), how the three methodologies are related.

2.1 The baseline model

Consider the vector autoregression of order p :

$$x_{t_q} = A_0 + A_1 x_{t_q-1} + \dots + A_p x_{t_q-p} + \varepsilon_{t_q}, \quad (1)$$

where x_{t_q} is an $n \times 1$ vector of endogenous variables, ε_{t_q} is a Normally-distributed multivariate white noise process with covariance matrix Σ_ε , and A_i for $i = 1, \dots, p$ are matrices of suitable dimension containing the model's parameters. When all variables in the vector x_{t_q} are available, the model can be readily estimated with standard Bayesian methods, reviewed for example in [Karlsson \(2013\)](#), which combine the likelihood with some informative priors.

We estimate the baseline quarterly model with 5 lags ($p = 5$), using the Normal-Inverse Wishart prior, which belongs to the class of natural conjugate priors. For Σ_ε , the covariance matrix of the residuals, we use an inverse Wishart with scale parameter given by a diagonal matrix Ψ and $d = n + 2$ degrees of freedom, which is the minimum number of degrees of freedom that guarantees the existence of the prior mean of Σ_ε (equal to $\frac{\Psi}{(d-n-1)} = \Psi$). We take Ψ to be a diagonal matrix with an $n \times 1$ vector ψ on the main diagonal, which we treat as a vector of hyperparameters.

For the constant A_0 term, we use a flat prior, while for the autoregressive coefficients (A_1, \dots, A_p), we combine the Minnesota prior, originally proposed by [Litterman \(1979\)](#), with the sum-of-coefficients prior proposed by [Doan, Litterman, and Sims \(1984\)](#), which is intended to limit the explanatory power of the VAR's deterministic component. As regards the Minnesota prior, conditional on the covariance matrix of the residuals, the prior distribution of the autoregressive coefficients is Normal with the following means and variances:

$$E(A_1) = \text{diag}(\mathbf{d}), E(A_2) = \dots = E(A_p) = \mathbf{0}_n, \quad (2)$$

$$\text{Cov}[(A_s)_{ij}, (A_r)_{hm} | \Sigma_\varepsilon] = \lambda^2 \frac{\Sigma_{\varepsilon, ih}}{s^2 \Psi_{ii}} \quad \text{if } m = j \text{ and } r = s, \text{ zero otherwise.} \quad (3)$$

where \mathbf{d} is a vector whose elements are either ones, to center the prior on a random walk, or zeros, to center the prior on a white noise. The former is appropriate for variables that are likely to be persistent or show trends, the latter for time series that are more likely to be

mean-reverting. The key hyperparameter is λ , which controls the scale of all prior variances and covariances, and effectively determines the overall tightness of the prior. For $\lambda = 0$ the posterior equals the prior and the data do not influence the estimates. If $\lambda \rightarrow \infty$, on the other hand, posterior expectations coincide with the Ordinary Least Squares (OLS) estimates. The factor $\frac{\Sigma_{\varepsilon,ii}}{\Psi_{jj}}$ accounts for the different scale and variability of the data while $\frac{1}{s^2}$ is the rate at which the prior variance decreases with increasing lag length⁸. The “sum-of coefficients” prior instead postulates that the sum of the coefficients associated with the own lags of each variable in the VAR equals one, while the sum of the coefficients associated with the lags of the other variables equals zero. This prior is imposed by means of “dummy observations” and the intensity by which it is enforced is described by the parameter μ .

Summing up, the setting of these priors depends on the hyperparameters λ , ψ and μ , which reflect the informativeness of the prior distribution for the model’s coefficients. As in [Giannone, Lenza, and Primiceri \(2015\)](#), we treat these hyperparameters as random variables, and we draw them from their posterior distributions. For the hyperparameters, we choose the same rather diffuse priors described in [Giannone, Lenza, and Primiceri \(2015\)](#). The only remaining parameter to set is the number of lags p . Since longer lags are shrunk more, inference tends to be robust to the specific value of p , provided that it is large enough. We set the number of lags p equal to 5 quarters.⁹

In the next three subsections we discuss alternative approaches to adapt the BVAR to handle mixed frequencies and jagged edges, which is necessary to incorporate macroeconomic information as soon as it gets released.

2.2 L-BVAR: Low frequency variables as latent processes

The first approach for dealing with mixed-frequency treats the quarterly variables as monthly variables, with missing observations in the first two months of the quarter. The VAR model is thus defined at monthly frequency, and Kalman filtering techniques are employed to estimate the latent monthly processes. [Zadrozny \(1990\)](#), [Mittnik and Zadrozny \(2004\)](#), [Giannone, Reichlin, and Simonelli \(2009\)](#), [Mariano and Murasawa \(2010\)](#), [Kuzin, Marcellino, and Schumacher \(2011\)](#), [Foroni, Guérin, and Marcellino \(2015\)](#) have exploited this approach in a frequentist setting, while [Eraker, Chiu, Foerster, Kim, and Seoane \(2014\)](#), [Schorfheide and Song \(2015\)](#), [Brave, Butters, and Justiniano \(2019\)](#) and [Cimadomo and D’Agostino \(2016\)](#) have explored a similar approach using Bayesian methods.

We assume that the (*log-*)*levels* of our n variables (collected in the n -dimensional vector x_{t_m}) are described by a *monthly* vector autoregressive process, but otherwise similar the one in equation (1), with $p = 17$ lags.¹⁰ The time subscript t_m indicates that the model is specified at monthly

⁸As it is standard in the BVAR literature, we set the parameter governing this decay, s , to 2.

⁹Results with $p = 10$ are qualitatively similar. We decided to use $p = 5$ in the baseline specification to ensure that the monthly models that are consistent with it, and the L-BVAR in particular, are not too computationally burdensome.

¹⁰17 *monthly* lags ensure consistency with the information sets of the B-BVAR and C-BVAR models, which

frequency.¹¹

We adopt a Normal-Inverse Wishart prior with the same parametrisation as the baseline case, which combines the Minnesota prior with the sum-of-coefficients prior. The prior for Σ_ε is an Inverse Wishart with scale Ψ and $d = n + 2$ degrees of freedom and, conditional on Σ_ε , the prior distribution of the autoregressive coefficients is Normal with means and variances reported in equations (2) and (3). The priors depend on the hyperparameters λ , ψ and μ , whose posterior distributions are obtained as part of our estimation algorithm.¹²

Tackling the issue of missing data due to irregular data releases and mixed frequencies is straightforward using Markov Chain Monte Carlo methods. We interpolate quarterly data using splines to obtain a preliminary complete monthly dataset, which we use to specify the initial conditions. The latter are assumed to be Normally-distributed with mean equal to the first p months in the complete dataset, and with variance equal to zero or equal to the prior variance Ψ_{ii} depending on whether the data is observed or estimated. Starting with the parameters set at their prior mean, we iterate the following steps: Using the simulation smoother of [Durbin and Koopman \(2001\)](#), we draw the complete monthly dataset (i.e., including draws of the latent missing values) conditional on the model parameters A_m 's and Σ_m ; then, using the posterior sampler of [Giannone, Lenza, and Primiceri \(2015\)](#), we draw the hyperparameters λ , μ and ψ conditional on the complete monthly dataset, and finally, we draw the model parameters conditional on the hyperparameters and the complete monthly dataset. This process naturally also yields draws of the nowcast/forecast conditional on the dataset used for estimation.

2.3 B-BVAR: Blocking or Stacking

The idea behind blocking is to align all frequencies to the lowest sampling frequency by treating the higher frequency (monthly) variables as multiple lower frequency (quarterly) variables. Similar methods have been developed for periodic systems in the control engineering literature (see [Bittanti, 1986](#); [Bittanti and Colaneri, 2009](#); [Chen, Anderson, Deistler, and Filler, 2011](#); [Zamani, Chen, Anderson, Deistler, and Filler, 2011](#)), and have been recently applied in macroeconometrics by [Carriero, Clark, and Marcellino \(2015\)](#), [McCracken, Owyang, and Sekhposyan \(forthcoming\)](#), [Ghysels \(2016\)](#).

are estimated with 5 *quarterly* lags. For example, with data available until the end of March, i.e. the first quarter of the current year, the B-BVAR and C-BVAR include lagged monthly information up until October of the year before the last (the former because of its block structure, the latter because monthly variables enter as three-month moving averages). To ensure that this is also the case with the L-BVAR, we need 17 monthly lags.

¹¹We treat quarterly data as monthly data available only in the last month of the quarter. Hence, the latent variable we estimate inherits the features of the quarterly variable (e.g. in the case of GDP it is still defined approximately as the sum of three consecutive monthly levels). Our modelling choice implies that, in practice, we might have a richer autoregressive structure in the latent variable process. An alternative path is to also approximately model the quarterly variables as the sum, within the quarter, of a latent monthly series. Enforcing these restrictions is important in factor models where the lag structure is typically less general than in our BVARs. Indeed, when we specify the restrictions in our state space, we do not find improvements given the very general lag structure of the model.

¹²By contrast, [Schorfheide and Song \(2015\)](#) and [Brave et al. \(2019\)](#), who deal with mixed frequency following a similar approach, resort to empirical Bayes methods to select the prior hyperparameters.

We specify the VAR at quarterly frequency and define the monthly variables as three separate series, one for each month of the quarter. For example, let x_{t_m} with $t_m = 1, 2, 3, \dots$ be a vector of monthly variables. We derive from it three quarterly variables by treating data from the first, second and third months of the quarter, respectively, as three individual series:

$$x_{t_q}^q \equiv [x'_{t_m-2} \ x'_{t_m-1} \ x'_{t_m}]' \equiv [x'_{t_q-2/3} \ x'_{t_q-1/3} \ x'_{t_q}]'$$

where $t_q = t_m/3$ for $t_m = 3, 6, 9, \dots$. These three (quarterly) series can now simply be stacked with other quarterly variables y_{t_q} in a vector $x_{t_q} = [y'_{t_q} \ x_{t_q}^q]'$. x_{t_q} is a vector of length $n = q + 3m$, where q is the number of quarterly variables and m is the number of monthly variables in our system. In our empirical application, we simply model this vector as a VAR(p), just like (1), and set the number of quarterly lags to $p = 5$.

