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

Nowcasting with large Bayesian
vector autoregressions

Abstract

Monitoring economic conditions in real time, or nowcasting, is among the key tasks routinely performed by economists. Nowcasting entails some key challenges, which also characterise modern Big Data analytics, often referred to as the three "Vs": the large number of time series continuously released (Volume), the complexity of the data covering various sectors of the economy, published in an asynchronous way and with different frequencies and precision (Variety), and the need to incorporate new information within minutes of their release (Velocity). In this paper, we explore alternative routes to bring Bayesian Vector Autoregressive (BVAR) models up to these challenges. We find that BVARs are able to effectively handle the three Vs and produce, in real time, accurate probabilistic predictions of US economic activity and, in addition, a meaningful narrative by means of scenario analysis.

JEL Classification: E32, E37, C01, C33, C53.
Keywords: Big Data, Scenario Analysis, Mixed Frequencies, Real Time, Business Cycles, Forecasting.

Non-Technical Summary

Policymakers and market participants need to extract timely and accurate signals about the state of the economy, that is, to "nowcast", and the economic fallout of the Covid-19 pandemic has highlighted, once again, the importance of monitoring macroeconomic conditions in real time. Nowcasting entails some key challenges related to Big Data analytics, which are generally referred to as the three "Vs": the large number of time series continuously released (Volume), the complexity of the data covering various sectors of the economy, published in an asynchronous way and with different frequencies and precision (Variety), and the need to incorporate new information within minutes of their release (Velocity). Such challenges have been traditionally addressed by means of Dynamic Factor Models (DFMs), which exploit the pervasiveness of business cycle fluctuations to provide a representation of macroeconomic dynamics that is, at the same time, accurate and parsimonious. Thanks to these features, factor models have been, so far, the tool of choice for monitoring macroeconomic conditions in real time.

This paper instead presents three approaches to nowcasting based on Bayesian Vector Autoregressive (BVAR) models, which can also handle a large set of variables sampled at different frequencies. It shows that, in a fully real-time setting, the proposed mixed-frequency BVARs can match the performance of a state-of-the-art DFM, recently developed by the New York Fed, in nowcasting US GDP growth. Moreover, to showcase the use of mixed-frequency BVAR models in particularly challenging circumstances, the paper reports a very real-time nowcast of US GDP for the first quarter of 2020, and also shows how the joint forecast of GDP growth and PCE inflation for 2020 evolves with the (weekly) arrival of new information.

From a methodological point of view, BVAR models have a more general structure than DFMs, do not require the data to be made stationary, as in standard DFM applications, and, crucially, are able to appropriately account for the uncertainty surrounding most of the specification choices such as, for example, the informativeness of prior distributions. At the same time, BVARs are already routinely used in central banks and other policy institutions for forecasting and for building a narrative underlying the economic outlook and its policy implications. The finding that timely information at different frequencies can also be successfully exploited with BVARs thus opens up the possibility to also enrich such analyses.

To showcase the ability of mixed-frequency BVAR models to deliver policy-relevant outputs, the paper presents two additional exercises. First, the generalized impulse response functions to an exogenous shift in GDP are shown to be qualitatively and quantitatively similar to those from traditional quarterly VAR models, highlighting that a mixed-frequency BVAR still provides a reliable account of the transmission mechanism of shocks Second, a counterfactual exercise tracking the 2008Q4 Fed Funds "shadow rate" shows that mixed-frequency BVARs would have provided a more timely warning of the Fed Funds rate reaching the effective zero lower bound than a quarterly model.

1 Introduction

Vector autoregressions (VARs) 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 can be 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, namely, 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 what has more recently become known as Big Data, a phenomenon that has spread in many disciplines over the last two decades. It is no coincidence that the first appearance of the term Big Data in an academic work was by Frank Diebold. His discussion, in 2000, of two papers by Reichlin (2003) and Watson (2003), respectively, presenting a new approach to dynamic factor models – titled "'Big Data' Dynamic Factor Models for Macroeconomic Measurement and Forecasting" – ushered the term Big Data into macroeconometrics and further stoked interest in methods apt to deal with growing amounts of data.[1]

Nowcasting – which can be defined as the prediction of the present, the very near future, and the very recent past (for a survey see Banbura, Giannone, Modugno, and Reichlin, 2013) – is inherently a Big Data problem, as it involves monitoring a multitude of macroeconomic time series with different frequencies, different release dates, and various data irregularities. Indeed, the three Vs of "volume, velocity and variety," the defining properties of Big Data, play an important role in nowcasting. In terms of *volume*, not only macroeconomic datasets grew 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.[2] Moreover, several methods have been proposed for exploiting, efficiently and in real time, the *velocity* of macroeconomic series, i.e. their release at different points in time and often with missing data, ragged edges and various other data irregularities. Finally, the data used for nowcasting are available at different frequencies, and vary substantially in terms of their sources (e.g., hard data versus soft data based on qualitative information such as survey and polls) and precision (e.g., revised versus unrevised data). These features relate to the notion of *variety* in Big Data.

---

[1]Diebold (2012) provides an insightful discussion of the origins of the term Big Data.

[2]Statisticians also refer to this connotation of 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). Note that 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.

Dynamic factor models (DFMs) were among the first and most successful methods for now-casting. The 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 (2012a).[3] DFMs exploit the pervasiveness of business cycle fluctuations to provide a representation of macroeconomic dynamics that is, at the same time, accurate and parsimonious. Thanks to these features, factor models have been, so far, the tool of choice for monitoring macroeconomic conditions in real time, starting from the contribution of Giannone, Reichlin, and Small (2008), and are nowadays used extensively by policy institutions and market participants (for a survey see Stock and Watson, 2017; Bok, Caratelli, Giannone, Sbordone, and Tambalotti, 2018).

