Handelshøyskolen BI

GRA 19703 Master Thesis

Thesis Master of Science 100% - W

Predefinert informasjon

Startdato:	16-01-2022 09:00	**Termin:**	202210
Sluttdato:	01-07-2022 12:00	**Vurderingsform:**	Norsk 6-trinns skala (A-F)
Eksamensform:	T		
Flowkode:	202210||10936||IN00||W||T		
Intern sensor:	(Anonymisert)		

Deltaker

Navn:	Anna Holden Rotheim og Stine Thomle Karlsen

Informasjon fra deltaker

Tittel *:	What Drives the Output Gap in Norway? - Evidence from a multivariate Beveridge-Nelson decomposition
Navn på veileder *:	Jamie Cross

Inneholder besvarelsen konfidensielt materiale?:	Nei	**Kan besvarelsen offentliggjøres?:**	Ja

Gruppe

Gruppenavn:	(Anonymisert)
Gruppenummer:	196
Andre medlemmer i gruppen:	

What Drives the Output Gap in Norway?

Evidence from a multivariate Beveridge-Nelson decomposition

Master Thesis

by

Anna Holden Rotheim & Stine Thomle Karlsen

MSc in Business, Major in Economics

June 2022

Supervisor: Jamie Cross

Abstract

We estimate the trend and cycle of real GDP in Norway using a multivariate Beveridge–Nelson decomposition with a large information set. This method allows us to identify a small collection of variables that are major business cycle drivers. Using a dataset of 76 variables covering various sectors of the economy, we find that the implied output gap measure accounts for all historical recessions between 1983 and 2021 and reveals a minimal set of variables that are important drivers of the business cycle: Unemployment, Total Reserves, GDP growth, Government Final Consumption Expenditure, Wages, Sight Deposit Rate, Income, Hours Worked, and Private Final Consumption Expenditure.

Acknowledgements

We wish to thank our supervisor Jamie Cross for the exceptional guidance throughout the entire process, as well as his knowledge and encouragement within this field of research. We would also like to thank Kjersti Haugland, Chief Economist at DnB Markets, for her comments on historical business cycles in Norway. Furthermore, we wish to thank Statistics Norway for allowing us to present our model and for posing questions that allowed us to broaden our perspective. Finally, we wish to extend our thanks to Benjamin Wong for clarifications on calculations in his 2020 paper with James Morley.

Table of Contents

1. Executive Summary

As the economy continues to grow and becomes more vulnerable to international and domestic macroeconomic developments, it would be highly beneficial to be able to predict the future economic path. The business cycles of a country can measure the economic situation and have implications for fiscal and monetary policy, and our findings may give valuable insight to policymakers. Identifying which variables drive the economy, in order to forecast business cycles, provides an opportunity for new knowledge of the path of output and its deviations from trend. In order to measure the business cycle, we can use the output gap to estimate the most important variables that drive the output growth. The output gap describes the difference between actual and potential output, where potential output is the level of output that is consistent with stable inflation. A positive output gap, meaning that actual output is higher than potential output, indicates pressures in the economy, usually accompanied by rising inflation. Contrariwise, a negative output gap indicates spare capacity and falling inflation (Bjørnland, Brubakk, & Jore, 2005). This output gap may be informative for the central bank to quantify macroeconomic indicators.

The aim of our thesis is twofold. The first objective is to estimate the output gap using a multivariate Beveridge-Nelson decomposition with large Bayesian vector autoregressions. The second objective is to establish which conditioning variables are most influential in generating output gap fluctuations. For both of these objectives, we wish to utilize the methods proposed by Morley and Wong (2020) in their paper "Estimating and accounting for the output gap with large Bayesian vector autoregressions" and apply this to Norwegian data. In doing so, we must consider that, in contrast to the US, Norway is a small open economy with a more considerable reliance on oil exports (Bjørnland, 2000), which might have particular implications for our second objective.

As a first step, we collect a dataset containing 76 macroeconomic variables informed by a number of sources that may, in totality, contain sufficient information to describe the Norwegian economy. Next, we shrink the model towards a smaller collection of relevant variables and establish a 26-variable benchmark model. We estimate the output gap using our benchmark model to assess our first objective and show that this collection of variables accurately represents the large dataset. To investigate our second objective, we calculate the standard deviation of informational contribution of the variables to the estimated output gap to identify the variables that contribute most to output gap fluctuations. We then apply this insight to estimate the output gap using a 9-variable model to examine objective two further. To check for robustness, we compare all three various-sized models and find that our results are robust in estimating the output gap fluctuations. We find that the nine variables: unemployment, total reserves, GDP growth, government final consumption expenditure (GFCF), wages, income, sight deposit rate, hours worked, and private final consumption expenditure (PFCE) provide the minimum set of variables that can accurately estimate the output gap fluctuations, and as such can, in future research, be informative for predicting future business cycles.

2. Literature Review

2.1 Previous research on Norwegian business cycles

In order to shed light on our research question, we want to examine previous literature on Norwegian business cycles for comparison. For this reason, we will briefly discuss the papers by; Bjørnland (2000), Aastveit, Jore, and Ravazzolo (2016), and Bjørnland, Brubakk, and Jore (2005).

Bjørnland (2000) analyzes the stylized facts of the Norwegian business cycles by comparing different detrending methods such as the Beveridge-Nelson (BN) decomposition and the Hodrick-Prescott (HP) filter. Estimating the trend and cycle using the BN decomposition, Bjørnland treats the trend as the long-term forecast of the series after adjusting for the average rate of change. This method was applied to establish the initial trend component and determine the trend and cycle components. The paper highlights some advantages but also explains some disadvantages of fitting low order Autoregressive Integrated Moving Average (ARIMA) models in a univariate setting, which seemingly tend to overestimate the random walk components of the data.

Interpreting the results, which emphasize a structural break in the trend in some of the variables, Bjørnland found some stylized facts. For all decompositions, the variables of investments and imports in Norway are highly volatile, while consumption and productivity have lower volatility than GDP. In contrast, Bjørnland found that when using a BN decomposition, the variable 'money' also ranks as the most volatile. Furthermore, the paper shows that GDP is procyclical while unemployment is persistently countercyclical and leads the cycle by one quarter, except when using the BN decomposition. Overall, the correlation coefficients are much lower than in other OECD countries. According to the paper, there is evidence that investment is leading GDP, whereas imports and productivity are lagging GDP. Furthermore, Norwegian output has a small correlation coefficient compared to other industrial countries, explained by the "fact that Norway is a small oil-producing country, which has experienced much more idiosyncratic shocks than many other OECD countries" (Bjørnland, 2000, p. 390). Finally, the results of Bjørnland's paper show that using the HP filter on difference-stationary or trend-stationary time series may create spurious cycles.

Another paper that has aimed to forecast Norwegian business cycles is the article by Aastveit, Jore, and Ravazzolo (2016). Their article defines and forecasts Norwegian business cycle turning points using a univariate and a multivariate Bry-Boschan (BB) approach. For technical details on the BB approach, see Bry and Boschan (1971). The authors use real-time out-of-sample forecasting to show that univariate Markov-switching models are timely and accurate in calling the last peak in real-time but not too accurate and timely in calling the through in real-time.

The abovementioned paper yields interesting results that may provide valuable insights into our two objectives. The methods used in their study differ from ours, but the article defines the historical business cycles of Norway in the period 1978Q1-2011Q4 and investigates which variables to use in real-time forecasting of Norwegian business cycles. The authors argue that one should use data for mainland Norway as the measure of economic activity when investigating economic conditions in Norway. In doing so, one excludes oil and gas extraction and international shipping by omitting offshore activity. The reasoning behind their choice is that the large fluctuations of offshore activity may have minor short-term effects on the Norwegian labor market and domestic production. Further, in the short-term, the mainland economy is insulated from fluctuating revenue from the petroleum sector, and all revenues are transferred to the sovereign wealth fund with fiscally determined withdrawals each year. In their choice of variables for inclusion in the models, the authors find that particularly important variables in the Norwegian economy are the Brent blend oil price, employment in mainland Norway, household consumption, private real investments in Norway, exports of traditional goods, and GDP for mainland Norway.

Aastveit et al. focus on predicting peaks and troughs of the business cycle, and they do not detrend the data as is often done when working with business cycles. In analyzing the dates for troughs and peaks in the Norwegian economy, the authors find that taking durations and amplitude into account while using a so-called "triangular approach," the size of the cumulative change from peak to trough in Norway is quite similar to that in the US but smaller than in the other European countries. On the other hand, for the cumulative change from trough to peak, the numbers are much larger, and still, the statistics of the Norwegian economy are closer to the US statistics than the other countries. While this does not mean that the Norwegian and US economies are the same, it is interesting to note the similarity in movement.

There are many ways to estimate the output gap in an economy, and the paper by Bjørnland, Brubakk, and Jore (2005) investigates some methods of estimation in the Norwegian economy. Different estimation methods may yield different results, and historical estimates might also change when the data is revised. The problem of data revision is primarily disregarded in the paper but is still a concern worthy of mention. For the comparison, the authors have decided to group the methods into two main categories: univariate methods and multivariate methods.

Bjørnland et al. propose several methods of detrending in a univariate case, such as the HP filter, band-pass, and univariate "unobserved component" methods (UC). The HP filter and band-pass decomposition have several drawbacks, like the end-of-sample problem for the HP filter and the exclusion of the estimated output gap at the beginning and end of the sample for band-pass when not extending the dataset with forecasting. For the multivariate case, the paper introduces the production function method (PF), which emphasizes the supply-side factors like resources and technology on the potential output. Considering the use of HP-filtering in this method makes us aware of the potential problems of PF. A second method in the multivariate case is the multivariate "unobserved component" method (MVUC) which uses unemployment, domestic inflation, and the relationship between these two variables and the output gap. Although this method has several advantages, it is highly dependent on assumptions that must be placed on the relationship between the variables. Finally, the structural vector autoregression (SVAR) model in the paper introduces three variables to estimate the output gap and has the advantage of imposing few constraints. However, the constraints must be highly consistent with economic theory to avoid misleading results.

2.2 Previous literature on the output gap

The literature on the output gap is too vast and we do not attempt to cover its entirety in this section. For a more complete overview of different estimation methods for the output gap, see Bjørnland and Thorsrud (2015). In this section, we focus on a paper by Morley and Wong (2020) which estimates and forecasts the output gap and finds the most impactful variables using US data.

The paper uses the BN decomposition based on a Bayesian vector autoregression to estimate the trend and cycle of a large information set. The article tackles two considerations: firstly, they examine which conditioning variables contain the relevant information pertaining to the output gap by solving for the BN trend as a function of forecast errors for different variables, and in doing so, they are accounting for the variables' contribution. Secondly, they utilize the Bayesian shrinkage method to mitigate the problem of overfitting in finite samples when working with models of large information sets. Morley and Wong used an empirical application of up to 138 different variables covering the US economy and found that the most important variables containing information beyond that in output growth for estimating the output gap are the unemployment rate and inflation. Further, they found six additional variables that to a lesser extent contain relevant information, which in total makes for an 8-variable model that accurately estimates the output gap.

In order to determine the variables that span the relevant information, Morley and Wong account for the BN trend and cycle in terms of contributions from different forecast errors in the vector autoregression (VAR). These forecast errors are used to define a relevant information set and to interpret which variables are most important for estimating the trend and cycle of a target variable. To avoid overfitting the model, the authors applied Bayesian shrinkage, using a Minnesota prior with a key hyperparameter is considered to minimize the pseudo-out-of-sample forecast error variance for the target variable, mitigating the effects of sampling error for larger systems.

2.3 Research on our choice of methodology

As we have chosen to apply the Beveridge-Nelson decomposition and utilize Bayesian estimation, we find it necessary to evaluate previous literature on related subjects to validate our choice of methodology. Firstly, we examine the original paper by Beveridge and Nelson (1981) as the baseline for our approach, followed by an introduction to the proposed approach to avoid overestimation by Banbura, Giannone, and Reichlin (2010), using Bayesian shrinkage. Finally, we discuss sources on the advantages and disadvantages of different detrending methods.

The well-known paper by Beveridge and Nelson (1981) introduced a new method to decompose data and shed light on the advantages of using their proposed approach. The method decomposes the non-stationary time series into a permanent and a transitory component, in which the permanent component is a random walk with drift and the transitory one is a stationary process with a zero mean.