The system can then be readily estimated with Bayesian methods. The use of Bayesian shrinkage allows us to handle large systems like the one implied by the blocking approach. In contrast, classical inference, as in Ghysels (2016), is not appropriate in this context due to the high number of free parameters. We adopt the same prior that we use for the quarterly model, namely a Normal-Inverse-Wishart prior for the coefficients of the VAR centred around a random walk/white noise model, combined with a “sum-of-coefficients” prior. As for the L-BVAR, we conduct posterior inference on the hyperparameters describing the informativeness of the priors, following Giannone, Lenza, and Primiceri (2015).¹³

Given the model parameters, the nowcasts can be viewed as forecasts conditional on different information sets. We compute these using the Kalman filtering techniques described in Bańbura, Giannone, and Lenza (2015).¹⁴

2.4 C-BVAR: Cube root

This section, with further details provided in Appendix A, reflects and expands the results previously derived for DSGE models by Giannone, Monti, and Reichlin (2016). For this approach, we assume, as in the L-BVAR case, that all variables exist at higher frequency, but some are only sampled at quarterly frequency, so for these variables we only have observations in March, June, September and December. We transform all variables to correspond to a quarterly quantity when observed at end of the quarter, following Giannone, Reichlin, and Small (2008). Let us again denote by $x_{t_m} = (x_{1,t_m}, \dots, x_{n,t_m})'$ the vector of (possibly latent) monthly counterparts

¹³These priors do not explicitly take into account that, in the B-BVAR, some variables reflect the observations for three consecutive months of the same monthly time series. We maintain the same priors used in other models to preserve comparability in our empirical application. We leave the exploration of other more sophisticated priors for future research. These alternative priors could either explicitly acknowledge the restrictions across equations implied by the B-BVAR (highlighted in 2), breaking the conjugacy of the setup, or take the form of weaker shrinkage for specific linear combinations of the variables, in line with the long run priors of Giannone et al. (2019a).

¹⁴The Kalman filter handles the jagged edges in a reduced-form VAR. A special case is the model by McCracken, Owyang, and Sekhposyan (forthcoming), which requires a recursive identification structure with the variables ordered according to the time in which the data are released by the statistical office.

to the variables that enter the quarterly model.

Consistent with our definition of the monthly variables, the vector $X_{t_m} = (x'_{t_m}, \dots, x'_{t_m-3p+3})'$ corresponds to its quarterly model-based concept X_{t_q} when observed in the last month of each quarter, where $t_q = t_m/3$ for $t_m = 3, 6, 9, \dots$

Consider the quarterly $VAR(p)$ model of equation (1), rewritten in companion form:

$$X_{t_q} = \Phi X_{t_q-1} + \nu_{t_q}, \quad (4)$$

with $\nu_{t_q} = (\varepsilon'_{t_q}, \mathbf{0}_{1 \times n(p-1)})'$, which can also be rewritten in terms of monthly quantities as

$$X_{t_m} = \Phi X_{t_m-3} + \nu_{t_m} \quad (5)$$

when t_m corresponds to the last month of a quarter. We assume that the *monthly* counterpart of model (4) can be written as¹⁵

$$X_{t_m} = \Phi_m X_{t_m-1} + \nu_{m,t_m}, \quad (6)$$

with $\nu_{m,t_m} = (\varepsilon'_{m,t_m}, \mathbf{0}_{1 \times n(p-1)})'$, $\nu_m \sim \mathcal{N}(0, \Omega_m)$. We also assume that Φ_m is full, real and stable. Iteration of (6) implies that

$$X_{t_m} = \Phi_m^3 X_{t_m-3} + \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}, \quad (7)$$

which together with our previous assumptions entails the following relationships between the quarterly model (5) and the monthly model (6):

$$\Phi_m = \Phi^{\frac{1}{3}} \quad (8)$$

$$\nu_{t_m} = \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}. \quad (9)$$

From (8) it is clear that an essential part of finding a suitable mapping between the two models is finding a cube root of Φ , which raises the issue of multiple solutions. We follow the procedure proposed in [Giannone, Monti, and Reichlin \(2016\)](#) to select among multiple cube roots of Φ .¹⁶ On the other hand, equation (9) implies that the monthly covariance matrix Σ_{ε_m} , and therefore

¹⁵If the variables considered are stocks, the formulation (6) implies no approximation, because selecting a higher frequency just means sampling at a different frequency. If instead the variables considered are flows, then our definition of the monthly variables as an average over the quarter implies that we are introducing a non-invertible moving average in the growth rates. Therefore modeling this monthly concept as autoregressive introduces some mis-specification.

¹⁶We can also evaluate the likelihood of all solutions using the Kalman filter and pick the one with the highest likelihood ([Anderson et al. \(2016b\)](#) show g-identifiability when (enough) high frequency data is available), though this is more computationally intensive. In the cases where we have tried it, the solution corresponds to the one with the roots with the smallest argument, as in [Giannone, Monti, and Reichlin \(2016\)](#).

Ω_m , can be recovered from

$$vec(\Sigma_{\varepsilon_m}) = (I_{n^2} + A \otimes A)^{-1} vec(\Sigma_\nu), \quad (10)$$

with $A = [\Phi_{m11}^2 - \Phi_{m11}(J'J)^{-1}J'\Phi_{m\cdot 1}]$, $J = [I_n \dots I_n]'$ and $\Phi_{m\cdot 1} = [\Phi'_{m21} \dots \Phi'_{mp1}]'$.

In summary, the first step to obtain the C-BVAR is to estimate a quarterly $VAR(p)$ model, like the one in Section 2.1. Given estimates of the parameters of the quarterly model (4), Φ and Ω , we define a monthly model (6) with parameters Φ_m and Ω_m , which can be recovered from equations (8) and (10). Finally, as for the B-BVAR, we compute the distributions of forecasts conditional on the real-time data flow, exploiting the Kalman filtering methods.

2.5 Mapping across methodologies

We now show analytically, in the context of a VAR(1), how the different methods presented above relate to each other, and the restrictions on economic dynamics they imply, extending the analysis in Anderson et al. (2016a) to the case of monthly and quarterly variables.

Consider a VAR(1) for vector $x_t = [x_t^f \ x_t^s]'$. where x_t^f is an $n^f \times 1$ vector of high-frequency (or fast) variables and x_t^s is an $n^s \times 1$ vector of low-frequency (or slow) variables. For simplicity we assume that $n^f = n^s = 1$. The high-frequency variables are available at each point in time, while the slow variables are available only at $t, t-3, t-6, \dots$. All variables are stocks:

$$\begin{bmatrix} x_t^f \\ x_t^s \end{bmatrix} = \underbrace{\begin{bmatrix} a_{ff} & a_{fs} \\ a_{sf} & a_{ss} \end{bmatrix}}_A \begin{bmatrix} x_{t-1}^f \\ x_{t-1}^s \end{bmatrix} + \nu_t \quad (11)$$

where $\nu_t = [\nu_t^f \ \nu_t^s]'$ is iid with mean zero and variance $\Sigma = \begin{bmatrix} \sigma_{ff} & \sigma_{fs} \\ \sigma_{sf} & \sigma_{ss} \end{bmatrix}$.

In order to write the system in terms of observed variables only, we define x_{t-1}^f and x_{t-2}^f , say the monthly variables in April and May, as a function of $[x_{t-3}^f \ x_{t-3}^s]'$, the monthly and quarterly variables in March, when both variables are observable.

$$\begin{aligned} x_{t-2}^f &= a_{ff}x_{t-3}^f + a_{fs}x_{t-3}^s + \nu_{t-2}^f \\ x_{t-1}^f &= a_{ff}(a_{ff}x_{t-3}^f + a_{fs}x_{t-3}^s + \nu_{t-2}^f) + a_{fs}(a_{sf}x_{t-3}^f + a_{ss}x_{t-3}^s + \nu_{t-2}^s) + \nu_{t-1}^f \end{aligned}$$

Stacking these expressions with (11), we obtain:

$$\begin{bmatrix} x_t^f \\ x_t^s \\ x_{t-1}^f \\ x_{t-2}^f \end{bmatrix} = \begin{bmatrix} & A^3 & & \mathbf{0}_{2,2} \\ a_{ff}^2 + a_{fs}a_{sf} & a_{ff}a_{fs} + a_{fs}a_{ss} & & \mathbf{0}_{1,2} \\ & a_{ff} & & a_{fs} \\ & & & \mathbf{0}_{1,2} \end{bmatrix} \begin{bmatrix} x_{t-3}^f \\ x_{t-3}^s \\ x_{t-4}^f \\ x_{t-5}^f \end{bmatrix} + \underbrace{\begin{bmatrix} A^2\nu_{t-2} + A\nu_{t-1} + \nu_t \\ a_{ff}\nu_{t-2}^f + a_{fs}\nu_{t-2}^s + \nu_{t-1}^f \\ \nu_{t-2}^f \end{bmatrix}}_{\tilde{\nu}_t} \quad (12)$$

where

$$\begin{aligned} \Sigma_{\tilde{\nu}} = & \begin{bmatrix} \sigma_{ff} & \sigma_{fs} & \mathbf{0}_{1,2} \\ \sigma_{sf} & \sigma_{ss} & \mathbf{0}_{1,2} \\ \mathbf{0}_{2,1} & \mathbf{0}_{2,1} & \mathbf{0}_{2,2} \end{bmatrix} + \begin{bmatrix} A \\ 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \sigma_{ff} & \sigma_{fs} \\ \sigma_{sf} & \sigma_{ss} \end{bmatrix} \begin{bmatrix} A' & 1 & 0 \\ & 0 & 0 \end{bmatrix} + \\ & + \begin{bmatrix} A^2 \\ a_{ff} & a_{fs} \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \sigma_{ff} & \sigma_{fs} \\ \sigma_{sf} & \sigma_{ss} \end{bmatrix} \begin{bmatrix} A^{2'} & a_{ff} & 1 \\ & a_{fs} & 0 \end{bmatrix} \end{aligned}$$

We can relate models (11) and (12) to the three methodologies described in the previous subsections:

- **L-BVAR.** This approach simply corresponds to (11), treating slow variables as latent processes. For estimation, we rely on standard filtering and smoothing techniques.
- **B-BVAR.** System (12) has the form of a blocked VAR, but with some additional restrictions on the covariance matrix of the residuals and on the autoregressive matrix. Note that we do not impose such restrictions, and instead conduct inference on an unrestricted VAR. In this sense, our B-BVAR therefore encompasses the L-VAR.
- **C-BVAR.** The cube root C-BVAR simply corresponds to the top two rows of system (12). We estimate A^3 and the corresponding covariance matrix at quarterly frequency, and then take advantage of the relationships implied by the top rows of (12) to obtain A and the covariance matrix in monthly space.

