

CAMA

Centre for Applied Macroeconomic Analysis

Estimating and Accounting for the Output Gap with Large Bayesian Vector Autoregressions

CAMA Working Paper 46/2017
July 2017

James Morley

University of New South Wales and
Centre for Applied Macroeconomic Analysis, ANU

Benjamin Wong

Reserve Bank of New Zealand and
Centre for Applied Macroeconomic Analysis, ANU

Abstract

We demonstrate how Bayesian shrinkage can address problems with utilizing large information sets to calculate trend and cycle via a multivariate Beveridge-Nelson (BN) decomposition. We illustrate our approach by estimating the U.S. output gap with large Bayesian vector autoregressions that include up to 138 variables. Because the BN trend and cycle are linear functions of historical forecast errors, we are also able to account for the estimated output gap in terms of different sources of information, as well as particular underlying structural shocks given identification restrictions. Our empirical analysis suggests that, in addition to output growth, the unemployment rate, CPI inflation, and, to a lesser extent, housing starts, consumption, stock prices, real M1, and the federal funds rate are important conditioning variables for estimating the U.S. output gap, with estimates largely robust to incorporating additional variables. Using standard identification restrictions, we find that the role of monetary policy shocks in driving the output gap is small, while oil price shocks explain about 10% of the variance over different horizons.

Keywords

Beveridge-Nelson decomposition, output gap, Bayesian estimation, multivariate information

JEL Classification

C18, E17, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Estimating and Accounting for the Output Gap with Large Bayesian Vector Autoregressions ^{*†}

James Morley^{1,3} and Benjamin Wong^{2,3}

¹University of New South Wales

²Reserve Bank of New Zealand

³Centre for Applied Macroeconomic Analysis, The Australian National University

Abstract

We demonstrate how Bayesian shrinkage can address problems with utilizing large information sets to calculate trend and cycle via a multivariate Beveridge-Nelson (BN) decomposition. We illustrate our approach by estimating the U.S. output gap with large Bayesian vector autoregressions that include up to 138 variables. Because the BN trend and cycle are linear functions of historical forecast errors, we are also able to account for the estimated output gap in terms of different sources of information, as well as particular underlying structural shocks given identification restrictions. Our empirical analysis suggests that, in addition to output growth, the unemployment rate, CPI inflation, and, to a lesser extent, housing starts, consumption, stock prices, real M1, and the federal funds rate are important conditioning variables for estimating the U.S. output gap, with estimates largely robust to incorporating additional variables. Using standard identification restrictions, we find that the role of monetary policy shocks in driving the output gap is small, while oil price shocks explain about 10% of the variance over different horizons.

JEL Classification: C18, E17, E32

Keywords: Beveridge-Nelson decomposition, output gap, Bayesian estimation, multivariate information

*Email: Morley: james.morley@unsw.edu.au Wong: benjamin.wong@rbnz.govt.nz

†We thank participants at the 25th Symposium of the Society for Nonlinear Dynamics and Econometrics, 11th RCEA Bayes Workshop, and seminar audiences at the University of Melbourne and the Australian National University for helpful comments and suggestions. The views expressed in this paper do not necessarily represent the views of the Reserve Bank of New Zealand. Any remaining errors are ours.

1 Introduction

Interpretation of macroeconomic data often involves decomposing time series into trend and cycle, especially as related concepts such as the neutral rate of interest, the output gap, and trend inflation are crucial inputs into macroeconomic policy decision making. The macroeconomic literature is replete with statistical methods to conduct such decompositions (e.g., Hodrick and Prescott, 1997; Christiano and Fitzgerald, 2003). These methods are often univariate and so only rely on a single time series (i.e., the series being detrended) for implementation. One challenge with univariate detrending is that the interpretation of the estimated trend and cycle from a statistical filter often needs to be corroborated “off-model” with other sources of information. While one could allow an explicit role for multivariate information to help conduct and interpret trend-cycle decomposition (e.g., Kozicki, 1999; Garratt, Robertson, and Wright, 2006; Sinclair, 2009; Garratt, Lee, and Shields, 2016; Chan and Grant, 2017), practical challenges remain in terms of which variables should be included in the information set or even with how large the information set can be to keep estimation tractable.

We address these issues associated with specifying and interpreting multivariate information within the context of a particular approach to calculating trend and cycle, namely the Beveridge and Nelson (BN) (1981) decomposition. The BN decomposition provides estimates of the trend and cycle for a time series by looking at its long-horizon conditional forecast. Meanwhile, vector autoregressions (VARs) are widespread and well developed models for forecasting and provide a natural starting point for incorporating multivariate information to calculate trend and cycle via the BN decomposition (see, for example, Evans and Reichlin, 1994). However, their use can give rise to practical concerns such as overfitting or a mechanical and possibly spurious decrease in the signal-to-noise ratio given more information, the latter a theoretical property of the multivariate BN decomposition as proven by Evans and Reichlin (1994). We show that Bayesian shrinkage can circumvent these practical challenges and that estimation of trend and cycle utilizing information sets containing well over a hundred variables is both feasible and can avoid overfitting.

We also show how to infer which sources of information are the most important for estimating trend and cycle, providing a guide for variable selection and setting the appropriate size of the information set in practice. In particular, we demonstrate that, because the estimated trend and cycle from a multivariate BN decomposition are linear functions of the historical forecast errors, the contribution of various sources of information to the estimated trend and cycle can be easily determined. Given the forecast errors and identification restrictions, an extension to structural VAR (SVAR) analysis is also straightforward when the objective is to infer which economic shocks drive the trend and cycle.

A key empirical finding from our analysis is that the U.S. output gap estimated using a large Bayesian VAR (BVAR) with 138 variables is similar to estimates for smaller BVARs with 23 or even just 8 variables, if appropriately selected. The reason for the robust results, despite information sets of very different sizes, is because of the Bayesian shrinkage, as well as our ability to determine which variables are the most important sources of information for the output gap. In particular, the most important variables for estimating the U.S. output gap in addition to output growth itself are the unemployment rate, CPI inflation, and, to a lesser extent, housing starts, consumption, stock prices, real M1, and the federal funds rate, with estimates largely robust to incorporating additional variables. Identifying monetary policy shocks and oil price shocks using standard restrictions in our 23 variable benchmark system, we find they respectively account for approximately 4% and 10% of the unconditional variance of the U.S. output gap. We also find that neither type of shock explains very much of the variance of trend growth, consistent with monetary neutrality and traditional theories of growth that assume technology shocks are the main determinant of the long-run level of output.

Our approach and analysis are heavily influenced by Evans and Reichlin (1994), who explored the multivariate generalization of the original univariate BN decomposition in Beveridge and Nelson (1981). Figure 1 presents an example to illustrate the two key insights from Evans and Reichlin (1994) that we build on. First, Evans and Reichlin show that adding relevant multivariate information into the forecasting model significantly alters the profile of the estimated output gap. This can be seen by comparing the output gap obtained from a BN decomposition based on a univariate AR(4) model of U.S. output growth plotted in the top panel of the figure with various multivariate BN decompositions based on VARs containing larger information sets, ranging from 2 to 7 variables, plotted in the bottom panel.¹ The univariate BN estimate of the output gap lacks both persistence and amplitude and also moves counter-cyclically with the NBER chronology of expansions and recessions, while the multivariate ones are more persistent, have larger amplitude, and move more pro-cyclically relative to the NBER chronology. Second, Evans and Reichlin show that the larger the information set, the lower the signal-to-noise ratio. This effect can be easily seen as the estimated output gap in the bottom panel clearly becomes larger in amplitude when the number of variables increases, with the 7 variable VAR producing the largest amplitude estimate of the output gap. Evans and Reichlin’s insights are powerful because they suggest that multivariate information should serve some role in making inference about the output gap, yet increasing the amount of information needs to be balanced with concerns about overfitting and specifically what is the appropriate

¹The data are for the U.S. economy and are described in full detail in Section 3. The 2 variable VAR includes output growth and the unemployment rate. The 3 variable VAR includes output growth, CPI inflation, and the federal funds rate. The 7 variable VAR includes all of the variables in the 2 and 3 variable system, as well as capacity utilization, the growth of industrial production, and the growth of real personal consumption expenditure. All of the VARs were estimated with 4 lags using least squares.

signal-to-noise ratio. Our goal in this paper is to provide a practical solution to balance these tradeoffs and address other challenges with large multivariate models. In particular, we suggest how to specify and estimate a multivariate forecasting model to handle various information sets, as well as how to determine the role of other variables and underlying structural shocks in explaining the estimated trend and cycle.

The remainder of this paper proceeds as follows. Section 2 discusses the BN decomposition in a multivariate setting and how to infer the role of different sources of information and underlying structural shocks given identification restrictions. Section 3 describes how to use Bayesian shrinkage to utilize large information sets and describes the data to be included in the large BVARs. Section 4 reports the resulting estimates of the U.S. output gap and considers the role of multivariate information in determining these estimates. Section 5 employs SVAR analysis to examine possible causal determinants of the U.S. output gap and trend growth. Section 6 considers Monte Carlo analysis to show how well our approach works when possibly misspecifying the size of the BVAR. We conclude by summarizing our key findings and also suggesting some possible extensions.

2 The BN Decomposition in a Multivariate Setting

Beveridge and Nelson (1981) define the trend of a time series as its long-horizon conditional expectation minus any deterministic future movements in the time series. In particular, let $\{y_t\}$ be a time series with a trend component that follows a random walk with a constant drift, μ , then the BN trend, τ_t , at time t is

$$\tau_t = \lim_{j \rightarrow \infty} \mathbb{E}_t[y_{t+j} - j \cdot \mu]. \quad (1)$$

The intuition behind the BN decomposition is that the long-horizon the expectation of the cycle is zero, meaning that the long-horizon conditional expectation of the time series will just reflect the trend component. Therefore, one only needs to specify a forecasting model for the first difference of the time series, $\{\Delta y_t\}$, to generate these forecasts. For univariate time series, ARMA models provide a natural choice for the forecasting model, as considered in Beveridge and Nelson (1981). In the multivariate context, linear VARs provide a natural choice, as considered in Evans and Reichlin (1994).

