



Australian
National
University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Trend-Cycle Decomposition After COVID

CAMA Working Paper 24/2024
March 2024

Güneş Kamber

International Monetary Fund

James Morley

The University of Sydney

Centre for Applied Macroeconomic Analysis, ANU

Benjamin Wong

Monash University

Centre for Applied Macroeconomic Analysis, ANU

Abstract

We revisit some popular univariate trend-cycle decomposition methods given the Covid-era data and find that only the output gap estimates from the Beveridge-Nelson filter remain both intuitive and reliable throughout the crisis and its aftermath. The real-time Hodrick-Prescott filter estimates for the output gap just prior to the pandemic are highly unreliable, although the estimated gap during the pandemic is reasonably similar to that of the Beveridge-Nelson filter. The Hamilton filter produces reliable estimates, but suffers from base effects that imply a purely mechanical spike in the output gap exactly two years after the onset of the crisis, in line with the filter horizon. Notably, unlike with the Beveridge-Nelson and Hodrick-Prescott filters, forecasts of the output gap for the Hamilton filter do not settle down to zero given plausible projected values of future output growth and display large spurious dynamics due to base effects given a simulated Covid-like shock in the projection. We also provide some refinements to the original Beveridge-Nelson filter that produce even more intuitive estimates of the output gap, while retaining the same strong revision properties.

Keywords

Beveridge-Nelson decomposition, output gap, real-time reliability

JEL Classification

C18, E17, E32

Address for correspondence:

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

ISSN 2206-0332

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

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

Trend-cycle decomposition after COVID*[†]

Güneş Kamber¹, James Morley^{2,4}, and Benjamin Wong^{3,4}

¹International Monetary Fund

²The University of Sydney

³Monash University, Australia

⁴CAMA

March 26, 2024

Abstract

We revisit some popular univariate trend-cycle decomposition methods given the Covid-era data and find that only the output gap estimates from the Beveridge-Nelson filter remain both intuitive and reliable throughout the crisis and its aftermath. The real-time Hodrick-Prescott filter estimates for the output gap just prior to the pandemic are highly unreliable, although the estimated gap during the pandemic is reasonably similar to that of the Beveridge-Nelson filter. The Hamilton filter produces reliable estimates, but suffers from base effects that imply a purely mechanical spike in the output gap exactly two years after the onset of the crisis, in line with the filter horizon. Notably, unlike with the Beveridge-Nelson and Hodrick-Prescott filters, forecasts of the output gap for the Hamilton filter do not settle down to zero given plausible projected values of future output growth and display large spurious dynamics due to base effects given a simulated Covid-like shock in the projection. We also provide some refinements to the original Beveridge-Nelson filter that produce even more intuitive estimates of the output gap, while retaining the same strong revision properties.

JEL Classification: C18, E17, E32

Keywords: Beveridge-Nelson decomposition, output gap, real-time reliability

*Email: Kamber: gkamber@imf.org Morley: james.morley@sydney.edu.au Wong: benjamin.wong@monash.edu

[†]The BN filter described in this paper can be implemented for any time series at <https://bnfiltering.com> and we thank Cristian Corrado for excellent work in developing the website. We also thank Abby Nguyen for helpful research assistance. Morley and Wong acknowledge funding support from the Australian Research Council (DP190100202, DP240100970 and DE200100693). The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. The other usual disclaimers also apply.

1 Introduction

The COVID-19 pandemic led to extreme outliers in macroeconomic data and created the need for important modifications to standard time series techniques, such as how to estimate vector autoregressions with Covid-era data, as discussed in Lenza and Primiceri (2022). However, in the case of trend-cycle decomposition, real-time reliability is a key attribute to consider when choosing a particular approach, and so it is of strong interest to know how different methods that were developed before the pandemic performed throughout the crisis and its aftermath.

In this paper, we revisit three popular univariate trend-cycle decomposition methods to examine their performance, including in terms of reliability, when faced with Covid-era data. These methods are the Hodrick and Prescott (1997) (HP) filter, the Hamilton (2018) filter, and Beveridge and Nelson (1981) (BN) filter from Kamber, Morley, and Wong (2018). In revisiting the BN filter, we also develop some refinements to the approach that are designed to be useful for practitioners.

After presenting our proposed refinements to the BN filter, we show that both the original and refined BN filters are very reliable, including during the pandemic. The refined BN filter selects a signal-to-noise ratio based on minimizing a non-zero stochastic trend volatility instead of maximizing the amplitude-to-noise ratio, as was considered in Kamber, Morley, and Wong (2018). In practice, this leads to a lower signal-to-noise ratio and produces an estimated output gap that displays even more coherence with the output gap for the HP filter. Although it is arguably even more intuitive by being more in line with the HP filter, we find that the refined BN filter output gap retains the same strong revision properties as the original BN filter in comparison with the less reliable HP filter.

In terms of the crisis, the real-time HP filter estimates of the output gap just prior to the pandemic turn out to be particularly unreliable. Specifically, the estimates at the time in 2019 suggested that the economy was very close to trend, while predictably given the large drop in output in 2020Q2 and the trend-smoothing criterion of the HP filter, the estimates of the output gap for 2019 have been subsequently revised upwards to imply the economy was in a huge boom at the time. These revisions are entirely mechanical and spurious, as we demonstrate with projections of future output growth that include a large one-time simulated Covid-like shock.

As hinted at by the criticism of the revision properties of the HP filter in Hamilton (2018), the Hamilton filter is much more reliable than the HP filter and this holds during the pandemic. However, the Hamilton filter suffers from base effects that produce a predictable spike in the estimated output gap exactly two years after the onset of the crisis, in line with the filter horizon. This sudden jump in the estimated output gap in 2022 is unintuitive and completely mechanical. The modified Hamilton filter due to Quast and Wolters (2022) is, as advertised, even more reliable and avoids such a large discrete jump in the estimated output gap in 2022 given that it averages over different horizons. However, both the original and modified Hamilton filters mechanically generate future spurious offsetting positive output gap estimates given the projections of future output growth that include a large one-time simulated Covid-like shock.

In addition, we find that both versions of the Hamilton filter produce very unintuitive forecasts of the output gap given the projected values of future output growth at a constant level. Specifically, the forecasts of the output gap do not settle down to zero, but grow on a trend path with a slope that depends on the assumed trend growth rate. Other than around the simulated Covid-like shock, the HP filter performs better in terms of the projections of future output growth, with forecasts of the output gap converging to zero when the economy is projected to grow at a constant trend rate. However, the refined BN filter performs well in terms of all of the projected data, while also being highly reliable in real time. Thus, we argue for the use of the refined BN filter when seeking both intuitive and reliable estimates of the output gap, including when considering real GDP data augmented with long-term projections as is often done in policy settings.

The rest of this paper is organised as follows: First, we recap the details of the BN filter and present some refinements designed to be useful for practitioners. Second, we consider the real-time reliability of the original and refined BN filters when including Covid-era data and compare revision properties to those of the HP filter and the original and modified Hamilton filters. Third, we consider what the different methods imply about forecasts of the future output gap given a projection of real GDP that reflects a plausible assumption about the long-run growth rate of the economy and also allows for a future Covid-like shock. We end with some brief conclusions.

2 The BN filter with refinements

We begin with a high-level description of the original BN filter developed in Kamber, Morley, and Wong (2018), before presenting our proposed refinements.