Comparing the three methods, it is clear that the B-BVAR imposes the least restrictions on the dynamics of the monthly model. This flexibility may be useful since the finite autoregression is to be seen as an approximation of the underlying data generating process. Moreover, in general, the mapping described above is less clear-cut, for example if there is a mix of variables with a stock or flow nature, or if the data are better approximated by a model with more lags. The cost of such flexibility is the larger number of free parameters, which is handled by means of Bayesian shrinkage. The C-BVAR is instead very parsimonious, but implies many restrictions on the monthly model. Unlike the C-BVAR, which backs out the parameters governing the evo-

lution of the quarterly variables at monthly frequency analytically, the L-BVAR estimates them iteratively with the latent variables, and is somewhere in between the two previous approaches in terms of how much structure is imposed on the monthly model.

3 Nowcasting

The mixed-frequency BVARs discussed in Section 2 can be used to nowcast the economy, taking advantage of the real-time information flow, while still accounting for all the sources of uncertainty inherent in producing a forecast. We compare the different mixed-frequency methods outlined in Section 2 by assessing their performance in a fully real-time nowcasting exercise. We compare the models' point nowcasts of US real GDP growth with the New York Fed Staff Nowcasts (see [Bok et al., 2018](#)), a naïve quarterly AR(2) model and the quarterly BVAR model presented in Section 2.1. We also assess the properties of the mixed-frequency BVARs' nowcast densities.

3.1 Data

The models are estimated on key macro variables (real GDP, real consumption, real investment and a measure of real disposable income), labour market indicators (a measure of real wage inflation based on compensation per hour, employment, the unemployment rate and average weekly hours), financial market variables (the Federal Fund rate, the spread between the annualized Moody's Seasoned Baa corporate bond yield and the 10-Year Treasury note yield at constant maturity), real indicators (such as industrial production and house starts), price data (CPI and PCE price indices, as well as the GDP deflator), a credit variable (business loans), a measure of uncertainty ([Baker, Bloom, and Davis \(2016\)](#)'s economic policy uncertainty index) and the manufacturing Purchasing Managers' Index (PMI). GDP, investment, the GDP deflator and compensation per hour are available at quarterly frequency only, while the other variables are available at monthly frequency, or higher (in which case, we take their monthly averages). We reconstruct real-time weekly vintages of data that replicate the exact data availability as of each Friday between the beginning of 2005 and the end of 2019, the same convention used for the weekly updates of the New York Fed Staff Nowcasts. In each vintage, all variables are available from October 1986.

The variables enter the different models in log-levels, except the PMI and those already defined in terms of (annualized) rates, such as the unemployment rate, which enter in levels.¹⁷¹⁸

To obtain real quantities, investment and compensation per employee are deflated with the GDP

¹⁷As discussed in Section 2.4, for the C-BVAR monthly variables are transformed so as to correspond to a quarterly quantity when observed in the final month of each quarter before taking logs (see [Giannone et al., 2008](#)) With our data, that means taking 3-months moving averages of all monthly variables.

¹⁸As articulated clearly by [Giannone et al. \(2019a\)](#), low frequency trends in the data combined with small samples can lead to very poor forecasts at long horizon. This problem can be corrected by using appropriate priors for the long run or by using alternative data transformations. Since the focus of the paper is on short and medium horizons, we leave this issue for future analysis.

Table 1: Data and timing of releases

Variable	Frequency	Publication timing	Delay (days)	Transformation	Prior	FRED id
Economic Policy Uncertainty Index	m	1 st bus. day of the month	3	level	wn	USEPUINDXM
Purchasing Managers' Index	m	1 st bus. day of the month	3	level	wn	NAPM ^a
Employment	m	1 st Friday of the month	7	log-level	rw	PAYEMS
Unemployment rate	m	1 st Friday of the month	7	level	rw	UNRATE
Avg. weekly hours	m	1 st Friday of the month	7	log-level	rw	AWHNONAG
Industrial production	m	middle of the month	17	log-level	rw	INDPRO
CPI inflation	m	middle of the month	18	log-level	rw	CPIAUSL
Loans	m	3 rd week of the month	26	log-level	rw	BUSLOANS
Housing starts	m	3 rd week of the month	27	log-level	rw	HOUST
Real GDP	q	last week of the month	28	log-level	rw	GDPC1
Business investment	q	last week of the month	28	log-level	rw	FPI
GDP deflator	q	last week of the month	28	log-level	rw	GDPDEF
Compensation per hour	q	last week of the month	28	log-level	rw	COMPNFB
Private consumption	m	last week of the month	30	log-level	rw	PCE
PCE price index	m	last week of the month	30	log-level	rw	PCEPI
Real Disp. Personal Income	m	last week of the month	30	log-level	rw	DSPIC96
Fed funds rate	m	last week of the month	0	level	rw	FEDFUNDS
Credit spread	m	last week of the month	0	level	rw	BAA10YM

^aNow on Haver Analytics

Note: Data series are ordered based on the release timing within the calendar month. The (indicative) delay of each release in the fourth column is relative to the end of the reference period and based on the 2017 calendar.

deflator, while consumption is deflated with its own price index. For the sake of parameterizing the Minnesota prior, the uncertainty indicator and the PMI are assumed to be stationary and hence the coefficient on their first lag is centered around zero rather than unity. Table 1 reports all variables used, their frequency and publication lag, whether they enter the model in levels or log-levels and whether their Minnesota prior is assumed to be centered around a random walk or a white noise, as well as their FRED id.

3.2 Nowcasting performance

We start by comparing the point nowcasting performance of the BVARs, the New York Fed Staff Nowcasts and the two quarterly benchmarks over a sample period that ranges from the beginning of 2005 to the end of 2019. Figure 1a reports, for every week in the quarter,¹⁹ the root mean square errors (RMSEs) for the point nowcasts of real GDP produced by the New York Fed Staff (NY Fed DFM), a BVAR only using quarterly versions of our data (Q-BVAR),²⁰ a simple AR(2) for real GDP (AR-2), and the three approaches for mixed-frequency BVARs described in Section 2, which are labelled B-BVAR, C-BVAR and L-BVAR, respectively. For the DFM, we take the historical nowcasts available on the NY Fed website, while for all other models, we take as point forecasts the medians of the respective predictive densities at the nowcast horizon.²¹

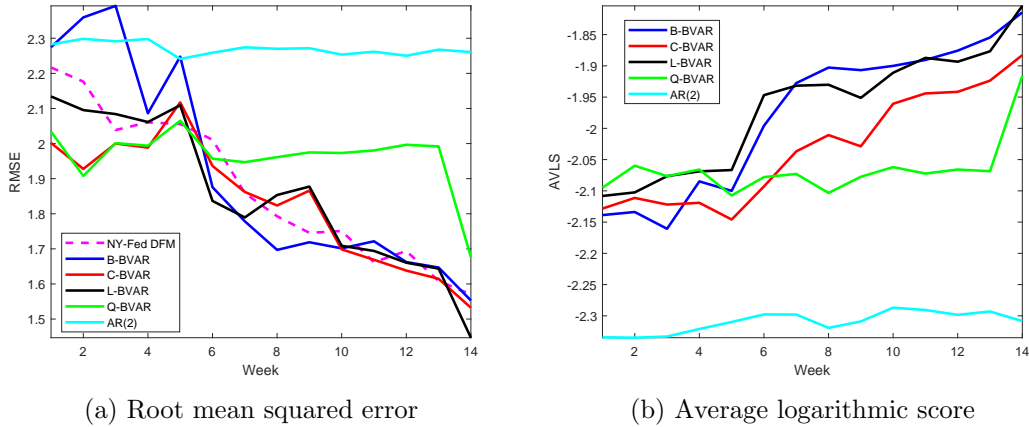
The performance of the three mixed-frequency models differs until week 6, when the first information on the current quarter becomes available. After that, the performance of all three

¹⁹Week 1 indicates the first week of a quarter, i.e., the one that contains the first Friday of that quarter; week 14 is the 14th week since the beginning of a quarter, and corresponds to Week 1 of the *following* quarter, i.e., contains its first Friday. So for example, the data vintage as of 7 January 2005 corresponds to week 1 of the 2005Q1 nowcast, while 1 April 2005 corresponds to both week 14 of the 2005Q1 nowcast and week 1 of the 2005Q2 nowcast.

²⁰The Q-BVAR corresponds to the first step needed to obtain the C-BVAR, see Section 2.4.

²¹Appendix B reports the mode of the posterior distribution of some selected prior hyperparameters, over time.

Figure 1: Nowcasting performance per week of the quarter



Note: The left panel compares the accuracy of point forecasts across models – measured by their Root Mean Squared Errors – while the right panel compares the accuracy of density forecasts – measured by their Average Logarithmic Scores – as more information becomes available in each week of the quarter. The forecast evaluation sample is 2005-2019.

models is comparable and shows monotonic improvements in accuracy as the information flows through the quarter, when the models are able to exploit the more stable contemporaneous correlations across variables, rather than the lead/lag correlations that longer-term forecasts rely on. The B-BVAR presents a somewhat more marked kink in week 5, due to a loss of accuracy in two specific episodes, namely 2009Q1 and 2009Q2.

The informational advantage that comes from being able to process higher-frequency information is evident when comparing the behaviour of the three mixed frequency BVARs to the quarterly benchmarks: while at the beginning of the quarter the performance of the Q-BVAR is comparable to or better than those of other models, by the middle of the second month, the quarterly model is clearly lagging behind, catching up to a certain extent only in week 14, when, at the close of the quarter, financial variables and the PMI and uncertainty indices for the full quarter become available.

In order to assess the ability of the models to characterize the uncertainty surrounding GDP nowcasts,²² we compute average log predictive scores for the nowcast densities at the end of each week of the quarter (Figure 1b). The average log predictive score is a common scoring rule, used to evaluate the quality of probabilistic forecasts given a set of outcomes, and is defined as:

$$S_h(\mathcal{M}) = \frac{1}{N_h} \sum_{t=R}^{T-h} \ln p(y_{t+h}|y_{1:t}, \mathcal{M}), \quad (13)$$

where h is the forecast horizon, R is the beginning of the forecast evaluation period, T is the latest period for which data are available, N_h is the number of forecast origins, and $p(y_{t+h}|y_{1:t}, \mathcal{M})$ is the predictive density from model \mathcal{M} estimated at time t and evaluated at the actual data

²²Historical density nowcasts for the NY Fed’s DFM are not publicly available, so the model is omitted from this comparison.

outturn.

The improvements of log scores with the arrival of new information throughout the quarter mostly mirror the corresponding reductions in the RMSEs. As for point forecasts, the density forecasts of the mixed-frequency BVARs perform similarly, while those of the Q-BVAR are only ‘competitive’ at the beginning and then again at the close of the quarter.

