Journal of Econometrics 232 (2023) 18–34

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Nowcasting the output gap

Tino Berger [a], James Morley [b,*], Benjamin Wong [c]

[a] University of Göttingen, Germany
[b] University of Sydney, Australia
[c] Monash University, Australia

ARTICLE INFO

Article history:
Received 29 May 2020
Received in revised form 20 August 2020
Accepted 24 August 2020
Available online 5 November 2020

JEL classification:
C32
C55
E32

Keywords:
Nowcasting
Output gap
COVID-19

ABSTRACT

We propose a way to directly nowcast the output gap using the Beveridge–Nelson decomposition based on a mixed-frequency Bayesian VAR. The mixed-frequency approach produces similar but more timely estimates of the U.S. output gap compared to those based on a quarterly model, the CBO measure of potential, or the HP filter. We find that within-quarter nowcasts for the output gap are more reliable than for output growth, with monthly indicators for a credit risk spread, consumer sentiment, and the unemployment rate providing particularly useful new information about the final estimate of the output gap. An out-of-sample analysis of the COVID-19 crisis anticipates the exceptionally large negative output gap of −8.3% in 2020Q2 before the release of real GDP data for the quarter, with both conditional and scenario nowcasts tracking a dramatic decline in the output gap given the April data.

1. Introduction

Central banks often rely on the output gap as a measure of the overall degree of slack in the economy. However, the output gap is unobserved and must be estimated. Recent research suggests that multivariate models designed to span the relevant information for aggregate shocks and capture reduced-form dynamics in the macroeconomic variables can be used to produce reasonable estimates of the output gap (Luciani and Barigozzi, 2019; Morley and Wong, 2020). Because these models are based on quarterly data that only become available with some delay, typically at least a month after the end of a quarter, there is concomitant delay in when the output gap for a given quarter can be estimated. This delay is problematic for central banks and other policy organizations wanting to make policy decisions informed by a quantitative measure of current overall slack in the economy, especially in the face of a large and sudden change in economic conditions such as has occurred with the COVID-19 pandemic.

We propose directly nowcasting the output gap using a multivariate Beveridge–Nelson (BN) decomposition based on a mixed-frequency Bayesian VAR. Our approach applies the Bayesian shrinkage suggested in Morley and Wong (2020) for estimating the output gap to a mixed-frequency VAR along the lines of Ghysels (2016) and McCracken et al. (2020). The model incorporates monthly indicators for interest rates, stock returns, consumer sentiment, the unemployment rate, inflation, industrial production, and housing starts to help predict quarterly real GDP. Following McCracken et al. (2020), we are able to update conditional expectations for the model as each monthly indicator is released and conduct scenario analysis using an approach based on Waggoner and Zha (1999).

* Corresponding author.
E-mail addresses: tino.berger@wiwi.uni-goettingen.de (T. Berger), james.morley@sydney.edu.au (J. Morley), benjamin.wong@monash.edu (B. Wong).

https://doi.org/10.1016/j.jeconom.2020.08.011

For different possible assumptions about what constitutes the "best" measure of the U.S. output gap, including the output gap based on Bayesian VAR with quarterly data following Morley and Wong (2020), the Congressional Budget Office's (CBO) measure of potential, and the Hodrick–Prescott (HP) filter, a key finding is that the estimated output gap based on the mixed-frequency Bayesian VAR is highly correlated with the other measures, with real-time estimates from our approach predicting revised values of the implied CBO and HP filter output gaps more accurately than their respective real-time counterparts. We also find that the mixed-frequency Bayesian VAR is able to track the output gap closely within the quarter and well before quarterly real GDP is released. In particular, we find that monthly indicators for a credit risk spread, consumer sentiment, and the unemployment rate provide particularly useful new information about the final estimate of the output gap, with the mean absolute error of the output gap nowcasts falling to below 0.1 percentage points following the release of data for the unemployment rate in the second month of a quarter.

Our estimation is for data from 1967–2019. However, to examine the economic crisis caused by the COVID-19 pandemic, we conduct an out-of-sample analysis that anticipates the exceptionally large negative output gap of −8.3% for 2020Q2 before the release of real GDP data for the quarter, with both conditional and scenario nowcasts tracking a dramatic decline in the output gap given the April data.

The rest of this paper is organized as follows: Section 2 presents a brief overview of the literature on mixed-frequency analysis to help motivate the approach taken in this paper. Section 3 presents our proposed approach for directly estimating the output gap using a multivariate BN decomposition based on a mixed-frequency Bayesian VAR, including constructing nowcasts given partial information from within a quarter and scenarios for the values of certain monthly indicators. Section 4 reports empirical results for our proposed approach when applied to U.S. data from 1967 to 2019, with a variety of robustness checks. Section 5 considers implications of the COVID-19 crisis for the U.S. output gap in the first half of 2020. Section 6 concludes the paper.

2. Mixed-frequency analysis

Various mixed-frequency approaches have been proposed to nowcast real GDP growth, but none to our knowledge to directly nowcast the output gap.[1] In terms of existing approaches, a useful classification by Bańbura et al. (2013) distinguishes between *partial* models, i.e., single-equation models that do not specify a joint model for the target variable and all predictors, and *full* models, i.e., multi-equation models that explicitly model the joint dynamics of all variables.

Even though partial models suffer from some drawbacks, they are popular tools at central banks to obtain early estimates of real GDP growth (Bańbura et al., 2013). Partial models include bridge equations and mixed-data sampling (MIDAS), both extensions of distributed lag models to mixed-frequency data.

A bridge equation with a single indicator can be defined as

$$\Delta y_t = c + \phi(L)\Delta y_{t-1} + \beta(L)q_t(m_t, m_{t-1/3}, m_{t-2/3}) + \varepsilon_t, \tag{1}$$

where Δy_t is real GDP growth in quarter t and $q_t(\cdot)$ is a quarterly indicator aggregated from higher-frequency monthly indicators, $m_{t-2/3}$, $m_{t-1/3}$, and m_t. The specific aggregation rule depends on the nature of the indicator variable.[2] The nowcast $\Delta y_{T+1|T+v}$ is the expectation conditional on information in period $T+v$ (the final observation of the higher-frequency variable), including real GDP growth in T (the final quarter for which GDP is available), i.e.,

$$\Delta y_{T+1|T+v} = \hat{c} + \hat{\phi}(L)\Delta y_T + \hat{\beta}(L)q_{T+1|T+v}, \tag{2}$$

where the hats denote OLS estimates of the parameters. Note that the nowcast requires higher-frequency observations or forecasts of the predictor and a time aggregation step. For forecasting the monthly predictor, simple autoregressive models are often used. Lags of the quarterly indicator are observed at time T (i.e., $q_{T-j|T+v-j} = q_{T+1-j}(\cdot)$ for $j > 0$). Consider, as an example, the nowcast for real GDP growth for Q2. Suppose industrial production (IP) is available for April and May (i.e., $v = 2/3$) and real GDP for Q1. A forecast of IP growth for June, together with the observed values for April and May, allows aggregation to a forecast of IP growth at a quarterly frequency. This projected value of the quarterly indicator can be used in Eq. (2) to nowcast real GDP growth in Q2. We note here that our scenario nowcast presented in Section 5 is much like the bridge equation approach in the sense that forecasts of monthly indicators are taken from external considerations and used to calculate the conditional expectation of real GDP given a dynamic model.

For the MIDAS approach, the predictors are directly included at their original frequency:

$$\Delta y_t = c + \phi(L)\Delta y_{t-1} + \beta B(L^{1/3};\theta)m_{t-1+v} + \varepsilon_t, \tag{3}$$

where $B(L^{1/3};\theta)$ is a lag polynomial, typically specified as an exponential Almon polynomial. The conditional expectation of Eq. (3) is given by

$$\Delta y_{T+1|T+v} = \hat{c} + \hat{\phi}(L)\Delta y_T + \hat{\beta}B(L^{1/3};\hat{\theta})m_{T+v}. \tag{4}$$

[1] Garratt et al. (2014) consider "nowcasting" the output gap using an ensemble of (one-step-ahead) forecast densities from quarterly VARs that include different univariate output gap measures and inflation. Their main finding is that there is a high degree of uncertainty about the output gap across different detrending methods. However, they do not consider higher-frequency data or within-quarter nowcasts allowed for by a mixed-frequency approach.

[2] See Stock and Watson (2002) for a discussion of different aggregation rules.

Using the example of nowcasting real GDP growth in Q2 given its value in Q1 and observing a monthly indicator in April and May, then, as in the previous example, $v = 2/3$. Foroni et al. (2015) propose a variant of MIDAS with unrestricted lag polynomials, which they label "U-MIDAS". In the quarterly/monthly case, a U-MIDAS model is given by

$$\Delta y_t = c + \phi(L)\Delta y_{t-1} + \beta(L^{1/3})m_{t-1+v} + \varepsilon_t. \tag{5}$$

That is, the individual lags are estimated separately. In contrast to standard MIDAS, the U-MIDAS model can be estimated using OLS. The U-MIDAS approach is particularly suitable when the difference in sampling frequencies is small, such as quarterly/monthly, as considered in our analysis. Thus, as discussed below, we take an approach closely related to U-MIDAS when estimating our multi-equation mixed-frequency model.

Full multi-equation mixed-frequency models in the literature include dynamic factor models and VARs. The key difference among the existing approaches is the baseline frequency, i.e., the frequency of observations at which the economy is assumed to evolve. Both approaches have been considered using Bayesian and frequentist estimation.

Models taking the high-frequency as the baseline frequency, referred to as *parameter-driven* models, treat the high-frequency observations of the low-frequency variables as missing observations. They rely on a state-space representation of the model and obtain the missing observations using the Kalman filter and smoother. A notable example is Schorfheide and Song (2015), who estimate a Bayesian mixed-frequency VAR with missing observations to nowcast real GDP growth in real time. Despite some attractive features, these models are computationally costly to work with and we do not take this approach in our analysis.