The methodology proposed by Beveridge and Nelson is based on the observation "that any time series which exhibits the kind of homogeneous non-stationarity typical of economic time series can be decomposed into two additive components, a stationary series and a pure random walk" (Beveridge & Nelson, 1981, p. 153). Their paper shows that the permanent component, i.e., the trend, of their decomposition will always be a random walk with drift. The transitory component is a stationary process with a zero mean, which perfectly correlates with the permanent component. For this reason, we wish to use the BN decomposition so as not to infer spurious cycles when dealing with random walks with drift. This decomposition method also allows the time series to contain a unit root.

Another problem we wish to mitigate is that of the end-of-sample problem. This problem presents severe concerns for studying real-time developments in indicator series since future observations are unavailable, and a moving average method would have to fill this gap, typically with the latest observations of the sample. The way in which Beveridge and Nelson proposed to decompose time series allows us to mitigate this problem. The procedure of cycle measurement is as follows; firstly, there is an identification of an ARIMA model for the first differences of the non-stationary series of interest, and secondly, a numerical evaluation of the cycle component using a practical equivalent of the uncorrelated error terms and their constants, lambda. In doing so, the computed value of the cyclical component will only contain past values of the observed series, thus mitigating the end-of-sample problem.

Morley and Wong (2020) express a concern that when using a multivariate VAR, the many sources of information in the data may disturb the estimating model and generate an overestimation of the variables, which may impact which variables drive the output gap. Banbura et al. (2010) solve this problem and show how Bayesian shrinkage in a vector autoregression model is an appropriate tool for structural analysis in large dynamic models. The paper introduces parameter restrictions to account for overparameterization when the model contains a larger set of variables. According to the paper, Bayesian shrinkage in dynamic systems suggests a better forecast performance and solves the problem of over-parameterization. The paper draws further on a discovery by Litterman, who introduced the Minnesota prior, which is a prior belief on the parameters, and found that setting the degree of shrinkage in relation to the model size controls for an over-fitting of the model in the case of collinearity in large systems. This Minnesota prior is a principle that all equations are 'centered' around the random walk with drift for some variables and white noise for others. This prior is a natural conjugate and "shrinks all the VAR coefficients towards zero except for coefficients on the first lags of the dependent variable in each equation" (Koop, 2013, p. 177). The findings show that the overall shrinkage should increase as the model itself increases.

Furthermore, Banbura et al. introduce three sizes of models from previous literature; a 7-variable model, a 20-variable, and a 131-variable model, and finds a sufficient size of model for the purpose of forecasting. The process of selecting variables starts with making the series stationary using the first difference. Then, a random walk prior is selected, and the paper concludes that a 20-variable BVAR model is feasible and remains robust. The paper by Banbura et al. is fundamental for selecting our model size and the forecast performance in terms of including a shrinkage prior.

Previous literature on decomposition methods by Christiano and Fitzgerald (2003) and Hodrick and Prescott (1997), namely the band-pass and HP decomposition, respectively, are typically used in a univariate setting. However, according to Morley and Wong (2020), these two decomposition methods often must be validated with information outside the model with other sources of information, and it can be difficult in a multivariate setting to determine which variables to include in the information set and how large it can be. The BN decomposition addresses this problem in combination with a vector autoregression.

As proposed in Evans and Reichlin (1994), the BN decomposition handles the cyclical component of output with the economically beneficial interpretation that the growth of output that exceeds the trend of output is forecasted as the return to trend. That is, the deviation between the output trend over time and the current level of output is the cycle, often referred to as the output gap. In contrast to other detrending methods, such as the HP filter, the BN decomposition proposes advantages in mitigating problems that can occur in time series data. In a critique by Nelson and Kang (1981), using the fact that the unit root of a time series cannot be rejected and that stochastic shocks to the output permanently affect the time series, they found that detrending data that is following a random walk will infer spurious cycles in the data. Spurious cycles in the series can come from detrending a deterministic trend and may cause the persistence of the cyclical component to be over-estimated while under-estimating the trend component. Using the BN decomposition, thus allowing the series to contain a highly volatile unit root, we mitigate the problem of spurious cycles when dealing with random walks with drift. Following the critique from Nelson and Kang, the detrending mechanisms in the HP filter may generate spurious cycles even if they are not present in the data (Bjørnland, 2000).

According to the critique by Hamilton (2018), the HP-filtering also proposes the disadvantage of the end-of-sample problem. As the detrending method penalizes any temporary shocks in the trend with its smoothing parameter using data ex-ante and ex-post its current level, any penalizing at the end of the sample will be absent. The lack of smoothing results in more sensitivity at the end of the sample, thus allowing for spurious cycles and higher volatility.

Although the BN detrending of time series may give different results when choosing different levels of ARMA, Morley and Wong (2020) find that the results are relatively robust when including more lags, given Bayesian shrinkage. A summary of the advantages and disadvantages of the most common detrending method is presented in Table 1 below.

Table 1: Detrending methods

Detrending	Advantages	Disadvantages
Hodrick-Prescott	Easy to understand Easy to compute	End-of-sample problem Spurious cycles Typically used in a univariate setting Validates outside the model Choosing smoothing parameter, a prioi Results are not robust to the value of the smoothing parameter
Band-Pass	Can easily change frequency, which expands the range of questions we can explore	End-of-sample problem Typically used in univariate setting Validated outside the model Selection of preferred frequencies a priori
Beveridge-Nelson	Easy to understand Relatively easy to implement Not inferring spurious cycles Mitigates end-of- sample problem Multivariate VAR	Overestimation of variables Time-consuming Selection of ARMA may give different results

3. Methodology and estimation

3.1 Theory and Model

3.1.1 BN Decomposition for a single variable

In the original article by Beveridge and Nelson (1981), they propose a univariate method of decomposing the trend and cycle. Generally, in order to identify the permanent and cyclical components, it is necessary to specify models that can be written as a stationary moving average process. For this reason, BN decompositions typically utilize ARIMA models to numerically evaluate the cycle using the uncorrelated error terms and their constants.

It has, however, later been shown that the implied trend of the BN method can be defined as the minimum mean squared error (MSE) forecast of the long-run level of the series minus any deterministic drift. Alternatively, one can formulate the BN trend as the present level of the series plus the infinite sum of the minimum MSE j-period ahead first difference forecasts (Bjørnland & Thorsrud, 2015, p. 142). Specifically, let y_t be a time series process with a trend component that follows a random walk with a constant drift μ, the BN trend τ_t at time t is:

$$\tau_t = \lim_{j \to \infty} [\, y_{t+j} - j * \mu \,] \qquad (1)$$

The BN cycle, c_t, is then simply the difference between the observed time series and trend at time t:

$$c_t = y_t - \tau_t \qquad (2)$$

3.1.2 BN Decomposition for multiple variables

To address objective two, we would need the use of a multivariate setting to investigate the contribution of each of the benchmark variables to output gap fluctuations. For this reason, we will move away from a univariate setting as described in Beveridge and Nelson (1981) and consider a multivariate setting as described in a paper by Evans and Reichlin (1994). Evans and Reichlin show that when working with a multivariate setting, we will obtain different estimates than those of a univariate setting using the BN decomposition. It is important to note that, while univariate BN decompositions typically imply ARIMA models, it is more common to consider a linear VAR model when moving to the multivariate case. Intuitively, output growth can be more accurately forecasted in multivariate models, resulting in a larger proportion of fluctuations in output growth being attributed to the cyclical component compared to the univariate case (Evans & Reichlin, 1994). Following the arguments of Morley and Wong (2020), we also know that the inclusion of variables that do not necessarily contain marginally relevant information for the forecast of the target variable will not yield different results than if we merely included the marginally relevant information sets. However, omitting one or more variables that span the relevant information, the BN cycle would have a strictly smaller variance. The key insight here is that the exclusion of any variable from the VAR representation that alters the BN cycle for output belongs in the model.

If the vector of variables of interest, Δx_t, has a finite-order VAR(p) representation, then it has the following companion form (Morley & Wong, 2020):

$$(\Delta X_t - \mu) = F(\Delta X_{t-1} - \mu) + He_t \qquad (3)$$

Where $\Delta X_t = \{\Delta x_t', \Delta x_{t-1}', \dots, \Delta x_{t-p+1}'\}'$, μ is a vector of unconditional means, F is the companion matrix, H is the VAR forecast errors of the companion form, and e_t is a vector of serially uncorrelated forecast errors, such that $He_t \sim N(0, \Sigma)$, and the eigenvalues of the companion matrix are less than one in absolute value. Consequently, given stationarity, we assume $(I - F)^{-1}$ exists.

Next, we wish to solve for the BN trend and cycle. We denote τ_t and c_t as vectors of the trend and cycle, respectively, and we calculate the trend to be the present level of the series, X_t, plus the infinite sum of the minimum MSE j-period ahead first difference forecasts (Morley, 2002):

$$\tau_t = X_t + \lim_{j \to \infty} \sum_{j=1}^{\infty} E\left(\Delta X_{t+j} - \mu\right) \qquad (4)$$

Analytical derivation from Equation (4) to Equation (5) can be found in section B of the Appendix. The trend is further defined by using the infinite sum of the geometric sequence of the companion matrix, F. The infinite sum of geometric series/sequence $\sum_{t=0}^{\infty} F^t = \frac{1}{1-F} = (1-F)^{-1}$, where $|F| < 1$, given the stationary time series. Thus, we can define the trend as follows:

$$\tau_t = X_t + F(I-F)^{-1}(\Delta X_t - \mu) \qquad (5)$$

The observed time series vector X is defined by the BN trend and cycle, $X_t = \tau_t + c_t$, and the cycle is then defined by:

$$c_t = -F(I-F)^{-1}(\Delta X_t - \mu) \qquad (6)$$

The original paper by Beveridge and Nelson examined the univariate case, and since we are in a multivariate setting, we turn to Evans and Reichlin (1994) to determine how to find the relevant conditioning variables for the model. In this setting, the relative importance of the cyclical component depends on the size of the information set and is necessarily higher with multivariate BN decompositions. Intuitively, a larger information set may result in better forecasts of output growth, thus leading to a bigger attribution of the output fluctuations to the cyclical component. The variance of the trend component is invariant to the size of the information set, but the variance of the cyclical component is not. Generally, a larger information set leads to a higher variance of the cyclical component. However, should we have variables in the information set that are not marginally relevant for the forecasting of the target variable, then we would not see an increase in the cycle variance as the size of the information set increases. Following the notation of Morley and Wong (2020), let w_t denote the vector of variables that span the relevant information, which follows a $VAR(p^*)$ process where p^* is the true lag length of the VAR process. Then, with population values for F and μ, and with $\Delta X_t = W_t$ where $W_t = \{w_t', \ w_{t-1}', \dots, \ w_{t-p^*+1}'\}'$, Equations (5) and (6) would recover the true trends and cycles for x_t.

Turning back to Evans and Reichlin (1994), we know that a smaller information set, denoted V_t, would yield a cycle with strictly smaller variance. Should we, however, have a larger information set that includes extraneous variables that do not contain marginally relevant information for forecasting the target variable, the variance of the cycle would be unchanged compared to the case of W_t. Denoting this larger information set as Z_t, we then have $V_t \subset W_t \subset Z_t$ in which the variance of the cycle would only change if we moved from W_t to V_t, given population values for F and μ. In practice, however, the cycle variance will change also when going from W_t to Z_t due to sampling error, with a strictly larger variance for Z_t.

Morley and Wong (2020) propose a practical way, based on the observations of Evans and Reichlin (1994), to determine which conditioning variables span the relevant information for forecasting the target variable. Inserting Equation (3) into Equation (6), where $\Gamma_i \equiv F^i(I-F)^{-1}$, they obtain a definition of c_t as a function of historical forecast errors:

$$\begin{aligned} c_t = -\Gamma_1(\Delta X_t - \mu) &= -\Gamma_1\{F(\Delta X_{t-1} - \mu) + He_t\} \\ &= -\Gamma_1 He_t - \{F\Gamma_1(\Delta X_{t-1} - \mu)\} \\ &= -\Gamma_1 He_t - \{\Gamma_2(\Delta X_{t-1} - \mu)\} \\ &= -\sum_{i=0}^{t-1} \Gamma_{i+1} He_{t-i} - \Gamma_{t+1}(\Delta X_0 - \mu) \approx -\sum_{i=0}^{t-1} \Gamma_{i+1} He_{t-i} \qquad (7) \end{aligned}$$

This multivariate method allows us to distinguish the variables most relevant for the cycle from the less relevant variables, thus establishing a unique minimum set of conditioning variables that span the relevant information for forecasting the target variable. Therefore, in contrast to the univariate BN decomposition, we can examine which conditioning variables contain the most important information for predicting output gap fluctuations to investigate our research question. The proposed approach by Morley and Wong (2020) is to start with a VAR based on the entire dataset and drop the variables that do not contain marginally relevant information until w_t is found. The key insight of this approach, and the paper by Evans and Reichlin (1994), is that any variable whose removal alters the cycle for the target variable belongs in the model.