4 Policy Analysis

In this section, we present three policy exercises. First, we report a GDP nowcast, as well as a joint forecast of the annual growth rate of real GDP and of annual PCE inflation, based on data up to 2020Q1, i.e., the start of the Covid-19 pandemic crisis. Second, we show that mixed-frequency BVAR models can be used for structural analysis, just like their quarterly counterparts and, as an example, we present the generalized impulse response functions to a GDP shock. Besides describing the dynamics of the US economy, by comparing impulse responses across our BVAR approaches, we can also draw useful insights on the reasons why considering mixed-frequency data is important: whether it’s mainly for their timeliness, or because they also provide information useful for identifying the dynamic relationships among variables. Finally, we show a counterfactual exercise aimed at forecasting the 2008Q4 Fed Funds “shadow rate” in real time, a common practice in central banks, used to define “benchmark” paths for their policy rates.

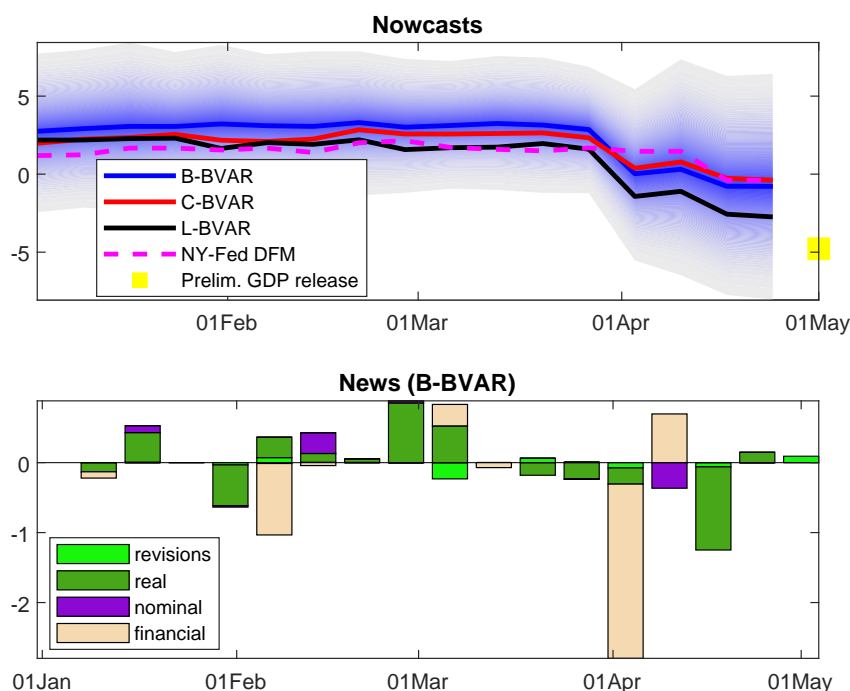
4.1 The current conjuncture: the Covid-19 crisis

The Covid-19 pandemic has triggered a dramatic contraction in economic activity worldwide, and has also strongly impacted the US economy. It seems therefore natural to apply the methods discussed in this paper to a situation in which the data flow received by the forecasters shows a progressive deterioration of the economic environment, which was indeed the case at the beginning of the 2020 pandemic crisis.

The top panel of Figure 2 reports the distributions of the nowcasts of real GDP in Q1 2020 produced by the blocked BVAR model (B-BVAR) at the end of weeks 1 through 18 of 2020, with the other models’ point nowcasts shown as lines, while the bottom panel relates the changes in the point B-BVAR nowcasts to various categories of data releases. As it is apparent, all nowcasts dropped considerably once March data started to become available in early April, and continued to deteriorate with the weekly data flow; the uncertainty around the B-BVAR nowcast also increased. Nevertheless, the preliminary GDP release on 29 April still surprised to the downside, although it fell within the range of plausible outcomes.

VARs also allow us to analyse the joint densities of two or more variables and how they evolve as more information becomes available over time. Figure 3 plots the joint distribution of the B-BVAR forecasts of annual real GDP growth and PCE inflation at different dates in the first quarter of 2020, together with their marginals. This figure describes how information about the

Figure 2: Nowcast for real GDP growth in Q1 2020



Note: The top panel shows the probability distribution of B-BVAR nowcasts in each week from the beginning of 2020 until the preliminary release of Q1 GDP on 29 April, the median nowcasts from the C-BVAR and L-BVAR, as well as the NY Fed’s DFM nowcast. The fan chart bands cover 99% of the support around the median: the darkest shade of blue corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The bottom panel imputes weekly changes in the B-BVAR’s (point) nowcast to existing data revisions and new data releases, grouped by type of variable.

economic fallout from the Covid-19 crisis is reflected in the forecasts for these two variables, both in terms of location and dispersion. Indeed, our latest forecast (as of 22 May 2020) is much more pessimistic than those made in January and even early April, and points to a median real GDP contraction of about 5.3% this year, while the bulk of the predictive distribution of PCE inflation is in negative territory. The uncertainty surrounding the May forecast is also much larger compared to the two earlier forecasts.²³


4.2 Impulse response functions

Mixed-frequency BVARs can also be used to identify shocks and investigate their transmission mechanism, thus retaining one of the most appealing features of VAR models, with the added benefit that the analysis can potentially be also carried out at monthly frequency.

Rather than engaging in a full-fledged structural identification exercise, which would rely on potentially debatable identification assumptions, for illustrative purposes, we present a general-

²³An additional avenue to improve in- and out-of-sample density forecasts is to account for stochastic volatility, which can be easily introduced in large VARs, as shown by [Carriero, Clark, and Marcellino \(2016\)](#). [Pettenuzzo, Timmermann, and Valkanov \(2016\)](#) and [Carriero, Clark, and Massimiliano \(2020\)](#) are example of models with stochastic volatility that exploit mixed-frequency data for nowcasting.

Figure 3: Evolution of the joint distribution of the forecasts for GDP growth and PCE inflation



Note: The scatter plot shows draws from the B-BVAR’s joint predictive densities for annual GDP growth and PCE inflation in 2020 in three different weekly vintages. The two plots along the axes show kernel-smoothed estimates of the marginal predictive densities for the two variables in the same three vintages. Annual growth rates are computed from the underlying projections in log levels for the corresponding variables.

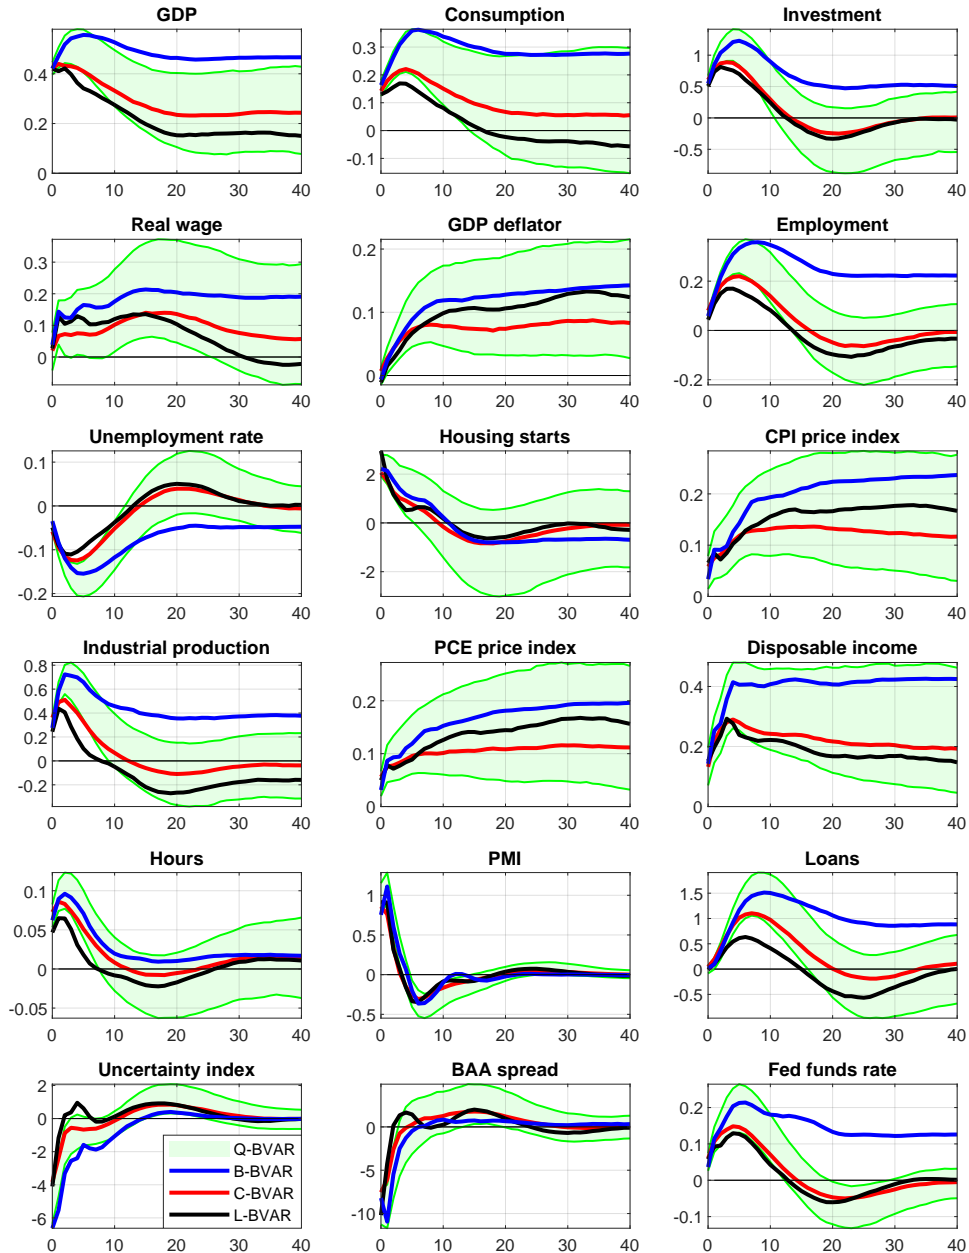
ized impulse response function to a one standard deviation GDP shock (Figure 4). Generalized impulse response functions to GDP capture the responses of the variables in the model to a linear combination of the structural shocks that have been the main historical drivers of innovations in GDP fluctuations (see, e.g., Pesaran and Shin, 1998; Bańbura, Giannone, and Lenza, 2015) and are helpful tools to characterize the dynamics of the US economy over a “typical” business cycle.²⁴ Other setups, both in terms of more elaborate identification schemes and of shocks occurring in different months of the quarter, can be easily accommodated within all three models. As an example, we report in Appendix C the impulse response functions to an uncertainty shock identified as in Bloom (2009).