Instead, we take the approach of a model specified at the *lowest* observed frequency, as in Ghysels (2016), Brave et al. (2019), and McCracken et al. (2020). As discussed in the next section, we are particularly guided by McCracken et al. (2020), who consider a mixed-frequency VAR model motivated by Ghysels (2016) with one quarterly variable (real GDP growth) and 12 monthly indicators (suitably transformed to induce stationarity). They consider Bayesian estimation and real-time data from the ALFRED database, which we also consider in our robustness analysis. In their approach, monthly data releases are treated as separate observations within a quarter. That is, the approach is much like a U-MIDAS model, as noted in Ghysels (2016), but is multi-equation and estimated using Bayesian methods to deal with parameter proliferation.

3. Our proposed approach

Beveridge and Nelson (1981) define the trend of a time series as its long-horizon conditional expectation minus any future deterministic drift. For a time series process $\{y_t\}$ with constant drift μ, the BN trend at time t, τ_t, is

$$\tau_t = \lim_{h\to\infty} \mathbb{E}_t\left[y_{t+h} - h\cdot\mu\right]. \tag{6}$$

The corresponding BN cycle of the time series at time t, c_t, is then given by

$$c_t = y_t - \tau_t. \tag{7}$$

Following Morley and Wong (2020), we specify a VAR model to evaluate the conditional expectation in Eq. (6) in order to calculate (7). The VAR in companion form is

$$\boldsymbol{X}_t = \boldsymbol{F}\boldsymbol{X}_{t-1} + \boldsymbol{H}\boldsymbol{\varepsilon}_t, \tag{8}$$

where $\boldsymbol{X}_t$ is a vector of demeaned variables, $\boldsymbol{F}$ is the companion matrix, $\boldsymbol{H}$ is a matrix that maps the forecast errors to the companion form, and $\boldsymbol{\varepsilon}_t$ is a vector of forecast errors. Let $\mathbf{s}_k$ be a selector row vector which consists of 1 as its kth element and zeros otherwise. Suppose also that real GDP growth Δy_t is included as the kth element of the vector $\boldsymbol{X}_t$ in Eq. (8). Following Morley (2002), the BN cycle of y_t can then be calculated as

$$c_t = -\mathbf{s}_k\boldsymbol{F}(\boldsymbol{I} - \boldsymbol{F})^{-1}\boldsymbol{X}_t. \tag{9}$$

Our objective, therefore, is to cast a mixed-frequency system into the form of Eq. (8) in order to obtain an estimate of the output gap based on the BN cycle of log real GDP from Eq. (9).

3.1. A mixed-frequency VAR system

We specify a mixed-frequency VAR (MF-VAR) along the lines of Ghysels (2016) and McCracken et al. (2020). The MF-VAR includes variables observed at both monthly and quarterly frequencies. Let $m_{j,t-1+v}$ be the jth variable observed at monthly frequency in quarter t, where $v \in \{1/3, 2/3, 1\}$ corresponds to the month within the quarter. That is, $m_{j,t-2/3}$ is observed first, then $m_{j,t-1/3}$, and finally $m_{j,t}$.

First, stack the k monthly variables as

$$\boldsymbol{m}_{t-v} = \begin{bmatrix} \tilde{m}_{1,t-1+v} \\ \tilde{m}_{2,t-1+v} \\ \vdots \\ \tilde{m}_{k,t-1+v} \end{bmatrix},$$

where $\tilde{m}_{j,t-1+v} \equiv m_{j,t-1+v} - \mu_j$, and μ_j is the mean of the jth variable observed at monthly frequency. Then, denoting $\Delta\tilde{y}_t \equiv \Delta y_t - \mu_{\Delta y}$, where $\mu_{\Delta y}$ is the mean of real GDP growth, stack all of the demeaned variables observed at monthly frequency within the quarter, along with demeaned real GDP growth, which is observed at quarterly frequency, as follows:

$$\boldsymbol{Y}_t = \begin{bmatrix} \boldsymbol{m}_{t-1+1/3} \\ \boldsymbol{m}_{t-1+2/3} \\ \boldsymbol{m}_t \\ \Delta\tilde{y}_t \end{bmatrix}.$$

For example, for 1995Q1 and previous quarters,

$$\boldsymbol{Y}_{1995Q1} = \begin{bmatrix} \boldsymbol{m}_{\text{Jan1995}} \\ \boldsymbol{m}_{\text{Feb1995}} \\ \boldsymbol{m}_{\text{Mar1995}} \\ \Delta\tilde{y}_{1995Q1} \end{bmatrix}, \boldsymbol{Y}_{1994Q4} = \begin{bmatrix} \boldsymbol{m}_{\text{Oct1994}} \\ \boldsymbol{m}_{\text{Nov1994}} \\ \boldsymbol{m}_{\text{Dec1994}} \\ \Delta\tilde{y}_{1994Q4} \end{bmatrix}, \boldsymbol{Y}_{1994Q3} = \begin{bmatrix} \boldsymbol{m}_{\text{Jul1994}} \\ \boldsymbol{m}_{\text{Aug1994}} \\ \boldsymbol{m}_{\text{Sept1994}} \\ \Delta\tilde{y}_{1994Q3} \end{bmatrix}, \quad \text{etc...}$$

The vector process $\boldsymbol{Y}_t$ is assumed to have a VAR(p) structure at the quarterly frequency:

$$\boldsymbol{Y}_t = \boldsymbol{\Phi}_1\boldsymbol{Y}_{t-1} + \boldsymbol{\Phi}_2\boldsymbol{Y}_{t-2} + \ldots + \boldsymbol{\Phi}_p\boldsymbol{Y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma}). \tag{10}$$

The stacking of the monthly and quarterly variables implies that fitting a VAR structure on $\boldsymbol{Y}_t$ suffices in specifying the companion-form system in Eq. (8), where $\boldsymbol{X}_t = \left[\boldsymbol{Y}_t' \;\; \boldsymbol{Y}_{t-1}' \ldots \boldsymbol{Y}_{t-p}'\right]'$. By not tying the parameters in Eq. (10) to a temporal aggregation rule or an exponential Almon polynomial, the mixed-frequency VAR is much like a U-MIDAS model, as noted in Ghysels (2016).

3.2. Estimation

We use Bayesian methods to implement shrinkage when estimating equation (10) due to the parameter proliferation in the MF-VAR implying a non-trivial risk of overfitting the data in sample. Our estimation procedure closely mimics Morley and Wong (2020) by using a natural conjugate prior with a standard Minnesota structure to implement shrinkage in a Bayesian VAR (BVAR). The natural conjugate prior implies an analytical solution to the posterior distribution. To motivate the prior, expand the VAR in Eq. (10):

$$\boldsymbol{Y}_t = \begin{bmatrix} \phi_1^{1,1} & \ldots & \phi_1^{1,3k+1} & \phi_2^{1,1} & \ldots & \phi_2^{1,3k+1} & \ldots & \ldots & \phi_p^{1,3k+1} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \phi_1^{3k+1,1} & \ldots & \phi_1^{3k+1,3k+1} & \phi_2^{3k+1,1} & \ldots & \phi_2^{3k+1,3k+1} & \ldots & \ldots & \phi_p^{3k+1,3k+1} \end{bmatrix} \begin{bmatrix} \boldsymbol{Y}_{t-1} \\ \boldsymbol{Y}_{t-2} \\ \vdots \\ \boldsymbol{Y}_{t-p} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \vdots \\ \varepsilon_{3k+1,t} \end{bmatrix}.$$

Letting $\phi_l^{j,k}$ denote the slope coefficient of the lth lag of variable k in the jth equation of the VAR in Eq. (10), we set the prior means and variances of the slope coefficients as follows:

$$\mathbb{E}[\phi_l^{j,k}] = 0,$$
$$Var[\phi_l^{j,k}] = \begin{cases} \frac{\lambda^2}{l^2}, & j = k \\ \frac{\lambda^2}{l^2}\frac{\sigma_j^2}{\sigma_k^2}, & \text{otherwise.} \end{cases}$$

The degree of shrinkage in the BVAR is governed by the hyperparameter λ, with $\lambda \to 0$ shrinking towards the assumption that the variables in the model are independent white noise processes or, equivalently, all of the first-differenced variables in the model follow random walk processes in levels. As long as $\lambda > 0$, the posterior will converge asymptotically to the population parameters. Following a standard Minnesota prior structure, the factor $1/l^2$ shrinks coefficients at longer lags closer to zero. The hyperparameter variances σ_j^2 and σ_k^2 are set to the variances of residuals from autoregressive models estimated using least squares for the corresponding variables, as is the usual practice (e.g., Bańbura et al., 2010; Koop, 2013).[3]

Following Morley and Wong (2020), we set λ by optimizing the one-step-ahead out-of-sample forecast of output growth. The focus on an out-of-sample forecast is aimed at not overfitting output growth with our model. Specifically, we conduct numerical optimization to find the λ that minimizes the one-step-ahead root mean squared forecast error (RMSFE) for output growth over an evaluation sample using pseudo-real-time estimation based on an expanding window starting with a particular initial fraction of the full sample. For our application, we start our recursive estimation with the initial 20 years of data and use the remaining 30+ years of data for evaluation of the RMSFE. Conditional on λ

[3] The hyperparameter variance for output growth is from an AR(4) regression. The mixed-frequency setup leads to a slightly different approach for the monthly variables. We use AR(12) regressions for the monthly variables, but dropping the first lag for variables corresponding to the second month of a quarter and the first two lags for variables corresponding to the third month of a quarter. Also, we set the ratio of hyperparameter variances to unity for the same jth variable across different months. For a shrinkage hyperparameter of $\lambda = 0.05$, this approach leads to very similar one-step-ahead forecasts for real GDP compared to quarterly BVARs with Minnesota priors estimated either in levels or differences.

and the assumption of normality for the variables, the calculation of posterior moments for the slope coefficients is straightforward, as the natural conjugacy of the prior implies that we can implement estimation using least squares with dummies observations (see, for example, Del Negro and Schorfheide, 2011; Woźniak, 2016).