Let $\{\Delta \mathbf{x}_t\}$ represent a vector of stationary variables that includes $\{\Delta y_t\}$.² We can write a finite-order VAR in companion form as

$$(\Delta \mathbf{x}_t - \boldsymbol{\mu}) = \mathbf{F}(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu}) + \mathbf{H}\boldsymbol{\nu}_t, \quad (2)$$

²By framing the stationary variables as being in differences, we can apply the BN decomposition to the integrated levels of these variables, $\{\mathbf{x}_t\}$, which importantly includes $\{y_t\}$, although variables that are naturally stationary in their levels could be included in $\{\Delta \mathbf{x}_t\}$ and the BN decomposition would implicitly be applied to the accumulation of their levels.

where \mathbf{F} is the companion matrix, $\boldsymbol{\mu}$ a vector of unconditional means, \mathbf{H} maps the forecast errors to the companion form, and $\boldsymbol{\nu}_t$ is a vector of serially uncorrelated forecast errors with covariance matrix $\boldsymbol{\Sigma}_\nu$. Given stationarity, $(\mathbf{I} - \mathbf{F})^{-1}$ exists and, from equation (2), the cumulative sum at time t of expected future deviations of the vector process from its unconditional mean can be written as

$$\mathbb{E}_t \sum_{j=1}^{\infty} (\Delta \mathbf{x}_{t+j} - \boldsymbol{\mu}) = \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}(\Delta \mathbf{x}_t - \boldsymbol{\mu}). \quad (3)$$

Then, denoting $\boldsymbol{\tau}_t$ and \mathbf{c}_t respectively as the vector of BN trends and BN cycles of \mathbf{x}_t at time t , we can solve for these following Morley (2002):

$$\boldsymbol{\tau}_t = \mathbf{x}_t + \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}(\Delta \mathbf{x}_t - \boldsymbol{\mu}) \quad (4)$$

$$\mathbf{c}_t = -\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}(\Delta \mathbf{x}_t - \boldsymbol{\mu}). \quad (5)$$

The BN cycles, \mathbf{c}_t , can be written as a function of the history of forecast errors through recursive substitution. Letting $\boldsymbol{\Gamma}_i \equiv \mathbf{F}^i(\mathbf{I} - \mathbf{F})^{-1}$ for notational convenience and repeatedly lagging and substituting equation (2) into equation (5), we get the following representation:

$$\begin{aligned} \mathbf{c}_t &= -\boldsymbol{\Gamma}_1(\Delta \mathbf{x}_t - \boldsymbol{\mu}) \\ &= -\boldsymbol{\Gamma}_1 \{ \mathbf{F}(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu}) + \mathbf{H}\boldsymbol{\nu}_t \} \\ &= -\boldsymbol{\Gamma}_1 \mathbf{H}\boldsymbol{\nu}_t - \{ \mathbf{F}\boldsymbol{\Gamma}_1(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu}) \} \\ &= -\boldsymbol{\Gamma}_1 \mathbf{H}\boldsymbol{\nu}_t - \{ \boldsymbol{\Gamma}_2(\Delta \mathbf{x}_{t-1} - \boldsymbol{\mu}) \} \\ &= - \left\{ \sum_{i=0}^{t-1} \boldsymbol{\Gamma}_{i+1} \mathbf{H}\boldsymbol{\nu}_{t-i} \right\} - \boldsymbol{\Gamma}_{t+1}(\Delta \mathbf{x}_0 - \boldsymbol{\mu}) \\ &\approx - \left\{ \sum_{i=0}^{t-1} \boldsymbol{\Gamma}_{i+1} \mathbf{H}\boldsymbol{\nu}_{t-i} \right\}. \end{aligned} \quad (6)$$

In practice, the initial condition (i.e., $\Delta \mathbf{x}_0 - \boldsymbol{\mu}$) is either treated as fixed when conducting conditional maximum likelihood estimation or set to zero according to its unconditional expectation. Note that, even if the initial condition were treated as fixed, $\boldsymbol{\Gamma}_{t+1}$ exponentially decays to zero as t increases for a stationary vector process. So the approximation in the last line should be highly accurate for all but the first few time periods given conditional maximum likelihood estimation. However, in our application, we use least squares estimation with backcast observations at the sample average. Thus, $\Delta \mathbf{x}_0 - \boldsymbol{\mu}$ is zero and there is no approximation, although we also note it would be straightforward to

include fixed values instead when accounting for the role of different sources of information. Meanwhuile, in terms of representing the changes in BN trends, $\Delta\tau_t$, as a function of the forecast errors, we can work off equation (4) to get the following:

$$\begin{aligned}\Delta\tau_t &= \{\mathbf{x}_t + \Gamma_1(\Delta\mathbf{x}_t - \boldsymbol{\mu})\} - \{\mathbf{x}_{t-1} + \Gamma_1(\Delta\mathbf{x}_{t-1} - \boldsymbol{\mu})\} \\ &= \boldsymbol{\mu} + \Gamma_0\mathbf{H}\boldsymbol{\nu}_t.\end{aligned}\tag{7}$$

To isolate a given BN cycle or trend as a function of the forecast errors, we can define a selection vector, \mathbf{e}_i , as a vector of zeros and one as its i^{th} element. In particular, to account for the contribution of the forecast error for the j^{th} variable to the BN cycle or trend of the s^{th} ordered variable, we pre-multiply \mathbf{c}_t or $\Delta\tau_t$ in equations (6) or (7) respectively by \mathbf{e}_s and post-multiply by \mathbf{e}_j .

2.1 Interpretation

Equations (6) and (7) provide a natural starting point for two different, but related, ways of looking at the BN trends and cycles in large systems. We briefly discuss each of these in turn.

Sources of Information Because equations (6) and (7) express the BN trends and cycles as functions of all historical forecast errors within the multivariate forecasting model, it is possible to relate these trends and cycles to different sources of information by relating them to forecast errors for each variable in the VAR. Thus, it can be easily determined what multivariate information is important for inferences about the trend and cycle of a given target variable such as log real GDP. If the trend and cycle of the target variable do not depend on the forecast errors for another variable, then there will be very little effect of including or excluding the other variable in the VAR when conducting a multivariate BN decomposition for the target variable. Importantly, though, this is not the basis for a variance decomposition of the trend and cycle because the forecast errors could be correlated across equations. However, as shown in our application, it is extremely useful for understanding how many and which variables should be included when estimating the trend and cycle of a target variable.

The Role of Structural Shocks As just noted, to the extent that the forecast errors in $\boldsymbol{\nu}_t$ are correlated, variance decompositions of trends and cycles based on the forecast errors are not possible. However, given the forecast errors and identification restrictions, causal inferences such as a variance decompositions based on structural shocks are possible. Recall that $\boldsymbol{\Sigma}_\nu$ is the covariance matrix associated with the forecasting model presented in equation (2). Let $\boldsymbol{\epsilon}_t$, where $\mathbf{A}\boldsymbol{\epsilon}_t = \boldsymbol{\nu}_t$, represent the vector of structural shocks with covariance matrix \mathbf{I} . The matrix \mathbf{A} maps the structural shocks to the reduced form

forecast errors by a series of identification restrictions that satisfy $\Sigma_{\nu} = \mathbf{A}\mathbf{A}'$ (see, for example, Kilian, 2013; Stock and Watson, 2016, for textbook treatments and details). In practice, \mathbf{A} could be a lower triangular Cholesky decomposition of Σ_{ν} , although numerous other schemes for identifying the structural shocks are also widely used and available. Then, once an identification scheme is specified, it is straightforward to plug $\mathbf{A}\epsilon_t$ in place of ν_t in equations (6) and (7) such that the BN trends and cycles become functions of the structural shocks. Because the structural shocks are orthogonal, this allow direct inference using variance decompositions or related measures about how much, say, monetary policy shocks, oil price shocks, or other structural shocks contribute to the trend and cycle of a given target variable.

As with any SVAR analysis, causal interpretations of the drivers of the trend and cycle directly rely on the plausibility of the identification of the structural shocks. Therefore, it is crucial to note that, while inferences about sources of information are invariant to identification restrictions, the causal analysis will necessarily be dependent on a particular identification scheme. At the same time, it should be noted that, again given orthogonality, it is possible to examine the causal effects of a given structural shock without needing to identify all of the other structural shocks in a system. This is convenient in large systems and is the approach we take in our application.

3 Bayesian Shrinkage and the Data

In this section, we first specify a VAR as the forecasting model in order to use Bayesian shrinkage to address problems with utilizing large information sets to calculate trend and cycle via a multivariate BN decomposition. Then we discuss the data we will use in a large BVAR to estimate the U.S. output gap.

Bayesian shrinkage methods for VARs are well developed (e.g., Litterman, 1986; Robertson and Tallman, 1999), including recently in a large BVAR context (see Banbura, Giannone, and Reichlin, 2010). Letting $\mathbf{Y}_t = (y_{1,t} \dots y_{n,t})'$ be a vector of n random variables, we model \mathbf{Y}_t as a VAR of lag order p :

$$\begin{aligned}
\mathbf{Y}_t &= \beta_1 \mathbf{Y}_{t-1} + \dots + \beta_p \mathbf{Y}_{t-p} + \mathbf{u}_t \\
&= \begin{bmatrix} \beta_1 & \beta_2 & \dots & \beta_p \end{bmatrix} \begin{bmatrix} \mathbf{Y}_{t-1} \\ \mathbf{Y}_{t-2} \\ \vdots \\ \mathbf{Y}_{t-p} \end{bmatrix} + \mathbf{u}_t \\
&= \begin{bmatrix} \beta_1^{11} & \dots & \beta_1^{1n} & \beta_2^{11} & \dots & \beta_2^{1n} & \dots & \dots & \beta_p^{1n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \beta_1^{n1} & \dots & \beta_1^{nn} & \beta_2^{n1} & \dots & \beta_2^{nn} & \dots & \dots & \beta_p^{nn} \end{bmatrix} \begin{bmatrix} \mathbf{Y}_{t-1} \\ \mathbf{Y}_{t-2} \\ \vdots \\ \mathbf{Y}_{t-p} \end{bmatrix} + \begin{bmatrix} u_{1,t} \\ \vdots \\ u_{n,t} \end{bmatrix}, \quad (8)
\end{aligned}$$

where $\mathbb{E}(\mathbf{u}'_t \mathbf{u}_t) = \Sigma_u$ and $\mathbb{E}(\mathbf{u}'_t \mathbf{u}_{t-i}) = \mathbf{0} \forall i > 0$. Assuming the variables in \mathbf{Y}_t are stationary and have been demeaned prior to estimation, the VAR in equation (8) can be cast into the companion form presented in equation (2) in a straightforward manner (see, for example, Hamilton, 1994) and the BN trends and cycles can be calculated and expressed as functions of historical forecast errors, as discussed in the previous section.

As motivated above, while Evans and Reichlin (1994) show that additional relevant multivariate information mechanically lowers the signal-to-noise ratio, their result leaves unanswered the question of how large the information set should be or how one can practically estimate a very large system in case many variables appear to be relevant. Meanwhile, it is a virtual certainty that parameter proliferation within a VAR framework will lead to overfitting as more variables with any estimated degree of Granger causality are added into the VAR system. The mechanical result in terms of a decrease in the signal-to-noise ratio was illustrated earlier in Figure 1 with the larger VAR system using least squares estimation producing the estimated U.S. output gap with the largest amplitude.

Here, we consider how to build on the insights of the BVAR literature in order to best utilize Bayesian shrinkage (e.g., see Litterman, 1986; Robertson and Tallman, 1999) to make estimation of BN trends and cycles with larger information sets tractable and not subject to a mechanical but possibly spurious decrease in the signal-to-noise ratio as more variables are included in the system. A basic insight from the large BVAR literature (see Banbura, Giannone, and Reichlin, 2010) is that one should apply more shrinkage as the system gets larger. We follow this approach, although we highlight that our proposed approach to shrinkage differs from Banbura, Giannone, and Reichlin's in three key ways.

First, while Banbura, Giannone, and Reichlin (2010) shrink variables towards random walk or noise processes depending on their persistence, we shrink the target variable (i.e., output growth when estimating the output gap) towards a process with a pre-specified signal-to-noise ratio, as discussed below, and the other variables to a noise process (if not differenced) or implicitly to a random walk process in levels (if differenced in the VAR).