Figure 4 reports the 68% credible intervals for the quarterly Q-BVAR model and shows the median responses for the three mixed-frequency approaches.²⁵ All models produce broadly similar IRFs. Consistent with the established VAR literature, a shock to GDP triggers a positive reaction of consumption, investment, and compensation per hour, while the unemployment rate decreases for about 12 quarters after the shock. The shock is inflationary, as shown by the positive reactions of the GDP deflator, CPI and PCE price indices, suggesting that demand shocks are important drivers of GDP in the US, and this is accompanied by a tightening of

²⁴In practice, the generalized impulse responses to GDP are equivalent to a perturbation of the forecast error of GDP in a recursively-identified VAR, with GDP ordered first, and yield results that are very similar to the responses to a typical business cycle shock, defined as the linear combination of structural shocks that have been the main historical drivers of innovations of GDP variation at business cycle frequencies (see Giannone, Lenza, and Reichlin, 2019b; Angeletos, Collard, and Dellas, 2020). The similarity between the two approaches was recently documented also by Del Negro, Lenza, Primiceri, and Tambalotti (2020).

²⁵The responses have been scaled to match the Q-BVAR’s impact real GDP response.

Figure 4: Generalised impulse response function to a GDP impulse



Note: Generalised impulse response function (GIRF) to a one standard deviation shock to GDP, shown at quarterly frequency. The green areas represent the 68% credible intervals for a one standard deviation impulse to the Q-BVAR model; the red, blue and black lines show the median GIRFs for the C-BVAR, B-BVAR and L-BVAR models, respectively, scaled to deliver the same impact GDP response as the Q-BVAR model. The dynamics are in line with standard results from quarterly VARs for the US economy (see for example [Del Negro et al., 2020](#)).

the Federal Funds rate for about three years, which reflects the systematic component of US monetary policy. As for other variables, there is a short-lived positive spike in the PMI index, business loans increase rather persistently after the shock, while the BAA spread and the uncertainty index drop, but only for a few quarters.

Besides describing the dynamics of the US economy, this exercise also allows us to draw some insights on the relevance of mixed-frequency data. In VAR models, the individual estimated parameters, especially in high-dimensional models such as those we consider in this paper, cannot be easily used to assess the similarity across models. However, our generalized impulse responses convolve the estimated VAR parameters and thus greatly facilitate such comparisons across models. Figure 4 shows that most impulse responses are very similar across the three mixed-frequency methodologies, including those of the C-BVAR, which is estimated exclusively on quarterly variables and, moreover, that they are also similar to those of a quarterly model. This suggests that the potential bias implied by temporal aggregation on the analysis of economic dynamics (e.g. Sims (1971), Hansen, Sargent et al. (1981), Christiano and Eichenbaum (1986)) is negligible, and mixed-frequency data do not help to uncover dynamic relations among variables that would be otherwise obscured by temporal aggregation. Hence, mixed-frequency data are mainly important for their timeliness.²⁶

4.3 The real-time evolution of the 2008Q4 Fed Funds shadow rate


Central banks routinely use counterfactual interest rate paths as benchmarks to gauge whether their policy rates, and the closely related short-term money market rates, are at reasonable levels given prevailing and expected economic conditions.²⁷ In this vein, we use our mixed frequency VAR framework to estimate the level of the Fed Funds rate compatible with US economic conditions. We focus on the level of the Fed Funds rate, a measure of the Fed's monetary policy stance, in the fourth quarter of 2008 because that was the first quarter in which the actual Fed Funds rate hit the zero lower bound due to the intensification of the 2007-2009 global financial crisis. Specifically, we ask the question at which point, in the course of 2008, a VAR analysis would have revealed that the Fed Funds rate was going to head decisively toward or even below zero. The assessment is carried out for each weekly data vintage of 2008 included in our real-time database.

Traditionally, the analysis of benchmark counterfactual rates has been based on the Taylor rule framework (see Taylor, 1993), which relates the level of the short-term interest rate to inflation and a measure of real economic activity (for recent examples, see Bernanke, 2015; Nechio, 2011; Hartmann and Smets, 2018). At the same time, in their monetary policy briefings, central

²⁶This result is corroborated also by re-running the nowcasting evaluation exercise of the previous Section, but only feeding the mixed-frequency models information that could also be incorporated in a quarterly model, i.e. only full quarterly data. When doing so, the mixed-frequency models' improvement in forecasting performance from weeks 5-6 onward all but disappears, and all three models perform very similarly to the Q-BVAR.

²⁷Since the zero lower bound has been reached in many countries, these exercises have taken a different twist, whereby a very negative counterfactual policy rate is taken as an indication that additional accommodation by means of non-conventional policy tools may be warranted (for example Giannone et al., 2019c).

Figure 5: Counterfactual Fed Funds Rate for 2008Q4



Note: Horizontal axis: data vintages used to compute the forecast of the 2008Q4 Fed Funds rate. Vertical axis: value of the (counterfactual) Fed Funds rate in percentage points. The fan chart bands cover 99% of the support around the median: the darkest shade of blue corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The black dashed line is the median of the Q-BVAR counterfactual Fed Funds rate estimated on the same vintages.

banks rely on many different sources of information, so that their assessment of economic conditions can be well-characterised as a Big Data problem (see, for example, [Giannone et al., 2005](#); [Bernanke et al., 2005](#)). Our VAR models are well equipped to capture this idea, given that they include a relatively large amount of information. Moreover, their ability to deal with mixed-frequency data and, hence, to account in a more timely fashion for incoming information potentially relevant for the setting of the Fed Funds rate, allows the assessment of the benchmark policy rate to be based on the latest news on US economic conditions.

The counterfactual path of the short-term interest rate estimates we derive for 2008Q4 is one that would be compatible with the developments in the US economy and the historical monetary policy rule implicit in our VAR estimates, which may be thought as a generalization of the Taylor rule. To derive this path, we assume that the data on the Fed Funds rate for 2008Q4 were missing even when, over the course of the fourth quarter of 2008, such data started to become available (at monthly frequency). Figure 5 reports the level of the counterfactual Fed Funds rate for 2008Q4, conditional on the information available at the time of the analysis (over the 52 weeks in 2008, with dates reported on the horizontal axis). The results are presented in the form of a fan chart, where areas closer to the median are indicated by a darker colour. For the sake of brevity, the results refer to the B-BVAR, but the L-BVAR and C-BVAR results are very similar. As a term of comparison, we also report the median estimates of the counterfactual Fed Funds rate from the Q-BVAR.

The mixed-frequency VAR would have led to a more timely assessment of the deterioration of economic conditions in the US economy and, consequently, suggested that the level of the Fed Funds rate would head toward very low levels ahead of the quarterly VAR.

5 Conclusions

This paper has shown that BVARs can be successfully used to handle Big Data – i.e., a large set of macroeconomic time series with different frequencies, staggered release dates, and various other irregularities – for the purpose of real-time nowcasting.

BVARs are more tractable and have several other advantages compared to competing nowcasting methods, most notably Dynamic Factor Models. For example, they have a more general structure and do not assume that shocks affect all variables in the model at the same time, face the econometrician with less stark modelling choices (e.g., related to the number of lags, the block-structure, etc.), or do not require the data to be made stationary.

We present three strategies for dealing with mixed-frequency data in the context of VARs: first, a model – labelled “latent BVAR” – which assumes that all variables are high-frequency time series, but that some of them are observed only at low frequency. Second, a methodology known as “blocking,” which treats higher-frequency data as multiple lower-frequency variables. Third, we use the estimates of a standard low-frequency VAR to update a higher-frequency model, and refer to this latter approach as “cube-root BVAR.”

Based on a sample of real-time data from the beginning of 2005 to the end of 2019, we show that these models would have nowcasted U.S. GDP growth as well as established benchmarks such as the New York Fed’s Dynamic Factor Model, displaying a clear improvement in forecast accuracy as the quarter progresses and more information becomes available. We also find that all the models produce similar predictions and impulse response functions, which indicates that they all capture genuine data features.

We also discuss the advantages and drawbacks of each of these approaches. Comparing the impulse response functions from these three models, which are very similar, we also infer that the importance of incorporating mixed-frequency data derives mainly from their timeliness. Indeed, models which do not rely on mixed-frequency data capture similar economic dynamics as the mixed-frequency approaches, implying that the potential bias due to temporal aggregation is negligible in practice. Finally, we show that mixed-frequency BVARs are also powerful tools for policy analysis, and can be used to evaluate the dynamic impact of shocks and to construct counterfactual scenarios, which increases their appeal as operational tools in central banks and international organisations.

References

- ALTAVILLA, C., D. GIANNONE, AND M. LENZA (2016): “The Financial and Macroeconomic Effects of the OMT Announcements,” *International Journal of Central Banking*, 12, 29–57.
- ANDERSON, B. D., M. DEISTLER, E. FELSENSTEIN, B. FUNOVITS, L. KOELBL, AND M. ZAMANI (2016a): “Multivariate AR systems and mixed frequency data: G-identifiability and estimation,” *Econometric Theory*, 32, 793–826.
- ANDERSON, B. D., M. DEISTLER, E. FELSENSTEIN, AND L. KOELBL (2016b): “The structure of multivariate AR and ARMA systems: Regular and singular systems; the single and the mixed frequency case,” *Journal of econometrics*, 192, 366–373.
- ANGELETOS, G. M., F. COLLARD, AND H. DELLAS (2020): “Business Cycle Anatomy,” TSE Working Papers 20-1065, Toulouse School of Economics (TSE).
- ANGELINI, E., M. LALIK, M. LENZA, AND J. PAREDES (2019): “Mind the gap: A multi-country BVAR benchmark for the Eurosystem projections,” *International Journal of Forecasting*, 35, 1658–1668.
- ANTOLIN-DIAZ, J., T. DRECHSEL, AND I. PETRELLA (2017): “Tracking the slowdown in long-run GDP growth,” *Review of Economics and Statistics*, 99, 343–356.
- ARUOBA, S. B., F. X. DIEBOLD, AND C. SCOTTI (2009): “Real-time measurement of business conditions,” *Journal of Business & Economic Statistics*, 27, 417–427.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring economic policy uncertainty,” *The quarterly journal of economics*, 131, 1593–1636.
- BANBURA, 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*, ed. by G. Elliott and A. Timmermann, Elsevier, vol. 2, 195–237.
- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): “Large Bayesian vector auto regressions,” *Journal of Applied Econometrics*, 25, 71–92.
- (2011): “Nowcasting,” in *The Oxford Handbook of Economic Forecasting*, ed. by M. P. Clements and D. F. Hendry, Oxford University Press, 193–224.
- BARIGOZZI, M., M. LIPPI, AND M. LUCIANI (2016): “Non-Stationary Dynamic Factor Models for Large Datasets,” Finance and Economics Discussion Series 2016-024, Board of Governors of the Federal Reserve System (U.S.).