Finally, we note that our use of demeaned variables in $\boldsymbol{Y}_t$ is equivalent to setting a flat prior on the unconditional means of the variables in our model. In the case of the same jth variable across different months, we use the common sample mean $\hat{\mu}_j$ for all months.[4]

3.3. Nowcasting structure

For U.S. macroeconomic variables, data are often released with a one month lag. For the monthly indicators, this means we get observations for December in January, observations for November in December, and so on. We also get the advance release of the fourth quarter real GDP at the end of January, the advance release of the third quarter real GDP at the end of October, and so forth. Going back to Eq. (9), given the entire vector of observations in $\boldsymbol{X}_t$, we can obtain c_t as our final estimate of the output gap. The nowcasting problem that we aim to address is how to also obtain an estimate of c_t within a quarter when we only partially observe the vector, $\boldsymbol{X}_t$. In particular, as monthly data are released within the quarter, how should the estimate of the output gap be updated?

To begin, using Eq. (9) to define the BN cycle at $T+1$ and then substituting in Eq. (8) shifted to $T+1$, we obtain the following expression:

$$\begin{aligned} c_{T+1} &= -\boldsymbol{s}_k \boldsymbol{F}(\boldsymbol{I}-\boldsymbol{F})^{-1}\boldsymbol{X}_{T+1} \\ &= -\boldsymbol{s}_k \boldsymbol{F}(\boldsymbol{I}-\boldsymbol{F})^{-1}\left[\boldsymbol{F}\boldsymbol{X}_T + \boldsymbol{H}\boldsymbol{\varepsilon}_{T+1}\right]. \end{aligned} \tag{11}$$

With this expression, first consider the case of obtaining a nowcast for the output gap, $c_{T+1|T}$, near the beginning of quarter $T+1$ conditional on observing all of the variables from the previous quarter T but no information yet about the current quarter.[5] Conditional on observing the vector $\boldsymbol{X}_T$, the nowcast for the output gap is simply the first part of the expression, i.e., $-\boldsymbol{s}_k \boldsymbol{F}(\boldsymbol{I}-\boldsymbol{F})^{-1}\boldsymbol{F}\boldsymbol{X}_T$, given $\boldsymbol{\varepsilon}_{T+1|T} = \boldsymbol{0}$.

As monthly indicators become available throughout the quarter, the nowcast can be updated because the information set has changed. Thus, we might want to calculate $c_{T+1|T+1/3}$. However, when $\boldsymbol{X}_{T+1/3}$ is observed, $\boldsymbol{\varepsilon}_{T+1}$ is only partially observed and $\boldsymbol{\varepsilon}_{T+1|T+1/3} \neq \boldsymbol{0}$. Our objective with nowcasting is to update the conditional expectation of $\boldsymbol{\varepsilon}_{T+1}$ in order to estimate the output gap using Eq. (11). Our approach closely follows McCracken et al. (2020) in calculating condition expectations, where the already-released intra-quarter data are accounted for using the Waggoner and Zha (1999) approach of conditional forecasting.[6]

To understand how conditional forecasting works, recall that the covariance matrix for the VAR forecast errors is $\boldsymbol{\Sigma}$. Intuitively, what we are doing is that, conditional on partially observing $\boldsymbol{\varepsilon}_{T+1}$, we update the rest of the $\boldsymbol{\varepsilon}_{T+1}$ vector to be consistent with the structure implied by $\boldsymbol{\Sigma}$. More specifically, our approach takes a lower-triangular Cholesky factor of $\boldsymbol{\Sigma}$, $\boldsymbol{B}$, where $\boldsymbol{\Sigma} = \boldsymbol{B}\boldsymbol{B}'$. Then, from the covariance structure, we define the following:

$$\boldsymbol{\varepsilon}_{T+1} = \boldsymbol{B}\boldsymbol{z}_{T+1}, \quad \text{where} \quad \boldsymbol{z}_{T+1} \sim (\boldsymbol{0}, \boldsymbol{I}), \tag{12}$$

---

4 Following Morley and Wong (2020), we can address structural change in the unconditional means of the variables in the BVAR, including the trend growth rate of real GDP, $\mu_{\Delta y}$, by using time-varying estimates of unconditional means when demeaning variables. In particular, Morley and Wong (2020) allow for a break in trend growth and other variables in 2006Q1 and find that the estimates of the U.S. output gap are quite robust. One reason for the robustness relative to allowing for structural change in univariate trend-cycle decompositions is that common changes in unconditional means can be captured as low-frequency movements in the variables included in the model. As we report in our robustness analysis, our estimated output gap appears largely robust to allowing for time variation in $\mu_{\Delta y}$ via dynamically demeaning output growth following Kamber et al. (2018).

5 It is somewhat a matter of semantics about what exactly constitutes a "nowcast" versus a "forecast" or even a "backcast". We consider an estimate of c_{T+1} to be a "nowcast" if it is conducted anytime between the onset of the $T+1$ quarter and when all of the data for the quarter in $\boldsymbol{X}_{T+1}$ become available, at which point c_{T+1} from Eq. (11) is determined. By contrast, we consider any changes to c_{T+1} after all of the monthly data and the advance release of real GDP are available (due to either data revisions or changes in parameter estimates) to be revisions to the "final" estimate of the output gap. Note that we are implicitly defining the output gap as corresponding to the baseline quarterly frequency of the mixed-frequency VAR. That is, we are not thinking of the nowcast as the actual BN cycle of $\{y_t\}$ at a higher frequency than real GDP is observed even though it would clearly satisfy a definition of, say, c_{t-1+v} for $v \in \{1/3, 2/3, 1\}$ based on a shift in timing for Eqs. (6) and (7), although one would have to be careful about units and the higher-frequency variables would have an odd VAR structure with missing lags. Thus, our notion of the output gap is as a percentage measure of how much the flow of production for the economy was above or below its potential flow over the whole quarter, not as a higher-frequency measure.

6 For our nowcasts and as described below, we consider the specific case of calculating conditional expectations as data are sequentially released. The Waggoner and Zha (1999) conditional forecasting approach also allows for more general scenarios where any subset of variables can be fixed and conditional expectations for the remaining variables calculated. Antolín-Díaz et al. (2020b) build on the Waggoner and Zha (1999) approach to allow for conditional density forecasts based on a subset of shocks identified from a structural VAR. We leave such extensions to our empirical application for future research.

or, in expanded form,

$$\begin{bmatrix} \varepsilon_{1,T+1/3} \\ \vdots \\ \varepsilon_{k,T+1/3} \\ \varepsilon_{1,T+2/3} \\ \vdots \\ \varepsilon_{k,T+2/3} \\ \varepsilon_{1,T+1} \\ \vdots \\ \varepsilon_{k,T+1} \\ \varepsilon_{\Delta y,T+1} \end{bmatrix} = \begin{bmatrix} b_{1,1} & 0 & \dots & 0 \\ b_{2,1} & b_{2,2} & 0 & \vdots \\ \vdots & \ddots & \vdots & \vdots \\ b_{3k+1,1} & \dots & \dots & b_{3k+1,3k+1} \end{bmatrix} \begin{bmatrix} z_{1,T+1} \\ \vdots \\ \vdots \\ z_{3k+1,T+1} \end{bmatrix}.$$

We then update the nowcast using the structure implied by Eq. (12). For example, Eq. (12) implies

$$\varepsilon_{1,T+1/3} = b_{1,1} z_{1,T+1}, \tag{13}$$

$$\varepsilon_{2,T+1/3} = b_{2,1} z_{1,T+1} + b_{2,2} z_{2,T+1}, \quad \text{etc...} \tag{14}$$

Conditional on observing $\varepsilon_{1,T+1/3}$, we can use Eq. (13) to solve for $z_{1,T+1}$. Then, a forecast for the rest of the vector $\varepsilon_{2,T+1/3}$ all the way to $\varepsilon_{\Delta y,T+1}$ can be formed conditional on the value of $z_{1,T+1}$ and the remaining elements of the $\boldsymbol{z}_{T+1}$ set to zero. From Eq. (14), conditional on observing both $\varepsilon_{1,T+1/3}$ and $\varepsilon_{2,T+1/3}$, we can solve for $z_{1,T+1}$ and $z_{2,T+1}$, and, thereafter, form a forecast for the rest of the vector $\varepsilon_{3,T+1/3}$ all the way to $\varepsilon_{\Delta y,T+1}$. As the $\boldsymbol{X}_{T+1}$ vector sequentially reveals itself, Eq. (12) provides a structure to update the vector $\boldsymbol{\varepsilon}_{T+1}$, and so updates the nowcast for the output gap through equation (11).

More formally, letting $\omega \in (0, 1)$ correspond to the fraction of the interval of time in which all data for a quarter are released that a particular monthly indicator becomes available, there will be a unique mapping to the variable by release order i assuming no exact simultaneity of availability (perfect simultaneity of a subset of variables means that the ordering for the variables within the subset does not matter). Given the first i-ordered elements of $\boldsymbol{\varepsilon}_{T+1}$ denoted $\boldsymbol{\varepsilon}^i_{T+1}$ and the Cholesky factor $\boldsymbol{B}$, we can solve for $\boldsymbol{z}^i_{T+1} = \boldsymbol{B}_{\boldsymbol{i}}^{-1}\boldsymbol{\varepsilon}^i_{T+1}$ based on Eq. (12) and where $\boldsymbol{B_i}$ corresponds to the first $i \times i$ elements of $\boldsymbol{B}$. Then, given the orthogonality of the elements of $\boldsymbol{z}_{T+1}$, we can set the remaining elements of $\boldsymbol{z}_{T+1}$ to zero to obtain $\boldsymbol{z}_{T+1|T+\omega} = \left[\boldsymbol{z}^i_{T+1} \;\; \boldsymbol{0}\right]'$ and calculate $\boldsymbol{\varepsilon}_{T+1|T+\omega} = \boldsymbol{B}\boldsymbol{z}_{T+1|T+\omega}$. The nowcast for the output gap is then

$$c_{T+1|T+\omega} = -\boldsymbol{s}_k \boldsymbol{F}(\boldsymbol{I} - \boldsymbol{F})^{-1}\left[\boldsymbol{F}\boldsymbol{X}_T + \boldsymbol{H}\boldsymbol{\varepsilon}_{T+1|T+\omega}\right]. \tag{15}$$

4. Empirical application

4.1. Data

In addition to U.S. quarterly real GDP growth, we also consider the federal funds rate in first differences, the 10-year-minus-1-year term spread, the BAA-minus-AAA credit risk spread, S&P500 stock returns, the consumer sentiment index, the civilian unemployment rate, CPI inflation, IP growth, and the growth rate of housing starts. All data series other than real GDP are available at both a monthly and a quarterly frequency and were mostly obtained from FRED to cover the sample period of 1967 to 2019, with observations also collected for 2020 for our out-of-sample analysis in Section 5.[7]

Growth rates, including stock returns, are measured as 100 times first differences of natural logarithms of corresponding levels. Quarterly consumer sentiment and the unemployment rate are averages of the monthly values, while the federal funds rate and interest rate spreads are averages of the daily values (annualized) for a given month or quarter. CPI inflation is measured as a monthly or quarterly rate, not year-on-year. We consider $p = 4$ quarterly lags (i.e., 12 lags of the monthly indicators in the MF-BVAR).