Second, we optimize our shrinkage hyperparameter relative to an out-of-sample one-step-ahead forecast, rather than targeting the in-sample fit of a smaller system. Third, as mentioned above, we demean all of our data prior to estimation and thus we do not estimate a vector of constants as part of the model. Implicitly, this means we are using the sample average to estimate the unconditional mean of a given time series, which would be equivalent to placing a highly diffuse prior on the constants if they were included in estimation.

Because it is the most nonstandard feature of our procedure relative to the wider BVAR literature, our choice to shrink the target variable towards a process with a particular signal-to-noise ratio merits some discussion. Typical BVAR methods would shrink a variable like output towards to a random walk process. The underlying idea is that because a random walk provides a competitive forecast for many macroeconomic variables, shrinking towards a random walk balances overfitting, which worsens the forecasting performance of the model, with a more parsimonious and accurate forecasting model. If forecasting were the sole objective, shrinkage towards a random walk might be a good choice. However, forecasting is not our sole objective. In particular, while we hope for a competitive forecasting model, hence our choice to optimize based on out-of-sample forecasting performance, the ultimate aim of our analysis is to estimate a variable such as the output gap with a multivariate BN decomposition. By definition, if the forecasting model for the level of the target variable $\{y_t\}$ were a pure random walk with drift, the BN decomposition would imply no cycle (i.e., the time series would be equivalent to the trend, as its long horizon forecast minus drift is equal to the current level of the time series). In other words, the tighter we shrink towards the random walk, the smaller would be the BN cycle, by construction. This is not a desirable property as it could conflict with our use of more shrinkage as the number of variables increase. Put differently, applying shrinkage in the usual fashion in the BVAR literature would mechanically shrink the size of the BN cycle as the number of variables increases, with there being no cycle in the limit as the number of variables goes to infinity.

Our solution, then, is to shrink the target variable equation towards a pre-specified signal-to-noise ratio, which we label as δ . To interpret this signal-to-noise ratio, $\delta = 0.01x$ implies $x\%$ of the variance in the forecast error for $\{\Delta y_t\}$ is due to permanent shocks to $\{y_t\}$. Kamber, Morley, and Wong (2017) demonstrate that one can perform a univariate BN decomposition with a pre-specified δ because there is a direct mapping from δ to the AR coefficients in an AR(p) model. In particular, letting ρ be the sum of AR coefficients in an AR(p) regression of output growth, the mapping between the two is $\rho = 1 - 1/\sqrt{\delta}$. In Kamber, Morley, and Wong (2017), the estimation of the output gap from a univariate AR(p) model of output growth treats ρ as being fixed and so can be viewed as a dogmatic prior on the signal-to-noise ratio. Here, in the multivariate environment, we place a prior on δ , but it is not dogmatic because the multivariate information can move the posterior

away from the prior depending on how well the multivariate information helps to forecast $\{\Delta y_t\}$. A prior on δ amounts to placing a prior on the sum of the autoregressive coefficients in the target variable equation, which we label $\rho(\delta)$.

To be specific about our shrinkage priors, let β_l^{ij} denote the VAR coefficient of the l^{th} lag of variable j in the i^{th} equation of the VAR, which are elements of matrix of VAR slope coefficients β which we introduced in equation (8). Let the target variable, $\{\Delta y_t\}$, be the s^{th} variable in our BVAR, and $\rho(\bar{\delta})$ be the sum of the autoregressive coefficients in the target variable equation consistent with a pre-specified $\bar{\delta}$. We set the prior means and variances as follows:

$$\mathbb{E}[\beta_l^{ij}] = 0 \quad (9)$$

$$\mathbb{V}[\beta_l^{ij}] = \begin{cases} \frac{\lambda^2}{l^2}, & i = j \\ \frac{\lambda^2}{l^2} \frac{\sigma_i^2}{\sigma_j^2}, & \text{otherwise} \end{cases} \quad (10)$$

$$\mathbb{E}\left[\sum_{l=1}^p \beta_l^{ss}\right] = \rho(\bar{\delta}) \quad (11)$$

$$\mathbb{V}\left[\sum_{l=1}^p \beta_l^{ss}\right] = \chi^2. \quad (12)$$

Equations (9) and (10) imply all of the differenced variables in the VAR will shrink towards random walk processes in their levels or white noise processes for stationary level variables in the VAR. The variances σ_i^2 and σ_j^2 are set by taking the variance of the residuals on an AR(4) estimated using least squares on each of the individual time series as per the usual practice (e.g., Banbura, Giannone, and Reichlin, 2010; Koop, 2013). The factor $1/l^2$ shrinks coefficients at longer lags closer to zero, embedding the Minnesota prior structure that shorter lags are more important than longer lags in modeling the dynamics of a time series. Equations (11) and (12) implement the prior on the signal-to-noise ratio for the target variable. As discussed, we shrink the sum of the AR coefficients in the target variable equation to $\rho(\bar{\delta})$. In our application, we set $\bar{\delta} = 0.25$, consistent with Kamber, Morley, and Wong (2017).³ In order to shrink towards $\bar{\delta}$, rather than a random walk in the limit, we apply more shrinkage to the sum of the autoregressive coefficients than for the individual coefficients – i.e., we require $\chi \ll \lambda$. Thus, we set $\chi = \lambda/10$.

Our main estimation strategy centers on the shrinkage hyperparameter λ . Intuitively, λ serves as an information processing parameter that controls the degree of multivariate information entering into the estimation of the BN trend and cycle for the target variable. As discussed previously, our approach is to shrink λ closer to zero as the number of time

³Kamber, Morley, and Wong (2017) employ an explicit objective of maximizing the amplitude-to-noise ratio of the estimated output gap in their univariate model, which implies $\bar{\delta} = 0.25$ for U.S. real GDP data. See Kamber, Morley, and Wong (2017) for more details.

series in the model increases so as to first prevent overfitting, but also to control the degree of the multivariate information entering into the estimate of the BN trend and cycle for the target variable and so prevent a mechanical decrease of the signal-to-noise ratio. Prior work suggests a number of methods to choose the degree of shrinkage. One could optimize based on the marginal data density of the model (see Carriero, Clark, and Marcellino, 2015) or even estimate the shrinkage hyperparameter through the treatment of the prior in a hierarchical fashion (Giannone, Lenza, and Primiceri, 2015). Banbura, Giannone, and Reichlin (2010) suggest targeting the in-sample fit of a smaller 3 variable system with that of the larger system. All of these suggested procedures are based on the basic principle of increasing the degree of shrinkage as the risk of overfitting increases with the size of the system.

We adopt the same basic principle of applying shrinkage to mitigate overfitting, but deviate slightly in how we choose λ . In particular, we set λ to minimize the one-step-ahead out-of-sample root mean squared forecast error (RMSFE) of the target variable $\{\Delta y_t\}$. This approach is still based on the principle that overfitting compromises out-of-sample forecasting performance. Intuitively, this approach means that larger systems with a greater potential to overfit will be subject to a larger degree of shrinkage in order to achieve similar or better out-of-sample forecasting performance.⁴ However, we highlight that our choice to minimize the out-of-sample RMSFE of a target variable serves another purpose within the context of performing a trend-cycle decomposition with the BN decomposition. In particular, Nelson (2008) argues that an AR(1) is a sensible forecasting model for performing a BN decomposition of log real GDP because, while parsimonious, it produces comparatively good forecasts of output growth. Within our context, our priority is not so much to optimize the forecasting performance of the entire BVAR. However, we do view requiring the BVAR to produce competitive forecasts for output growth relative to an AR(1) as an important step to mitigate Nelson’s critique. Therefore, while we retain the principle of applying shrinkage to deal with overfitting, our main goal is to perform a BN decomposition of log real GDP and we view this deviation from the BVAR literature as being sensible within the context of trend-cycle decomposition of a given target variable.

To conduct Bayesian estimation of the model, we cast the VAR in equation (8) into a system of multivariate regressions:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}, \tag{13}$$

where $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_T]'$, $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_T]'$ with $X_t = [\mathbf{Y}'_{t-1}, \dots, \mathbf{Y}'_{t-p}]'$, and $\mathbf{u} =$

⁴For the out-of-sample forecast evaluation, we consider an expanding window for the estimation sample, starting with the initial 20 years of data (roughly one third of the 230 quarterly observations). However, we note that degree of shrinkage appears largely robust to an initial sample as short as 10 years.

$[\mathbf{u}_1, \dots, \mathbf{u}_T]'$. We then employ a Normal-Inverse Wishart prior, which has the form

$$\text{vec}(\boldsymbol{\beta})|\Sigma \sim N(\text{vec}(\boldsymbol{\beta}_0), \Sigma \otimes \Omega_0) \quad \text{and} \quad \Sigma \sim IW(S_0, \alpha_0), \quad (14)$$

where the prior parameters $\boldsymbol{\beta}_0$, Ω_0 , S_0 , and α_0 are set to be consistent with equations (9), (10), (11), and (12) and the expectation of Σ being $\text{diag}(\sigma_1^2, \dots, \sigma_n^2)$. The prior from (14) can then be implemented by choosing the following dummy observations in order to match the moments of the prior (see, e.g., Del Negro and Schorfheide, 2011; Woźniak, 2016):

$$Y_d = \begin{pmatrix} & & 0_{k,n} \\ & \text{diag}(\sigma_1 \dots \sigma_n) & \\ 0_{1,s-1} & \rho/\chi & 0_{1,n-s} \end{pmatrix}, X_d = \begin{pmatrix} J_p \otimes \text{diag}(\sigma_1 \dots \sigma_n)/\lambda & & \\ & 0_{n,k} & \\ 1_{1,n} \otimes \begin{bmatrix} 0_{1,s-1} & 1/\chi & 0_{1,n-s} \end{bmatrix} & & \end{pmatrix}, \quad (15)$$

where Y_d and X_d are dummy observations, $J_p = \text{diag}(1, \dots, p)$, $S_0 = (Y_d - X_d B_0)'(Y_d - X_d B_0)$, $B_0 = (X_d' X_d)^{-1} X_d' Y_d$, $\Omega_0 = (X_d' X_d)^{-1}$, and $\alpha_0 = T_d - k$, where T_d is the number of rows for both Y_d and X_d and $k = n \times p$.⁵ The first block of dummy observations places the prior on all of the individual VAR slope coefficients, the second block imposes the priors on the covariance matrix, and the third block implements the prior on the signal-to-noise ratio.

Augmenting the regression in equation (13) with the dummy observations gives the following:

$$\mathbf{Y}^* = \mathbf{X}^* \boldsymbol{\beta} + \mathbf{u}^*, \quad (16)$$

where $\mathbf{Y}^* = [\mathbf{Y}', \mathbf{Y}_d']'$, $\mathbf{X}^* = [\mathbf{X}', \mathbf{X}_d']'$ and $\mathbf{u}^* = [\mathbf{u}', \mathbf{u}_d']'$. Estimating the BVAR then simply amounts to conducting least squares regression of Y^* on X^* . Therefore, the posterior distribution has the form

$$\text{vec}(\boldsymbol{\beta})|\Sigma, \mathbf{Y} \sim N(\text{vec}(\tilde{\boldsymbol{\beta}}, \Sigma \otimes (\mathbf{X}^* \mathbf{X}^*)^{-1}) \quad (17)$$

$$\Sigma|\mathbf{Y} \sim IW(\tilde{\Sigma}, T_d + T - k + 2), \quad (18)$$

where $\tilde{\boldsymbol{\beta}} = (\mathbf{X}^* \mathbf{X}^*)^{-1} \mathbf{X}^* \mathbf{Y}^*$ and $\tilde{\Sigma} = (\mathbf{Y}^* - \mathbf{X}^* \tilde{\boldsymbol{\beta}})'(\mathbf{Y}^* - \mathbf{X}^* \tilde{\boldsymbol{\beta}})$.