- BERNANKE, B. (2015): “The Taylor Rule: A benchmark for monetary policy?” *Blog post available at <https://www.brookings.edu/blog/ben-bernanke/2015/04/28/the-taylor-rule-a-benchmark-for-monetary-policy/>*.
- BERNANKE, B. S., J. BOIVIN, AND P. ELIASZ (2005): “Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach,” *The Quarterly Journal of Economics*, 120, 387–422.
- BITTANTI, S. (1986): “Deterministic and stochastic linear periodic systems,” in *Time series and linear systems*, Springer, 141–182.
- BITTANTI, S. AND P. COLANERI (2009): *Periodic systems: filtering and control*, vol. 5108985, Springer Science & Business Media.
- BLOOM, N. (2009): “The impact of uncertainty shocks,” *econometrica*, 77, 623–685.
- BOK, B., D. CARATELLI, D. GIANNONE, A. M. SBORDONE, AND A. TAMBALOTTI (2018): “Macroeconomic nowcasting and forecasting with big data,” *Annual Review of Economics*, 10, 615–643.
- BRAVE, S. A., R. A. BUTTERS, AND A. JUSTINIANO (2019): “Forecasting economic activity with mixed frequency BVARs,” *International Journal of Forecasting*, 35, 1692–1707.
- BURNS, A. F. AND W. C. MITCHELL (1946): *Measuring Business Cycles*.
- CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2015): “Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178, 837–862.
- (2016): “Common Drifting Volatility in Large Bayesian VARs,” *Journal of Business & Economic Statistics*, 34, 375–390.
- CARRIERO, A., T. E. CLARK, AND M. MASSIMILIANO (2020): “Nowcasting Tail Risks to Economic Activity with Many Indicators,” Working Papers 202013R2, Federal Reserve Bank of Cleveland.
- CHEN, W., B. D. ANDERSON, M. DEISTLER, AND A. FILLER (2011): “Properties of Blocked Linear Systems,” *IFAC Proceedings Volumes*, 44, 4558 – 4563, 18th IFAC World Congress.
- CHRISTIANO, L. J. AND M. EICHENBAUM (1986): “Temporal aggregation and structural inference in macroeconomics,” *NBER Working paper*.
- CIMADOMO, J. AND A. D’AGOSTINO (2016): “Combining Time Variation and Mixed Frequencies: an Analysis of Government Spending Multipliers in Italy,” *Journal of Applied Econometrics*, 31, 1276–1290.

- DE MOL, C., D. GIANNONE, AND L. REICHLIN (2008): “Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components?” *Journal of Econometrics*, 146, 318–328.
- DEL NEGRO, M., M. LENZA, G. E. PRIMICERI, AND A. TAMBALOTTI (2020): “What’s up with the Phillips Curve?” *Brookings Papers on Economic Activity*.
- DIEBOLD, F. (2003): “Macroeconomic forecasting using many predictors: A Discussion of the Papers by Reichlin and Watson,” in *Advances in Economics and Econometrics Theory and Applications, Eighth World Congress*, Cambridge University Press, 115–122.
- DIEBOLD, F. X. (2012): “On the Origin (s) and Development of the Term Big Data,” PIER Working Paper No. 12-037.
- DOAN, T., R. LITTERMAN, AND C. SIMS (1984): “Forecasting and conditional projection using realistic prior distributions,” *Econometric reviews*, 3, 1–100.
- DOMIT, S., F. MONTI, AND A. SOKOL (2019): “Forecasting the UK economy with a medium-scale Bayesian VAR,” *International Journal of Forecasting*, 35, 1669–1678.
- DOZ, C. AND P. FULEKY (2019): “Dynamic Factor Models,” Working Papers 2019-4, University of Hawaii Economic Research Organization, University of Hawaii at Manoa.
- DOZ, C., D. GIANNONE, AND L. REICHLIN (2012): “A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models,” *The Review of Economics and Statistics*, 94, 1014–1024.
- DURBIN, J. AND S. J. KOOPMAN (2001): *Time Series Analysis by State Space Methods*, no. 9780198523543 in OUP Catalogue, Oxford University Press.
- D’AGOSTINO, A., D. GIANNONE, M. LENZA, AND M. MODUGNO (2016): “Nowcasting Business Cycles: A Bayesian Approach to Dynamic Heterogeneous Factor Models,” in *Dynamic Factor Models*, ed. by E. Hillebrand and S. J. Koopman, Emerald Publishing Ltd, vol. 35 of *Advances in Econometrics*, 569–594.
- ELLAHIE, A. AND G. RICCO (2017): “Government purchases reloaded: Informational insufficiency and heterogeneity in fiscal VARs,” *Journal of Monetary Economics*, 90, 13–27.
- ENGLE, R. AND M. WATSON (1981): “A one-factor multivariate time series model of metropolitan wage rates,” *Journal of the American Statistical Association*, 76, 774–781.
- ERAKER, B., C. W. CHIU, A. T. FOERSTER, T. B. KIM, AND H. D. SEOANE (2014): “Bayesian mixed frequency VARs,” *Journal of Financial Econometrics*, 13, 698–721.
- FERNANDES, P., B. PLATEAU, AND W. J. STEWART (1998): “Efficient Descriptor-Vector Multiplications in Stochastic Automata Networks,” *J. ACM*, 45, 381–414.

- FORNI, M., M. HALLIN, M. LIPPI, AND L. REICHLIN (2000): “The Generalized Dynamic-Factor Model: Identification And Estimation,” *The Review of Economics and Statistics*, 82, 540–554.
- FORONI, C., P. GUÉRIN, AND M. MARCELLINO (2015): “Markov-switching mixed-frequency VAR models,” *International Journal of Forecasting*, 31, 692–711.
- GEWEKE, J. (1977): “Dynamic Factor Analysis of Economic Time Series. In Latent Variable in Socio! Economic Models, eds. by DJ Aigner and AS Goldberger,” .
- GHYSELS, E. (2016): “Macroeconomics and the reality of mixed frequency data,” *Journal of Econometrics*, 193, 294–314.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): “Prior Selection for Vector Autoregressions,” *The Review of Economics and Statistics*, 97, 436–451.
- (2019a): “Priors for the Long Run,” *Journal of the American Statistical Association*, 114, 565–580.
- GIANNONE, D., M. LENZA, AND L. REICHLIN (2019b): “Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis?” *International Journal of Central Banking*, 15, 137–173.
- (2019c): “Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis?” *International Journal of Central Banking*, 15, 137–173.
- GIANNONE, D., F. MONTI, AND L. REICHLIN (2016): “Exploiting the monthly data flow in structural forecasting,” *Journal of Monetary Economics*, 84, 201–215.
- GIANNONE, D., L. REICHLIN, AND L. SALA (2005): “Monetary Policy in Real Time,” in *NBER Macroeconomics Annual 2004, Volume 19*, National Bureau of Economic Research, Inc, NBER Chapters, 161–224.
- GIANNONE, D., L. REICHLIN, AND S. SIMONELLI (2009): “Nowcasting Euro Area Economic Activity In Real Time: The Role Of Confidence Indicators,” *National Institute Economic Review*, 210, 90–97.
- GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics*, 55, 665–676.
- HANSEN, L. P., T. J. SARGENT, ET AL. (1981): “Exact linear rational expectations models: Specification and estimation,” Tech. rep., Federal Reserve Bank of Minneapolis.
- HARTMANN, P. AND F. SMETS (2018): “The first twenty years of the European Central Bank: monetary policy,” *Brookings Papers on Economic Activity*.

- KARLSSON, S. (2013): “Forecasting with Bayesian Vector Autoregression,” in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, Elsevier, vol. 2 of *Handbook of Economic Forecasting*, chap. 0, 791–897.
- KILIAN, L. AND H. LÜTKEPOHL (2018): *Structural Vector Autoregressive Analysis*, no. 9781107196575 in Cambridge Books, Cambridge University Press.
- KOOP, G. (2017): “Bayesian Methods for Empirical Macroeconomics,” *Review of Economic Analysis*, 9, 33–56.
- KUZIN, V., M. MARCELLINO, AND C. SCHUMACHER (2011): “MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area,” *International Journal of Forecasting*, 27, 529–542.
- LITTERMAN, R. B. (1979): “Techniques of forecasting using vector autoregressions,” Tech. rep.
- LUCIANI, M. (2017): “Large-Dimensional Dynamic Factor models in Real-Time: A Survey,” in *Handbook on Cyclical Composite Indicators*, ed. by G. L. Mazzi and A. Ozyildirim, Eurostat, 429–451.
- MARIANO, R. S. AND Y. MURASAWA (2010): “A Coincident Index, Common Factors, and Monthly Real GDP,” *Oxford Bulletin of Economics and Statistics*, 72, 27–46.
- MCCRACKEN, M. W., M. T. OWYANG, AND T. SEKHOSYAN (forthcoming): “Real-Time Forecasting with a Large, Mixed Frequency, Bayesian VAR,” Tech. rep.
- MIRANDA-AGRIPPINO, S. AND H. REY (2020): “U.S. Monetary Policy and the Global Financial Cycle,” *The Review of Economic Studies*, 87, 2754–2776.
- MIRANDA-AGRIPPINO, S. AND G. RICCO (2018): “Bayesian vector autoregressions,” Bank of England working papers 756, Bank of England.
- MITTNIK, S. AND P. A. ZADROZNY (2004): “Forecasting Quarterly German GDP at Monthly Intervals Using Monthly IFO Business Conditions Data,” CESifo Working Paper Series 1203, CESifo Group Munich.
- NECHIO, F. (2011): “Monetary policy when one size does not fit all,” *FRBSF Economic Letter*.
- PESARAN, H. H. AND Y. SHIN (1998): “Generalized impulse response analysis in linear multivariate models,” *Economics Letters*, 58, 17–29.
- PETTENUZZO, D., A. TIMMERMANN, AND R. VALKANOV (2016): “A MIDAS approach to modeling first and second moment dynamics,” *Journal of Econometrics*, 193, 315–334.
- SARGENT, T. AND C. SIMS (1977): “Business Cycle Modeling Without Pretending to Have Too Much a Priori Economic Theory,” in *New Methods in Business Cycle Research: Proceedings From a Conference*, Federal Reserve Bank of Minneapolis, 45–109.

