

Quast, Josefine; Wolters, Maik H.

Conference Paper

Reliable Real-time Output Gap Estimates Based on a Modified Hamilton Filter

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Applied Macroeconometrics I, No. A13-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Quast, Josefine; Wolters, Maik H. (2019) : Reliable Real-time Output Gap Estimates Based on a Modified Hamilton Filter, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Applied Macroeconometrics I, No. A13-V1, ZBW - Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/203535>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Reliable Real-time Output Gap Estimates Based on a Modified Hamilton Filter

Josefine Quast*

University of Jena

Maik H. Wolters[†]

University of Jena, Kiel Institute for the World Economy, IMFS at Goethe University Frankfurt

February 26, 2019

Preliminary Draft.

Please do not quote or distribute without the authors' permission.

Abstract

We propose a simple modification of the time series filter by Hamilton (2018b) that yields reliable and economically meaningful real-time output gap estimates. The original filter relies on 8-quarter ahead forecasts errors of an autoregression. While this approach yields a cyclical component of GDP that is hardly revised with new incoming data due to the one-sided filtering approach, it does not cover typical business cycle frequencies evenly, but short business cycles are muted and medium length business cycles are amplified. Further, the estimated trend is as volatile as GDP and can thus hardly be interpreted as potential GDP. A simple modification that is based on the mean of 4- to 12-quarter-ahead forecast errors shares the favorable real-time properties of the Hamilton filter, but leads to a much better coverage of typical business cycle frequencies and a smooth estimated trend. Based on output growth and inflation forecasts and a comparison to revised output gap estimates from policy institutions, we find that real-time output gaps based on the modified Hamilton filter are economically much more meaningful measures of the business cycles than those based on other simple statistical trend-cycle decomposition techniques such as the HP or the Bandpass filter.

Keywords: Output gap, potential output, trend-cycle decomposition, Hamilton filter, real-time data, inflation forecasting

JEL-Codes: C18, E32, E37

*josefine.quast@uni-jena.de

[†]maik.wolters@uni-jena.de

1 Introduction

Ever since Orphanides and van Norden (2002) have provided evidence for the poor real-time performance of commonly used output gap estimation methods, the debate regarding the reliability of output gap estimates has been vivid.¹ Edge and Rudd (2016) and Champagne et al. (2018) have shown that the reliability of Federal Reserve and Bank of Canada staff output gap estimates has increased since the mid-1990s, while the reliability of purely statistical detrending procedures, like, for example, the Hodrick-Prescott (HP) filter, continues to be poor. In a recent article, Hamilton (2018b) has proposed a new regression based filter for detrending time series as an alternative to the HP filter. Hamilton proposes to use the 8-quarter forecast error of a projection based on an AR(4) model as the cyclical component of a macroeconomic time series. The 8-quarter horizon is chosen, because the primary reason for forecast errors over such a horizon is cyclical factors such as whether a recession occurs over the next two years rather than large trend changes. The filter produces a stationary cycle for a wide range of time series, but suffers much less from end-of-sample bias than the HP filter and the other trend-cycle decomposition methods considered by Orphanides and van Norden (2002), because it is a one-sided filter.²

Hamilton applies the filter to GDP and shows that its cyclical component turns negative during NBER defined recessions and positive during expansions. Hence, it is a potentially useful measure of the business cycle and the approach can possibly solve the long-standing problem of the unreliability of real-time output gap estimates. Based on spectral density analysis we show, however, that the filter does not cover typical business cycle frequencies from 6 to 32 quarters evenly. Cycles of lengths between 10 and 20 quarters are amplified substantially relative to longer and shorter cycles. The latter are muted almost completely. Further, the extracted GDP trend is not smooth, but is as volatile as GDP itself, so that it cannot be interpreted as measuring potential GDP.

We propose a simple modification of the Hamilton filter that shares its favorable real-time properties, but leads to a more even coverage of typical business cycle frequencies. Rather than using a fixed 8-quarter forecast horizon, we take a simple average of filtered values based on forecast horizons ranging from 4 to 12 quarters. In this way short, medium and long business cycles are covered more evenly, so that the original detrending method by Hamilton is turned into one that is useful for detrending and business cycle filtering.³ The modified filter further avoids unreasonable spikes and an unreasonably large amplitude of the cyclical component of GDP and yields a smooth trend of GDP. Hence, an interpretation of the trend as potential GDP and of the cyclical component as an output gap is possible. The modified filter is still centered around the 8-quarter horizon proposed by Hamilton and it is similarly easy to compute as Hamilton's original approach.

We analyze the reliability of US output gaps computed with the modified Hamilton filter and compare it to output gaps computed with the original one and the HP and the Bandpass (BP) filter as examples

¹Other papers have confirmed the unreliability of real-time output gap estimates produced with common statistical methods for other economies (see, e.g., Bernhardsen et al., 2005; Cayen and van Norden, 2005; Marcellino and Musso, 2011).