VAR models are standard components of the macroeconomist's toolkit since the pioneering work of Sims 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. Bayesian shrinkage, however, is a powerful alternative to factor models for controlling the high estimation uncertainty due to the proliferation of parameters in a high-dimensional setting.

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 real-time policy analysis. Bayesian VARs are regularly used in policy institutions for forecasting and policy analysis,[4] and play an important role as the empirical counterpart to the general equilibrium models used for policy analysis. Since it provides the ability to monitor the economy in real time within the same framework, nowcasting with BVARs opens up the possibility of a richer and timelier narrative through scenario analysis.

BVARs also offer other 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

---

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

[4]Recent work includes Miranda-Agrippino and Rey (2015); Altavilla et al. (2016); Giannone et al. (2019a); Angelini et al. (2019); Domit et al. (2019); Del Negro et al. (2020). For a survey of the literature see Koop and Korobilis (2010); Miranda-Agrippino and Ricco (2018).

[5]D'Agostino, Giannone, Lenza, and Modugno (2016) and Antolin-Diaz, Drechsel, and Petrella (2017) allow for some degree of dynamic heterogeneity, but this is not common practice.

[6]Formally, large BVARs encompass DFMs, in the sense that if the data being analyzed actually have a factor structure, the Bayesian VAR would capture it, as shown in De Mol et al. (2008), and the bias introduced by the imposition of priors would disappear asymptotically as the number of variables increases.

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

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. In this paper, we adopt the approach of Giannone, Lenza, and Primiceri (2015) to produce probabilistic forecasts that reflect all sources of uncertainty, including that coming from the setting of hyperparameters underlying the prior distributions.

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 the higher frequency but were only observed at a lower frequency. We label this method "SS-BVAR," where SS stands for "state-space." The estimates for the latent processes and the uncertainty surrounding them are obtained by means of Kalman filtering techniques. A similar approach has been exploited for handling mixed frequencies, for example, by 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) in a frequentist setting, and by Eraker, Chiu, Foerster, Kim, and Seoane (2014), Schorfheide and Song (2015), Brave, Butters, and Justiniano (2016) and Cimadomo and D'Agostino (2016) using Bayesian methods.

An alternative approach to conduct 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". 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 (e.g., Carriero, Clark, and Marcellino, 2015b; McCracken, Owyang, and Sekhposyan, 2015; Ghysels, 2016). This approach implies, quite conveniently, that the mixed frequency VAR model can be estimated and analyzed as a standard VAR, without resorting to optimal filtering, therefore substantially cutting the computational burden compared to the methodology described above. Moreover, while it is less parsimonious than the SS-BVAR, 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, i.e. those estimated on balanced quarterly datasets. The

method – labelled "cube root", hereafter "CR-BVAR" – was first discussed in Giannone, Monti, and Reichlin (2016) in the context of a DSGE model, and involves estimating the model at a lower frequency (quarterly) and then mapping it into a corresponding model at a higher frequency (monthly). The appeal of this method is that it efficiently exploits the insight from the real-time data flow even in the context of the existing VAR models that most policy institutions have been developing and maintaining over the last decades.

We evaluate the three approaches in the context of a real-time exercise, based on the US data which would have been available to an econometrician in each week from the beginning of 2005 to the end of the first quarter of 2018. The dataset comprises eighteen variables: a first block includes some key macro variables used in most structural macroeconomic models, such as GDP, consumption, investment, labour market variables and factor prices. The dataset also features 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 the US GDP. The results indicate that these tools are valid nowcasting devices. Indeed, all three BVAR approaches produce forecasts that are highly correlated with and as accurate as the New York Fed Staff Nowcasts (see Bok et al., 2018). We also find that the nowcasts of the mixed-frequency approaches are uniformly superior to those of a standard quarterly VAR. This suggests that the mixed-frequency techniques are particularly efficient in seeing through the volatility of high frequency information and extracting the news content of the latter, despite the fact that the mixed-frequency models, with the exception of the CR-BVAR model, are less parsimonious than quarterly VARs.

The real-time nowcasting accuracy of the models should already speak in favour of their usefulness for policy analysis. To gain further insight into the ability of such models to capture the complex dynamic interactions among macroeconomic variables, we propose three additional policy exercises. First, we report a *very* real-time nowcast and a joint forecast of GDP growth and PCE inflation for 2020 – all based on the latest available information for the weeks just before and entering the current Covid-19 crisis period. Second, we show that our mixed-frequency models can be used to trace out the transmission mechanisms of shocks to the US economy, by looking at the generalized impulse response functions in response to an exogenous shift in GDP. Finally, we report a counterfactual exercise aimed at tracking the 2008Q4 Fed Funds "shadow rate".

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

2   Methodology

We present three alternative methods for dealing with mixed-frequency data in a BVAR setting, and provide details about their implementation in the context of our empirical application. The choices regarding the specification and prior distributions are consistent across the models, to the extent possible. In particular, throughout the paper, we maintain the view that the selection of prior distributions is surrounded by uncertainty and, hence, we treat the hyperparameters characterizing the informativeness of the prior distributions as random variables, as in Giannone, Lenza, and Primiceri (2015).

2.1   SS-BVAR: Low-Frequency Variables as Latent Processes

This model assumes that quarterly variables are monthly variables, with missing observations in the first two months of the quarter. The VAR is defined at monthly frequency, and Kalman filtering techniques are employed to estimate the latent monthly processes. More in detail, we assume that the *(log-)levels* of our N variables (collected in the N-dimensional vector X_{t_m}) are described by the following monthly vector autoregressive process (with $p = 17$) lags[8]:

$$X_{t_m} = A_0 + A_1 X_{t_m} + ... + A_p X_{t_m-p} + e_{t_m}, \tag{1}$$

where A_p is the $N \times N$ matrix collecting the coefficients associated with the p-th lag and e_{t_m} is a Normally-distributed multivariate white noise process, with covariance matrix Σ. For what concerns the prior distributions, we choose them in the class of natural conjugate priors (as for the other two VAR approaches) and, more specifically, we use the Normal-Inverse Wishart prior. To parameterise the prior distributions, we proceed as follows.