2.1 The Kamber, Morley, and Wong (2018) BN Filter

In Kamber, Morley, and Wong (2018), we employed the definition of trend from Beveridge and Nelson (1981) as the long-horizon conditional expectation of a time series minus any *a priori* known (i.e., deterministic) future movements in the time series. Denoting $\{y_t\}$ as a time series with a trend component that follows a random walk with constant drift, the BN trend at time t , τ_t^{BN} , is

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

The basic intuition behind the BN decomposition is that the long-horizon conditional expectation of a time series is the same as the long-horizon conditional expectation of the trend component under the assumption that the conditional expectation of the remaining cyclical component goes to zero at long horizons. By removing the deterministic drift, $\mathbb{E} [\Delta y_t]$, the conditional expectation in equation (1) remains finite and becomes an optimal (minimum mean squared error) estimate of the current trend component (see Watson, 1986; Morley, Nelson, and Zivot, 2003).

The BN filter in Kamber, Morley, and Wong (2018) implements the BN decomposition by applying the definition in equation (1) given a restricted AR(p) forecasting model for $\{\Delta y_t\}$ to calculate conditional expectations, where the restriction allows the imposition of a signal-to-noise ratio to imply a smooth trend, similar to the HP filter. More precisely, we specify an AR(p) model:

$$\Delta y_t = \mu + \sum_{j=1}^p \phi_j (\Delta y_{t-j} - \mu) + e_t, \quad e_t \sim N(0, \sigma_e^2), \quad (2)$$

where μ is equal to $\mathbb{E}[\Delta y_t]$ and, therefore, denotes the deterministic drift in $\{y_t\}$.¹ Modelling a time series as an AR(p) process in first differences allows for a stochastic trend in the level because a forecast error for the first differences will have a *permanent* effect on the long-horizon conditional expectation of $\{y_t\}$. Using the state-space approach to calculating the BN decomposition in Morley (2002), the BN cycle at time t , c_t^{BN} , for this model is

$$c_t^{BN} = -[1 \quad 0 \quad \dots \quad 0]F(I - F)^{-1}X_t, \quad (3)$$

where $X_t = (\Delta y_t - \mu, \Delta y_{t-1} - \mu, \dots, \Delta y_{t-p+1} - \mu)'$, and F is the companion matrix for the AR(p) model:

$$F = \begin{bmatrix} \phi_1 & \phi_2 & \dots & \phi_p \\ 1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & 1 & 0 \end{bmatrix}.$$

Defining the signal-to-noise ratio, δ , in terms of the variance of the change in trend as a fraction of the overall forecast error variance (i.e., $\delta \equiv \sigma_{\Delta\tau}^2/\sigma_e^2$), Kamber, Morley, and Wong (2018) note that the signal-to-noise ratio δ maps directly to the coefficients of the AR(p) model, so that one can impose the signal-to-noise ratio directly on the model in equation (2) by specifying a restriction on the AR coefficients. Specifically, we transform the AR(p) model in equation (2) into its Dickey-Fuller representation:

$$\Delta y_t = \mu + \rho(\Delta y_{t-1} - \mu) + \sum_{j=1}^{p-1} \phi_j^* (\Delta^2 y_{t-j} - \mu) + e_t, \quad (4)$$

where $\rho \equiv \phi_1 + \phi_2 + \dots + \phi_p$ and $\phi_j^* \equiv -(\phi_{j+1} + \dots + \phi_p)$. There is a direct mapping from the signal-to-noise ratio δ to the sum of the AR coefficients ρ in the Dickey-Fuller representation, with $\delta = (1 - \rho)^{-2}$. This mapping underpins a key motivation of Kamber, Morley, and Wong (2018), given that freely estimating an AR(p) model for real GDP growth often yields a signal-to-noise ratio in excess of 1, whereas common methods of removing the trend in macroeconomic time series data such as the HP filter explicitly or implicitly impose a much lower signal-to-noise

¹Our choice of exposition is in terms of the demeaned form of the AR(p) model rather than specifying an intercept form in order to make it clear where the drift that enters into the calculation of the BN decomposition comes from. If one were to estimate an intercept, denoting c as the intercept, the drift would simply correspond to $c/(1 - \sum_i \phi_i)$.

ratio. The procedure developed in Kamber, Morley, and Wong (2018) thus allows one to impose a low signal-to-noise ratio by specifying $\bar{\delta}$. Then, equation (4) can be estimated imposing a particular signal-to-noise ratio $\bar{\delta}$ by fixing $\bar{\rho}$ as follows:

$$\bar{\rho} = 1 - 1/\sqrt{\bar{\delta}}. \quad (5)$$

The BN decomposition can be applied imposing the particular signal-to-noise ratio $\bar{\delta}$ by first solving for the restricted estimates of $\{\phi_j\}_{j=1}^p$ by inverting the Dickey-Fuller transformation given $\bar{\rho}$ and estimates of $\{\phi_j^*\}_{j=1}^{p-1}$ and then calculating the BN cycle following equation (3).²

Implementation of the BN filter procedure could be as straightforward as setting $\bar{\delta}$ to a particular low value such as, for example, $\bar{\delta} = 0.05$, which would correspond to the assumption that only 5% of the quarterly forecast-error variance for output growth is due to trend shocks, and then carrying out estimation of the restricted model imposing the corresponding $\bar{\rho}$ and applying the BN decomposition.

Because making a particular choice for $\bar{\delta}$ might appear somewhat arbitrary in practice, we proposed an automatic selection of $\bar{\delta}$ based on maximizing the amplitude-to-noise ratio in Kamber, Morley, and Wong (2018). The intuition of this approach is that, given practitioners often have a low signal-to-noise ratio (and thus high amplitude cycles) in mind, this criterion allows for a trade-off of a larger amplitude cycle against the poorer fit of a corresponding restricted forecasting model.

Time-varying drift While the key innovation with of our original BN filter was to impose a low signal-to-noise ratio in an AR(p) model when performing the BN decomposition, we also explicitly allowed for the possibility of time-varying drift. This accommodation was important because empirical evidence suggests the possibility of breaks in drift for many time series, including US real GDP, and so allowing for the possibility of time-varying drift broadens the utility of the BN filter for practitioners. Generally, the drift term (i.e. μ in equation (2)) needs to be estimated. We originally considered two possibilities for how to deal with time-varying drift in applied work. We first considered testing for structural breaks in μ , using something like the Bai and Perron (2003) approach, and if breaks in μ appear to have occurred, we then adjust for these breaks by demeaning $\{\Delta y_t\}$ using subsample averages given estimated break dates and then proceeding with the rest of the procedure. In an output gap application, this amounts to finding breaks in the mean growth rate of real GDP and demeaning the growth rate based on the breaks. While such an approach is useful for *ex-post* analysis, it is less useful for real-time analysis. This is necessitated by the fact that breaks are, by their nature, two-sided, and one needs to know the future growth rate before confidently dating the timing of a break *ex post*, a luxury that practitioners who need to estimate the output gap in real time do not enjoy. We therefore proposed a second approach of using a rolling 40-quarter window to estimate a

²Equation (4) can be easily estimated via classical or Bayesian methods. For the BN filter in Kamber, Morley, and Wong (2018), the model is estimated using Bayesian methods with Minnesota-like priors to mitigate possible overfitting when specifying a high-order AR(p) model that can allow for a relatively complicated autocovariance structure for $\{\Delta y_t\}$.