Although our model includes growth rates measured using log differences, we report the output gap in basic percentage deviation terms for the level of real GDP. That is, the reported output gap is $100\% \times \left(e^{\frac{c_t}{100}} - 1\right)$, where c_t is the BN cycle for y_t measured as 100 times the natural logarithm of real GDP. Normally, with relatively "small" values for the output

[7] FRED includes all of the data series except for S&P500 returns, which were obtained from Yahoo Finance. We note that consumer sentiment was only sampled at a lower frequency before 1977. Thus, for the pre-1977 sample period, we interpolated monthly and quarterly values of this variable from the available lower-frequency data. For most of our analysis, the data are from August 2020 vintages. However, for the real-time analysis, we obtained earlier vintages for real GNP/GDP from the Philadelphia Fed Real-Time Data Set and the unemployment rate, IP, and housing starts from ALFRED. We note that, while real GNP/GDP has sizeable revisions, most of the other data series are revised much less or effectively not at all. IP growth can also undergo sizeable revisions, but we find it is not a major source of information in terms of nowcasting the output gap. Possibly related, Barbarino et al. (2020) find that models which rely of labor market data, including the unemployment rate, produce relatively reliable estimates of the output gap in real time. So real-time data issues may be less relevant for nowcasting the output gap with our model given that it includes the unemployment rate. We examine the impact of real-time data issues on the estimated output gap in our robustness analysis.

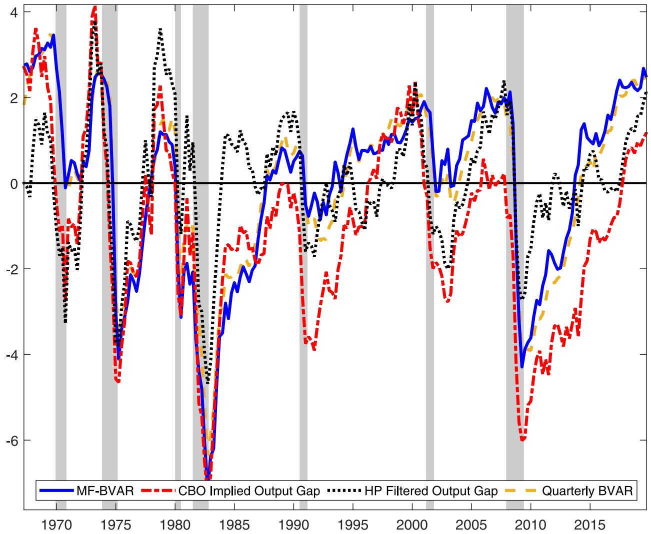

Fig. 1. Comparison of U.S. output gap estimates. Note: The shaded areas denote NBER recession dates.

gap, there is little difference between the units. However, given some large negative nowcasts for the output gap in the out-of-sample analysis of the COVID-19 crisis in Section 5, we believe it is useful to be clear that the measure of the output gap is in percentage deviation terms, not 100 times log differences.

4.2. The estimated output gap

The 9 variables that we consider are similar to those included in the 8-variable BVAR in Morley and Wong (2020).[8] Reflecting this similarity, the output gaps based on the quarterly and mixed-frequency BVARs displayed in Fig. 1 are quite comparable to the reported output gap using the 8-variable (as well as the benchmark 23-variable) quarterly BVAR in Morley and Wong (2020) for the common sample period of 1967 to 2016.

For comparison to other measures of the output gap, Fig. 1 also displays the implied CBO output gap and an output gap based on the HP filter applied to 100 times log real GDP.[9] There is a considerable degree of similarity between these other measures and the BVAR-based output gaps, especially in the case of the implied CBO output gap. The correlations with our estimated output gap based on the MF-BVAR are 80% for the implied CBO output gap and 60% for the HP filter output gap. Thus, nowcasting the output gap based on the MF-BVAR could potentially be useful in tracking changes in the CBO or HP filter measures of the output gap. We discuss this possibility in more detail when looking at real-time data issues in our robustness analysis.

The key thing to notice in Fig. 1 is how similar the estimated output gaps are between the quarterly BVAR and the MF-BVAR that includes monthly indicators of all of the variables other than real GDP growth. The quarterly BVAR can be thought of as restricted version of the MF-BVAR that assumes the same coefficient on a given monthly indicator within each quarter. The fact that the estimated output gaps are so similar, with a correlation of 98%, suggests that this restriction is not particularly problematic and, indeed, could potentially lead to more precision given fewer parameters to estimate. However, even if the output gap for a quarterly BVAR were thought of as a "best" estimate, allowing different coefficients on the within-quarter monthly variables in the MF-BVAR opens up the possibility that data from different months can be more or less informative about the estimated output gap. Thus, we consider the output gap from the MF-BVAR as the "final" estimate in our remaining analysis, but note its very close correspondence to the estimate from the quarterly BVAR in Fig. 1.

8 The 8-variable BVAR in Morley and Wong (2020) included real GDP growth, real PCE growth, the unemployment rate, the growth rate of housing starts, CPI inflation, the first difference of the federal funds rate, real M1 growth, and a different measure of stock returns. See Morley and Wong (2020) for more details.

9 The implied CBO output gap is constructed as the percentage deviation of real GDP from the CBO estimate of potential obtained from FRED. The HP filter output gap is based on the smoothing parameter λ=1,600. For real-time estimates in our robustness analysis, we use earlier vintages of real GNP/GDP and the CBO's estimate of potential available at the time of the initial real GNP/GDP release if it has the same base-year units or the next available estimate of potential if there was an interim change in base-year units of real GNP/GDP.

Table 1
Correlations of nowcasts with the output gap and output growth.

Within-quarter information	Output gap	Output growth
None	0.980	0.700
First month	0.994	0.794
Second month	0.999	0.822
Third month	1.000	0.837

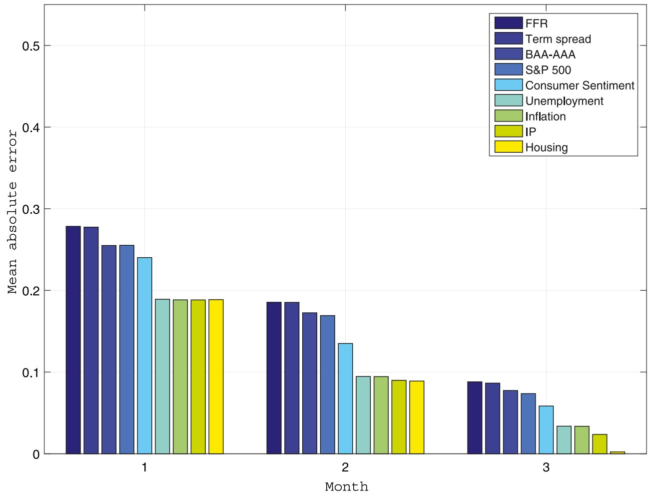

Fig. 2. Average percentage point deviation from final estimate with each monthly data release.

4.3. What monthly information is most relevant?

The data for the MF-BVAR are available in a more timely manner than for the quarterly BVAR. Here we examine what information the release of each monthly indicator contains for the final estimate from the MF-BVAR. Because the BN decomposition effectively acts as a one-sided filter given its reliance on conditional expectations, the release of real GDP for a given quarter completes the information needed to get a final estimate, at least assuming minimal data revisions and stable parameter values. Thus, we can get a direct sense from this analysis of how generally reliable within-quarter nowcasts actually are.

Table 1 reports the correlations between within-quarter nowcasts and the final estimate of the output gap. For comparison, we also report the correlations between the model-implied within-quarter nowcasts for output growth and realized output growth. An immediately striking result is that even at the beginning of the quarter with no within-quarter monthly information available, the correlation for the output gap is already 98%. That is, nowcasts for the output gap will generally be quite reliable due to the persistence of the output gap and many of the variables, such as the unemployment rate, that help forecast output growth. By contrast, the nowcasts for output growth appear much less reliable, presumably due to less persistence in output growth than the output gap from one quarter to the next.

The second thing to notice in Table 1 is that the release of data for the first month within a quarter increases the output gap correlation. Data for the second and third months continue to increase the correlation, with perfect correlation (to the third decimal) given the data from the third month. Thus, with all of the monthly data for a quarter at hand, it appears to be unnecessary to see the real GDP data for a particular quarter in order to get a highly accurate reading of the final estimate of the output gap. In comparison, the output growth correlation also increases with the release of monthly data, but even with all of the monthly data at hand, the nowcast for output growth is still far less reliable and there is considerable new information about output growth in the actual release of real GDP.