For Σ, 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 (equal to $\frac{\Psi}{(d-N-1)} = \Psi$). As it is customary in the BVAR literature, the diagonal of Ψ is set equal to the variances of the residuals of estimated autoregressive processes for each variable. For the constant A_0 term, we use a flat prior. For the autoregressive coefficients $(A_1, \ldots, A_p)$, we combine the Minnesota prior, as originally proposed by Litterman (1979), with a 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:

---

[8]17 *monthly* lags ensure full consistency with the information sets of the B-BVAR and CR-BVAR models, which are estimated with 5 lags of *quarterly* data.

$$E(A_1) = I_N, E(A_2) = ... = E(A_p) = 0_{N,N}, \tag{2}$$

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

i.e. it is centered on the random walk model for all the variables. The key hyperparameter is λ, which controls the scale of all the 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 $\lambda \to \infty$, on the other hand, posterior expectations coincide with the Ordinary Least Squares (OLS) estimates. The factor $\frac{1}{s^2}$ is the rate at which the prior variance decreases with increasing lag length[9] and $\frac{\Sigma_{ii}}{\Psi_{jj}}$ accounts for the different scale and variability of the data. 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 and μ, which reflect the informativeness of the prior distribution for the model's coefficients. As in Giannone, Lenza, and Primiceri (2015), we treat the two hyperparameters as random variables and we draw them from their posterior distributions, adapting the estimation algorithm developed in Giannone, Lenza, and Primiceri (2015) to the case in which there are missing data. Schorfheide and Song (2015) and Brave et al. (2016) resort to empirical Bayes methods to optimally select the prior hyperparameters, as in Carriero et al. (2015a). For the hyperparameters, we choose the same rather diffuse priors described in Giannone, Lenza, and Primiceri (2015).

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 prior variance Ψ and 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's and Σ; then, using the posterior sampler of Giannone, Lenza, and Primiceri (2015), we draw the hyperparameters μ and λ conditional on the complete monthly dataset, and finally, we draw the model parameters conditional on the hyperparameters and the complete monthly dataset.

---

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

2.2 B-BVAR: Blocking or Stacking

The idea behind blocking is to align all frequencies to the lowest sampling frequency by defining the higher frequency (monthly) variables as multiple lower frequency (quarterly) variables. We therefore 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 $t_m = 1, 2, 3...$ be a monthly variable. 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 $t_q = 1, 2, 3...$ These three series can now simply be stacked in a VAR(p) with other quarterly variables y_{t_q}. Let $Y_{t_q} = \begin{bmatrix} y_{t_q} & x_{t_q}^q \end{bmatrix}$ and:

$$Y_{t_q} = A_0 + A_1 Y_{t_q-1} + A_2 Y_{t_q-2} + \cdots + A_p Y_{t_q-p} + \varepsilon_{t_q}, \tag{4}$$

where $\varepsilon \ i.i.d \ \mathcal{N}(0, \Sigma)$. Y_{t_q} is a vector of size $N = Q + 3M$, where Q is the number of quarterly variables and M is the number of monthly variables in our system. In the empirical application, we set the number of quarterly lags to $p = 5$.

The system can then be readily estimated with Bayesian methods. We adopt a Normal-Inverse-Wishart prior for the coefficients of the VAR centred around a random walk model. The prior for Σ is an Inverse Wishart with scale Ψ and $d = N + 2$ degrees of freedom. Conditional on Σ, the prior distribution of the autoregressive coefficients is Normal with the same means and variances reported in equations (2) and (3). The use of Bayesian shrinkage allows us to handle large systems. Classical inference, as in Ghysels (2016), is not appropriate in this context due to the high number of free parameters.

Following Giannone, Lenza, and Primiceri (2015), we treat the parameter λ, which controls the tightness of the prior distribution for the model's coefficients, as a random variable and draw it from its posterior distribution, and set the parameter that determines how fast the prior variance decreases with increasing lag length, s, to 2. We also set the diagonal elements of Ψ to the variances of the residuals of estimated autoregressive processes for each variable, in line with the treatment of priors in the previous subsection. As for the SS-BVAR, we also implement the "sum-of-coefficients" prior with dummy observations and treat the hyperparameter μ as a random variable, drawing it from its posterior distribution. Moreover, we also add the "dummy-initial-observation" prior advocated by Sims (1993) to ensure that the prior is consistent with cointegration. The parameter governing the tightness of this prior is simply set to 1.

Given the model parameters, the weekly 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) and based on Durbin and Koopman (2001).[10]

2.3 CR-BVAR: Cube Root

This section shows how to obtain a monthly specification for a $VAR(p)$ defined and estimated at quarterly frequency, reflecting and expanding the results previously derived for DSGE models by Giannone, Monti, and Reichlin (2016).

The idea is that all variables exist at higher frequency, but some are only sampled at quarterly frequency, so we only have observations for these variables in March, June, September or December. Let us define t_m as the time in months and denote by $y_{t_m} = (y_{1,t_m}, ..., y_{n,t_m})'$ the vector of (possibly latent) monthly counterparts to the variables that enter the quarterly model. We transform all variables to correspond to a quarterly quantity when observed at end of the quarter, i.e. when t_m corresponds to March, June, September or December, following Giannone, Reichlin, and Small (2008).