Data For our empirical application, we consider data from 1959Q2 to 2016Q4. We use a 23 variable BVAR for our benchmark specification. The raw data are the oil price, real GDP, the Consumer Price Index (CPI), the unemployment rate, hourly earnings, the

⁵Note that because we demean all the variables prior to estimation, we do not include a constant in our BVAR. Thus the number of parameters in each equation is $n \times p$, not $n \times (p + 1)$.

fed funds rate, stock prices, the slope of the yield curve, the GDP deflator, employment, income, real personal consumption expenditure (PCE), industrial production, capacity utilization, housing starts, producer price index for all commodities, the PCE deflator, hours worked, nonfarm real output per hour, total reserves, non-borrowed reserves, real M1, and real M2. Much of the 23 variable system is informed by Banbura, Giannone, and Reichlin’s 20 variable system, which in turn is informed by an influential monetary VAR study by Christiano, Eichenbaum, and Evans (1999). In particular, Banbura, Giannone, and Reichlin (2010) suggest their medium-sized 20 variable BVAR system contains a sufficiently broad information set for macroeconomic forecasting purposes and so we believe it serves as a reasonable starting point for a model which should contain most, if not all, of the relevant information for estimating the output gap via the BN decomposition. However, we also consider a 138 variable BVAR, many of the additional variables of which are just subcomponents of the 23 variable system, to make it comparable to the large BVAR (see Banbura, Giannone, and Reichlin, 2010) or FAVAR (e.g., Bernanke, Boivin, and Elias, 2005) studies employed within the wider empirical literature. Finally, we consider an 8 variable system that contains a subset of variables of the baseline model, as discussed in the next section. All of the raw data are sourced from FRED and IFS. We leave definitions and details of the data to the appendix.

We take natural logarithms of series when appropriate and then take first differences of any series if either a unit root test cannot reject at a 5% level of significance or a simple t -test can reject an equal sample mean for the first half and latter half of the sample at a 10% level of significance.⁶ We transform the data in this way because we have assumed stationarity of the variables in our VAR in order to construct BN trends and cycles, as noted in Section 2.⁷ All series, once rendered stationary, are backcast with their sample average so as to keep the initial observations as part of the estimation sample. We conduct all of our estimation with 4 lags, as is standard for quarterly data.

4 U.S. Output Gap Estimates from BVARs

Figure 2 plots the estimated BN output gap based on our benchmark 23 variable BVAR. The estimated output gap moves with the NBER chronology of business cycle peaks and troughs reasonably well. In comparison to Figure 1, it appears that the use of Bayesian shrinkage has been successful in processing multivariate information, without necessarily producing an output gap that has a mechanically oversized amplitude.

⁶The simple sample splitting approach for testing for a break in mean is for convenience given the large number of series to process. However, it would certainly be possible to consider more formal tests for structural breaks in mean at unknown breakpoints.

⁷Preliminary analysis showed that, despite the shrinkage, incorporating very persistent series in the BVAR, including those with apparent shifts in the mean based on the split sample test, results in BN cycles that appear to drift up or down over rather than reverting to a mean of zero.

As discussed previously, equation (6) shows that the BN cycle can be represented as a function of the forecast errors for the various n variables, where the informational contribution of the i^{th} variable, S^i , to the BN cycle of output for output growth as the s^{th} variable can be expressed as

$$S^i = \mathbf{e}_s \mathbf{c}_t \mathbf{e}_i. \quad (19)$$

We calculate the standard deviations of the various n informational contributions to understand which sources of multivariate information are most important for the estimation of the output gap. These standard deviations are presented as unnormalized shares in Figure 3. It can be seen from the figure that the two most important sources of multivariate information for the output gap in our 23 variable system are the unemployment rate and CPI inflation. We also observe that five other variables, personal consumption expenditure, housing starts, the federal funds rate, real M1, and stock prices are also somewhat important.

To confirm the importance of these particular variables as sources of information, we consider a more parsimonious 8 variable BVAR, which includes the seven most important variables from our 23 variable benchmark BVAR, along with output growth itself, in order to estimate the output gap. We also consider a large 138 variable BVAR to see if other macroeconomic variables might provide useful information.

The top panel of Figure 4 plots the output gap estimates for three different BVARs containing 8, 23, and 138 variables, respectively. The output gap estimates for these three models are largely similar. The reason for the similarity is that Bayesian shrinkage has suppressed all but the most important sources of multivariate information for output gap estimation. This is even true for the 8 variable BVAR that produces a smaller amplitude output gap than the estimate based on a 7 variable VAR estimated by least squares in Figure 1. Meanwhile, the bottom panel of Figure 4 compares the output gap estimate from our benchmark BVAR with the estimate for a 22 variable BVAR that omits the unemployment rate, which was the most important source of information according to the results in Figure 3. The output gap estimate now differs substantially from the benchmark estimate, supporting an approach that determines the main sources of information for a given BN cycle, rather than, say, using a factor model approach to process the information in a large dataset. In particular, it is the inclusion of key variables that matters, not necessarily the general size of the information set.

Figure 5 plots the out-of-sample root mean square error of output growth as a function of our shrinkage hyperparameter, λ , for the three different cases in terms of the number of variables in the BVAR. As expected, our procedure produces more shrinkage as the system gets larger. The shrinkage hyperparameter, λ , minimizes the one-step-ahead RMSFE at 0.49, 0.11, and 0.04, for our 8, 23, and 138 variable BVAR models, respectively. Therefore,

while our procedure of optimizing the degree of shrinkage differs from Banbura, Giannone, and Reichlin (2010), targeting out-of-sample fit achieves the same goal of tightening the degree of shrinkage as we increase the size of the BVAR, suggesting that this is a viable alternative option for determining the degree of shrinkage in large BVARs.

We also note that the out-of-sample fit improves as we move away from $\lambda = 0$. Recall that the prior implies output growth is shrunk towards an AR(4) with a pre-specified sum of coefficients that coincides with a fixed signal-to-noise ratio, $\bar{\delta}$. As shown in Kamber, Morley, and Wong (2017), maximum likelihood estimates of δ are far in excess of 1, so it is not surprising that the fit when $\delta = 0.25$, as implied by the prior, can be improved upon. However, we also note that, due to likely overfitting, the out-of-sample fit deteriorates much quicker for the larger models as we loosen the degree of shrinkage. Therefore, our objective of targeting the out-of-sample forecasting performance balances out the improvement of in-sample fit from being less restrictive on the signal-to-noise ratio and adding multivariate information against possible overfitting.

Interestingly, the minimum achievable out-of-sample RMSE does fall marginally as we increase the size of the information set. This suggests that additional variables do contain more information to improve out-of-sample forecasts, which one could exploit with judicious choices about the degree of shrinkage. For comparison, we also plot a horizontal line to indicate the out-of-sample forecasting performance of an AR(1) model. As can be seen, our BVARs produce very competitive forecasts relative to an AR(1) model, even beating the AR(1) model when we increase the information set to 138 variables. Our benchmark BVAR does only marginally worse than an AR(1) model in terms of its out-of-sample forecasting performance. This addresses a key critique by Nelson (2008), who argues that the BN decomposition based on an AR(1) model is relevant because an AR(1) matches the autocovariance structure of the output growth data. We show that with shrinkage, it is possible to specify other models that do as well as AR(1) model in out-of-sample forecasts, supporting our approach as possible way to conduct multivariate BN decompositions without necessarily generating spurious cycles.

We draw three general conclusions from the preceding analysis. First, Evans and Reichlin (1994) show that multivariate information will lead the BN decomposition to produce larger cycles relative to univariate BN decompositions. However, given least squares estimation of a VAR, the addition of multivariate information mechanically lowers the signal-to-noise ratio, implying a larger amplitude cycle. That the cycle is a mechanical function of the size of the information set is a relevant obstacle to using such multivariate BN decompositions for practical analysis. However, we show that Bayesian shrinkage is helpful in filtering information, so the general issues that Evans and Reichlin (1994) discuss are no longer impediments to using multivariate BN decompositions for practical analysis. As long as a researcher specifies a set of relevant multivariate information, the output gap estimate should be robust to varying the exact size of the information

set. Second, we show that it is crucial to determine what is important information. We do this by examining the informational contributions of different forecast errors in driving the BN cycle. Not surprisingly, discarding or omitting important information has a noticeable impact on the output gap estimate. Third, we show shrinkage is an important tool to be able to consider very large datasets. That we have successfully estimated an output gap using 138 variables suggest that we may have a feasible solution to a scenario not dissimilar to that observed in the policy environment. In particular, policy institutions such as central banks typically monitor and observe large volumes of data, which they use to inform their view of the degree of economic slack. We present a method to directly incorporating a vast amount of information directly into formal trend-cycle decomposition aimed at estimating the output gap.

In sum, we have demonstrated that it is possible to utilize Evans and Reichlin's insights about using multivariate information to estimate the output gap, but without the disadvantage of a mechanical lowering of the signal-to-noise ratio or overfitting. We have also showed that it is possible to produce a robust estimate of the U.S. output gap in particular using as few as 8 variables, with the unemployment rate and CPI inflation being particularly important sources of multivariate information in this case.

Why was the estimated output gap less deep during the Great Recession than in the early 1980s?

The estimates in Figures 2 and 4 suggest that the output gap was more negative in the early 1980s than during the Great Recession. Given that output fell by more with the Great Recession, this directly suggests that the decline in trend output during the Great Recession must have been more substantial than in the early 1980s. This result is consistent with the assessment by the President of the Federal Reserve Bank of St. Louis, James Bullard, who suggested that the Great Recession resulted in large permanent decreases in output that cannot be expected to be reversed (see Bullard, 2012). A lower level of trend output also means that the estimated output gap was back to zero by 2014Q1 and was as positive as 1% by early 2016.

Because we consider multivariate information, we can also study which source of multivariate information help explain trend growth. Our analysis suggest that the information from consumption during and after the Great Recession played the main role in explaining the lower estimate of trend output. In particular, the top panel of Figure 6 presents the change in trend output (cumulated because of the stochastic trend) that is accounted for by the forecast errors of consumption, comparing two periods, 1980Q1-1983Q1 and 2008Q1-2013Q4. The former period encompasses the twin recessions in the early 1980s and serves as a contrast as we try to understand what happened during the Great Recession. Recall that we estimate a large negative output gap in the 1981-1983 recession,

but less so in the Great Recession. Evidently, the forecast errors of consumption only lower the estimate of trend growth during the twin recessions in the early 1980s by less than -0.5%. By contrast, the forecast errors of consumption lower the estimate of trend output by about 2% during the period between 2008Q1 to 2013Q4. Therefore, because the estimate of trend output was substantially lower during and after the Great Recession, this corresponds to a less deep output gap as our BVAR suggests much of the decline in real GDP with the Great Recession was permanent.