Fig. 2 displays the mean absolute error of the within-quarter output gap nowcast for each data release compared to the final estimate. This allows us to see specifically which variables from each month are responsible for the improvement in the correlation between within-quarter nowcasts and the final estimate of the output gap evident in Table 1. To get a sense of statistical significance, Table 2 reports corresponding p-values for two-tailed (Diebold and Mariano, 1995) tests of no change in predictive accuracy under a "lin-lin" (absolute error) loss function with each data release, where the test for federal funds rate in the first month is relative to the nowcast for the output gap when there is no within-quarter monthly information available.

Given a mean absolute error of 0.29 percentage points for the nowcast before any within-quarter monthly information becomes available, we can see that observing the federal funds rate and the term spread in the first month has relatively little impact in terms of improving the nowcast for the output gap. However, the risk spread appears to help, as possibly

Table 2
Diebold–Mariano test *p*-values.

Monthly indicator	Month 1	Month 2	Month 3
Federal funds rate	0.21	0.24	0.13
Term spread	0.30	0.91	0.22
Risk spread	0.02	0.01	0.01
Stock returns	0.67	0.14	0.03
Consumer sentiment	0.11	0.00	0.00
Unemployment rate	0.07	0.00	0.00
CPI inflation	0.31	0.65	0.57
IP growth	0.85	0.01	0.00
Housing starts growth	0.98	0.48	0.00

do consumer sentiment and the unemployment rate, with the mean absolute error reduced by about a third from its initial value given the release of data for the first month. Looking back at Table 1, this reduction in mean absolute error corresponds to the comparatively large increase in the correlation of the nowcast given data from the first month with the final estimate. Meanwhile, the remaining variables appear to have little impact on the reliability of the nowcasts, with the visual impressions in Fig. 2 generally confirmed by the Diebold–Mariano test results in Table 2.

For the second month, the pattern is similar, with a reduction of another third of the mean absolute error from its initial value at the beginning of the quarter driven by information in consumer sentiment and the unemployment rate in similar portions, but also by the risk spread, as in the first month, as well as possibly by IP growth, at least according to the Diebold–Mariano test.

For the third month, most of the variables appear to help improve the reliability of the nowcasts, with the risk spread, stock returns, consumer sentiment, the unemployment rate, IP growth, and housing starts growth all significant at the 5% level for the Diebold–Mariano tests. With the release of housing starts and all of the data from the third month at hand, the remaining mean absolute error is effectively zero, corresponding to the perfect correlation with the final estimate of the output gap in Table 1.

The apparent lack of information in CPI inflation for nowcasting the output gap is notable given that many procedures to estimate the output gap involve imposing a Phillips Curve relationship between the output gap and inflation (see, for example, Kuttner, 1994). Our results suggest that CPI inflation is not particularly informative from an econometric point of view about the output gap once other variables are included in conducting the multivariate trend-cycle decomposition. However, it should be noted that our results are still consistent with a Phillips Curve relationship between the output gap and inflation. In particular, we find a significant correlation of 31% between the quarterly changes in the output gap and CPI inflation.

4.4. Robustness

The mean absolute errors in Fig. 2 were calculated using observations from the full sample period from 1967 to 2019. An immediate question is whether different variables are more or less important at different points of time, including possible structural change in the relevance of variables over the course of the sample period, as well as with the onset of a recession, as would be relevant for the COVID-19 crisis considered in the next section. To address this question, we look at mean absolute errors of output gap nowcasts in the first and second halves of the sample period separately and then in two admittedly small-sample cases: (i) the first two quarters of recession (including the business cycle peak in the first quarter) and ii) the two quarters *after* the onset of a recession. These two cases address the fact that the timing of an onset of recession within a quarter has varied historically, so it is important to check how robust results about information content are to these two related subsamples.

Fig. 3 displays mean absolute errors in the first and second halves of the sample period. The results are similar to those for the full sample period in Fig. 2. The risk spread, consumer sentiment, and the unemployment rate are the key informational variables in both subsamples, with more of the other variables appearing to be informative in the third month. Consumer sentiment appears relatively more informative in the second half of the sample period, while the risk spread is somewhat less informative, especially in the first month. Overall, though, the implication is that the information content in different variables appears relatively stable across the full sample period and that we might still expect the risk spread, consumer sentiment, and the unemployment rate to remain informative out of sample.

Fig. 4 displays mean absolute errors in the two cases around the onset of recessions. The first thing to see is that the results are qualitatively similar to the full-sample results in Fig. 2 despite the small sample of observations in the recession cases (12 observations for the 6 recessions in the 1967–2019 sample period). However, the main quantitative difference, perhaps not surprisingly, is that magnitude of errors starts out higher at the beginning of a quarter with the onset of a recession closer to 0.5 percentage points compared to 0.3 for the full sample period. The risk spread remains an important source of information in the first month, with the mean absolute error dropping around 0.15 percentage points just with the releases of this variable in the first case. Meanwhile, consumer sentiment and the unemployment rate generally appear informative in all three months. In the second case, there is generally a bit more uncertainty about

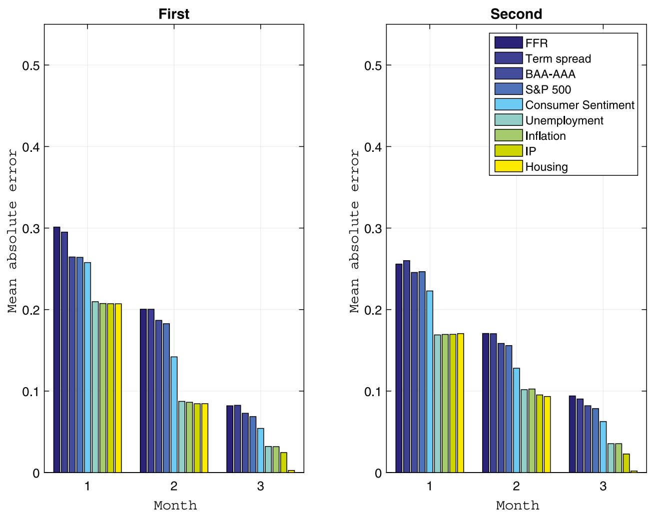

Fig. 3. Average percentage point deviation from final estimate with each monthly data release in the first and second halves of sample period.

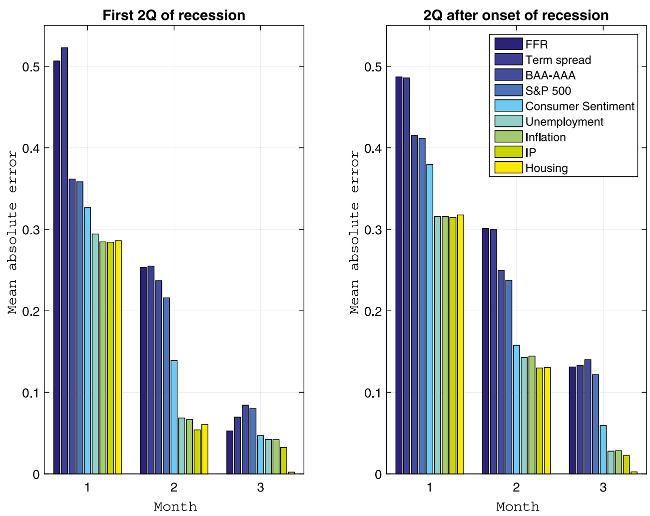

Fig. 4. Average percentage point deviation from final estimate with each monthly data release around the onset of recessions.

the output gap compared to the full sample, but the reliability of the nowcasts is eventually as good with the release of consumer sentiment in the third month. In the first case, uncertainty about the output gap actually becomes lower than for the full sample with the release of the unemployment rate in the second month.

In addition to possible sample sensitivity, another question is how sensitive our results are to including other monthly indicators such as initial claims and non-farm payroll employment growth that are often considered in mixed-frequency analysis (see, for example, McCracken et al., 2020), but which we do not include due to extreme outliers during the COVID-19 crisis. We find that the correlations with our output gap based on the 9-variable MF-BVAR are 97% for a model adding in initial claims and employment growth and 95% for a model adding in log initial claims and employment growth.[10] This high degree of similarity suggests that our benchmark 9-variable model is sufficient to capture relevant information about the output gap. In addition, our benchmark model is effectively as timely given that our most informative monthly indicators are the risk spread, consumer sentiment, and the unemployment rate, which are available before or around the same time as the additional variables, although initial claims can be partly tracked at a higher weekly frequency within a month. Given legislative changes with the COVID crisis, we find that the relationship between initial claims and the other variables in our BVAR breaks down in our out-of-sample analysis, while parameter estimates and nowcasts are more robust for models that exclude initial claims (also see Larson and Sinclair, 2020). We discuss this issue in more detail in our analysis of the COVID crisis in the next section.

10 The data series for initial claims and nonfarm payroll employment were also obtained from FRED.

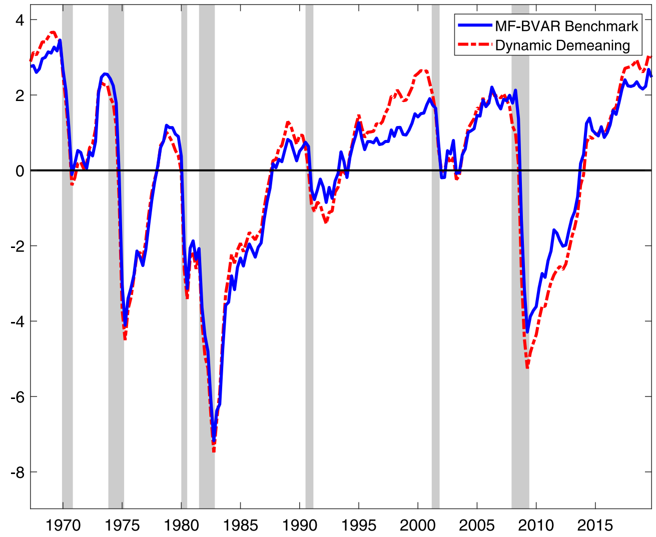

Fig. 5. The U.S. output gap allowing for structural change in trend growth. Note: The shaded areas denote NBER recession dates.