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

Consider a VAR model of order $p > 1$:

$$y_{t_q} = A_1 y_{t_q-1} + A_2 y_{t_q-2} + \cdots + A_p y_{t_q-p} + \varepsilon_{t_q}, \tag{5}$$

where t_q is time in quarters, y is an n-dimensional vector of observable variables and $\varepsilon_{t_q} \sim \mathcal{N}(0, \Sigma_\varepsilon)$. This model can be rewritten in companion form as

$$Y_{t_q} = \Phi Y_{t_q-1} + \nu_{t_q}, \tag{6}$$

with $Y_{t_q} = (y'_{t_q}, ..., y'_{t_q-p+1})'$, $\nu \sim \mathcal{N}(0, \Omega)$ and

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

Model (6) can be rewritten in terms of monthly quantities as

$$Y_{t_m} = \Phi \, Y_{t_m-3} + \nu_{t_m} \tag{7}$$

---

[10]The Kalman filter handles the jagged edges in a reduced-form VAR. A special case is the model by McCracken, Owyang, and Sekhposyan (2015), which required a recursive identification structure with the variables ordered according to the time in which the data are released by the statistical office.

when t_m corresponds to the last month of a quarter.

Assume that the *monthly* counterpart of model (6) can be written in state-space form as[11]

$$Y_{t_m} = \Phi_m Y_{t_m-1} + \nu_{m,t_m}, \tag{8}$$

with $\nu_m \sim \mathcal{N}(0, \Omega_m)$ and

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

Also assume that Φ_m is real and stable.

The first n rows of system (8) correspond to a restricted monthly VAR of the following form:

$$y_{t_m} = \Phi_{m11} y_{t_m-1} + \Phi_{m12} y_{t_m-4} + \cdots + \Phi_{m1p} y_{t_m-3p+2} + \varepsilon_{m,t_m}, \tag{9}$$

The restriction is that current (monthly) values only depend on one month within each lagged quarter. On the other hand, the remaining rows impose restrictions on how the (possibly latent) lagged monthly states are updated each month with the arrival of new information.

$$\begin{bmatrix} y_{t_m-3} \\ \vdots \\ y_{t_m-3p+3} \end{bmatrix} = \begin{bmatrix} \Phi_{m21} y_{t_m-1} + \Phi_{m22} y_{t_m-4} + \cdots + \Phi_{m2p} y_{t_m-3p+2} \\ \vdots \\ \Phi_{mp1} y_{t_m-1} + \Phi_{mp2} y_{t_m-4} + \cdots + \Phi_{mpp} y_{t_m-3p+2} \end{bmatrix} \tag{10}$$

One interesting feature of the relationships in (10) is 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. Indeed, iteration implies that

$$Y_{t_m} = \Phi_m^3 Y_{t_m-3} + \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}, \tag{11}$$

which together with our previous assumptions entails the following relationships between the

---

[11]If the variables considered are stocks, the formulation (8) 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. Doz, Giannone, and Reichlin (2012b) show the effect of such mis-specification to be small.

quarterly model (7) and the monthly model (8):

$$\Phi_m = \Phi^{\frac{1}{3}} \tag{12}$$

$$\nu_{t_m} = \nu_{m,t_m} + \Phi_m \nu_{m,t_m-1} + \Phi_m^2 \nu_{m,t_m-2}. \tag{13}$$

From (12) it is clear that an essential part of finding such mapping is finding the cube root of Φ, which raises the issue of multiple solutions. If the autoregressive matrix of the transition equation is diagonalizable,[12] *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. We can also evaluate the likelihood of solutions using the Kalman filter and pick the one with the highest likelihood, though this is more computationally intensive. In the cases we have tried it, it corresponds to the one with the roots with the smallest argument, as in Giannone, Monti, and Reichlin (2016).[13]

Equation (13) implies that the monthly covariance matrix Σ_{ε_m}, and therefore Ω_m, can be recovered from

$$vec\left(\Sigma_{\varepsilon_m}\right) = \left(I + A \otimes A\right)^{-1} vec\left(\Sigma_\varepsilon\right). \tag{14}$$

with $A = \left[\Phi_{m11}^2 - \Phi_{m11}\left(\sum_{i=2}^p \Phi_{mi1}\right)^{-1}\left(\sum_{i=2}^p \Phi_{mi1}^2\right)\right]$. The solution of (14) 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 as long as A is diagonalizable, so that $A = P\Lambda P^{-1}$, the properties of the Kronecker product imply that the inverse can be computed as

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

which is much more appealing, since $(I + \Lambda \otimes \Lambda)$ is diagonal and thus its inverse is trivial to compute directly.

In summary, to obtain the CR-BVAR the first step is to estimate a quarterly $VAR(p)$ model.[14]

---

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

[13]Anderson, Deistler, Felsenstein, and Koelbl (2016) show g-identifiability when (enough) high frequency data is available.

[14]For our empirical exercise, we assume the same lag length and priors that we have assumed for the B-BVAR.

Table 1: Data and timing of releases

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

[a]Now 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.

Given estimates of the parameters of the quarterly model (6), Φ and Ω, we define a monthly model (8) with parameters Φ_m and Ω_m, which can be recovered from equations (12) and (14). Finally, we compute the distributions of forecasts conditional on the real-time data flow, exploiting the Kalman filtering methods proposed by Bańbura, Giannone, and Lenza (2015) and based on the simulation smoother of Durbin and Koopman (2001).

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 the forecast. We compare the different methods outlined in Section 2 by assessing their performance in a fully real-time nowcasting exercise. In particular, we compare the models' point nowcasts of US real GDP growth with the the New York Fed Staff Nowcasts (see Bok et al., 2018) and with selected quarterly benchmarks, and we also discuss the properties of their nowcast densities.