time-varying drift. That is, we set $\hat{\mu}_t = \frac{1}{40} \sum_{i=1}^{40} \Delta y_{t+1-i}$ and thereafter dynamically demean the first differences of the time series being detrended. The dynamic demeaning procedure has two advantages. First, because it is backward-looking, one can adjust for a break in real time before there is evidence of a break in mean. Second, if the drift did not change or was largely unchanged, the adjustment will make little practical difference to the estimated cyclical component, while if there were indeed large breaks, one would not require strong statistical evidence in the form of a Bai and Perron (2003) test, which could take a decade or more to statistically date a break, before accounting for the break. This is because a sufficiently large break in real time should induce a discernible change on the estimates of the mean growth rate over the past 40 quarters.³

2.2 Refinements to the BN filter

We propose four separate refinements to the BN filter. At a high level, we view only one of these as a major refinement, with the other three being relatively more minor, but still potentially useful. The main refinement pertains to a trend-smoothness loss function when selecting a particular signal-to-noise ratio $\bar{\delta}$. For the minor refinements, we first discuss an update to our dynamic-demeaning procedure to address outliers that can affect the average change in the cycle over the window used to estimate time-varying drift. Then we consider construction of error bands for the estimated cycle to allow for time-varying volatility. Finally, we update our backcasting procedure to provide an estimated cycle for the full sample of the variable being detrended. We now present the details of these refinements in turn.

2.2.1 A trend-smoothness loss function when selecting $\bar{\delta}$

As an alternative to maximizing the amplitude-to-noise ratio, we propose choosing a positive $\bar{\delta}$ that minimizes the variance of the change in trend. Specifically, our optimization problem for this new loss function is given by

$$\min_{\delta \in \mathbb{R}^+} \sigma_{\Delta\tau}^2(\delta), \quad (6)$$

where we explicitly consider an interior solution to this problem. Obviously, imposing the limiting case of $\delta = 0$ would truly minimize the implied variance of trend shocks by implying there is no stochastic trend at all – i.e., $\sigma_{\Delta\tau}^2(0) = 0$. However, that limiting case corresponds to a unit MA root for $\{\Delta y_t\}$ and is, therefore, incompatible with using a finite-order AR(p) model, for which the δ implied by ρ is strictly positive. At the same time, near zero values of δ would not be close to the optimum if the fit of the model deteriorates as $\delta \rightarrow 0$ – i.e., $\lim_{\delta \rightarrow 0} \sigma_e^2(\delta)' < 0$. Specifically, noting that $\sigma_{\Delta\tau}^2(\delta) = \delta \sigma_e^2(\delta)$ and, therefore, $\sigma_{\Delta\tau}^2(\delta)' = \delta \sigma_e^2(\delta)' + \sigma_e^2(\delta)$ and

³We can allow for different possible windows for estimating a time-varying drift, but we use 40 quarters when considering estimation of the output gap given that the window is long enough that temporary effects of business cycles on growth rates should average out and the estimated drift will also be reasonably precise, while it is short enough that it allows for substantial low frequency changes in long-run growth. One could start with a longer window and shorten it until results for implied time-varying drift become either relatively unchanged or very imprecise with a further shortening of the window.

$\sigma_{\Delta\tau}^2(\delta)'' = \delta\sigma_e^2(\delta)'' + 2\sigma_e^2(\delta)'$, it is easy to see that $\lim_{\delta \rightarrow 0} \sigma_{\Delta\tau}^2(\delta)' > 0$ and $\lim_{\delta \rightarrow 0} \sigma_{\Delta\tau}^2(\delta)'' < 0$, meaning that a local minimum would be interior as it requires $\sigma_{\Delta\tau}^2(\bar{\delta})' = 0$ and $\sigma_{\Delta\tau}^2(\bar{\delta})'' > 0$.⁴

We see this trend-smoothness loss function as being in the spirit of trend-cycle decomposition methods that impose explicit smoothness priors on the trend, such as Harvey, Trimbur, and Van Dijk (2007), or impose a very low signal-to-noise, such as the HP filter. This loss function should be well suited to events such as the COVID-19 pandemic that involve extreme movements in the data without such an obvious of a change in how much the trend moves. At the same time, smoothing the trend in this way has the potential to alter the reliability properties of the BN filter given other methods that smooth the trend, such as the HP filter, are notably less reliable and can also generate spurious cycles (Cogley and Nason, 1995). We also note that our proposed loss function has also been useful beyond considering the univariate BN filter. In related work estimating the euro area output gap using a multivariate BN decomposition, minimizing the variance of the change in trend also helped in estimating the depth of the output gap during the Covid-19 recession for the euro area (see Morley, Rodriguez-Palenzuela, Sun, and Wong, 2023).

2.2.2 Iterative dynamic demeaning

To understand our proposed iterative dynamic-demeaning approach, it is helpful to revisit the basic assumption that the trend of $\{y_t\}$ is a random walk with drift. That is, the non-zero unconditional expectation $\mathbb{E}[\Delta y_t] = \mu \neq 0$ is associated with the trend, not the cycle. Recall the basic trend-cycle identity:

$$y_t = \tau_t + c_t. \quad (7)$$

Then, it directly follows that

$$\Delta y_t = \Delta \tau_t + \Delta c_t \quad (8)$$

and

$$\mathbb{E}[\Delta y_t] = \mathbb{E}[\Delta \tau_t] + \mathbb{E}[\Delta c_t]. \quad (9)$$

The assumption that drift is associated with the trend is equivalent to assuming $\mathbb{E}[\Delta \tau_t] = \mu$ and $\mathbb{E}[\Delta c_t] = 0$, which directly follows from the maintained assumption $\mathbb{E}[c_t] = 0$ when applying the BN decomposition.

When we use the sample mean of $\{\Delta y_t\}$ to estimate μ , we are implicitly assuming that the sample mean of $\{\Delta c_t\}$ is essentially zero, consistent with its unconditional expectation. That

⁴If the model fit also deteriorates for large enough δ – i.e., $\lim_{\delta \rightarrow \infty} \sigma_e^2(\delta)' > 0$, this implies $\sigma_e^2(\delta^*)' = 0$ and $\sigma_e^2(\delta^*)'' > 0$ for some $\delta^* > 0$, which, in turn, implies $\sigma_{\Delta\tau}^2(\delta^*)' > 0$ and $\sigma_{\Delta\tau}^2(\delta^*)'' > 0$. This further implies $\sigma_{\Delta\tau}^2(\underline{\delta})'' = 0$ for some smaller value of $0 < \underline{\delta} < \delta^*$ for which $\sigma_e^2(\underline{\delta})' < 0$ and $\sigma_e^2(\underline{\delta})'' > 0$. Then solving for an optimum by finding $\bar{\delta} > \underline{\delta}$ such that $\sigma_{\Delta\tau}^2(\bar{\delta})' = 0$ would correspond to a local interior minimum because $\sigma_{\Delta\tau}^2(\delta)'' > 0$ for $\delta > \underline{\delta}$. An exact analytical solution depends on the noise function $\sigma_e^2(\delta)$. In practice, we consider a numerical grid to find the $\bar{\delta}$ that minimizes the $\sigma_{\Delta\tau}^2(\delta)$. It is straightforward, then, with the grid to confirm existence and uniqueness of the local minimum in any given setting. We always find an interior minimum in practice.

is, the trend-cycle identity in equation (7) also implies

$$\frac{1}{T} \sum_{t=1}^T \Delta y_t = \frac{1}{T} \sum_{i=t}^T \Delta \tau_t + \frac{1}{T} \sum_{t=1}^T \Delta c_t \quad (10)$$

and

$$\frac{1}{T} \sum_{t=1}^T \Delta c_t \approx 0 \Rightarrow \frac{1}{T} \sum_{i=t}^T \Delta \tau_t \approx \frac{1}{T} \sum_{t=1}^T \Delta y_t. \quad (11)$$

Thus, the sample mean of $\{\Delta y_t\}$ approximates the sample mean of $\{\Delta \tau_t\}$, and thus provides an easy way to estimate the drift associated with the trend.