Looking at the profile of consumption growth during the two time periods in the bottom panel, one striking observation is that consumption stayed sluggish throughout the Great Recession. Compare this with the early 1980s, where there was a sharp fall in the consumption growth rate in 1980Q2, but which was immediately reversed in the following quarter. Meanwhile, during the 1982-83 recession, consumption growth did not collapse, but was comparatively strong throughout most of the recession. For the Great Recession, consumption growth stayed largely negative during the whole of the NBER dated recession and remained weak afterwards. In fact, consumption growth in the period 2011-2013 was weaker than in the period from 1982-1983 when NBER still considered the U.S. economy to be in a recession.

We can conjecture why consumption stayed sluggish in the aftermath of the Great Recession. To the extent the permanent income hypothesis holds and households perceived slower growth the permanent income, consumption growth should be sluggish. Our estimates suggesting that much of the fall of output was permanent in the Great Recession and was related to the behaviour of consumption is consistent with this conjecture. Meanwhile, we do not wish to give the impression that the Great Recession only produced a small negative output gap. Looking across the whole sample period, our estimated output gap was still relatively deep, but this is despite our concurrent estimates of a large fall in trend output.

What does multivariate information really add?

While it is apparent from the previous analysis that the addition of multivariate information substantially impacts our estimates, we now explore exactly what the multivariate information adds beyond a univariate approach in the vein of Kamber, Morley, and Wong (2017). A natural way to do this is to plot the univariate estimate from such an approach against our multivariate estimate. In particular, we can compare the benchmark estimate of the output gap against the “BN Filter” estimate based on Kamber, Morley, and Wong (2017).⁸ For completeness, we can also compare our multivariate estimate against the prior mean, which can be captured by estimating the output gap with a very tight prior

⁸We conduct the Kamber, Morley, and Wong (2017) procedure of maximizing the amplitude-to-noise ratio for an AR(4) model to make it directly comparable to our BVAR. We find $\bar{\delta} = 0.25$, consistent with the finding in Kamber, Morley, and Wong (2017) for an AR(12) model.

by setting λ effectively to zero.⁹ Recall that this recovers the prior on the signal-to-noise ratio of $\bar{\delta} = 0.25$. These comparisons are plotted in the top panels of Figure 7 and two results are immediately noticeable. First, the estimated output gap that incorporates multivariate information tends to be of larger amplitude relative to the univariate “BN Filter” approach or the prior. Second, the estimated output gap incorporating multivariate information appears to be capturing booms in the sense that there is always a large positive output gap near the peak of NBER expansions. The BN Filter or the estimated output gap based on the prior mean only appear to be capturing troughs, with the estimated output gap often falling before the start of NBER recessions. A suggestive interpretation is that multivariate information could be more helpful in capturing booms, but less so for capturing troughs.

To investigate this possibility in greater detail, we track the shrinkage hyperparameter, λ , in pseudo-real-time. Recall λ determines how much multivariate information is allowed into the detrending problem and is chosen based on minimizing the one-step-ahead out-of-sample RMSFE. We can interpret a larger λ as implying a greater role for multivariate information. We thus track λ for our 23 variable benchmark and 138 variable models by increasing the sample, one observation at a time.¹⁰ The pseudo-real-time values of λ are plotted in the bottom panels of Figure 7. Looking at large movements in λ , there is some evidence that multivariate information appears to be overfitting in recessions, but appears to be more useful in expansions. For example, the value of the shrinkage hyperparameter collapses at the trough of the Great Recession, but appears to be rising steadily in the current expansion.

We also know from observing the raw data that there is some negative skewness in the sense of abnormally large negative observations for output growth in recessions, but not so many abnormally large positive growth rates in expansions. Instead, expansions tend to be dominated by realizations of output growth closer to the mean. One possible interpretation is that, because we get a large negative growth rate in recession, one may not require multivariate information in recessions to understand the permanent and transitory effects of recessions. By contrast, because output growth tends to be closer to its average rate in expansions, multivariate information can help provide better identification of permanent and transitory movements in output.

We also note that the apparent asymmetry in output growth has given rise to the idea in the literature that output growth may be driven by nonlinear dynamics and/or the output gap is asymmetric (e.g., see Hamilton, 1989; Morley and Piger, 2012). Despite

⁹We set $\lambda = 1e - 15$. This produces slightly different λ results, although a similar shape, compared to the BN filter. The difference is because of the prior on lagged differences is not exactly the same for the two approaches.

¹⁰Note that whatever the size of the pseudo-real-time sample, we maintain an initial window size of 20 years to estimate the coefficients—i.e., the out-of-sample root mean square error is always only calculated from the 81st observation onwards.

the multivariate information being informative for estimating a sizeable positive output gap, our output gap still appears to be somewhat asymmetric. Because our BVAR model is linear and our estimated BN cycle is a function of the historical forecast errors, this asymmetry must, by construction, reflect asymmetry in the estimated forecast errors, which we argue should be quite prevalent given our observation about output growth asymmetry.

These findings may be related to Harding and Pagan (2010), who document that linear univariate models struggle to identify business cycle peaks. We show that by incorporating multivariate information, our BVAR models are better able to estimate significantly positive output gaps, and thus may provide some support for the idea that multivariate information is more useful in identifying the output gap in expansions than in recessions.

5 Causal Determinants of the U.S. Output Gap and Trend Growth

So far, our analysis has largely abstracted from causality given that all we have been doing is associate movements in the estimated output gap with information embedded in various forecast errors. In this section, we turn to conducting more structural analysis in order to examine possible causal determinants of those movements. In particular, we use SVAR analysis to identify two widely-considered structural shocks: a monetary policy shock and an oil price shock. The monetary policy shock is identified by ordering the federal funds rate after “slow moving” variables, but before “fast moving” ones in a Cholesky decomposition. This identification strategy is similar in spirit to work by, *inter alia*, Christiano, Eichenbaum, and Evans (1999) and Bernanke, Boivin, and Elias (2005), where the idea is that financial market variables are in the fast moving block because they can respond contemporaneously to monetary policy shocks, while slow moving variables take at least a quarter to respond. The “fast moving” variables in our benchmark 23 variable specification are real M1 and M2, stock prices, non-borrowed reserves, total reserves, and the slope of the yield curve. The oil price shock is identified by drawing from Kilian and Vega (2011), who show that oil prices do not appear responsive to macroeconomic news and thus can be taken to be pre-determined. This in essence orders the oil price first in a Cholesky decomposition and also has precedence in the wider SVAR literature studying oil price shocks (e.g, see Edelstein and Kilian, 2009; Wong, 2015). Our system is partially identified, in the sense that we only identify two out of 23 potential structural shocks in our benchmark system and we do not attempt to disentangle any of the remaining 21 unidentified shocks.

We first examine how much a given structural shock has driven the historical BN trend

and cycle by performing a variance decomposition. To set up a variance decomposition of the BN cycle, we first note that $\mathbb{E}\boldsymbol{\nu}_t = \mathbf{0}$. Working off Equation (6), it can be verified that the difference between the actual h -step-ahead BN cycle and the conditional expectation of the BN cycle at time $t - 1$ is

$$\mathbf{c}_{t+h} - \mathbb{E}_{t-1}\mathbf{c}_{t+h} = \sum_{i=0}^h \boldsymbol{\Gamma}_{i+1} \mathbf{H} \boldsymbol{\nu}_{t+h-i} \quad (20)$$

$$= \sum_{i=0}^h \boldsymbol{\Gamma}_{i+1} \mathbf{H} \mathbf{A} \boldsymbol{\epsilon}_{t+h-i}, \quad (21)$$

where the second equality follows from the identification associated with the structural shocks from the SVAR. Because $\mathbb{E}(\boldsymbol{\nu}'_t \boldsymbol{\nu}_{t-i}) = 0, i > 0$, the total variance can therefore be written as

$$\text{Var}(\mathbf{c}_{t+h} - \mathbb{E}_{t-1}\mathbf{c}_{t+h}) = \sum_{i=0}^h \boldsymbol{\Gamma}_{i+1} \mathbf{H} \boldsymbol{\Sigma}_{\nu} \mathbf{H}' \boldsymbol{\Gamma}_{i+1}' \quad (22)$$

It follows, then, that a variance decomposition of the h -step-ahead variation in the BN cycle can be calculated using equations (21) and (22):

$$FEVD_{j,h}^c = \frac{\left[\sum_{i=0}^h \mathbf{e}_s \boldsymbol{\Gamma}_{i+1} \mathbf{H} \mathbf{A} \mathbf{e}'_j \right]^2}{\mathbf{e}_s \left[\sum_{i=0}^h \boldsymbol{\Gamma}_{i+1} \mathbf{H} \boldsymbol{\Sigma}_{\nu} \mathbf{H}' \boldsymbol{\Gamma}_{i+1}' \right] \mathbf{e}'_s}, \quad (23)$$

where $FEVD_{j,h}^c$ is the h -step-ahead share of the variance of the output gap due to the j^{th} structural shock, where output growth is once again the s^{th} variable in the system. Similarly, to perform a variance decomposition of trend growth, it is straightforward to verify from equation (7) that the variance of the change in trend can be written as

$$\text{Var}(\Delta \boldsymbol{\tau}_t - \mathbb{E}_{t-1} \Delta \boldsymbol{\tau}_t) = \boldsymbol{\Gamma}_0 \mathbf{H} \boldsymbol{\Sigma}_{\nu} \mathbf{H}' \boldsymbol{\Gamma}'_0 \quad (24)$$

and the share of the variance can be similarly decomposed as

$$FEVD_j^{\tau} = \frac{\left[\sum_{i=0}^h \mathbf{e}_s \boldsymbol{\Gamma}_0 \mathbf{H} \mathbf{A} \mathbf{e}'_j \right]^2}{\mathbf{e}_s \left[\boldsymbol{\Gamma}_0 \mathbf{H} \boldsymbol{\Sigma}_{\nu} \mathbf{H}' \boldsymbol{\Gamma}'_0 \right] \mathbf{e}'_s}. \quad (25)$$

Note that due to the random walk trend, the variance of trend is unbounded as the time horizon goes to infinity. Consequently, a decomposition of the contemporaneous variance of the change in the trend is sufficient to provide insight into how much of the variation of trend growth is due to the various identified structural shocks given the random walk trend implies a one-off permanent shift in the random walk trend.

Figure 8 presents a variance decomposition of the output gap and trend growth. For

the output gap, we present the share of monetary policy shocks and oil price shocks at horizons $h = 0$, $h = 4$, and $h = \infty$. Neither the monetary policy shock nor the oil price shock explain more than 10% of the variance of the output gap at any horizon. While the monetary policy shock explains about 7% of the variance of the output gap contemporaneously, its share quickly dissipates and it only explains about 4% of the unconditional variance. Therefore, it appears that the role of the monetary policy shock in driving the output gap is limited and relatively short lived. This finding is consistent with the wider SVAR literature, which often reports that monetary policy shocks explain only a small part of real economic activity. The oil price shock explains a somewhat larger share at about 10% of the variance of the output gap over all horizons. Meanwhile, consistent with traditional theories of growth that assume technology shocks are the main determinant of the long-run level of output, neither of these shocks explains much of trend growth, with shares of about 5% for the oil price shock and less than 4% for the monetary policy shock. Notably, the latter result is reflective of the idea of money neutrality, which suggests monetary policy should not have any permanent effects on the level of output.