3.1 Data

The models are estimated on key macro variables (real GDP, real consumption, real investment, the GDP deflator), labour market indicators (a measure of real wage inflation based on compensation per employee, employment, the unemployment rate), financial market variables (the policy 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, house starts, etc...), price data (CPI and PCE price indices), a credit variables (business loans), a measure of uncertainty (Baker, Bloom, and Davis (2016)'s economic policy un-

Figure 1: Nowcasting performance per week of the quarter

| (a) Root mean squared error | (b) Average logarithmic score |
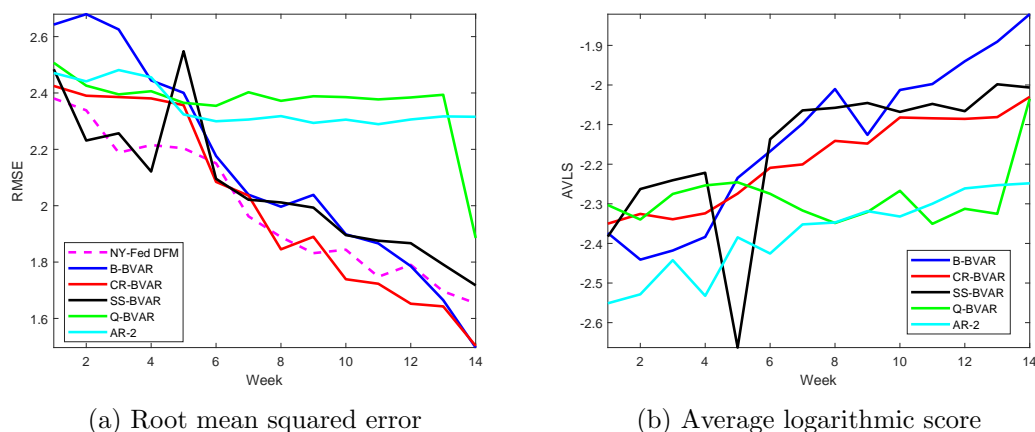
Note: The left panel compares the accuracy of point forecasts across models – measured by their Root Mean Squared Errors – while the right panel compares the accuracy of density forecasts – measured by their Average Logarithmic Scores – as more information becomes available in each week of the quarter. The forecast evaluation sample is 2005-2018Q1.

certainty index) and the manufacturing Purchasing Managers' Index (PMI). GDP, investment, the GDP deflator and compensation per employee are available at a quarterly frequency only, while the other variables are available at monthly frequency, or higher (in which case, we take their monthly averages). Table 1 reports all the variables used, their frequency, their publication lag, whether they enter the model in levels or log-levels, and their FRED id. 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 first quarter of 2018, the same convention used for the weekly updates of the New York Fed Staff Nowcasts. in each vintage, all variables are available since 1986m10.

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.[15] To obtain real quantities, investment and compensation per employee are deflated with the GDP 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.

3.2 Nowcasting Performance

We start by comparing the point nowcasting performance of the BVARs, the New York Fed Staff Nowcasts and two quarterly benchmarks over the sample period which ranges from the beginning

---

[15]As discussed in Section 2.3, for the CR-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.

of 2005 to the first quarter of 2018. Figure 1a reports, for every week in the quarter,[16] 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),[17] 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, CR-BVAR and SS-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.

The mixed-frequency BVARs have RMSEs that are comparable to the DFM and display the usual reduction in RMSE as the quarter progresses and more information becomes available. The SS-BVAR presents a 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 performances of the Q-BVAR and the AR-2 are comparable to those of all other models, by the middle of the second month, the quarterly models are clearly lagging behind, with the Q-BVAR only catching up to a certain extent 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 the GDP nowcasts[18], we compute the average log predictive scores for the nowcasts 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:

$$S_h(\mathcal{M}) = \frac{1}{N_h} \sum_{t=R}^{T-h} \ln p\left(y_{t+h}|y_{1:t}, \mathcal{M}\right), \tag{15}$$

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 $p\left(y_{t+h}|y_{1:t}, \mathcal{M}\right)$ is the predictive density from model $\mathcal{M}$ estimated at time t and evaluated at the actual data outturn.

The improvements in log scores with the arrival of new information throughout the quarter mainly mirror the fall in the RMSEs, although the average variances of the nowcast predictive densities also tend to become somewhat smaller throughout the quarter. As for point forecasts,

---

[16]Week 1 indicates the first week of a quarter, i.e. the one that contains the first Friday of that quarter; week 14 is the 14^{th} 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.

[17]The Q-BVAR corresponds to the first step needed to obtain the CR-BVAR, see Section 2.3.

[18]Historical density nowcasts for the NY Fed's DFM are not publicly available, so the model is omitted from this comparison.

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 the joint forecast of the annual growth rate of real GDP and of annual PCE inflation based on the latest available information, which relates to the current Covid-19 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. Finally, we illustrate a counterfactual exercise aimed at forecasting the 2008Q4 Fed Fund rate in real time, a common practice in central banks 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 U.S. economy. It seems therefore natural to apply the methods discussed in this paper to the current conjuncture, using the latest available (monthly) information.

The top panel of Figure 2 reports the distributions of the nowcasts for 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 nowcast to various categories of data releases. As can be seen, 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 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 for PCE inflation is in negative territory. The uncertainty surrounding the 22-May 2020 forecast is also much larger compared to the two earlier forecasts.