However, when we consider dynamic demeaning to account for time-varying drift, the sample for our rolling-window estimation can become small enough that the average change in the cycle over the rolling window might not be zero. Of course, when estimating the cycle, we obviously do not observe it *a priori*, so we cannot simply remove the sample mean of the change in the cycle to estimate the drift associated with trend growth. But if our estimated cycle using $\{\Delta y_t\}$ to estimate time-varying drift implies a non-zero sample mean of the change in the cycle, then we can consider iterative estimation of trend growth until the estimated cycle is consistent with the sample mean of the estimated change in the cycle. Specifically, initially setting $c_t^{\{0\}} = 0$ and $j = 1$, we repeatedly estimate the cycle using dynamic demeaning according to $\mu_t^{\{j\}} = \frac{1}{40} \sum_{i=1}^{40} \Delta y_{t+1-i} - \frac{1}{40} \sum_{i=1}^{40} \Delta c_{t+1-i}^{\{j-1\}}$ for j iterations until $c_t^{\{j\}} \approx c_t^{\{j-1\}}$ up to some arbitrary level of precision. Note that $c_t^{\{1\}}$ is just our original dynamic-demeaning estimate using $\{\Delta y_t\}$ to estimate time-varying drift. The iterative approach helps address unusual movements in a cycle that can result in the average change in the cycle over the rolling window being significantly different than zero. We are motivated by the large outliers associated with COVID-19 to consider this approach, although we find relatively little difference in estimates during the pandemic using the original or iterative approaches given that the change in the cycle, while large, was not persistent. We find larger changes in estimates at other points of the sample period.

2.2.3 Time-varying error bands

In the appendix of Kamber, Morley, and Wong (2018), we presented a method to assess the uncertainty associated with the BN filter estimates. In particular, based on equations (3) and (4), we solved for the variance of the BN cycle, σ_c^2 , as follows:

$$\sigma_c^2 = [1 \ 0 \ \dots \ 0] F(I - F)^{-1} \Sigma_X ((I - F)')^{-1} F' [1 \ 0 \ \dots \ 0]', \quad (12)$$

where Σ_X is the variance of X_t and $vec(\Sigma_X) = (I - F \otimes F)^{-1}vec(Q)$, with Q being the variance-covariance matrix for the innovation vector of the companion form for the AR(p) model:

$$Q = \begin{bmatrix} \sigma_e^2 & 0 & \cdots \\ 0 & 0 & \ddots \\ \vdots & \ddots & \ddots \end{bmatrix}. \quad (13)$$

Because the estimated trend provides an unbiased estimator of the true trend, we proposed constructing 95% confidence bands for the BN cycle by inverting a simple z -test for different values of c_t as follows:

$$c_t^{BN} \pm 1.96\sigma_e. \quad (14)$$

While useful, the bands based on this approach have a constant width. However, if the conditional volatility of $\{\Delta y_t\}$ were actually time varying, the originally proposed approach could incorrectly estimate the degree of estimation uncertainty at different points of the sample, which seems particularly relevant when faced with events such as the COVID-19 pandemic when there was unprecedented volatility in real GDP growth.⁵ Thus, the refinement we propose for the construction of error bands is to allow for the possibility of time-varying volatility. In particular, our simple proposal is to estimate time-varying conditional volatility of $\{\Delta y_t\}$ using a rolling window, although one could certainly consider a more complicated approach. Specifically, given window size k , we replace the estimate of conditional volatility σ_e^2 in equation (13) based on least squares residuals from equation (2) with estimates of $\sigma_{e,t}^2$ using a rolling window from $t - k + 1$ to t . We consider $k = 40$ quarters as the choice for the rolling window to align with our choice for dynamic demeaning.⁶

2.2.4 Iterative backcasting to generate trend-cycle estimates for y_1

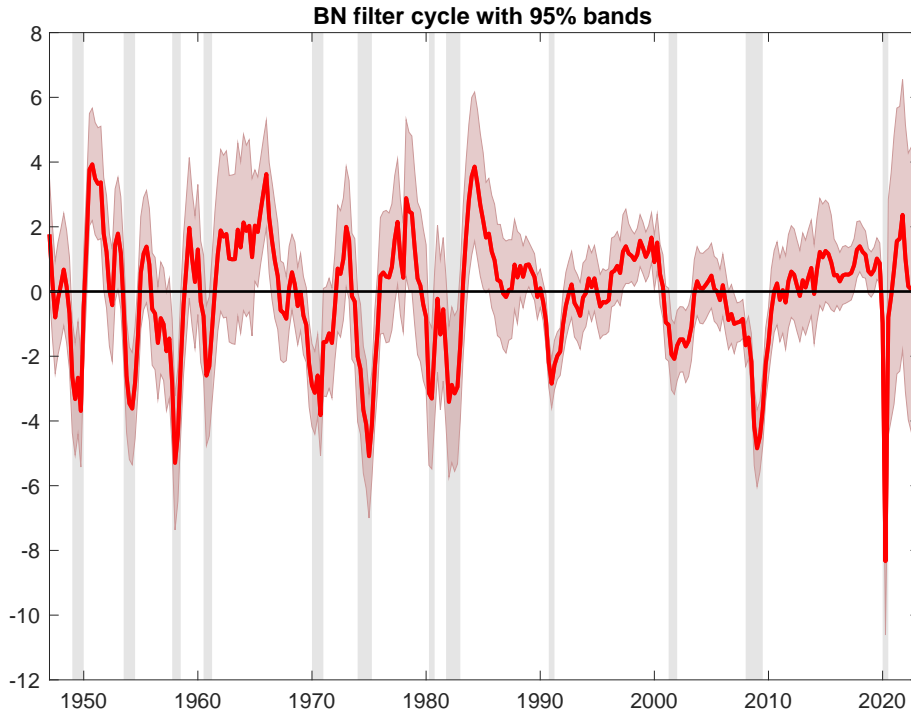
Our last refinement involves iterative backcasting of initial observations for estimation by utilizing the time reversibility of a linear time series process that can be described in terms of ARMA dynamics (see, for example, Ramsey and Rothman, 1996). Our motivation of doing so is because trend-cycle decompositions such as the HP filter produce an estimate of trend and cycle for the first observation of the level of the time series being decomposed, y_1 , while the original BN filter only does so for the first available observation of the first differences, which corresponds to the second observation in levels, y_2 . In order to obtain a trend and cycle estimate from the BN filter for the first observation in levels, we need a better estimate of the first difference associated with the initial observation, Δy_1 , than just using the estimated drift, $\hat{\mu}$, which is what we used for backcasting initial observations when estimating the restricted AR(p) model in the original BN filter. We find estimates of initial observations by using the

⁵We note that accounting for changes in volatility is less relevant for the estimated cycle based on equation (3) because modelling heteroskedasticity would have minimal impact on estimates of the autoregressive coefficients in F , restricted or otherwise. Thus, it is not such an issue when conducting the BN decomposition to ignore potential heteroskedasticity. It is only when constructing confidence bands that this heteroskedasticity matters.