²Hamilton (2018b) argues that that his approach does not only solve the end-of-sample bias of the HP filter, but also two additional drawbacks: the creation of spurious cycles and an ad hoc smoothing parameter that is at odds with the one chosen by a formal statistical procedure (see also Baxter and Stockman, 1989; Baxter, 1991; King and Rebelo, 1993; Harvey and Jaeger, 1993; Cogley and M.Nason, 1995; Canova, 1998, for problems with the HP filter).

³The Hamilton filter has already been used in the literature for detrending various macroeconomic variables (see, e.g., Ahn and Hamilton, 2016; Bordo and Siklos, 2017; López-Salido et al., 2017; Richter et al., 2017; Van Zandweghe, 2017; Danielsson et al., 2018; Richter et al., 2018). In all these studies the obtained trends and gaps are interpreted economically. For this, an even coverage of the relevant frequencies seems to be important. Hamilton himself uses the filter to construct a cyclical measure of global economic activity (Hamilton, 2018a).

of commonly used simple statistical trend-cycle decomposition techniques. In particular, we answer two research questions. First, are the different real-time output gap estimates reliable, i.e. are subsequent revisions small? Second, are they economically meaningful measures of the business cycle? It goes without saying that a positive answer to the first question needs not necessarily entail a positive answer to the second one. Fortunately, we find that the answer to both questions is yes for the modified Hamilton filter.

To evaluate the revision properties, we compute output gaps using real-time data vintages and compare them to those based on revised data. We find that the revisions of output gaps based on the modified and the original Hamilton filter are small and mainly due to revisions in the underlying data. There are two reasons for the small revisions. First, the Hamilton filter is primarily a one-sided filter, though Hamilton recommends to use the whole available sample to estimate the AR(4) parameters. Second, parameters of univariate AR models for log real GDP are particularly stable in comparison to parameters of more complicated multivariate models. Contrary to that, HP and BP filter estimates suffer from large end-point problems due to their two-sided nature and exhibit revisions as large as the gaps themselves.

Evaluating the meaningfulness of output gap estimates and comparing competing output gap estimates in this regard is difficult, because there is no clear benchmark and the true cycle is unknown. We analyze the meaningfulness of output gaps from different angles to achieve nevertheless convincing results.

First, we compare real-time output gap estimates to revised output gap estimates of the Federal Reserve, the Congressional Budget Office, the IMF and the OECD. While there is no true cycle that can be used as a benchmark, important policy institutions should have at least in retrospect expert knowledge on the size and length of past business cycle phases. At the very least such expert benchmarks can be helpful in measuring business cycle characteristics that matter from a practitioner's perspective. We find that the correlation relationship among the multiple ex post expert gaps and the real-time output gap based on the modified Hamilton filter is significantly stronger than that of those based on the HP or the BP filter.

Second, we test whether an output gap estimate has predictive content for output growth and inflation. For example, if output is below potential, this implies that output growth should increase in the future so that output reverts back to potential. Further, output gaps should have according to Phillips curve models implications for future inflation.

While we do not find significant differences in output growth forecasting accuracy for the full sample, the output gap based on the modified Hamilton filter leads to significantly better forecasts compared to the HP and BP filter when excluding the Great Recession. Evaluating out-of-sample inflation forecasts by means of a standard Phillips curve forecasting equation manifests the finding in the literature that no statistically meaningful distinction can be made between differently filtered gap measures (see e.g. Edge and Rudd, 2016; Champagne et al., 2018; Kamber et al., 2018). This reflects the general difficulty of beating univariate inflation forecast models with output gap based models (Atkeson and Ohanian, 2001; Fisher et al., 2002; Orphanides and van Norden, 2005; Stock and Watson, 2007, 2008), rather than output gap measurement problems. It is, however, noteworthy that the modified Hamilton filtered real-time output gap shows small though insignificant gains in inflation forecasting accuracy compared to the HP and BP filtered output gaps.

Overall, we find that output gap estimates based on the modified Hamilton filter have favorable real-time properties and are meaningful measures of the business cycles with much fewer drawbacks than

those based on other simple statistical trend-cycle decomposition methods.

In addition to the literature on the real-time reliability of output gaps, the paper by Kamber et al. (2018) is closely related to our work. They use a modified Beveridge-Nelson decomposition to estimate output gaps. The method is related to the Hamilton filter, because it also uses an autoregression to estimate the trend. Output gaps computed with that method have a similarly good real-time reliability as the Hamilton filtered output gaps, because of the stability of estimated autoregressive parameters. The approach of Kamber et al. (2018) is on the one hand less ad hoc, because they estimate the trend based on long-horizon conditional expectation, rather than imposing a fixed horizon of 8 quarters. Further, the long-horizon conditional expectations are measured using data until period t rather than $t - 8$, so that trend changes between $t - 8$ and t are accounted for. By contrast, when using the Hamilton filter one implicitly assumes that there are no relevant trend changes during this period. On the other hand, the computation of the Hamilton filter and also the modified version proposed in this paper is much simpler and more intuitive. The standard Beveridge-Nelson decomposition yields an output gap that is completely at odds with standard business cycle facts and therefore an algorithm with several steps has to be run to get the modified Beveridge-Nelson decomposition that makes sure that most GDP dynamics are attributed to the cycle rather than trend changes. For the Hamilton filter the computation of forecast errors based on simple autoregressions is instead sufficient.⁴