Following from Equation (7) and using the fact that Γ_{i+1} converges to zero when t increases since the series is stationary, we can define a new vector $s_{r,q}$. This is a vector of zeros in all rows, r, but equals 1 in the qth row, assuming n variables and p lags in the VAR. The contribution of the forecast error in the kth variable is defined as:

$$c_{k,t} = -\sum_{i=0}^{t-1} s'_{np,1}\Gamma_{i+1}Hs_{n,k}s'_{n,k}e_{t-i} \qquad (8)$$

This accounts "for the contribution of the forecast errors for the kth variable to the BN cycle of the lth-ordered target variable" (Morley & Wong, 2020, p. 4). For completeness, Morley and Wong define the BN trend growth as a function of forecast errors, which can be found by taking the first difference of Equation (5):

$$\Delta\tau_t = X_t + \Gamma_1(\Delta X_t - \mu) - \{X_{t-1} + \Gamma_1(\Delta X_{t-1} - \mu)\} = \mu + \Gamma_0 He_t \qquad (9)$$

3.2 Estimation

As the Norwegian economy is comprised of a vast set of macroeconomic variables, the need for a large dataset makes it hard to accurately estimate the output gap using ordinary least squares (OLS) estimation. Although OLS has the advantage of being the best linear unbiased estimator (BLUE), it can be hard to estimate large datasets as the estimates become imprecise with large standard deviations and imprecise coefficients. This comes from the bias-variance tradeoff, where OLS suggests a bias of zero, but the variance might be very high. It is worth noticing that if the data has enough information, it will always be pushed to the OLS estimate. A potential solution to this problem in large datasets is to introduce some bias to lower the variance and thus lower the mean squared errors to produce a better fit. This motivates us to introduce a ridge (penalized) regression to increase the bias slightly in order to decrease the variance. The connection between a penalized regression and a Bayesian prior allows us to use the latter prior to obtain the same point estimates, with the advantage of the probabilistic interpretation and the easily quantified methodology.

Bayesian estimation is a probabilistic approach to estimating model parameters founded on Bayes theorem, with the idea being that we want to estimate parameters θ given data Y:

$$p(\theta|Y) = \frac{p(\theta, Y)}{p(Y)}$$

The left-hand side is the posterior probability, and the right-hand side is the joint distribution between the parameters and the data. This joint distribution can be factorized into a product of the conditional distribution of the data and the marginal distribution of the parameters:

$$p(\theta, Y) = p(Y|\theta)p(\theta)$$

Where $p(Y|\theta)$ is the likelihood function and $p(\theta)$ is the prior probability distribution of the parameters. The likelihood function emphasizes a fundamental part of the Bayesian perspective, namely that the observed data is given and thus non-random, as opposed to a frequentist perspective in which the observed data, Y, is just a single realization of a data generating process. Accordingly, by the Bayesian perspective, the parameters are treated as random and consequently are characterized by a probability distribution. This prior distribution characterizes the uncertainty about the parameters before observing the data. Finally, in order to obtain the posterior probability, the joint distribution is divided by a marginal likelihood of the data, $p(Y)$. In total, Bayes' rule is simply represented in the following equation:

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)}$$

The relationship between the joint distribution of the data and the parameters also has the advantageous interpretation that Bayes' rule is a learning mechanism in which prior beliefs about the parameters are updated using information from the data in order to obtain the posterior distribution:

$$p(\theta|Y)p(Y) = p(\theta, Y) = p(Y|\theta)p(\theta)$$

As the parameters are treated as random, some prior belief must be specified by the researcher, and there is a wide variety of prior distributions based on different arguments, such as general properties of macroeconomic time series, theoretical models, and ease of computation (Woźniak, 2016). One such prior distribution is the so-called Minnesota prior, motivated by the observation that macroeconomic time series typically are unit-root non-stationary, meaning they can be interpreted as multivariate random walks. The key hyperparameter of the Minnesota prior, calibrated to minimize the pseudo out-of-sample forecast error variance of the target variable, determines the degree of shrinkage (Morley & Wong, 2020), which shrinks the parameters with weak signals to zero in order to enable more precise estimates of the remaining parameters. The general intuition behind the specification of priors is illustrated in Figure 1, in which prior 3 represents a flat prior that is uninformed, prior 2 represents an informative prior, and prior 1 is the case in between. Prior 3 does not impose any shrinkage on the parameters, whereas prior 2 imposes a high degree of shrinkage. We want to minimize the mean squared estimate (MSE) consisting of the bias and variance of the model fit, and prior 1 illustrates the middle case in which we impose some bias in order to decrease the variance.

Figure 1: Illustration of an informative and uninformative prior distribution

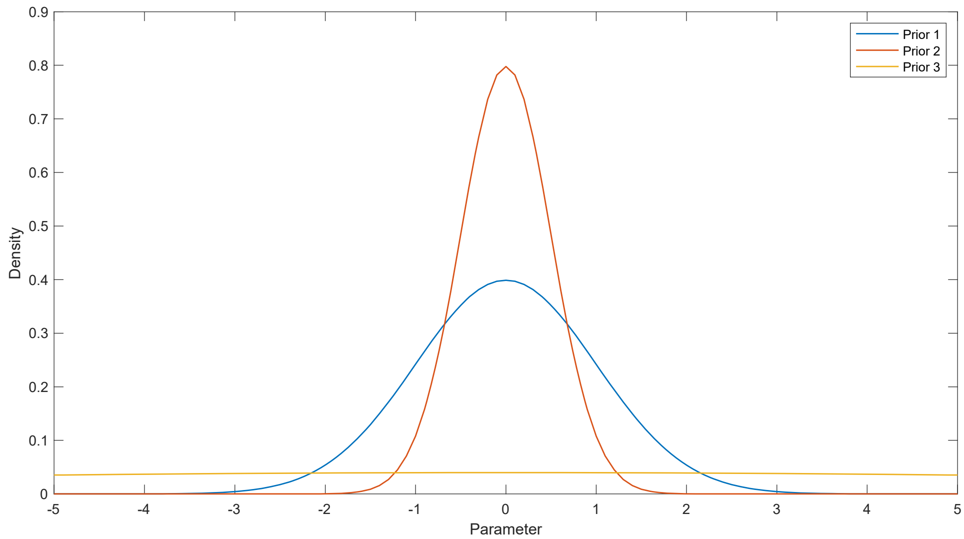

When estimating a finite dataset, Evans and Reichlin (1994) find that the BN trend and cycle estimates are sensitive to sampling errors compared to a univariate setting. The estimation of the large multivariate VAR model will lead to large standard errors, meaning a high variance. In accordance with the bias-variance tradeoff, including more variables or too many variables compared to what is given by the true data generating process (DGP), we will over-parameterize the model (Bjørnland & Thorsrud, 2015, p. 39). This can come from the fact that the model includes many possible sources of information. When introducing a prior, we influence the posterior distribution according to the degree of shrinkage, resulting in a posterior distribution that is slightly different from the likelihood function from the data, which is the classical estimator obtained using OLS. The abovementioned learning mechanism of Bayes' rule is also illustrated in Figure 2, in that the prior distribution is updated using information from the likelihood function.

Figure 2: Illustration of prior distribution, likelihood function and the resulting posterior distribution

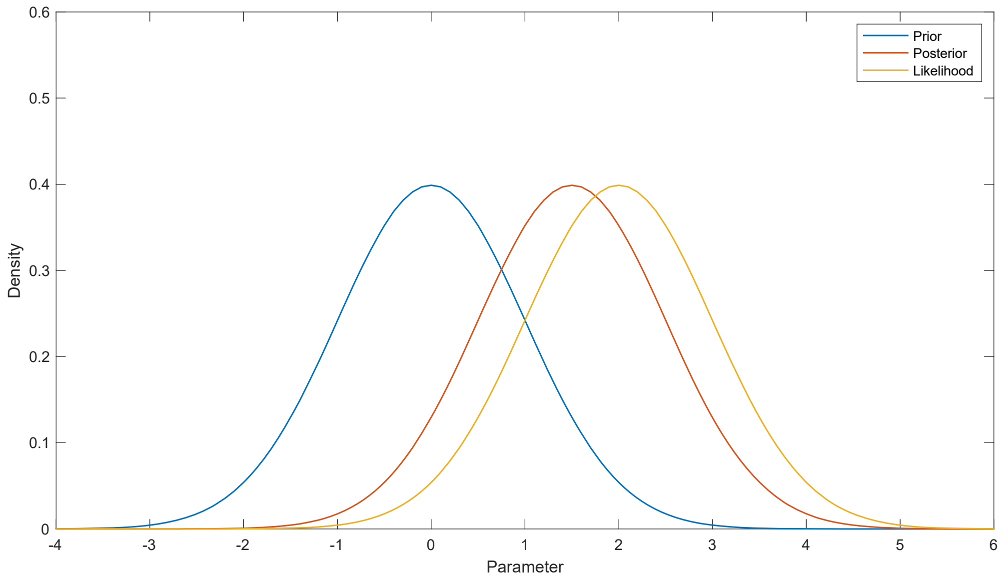

Although the Beveridge-Nelson decomposition contributes to several advantages for filtering time series, we must also consider the disadvantage of different results that can occur when selecting different ARMA models. Selecting the lag length can impose various results, and the main finding from Cochrane (1988) shows that selecting a low-ordered ARIMA model can overestimate the permanent component, namely the trend. Furthermore, the large increase in coefficients that follows from an increase in endogenous variables in BVARs will, without regulation, lead to high parameter uncertainty and the previously mentioned over-parametrization problem, resulting in unreliable estimates (Cross, Hou, & Poon, 2020). The introduction of a prior on the lagged coefficients allows us to overcome this problem by shrinking the parameters with weak signals to zero, albeit at the cost of introducing some bias to the model. Morley and Wong (2020) find that using the Bayesian VAR gives robust results when introducing more lags, but the sensitivity increases in least squares estimation in accordance with Evans and Reichlin (1994).

In order to mitigate the problem of sampling error in larger systems, where the large finite dataset contains many possible sources of information, Morley and Wong (2020) propose a practical way to employ a Minnesota-type shrinkage prior. For simplicity, we can denote the estimated vector of variables as the vector of variables minus the unconditional means, $\Delta\tilde{x}_t \equiv \Delta x_t - \mu$, and we consider the following $VAR(p)$:

$$\begin{aligned}\Delta\tilde{x}_t &= \Phi_1\Delta\tilde{x}_{t-1} + \cdots + \Phi_p\Delta\tilde{x}_{t-p} + e_t \\ &= \begin{bmatrix} \phi_1^{11} & \cdots & \phi_1^{1n} & \phi_2^{11} & \cdots & \phi_2^{1n} & \cdots & \cdots & \phi_p^{1n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \phi_1^{n1} & \cdots & \phi_1^{nn} & \phi_2^{n1} & \cdots & \phi_2^{nn} & \cdots & \cdots & \phi_p^{nn} \end{bmatrix} \begin{bmatrix} \Delta\tilde{x}_{t-1} \\ \Delta\tilde{x}_{t-2} \\ \vdots \\ \Delta\tilde{x}_{t-p} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ \vdots \\ e_{n,t} \end{bmatrix}\end{aligned} \tag{11}$$

where $E(e_t', e_t) = \Sigma$ and $E(e_t', e_{t-i}) = 0$ for all $i > 0$. Demeaning the estimated observation in $\Delta\tilde{x}_t$ is the equivalent to setting a flat prior on the unconditional means, μ, meaning there is no need for an intercept in the model. The Minnesota shrinkage prior is represented in Φ_i, specifically in the slope coefficient ϕ_i^{jk}. The shrinkage prior calibrates a key hyperparameter that minimizes the pseudo out-of-sample forecast error for the target variable as seen in the variance of the slope coefficient. Following Morley and Wong (2020), the prior means and variances of the slope coefficients are defined as:

$$\mathbb{E}[\phi_i^{jk}] = 0 \quad (12)$$

$$var[\phi_i^{jk}] = \begin{cases} \frac{\lambda^2}{i^2}, & j = k \\ \frac{\lambda^2 \sigma_j^2}{i^2 \sigma_k^2}, & otherwise, \end{cases} \quad (13)$$

The shrinkage hyperparameter λ is defined in the variance of the slope coefficient ϕ_i^{jk}, where ϕ is the slope coefficient in the kth variable of the ith lag in the jth equation of the VAR. Equation (12) defines the prior means of the slope coefficients, while Equation (13) defines the variances of the slope coefficients. When the jth equation of the VAR equals the kth variable, we get the variance of a variable's own lag as seen in the upper expression of Equation (13), while the lower expression determines the variance of the cross lags. The variables, σ_j^2 and σ_k^2, are the variances of the residuals from the AR(4) model estimated using OLS in accordance with Banbura et al. (2010) and Koop (2013), and the term i^2 in the denominator implies that coefficients shrink towards zero at longer lags, following the Minnesota prior (Morley & Wong, 2020). The shrinkage hyperparameter $\lambda \rightarrow 0$, thus shrinking the variables following the assumption that they are independent white noise.