⁶As with dynamic demeaning, the estimates for the first 40 quarters are constant and based on the first 40 quarters.

restricted AR(p) model to forecast $p + 1$ observations in first differences out of sample and then reversing the time series in differences including these forecasted observations in order to backcast $p + 1$ initial observations in first differences. We do this iteratively until the estimated initial observations and out-of-sample forecasts converge to an arbitrarily small tolerance. In practice, we find this convergence is almost immediate. Then, with $p + 1$ backcast initial observations, $\Delta\hat{y}_1, \Delta\hat{y}_0, \dots, \Delta\hat{y}_{-p+1}$, we can use equation (3) to calculate the cycle (and, therefore, trend) for y_1 .

Figure 1: Estimates of the US output gap using the refined BN filter



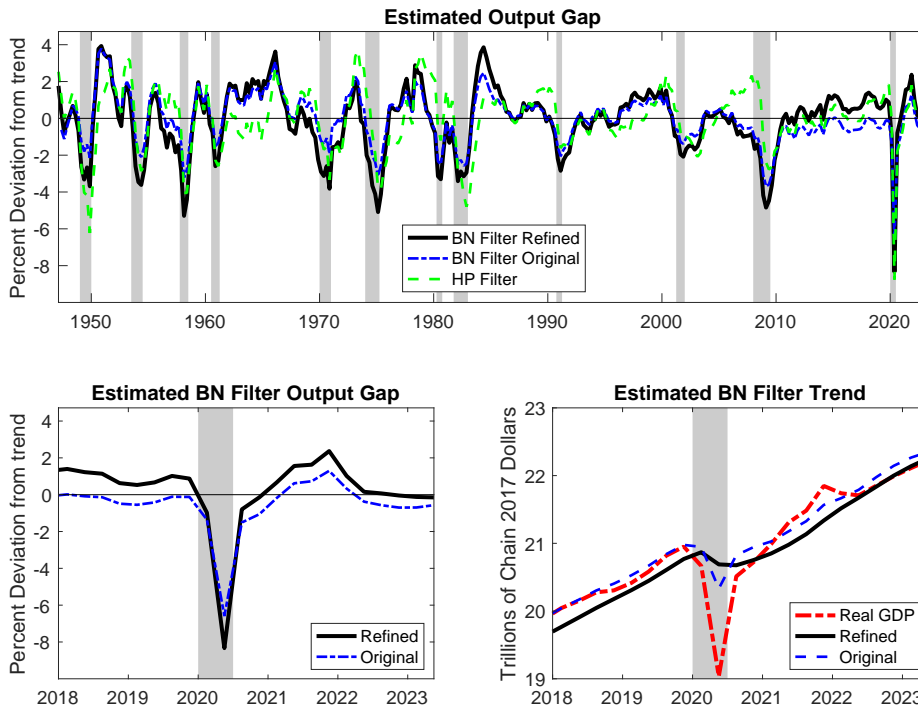
Notes: Units are 100 times natural log deviation from trend. Sample period is 1947Q1 to 2023Q2. Shaded bands around the estimate correspond to 95% confidence intervals based on inverting a z -test that the true output gap is equal to a hypothesized value using the standard deviation of the BN cycle estimated with rolling-window estimates of time-varying conditional volatility for output growth. Shaded bars correspond to NBER recession dates.

2.3 Estimates of the output gap using the refined BN filter

We first present estimates of the output gap using US real GDP data from 1947Q1 to 2023Q2 when considering all of the refinements. Figure 1 plots these estimates, along with 95% confidence intervals. Similar to what we found in Kamber, Morley, and Wong (2018), the refined BN filter produces intuitive estimates of the output gap, with movements in it being well aligned to the NBER reference cycle. We note that there is a larger degree of estimation uncertainty in the earlier part of the sample than during the Great Moderation, but a higher level of uncertainty has returned since the COVID-19 pandemic given the inclusion of outlier growth rates associated with the pandemic in the rolling window when estimating time-varying conditional volatility of output growth. The $\bar{\delta}$ selected given the new loss function is 0.02, implying that only 2% of the quarterly forecast-error variance for output growth corresponds to trend shocks.

This contrasts with the $\bar{\delta} = 0.24$ found in Kamber, Morley, and Wong (2018), which while yielding intuitive and reliable estimates of the output gap, may have been larger than believed by many practitioners seeking to estimate a smooth trend.

Figure 2: Comparison of the original and refined BN filter



Notes: Units are 100 times natural log deviation from trend for the first two panels and trillions of 2017 dollars for the third panel. Sample period is 1947Q1 to 2023Q2 for the top panel and 2018Q1 to 2023Q2 for the bottom panels. Shaded bars correspond to NBER recession dates.

To help understand the role of the refinements, Figure 2 compares the results for the refined BN filter with those for the updated sample based on the original BN filter proposed in Kamber, Morley, and Wong (2018). The top panel of Figure 2 plots the estimated output gap for the refined BN filter together with estimated output gap for the original BN filter, while also including the output gap based on the HP filter using the typical smoothing parameter of 1600 as a point of reference. In general, the refined BN filter produces an estimated output gap with fluctuations that are very similar to those for the original BN filter. However, there are some key differences. The refined BN filter estimates display a more similar amplitude to that of the HP filter estimates than the original BN filter estimates, with the refined BN filter and HP filter estimates being reasonably similar to each other around the pandemic. The differences in amplitude between the original and refined BN filters are particularly notable during the 1970s recessions, the mid-1980s, the Great Recession, and the COVID-19 pandemic. Specifically, because the new loss function is explicitly designed to minimize the change in trend, it mechanically attributes a larger proportion of the fluctuations in real GDP to the cycle.⁷ All in all, the comparison with the HP filter suggests that the new loss function for the

⁷We note that the original procedure selects a $\bar{\delta}$ of 0.15 for the updated sample, which is lower than the selected value of 0.24 for the original sample from 1947Q1 to 2016Q2 used in Kamber, Morley, and Wong (2018). Nonetheless, this is still much higher than the $\bar{\delta}$ of 0.02 selected with the proposed new loss function.

refined BN filter comes closer to the implicit loss function of practitioners who continue to use the HP filter despite repeated warnings not to do so (e.g., Cogley and Nason, 1995; Hamilton, 2018).