A further issue of robustness that we consider is whether allowing for time variation in the unconditional mean of output growth affects our inferences. Perron and Wada (2009, 2016) argue that it is important to account for structural change in trend growth when conducting trend-cycle decomposition, while a growing number of studies find evidence of a structural break in U.S. trend growth in the mid 2000s, including Antolín-Díaz et al. (2017), Grant and Chan (2017), and Kamber et al. (2018). To account for the possibility of structural change, we follow Kamber et al. (2018) by dynamically demeaning output growth using a backward-looking rolling 40-quarter window. Fig. 5 displays the estimated output gap allowing for structural change and the estimated output gap assuming constant trend growth from Fig. 1 for comparison. The estimates are reasonably similar, especially in periods of recession, consistent with what was found when allowing for structural change in trend growth in Kamber et al. (2018) and Morley and Wong (2020). Thus, our inferences appear robust to allowing for structural change in trend growth.[11]

A final issue of robustness that we consider is the impact of data revisions on our findings. To isolate this impact, we re-calculate the BN cycle for output at each point of time using real-time data vintages. Specifically, the output gap is estimated for each quarter using the MF-BVAR parameters for the full sample period, but with the first available vintage of complete quarterly data for real GDP growth and monthly data for the unemployment rate, IP growth, and housing starts growth for the corresponding quarter. All of the other data, such as for the financial variables, are assumed not be revised and are taken from the final vintage used in our main analysis. Fig. 6 displays the output gap estimated using real-time vintages of data and the estimated output gap using the final vintage of data from Fig. 1 for comparison. The estimated output gaps are strikingly similar with near perfect correlation, suggesting that data revisions play virtually no role in the estimates of the output gap. This is in line with the findings in Orphanides and van Norden (2002) that data revisions play a comparatively small role in the reliability of output gap estimates. However, an immediate reason for the near identical results in our case is that the most informationally-relevant variables for the output gap are not revised very much, if at all. Meanwhile, given that the BN decomposition effectively acts as a one-sided filter, this result means that the only possible remaining source of real-time revisions to the output gap would be parameter uncertainty.

A full analysis of real-time revisions due to parameter uncertainty is challenging given the need for a substantial sample period to estimate our highly-parameterized VAR model, including due to the use of out-of-sample forecast evaluation to optimize the Bayesian shrinkage hyperparameter λ, as discussed in Section 3. However, we discuss possible sensitivity of parameter and output gap estimates to sample period in the next section. We also consider here how well a real-time version of our output gap based on parameter estimates using data only up to 1990Q4 predicts out-of-sample revisions in the implied CBO and HP filter output gaps for an evaluation period of 1990Q4-2019Q4 that was determined by the availability of real-time data for the CBO's estimate of potential in ALFRED. Strikingly, we find that this real-time output gap based on the MF-BVAR and parameters estimated using data only up to 1990Q4 predicts the final-vintage versions of the implied CBO and HP filter output gaps more accurately than their respective real-time counterparts. The correlations are high and similar for the final-vintage implied CBO estimate, with our real-time estimate having a 95% correlation and the real-time implied CBO estimate having a 92% correlation, while the correlations are not as high and more different from each other for the final-vintage/two-sided HP filter output gap, with our real-time estimate having a 68% correlation

[11] Somewhat related, we have also considered robustness of our estimated output gap to consideration of housing starts in dynamically-demeaned log levels instead of growth rates. We found that the estimated output gap remained largely unchanged.

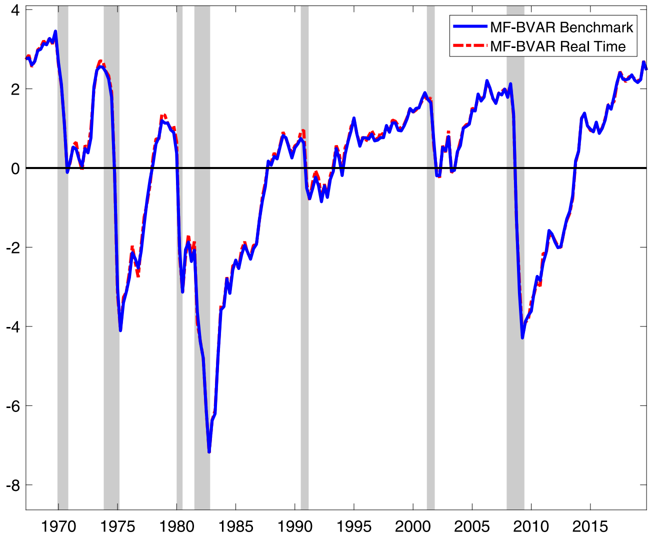

Fig. 6. The U.S. output gap using different vintages of data. Note: The shaded areas denote NBER recession dates.

and the real-time/one-sided HP filter estimate having only a 45% correlation. In terms of revisions, the difference between our real-time estimate and the real-time CBO estimate has a 43% correlation with the revisions in the CBO output gap, while the difference between our real-time estimate and the real-time HP filter estimate has a 53% correlation with the revisions in the HP filter output gap. Thus, our approach appears to provide not just a timely estimate of the output gap according the multivariate BN decomposition of Morley and Wong (2020), but also of the eventual revised output gap estimates for the CBO and the HP filter.

5. Analysis of the COVID-19 crisis

In this section, we apply our estimated model from the previous section to track the U.S. output gap following the onset of a global economic crisis due to the COVID-19 pandemic during the first half of 2020.[12] This analysis is not designed to predict the entire future path of the economy due to the crisis, an understanding of which would require knowing many different factors such as the full extent of policy responses beyond the initial emergency measures. Instead, we consider how our approach anticipates the behavior of the output gap immediately after the crisis began and before the release of real GDP data. The analysis here provides a proper out-of-sample application of our proposed approach to nowcasting the output gap because it considers the model estimated in the previous section with data up to 2019 only, although we also discuss the robustness of parameter estimates to updated data.[13]

5.1. Nowcasts for 2020Q1

The initial large decline in economic activity due to the COVID-19 crisis only occurred during the last few weeks of 2020Q1 (see Lewis et al., 2020), but the crisis led to substantial movements in the various monthly indicators in March compared to their January and February values. Given this, we focus on the data releases for March in tracking how the nowcasts for the output gap in 2020Q1 evolved over the quarter.

Table 3 reports the realized values of the monthly indicators in 2020Q1, as well as nowcasts for the output gap and output growth. The effects of the crisis show up in almost all of the indicators for March, with a clear deterioration of various financial and economic conditions. Related, the nowcast for the output gap only shows much of a change from its previous value of 2.4% in 2019Q4 with a fall to 2.1% given an increase in the risk spread in March that also led to a similar downward revision in the nowcast for output growth. The March data for stock returns and consumer sentiment then actually return the nowcast for the output gap back up to 2.5%, despite leading to a further deterioration of the nowcast

[12] There is a rapidly growing literature on the possible economic consequences of the COVID-19 crisis, including Antolín-Díaz et al. (2020a), Baker et al. (2020), Carriero et al. (2020), Coibion et al. (2020), Jordà et al. (2020b), Larson and Sinclair (2020), Lenza and Primiceri (2020), Lewis et al. (2020), Primiceri and Tambalotti (2020), Schorfheide and Song (2020), amongst many others. Most closely related to our analysis, Antolín-Díaz et al. (2020a), Carriero et al. (2020), and Schorfheide and Song (2020) look at nowcasting U.S. real GDP with mixed-frequency models, but they do not directly look at the output gap.

[13] Lenza and Primiceri (2020) suggest that a simple way to address undue influence of extreme observations with the COVID-19 crisis on VAR parameters is to avoid updating the sample period to include these observations. Likewise, Schorfheide and Song (2020) argue that forecasts based on VAR parameters estimated using only data before the crisis appear “more stable and reasonable” than those based on updated estimates.

Table 3
Monthly indicators and nowcasts in 2020Q1.

(a) Realized monthly indicators

Monthly indicator	January	February	March
Federal funds rate (%)	1.6	1.6	0.7
Term spread (%)	0.2	0.1	0.5
Risk spread (%)	0.8	0.8	1.3
Stock returns (%)	−0.2	−8.8	−13.4
Consumer sentiment (indx.)	99.8	101.0	89.1
Unemployment rate (%)	3.6	3.5	4.4
CPI inflation (%)	0.1	0.1	−0.4
IP growth (%)	−0.4	0.1	−4.5
Housing starts growth (%)	1.9	−3.1	−19.0

(b) Nowcasts for the output gap and output growth

Within-quarter information		Output gap (%)	Output growth (%)
None		2.5	0.5
January		2.4	0.6
February		2.5	0.5
March:	Federal funds rate	2.4	0.6
	Term spread	2.4	0.6
	Risk spread	2.1	0.2
	Stock returns	2.3	0.1
	Consumer sentiment	2.5	0.0
	Unemployment rate	2.2	−0.2
	CPI inflation	2.1	−0.4
	IP growth	2.4	−0.6
	Housing starts growth	2.4	−0.6
Final:	Real GDP (advance)	2.4	−1.2

for output growth. The March unemployment rate and CPI inflation lower the nowcast for the output gap again to 2.1% and imply further declines in the nowcast for output growth. Finally, March IP growth brings the nowcast for the output gap back to its original value of 2.4%, where the estimate ends up with the release of real GDP growth for 2020Q1, while again producing a further drop in the nowcast for output growth, although not all the way to its realized value of −1.2%.

It is notable that the March values for stock returns, consumer sentiment, and IP growth, which signaled deteriorating conditions in the economy, did not lead to downward revisions in the nowcast for the output gap, while they did so with the nowcast for output growth. This difference is because the nowcast for the output gap ultimately depends on what information each data release contains not just for current output growth, but also for future output growth. The long-horizon calculation involved in the BN decomposition is evidently more stable in the face of new information than the prediction of current-quarter output growth. Meanwhile, just as the exact match of the nowcast for the output gap given all monthly data with the final estimate is consistent with the findings in Table 1, so to is the difference between the nowcast for output growth given all the monthly data and the realized value of output growth. Furthermore, the almost perfect coherence between real-time and final vintage estimates of the output gap in Fig. 6 means that we should not expect any major change in the estimated output gap even when some observations, including the advance estimate of real GDP growth in 2020Q1, are revised over time.[14]

5.2. Nowcasts for 2020Q2

Next, we track how the nowcasts for the output gap in 2020Q2 evolved as the full severity of the COVID-19 crisis showed up in the various monthly indicators. To provide a baseline, we first consider the nowcast at the beginning of the quarter before any monthly data became available. Then we consider nowcasts given data for April and then for May and then for June. For these nowcasts, we use the Waggoner and Zha (1999) approach to consider both conditional expectations and a scenario based on the behavior of monthly indicators in previous recessions.