Previous studies on forecasting with large BVARs imply that the hyperparameter should be set closer to zero as the number of variables in the model increases, meaning that the overall shrinkage should increase with the size of the model. According to Banbura et al. (2010), Bayesian shrinkage in dynamic systems suggests a better forecast performance and solves the problem of over-parameterization. In contrast to the proposed method by Banbura, which focuses on choosing a hyperparameter to maximize the fit of the entire system, Morley and Wong (2020) propose to instead focus on point forecast accuracy. The reasoning behind this is that as the number of variables increases, the relative weight put on the target variable decreases, and so the BN cycle changes even as extraneous information is added to the model when focusing on the fit of the entire system. The specifics of the point forecast approach are that one conducts "numerical optimization to find the hyperparameter that maximizes the one-step-ahead root mean squared forecast error (RMSFE) for the target variable y_t over an evaluation sample using pseudo real-time estimation based on an expanding window starting with a particular initial fraction of the full sample” (Morley & Wong, 2020, p. 6).

4. Analysis

4.1 Data analysis

In order to determine which variables drive the Norwegian output gap, we have collected a large sample of Norwegian macroeconomic data. We used the variables included in the paper by Morley and Wong (2020) as a baseline and assessed and evaluated the relevance of the American data in comparison to Norwegian data. Additionally, a range of sources on business cycles in Norway have been informative in finding variables especially important for the Norwegian economy, as described in Section 2.1. In total, we retrieved 76 variables that together should span the relevant information of the Norwegian economy with a sufficient length to forecast the output gap. In addition, data availability has been a limiting factor, which has had implications for our collection of macroeconomic variables.

The majority of the collected data has been retrieved from the Federal Reserve Economic Data (FRED) and Statistics Norway (SSB), but in some instances, like the collection of historical rates, the data has been retrieved from Norges Bank. Additionally, exchange rates have been collected from Refinitiv Eikon, and total reserves have been collected from The World Bank. Table A1 in the appendix specifies the sources for each variable. The total dataset of 76 variables is retrieved from as early as possible on a quarterly basis and is seasonally adjusted. The dataset is mainly presented in growth rates and consists of information on our manipulation, code of variable or table, unit of original data, and the source of collection.

A sub-annual frequency makes for an easier analysis of tracking, understanding, and forecasting mechanisms in the economy (Miralles, Lladosa, & Vallés, 2003). Several economic time series relating to our thesis objectives are available at a quarterly frequency, but we necessarily had to transform the data to a quarterly frequency for those of an annual or monthly frequency. Monthly data reported in indices, rates, and levels have been transformed using a method of averaging to obtain a quarterly frequency. In the cases of additive units, the data has been summed over the period to obtain a quarterly frequency. The transformation of annual data into quarterly data can be done in a multitude of ways, and a much-used method is that proposed by Chow and Lin (1971). One caveat of the Chow-Lin method provides a quarterly series that is smoother than what the original series would be (Miralles et al., 2003), and as such, we have aimed to find time series with a higher frequency when possible. However, in order to include all possibly relevant data, we have transformed some variables using this method, keeping in mind that they may not be as reliable as quarterly reported data. In some special cases, we have obtained both annual and quarterly data of the same variables where the older data are only reported in an annual frequency. In these cases, we have applied the Chow-Lin method to the annual data and merged the two series to obtain a longer series. These cases make it apparent that the Chow-Lin provides a smoother time series than what is reported in the more recent quarterly data.

Analyses of economic time series generally use seasonally adjusted data, and for the trend and cycle decomposition, we wish to avoid spurious seasonal variation (Beveridge & Nelson, 1981), thus retrieving seasonally adjusted data when possible. However, several time series were not seasonally adjusted, and for these variables we have used a MATLAB function to deseasonalize the time series. The removal of seasonality in this manner is typically done to more accurately exert the trend and cycle.

Stationarity is necessary to construct the Beveridge-Nelson trend and cycle (Morley & Wong, 2020), meaning we must transform all data to be stationary. When appropriate, we have taken natural logarithms of the data as well as differences to obtain growth rates. However, natural logarithms are not possible for some variables due to negative values. In these cases, as well as for variables already in percent, we have instead differenced the data only when an ADF-test cannot reject a unit root or if a Durbin-Watson statistics test implies autocorrelation in the residuals. The Augmented Dickey-Fuller (ADF) test checks if the data has a unit root and thus tests if the series is a random walk or trend stationary (Thorsrud & Bjørnland, 2015). To ensure stationarity, we have also performed both tests on data transformed to their natural logarithm, differencing those who did not satisfy the conditions of the tests. Table A2 in the Appendix presents the variables that did not initially satisfy the ADF test, along with the transformations to ensure stationarity.

The full dataset contains roughly 150 quarters, starting in the second quarter of 1983. A recursive estimation uses the first 12.5 years (one-third of the sample) and the remaining 25 years in the evaluation of the root mean square forecast error estimation, and the shrinkage hyperparameter that minimizes this RMSFE determines the degree of shrinkage in our various-sized models. This method yields a lambda equal to 0.04, as illustrated in Figure A1 in the Appendix.

4.2 Benchmark model

To establish a benchmark model, containing what we believe to be the most relevant variables, a range of studies have been relevant for our choice of inclusion. According to Husebø and Wilhelmsen (2005) the stylized facts of Norway are fairly similar to those of the US. Consequently, the benchmark model established by Morley and Wong (2020) provides a good motivation for the selection of our benchmark model. Furthermore, we are motivated by Husebø and Wilhelmsen (2005) to include variables that are particularly important to the Norwegian economy. Banbura et al. (2010) finds that a small variable model is feasible and remains robust in estimating the output gap.

Importantly, it has been necessary to consider variables of particular relevance to the Norwegian economy. A paper by Aastveit et al. (2016) investigating the method of forecasting the Norwegian business cycles argues that one should use data for mainland Norway as the measure of economic activity when investigating economic conditions in Norway. The reasoning behind this choice is that the large fluctuations of offshore activity may have small short-term effects on the Norwegian labor market and domestic production. Furthermore, in the short term, the mainland economy is insulated from fluctuating revenue from the petroleum sector, and all revenues are transferred to the sovereign wealth fund with fiscally determined withdrawals each year. In their choice of variables for inclusion in the models, the authors find that particularly important variables in the Norwegian economy are the Brent Blend oil price, employment in mainland Norway, household consumption, private real investments in Norway, exports of traditional goods, and GDP for mainland Norway.

Since the results of Husebø and Wilhelmsen (2005) suggest that the stylized facts of Norway and the US are similar, we are motivated by Morley and Wong (2020) to include real GDP, private final consumption expenditure (PFCE), consumer price index (PCI), M1 and M2 for Norway, producer price index (PPI) for all industrial activities, total manufacturing production, GDP implicit price deflator, total work started: Dwellings/Residential buildings, unemployment rate, persons in the labor force, wages and salaries for mainland Norway, hourly earnings for manufacturing, household disposable income, hours worked, value added per hour worked (productivity), total reserves, and industrial share prices in our benchmark model. We are further motivated by Husebø and Wilhelmsen (2005) to include imports of goods and services, exports of goods and services, gross fixed capital formation (GFCF) for mainland Norway, GFCF dwelling service for households, government final consumption expenditure (GFCE), Brent crude price, EURNOK exchange rate, and the sight deposit rate. The total benchmark model thus consists of 26 macroeconomic explanatory variables of the Norwegian economy.

5. Results

5.1 Estimating the output gap

In order to evaluate our research question of which variables drive the Norwegian output gap, we will begin by estimating the Norwegian output gap using our predetermined benchmark model consisting of 26 variables.

Figure 3: Estimated Norwegian output gap for benchmark BVAR with 90% credible set

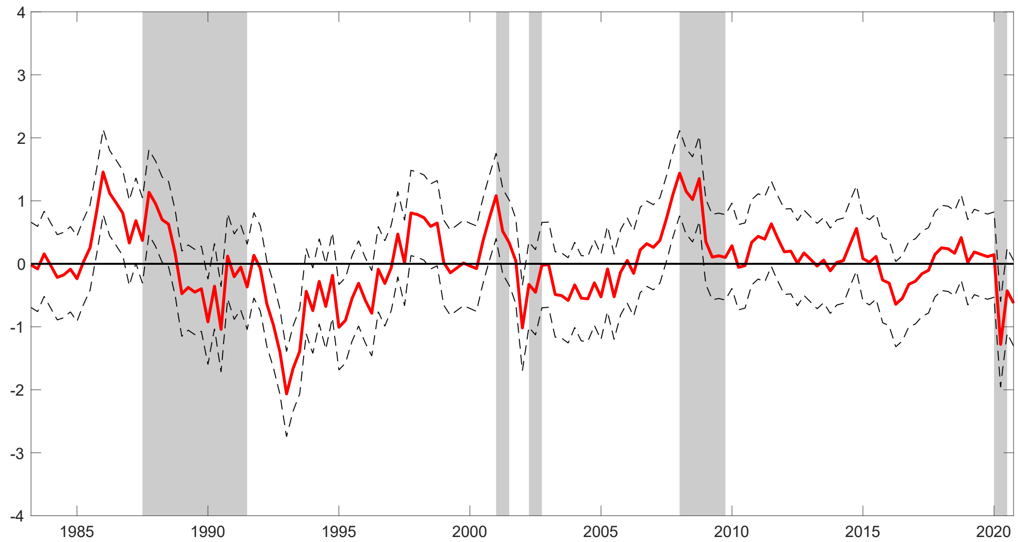

Figure 3 presents the estimated output gap using the Beveridge-Nelson decomposition on the 26-variable benchmark BVAR model with a 90% credible set. The Y-axis corresponds to 100 times natural log deviation from trend, and the output gap is calculated using Equation (6), which represents the analytical solutions of the BN cycle. The shaded areas are the recession periods in Norway, and the cycle corresponds reasonably well to these periods. In contrast to the officially published NBER reference cycles in the US, Norway has no official measure of recessions. For this reason, the choice of marked recessions in our figures is informed by Norges Bank (2016) and Aastveit et al. (2016). Notably, there is no marked recession for the period of 1993, yet we know from Norwegian economic history that the trough in 1993 is in accordance with the data. Starting in 1988, Norway was hit by a significant bank crisis, following considerable loan growth and deregulation in the bank sector. This crisis did not end until 1993, when key aspects of the Norwegian economy finally started showing signs of expansion (Norges Bank, n.d.). As we will show in Figure 6, the 26-variable benchmark model is sufficiently similar to the estimated output gap of the full 76-variable model, thus indicating that it contains all the relevant information for our objective. Naturally, one might wonder how the benchmark BVAR compares to other methods of estimating the output gap, and for this reason, we want to investigate the differences between our benchmark BVAR model, and the estimated output gap based on the BN decomposition for smaller models, as well as the HP filter.