Figure 2: Nowcast for real GDP growth in Q1 2020
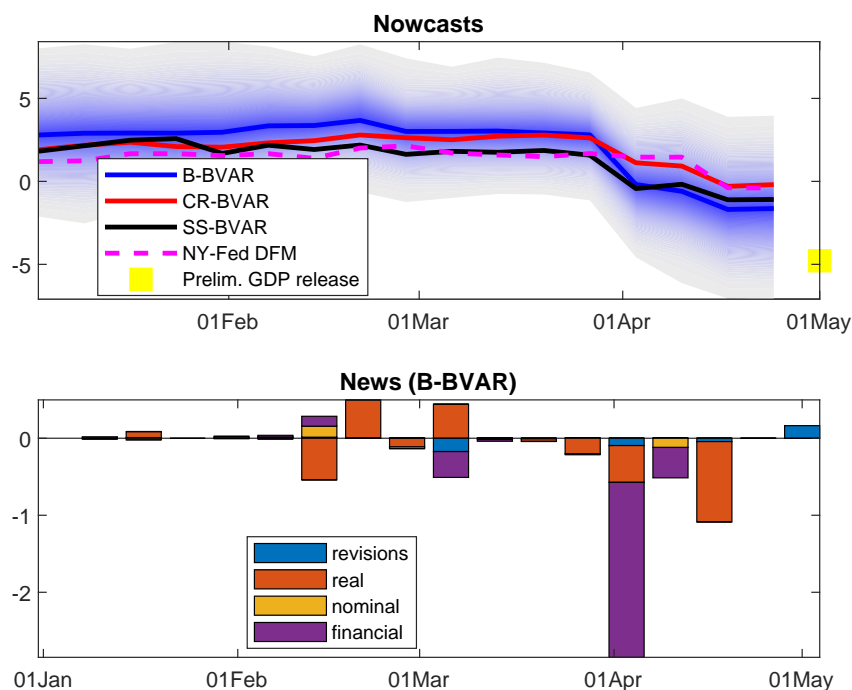
Note: The top panel shows the probability distribution of B-BVAR nowcasts in each week from the beginning of 2020 until the preliminary relase of Q1 GDP on 29 April, the median nowcasts from the CR-BVAR and SS-BVAR, as well as the NY Fed's DFM nowcast. The fan chart bands cover 99% of the support around 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.

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 a monthly frequency.

For illustrative purposes, we present a generalized impulse response function to a one standard deviation GDP shock (Figure 4), but other setups in terms of more elaborate identification schemes or shocks occurring in different months of the quarter, can be easily accommodated within all three models. Generalized impulse response functions to GDP are equivalent to a perturbation of the forecast error of GDP in a recursively identified VAR with GDP ordered first. These are the responses to the linear combination of structural shocks that have have been the main historical drivers of innovations in GDP fluctuations (see, e.g., Pesaran and Shin, 1998; Bańbura, Giannone, and Lenza, 2015).[19]

---

[19]The generalized impulse response functions 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 have been the main historical drivers of innovations of GDP variation at business cycle frequencies (see Giannone, Lenza, and Reichlin, 2019a;

Figure 3: Evolution of the joint distribution of the forecasts for GDP growth and PCE inflation
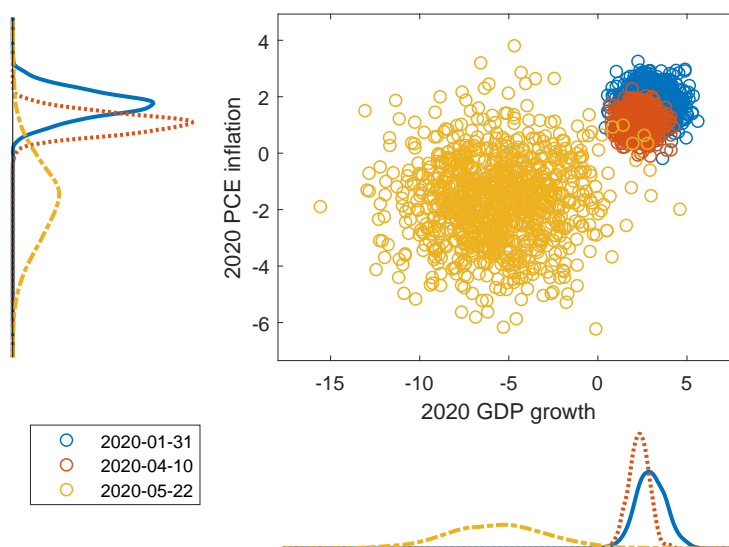
Note: The scatter plot shows draws from the B-BVAR's joint predictive densities for annual GDP growth and PCE inflation in 2020 in three different weekly vintages. The two plots along the axes show kernel-smoothed estimates of the marginal predictive densities for the two variables in the same three vintages. Annual growth rates are computed from the underlying projections in log levels for the corresponding variables.

Figure 4 reports the 68% credible intervals for the stacked B-BVAR model and shows the median responses for the other two mixed-frequency approaches.[20] All models produce broadly similar IRFs, though the ones from the SS-BVAR are at times out of the B-BVAR's credible intervals. Consistently with the related 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 reflected by the positive reactions of the GDP deflator, CPI and PCE price indices, suggesting that demand shocks are important driver of GDP in the US, and this is accompanied by a tightening of the Federal Funds rate for about three years, which reflects the systematic component of US monetary policy. As concerns 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 drops, but only for a few quarters.

4.3 The Real-Time Evolution of the 2008Q4 Fed Funds Rate

In this subsection, 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 US monetary policy stance, in the fourth quarter of 2008 because

Angeletos, Collard, and Dellas, 2020). The similarity between the two approaches was recently documented also by Del Negro, Lenza, Primiceri, and Tambalotti (2020).
    [20]The CR-BVAR and SS-BVAR responses have been scaled to match the B-BVAR's impact real GDP response.