To provide some sense of how our updated approach deals with the pandemic, the bottom panels of Figure 2 compare the refined BN filter with the original BN filter, zooming in on the COVID-19 period in particular. For the comparison, we present both the output gap and trend output estimates. The refined BN filter attributes more of the fall of real GDP during 2020Q2 as being transitory rather than permanent. The corresponding trend estimates show that the refined BN filter implies a less sharp reversal in trend output. The original BN filter, on the other hand, implies a sharper drop followed by increase in trend output in 2020Q2 and 2020Q3, respectively. Given what we know happened during that period, the refined BN filter produces estimates of the output gap and trend output that are arguably more consistent with the economic narrative of the time. In particular, the outlier fall in real GDP growth in 2020Q2 was clearly associated with temporary and massive restrictions on economic activity such as lockdowns that were largely reversed in 2020Q3, which suggests an economic narrative that most of the decline in activity due to the pandemic was transitory.⁸ The differences in estimates are largely about the new loss function, with some differences around Covid also reflecting dynamic demeaning given lower estimated trend growth in the 2010s implying a generally higher level of the output gap. However, the new iterative approach to dynamic demeaning plays a more limited role, with only some small changes in estimates such as during the Great Recession given a large persistent drop in the output gap, and almost no change in estimates during the pandemic given the large drop in the output gap only lasted for one quarter.⁹

3 Real-time reliability when including Covid-era data

We now compare the real-time reliability of the BN filter against some widely-used trend-cycle methods by examining their revision properties when including Covid-era data in the analysis. This evaluation is motivated by the well-documented general unreliability of output gap estimates in real time, as most prominently demonstrated by Orphanides and van Norden (2002). A key attraction of the original BN filter is that it appears to have revision properties which largely circumvent the Orphanides and van Norden (2002) critique. It is thus of interest to understand whether both the original and refined BN filters retain such good revision properties when confronted with Covid-era data and how they compare to each other and other widely-used methods.


In terms of the other methods, we consider the HP and Hamilton filters. The HP filter is a natural choice given its ubiquitous use in academia and policy environments, and we consider the standard implementation of the HP filter using a smoothing parameter of 1600, as in Figure

⁸It is possible that the disruption associated with the pandemic had some more persistent and possibly permanent negative implications for economic activity as well, consistent with finding of some reduction in trend when using our trend-smoothing loss function or even the HP filter.

⁹The iterative backcasting also plays a minimal role in the differences in output gap estimates between the original and refined BN filter, although it is still beneficial because it provides an estimate for 1947Q1.

2. The Hamilton filter is also an obvious choice given its recent emergence due to its ease of use and also its purported ability to address various shortcomings of the HP filter. Furthermore, since the publication Kamber, Morley, and Wong (2018), Quast and Wolters (2022) have shown that the Hamilton filter possesses similarly good revision properties when comparing to the BN filter. The Hamilton filter estimates the output gap by obtaining the residual from a least squares regression of h -step-ahead log real GDP on its contemporaneous value and three lags as well as a constant. We consider two versions of the Hamilton filter: the original version in Hamilton (2018), which sets $h = 8$, and the modified version in Quast and Wolters (2022), which averages over gaps obtained from $h = 4$ to $h = 12$.

Figure 3: Ex-post versus pseudo-real-time output gap estimates



Units are 100 times natural log deviation from trend. Sample period is 1947Q1 to 2023Q2. Shaded bars correspond to NBER recession dates.

Figure 3 plots *ex-post* and pseudo-real-time estimates of the US output gap for the different methods. By *ex post*, we are referring to the output gap estimated using the full 1947Q1-2023Q3 sample of data. For the pseudo-real-time estimates, we sequentially estimate the output gap starting with a shorter sample from 1947Q1 to 1970Q1 and then adding one observation of real GDP at a time, retaining the last estimated value of the output gap for each pseudo-real-time sample. This is analogous to the filtered estimate of the output gap if one were estimating the gap in real time using the Kalman filter, albeit with the final vintage of data.¹⁰

¹⁰We consider just a single vintage of data in order to isolate the role of the methods rather than data revisions for three reasons. First, Orphanides and van Norden (2002) show that most of the revisions for output gap estimates (about two-thirds) are due to the method rather than data revisions. Second, none of the methods are explicitly designed to address data revisions, so using a single vintage of data isolates the analysis to how well the various methods are able to deal with the so-called “end-point” problem. Third, because none of these methods are designed to address data revisions, Kamber, Morley, and Wong (2018) show it makes relatively little difference whether one uses real-time data or a single vintage if one were just seeking to understand the

Consistent with Orphanides and van Norden (2002) and Kamber, Morley, and Wong (2018), Figure 3 shows that the output gap based on the HP filter is often heavily revised, with the pseudo-real-time and *ex-post* estimates diverging considerably. Notably, despite being near the end of the sample, the real-time HP filter estimates of the output gap just prior to the pandemic are clearly unreliable, with the *ex-post* estimates revised upwards substantially in 2019 from being very close to zero to implying a huge boom. By contrast, all variants of the BN filter and the Hamilton filter have pseudo-real-time and *ex-post* estimates that are much more similar.

Despite the inclusion of Covid-era data in the analysis, it is clear from Figure 3 that the original BN filter retains the good revision properties reported in Kamber, Morley, and Wong (2018) and also confirmed in Quast and Wolters (2022) and Barigozzi and Luciani (2023). Meanwhile, the refined BN filter displays similarly good revision properties as the original BN filter. Note that because we estimate the output gaps in pseudo real time, we are also allowing $\bar{\delta}$ to change in the real-time setting. While the pandemic did change the selected $\bar{\delta}$, most prominently for the original BN filter from the originally published 0.24 to 0.15 at the final data point, this makes little difference to the *ex-post* versus pseudo-real-time estimates. In fact, it is precisely because of the BN filter’s reliability that $\bar{\delta}$ falls post pandemic. Given that the historical and *ex-post* estimates do not diverge post pandemic, the only way that the BN filter can retain its reliability when considering either the amplitude-to-noise ratio or the variance of the change in trend in the presence of large outlier shocks that boost the sample variance of the forecast errors is through a corresponding decrease in the signal-to-noise ratio δ given “noise” is being measured by σ_e^2 .


Turning to the Hamilton filter, estimates for both the original and modified versions are also little revised in Figure 3, consistent with what was documented by Quast and Wolters (2022). An interesting finding, though, is that, while the estimated output gaps between the original and modified versions were very similar pre pandemic, they are quite different post pandemic. In particular, the original Hamilton filter output gap displays a large mechanical spike in the estimated output gap in 2022Q2, exactly two years after the onset of the crisis. The spike is mechanical because it was perfectly predictable prior to 2022Q2 given base effects for the 8-quarter-ahead projection of the level of log real GDP for 2022Q2 due to the low value of log real GDP in 2020Q2 and then for 2022Q3 due to the largely recovered value of log real GDP in 2020Q3.¹¹ We will return to this issue of how and why the original and modified versions of

comparative revision properties of the methods under consideration.

¹¹The predictable base effects may be easiest to think about by considering the simplified version of the Hamilton filter also discussed in Hamilton (2018) that constructs the output gap as the eight-quarter difference in log real GDP – i.e., $y_t - y_{t-8} = \Delta y_t + \Delta y_{t-1} + \dots + \Delta y_{t-7}$. From 2020Q3 to 2022Q1, the eight-quarter difference will include the 2020Q2 and 2020Q3 quarterly growth rates, which were largely offsetting. So, if all other quarterly growth rates were much closer to the drift μ , the eight-quarter difference from 2020Q3 to 2022Q1 would be close to $6 \cdot \mu$. Then, predictably in 2022Q2, the output gap will jump up by about $-\Delta y_{2020Q2} + \mu$ as the very negative quarterly growth for 2020Q2 is dropped from the eight-quarter growth rate and the 2022Q2 growth rate, which is assumed to be closer to the drift μ , is added. Furthermore, predictably in 2022Q3, the output gap will fall back down by $-\Delta y_{2020Q3} + \mu$ to close to $8 \cdot \mu$ as the very positive growth rate for 2020Q3 is dropped and the 2022Q3 growth rate, which is again assumed to be closer to the drift μ , is added. In this sense, it is predictable changes in the “base” for the eight-quarter difference (i.e., y_{t-8}) that mechanically explain the positive spike exactly 8 quarters after a large negative spike.

the Hamilton filter are so different post pandemic in Section 4.