For the scenario nowcasts, we follow Berger and Vierke (2017) and calculate the average paths of our monthly indicators from their values at business cycle peaks, as dated by the NBER. According to the NBER, the most recent business cycle peak occurred in February 2020. However, we note that the recession due to the COVID-19 pandemic differs from previous recessions given the sudden stop of large sectors of business and production immediately in March causing an extraordinary speed and scale of decline in economic activity, especially in terms of the labor market. Thus, as described below, we adjust the scale and timing of the average responses of our monthly indicators to reflect the rapid and severe effects of the COVID-19 recession.

[14] For example, the second estimate of quarterly real GDP growth in 2020Q1 was −1.3%, which did not lead to any change in the estimated output gap.

Table 4
Monthly indicators and nowcasts in 2020Q2.

(a) Projected and realized monthly indicators

Monthly indicator	April	May	June
	Scenario/Realized	Scenario/Realized	Scenario/Realized
Federal funds rate (%)	0.1/0.1	0.1/0.1	0.1/0.1
Term spread (%)	0.5/0.5	0.5/0.5	0.5/0.6
Risk spread (%)	1.3/1.7	1.3/1.5	1.3/1.2
Stock returns (%)	0.0/11.9	0.0/4.4	0.0/1.8
Consumer sentiment (indx.)	84.8/71.8	65.0/72.3	62.7/78.1
Unemployment rate (%)	12.9/14.7	12.3/13.3	12.2/11.1
CPI inflation (%)	0.2/−0.8	0.2/−0.1	−0.3/0.6
IP growth (%)	−1.3/−13.6	−2.6/1.4	−3.3/5.3
Housing starts growth (%)	−3.9/−26.4	15.5/11.1	4.8/17.5

(b) Nowcasts for the output gap and output growth

Within-quarter information		Output gap (%)	Output growth (%)
None:	Conditional/Scenario (April–June)	1.2/−7.3	−0.5/−1.6
April:	Conditional/Scenario (May–June)	−10.1/−7.5	−7.4/−6.0
May:	Conditional/Scenario (June)	−9.3/−8.1	−5.3/−5.0
June:	Conditional	−8.3	−3.7
Final:	Real GDP (advance)	−8.3	−9.5

Federal funds rate. With the onset of the recession, the federal funds rate was reduced to its effective lower bound and it was immediately clear that the Federal Reserve would keep it there well beyond 2020Q2 and possibly for years. Thus, we keep the federal funds rate fixed at 0.1% in April, May, and June for our scenario nowcasts. We note that a different scenario regarding the federal funds rate could be used to determine what would happen if the Federal Reserve were to hypothetically change the federal funds rate within the forecasting horizon.[15] Furthermore, unconventional policy could be addressed, for example, by incorporating a shadow rate or other measures of these policies. However, we leave such analysis for future research.

Term spread, risk spread, and stock returns. For our financial indicators, we use a simple random walk assumption for our scenarios. Specifically, the term spread and risk spread are projected to remain at their March levels in April, May, and June, while stock returns are projected to be equal to zero.

Consumer sentiment, unemployment rate, CPI inflation, IP growth, and housing starts growth. With the exception of the unemployment rate, the scenario values for these monthly indicators are set based on their average responses in the second to fourth months after the business cycle peak, with the average responses multiplied by a factor of three to capture the severity of the COVID-19 recession. Because of the unusually rapid effect of the recession on the labor market, as was evident, for example, in the unprecedented rise in initial claims for unemployment insurance in the middle of March, we do not consider the average responses of the unemployment rate in the early months of recession, which is usually comparatively muted, but instead consider the average cumulative response to the peak unemployment rate in projecting a value for April and then the average responses in the two months after the peak in unemployment when projecting values for May and June. Again, we multiply these average responses by a factor of three to reflect the severity of the COVID-19 recession.

Table 4 reports the projected and realized values of the monthly indicators in 2020Q2, as well as nowcasts for the output gap and output growth. We note that many of the projected values for our scenario are in line with the values we found were implied by the conditional expectations of the MF-BVAR, although there are some key exceptions that motivate our consideration of the scenario projections in addition to conditional expectations. In particular, we found that the conditional expectations project negative values for the federal funds rate and generally quite optimistic values for consumer sentiment and the unemployment rate, at least at the start of the quarter. By contrast, the scenario projections keep the federal funds rate fixed at the effective lower bound and imply comparatively lower consumer sentiment and a higher unemployment rate according to their dynamic paths in previous recessions.

In terms of the nowcasts, even the conditional nowcast at the beginning of the quarter before any monthly data becomes available suggests a nontrivial one-quarter decline of 1.2 percentage points, while the scenario nowcast implies an unprecedented 9.7 percentage point decline, the difference largely due to the relatively pessimistic projection for the unemployment rate compared to the conditional nowcast. However, the conditional nowcast given the April data catches

[15] To be clear, however, these scenarios are conditional on observables, not structural shocks, which would require identification of structural shocks. See Antolín-Díaz et al. (2020b) on how to conduct structural scenario analysis with VARs. Also, see Morley and Wong (2020) for a discussion of how to decompose movements in trend and cycle into structural shocks from an identified SVAR model when using the multivariate BN decomposition. We note that such a structural assessment of the output gap is not the focus of our nowcasting application, which requires a MF-BVAR that spans the relevant information with timely indicators rather than variables that might be needed for structural identification.

up and exceeds the scenario nowcast in terms of the implied decline in the output gap. Again, this is due to a larger increase in the unemployment rate in April than was even anticipated given the seemingly pessimistic scenario. Then, with the release of the May data, the conditional and scenario nowcasts start to converge to each other and towards the final estimate of −8.3%. This convergence is not surprising given that only the projected June values could differ and we find that they do not differ as much as at the beginning of the quarter, with the projection for the unemployment rate in June actually more pessimistic for the conditional nowcast than the scenario nowcast. These patterns are evident in the conditional and scenario nowcasts for output growth as well, although the nowcast for output growth given the June data actually increases and ends up more optimistic than the advance release, while the nowcast for the output gap given the June data is the same as the final estimate.[16] As with the 2020Q1 nowcasts, the disparate performance of output gap and output growth nowcasts is consistent with our general findings about their comparative reliability presented in Table 1.

Regarding the scenario nowcasts, we emphasize that they are primarily for illustration of how the Waggoner and Zha (1999) approach can be used in conjunction with nowcasting the output gap. Other credible ex ante scenario projections could have been proposed and the scenario projections could have been updated month by month. The key point, however, is that scenario nowcasts will only be more reliable than the conditional nowcasts to the extent that the scenario projections for the monthly indicators end up closer to their actual realized values than projections based on conditional expectations. This was clearly the case for the scenario projections for the federal funds rate given that they avoided going negative, but also, importantly, ended up being the case for the unemployment rate, at least at the beginning of the quarter. However, any external considerations that lead to improved predictions of monthly indicators would be helpful in improving the scenario nowcasts, which can be thought of as akin to nowcasts based on a bridge equation with forecasts for monthly indicators determined externally.

5.3. Comparison with other estimates for the COVID-19 crisis

If we consider the larger BVARs also discussed in the robustness section that included initial claims and employment growth, we find that estimated output gap for 2020Q2 is −31.2% for the model with initial claims in levels and −14.4% for the model with initial claims in logs. Thus, despite a close correspondence between the output gaps for these models and our benchmark model during the 1967–2019 sample, there is much more sensitivity out of sample. This sensitivity can be explained by the extreme outlier observations for initial claims and employment growth starting in March that resulted from the COVID-19 crisis. One way to see this point is to consider what happens if we update the parameter estimates based on data up to 2020Q2. In this case, the estimated output gap for these other models change dramatically and converge towards the estimate for our benchmark model. In particular, the estimated output gap for 2020Q2 becomes −8.2% for the model with initial claims in levels and −9.2% for the model with initial claims in logs. Because the extended data series for initial claims and employment growth effectively act like dummy variables for 2020Q2 rather than data series that are strongly correlated with other variables such as the unemployment rate, the estimated coefficients on these variables change dramatically with an updated estimation sample and the estimated larger BVARs become much more like our benchmark model, leading to similar estimates of the output gap compared what we found for our benchmark model.

Notably, the estimated output gap in 2020Q2 for our benchmark model changes much less from −8.3% to −8.0% when updating the estimation sample, suggesting that parameter estimates for the model are much more stable even given large outliers in many of the remaining variables associated with the crisis. We note that it is not really feasible to throw out all variables with outliers in 2020Q2 from the model given that even real GDP growth took on an historically large negative value. The important point is to keep variables in the model that have similar scale movements as real GDP growth and also maintain their historical correlations with real GDP growth. Our benchmark MF-BVAR without initial claims and employment growth appears to do this.

To compare with other measures of the output gap, Fig. 7 displays the real-time and ex post estimates of the output gap based on our MF-BVAR, the CBO measure of potential, and the HP filter over the period from 2007 to 2020. We start the plot in 2007 to provide some context for the scale of the decline in the output gap with the COVID-19 recession in comparison to the Great Recession in 2007–2009. In particular, the estimated output gap for all three measures takes on a larger negative value in 2020Q2 than at any point during the Great Recession. Also, going back to the Great Recession, the relative reliability of the output gap based on the MF-BVAR is clear, with upward revisions in the implied CBO and HP filter output gaps towards our estimated output gap already occurring for the quarters just before the onset of the COVID-19 recession due to revisions in the CBO's estimate of potential in August 2020 and the difference between one-sided and two-sided estimates for the HP filter. This mirrors what happened to the estimates around the Great Recession, with the real-time implied CBO and HP filter estimates revised ex post towards the MF-BVAR output gap, almost perfectly so in the case of the HP filter. Indeed, the ability of the real-time MF-BVAR output gap to predict revisions in the implied CBO and HP filter estimates is evident throughout the period displayed in Fig. 7, especially with an overly pessimistic real-time implied CBO output gap and overly optimistic real-time HP filter output gap during the early years of the recovery

[16] Although it is unlikely to alter the estimate for the output gap much, we note that there could be a large revision in output growth for 2020Q2, as discussed in Jordà et al. (2020a).