Although variance decompositions are useful to gain an overall perspective on the importance of shocks, we can also calculate historical contribution of shocks to the output gap to understand specific episodes. This analysis is displayed in Figure 9. For historical monetary policy shocks, we can observe that they explain a large share of the positive output gap before the 1980 recession, consistent with anecdotal evidence that the Fed may have been overly heating the economy. Although we can see that monetary policy shocks contributed to some of the negative output gap in the early 1980s, consistent with the Volcker disinflation, the overall output gap in the early 1980s was estimated to be large and negative, with monetary policy shocks only contributing to part of the negative gap rather than being the dominant cause. Meanwhile, a recent interpretation of the events leading to the Great Recession suggest that Chairman Greenspan was perhaps running the economy too hot before 2008 (e.g., see Taylor, 2012). Our historical decomposition does not support this story. We find that, while monetary policy shocks did contribute modestly to a rising positive output gap in the early 2000s, this contribution largely turned negative by 2005, while the estimated output gap continued to increase up until the advent of the Great Recession.

We find that with historical oil price shocks, they tend to contribute positively to the output gap when oil prices are low and contribute negatively when oil prices are high. This can be seen from the negative contribution of oil price shocks throughout the 2000s and the positive contribution in the late 1990s. We also observe a positive contribution turning negative around 1990, consistent with the timing when the First Gulf War caused oil prices to rise from a low starting level. Furthermore, oil price shocks contributed negatively to the output gap around 1979 and 1980, consistent with the timing of the Iranian hostage crisis and the start of the Iraq-Iran War.

Overall, we find that the contribution of the historical monetary policy and oil price shocks line up with many well-understood events. This analysis thus provides support for using the tools developed in this paper as way to understand and interpret historical estimates of trend and cycle.

6 Robustness to Specifying the Size of the Model

We have provided a method to utilize Bayesian shrinkage and the BN decomposition in order to estimate the output gap with potentially large information sets. In this section, we conduct Monte Carlo analysis to better understand the robustness of our procedure to correctly specifying the size of the model. We consider three data generating process (DGPs), which are based on model estimates when applying our shrinkage procedure to U.S. output growth data (i.e., only one variable) and the variables in our 8 and 23 variable systems, respectively. We generate 1000 artificial datasets for each of the three DGPs of length $T = 230$, consistent with the length of the sample in our empirical analysis. The true output gap is defined as the BN decomposition of output based on the VAR coefficients which parametrize a given DGP. For each artificial dataset, we estimate VARs using either least squares or the shrinkage procedure introduced in this paper, considering respectively a 1 variable, 8 variable, and 23 variable system. In each case, we construct the output gap via the BN decomposition based on the estimated parameters. For each draw, we calculate the root mean squared error (RMSE) of the estimated output gap relative to the true cycle. For each DGP, there is therefore one correct specification of the number and set of variables and two other specifications that do not coincide with the number or set of variables in the underlying DGP.¹¹

Table 1 reports the results for our Monte Carlo analysis. The top panel (a) presents the mean and standard deviation of the RMSEs across draws in each case of estimated model size for a given DGP. The bottom panel (b) presents the proportion of Monte Carlo draws where our BVAR approach has a lower RMSE compared to estimates based on least squares.

In cases where the size of the model is correctly specified (i.e., the diagonals in both panels (a) and (b)), we find that, apart from the univariate case when shrinkage matters the least, our BVAR approach is more accurate than least squares. Even in a moderately-sized 8 variable system, our shrinkage procedure produces a more accurate estimate of the output gap over 60% of the time, with both a lower mean and lower standard deviation of the RMSE. The benefits of shrinkage are particularly noticeable when we consider a

¹¹When there are additional variables which do not feature in the underlying DGP, we use their historical sample values as the additional data used in our estimation. For example, for our 8 variable DGP, we use the 8 simulated series for the variables in the DGP and the 15 additional historical data series for the variables that are not in the DGP to estimate a 23 variable model.

larger 23 variable DGP. In such a situation, overfitting would likely be a concern with least squares even if one knew the 23 variables of the underlying DGP with certainty. For the 23 variable system, our shrinkage procedure produces a more accurate estimate of the output gap 99% of the time.

Notably though, our BVAR approach is not only useful when the size of the model is correctly specified, but also when it is not. In particular, we can investigate the effects of misspecifying the size of the model by looking at the off-diagonals in panels (a) and (b). When the model size is misspecified, our Monte Carlo results suggest that shrinkage helps keep the estimate of the output gap comparatively accurate. For example, we find that under misspecification, our BVAR approach always produces a more accurate output gap relative to least squares more than 75% of the time and as much as 99% of the time in multiple cases. Notably, this proportion becomes larger as either the DGP or the estimated model size becomes larger.

7 Conclusion

In this paper, we have shown how to apply the Beveridge-Nelson decomposition to obtain estimates of trend and cycle using very large multivariate models with Bayesian shrinkage. We have also shown how to account for and interpret the various sources of multivariate information. In our empirical application, we present estimates of the U.S. output gap, with information sets ranging from eight to well over a hundred variables. We find the U.S. unemployment rate and CPI inflation, together with, to a lesser extent, housing starts, real M1, real consumption, the federal funds rate, and stock prices, are important sources of information in estimating the U.S. output gap. We also show how to conduct causal analysis given identified structural shocks and find that neither monetary policy shocks nor oil price shocks are responsible for the bulk of movements in the U.S. output gap or trend growth.

Future Extensions We view two advantages of the tools developed in this paper that motivate future extensions and applications. The first advantage of our approach is that casting the detrending problem within a linear regression framework allows us to utilize complicated datasets for inferring sources of movements in trend and cycle. In particular, many time series problems can naturally be cast into the BVAR models considered in this paper. For example, policy institutions often construct an output gap by monitoring very broad set of data of differing frequencies. One could, with some extra work, cast the problem into a BVAR with mixed frequencies, and thus allow information from monthly data to directly enter the problem of nowcasting the output gap, even though real GDP is often only available at a quarterly frequency. Another potential extension is joint detrending. Although we only target a single variable to estimate the output gap in

our analysis, one could cast the problem in the form where one could jointly detrend multiple variables to obtain estimates about, say, the natural rate of unemployment, trend inflation, the natural rate of interest, and potential output within a unified and consistent framework. A second advantage of our approach is the ability to interpret trend and cycle by appealing further to tools from the well-developed SVAR literature. This allows us to be able to meaningfully discuss shocks driving the trend and cycle and to attribute causality. The standard frameworks of trend-cycle decomposition using time series methods like unobserved components models can struggle to attribute causality, in addition to being more difficult to estimate than the models we propose, especially given large information sets. For example, Kamber and Wong (2017) employ the methods introduced in this paper to estimate the role of foreign shocks in driving trend inflation and the inflation gap for a number of open economies. One could similarly use the tools we introduce in this paper to answer current relevant policy questions such as what drives low neutral interest rates or what drives financial cycles. However, we leave this analysis to future research.

Data Appendix

IFS in the mnemonic field refers to the time series being sourced from the International Financial Statistics. Otherwise, the time series is sourced from the Federal Reserve Economic Data (FRED), and the mnemonic field refers to the FRED mnemonic.

Series	Mnemonic
U.S.: Commodity Price: W Texas Interm Spot Price (US\$/Barrel)	IFS
Real Gross Domestic Product, 3 Decimal	GDPC96
Real Personal Consumption Expenditures	PCECC96
Personal Consumption Expenditures: Durable Goods	PCDGx
Personal Consumption Expenditures: Services	PCESVx
Personal Consumption Expenditures: Nondurable Goods	PCNDx
Real Gross Private Domestic Investment, 3 decimal	GPDIC96
Fixed Private Investment	FPIx
Gross Private Domestic Investment: Fixed Investment: Nonresidential: Equipment	Y033RC1Q027SBEAx
Private Nonresidential Fixed Investment	PNFIx
Private Residential Fixed Investment	PRFIx
Shares of gross domestic product: Gross private domestic investment: Change in private inventories	A014RE1Q156NBEA
Real Government Consumption Expenditures and Gross Investment	GCEC96
Real Government Consumption Expenditures and Gross Investment: Federal	A823RL1Q225SBEA
Federal Government Current Receipts	FGRECPTx
State and Local Consumption Expenditures & Gross Investment	SLCEx
Real Exports of Goods and Services, 3 Decimal	EXPGSC96
Real Imports of Goods and Services, 3 Decimal	IMPGSC96
Real Disposable Personal Income	DPIC96
Nonfarm Business Sector: Real Output	OUTNFB
Business Sector: Real Output	OUTBS
Industrial Production Index	INDPRO
Industrial Production: Final Products (Market Group)	IPFINAL
Industrial Production: Consumer Goods	IPCONGD
Industrial Production: Materials	IPMAT

Industrial Production: Durable Materials	IPDMAT
Industrial Production: Nondurable Materials	IPNMAT
Industrial Production: Durable Consumer Goods	IPDCONGD
Industrial Production: Durable Goods: Automotive products	IPB51110SQ
Industrial Production: Nondurable Consumer Goods	IPNCONGD
Industrial Production: Business Equipment	IPBUSEQ
Industrial Production: Consumer energy products	IPB51220SQ
Capacity Utilization: Manufacturing (SIC)	CUMFNS
All Employees: Total Nonfarm Payrolls	PAYEMS
All Employees: Total Private Industries	USPRIV
Civilian Employment Level	CE16OV
Civilian Labor Force Participation Rate	CIVPART
Civilian Unemployment Rate	UNRATE
Unemployment Rate: 16 to 19 years	LNS14000012
Unemployment Rate: 20 years and over, Men	LNS14000025
Unemployment Rate: 20 years and over, Women	LNS14000026
Number of Civilians Unemployed for Less Than 5 Weeks	UEMPLT5
Number of Civilians Unemployed for 5 to 14 Weeks	UEMP5TO14
Number of Civilians Unemployed for 15 to 26 Weeks	UEMP15T26
Number of Civilians Unemployed for 27 Weeks and Over	UEMP27OV
Employment Level: Part-Time for Economic Reasons, All Industries	LNS12032194
Business Sector: Hours of All Persons	HOABS
Nonfarm Business Sector: Hours of All Persons	HOANBS
Average Weekly Hours of Production and Nonsupervisory Employees: Manufacturing	AWHMAN
Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Manufacturing	AWOTMAN
Housing Starts: Total: New Privately Owned Housing Units Started	HOUST
Privately Owned Housing Starts: 5-Unit Structures or More	HOUST5F
Housing Starts in Midwest Census Region	HOUSTMW
Housing Starts in Northeast Census Region	HOUSTNE
Housing Starts in South Census Region	HOUSTS