Figure 4: Revision properties of output gap estimates



Standard deviation and root mean square of revisions to the pseudo-real-time estimate of the output gap are normalized by the standard deviation of the *ex-post* estimate of the output gap. “Correlation” refers to the correlation between the pseudo-real-time estimate and the *ex post* estimate of the output gap. “Same Sign” refers to the proportion of pseudo-real-time estimates that share the same sign as the *ex post* estimate of the output gap. The sample period for calculation of revision statistics is 1970Q1 to 2018Q2.

To round out the real-time reliability analysis, Figure 4 presents the revision statistics for the various methods, albeit dropping the last five years of the sample when calculating the statistics because more recent *ex-post* estimates near the end of the sample may end being heavily revised in the future. These statistics were originally proposed by Orphanides and van Norden (2002). The top panel plots the size of the revisions between the pseudo-real-time and *ex-post* estimates in terms of standard deviation and root mean square, both normalized by the standard deviation of the *ex-post* estimated gap, as suggested by Orphanides and van Norden (2002). The bottom panel plots the correlation between the pseudo-real-time and *ex-post* estimates and proportion of these estimates that share the same sign. As one might suspect from Figure 3, all variants of the BN filter and the Hamilton filter have good revision statistics, whereas the HP filter does poorly. On some level, this is perhaps not entirely surprising. As first pointed out by Kamber, Morley, and Wong (2018), the BN filter does well because it is a one-sided filter which does not rely on future information, unlike the HP filter and the various other filters considered by Orphanides and van Norden (2002). The Hamilton filter also falls in the category of a one-sided filter, so it naturally also does well on these metrics as long as the estimated regression parameters remain little changed with the addition of new data. Finally, we note that, even though the original Hamilton and BN filters do well on these metrics in an absolute sense, their updated versions seem to do better on at least some of the metrics, especially for the refined BN filter, which suggest that the updated versions can provide some

valuable enhancements without sacrificing real-time reliability. Again, this is a particularly notable finding for the refined BN filter given that the trend-smoothing criterion of the HP filter seems to be a source of its unreliability, while the trend-smoothness loss function for the BN filter perhaps surprisingly does not seem to lead to any deterioration in its reliability.

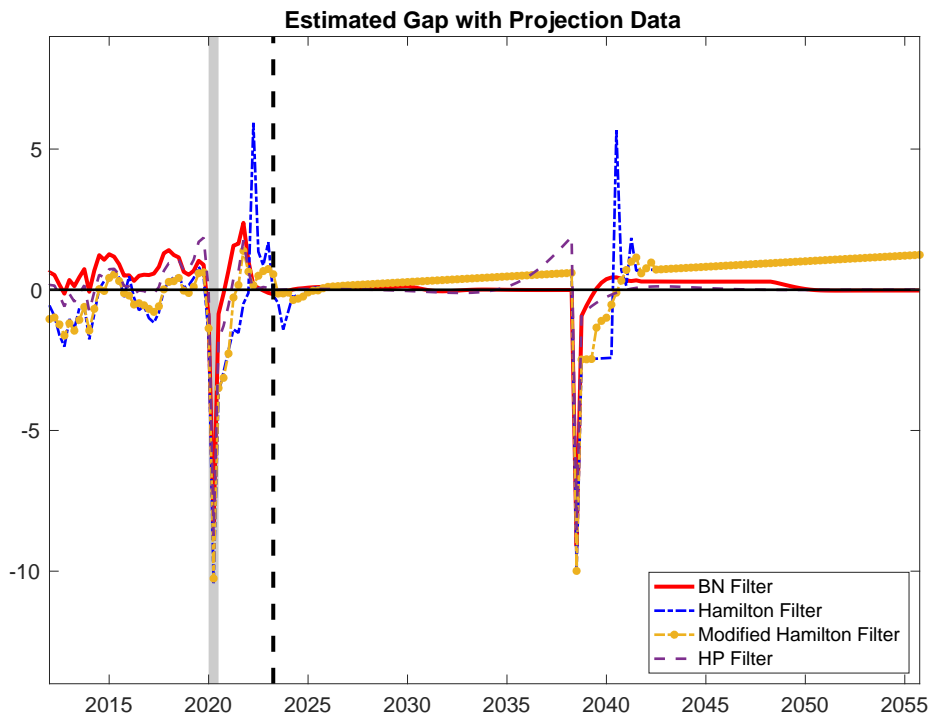
4 Trend-cycle decomposition with projected data

Next, we augment the US real GDP data from 1947Q1 to 2023Q2 with projections of future output growth from 2023Q3 to 2055Q4 that include a simulated Covid-like shock occurring 15 years in the future in 2038Q3 to see how the different methods would forecast the output gap and process an outlier shock for which we know the true impact on trend and cycle. We project that real GDP generally grows from the end of our sample by 0.6% in every quarter. This coincides with a drift of about 2.5% real GDP growth per annum, close to the historical norm. However, in 2038Q3, we also project a (i) a one-time permanent 2.5% reduction in the level of real GDP and (ii) an additional 7.5% transitory reduction in the level of real GDP. Our scenario thus assumes a overall 10% reduction in real GDP in 2038Q3, followed by a partial recovery in 2038Q4. That is, one-quarter of the overall Covid-like shock corresponds to a reduction in trend and the remaining three-quarters of the shock corresponds to a drop in the output gap. We then apply the different methods to the augmented data to forecast the output gap under this scenario. Note that because only three-quarters of the Covid-like shock is permanent, some of the projected reduction in real GDP should be attributed to the trend, not the cycle.

Figure 5 plots the estimated output gaps for the projection data, with the vertical dashed black line marking the beginning of the forecasted output gaps. Just before the start of the projection, it can be seen that the output gap estimates are reasonably consistent with the estimated output gaps reported in Figures 2 and 3 even though estimation now uses the augmented data up to 2055. Specifically, all the estimated output gaps suggest that the COVID-19 recession largely represented a transitory decline in output, consistent with the economic narrative that economic activity that was constrained by mitigation measures in 2020Q2 was largely restored in 2020Q3. Furthermore, as discussed previously, even though both versions of the Hamilton filter produced output gap estimates that were very similar pre pandemic, this does not appear to be the case post pandemic. The differences for the original and modified Hamilton filters are thus entirely driven by the large shock at the onset of the pandemic. To be sure, in the absence of an outlier shock, the original and modified Hamilton filter produce very similar estimates of the output gap, which they did for the pre-pandemic sample, as also shown by Quast and Wolters (2022). However, when large shocks such as the Covid crisis occur, the two versions of the Hamilton filter produce vastly different output gap estimates. In fact, because the modified version of the Hamilton filter averages the gap obtained from $h = 4$ to $h = 12$, the base effects of the outlier quarterly growth rates affect the modified filter starting at four quarters after the large shock (2021Q2), which we can see from Figure 5, is when the estimated gaps for the original and modified Hamilton filters start diverging. Of course, because the base

effects of the large outliers are averaged in the modified version, the modified version does not see the same large spike after eight quarters (2022Q2) as found with the original Hamilton filter.