- SCHORFHEIDE, F. AND D. SONG (2015): “Real-Time Forecasting With a Mixed-Frequency VAR,” *Journal of Business & Economic Statistics*, 33, 366–380.
- SIMS, C. A. (1971): “Discrete approximations to continuous time distributed lags in econometrics,” *Econometrica: Journal of the Econometric Society*, 545–563.
- (1980): “Macroeconomics and reality,” *Econometrica: journal of the Econometric Society*, 1–48.
- SIMS, C. A., J. H. STOCK, AND M. W. WATSON (1990): “Inference in Linear Time Series Models with Some Unit Roots,” *Econometrica*, 58, 113–144.
- STOCK, J. AND M. WATSON (2016): “Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2 of *Handbook of Macroeconomics*, chap. 0, 415–525.
- STOCK, J. H. AND M. W. WATSON (1999): “Forecasting inflation,” *Journal of Monetary Economics*, 44, 293–335.
- (2001): “Vector Autoregressions,” *Journal of Economic Perspectives*, 15, 101–115.
- (2017): “Twenty years of time series econometrics in ten pictures,” *Journal of Economic Perspectives*, 31, 59–86.
- TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- WEST, M. (2002): “Bayesian factor regression models in the “large p, small n” paradigm,” Tech. rep., Bayesian Statistics.
- ZADROZNY, P. (1990): “Estimating a Multivariate ARMA Model with Mixed-Frequency Data: An Application to Forecasting US GNP at Monthly Intervals,” Working Paper Series 90-6, Federal Reserve Bank of Atlanta.
- ZAMANI, M., W. CHEN, B. D. O. ANDERSON, M. DEISTLER, AND A. FILLER (2011): “On the zeros of blocked linear systems with single and mixed frequency data,” in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 4312–4317.

A C-BVAR: Detailed derivation

Start from the quarterly $VAR(p)$ model of equation (1), rewritten in companion form:

$$X_{t_q} = \Phi X_{t_q-1} + \nu_{t_q}, \quad (\text{A.1})$$

with $X_{t_q} = (x'_{t_q}, \dots, x'_{t_q-p+1})'$, $\nu_{t_q} = (\varepsilon'_{t_q}, \mathbf{0}_{1 \times n(p-1)})'$, $\nu \sim \mathcal{N}(0, \Omega)$ and

$$\Phi = \begin{bmatrix} A_1 & A_2 & \cdots & A_p \\ I_n & 0_n & \cdots & 0_n \\ 0_n & \ddots & \cdots & 0_n \\ 0_n & \cdots & I_n & 0_n \end{bmatrix} \quad \Omega = \begin{bmatrix} \Sigma_\varepsilon & 0_n & \cdots & 0_n \\ 0_n & \cdots & \ddots & 0_n \\ 0_n & 0_n & \cdots & 0_n \end{bmatrix}.$$

Model (A.1) can also be rewritten in terms of monthly quantities as

$$X_{t_m} = \Phi X_{t_m-3} + \nu_{t_m} \quad (\text{A.2})$$

when t_m corresponds to the last month of a quarter. Assume that the *monthly* counterpart of model (A.1) can be written in state-space form as

$$X_{t_m} = \Phi_m X_{t_m-1} + \nu_{m,t_m}, \quad (\text{A.3})$$

with $\nu_{m,t_m} = (\varepsilon'_{m,t_m}, \mathbf{0}_{1 \times n(p-1)})'$, $\nu_m \sim \mathcal{N}(0, \Omega_m)$ and

$$\Phi_m = \begin{bmatrix} \Phi_{m11} & \Phi_{m12} & \cdots & \Phi_{m1p} \\ \Phi_{m21} & \ddots & & \\ \vdots & & & \\ \Phi_{mp1} & & & \Phi_{mpp} \end{bmatrix} \quad \Omega_m = \begin{bmatrix} \Sigma_{\varepsilon_m} & 0_n & \cdots & 0_n \\ 0_n & \cdots & \ddots & 0_n \\ 0_n & 0_n & \cdots & 0_n \end{bmatrix}.$$

Also assume that the elements of Φ_m are real and stable.

The first n rows of system (A.3) correspond to a restricted monthly VAR of the following form:

$$x_{t_m} = \Phi_{m11}x_{t_m-1} + \Phi_{m12}x_{t_m-4} + \cdots + \Phi_{m1p}x_{t_m-3p+2} + \varepsilon_{m,t_m}, \quad (\text{A.4})$$

The restriction is that current (monthly) values only depend on one month within each lagged quarter. The remaining rows impose restrictions on how the (possibly latent) lagged monthly states are updated each month with the arrival of new information. They imply that the lagged states on the left-hand side also depend on future states on the right-hand side. Intuitively, this happens because our assumptions require the states of the monthly model to match those of the quarterly one at the end of each quarter, and thus all latent states within a quarter need to be updated with the arrival of new information.

Iteration of (A.3) implies that

$$X_{t_m} = \Phi_m^3 X_{t_m-3} + \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}, \quad (\text{A.5})$$

which together with our previous assumptions entails the following relationships between the quarterly model (A.2) and the monthly model (A.3):

$$\Phi_m = \Phi^{\frac{1}{3}} \quad (\text{A.6})$$

$$\nu_{t_m} = \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}. \quad (\text{A.7})$$

Equation (A.6) implies that an essential part of finding a suitable mapping is computing the cube root of Φ , which raises the issue of multiple solutions. If the autoregressive matrix of the transition equation is diagonalizable,²⁸ *i.e.* if there exist a diagonal matrix D and an invertible matrix V such that $\Phi = VDV^{-1}$, then the cube root of Φ can be obtained as $\Phi = VD^{\frac{1}{3}}V^{-1}$, where $D^{\frac{1}{3}}$ is a diagonal matrix containing the cube roots of the elements of D . The real elements of D , which are associated with real-valued eigenvectors, have a unique real cube root, which is the only one that gives rise to real values when combined with its associated eigenvector. Complex conjugate eigenvalues instead have three complex cube roots. When combined with their associated eigenvector, these still return a real-valued vector. Thus, if k is the number of complex conjugate couples of eigenvalues in D , then there will be 3^k real-valued cube roots for Φ . We follow the procedure proposed in [Giannone, Monti, and Reichlin \(2016\)](#) to select among these alternative cube roots of Φ : in the case of real eigenvalues, simply select their real cube root; in the case of complex conjugate couples, choose the cube root which is characterized by the least oscillatory behaviour, *i.e.*, the cube root with the smallest argument. An alternative is to evaluate the likelihood of solutions using the Kalman filter and pick the one with the highest likelihood²⁹, though this is more computationally intensive. In the cases where we have tried it, it corresponds to the one with the roots with the smallest argument, as in [Giannone, Monti, and Reichlin \(2016\)](#).

On the other hand, equation (A.7) imposes a series of restrictions on the behaviour of the monthly residuals ν_{m,t_m} . To see that, it's useful to write it out explicitly:

$$\begin{aligned} \varepsilon_{t_m} &= \varepsilon_{m,t_m} + \Phi_{m11}\varepsilon_{m,t_m-1} + \Phi_{m11}^2\varepsilon_{m,t_m-2} \\ \mathbf{0} &= \Phi_{m21}\varepsilon_{m,t_m-1} + \Phi_{m21}^2\varepsilon_{m,t_m-2} \\ \mathbf{0} &= \Phi_{m31}\varepsilon_{m,t_m-1} + \Phi_{m31}^2\varepsilon_{m,t_m-2} \\ &\vdots \\ \mathbf{0} &= \Phi_{mp1}\varepsilon_{m,t_m-1} + \Phi_{mp1}^2\varepsilon_{m,t_m-2} \end{aligned}$$

The last $n(p-1)$ rows constitute an over-determined system of linear equations that, if $\Phi_{m-1} \neq \mathbf{0}$,

²⁸For the non-diagonalizable case, see the discussion in [Giannone, Monti, and Reichlin \(2016\)](#)

²⁹[Anderson et al. \(2016b\)](#) show g-identifiability when (enough) high frequency data is available.

can be approximately solved for ε_{m,t_m-1} as follows:

$$\varepsilon_{m,t_m-1} = - (J'J)^{-1} J' \Phi_{m \cdot 1} \varepsilon_{m,t_m-2} \quad (\text{A.8})$$

with $J = [I_n \dots I_n]'$ and $\Phi_{m \cdot 1} = [\Phi'_{m21} \dots \Phi'_{mp1}]'$.

Substituting ε_{m,t_m-1} , as solved in (A.8), into the first n rows of (A.7) allows to recover the monthly covariance matrix Σ_{ε_m} , and therefore also Ω_m , from

$$\text{vec}(\Sigma_{\varepsilon_m}) = \left\{ I + \begin{bmatrix} \Phi_{m11}^2 - \Phi_{m11} (J'J)^{-1} J' \Phi_{m \cdot 1} \\ \Phi_{m11}^2 - \Phi_{m11} (J'J)^{-1} J' \Phi_{m \cdot 1} \end{bmatrix} \otimes \right\}^{-1} \text{vec}(\Sigma_{\varepsilon}). \quad (\text{A.9})$$

The solution of (A.9) can become computationally costly as the number of variables increases, as it involves the inversion of an $n^2 \times n^2$ matrix. However, it can be greatly simplified by noting that the inverse is of the form

$$(I + A \otimes A)^{-1}$$

with $A = \begin{bmatrix} \Phi_{m11}^2 - \Phi_{m11} (J'J)^{-1} J' \Phi_{m \cdot 1} \end{bmatrix}$. As long as A is diagonalizable, that is, as long as there is a diagonal matrix Λ and an invertible matrix P such that $A = P\Lambda P^{-1}$, the inverse can be computed as³⁰

$$(I + A \otimes A)^{-1} = (P \otimes P) (I + \Lambda \otimes \Lambda)^{-1} (P^{-1} \otimes P^{-1}) \quad (\text{A.10})$$

which is much more appealing, since $(I + \Lambda \otimes \Lambda)$ is diagonal and thus its inverse is trivial to compute directly, and some of the multiplications in (A.10) can be carried out without explicitly computing the Kronecker product (see [Fernandes et al., 1998](#)).