Figure 4: Estimated Norwegian output gap from univariate and multivariate BN decompositions, and using HP filter

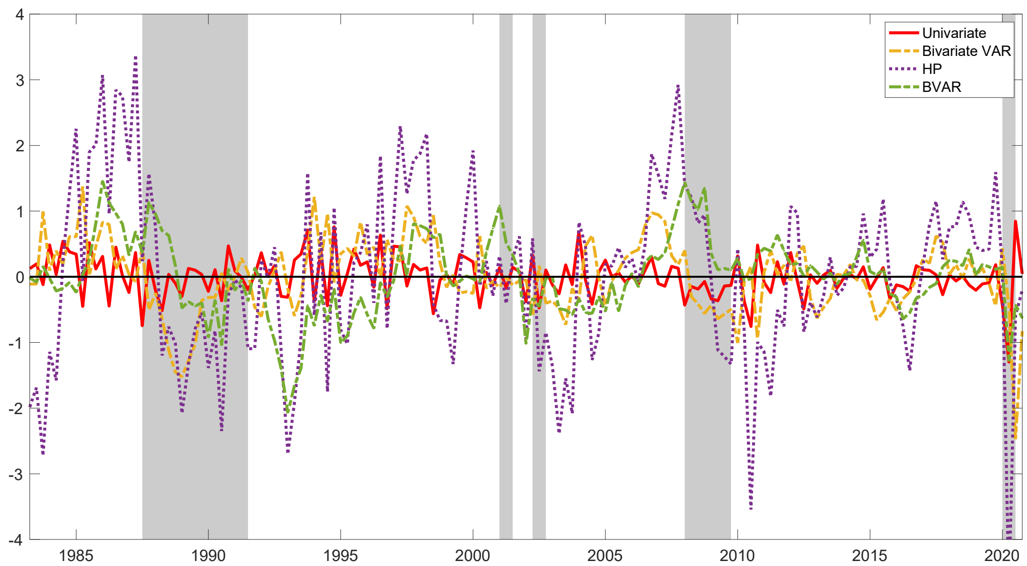

Figure 4 graphically illustrates the different methods, where the univariate model is based on an AR(4) for output growth, while the bivariate VAR model is based on a VAR(4) with output growth and the unemployment rate. The HP model is the estimated output gap using an HP filter with lambda set to 1600, and, finally, the BVAR is our 26-variable benchmark model estimated using the procedure of Morley and Wong (2020). The univariate and bivariate models, both of which are based on the BN decomposition, are obtained using least squares estimation. This estimation is equivalent to the maximum likelihood estimation under the assumption of Normality, which is also satisfied in Bayesian estimation. Consistent with the results from Evans and Reichlin (1994), larger information sets result in a higher proportion of output growth being forecastable and thus ascribed to the cyclical component. This is especially apparent when comparing the univariate and bivariate models to the BVAR model. The difference in results is partly due to the theoretical features differing between univariate and multivariate BN compositions, but part of the difference in results is also due to estimation issues. As we will see below, the unemployment rate is an essential conditioning variable for the output gap, so it is intuitive that the inclusion of this variable in the bivariate VAR will increase the amplitude of the estimated output gap.

Further, the inclusion of more variables will generally increase amplitude in a purely mechanical way, as the relative importance of the cyclical component depends on the size of the information set, and multivariate models can better forecast output growth (Evans & Reichlin, 1994). When using the HP filter to estimate the output gap, we face the issues of spurious cycles and the end-point problem, as can be seen at the beginning and end of our sample in Figure 4. The BN decomposition explicitly takes account of a random walk stochastic trend in the target variable, thus implicitly allowing for correlation between movements in trend and cycle. This is not the case for methods that assume trend stationarity, like the HP filter.

5.2 What drives the output gap

To further analyze which variables drive the Norwegian output gap, we consider the impact of the information set that explains and contributes to the output gap.

Figure 5A: Standard deviation of informational contributions

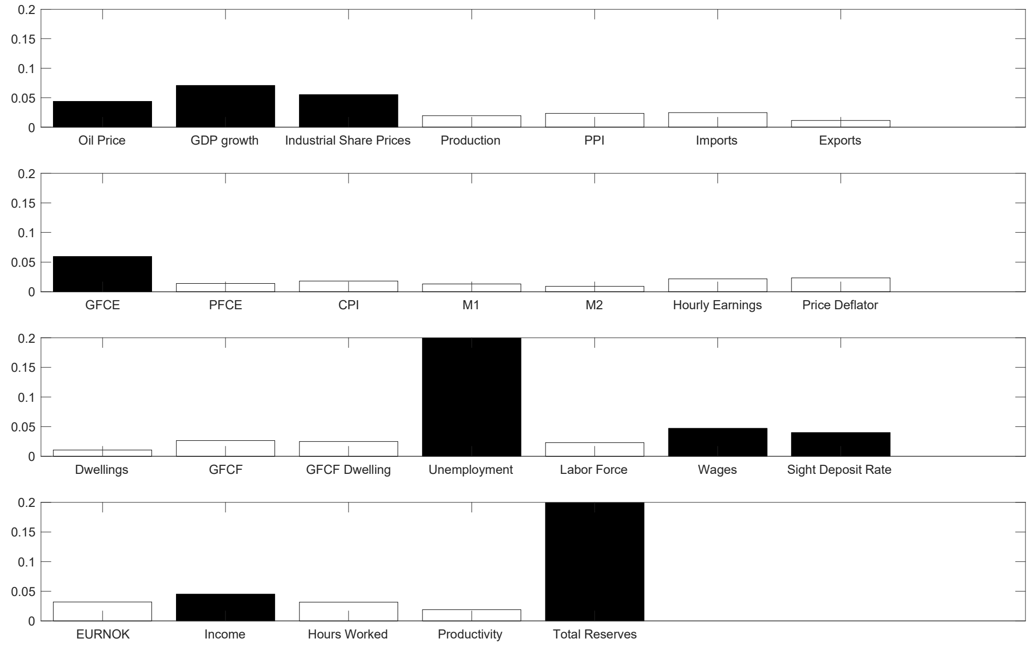

Figure 5A illustrates all 26 variables in the BVAR benchmark model and their standard deviation of informational contribution to the output gap. Shares are calculated using Equation (8) for the contribution to the BN cycle of the output gap and are presented as the standard deviations of each variable that span the most relevant information. The black bars are the nine variables that have the highest standard deviation, including output growth itself, and we find that these are Oil Price, Industrial Share Prices, Government Final Consumption Expenditure, Unemployment rate, Wages and Salaries, Sight Deposit Rate, Household Disposable Income, and Total Reserves. Unemployment rate and Total Reserves have the highest shares of informational contribution and thus drive the output gap the most, while the target variable, output growth, contributes much less than these other two variables. The low output growth contribution also provides some clarity as to why the bivariate and univariate cases differ in Figure 4, where Unemployment rate is included in the bivariate setting.

According to Christiano et al. (1999), the contributed shares of a 20-variable benchmark model cover a sufficient information set for forecasting and estimation purposes in a multivariate BN decomposition using US data. Moreover, Banbura et al. (2010) conclude that a 20-variable BVAR model is feasible and remains robust for the purpose of forecasting and estimating, and so we know that the benchmark model, in theory, contains sufficient information to estimate the output gap. We wish to analyze which variables drive the output gap and therefore try to eliminate and exclude the variables with the lowest standard deviation of informational contribution, following the method proposed by Morley and Wong (2020).

After evaluating and plotting the estimated output gap using only the nine variables of the highest share of informational contribution, we find that the variables marked in black bars in Figure 5A lead to an estimate that deviates significantly from the benchmark BVAR model and, accordingly, the large model as well. Figure A2 in the appendix displays that the magnitude of the estimated output gap is more than ten times as large in percentage change at the most extreme, indicating that the most impactful variables over-estimate the output gap. For this reason, it is evident that we must apply a different approach to obtain the small model, as illustrated in section E of the Appendix.

Figure 5B: Most impactful variables – small model

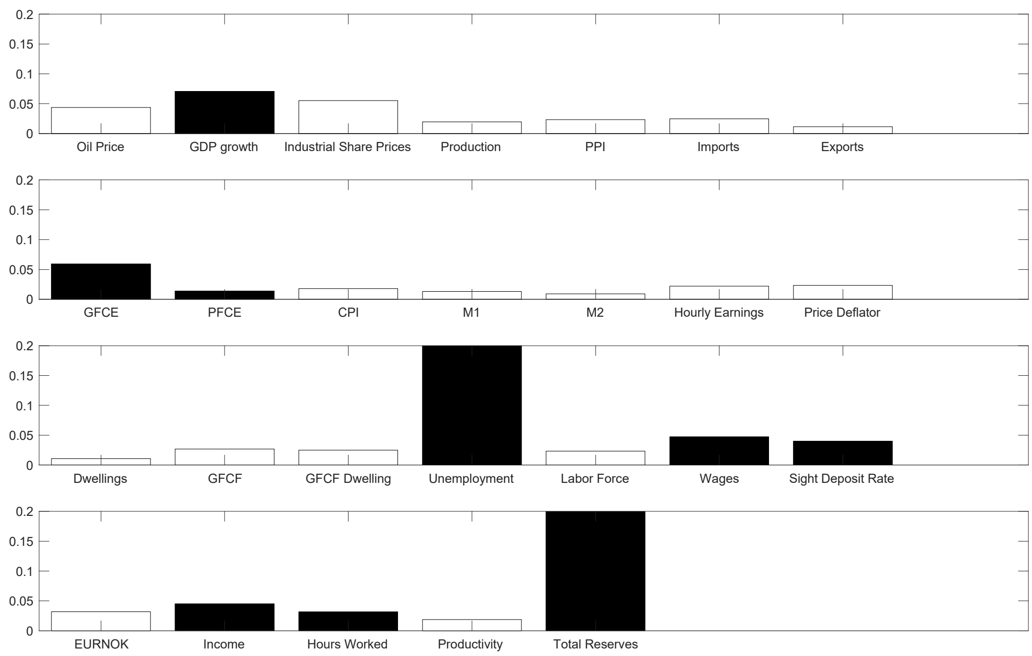

As the first approach proposed by Morley and Wong (2020) does not yield accurate estimates of the output gap, we find it necessary to instead apply their second proposed approach, namely dropping variables one by one, starting from the least impactful, in order to find the most relevant variables for the small model. Furthermore, the exclusion of some variables with low contributions, like the PFCE, significantly altered the cycle, pointing to their significance in forecasting the output gap. Notably, Oil price and Share prices are pro-cyclical, and the exclusion of these components reduces the amplitude compared to the large dataset and the benchmark model. As a result, the final combination set of the small 9-variable BVAR model, as seen in the black bars in Figure 5B, estimates and forecasts the benchmark model more adequately. The accuracy of our established 9-variable model is presented in Figure 6.

Figure 6: Estimated Norwegian output gap for the three different-sized BVAR models.

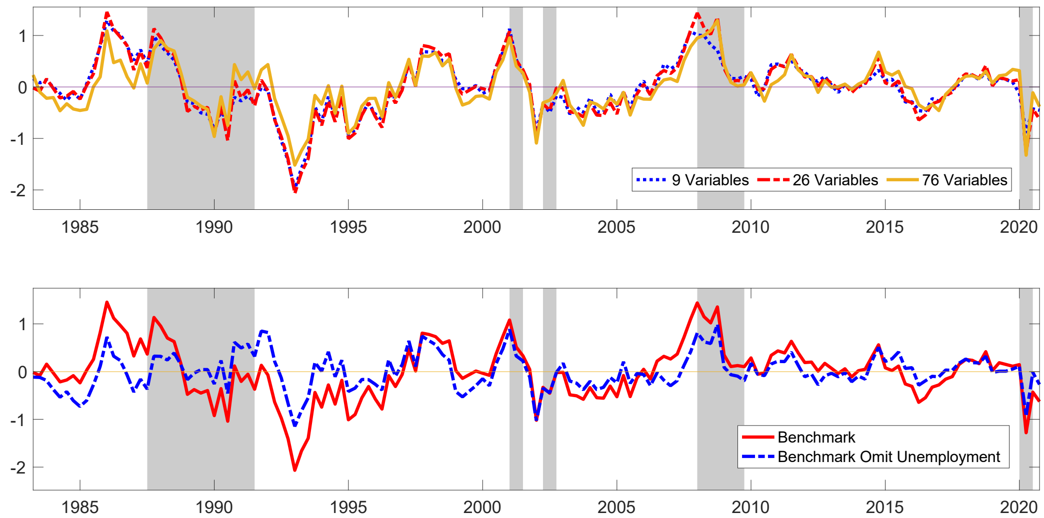

Our medium-sized benchmark model is informed by various sources on business cycles and the Norwegian economy, and accordingly, our small 9-variable model is informed by these sources as well. In order to check the robustness of the two smaller models, we plot the estimated output gap for all three models in the top panel of Figure 6, ensuring that both the benchmark and the small models consistently estimate output gaps that are approximately equal to the large full-sized model. In doing this, we can ensure that we have found which conditioning variables span the relevant information for predicting the fluctuations in output growth.