Housing Starts in West Census Region	HOUSTW
Personal Consumption Expenditures: Chain-type Price Index	PCECTPI
Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index)	PCEPILFE
Gross Domestic Product: Chain-type Price Index	GDPCTPI
Gross Private Domestic Investment: Chain-type Price Index	GPDICTPI
Business Sector: Implicit Price Deflator	IPDBS
Personal consumption expenditures: Goods (chain-type price index)	DGDSRG3Q086SBEA
Personal consumption expenditures: Services (chain-type price index)	DSERRG3Q086SBEA
Consumer Price Index for All Urban Consumers: All Items	CPIAUCSL
Producer Price Index for All Commodities	PPIACO
Producer Price Index by Commodity Industrial Commodities	PPIDC
Producer Price Index by Commodity for Fuels and Related Products and Power: Crude Petroleum (Domestic Production)	WPU0561
Average Hourly Earnings of Production and Nonsupervisory Employees: Construction	CES2000000008x
Average Hourly Earnings of Production and Nonsupervisory Employees: Manufacturing	CES3000000008x
Nonfarm Business Sector: Real Compensation Per Hour	COMPRNFB
Business Sector: Real Compensation Per Hour	RCPHBS
Nonfarm Business Sector: Real Output Per Hour of All Persons	OPHNFB
Business Sector: Real Output Per Hour of All Persons	OPHPBS
Business Sector: Unit Labor Cost	ULCBS
Nonfarm Business Sector: Unit Labor Cost	ULCNFB
Nonfarm Business Sector: Unit Nonlabor Payments	UNLPNBS
Producer Price Index by Commodity Metals and metal products: Primary nonferrous metals	PPICMM
Consumer Price Index for All Urban Consumers: Apparel	CPIAPPSL
Consumer Price Index for All Urban Consumers: Transportation	CPITRNSL
Consumer Price Index for All Urban Consumers: Medical Care	CPIMEDSL
Consumer Price Index for All Urban Consumers: Commodities	CUSR0000SAC
Consumer Price Index for All Urban Consumers: Durables	CUUR0000SAD
Consumer Price Index for All Urban Consumers: Services	CUSR0000SAS
Consumer Price Index for All Urban Consumers: All Items Less Food	CPIULFSL
Consumer Price Index for All Urban Consumers: All items less shelter	CUUR0000SA0L2

Consumer Price Index for All Urban Consumers: All items less medical care	CUSR0000SA0L5
Average Hourly Earnings of Production and Nonsupervisory Employees: Goods-Producing	CES0600000008
Consumer Motor Vehicle Loans Owned by Finance Companies, Outstanding	DTCOLNVHFNM
Effective Federal Funds Rate	FEDFUNDS
3-Month Treasury Bill: Secondary Market Rate	TB3MS
6-Month Treasury Bill: Secondary Market Rate	TB6MS
1-Year Treasury Constant Maturity Rate	GS1
10-Year Treasury Constant Maturity Rate	GS10
Moody's Seasoned Aaa Corporate Bond Yield	AAA
Moody's Seasoned Baa Corporate Bond Yield	BAA
Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity	BAA10YM
6-Month Treasury Bill Minus Federal Funds Rate	TB6SMFFM
10-Year Treasury Constant Maturity Minus Federal Funds Rate	T10YFFM
Real St. Louis Adjusted Monetary Base	AMBSLREALx
Real M1 Money Stock	M1REALx
Real M2 Money Stock	M2REALx
Real MZM Money Stock	MZMREALx
Commercial and Industrial Loans, All Commercial Banks	BUSLOANSx
Consumer Loans at All Commercial Banks	CONSUMERx
Total Nonrevolving Credit Owned and Securitized, Outstanding	NONREVSLx
Real Estate Loans, All Commercial Banks	REALLNx
Total Consumer Credit Owned and Securitized, Outstanding	TOTALSLx
Households and Nonprofit Organizations; Total Assets, Level	TABSHNOx
Households and Nonprofit Organizations; Total Liabilities, Level	TLBSHNOx
Households and Nonprofit Organizations; Credit Market Instruments; Liability, Level	CMDEBT
Households and Nonprofit Organizations; Net Worth, Level	TNWBSHNOx
Households and Nonprofit Organizations; Total Financial Assets, Level	TFAABSHNO
Households and nonprofit organizations; real estate at market value, Level	HNOREMQ027Sx
Households and Nonprofit Organizations; Total Financial Assets, Level	TFAABSHNOx
Shares of gross domestic product: Exports of goods and services	B020RE1Q156NBEA

Shares of gross domestic product: Imports of goods and services	B021RE1Q156NBEA
Industrial Production: Manufacturing (SIC)	IPMANSICS
Industrial Production: Residential utilities	IPB51222S
Industrial Production: Fuels	IPFUELS
Average (Mean) Duration of Unemployment	UEMPMEAN
Average Weekly Hours of Production and Nonsupervisory Employees: Goods-Producing	CES060000007
Total Reserves of Depository Institutions	TOTRESNS
Reserves of Depository Institutions, Nonborrowed	NONBORRES
5-Year Treasury Constant Maturity Rate	GS5
3-Month Treasury Bill Minus Federal Funds Rate	TB3SMFFM
5-Year Treasury Constant Maturity Minus Federal Funds Rate	T5YFFM
Moody's Seasoned Aaa Corporate Bond Minus Federal Funds Rate	AAAFFM
Total Consumer Loans and Leases Owned and Securitized by Finance Companies, Outstanding	DTCTHFNM
Securities in Bank Credit at All Commercial Banks	INVEST
Nikkei Stock Average, Nikkei 225	NIKKEI225
Nonfinancial Corporate Business; Total Liabilities, Level	TLBSNNCBx
Nonfinancial Corporate Business; Nonfinancial Assets, Level	TTAABSNNCBx
Nonfinancial Corporate Business; Net Worth, Level	TNWMVBBSNNCBx
Nonfinancial noncorporate business; total liabilities, Level	NNBTILQ027Sx
Nonfinancial noncorporate business; total assets, Level	NNBTASQ027Sx
Nonfinancial Noncorporate Business; Proprietors' Equity in Noncorporate Business (Net Worth), Level	TNWBSNNBx
Corporate Net Cash Flow with IVA	CNCFx
U.S.: Industrial Share Prices (2010=100)	IFS

References

- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): “Large Bayesian vector auto regressions,” *Journal of Applied Econometrics*, 25(1), 71–92.
- 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(1), 387–422.
- BEVERIDGE, S., AND C. R. NELSON (1981): “A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the ‘business cycle’,” *Journal of Monetary Economics*, 7(2), 151–174.
- BULLARD, J. (2012): “Inflation Targeting in the USA,” Speech at the Union League Club of Chicago, Breakfast@65West Chicago, Feb 6, 2012.
- CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2015): “Bayesian VARs: Specification Choices and Forecast Accuracy,” *Journal of Applied Econometrics*, 30(1), 46–73.
- CHAN, J. C. C., AND A. L. GRANT (2017): “Measuring the output gap using stochastic model specification search,” CAMA Working Papers 2017-02, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary policy shocks: What have we learned and to what end?,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 2, pp. 65–148. Elsevier.
- CHRISTIANO, L. J., AND T. J. FITZGERALD (2003): “The Band Pass Filter,” *International Economic Review*, 44(2), 435–465.
- DEL NEGRO, M., AND F. SCHORFHEIDE (2011): “Bayesian Macroeconometrics,” in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. Van Dijk, pp. 293–389. Oxford University Press.
- EDELSTEIN, P., AND L. KILIAN (2009): “How sensitive are consumer expenditures to retail energy prices?,” *Journal of Monetary Economics*, 56(6), 766–779.
- EVANS, G., AND L. REICHLIN (1994): “Information, forecasts, and measurement of the business cycle,” *Journal of Monetary Economics*, 33(2), 233–254.
- GARRATT, A., K. LEE, AND K. SHIELDS (2016): “Information rigidities and the news-adjusted output gap,” *Journal of Economic Dynamics and Control*, 70(C), 1–17.
- GARRATT, A., D. ROBERTSON, AND S. WRIGHT (2006): “Permanent vs transitory components and economic fundamentals,” *Journal of Applied Econometrics*, 21(4), 521–542.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): “Prior Selection for Vector Autoregressions,” *The Review of Economics and Statistics*, 97(2), 436–451.
- HAMILTON, J. (1994): *Time series analysis*. Princeton Univ. Press, Princeton, NJ.
- HAMILTON, J. D. (1989): “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 57(2), 357–384.

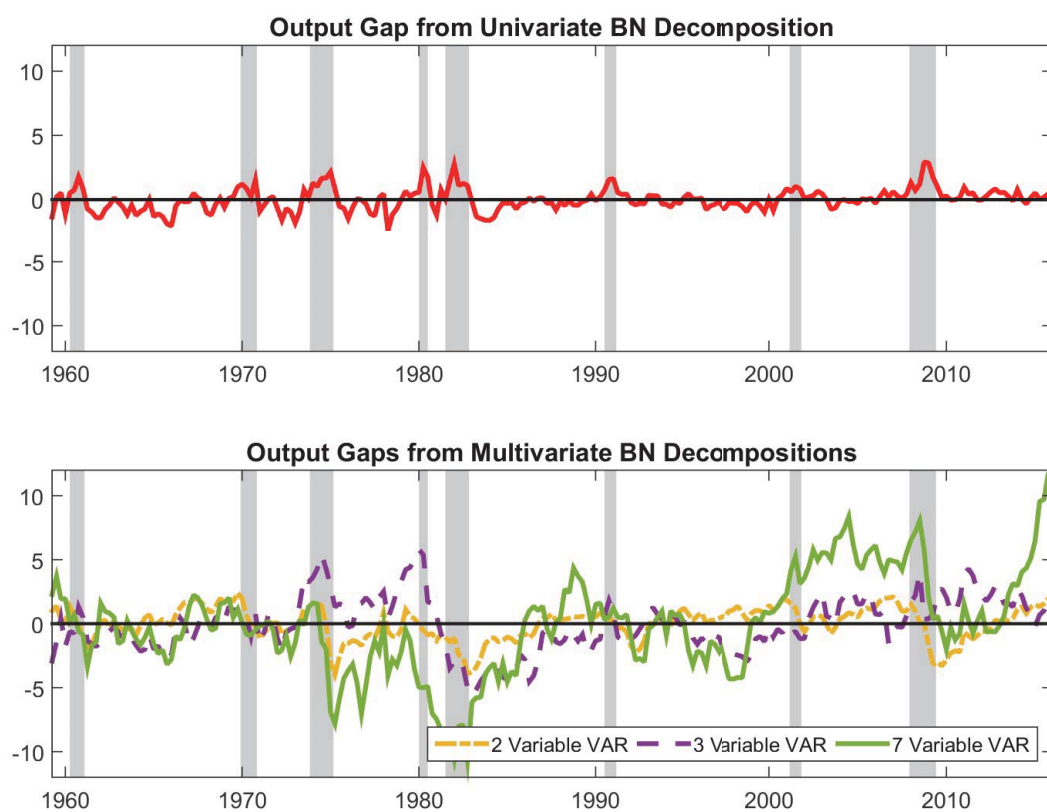
- HARDING, D., AND A. PAGAN (2010): “Can We Predict Recessions?,” NCER Working Paper 69.
- HODRICK, R. J., AND E. C. PRESCOTT (1997): “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking*, 29(1), 1–16.
- KAMBER, G., J. MORLEY, AND B. WONG (2017): “Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter,” *Review of Economics and Statistics*, forthcoming.
- KAMBER, G., AND B. WONG (2017): “An Open Economy Model of Trend Inflation,” *mimeo*.
- KILIAN, L. (2013): “Structural vector autoregressions,” in *Handbook of Research Methods and Applications in Empirical Macroeconomics*, Chapters, chap. 22, pp. 515–554. Edward Elgar Publishing.
- KILIAN, L., AND C. VEGA (2011): “Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices,” *The Review of Economics and Statistics*, 93(2), 660–671.
- KOOP, G. M. (2013): “Forecasting with Medium and Large Bayesian VARs,” *Journal of Applied Econometrics*, 28(2), 177–203.
- KOZICKI, S. (1999): “Multivariate detrending under common trend restrictions: Implications for business cycle research,” *Journal of Economic Dynamics and Control*, 23(7), 997–1028.
- LITTERMAN, R. B. (1986): “Forecasting with Bayesian Vector Autoregressions-Five Years of Experience,” *Journal of Business & Economic Statistics*, 4(1), 25–38.
- MORLEY, J., AND J. PIGER (2012): “The Asymmetric Business Cycle,” *The Review of Economics and Statistics*, 94(1), 208–221.
- MORLEY, J. C. (2002): “A state-space approach to calculating the Beveridge-Nelson decomposition,” *Economics Letters*, 75(1), 123–127.
- NELSON, C. R. (2008): “The Beveridge-Nelson decomposition in retrospect and prospect,” *Journal of Econometrics*, 146(2), 202–206.
- ROBERTSON, J. C., AND E. W. TALLMAN (1999): “Vector autoregressions: forecasting and reality,” *Economic Review*, (Q1), 4–18.
- SINCLAIR, T. M. (2009): “The Relationships between Permanent and Transitory Movements in U.S. Output and the Unemployment Rate,” *Journal of Money, Credit and Banking*, 41(2-3), 529–542.
- STOCK, J., AND M. WATSON (2016): “Factor Models and Structural Vector Autoregressions in Macroeconomics,” in *Handbook of Macroeconomics*. Elsevier.
- TAYLOR, J. B. (2012): “Monetary Policy Rules Work and Discretion Doesn’t: A Tale of Two Eras,” *Journal of Money, Credit and Banking*, 44(6), 1017–1032.
- WONG, B. (2015): “Do Inflation Expectations Propagate the Inflationary Impact of Real Oil Price Shocks?: Evidence from the Michigan Survey,” *Journal of Money, Credit and Banking*, 47(8), 1673–1689.
- WOŹNIAK, T. (2016): “Bayesian Vector Autoregressions,” *Australian Economic Review*, 49(3), 365–380.