Figure 5: Forecasted output gaps



Units are 100 times natural log deviation from trend. Estimates are reported for 2012Q1 to 2050Q4 based data from 1947Q1 to 2023Q2 augmented with projected data from 2023Q3 to 2050Q4. The vertical dashed black line denotes 2023Q3, which is the beginning of the forecast period. The shaded bar corresponds to NBER recession dates.

Turning to the forecasts in Figure 5, because the output gap estimates are all close to zero in 2023Q2 and output growth is set to trend growth for the 15 years before the outlier shock in 2038Q3, we would expect the forecasted output gap to converge to zero as the effects of previous shocks dissipate. This is the case for the BN and HP filters. However, for both versions of the Hamilton filter, the forecasted output gap slowly increases over the forecast horizon. This is a function of the fact that the actual growth in the projection of 0.6% per quarter is not identical to what would be predicted based on the regression model, which will largely reflect the average growth rate over the full sample (including the augmented data). In principle, the longer the projection sample, the closer the predicted growth would get. But, in practice, because there is no stochastic variation in projected real GDP for most of the projection (other than the Covid-like shock in 2038, which we discuss below), the regression for the Hamilton filter would become akin to a regression of a linear time trend on lagged linear time trends and constant, which would move towards a singularity in the limit and not actually identify coefficients summing to one with the constant converging to the simulated drift. It is possible to find a precise growth rate for which the Hamilton filter would imply a forecasted output gap that is a flat line close to zero. But this is a knife-edge case and there is no reason the corresponding growth rate would be the same projection of output growth that a policymaker would want to consider. Meanwhile, consistent with nearly identical estimates in the absence

of outliers, this issue with the drifting forecast is not addressed by the modified Hamilton filter, which produces an indistinguishable forecast of the output gap that drifts upwards on the same path as for the original Hamilton filter.

In terms of the simulated Covid-like shock in 2038 in Figure 5, we can see it has very similar effects on the forecasted output gaps as what actually happened with the estimated output gaps around the COVID-19 pandemic. Confirming the idea that the earlier spike in 2022Q2 was spurious, the original Hamilton filter again implies a big mechanical spike exactly two years after the shock, while the modified Hamilton filter smoothes this effect out, but both continue on their upward trajectory a few years after projected output returns to its trend growth rate. Both Hamilton filters also overstate the decline in the output gap in 2038Q3 at more than 10%, even though the actual transitory decline in output was 7.5% in the simulation. Given output growth was -10% in 2038Q3, the Hamilton filters imply a slight increase in trend when the true trend actually fell by 2.5%. Both the HP and BN filters capture an output gap of -7.5% in 2038Q3, but this actually implies the HP filter misses most of the actual decline in trend because the HP filter output gap falls by close to 10% from a spuriously positive level prior to the shock that is similar to what happened with revised estimates based on the HP filter prior to the COVID-19 pandemic. Both the HP and BN filters eventually return to zero in the long run, indeed the HP filter does so somewhat faster than the BN filter. But only the BN filter captures the movements in trend and cycle during the Covid-like shock correctly, while also generally having good long-run properties with its forecast.

5 Conclusion

We have proposed some refinements to the original BN filter in Kamber, Morley, and Wong (2018) and investigated whether the BN filter remains intuitive and reliable in the face of outliers such as occurred with the COVID-19 pandemic. We find that the BN filter remains reliable, with the proposed refinements enhancing how intuitive its output gap estimates are. Comparing against other popular methods, the HP filter also produces intuitive output gap estimates, but its well-known lack of real-time reliability remains a problem. The Hamilton filter, on the other hand, is clearly reliable, but may be ill-suited to deal with outlier type data such as occurred during the pandemic since it mechanically produces a future spike in the estimated output gap following an outlier shock exactly in line with the filter horizon. It also produces unintuitive forecasts of the output gap given plausible projected values of future output growth. From this perspective, especially when considering outlier data such as experienced during the COVID-19 pandemic and also data augmented with long-term projections as often considered in policy settings, the BN filter appears to be the only of the three univariate trend-cycle decomposition methods under consideration that is both intuitive and reliable.

References

- BAI, JUSHAN AND PIERRE PERRON (2003): “Computation and Analysis of Multiple Structural Change Models,” *Journal of Applied Econometrics*, 18 (1), 1–22.
- BARIGOZZI, MATTEO AND MATTEO LUCIANI (2023): “Measuring US aggregate output and output gap using large datasets,” *Review of Economics and Statistics*, 105 (6), 1500–1514.
- BEVERIDGE, STEPHEN AND CHARLES R. NELSON (1981): “A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the ‘Business Cycle’,” *Journal of Monetary Economics*, 7 (2), 151–174.
- COGLEY, TIMOTHY AND JAMES M. NASON (1995): “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research,” *Journal of Economic Dynamics and Control*, 19 (1–2), 253–278.
- HAMILTON, JAMES D. (2018): “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 100 (5), 831–843.
- HARVEY, ANDREW C., THOMAS M. TRIMBUR, AND HERMAN K. VAN DIJK (2007): “Trends and Cycles in Economic Time Series: A Bayesian Approach,” *Journal of Econometrics*, 140 (2), 618–649.
- HODRICK, ROBERT J AND EDWARD C PRESCOTT (1997): “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking*, 29 (1), 1–16.
- KAMBER, GÜNEŞ, JAMES MORLEY, AND BENJAMIN WONG (2018): “Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter,” *The Review of Economics and Statistics*, 100 (3), 550–566.
- LENZA, MICHELE AND GIORGIO E. PRIMICERI (2022): “How to Estimate a Vector Autoregression after March 2020,” *Journal of Applied Econometrics*, 37 (4), 688–699.
- MORLEY, JAMES, DIEGO RODRIGUEZ-PALENZUELA, YIQIAO SUN, AND BENJAMIN WONG (2023): “Estimating the Euro Area Output Gap Using Multivariate Information and Addressing the COVID-19 Pandemic,” *European Economic Review*, 153, 104385.
- MORLEY, JAMES C. (2002): “A State-Space Approach to Calculating the Beveridge-Nelson Decomposition,” *Economics Letters*, 75 (1), 123–127.
- MORLEY, JAMES C., CHARLES R. NELSON, AND ERIC ZIVOT (2003): “Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?” *The Review of Economics and Statistics*, 85 (2), 235–243.
- ORPHANIDES, ATHANASIOS AND SIMON VAN NORDEN (2002): “The Unreliability of Output-Gap Estimates in Real Time,” *The Review of Economics and Statistics*, 84 (4), 569–583.
- QUAST, JOSEFINE AND MAIK H. WOLTERS (2022): “Reliable Real-Time Output Gap Estimates Based on a Modified Hamilton Filter,” *Journal of Business & Economic Statistics*, 40 (1), 152–168.
- RAMSEY, JAMES B AND PHILIP ROTHMAN (1996): “Time Irreversibility and Business Cycle Asymmetry,” *Journal of Money, Credit and Banking*, 28 (1), 1–21.
- WATSON, MARK W. (1986): “Univariate Detrending Methods with Stochastic Trends,” *Journal of Monetary Economics*, 18 (1), 49–75.