The benchmark and small models are very similar across the whole sample, except for some more noticeable deviation in the 2008 recession, which suggests that the approach we have chosen to select the conditioning variables in the small model works well in practice. The difference between the large model and the two smaller models is more apparent in the first half of the sample and may, to some extent, reflect the structural changes in monetary policy in Norway. Norges Bank started conducting inflation targeting in 2001, and in the decades prior, there was first a target for low rates, followed by a target for stable exchange rates. Both of these targets eventually led to large fluctuations and periodically high inflation. The Norwegian government introduced a new mandate for the monetary policy authority in 2001 to maintain low and stable inflation, thus marking a shift to inflation targeting in Norway (Bergo, 2003). The targeting of inflation also helps stabilize employment, and interestingly, we can observe that the estimated output gap has lower fluctuations pre-2001 compared to the benchmark when we omit unemployment in the bottom panel of Figure 6. This further points to the fact that the Bayesian VAR approach can be quite sensitive to the omission of variables, especially when the variables contain highly relevant information, suggesting that it is the relevancy of the variables included that matters most for the accuracy of the estimated output gap, and not only the size of the information set.

6. Discussion

The objective of our thesis is twofold: We wish to estimate the output gap using a multivariate Beveridge Nelson decomposition and establish which of the conditioning variables are most influential in generating output gap fluctuations. The main result of our first objective is presented in Figure 3, where we have plotted the estimated output gap using our benchmark 26-variable model with its corresponding 90% credible set. An advantage of using a Bayesian approach is that we can interpret this credible set as the certainty with which the output gap falls within these bounds. As the trend and cycle are inherently unobservable, this probabilistic interpretation gives us a more intuitive understanding of the estimate. A key aspect of evaluating whether an estimated output gap is credible is to check that the fluctuations are in line with the recessions of an economy. Unfortunately, there is no official record of recession periods in Norway, but as our marked dates for recessions are informed by Norges Bank (2016) and Aastveit et al. (2016), we feel confident that this is well enough informed to evaluate the estimate of the output gap. We see from the plot that the multivariate BN decomposition provides an estimate that accounts for all marked recessions in Norway over the sample period, and more generally, it is in line with the movements of other estimates of the Norwegian output gap like the one produced by Statistics Norway (2014). Their particular estimate is produced using an HP filter, although it is essential to bear in mind that Statistics Norway uses a lambda of 40 000, and the movements of the business cycles look more similar to the benchmark model in Figure 3 than the HP cycle presented in Figure 4 with a lambda of 1600. Statistics Norway justify their choice of lambda by the fact that the Norwegian economy is relatively small, and so random fluctuations might have too much of an influence even on aggregate data (Statistics Norway, 2018, p. 16). Furthermore, the estimate using our benchmark model is robust compared to the estimate of the entire sample with 76 variables as presented in Figure 6, where we would see that the estimates of the different sized models are virtually identical had we included the 90% credible interval.

The second objective of our thesis is to establish which conditioning variables have the most informational contribution in accounting for the output gap fluctuations. Figure 5A highlights the variables with the highest share of informational contribution, with unemployment and total reserves explaining approximately 30%

of the fluctuations in output growth each, and they do so to a much larger extent than output growth itself. The following most informational variables are, in descending order, GDP growth, government final consumption expenditure (GFCE), industrial share prices, wages, income, Brent blend oil price, and the sight deposit rate. The exact values of the informational contributions of the various benchmark variables are presented in Table A3 in the appendix. Following the results of Morley and Wong (2020), we expected to see that the variables with the highest informational contribution would yield an estimate of the output gap that was sufficiently similar to the benchmark model. Figure A2 in the appendix makes it evident that this is not the case with Norwegian data. The amplitude of the fluctuations is vastly larger for most of the time series, but the instant we omit oil price, we see that the amplitude falls dramatically in Figure A3, albeit still with an overestimation problem. This could point to the fact that the mainland data for Norway is to some extent isolated from the more volatile offshore activity and oil-related variables. While the insight of Figure 5A is valuable, we wish to establish the minimum number of variables that can be used to accurately estimate the output gap. As our initial approach of using the most informational variables did not yield satisfactory results in this area, we instead opted for the approach of dropping and adding variables from the benchmark model until we reached the small model, consisting of the highlighted variables in Figure 5B. The procedure for obtaining the small model is illustrated in section E of the Appendix. Interestingly, we omit the Brent blend oil price and Industrial share prices to obtain this model and instead include hours worked and private final consumption expenditure (PFCE).

Following the results presented in Figure 5B, we found that the variables most influential in generating output gap fluctuations are dominated by the unemployment rate and total reserves. Although the findings on the impact of unemployment are foreseen, with stylized facts that the unemployment rate is countercyclical and leads the Beveridge-Nelson cycle by one quarter (Bjørnland, 2000), we found the results of total reserves quite astonishing. The World Bank defines total reserves as "the holdings of monetary gold, special drawing rights, reserves of IMF members held by the IMF, and holdings of foreign exchanges under the control of monetary authorities" (The World Bank Group, n.d.). Accordingly, Norges Bank reports its management of foreign exchange reserves, which is described as both a fixed income portfolio and an equity portfolio. A petroleum buffer portfolio is also included in the foreign exchange reserves and is used to convert currency for the Government Pension Fund Global (GPFG) (Norges Bank, n.d.). The correlation between both the fixed income and equity market with GDP might be an explanatory factor to the high impact on the output gap. Also, some cross-country variation factors may be incorporated in the total reserve variable where the performance of both the fixed income and equity market is highly dependent on international presentation.

While our analysis makes it evident that the use of large Bayesian vector autoregressions can provide a good estimate of the output gap as well as answer questions regarding the importance and contributions of macroeconomic variables, there are still some extensions that could yield even more informed estimates and results. One such extension could be to include the GDP growth from the main trading partners of Norway, as the country is a small open economy that can, to some extent, be highly influenced by international economic movements. Additionally, one could include international macroeconomic variables that may be explanatory, such as unemployment rates of key trading partners, CPIs, or key interest rates. The inclusion of exchange rates can partly proxy international relationships, but there may be even more informational value in including these variables directly. The model can also be extended to answer further questions about what drives other key variables of interest, such as inflation or unemployment. One can easily change the target variable of interest, and the model is thus a flexible tool for assessing relevant policy questions. We have not taken data revision into account in our analysis, which is a problem that can often be disregarded in estimations of business cycles. Accounting for this problem would therefore be a valuable extension.

Although our analysis focuses on estimating the output gap and establishing the most influential variables, another interesting extension could be to find the stylized facts using this method and compare them with those of Bjørnland (2000), who uses a range of estimation methods such as the HP filer and a univariate BN decomposition. Furthermore, one could extend our analysis to investigate the peaks and troughs of the business cycle as is done in Aastveit et al. (2016), allowing for a new method of establishing recession periods in Norway. Nonetheless, we leave all the above-mentioned expansions and modifications of the model to future research.

7. Conclusion

In this paper, we have applied a multivariate Beveridge-Nelson decomposition with Bayesian shrinkage to obtain an estimate of the Norwegian output gap as our first objective. We have further used the Bayesian VAR approach to investigate our second objective: which of the conditioning variables are most influential in generating output gap fluctuations. After establishing a full model containing 76 macroeconomic explanatory variables, we have found that a medium-sized benchmark model comprising 26 key macroeconomic variables effectively estimates the output gap with four lags, accounting for all marked recession periods over the sample period. We have utilized the standard deviation of informational contribution to establish the variables that have the highest power in explaining the output gap fluctuations, as well as a minimum set of 9 variables that can accurately estimate the output gap in accordance with the benchmark model and the full 76-variable model. Our findings are robust to the size of the models used to estimate the output gap, as all three models of different sizes are virtually identical. Using the standard deviation of informational contribution, we find that Unemployment and Total Reserves have the highest share of contribution to the output gap fluctuations. The advantage of using this proposed method is that it allows us to utilize large multivariate datasets to accurately estimate the trend and cycle while mitigating detrending issues, such as the end-of-sample problem. Furthermore, there are several ways to extend the model, for instance, by introducing cross-country variation factors or changing the target variable in order to analyze the driving factors of inflation or unemployment.

References

Aastveit, K. A., Jore, A. S., & Ravazzolo, F. (2016). Identification and real-time forecasting of Norwegian business cycles. *International Journal of Forecasting*, *32*(2), 283-292. https://doi.org/10.1016/j.ijforecast.2015.06.006

Banbura, M., Giannone, D., & Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25, 71–92. https://doi.org/10.1002/jae.1137

Bergo, J. (2003, February 7). *Litt om pengepolitikken og samspillet i den økonomiske politikken* [Press release]. https://www.norges-bank.no/aktuelt/nyheter-og-hendelser/Foredrag-og-taler/2003/2003-02-07/

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". *Journal of Monetary Economics*, 7, 151–174. https://doi.org/10.1016/0304-3932(81)90040-4

Bjørnland, H.C. (2000) Detrending methods and stylized facts of business cycles in Norway–an international comparison. *Empirical Economics,* 25(3), 369-392. https://doi.org/10.1007/s001810000023

Bjørnland, H.C., Brubakk, L., & Jore, A.S. (2005). The output gap in Norway- A comparison of different methods. *Norges Bank Economic Bulletin*, 5(2), 90-100. http://hdl.handle.net/11250/2504375

Bjørnland, H. C., & Thorsrud, L. A. (2015). *Applied time series for macroeconomics*. Gyldendal akademisk.

Bry, G., & Boschan, C. (1971). Cyclical Analysis of Time Series: Selected Procedures and Computer Programs. *National Bureau of Economic Research, Inc*, https://EconPapers.repec.org/RePEc:nbr:nberbk:bry_71-1.

Chow, G. C., & Lin, A. (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *The Review of Economics and Statistics,* 53(4), 372–375. https://doi.org/10.2307/1928739

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics, 1*, 65-148. https://doi.org/10.1016/S1574-0048(99)01005-8

Cochrane, J. H. (1988). How Big Is the Random Walk in GNP? *Journal of Political Economy*, *96*(5), 893–920. https://doi.org/10.1086/261569

Cross, J. L., Hou, C., & Poon, A. (2020). Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity. *International Journal of Forecasting, 36*(3)*,* 899-915. https://doi.org/10.1016/j.ijforecast.2019.10.002

Evans, G., & Reichlin, L. (1994). Information, forecasts, and measurement of the business cycle. *Journal of Monetary Economics*, *33*(2), 233–254. https://doi.org/10.1016/0304-3932(94)90002-7

Hamilton, J. D. (2018). Why You Should Never Use the Hodrick-Prescott Filter. *The Review of Economics and Statistics, 100(5), 831-843.* https://doi.org/10.1162/REST_a_00706

Husebø, T. A., & Wilhelmsen, B. R. (2005). *Norwegian business cycles 1982-2003* (No. 2005/2). Staff Memo. http://hdl.handle.net/11250/2508134

Koop, G. (2003). *Bayesian econometrics.* John Wiley & Sons.

Koop, G. M. (2013). Forecasting with Medium and Large Bayesian VARs. *Journal of Applied Econometrics*, *28*(2), 177–203. https://doi.org/10.1002/jae.1270

Litterman R. (1986). Forecasting With Bayesian vector autoregressions: five years of experience. *Journal of Business and Economic Statistics,* 4, 25–38. https://doi.org/10.1080/07350015.1986.10509491

Miralles, J. M. P., Lladosa, L. E. V. & Vallés, R. E. On the performance of the Chow-Lin procedure for quarterly interpolation of annual data: Some Monte-Carlo analysis. *Spanish Economic Review* 5, 291–305 (2003). https://doi.org/10.1007/s10108-003-0068-9

Morley, J. C. (2002). A state-space approach to calculating the Beveridge–Nelson decomposition. *Economics Letters*, 75, 123–127. https://doi.org/10.1016/S0165-1765(01)00581-X

Morley, J., & Wong, B. (2020). Estimating and accounting for the output gap with large Bayesian vector autoregressions. *Journal of Applied Econometrics*, 35(1), 1–18. https://doi.org/10.1002/jae.2733

Nelson C. R., & Kang H. (1981) Spurious periodicity in inappropriately detrended time series. *Econometrica: Journal of the Econometric Society*, 741-751. https://doi.org/10.2307/1911520

Norges Bank. (2016, March). *Pengepolitisk Rapport med vurdering av finansiell stabilitet 1/16* (No. 1/2016). 07 Media AS. https://www.norges-bank.no/contentassets/28a7f21533e74d0aa13e3ef888cae008/ppr_1_16.pdf?v=03/09/2017123501&ft=.pdf

Norges Bank. (2021, February 25). *Foreign exchange reserves*. Retrieved June 3, 2022, from https://www.norges-bank.no/en/topics/liquidity-and-markets/foreign-exchange-reserves/

Norges Bank. (n.d.). *Bankkrisen 1988–1993*. Retrieved May 27, 2022, from https://www.norges-bank.no/tema/Om-Norges-Bank/historien/Pengepolitikk-finansiell-stabilitet-og-kapitalforvaltning/Finansiell-stabilitet/1990-tallet/

Statistics Norway. (2014, September). *Economic Survey* (No. 3–2014). https://www.ssb.no/en/nasjonalregnskap-og-konjunkturer/artikler-og-publikasjoner/_attachment/196622?_ts=148699c5470

Statistics Norway. (2018, December). *Economic Survey* (No. 4-2018). https://www.ssb.no/en/nasjonalregnskap-og-konjunkturer/artikler-og-publikasjoner/_attachment/372056?_ts=167ad027a40

The World Bank Group. (n.d.). *Total reserves (includes gold, current US$) - Norway | Data*. Retrieved March 21, 2022, from https://data.worldbank.org/indicator/FI.RES.TOTL.CD?locations=NO

Woźniak, T. (2016). Bayesian vector autoregressions. *Australian Economic Review*, 49, 365–380. https://doi.org/10.1111/1467-8462.12179

Appendix

A. Data

Table A1 presents the 76 variables of the full dataset, with the source of the data in the last column. For variables sourced from SSB, the table is presented in column three. For variables sources from FRED, the corresponding FRED-code is presented in column 2. The remaining variables have been given their own individual codes. The fourth column, "Adjust", presents the manipulation applied to each variable, where "Log" is the natural logarithm of the variable, and "Δ^i" means the variable have been differenced *i* times. The variables marked with an "x" in the BM column means the variable is in the 26-variable benchmark model. The code used for obtaining growth rates have resulted in variables having a negative sign where intuition would have them with a positive sign, and vice versa. For our analysis, we have simply taken the inverse of the graphs to ensure that all estimates of the output gap are converted to the correct sign for interpretation.