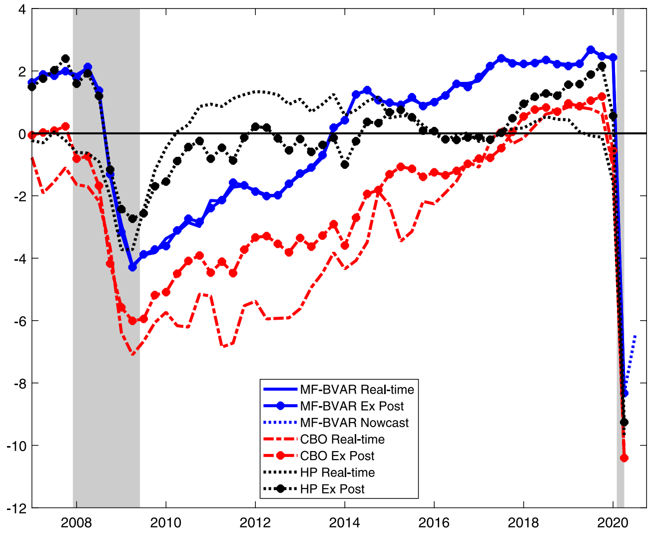

Fig. 7. Real-time and ex post U.S. output gap estimates from 2007 to 2020. Notes: The shaded areas denote NBER recession dates (assumes trough in 2020Q2).

following the Great Recession.[17] Finally, Fig. 7 also includes the conditional nowcast for 2020Q3 given data for the July monthly indicators. The nowcast of −6.4% suggests 2020Q2 was a trough, but with only a very partial recovery in 2020Q3.

6. Conclusion

We have proposed a way to produce direct and timely estimates of the output gap which can be updated as higher-frequency data, such as monthly indicators, become available. We find that our mixed-frequency approach produces very similar estimates of the U.S. output gap compared to those based on a quarterly BVAR in Morley and Wong (2020), as well as being highly correlated estimates with CBO and HP filter measures. Crucially, our approach does not require waiting until after a quarter ends to provide reliable estimates. Monthly indicators for a credit risk spread, consumer sentiment, and the unemployment rate contain particularly useful new information when nowcasting the output gap, although the information content varies by which month in the quarter the data are taken from. Our out-of-sample analysis of the COVID-19 crisis anticipates the exceptionally large negative output gap of −8.3% in 2020Q2 before the release of real GDP data for the quarter, with both conditional and scenario nowcasts tracking a dramatic decline in the output gap given the April data.

Acknowledgments

We thank an anonymous Associate Editor and an anonymous referee, as well as Giovanni Caggiano, Todd Clark, Mike McCracken, Warwick McKibbin, and participants in the Virtual Australian Macro Seminar and CAMA Webinar for helpful comments and suggestions. The latest nowcasts and data for the output gap based on this paper are available at http://outputgapnow.com and we thank Franz Weitkamp for help with the webpage. This research was supported by Australian Research Council grants DP190100202 and DE200100693 and by the Australia–Germany Joint Research Co-operation Scheme initiative of Universities Australia and the German Academic Exchange Service (DAAD). The usual disclaimers apply.

References

Antolín-Díaz, J., Drechsel, T., Petrella, I., 2017. Tracking the slowdown in long-run GDP growth. Rev. Econ. Stat. 99 (2), 343–356.

Antolín-Díaz, J., Drechsel, T., Petrella, I., 2020a. Advances in nowcasting economic activity: Secular trends, large shocks and new data. Working paper.

Antolín-Díaz, J., Petrella, I., Rubio-Ramírez, J.F., 2020b. Structural scenario analysis with SVARs. J. Monetary Econ. in press.

Baker, S.R., Bloom, N., Davis, S.J., Terry, S.J., 2020. COVID-induced economic uncertainty. NBER Working Paper 26983.

Bańbura, M., Giannone, D., Modugno, M., Reichlin, L., 2013. Now-casting and the real-time data flow. In: Elliott, G., Timmermann, A. (Eds.), Handbook of Economic Forecasting. In: Handbook of Economic Forecasting, 2, Elsevier, pp. 195–237, Chapter 4.

[17] The extent to which the real-time MF-BVAR output gap predicts revisions in the implied CBO and HP filter estimates for a particular quarter likely reflects the accuracy of the MF-BVAR forecasts for future output growth given that the CBO and the two-sided HP filter are clearly influenced by realizations of future output growth when making revisions.

Bańbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector auto regressions. J. Appl. Econometrics 25 (1), 71–92.

Barbarino, A., Berge, T.J., Chen, H., Stella, A., 2020. Are all output gap estimates unstable in real time? Working paper.

Berger, T., Vierke, H., 2017. Estimating the natural rate of hours. Macroecon. Dyn. 21 (6), 1426–1453.

Beveridge, S., Nelson, C.R., 1981. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. J. Monetary Econ. 7 (2), 151–174.

Brave, S.A., Butters, R.A., Justiniano, A., 2019. Forecasting economic activity with mixed frequency BVARs. Int. J. Forecast. 35 (4), 1692–1707.

Carriero, A., Clark, T.E., Marcellino, M., 2020. Nowcasting tail risks to economic activity with many indicators. Working paper.

Coibion, O., Gorodnichnenko, Y., Weber, M., 2020. Labor markets during the COVID-19 crisis: A preliminary view. Working paper.

Del Negro, M., Schorfheide, F., 2011. Bayesian macroeconometrics. In: Geweke, J., Koop, G., Van Dijk, H. (Eds.), The Oxford Handbook of Bayesian Econometrics. Oxford University Press, pp. 293–389.

Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. J. Bus. Econom. Statist. 13 (3), 253–263.

Foroni, C., Marcellino, M., Schumacher, C., 2015. Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. J. R. Stat. Soc.: Ser. A 178 (1), 57–82.

Garratt, A., Mitchell, J., Vahey, S.P., 2014. Measuring output gap nowcast uncertainty. Int. J. Forecast. 30 (2), 268–279.

Ghysels, E., 2016. Macroeconomics and the reality of mixed frequency data. J. Econometrics 193 (2), 294–314.

Grant, A.L., Chan, J.C.C., 2017. A Bayesian model comparison for trend-cycle decompositions of output. J. Money Credit Bank. 49 (2–3), 525–552.

Jordà, O., Kouchekinia, N., Merrill, C., Sekhposyan, T., 2020a. The fog of numbers. FRBSF Economic Letter 2020-20.

Jordà, O., Singh, S.R., Taylor, A.M., 2020b. Longer-run economic consequences of pandemics. Federal Reserve Bank of San Francisco Working Paper 2020-09.

Kamber, G., Morley, J., Wong, B., 2018. Intuitive and reliable estimates of the output gap from a Beveridge-Nelson filter. Rev. Econ. Stat. 100 (3), 550–566.

Koop, G.M., 2013. Forecasting with medium and large Bayesian VARs. J. Appl. Econometrics 28 (2), 177–203.

Kuttner, K.N., 1994. Estimating potential output as a latent variable. J. Bus. Econom. Statist. 12 (3), 361–368.

Larson, W.D., Sinclair, T.M., 2020. Nowcasting unemployment insurance claims in the time of COVID-19. FHFA Staff Working Paper Series 20-02.

Lenza, M., Primiceri, G.E., 2020. How to Estimate a VAR after March 2020. Working paper.

Lewis, D., Mertens, K., Stock, J.H., 2020. U.S. Economic activity during the early weeks of the SARS-COV-2 outbreak. NBER Working Paper 26954..

Luciani, M., Barigozzi, M., 2019. Measuring the output gap using large datasets. SSRN Working Paper 3217816.

McCracken, M., Owyang, M., Sekhposyan, T., 2020. Real-time forecasting and scenario analysis using a large mixed-frequency Bayesian VAR. Working paper, Federal Reserve Bank of St. Louis.

Morley, J.C., 2002. A state-space approach to calculating the Beveridge-Nelson decomposition. Econom. Lett. 75 (1), 123–127.

Morley, J., Wong, B., 2020. Estimating and accounting for the output gap with large Bayesian vector autoregressions. J. Appl. Econometrics 35 (1), 1–18.

Orphanides, A., van Norden, S., 2002. The unreliability of output-gap estimates in real time. Rev. Econ. Stat. 84 (4), 569–583.

Perron, P., Wada, T., 2009. Let's take a break: Trends and cycles in U.S. real GDP. J. Monetary Econ. 56 (6), 749–765.

Perron, P., Wada, T., 2016. Measuring business cycles with structural breaks and outliers: Applications to international data. Res. Econ. 70 (2), 281–303.

Primiceri, G.E., Tambalotti, A., 2020. Macroeconomic forecasting in the time of COVID-19. Working paper.

Schorfheide, F., Song, D., 2015. Real-time forecasting with a mixed-frequency VAR. J. Bus. Econom. Statist. 33 (3), 366–380.

Schorfheide, F., Song, D., 2020. Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic. Working paper.

Stock, J.H., Watson, M.W., 2002. Macroeconomic forecasting using diffusion indexes. J. Bus. Econom. Statist. 20 (2), 147–162.

Waggoner, D.F., Zha, T., 1999. Conditional forecasts in dynamic multivariate models. Rev. Econ. Stat. 81 (4), 639–651.

Woźniak, T., 2016. Bayesian vector autoregressions. Aust. Econ. Rev. 49 (3), 365–380.