A.1 AR(2) example

The simplest model our C-BVAR framework applies to is a quarterly AR(2) model³¹, which can be written in companion form as:

$$\begin{bmatrix} y_t \\ y_{t-3} \end{bmatrix} = \Phi \begin{bmatrix} y_{t-3} \\ y_{t-6} \end{bmatrix} + \begin{bmatrix} \varepsilon_{t_q} \\ 0 \end{bmatrix} \quad (\text{A.11})$$

³⁰The result follows from the properties of the Kronecker product. If A is diagonalizable, then

$$A \otimes A = (P\Lambda P^{-1}) \otimes (P\Lambda P^{-1}) = (P \otimes P) (\Lambda \otimes \Lambda) (P^{-1} \otimes P^{-1}).$$

Furthermore, because $(P \otimes P)^{-1} = P^{-1} \otimes P^{-1}$, it follows that

$$I + A \otimes A = (P \otimes P) (I + \Lambda \otimes \Lambda) (P^{-1} \otimes P^{-1})$$

and also that

$$(I + A \otimes A)^{-1} = (P \otimes P) (I + \Lambda \otimes \Lambda)^{-1} (P^{-1} \otimes P^{-1}).$$

³¹For models with only one quarterly lag, the results in [Giannone et al. \(2016\)](#) apply directly.

$$\Phi = \begin{bmatrix} \phi_1 & \phi_2 \\ 1 & 0 \end{bmatrix} \quad \varepsilon_{t_q} \sim \mathcal{N}(0, \sigma_\varepsilon)$$

Define the cube-root matrix Φ_m :

$$\Phi_m \equiv \begin{bmatrix} \phi_{m11} & \phi_{m12} \\ \phi_{m21} & \phi_{m22} \end{bmatrix} = \Phi^{\frac{1}{3}} \quad (\text{A.12})$$

For future reference:

$$\text{vec}(\Phi_m^2) = \begin{bmatrix} \phi_{m11}^2 + \phi_{m12}\phi_{m21} \\ \phi_{m21}(\phi_{m11} + \phi_{m22}) \\ \phi_{m12}(\phi_{m11} + \phi_{m22}) \\ \phi_{m12}\phi_{m21} + \phi_{m22}^2 \end{bmatrix}$$

$$\text{vec}(\Phi_m^3) = \begin{bmatrix} \phi_{m11}(\phi_{m11}^2 + \phi_{m12}\phi_{m21}) + \phi_{m12}\phi_{m21}(\phi_{m11} + \phi_{m22}) \\ \phi_{m21}(\phi_{m11}^2 + \phi_{m12}\phi_{m21}) + \phi_{m21}\phi_{m22}(\phi_{m11} + \phi_{m22}) \\ \phi_{m11}\phi_{m12}(\phi_{m11} + \phi_{m22}) + \phi_{m12}(\phi_{m12}\phi_{m21} + \phi_{m22}^2) \\ \phi_{m12}\phi_{m21}(\phi_{m11} + \phi_{m22}) + \phi_{m22}(\phi_{m12}\phi_{m21} + \phi_{m22}^2) \end{bmatrix} = \begin{bmatrix} \phi_1 \\ 1 \\ \phi_2 \\ 0 \end{bmatrix} = \text{vec}(\Phi)$$

Using Φ_m we can posit the following model:

$$\begin{bmatrix} x_t \\ x_{t-3} \end{bmatrix} = \begin{bmatrix} \phi_{m11} & \phi_{m12} \\ \phi_{m21} & \phi_{m22} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ x_{t-4} \end{bmatrix} + \begin{bmatrix} \varepsilon_{m,t_m} \\ 0 \end{bmatrix} \quad (\text{A.13})$$

This is a model where x_t behaves like a monthly VAR with some restrictions – namely that it only depends on one month within each lagged quarter – and x_{t-3} gets updated at each iteration in a way consistent with satisfying the cube root relationship. Iterating backwards we get

$$\begin{bmatrix} x_t \\ x_{t-3} \end{bmatrix} = \Phi_m^3 \begin{bmatrix} x_{t-3} \\ x_{t-6} \end{bmatrix} + \begin{bmatrix} \varepsilon_{m,t_m} \\ 0 \end{bmatrix} + \Phi_m \begin{bmatrix} \varepsilon_{m,t_m-1} \\ 0 \end{bmatrix} + \Phi_m^2 \begin{bmatrix} \varepsilon_{m,t_m-2} \\ 0 \end{bmatrix} \quad (\text{A.14})$$

Writing out the first row:

$$\begin{aligned} x_t &= [(\phi_{m11}^2 + \phi_{m12}\phi_{m21})\phi_{m11} + (\phi_{m11}\phi_{m12} + \phi_{m12}\phi_{m22})\phi_{m21}]x_{t-3} + \\ &+ [(\phi_{m11}^2 + \phi_{m12}\phi_{m21})\phi_{m12} + (\phi_{m11}\phi_{m12} + \phi_{m12}\phi_{m22})\phi_{m22}]x_{t-6} + \\ &+ \varepsilon_{m,t_m} + \phi_{m11}\varepsilon_{m,t_m-1} + (\phi_{m11}^2 + \phi_{m12}\phi_{m21})\varepsilon_{m,t_m-2} \\ &= \phi_1 x_{t-3} + \phi_2 x_{t-6} + \varepsilon_{m,t_m} + \phi_{m11}\varepsilon_{m,t_m-1} + (\phi_{m11}^2 + \phi_{m12}\phi_{m21})\varepsilon_{m,t_m-2} \end{aligned}$$

The quarterly residual ε_{t_q} is thus a moving average of the monthly residuals within the quarter.

Writing out the second row:

$$\begin{aligned}
x_{t-3} &= [\phi_{m11}^2 \phi_{m21} + \phi_{m11} \phi_{m21} \phi_{m22} + \phi_{m12} \phi_{m21}^2 + \phi_{m21} \phi_{m22}^2] x_{t-3} + \\
&+ [\phi_{m11} \phi_{m12} \phi_{m21} + \phi_{m12} \phi_{m21} \phi_{m22} + \phi_{m12} \phi_{m21} \phi_{m22} + \phi_{m22}^3] x_{t-6} + \\
&+ (\phi_{m11} \phi_{m21} + \phi_{m21} \phi_{m22}) \varepsilon_{m,t_m-2} + \phi_{m21} \varepsilon_{m,t_m-1} \\
&= x_{t-3} + (\phi_{m11} \phi_{m21} + \phi_{m21} \phi_{m22}) \varepsilon_{m,t_m-2} + \phi_{m21} \varepsilon_{m,t_m-1}
\end{aligned}$$

This expression entails a restriction on the monthly shocks (for $\phi_{m21} \neq 0$):

$$\varepsilon_{m,t_m-1} = -(\phi_{m11} + \phi_{m22}) \varepsilon_{m,t_m-2}$$

Substituting into the expression for x_t :

$$x_t = \phi_1 x_{t-3} + \phi_2 x_{t-6} + [\phi_{m12} \phi_{m21} - \phi_{m11} \phi_{m22}] \varepsilon_{m,t_m-2} + \varepsilon_{m,t_m}$$

this implies the restriction that

$$\varepsilon_{t_q} = \varepsilon_{m,t_m} + [\phi_{m12} \phi_{m21} - \phi_{m11} \phi_{m22}] \varepsilon_{m,t_m-2}$$

and that we can solve for $\text{var}(\varepsilon_{m,t_m})$ from

$$\text{var}(\varepsilon_{m,t_m}) = [1 + (\phi_{m12} \phi_{m21} - \phi_{m11} \phi_{m22})^2]^{-1} \text{var}(\varepsilon_{t_q}) \quad (\text{A.15})$$

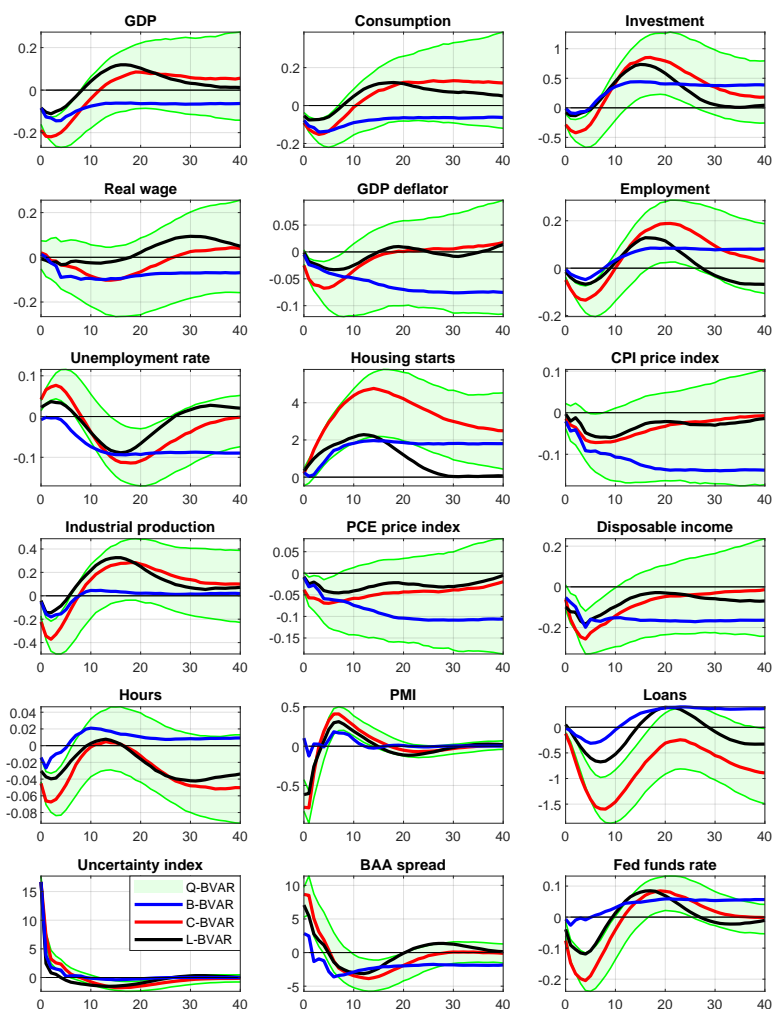
B Hyperparameter estimates

Table B.1: Selected hyperparameter estimates across vintages

	2005-2009	2010-2014	2016-2019	Whole sample
λ				
B-BVAR	0.96	0.61	0.68	0.75
C-BVAR	0.68	0.68	0.63	0.66
L-BVAR	0.63	0.58	0.56	0.59
μ				
B-BVAR	0.92	1.33	0.66	0.97
C-BVAR	0.65	1.35	1.40	1.13
L-BVAR	0.80	1.92	2.43	1.72

Note: The table shows averages across vintages of the posterior modes of the hyperparameters λ and μ , governing the tightness of the Minnesota and sum-of-coefficients priors, respectively.

C Impulse response functions to an uncertainty shock



Note: Impulse response function (IRF) to a one standard deviation shock to uncertainty (in the first month of the quarter for the mixed-frequency models), identified as in [Bloom \(2009\)](#), shown at quarterly frequency. The green areas represent the 68% credible intervals for a one standard deviation impulse to the Q-BVAR model; the red, blue and black lines show the median IRFs for the C-BVAR, B-BVAR and L-BVAR models, respectively, scaled to deliver the same impact uncertainty index response as the Q-BVAR model.