Figure 4: Generalised impulse response function to a GDP impulse
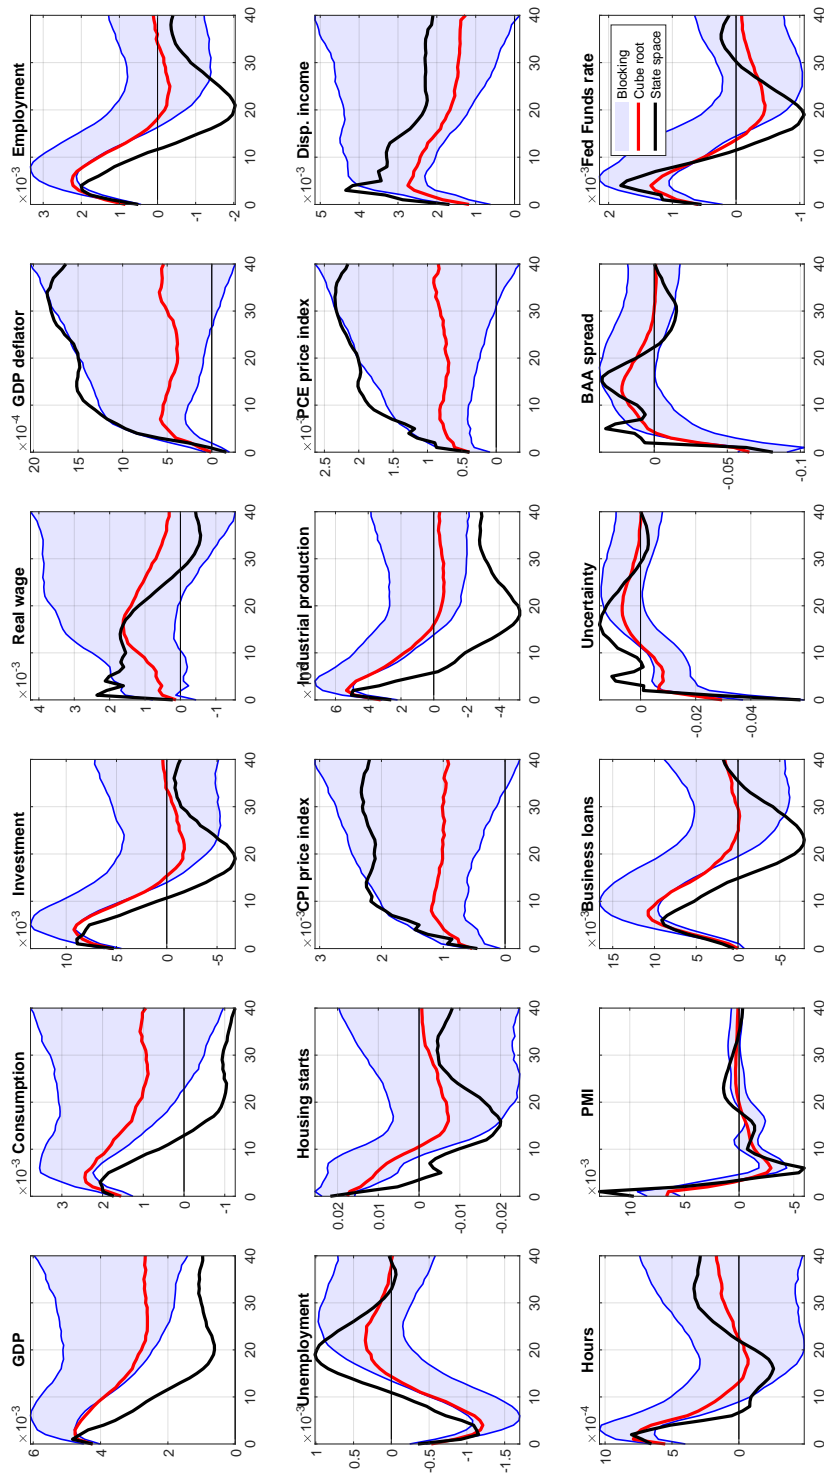
Note: Generalised impulse response function (GIRF) to a shock to GDP, shown at quarterly frequency. The blue area represents the 68% credible interval for a one standard deviation impulse to the B-BVAR model; the red and black lines show the median GIRFs for the CR-BVAR and SS-BVAR models, respectively, scaled to deliver the same impact GDP response as the B-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).

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 seek to answer 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 estimation is carried out for each weekly data vintage of 2008 available in our real-time database. Central banks routinely look at counterfactual interest rate paths, such as the one we derive in this exercise, as a benchmark to gauge whether their policy rates and the closely related short-term money market rates are at reasonable levels, given the prevailing and expected economic conditions.[21]

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 banks rely on many different sources of information, so that the 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 are missing also 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, dates reported on the horizontal axis) and the historical regularities captured by our VAR. The results are cast in terms of a heatmap, whereby the areas with higher probability density are characterized by a lighter colour. For the sake of brevity, the heatmap results refer to the B-BVAR, but the SS-BVAR and CR-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 described in the section on nowcasting results.

The mixed frequency VAR leads to a more timely assessment of the deterioration of economic conditions in the US economy and, consequently, suggests that the level of the Fed Funds rate would head toward very low levels much earlier than the quarterly VAR. Eventually, both models

---

[21]Since 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., 2019b).