Table 1: Monte Carlo Results

		Estimated Model Size			
		1 variable	8 variables	23 variables	
(a) Mean and standard deviation of RMSEs					
DGP	1 variable	Shrinkage	0.39 (0.212)	1.82 (0.916)	2.77 (1.833)
		MLE	0.28 (0.188)	2.08 (0.929)	4.18 (2.049)
	8 variables	Shrinkage	1.47 (0.176)	0.92 (0.401)	0.94 (0.348)
		MLE	1.52 (0.194)	1.07 (0.606)	1.59 (0.709)
	23 variables	Shrinkage	1.33 (0.191)	0.73 (0.378)	0.61 (0.297)
		MLE	1.36 (0.207)	0.90 (0.492)	1.20 (0.510)
(b) Proportion of Monte Carlo Draws $RMSE_{\text{Shrinkage}} < RMSE_{\text{MLE}}$					
		Estimated Model Size			
		1 variable	8 variables	23 variables	
DGP	1 variable	0.27	0.77	0.91	
	8 variables	0.99	0.62	0.96	
	23 variables	0.99	0.81	0.99	


Notes: Root mean square errors (RMSE) are calculated for estimates compared to the true output gap for a given simulation under the respective DGP. Estimated model size indicates the size of the variable set for which the VAR is estimated to derive the output gap. The three DGPs contain 1, 8, and 23 variables respectively. Panel (a) presents the mean and standard deviation (in parenthesis) of the RMSEs across 1000 Monte Carlo draws. Shrinkage and MLE refers to respectively using our procedure or MLE to estimate the VAR. Panel (b) counts the proportion of Monte Carlo draws where our shrinkage procedure produces an output gap that has a lower RMSE relative to the true output gap compared to estimating the VAR using conditional maximum likelihood.

Figure 1: Estimated Output Gap from Univariate and Multivariate BN Decompositions



Notes: Units are 100 times natural log deviation from trend. Shaded bars correspond to NBER recession dates. See footnote 1 for descriptions of the 2 variable, 5 variable and 7 variable VAR systems.

Figure 2: Estimated Output Gap for Benchmark BVAR Model



Notes: Units are 100 times natural log deviation from trend. Shaded bars correspond to NBER recession dates.

Figure 3: Standard Deviations of Informational Contributions


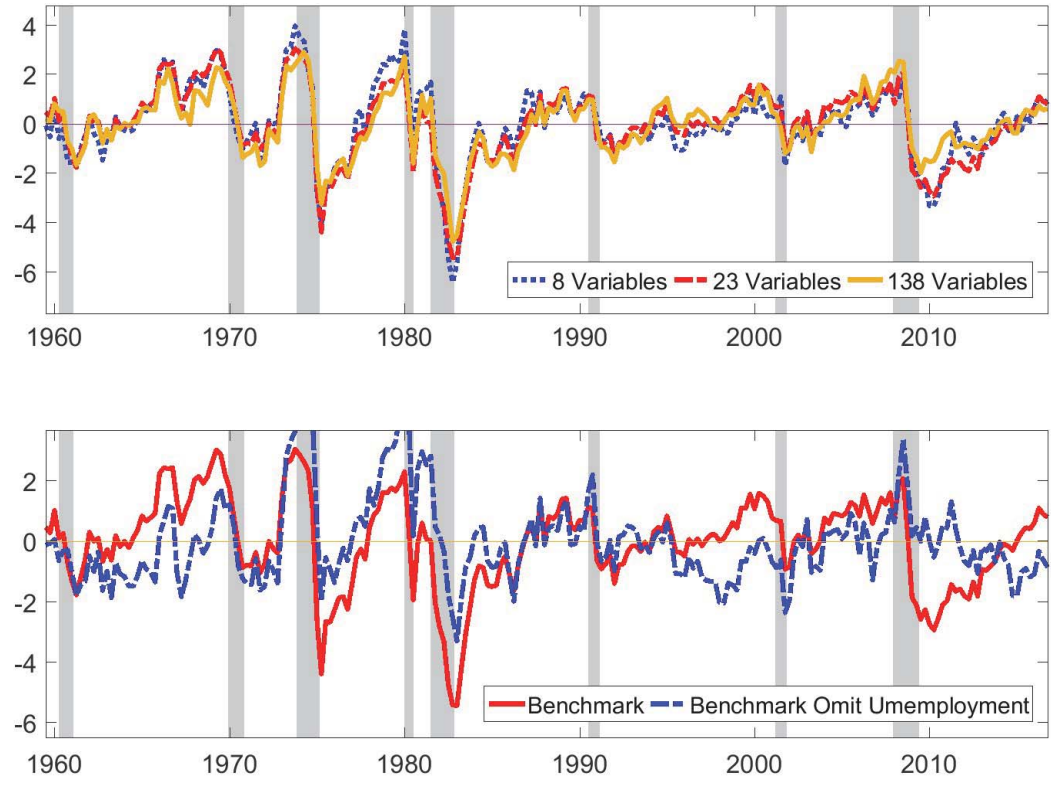



Figure 4: Estimated Output Gap for Various Sized BVAR Models




Notes: Units are 100 times natural log deviation from trend. Shaded bars correspond to NBER recession dates.

Figure 5: One-Step-Ahead Out-of-Sample Root Mean Square Error



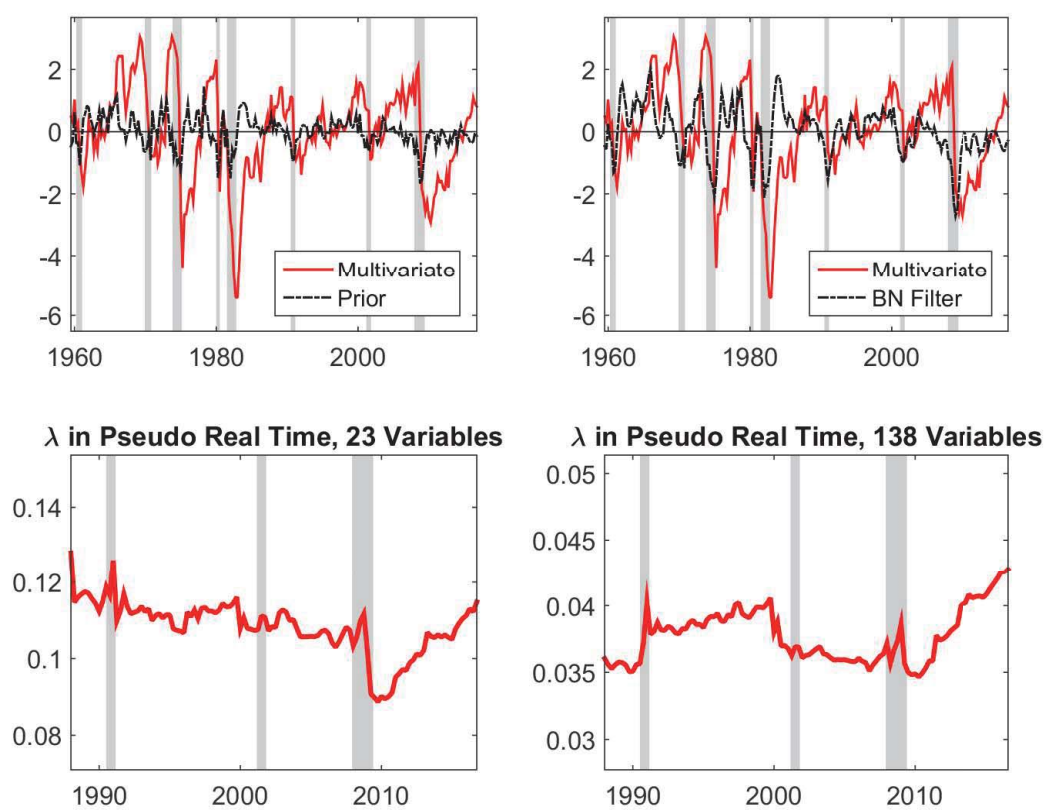
Notes: The horizontal axis represents the tightness of the prior on the hyperparameter λ . The vertical axis represents the one-step-ahead out-of-sample root mean square error. 8 variables, 23 variables, and 138 variables refer to the size of the various BVARs. The horizontal line is the one-step-ahead out-of-sample root mean square error forecasting output growth with an AR(1) estimated by least squares.

Figure 6: Role of Consumption




Notes: Cumulative change in trend is in terms of 100 times natural logs. Shaded bars correspond to NBER recession dates. Real personal consumption expenditure is plotted as a quarterly percent change.

Figure 7: Role of Multivariate Information




Notes: Units in top panels are 100 times natural log deviation from trend. Shaded bars correspond to NBER recession dates.

Figure 8: Variance Decompositions



Notes: Percentage of total variation

Figure 9: Historical Decomposition of the Estimated Output Gap



Notes: Units are 100 times natural log deviation from trend. Shaded bars correspond to NBER recession dates.