Table A1: Variable details

Variable name	Code	Table (SSB)	Adjust	BM	Source
Brent Crude price- Oil Prices	LCOc1		Δ	x	Refinitiv
Real Gross Domestic Product, 3 Decimal	CLVMNACSCAB1GQNO		Log, Δ,	x	FRED
Total Retail Trade	NORSARTQISMEI		Log, Δ		FRED
Balance of payments BPM6: Capital account	NORB6CATT00NCCUQ				FRED
US$ exchange rate for Norway	NORCCUSMA02GYQ				FRED
Total Industrial Share Prices	SPINTT01NOQ661N		Log, Δ	x	FRED
Credit to Private Non-Financial Sector by Banks	QNOPBMUSDA		Log, Δ		FRED
Credit to Private Non-Financial Sector by Banks	QNOPBM770A		Log, Δ		FRED
Total Credit to Private Non-Financial Sector	CRDQNOAPUBIS		Log, Δ		FRED
Residential Property Prices	QNON628BIS		Log, Δ		FRED
Total Manufacturing Production	PRMNTO01NOQ657S			x	FRED
Production: Total manufacturing	NORPRMNTO01GYSAQ				FRED
Gross Domestic Product by Expenditure: Exports of Goods and Services for Norway	NAEXKP06NOQ652S		Log, Δ		FRED
Producer Prices Index: Total Industrial Activities	PIEATI01NOQ661N		Log, Δ	x	FRED
Gross Domestic Product by Expenditure: Gross Fixed Capital Formation	NAEXKP04NOQ189S		Log, Δ		FRED
Gross Domestic Product by Expenditure: Imports of Goods and Services	NAEXKP07NOQ657S				FRED
Gross Domestic Product by Expenditure: Private Final Consumption Expenditure	NAEXKP02NOQ189S		Log, Δ		FRED
Imports of Goods and Services	NORIMPORTQDSMEI		Log, Δ	x	FRED

Exports of Goods and Services	NOREXPORTQDSMEI		Log, Δ	x	FRED
Government Final Consumption Expenditure	NORGFCEQDSMEI		Log, Δ	x	FRED
Private final consumption expenditure	NORPFCEQDSMEI		Log, Δ	x	FRED
Consumer Price Index: All Items for Norway	NORCPIALLQINMEI		Log, Δ	x	FRED
M1 for Norway	MANMM101NOQ189N		Log, Δ	x	FRED
M2 for Norway	M2NO		Log, Δ	x	FRED
Real Effective Exchange Rates Based on Manufacturing Consumer Price Index	CCRETT01NOQ661N		Log, Δ		FRED
Net Trade: Value Goods	XTNTVA01NOQ664S		Δ		FRED
Hourly Earnings: Manufacturing	NORHOUREAQISMEI		Log, Δ	x	FRED
GDP Implicit Price Deflator	NORGDPDEFQISMEI		Log, Δ	x	FRED
Narrow Effective Exchange Rate	NNNOBIS		Log, Δ		FRED
Work Started: Construction: Dwellings / Residential buildings: Total for Norway	NORWSCNDW01GPSAM			x	FRED
Final Consumption expenditure of local government	09190x1	09190	Log, Δ		SSB
Final consumption expenditure of central government	09190x2	09190	Log, Δ		SSB
Household final consumption expenditure: Goods	9173x1	9173	Log, Δ		SSB
Household final consumption expenditure: Services	9173x2	9173	Log, Δ		SSB
Final consumption expenditure of households: Non-durable goods	09173x3	09173	Log, Δ		SSB
Gross fixed capital formation (GFCF)	09190x3	09190	Log, Δ		SSB
GFCF Mainland Norway	09190x4	09190	Log, Δ	x	SSB
Dwelling service (households): GFCF	09190x5	09190	Log, Δ	x	SSB
GFCF Oil and gas extraction including services	09183x1	09183	Log, Δ		SSB
Production account and income generation, Total Industry	09171x1	09171	Log, Δ		SSB
Production account and income generation, Mainland Norway	09171x2	09171	Log, Δ		SSB
Unemployment rate: 15-74 years	08518x1	08518		x	SSB
Unemployed: 15-24 years	08518x2	08518	Δ		SSB
Persons in the labor force: 15-74 years	08518x3	08518	Δ	x	SSB
Household final consumption expenditure: Goods (price indices)	09190x6	09190	Log, Δ		SSB
Household final consumption expenditure: Services (price indices)	09109x7	09190	Log, Δ		SSB
Wages and salaries, mainland Norway (SPLIT series)	09175x1	09175	Log, Δ	x	SSB
Sum of Dwelling started (SPLIT series)	11006x10996	11006/10996	Log, Δ		SSB
Sight deposit rate (styringsrente, nominell)	SRATE		Δ	x	Norges Bank
Consumer price index: Transport	03013x1	03013	Log, Δ		SSB
Consumer price index: Health	03013x2	03013	Log, Δ		SSB
Consumer price index: Housing, water, electricity, gas and other fuels	03013x3	03013	Log, Δ		SSB

Consumer price index: Miscellaneous goods and services	03013x4	03013	Log, Δ		SSB
Consumer price index: Food and non-alcoholic beverages	03013x5	03013	Log, Δ		SSB
Consumer price index: Alcoholic beverages and tobacco	03013x6	03013	Log, Δ		SSB
Consumer price index: Communications	03013x7	03013	Log, Δ		SSB
Consumer price index: Recreation and culture	03013x8	03013	Log, Δ^2		SSB
Consumer price index: Education	03013x9	03013	Log, Δ		SSB
Consumer price index: Restaurants and hotels	03013x10	03013	Log, Δ		SSB
Consumer price index: Furnishings, household equipment and routine maintenance	03013x11	03013	Log, Δ		SSB
EURNOK	EURNOK		Δ	x	Refinitiv
USDNOK	USDNOK		Δ		Refinitiv
GBPNOK	GBPNOK		Δ		Refinitiv
Household disposable income	10799x10	10799	Log, Δ	x	SSB
Employed persons. Employees and self-employed	09175x10	09175	Log, Δ		SSB
Full time equivalent employment. Employees and self-employed	09175x11	09175	Log, Δ		SSB
Total hours worked for employees and self-employed	09175x12	09175	Log, Δ	x	SSB
Compensation of employees	09175x13	09175	Log, Δ		SSB
Value added at basic values per hour worked	09176x1	09176	Δ	x	SSB
Wages and salaries per hour	09176x2	09176	Δ		SSB
Full time equivalent employment	09176x3	09176	Log, Δ		SSB
Gross Domestic Product Per Capita for Norway	NORPFCEQDSMEI		Log, Δ		FRED
Total reserves incl. Gold	FI.RES.TOTL.CD		Log, Δ	x	World Bank
Total reserves ex. Gold	FI.RES.XGLD.CD		Log, Δ		World Bank
Direct purchases by non-residents	09190x8	09190	Δ		SSB
Share prices Index	SPIOS		Log, Δ		Oslo Børs

Table A2 presents the variables of our full dataset which did not satisfy the stationarity check. All variables of which an ADF test of five lags could not reject a unit root at the 5% critical value level were deemed to not be stationary, and Table A2 presents our adjustments to ensure stationarity. In the second column, >10% crit value means that the ADF test cannot reject a unit root at the 10% critical value level. If the ADF test could reject a unit root at 10% for only one or a few lags, the specific lag is specified in the second column as well. The third column describes the transformation of the data to ensure stationarity, where "Δ" means that the variable has been differenced, and log refers to taking the natural logarithm of the data. The fourth column then describes the level at which a unit root can be rejected after the transformations of the data, where all variables are stationary following the adjustment.

Table A2: Stationarity test

Variable	Rejected level	Adjust	Rejected level
Brent Crude price- Oil Prices	>10% crit value	Δ	1% crit value
Real Gross Domestic Product, 3 Decimal	>10% crit value	Log, Δ	1% crit value
M1 for Norway	>10% crit value	Log, Δ	1% crit value
M2 for Norway	>10% crit value	Log, Δ	1% crit value
Net Trade: Value Goods	>10% crit value	Δ	1% crit value
Unemployed: 15-24 years	10% on 2 lags, >10% on 3 lags	Δ	1% crit value
Persons in the labor force: 15-74 years	>10% crit value	Δ	1% crit value
Consumer price index: Recreation and culture	>10% at lag 3	Δ	1% crit value
EURNOK	>10% crit value	Δ	1% crit value
USDNOK	>10% crit value	Δ	1% crit value
GBPNOK	>10% at lag 1 and 2, 10% at lag 3	Δ	1% crit value
Value added at basic values per hour worked	10% at lag 3, >10% at lag 4	Δ	1% crit value
Wages and salaries per hour	10% crit value	Δ	1% crit value
Direct purchases by non-residents	>10% crit value	Δ	1% crit value

B. Derivation

In order to obtain Equation (5), we first start by inserting from Equation (3) into Equation (4):

$$(\Delta X_t - \mu) = F(\Delta X_{t-1} - \mu) + He_t \quad (3)$$

$$\tau_t = X_t + \lim_{j\to\infty} \sum_{j=1}^{\infty} E\left(\Delta X_{t+j} - \mu\right), \qquad (4)$$

This yields the following expression:

$$\tau_t = X_t + \mathrm{E_t} \sum_{j=1}^{\infty} (F\left(\Delta X_{t+j-1} - \mu\right) + He_t), \qquad (4.b)$$

Since we have that $He_t \sim N(0, \Sigma)$, the equation implies the following:

$$\tau_t = X_t + \sum_{j=1}^{\infty} F\left(\Delta X_{t+j-1} - \mu\right), \qquad (4.c)$$

Using the infinite sum of geometric series/sequence $\sum_{t=0}^{\infty} F^t = \frac{1}{1-F} = (1 - F)^{-1}$, where $|F| < 1$, we arrive at the following expression in Equation (5), with the term $F(I - F)^{-1}$ accounting for the fact that our geometric sum begins in period $j = 1$:

$$\tau_t = X_t + F(I - F)^{-1}(\Delta X_t - \mu) \qquad (5)$$

C. Informational contribution

Table A3 presents the shares of standard deviation of informational contribution as shown graphically in Figure 5A. Units are standard deviations, and the contributions to the estimated output gap are calculated using Equation (8).

Table A3: Shares of standard deviation of informational contributions

Variable	Shares
Unemployment	**0.3660**
Total Reserves	**0.2985**
GDP growth	**0.0705**
GFCE	**0.0593**
Industrial Share Prices	**0.0552**
Wages	**0.0471**
Income	**0.0451**
Oil Price	**0.0437**
Sight Deposit Rate	**0.0399**
EURNOK	**0.0319**
Hours Worked	**0.0316**
GFCF	**0.0266**
GFCF Dwelling	**0.0249**
Imports	**0.0246**
PPI	**0.0234**
Price Deflator	**0.0233**
Labor Force	**0.0231**
Hourly Earnings	**0.0218**
Production	**0.0195**
Productivity	**0.0188**
CPI	**0.0179**
PFCE	**0.0139**
M1	**0.0132**
Exports	**0.0114**
Dwellings	**0.0107**
M2	**0.0091**
	1.371

Table A4 presents the shares of standard deviation of informational contribution for the variables that constitute the small 9-variable model. The procedure of establishing the small model is described in Figure A2 through A6.