Figure 5: Counterfactual Fed Funds Rate
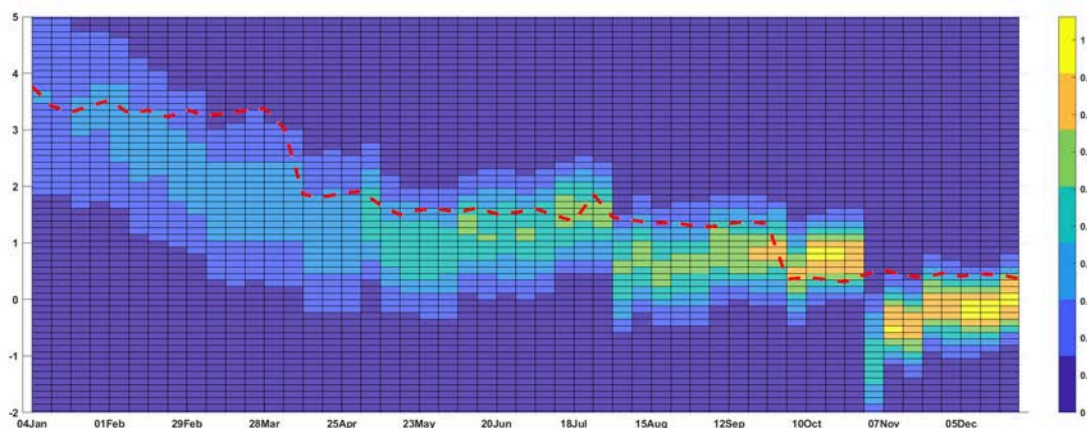
Note: Horizontal axis: data vintages used to compute the forecast of the 2008Q4 Fed Funds rate (52 weeks). Vertical axis: value of the (counterfactual) Fed Funds rate in percentage points. Areas with high probability density in the heatmap are characterized by lighter colours, according to the convention defined in the colorbar next to the plot. The red dashed line is the median of the Q-BVAR counterfactual Fed Funds rate estimated with the same vintages.

suggest that the Fed Funds rate would be constrained by the zero lower bound. The timeliness of the mixed frequency VAR is also reflected in a higher volatility of the assessment of the counterfactual policy rate. But since the analysis of its nowcasting performance has shown that the model does a good job at distilling news from noise as more and more high-frequency information is incorporated in the assessment, this volatility should not simply be viewed as noise, but rather as an inherent feature of the real-time information flow.

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 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, they require less modelling choices (e.g., related to the number of lags, the block-structure, etc.), and they do not require the data to be made stationary.

We present three strategies for dealing with mixed-frequency in the context of VARs: first, a model – labelled "state-space BVAR" – which assumes that all variables are high-frequency time series, but that some of them are observed only every quarter. Second, we adopt the methodology known as "blocking", which allows to treat 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. We 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 the first quarter of 2018, we show how these models would have nowcasted U.S. GDP growth. Moreover, we discuss the advantages and drawbacks of each of these approaches.

Our results suggest that these models have a nowcasting performance similar to the New York Fed's Dynamic Factor Model, and display a clear improvement in forecast accuracy as the quarter progresses and more information becomes available. 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 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, AND L. KOELBL (2016): "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.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring economic policy uncertainty," *The quarterly journal of economics*, 131, 1593–1636.

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

BANBURA, M., D. GIANNONE, M. MODUGNO, AND L. REICHLIN (2013): *Now-Casting and the Real-Time Data Flow*, Elsevier, vol. 2 of *Handbook of Economic Forecasting*, chap. 0, 195–237.

BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25, 71–92.

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.

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., R. A. BUTTERS, AND A. JUSTINIANO (2016): "Forecasting Economic Activity with Mixed Frequency Bayesian VARs," Working Paper Series WP-2016-5, Federal Reserve Bank of Chicago.

BURNS, A. F. AND W. C. MITCHELL (1946): *Measuring Business Cycles*.

CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2015a): "Bayesian VARs: Specification Choices and Forecast Accuracy," *Journal of Applied Econometrics*, 30, 46–73.

——— (2015b): "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.

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.

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?" NBER Working Papers 27003, National Bureau of Economic Research, Inc.

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 (2012a): "A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models," *The Review of Economics and Statistics*, 94, 1014–1024.

——— (2012b): "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.

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.

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.

Giannone, D., M. Lenza, and L. Reichlin (2019a): "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.

——— (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.

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.

HARTMANN, P. AND F. SMETS (2018): "The first twenty years of the European Central Bank: monetary policy," CEPR Discussion Papers 13411, C.E.P.R. Discussion Papers.

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. AND D. KOROBILIS (2010): "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends(R) in Econometrics*, 3, 267–358.

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.

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. SEKHPOSYAN (2015): "Real-Time Forecasting with a Large, Mixed Frequency, Bayesian VAR," Working Papers 2015-30, Federal Reserve Bank of St. Louis.

MIRANDA-AGRIPPINO, S. AND H. REY (2015): "World Asset Markets and the Global Financial Cycle," CEPR Discussion Papers 10936, C.E.P.R. Discussion Papers.

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.

REICHLIN, L. (2003): "Factor Models in Large Cross Sections of Time Series," in *Advances in Economics and Econometrics Theory and Applications, Eighth World Congress*, Cambridge University Press, 87–115.

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. (1980): "Macroeconomics and reality," *Econometrica: journal of the Econometric Society*, 1–48.

——— (1993): "A nine-variable probabilistic macroeconomic forecasting model," in *Business cycles, indicators, and forecasting*, University of Chicago press, 179–212.

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.

WATSON, M. (2003): "Macroeconomic forecasting using many predictors," in *Advances in Economics and Econometrics Theory and Applications, Eighth World Congress*, Cambridge University Press, 87–115.

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 Forecating 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.

Jacopo Cimadomo
European Central Bank, Frankfurt am Main, Germany; email: jacopo.cimadomo@ecb.europa.eu

Domenico Giannone
Amazon, Seattle, United States; email: dgiannon2@gmail.com

Michele Lenza (corresponding author)
European Central Bank, Frankfurt am Main, Germany; ECARES-ULB; email: michele.lenza@ecb.europa.eu

Francesca Monti
King's Business School, London, United Kingdom; Centre for Macroeconomics; email: francesca.monti@kcl.ac.uk

Andrej Sokol
European Central Bank, Frankfurt am Main, Germany; Bank of England; Centre for Macroeconomics;
email: andrej.sokol@ecb.europa.eu