Table A4: Shares of contribution on variables in the 9-variable model

Variable	**Shares**
Unemployment	**0.3660**
Total Reserves	**0.2985**
GDP growth	**0.0705**
GFCE	**0.0593**
Wages	**0.0471**
Income	**0.0451**
Sight Deposit Rate	**0.0399**
Hours Worked	**0.0316**
PFCE	**0.0139**
	0.9720

D. Calculation of lambda

Figure A1: One-step-ahead pseudo out-of-sample root mean square forecast error for various-sized BVARs

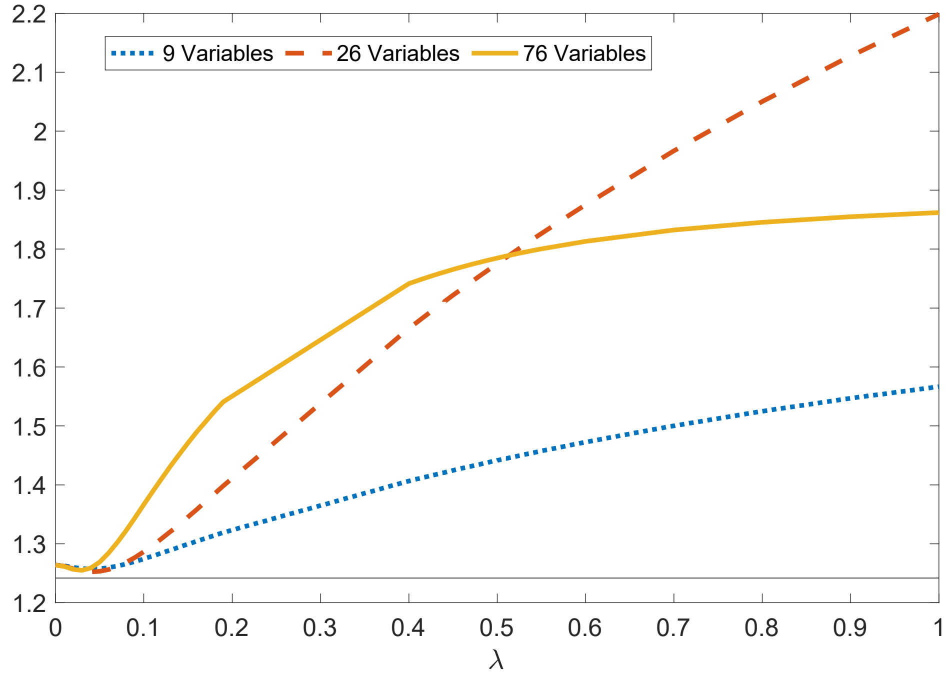

Figure A1 illustrates the one-step-ahead pseudo out-of-sample RMSFE for the 9-variable, 26-variable, and 76-variable BVAR models. The line of the shrinkage parameter λ is set to minimize the RMSFE and decreases when more variables are added to the model, meaning as more variables are added to the model, more shrinkage is added to the model through a lower λ. The estimation of the shrinkage hyperparameter λ is set using the approach presented in Morley and Wong (2020). Our data contains roughly 150 quarters, resulting in using the first 12.5 years (one-third of the sample) in the recursive estimation and the remaining 25 years in the evaluation of the root mean square forecast error estimation, resulting in $\lambda \approx 0.04$.

E. Selection of variables in the small model

The smallest model with nine variables with the highest shares of standard deviation of informational contributions: Unemployment, Total Reserves, GDP growth, GFCE, Industrial Share Prices, Wages, Income, Oil Price, Sight Deposit Rate.

Figure A2 displays the estimates output gap using the 9-variables with the highest share of informational contribution. These variables are as follows: Unemployment, Total Reserves, GDP growth, GFCE, Industrial Share Prices, Wages, Income, Oil Price, Sight Deposit Rate.

Figure A2: Estimated Norwegian output gap for 3 different sized BVAR models

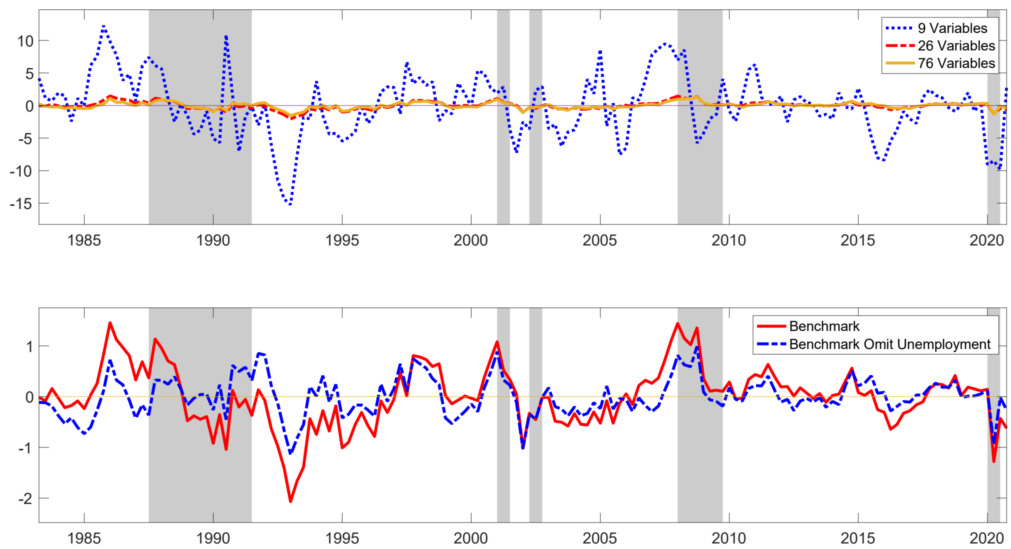

Following from Figure A2 using the highest contribution shares, we then omit the Oil Price component as this variable are the one of the lowest contributors to the output gap. Figure A3 represent the small VAR-sized model compared to the benchmark-model and the full-sized 76 variable model, and we see a significantly change in the amplitude of the cycle.

Figure A3: Estimated Norwegian output gap for 3 different sized BVAR models

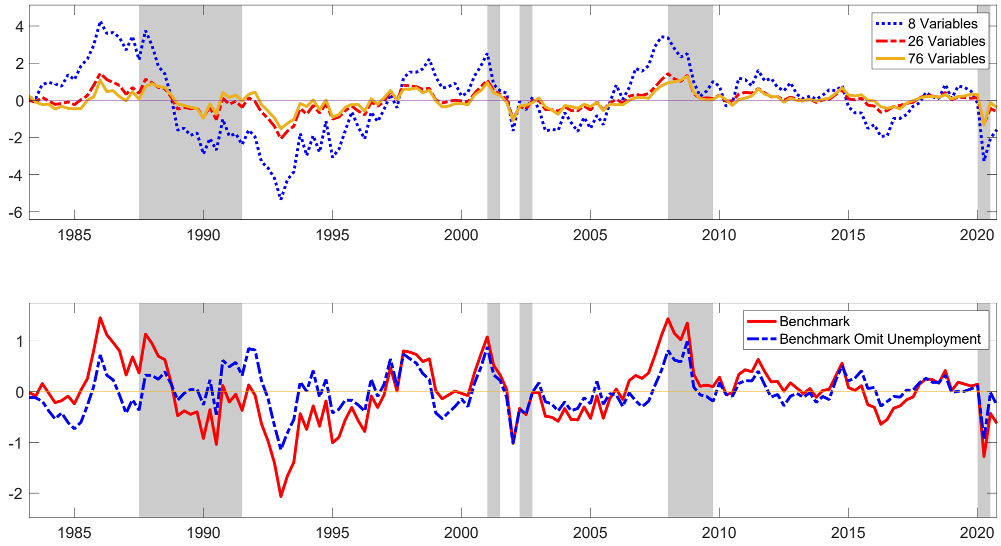

The selection continues from Figure A3 where we further omit the variable Industrial Share Prices. A representation is shown in Figure A4.

Figure A4: Estimated Norwegian output gap for 3 different sized BVAR models

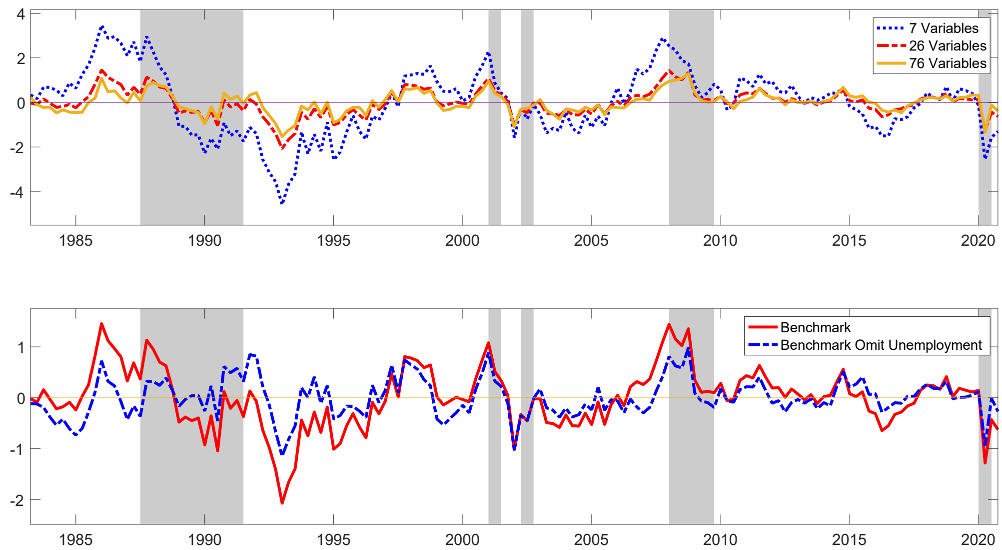

The small model is then modified by including Hours Worked after adding and dropping variables. Figure A5 shows that the amplitude of the model converges towards the benchmark model and the full-sized variable model.

Figure A5: Estimated Norwegian output gap for 3 different sized BVAR models

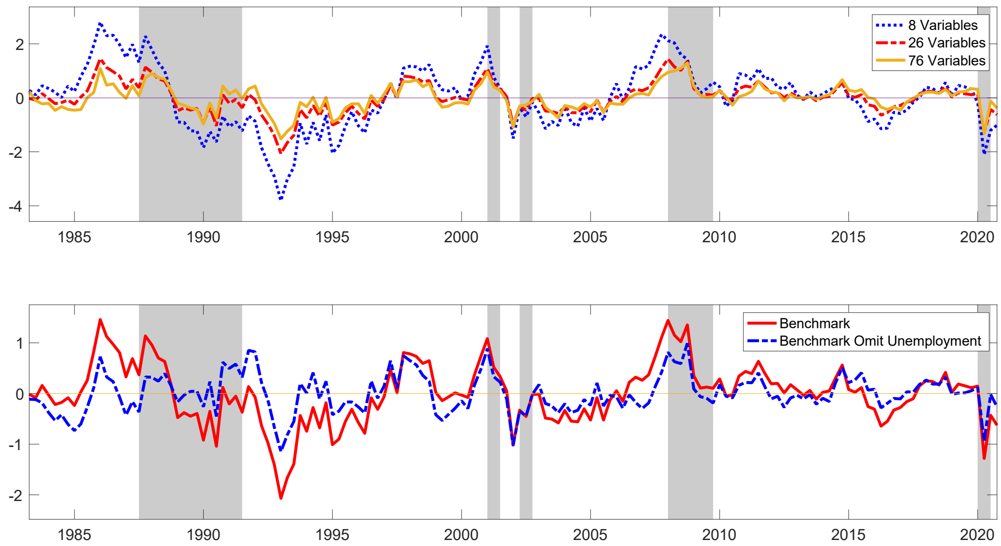

The final model selection is displayed in Figure A6 below, where we included the variable PFCE to the variable selection from Figure A5. Compared to the model selection in Figure A2, we have omitted the Oil Price and Industrial Share Prices, and added Hours worked and Private Final Consumption Expenditure (PFCF) from the benchmark model. The final 9-varibale model now includes Unemployment, Total Reserves, GDP growth, GFCF, Wages, Income, Sight Deposit Rate, Hours Worked and PFCE.

Figure A6: Estimated Norwegian output gap for 3 different sized BVAR models