2 A Simple Modification of the Hamilton Filter

Hamilton (2018b) proposed to use the 8-quarter forecast error of a projection based on an AR(4) model as the cyclical component of a macroeconomic time series. Hence, to detrend a macroeconomic time series, y_t , the following simple autoregression can be estimated by OLS:

$$y_t = \beta_0 + \beta_1 y_{t-8} + \beta_2 y_{t-9} + \beta_3 y_{t-10} + \beta_4 y_{t-11} + v_t. \quad (1)$$

The cyclical component is given by the residual \hat{v}_t :

$$\hat{v}_t = y_t - \hat{\beta}_0 - \hat{\beta}_1 y_{t-8} - \hat{\beta}_2 y_{t-9} - \hat{\beta}_3 y_{t-10} - \hat{\beta}_4 y_{t-11}. \quad (2)$$

Applied to log quarterly real GDP it is tempting to interpret \hat{v}_t as an output gap. However, Hamilton (2018b) and Schöler (2018) remark that cycles of 8, 4, 8/3 and 2 quarters are completely eliminated for the special case of the difference filter to which the filter reduces when being applied to a random walk. Hence, typical business cycle frequencies between 6 and 32 quarters (Burns and Mitchell, 1946; Stock and Watson, 1999a) are not covered evenly and especially short business cycle frequencies around two years are eliminated or considerably dampened.

To get an output gap that covers business cycle frequencies from 6 to 32 quarters more evenly, we propose a simple modification of the Hamilton filter. Rather than using a fixed 8-quarter horizon, we propose using an equally weighted average of forecast errors based on 4- to 12-quarter ahead projections

⁴In addition to the Hamilton filter and the Beveridge-Nelson filter by Kamber et al. (2018) also carefully specified Phillip curves can yield a reliable real-time output gap as demonstrated by Gruen et al. (2005) for Australian data.

to estimate the output gap, \tilde{y}_t :

$$\tilde{y}_t = 1/9 \sum_{i=4}^{12} \hat{v}_{t,i}, \quad \text{with} \quad (3)$$

$$\hat{v}_{t,i} = y_t - \hat{\beta}_{0,i} - \hat{\beta}_{1,i}y_{t-i} - \hat{\beta}_{2,i}y_{t-i-1} - \hat{\beta}_{3,i}y_{t-i-2} - \hat{\beta}_{4,i}y_{t-i-3}. \quad (4)$$

In the following we show based on analysis in the frequency domain that this indeed mitigates the problematic cyclical properties of the original filter.

The spectral density of a covariance stationary process y_t at frequency $\omega \in [0, 2\pi]$ is given by the Fourier transform of the autocovariance function (see, e.g., Canova, 2011):

$$SD_y(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} ACF_y(\tau) e^{-i\omega\tau}, \quad (5)$$

where $i = \sqrt{-1}$ and ω is measured in radians. A time-invariant filter with absolutely summable weights can be written as a two-sided moving average:

$$x_t = \sum_{j=-\infty}^{\infty} b_j y_{t-j}. \quad (6)$$

The Power Transfer Function (PTF) $|B(\omega)|^2 = |b(e^{-i\omega})|^2$ measures the squared gain of such a linear filter, i.e. the variance each frequency ω contributes to the filtered series x_t compared to its variance contribution in the original series y_t :

$$SD_x(\omega) = |B(\omega)|^2 SD_y(\omega). \quad (7)$$

To filter business cycles the PTF should take a value of 1 for the business cycle frequencies of cycles from 6 to 32 quarters, which correspond to radians $\omega = 2\pi/6$ to $\omega = 2\pi/32$ and a value of zero for all other frequencies.

The frequency response for linear filters as in equation (6) is given by:

$$B(\omega) = b_0 + \sum_{-j=-\infty, j \neq 0}^{\infty} b_j \cos(\omega j). \quad (8)$$

For the Hamilton filter the weights b_{-1}, \dots, b_{-4} are given by the estimated coefficients $\hat{\beta}_1, \dots, \hat{\beta}_4$ from equation (2), $b_0 = 1$ and all other weights equal zero. From this the Power Transfer Function can directly be computed as the square modulus of $B(\omega)$.

Figure 1 shows the PTFs for the original and the modified Hamilton filter. The regression coefficients of the AR(4) process (equation 1) have been computed based on quarterly U.S. log real GDP data from 1947Q1 to 2017Q4. The grey shaded areas indicate cycle lengths that are typically associated with business cycle fluctuations. We plot the PTFs in the frequency (panel A) as well as in the time domain (panel B) to facilitate the reading of the figure.⁵ Showing a time domain representation in addition to the usually shown frequency domain representation is particular important as in the latter high frequency fluctuations take up more than half of the graph which might give the false impression that these are particularly important. These are irregular fluctuations though, i.e. noise, with little practical relevance.

⁵A Fourier frequency ω is associated with cycle length $p = 2\pi/\omega$.

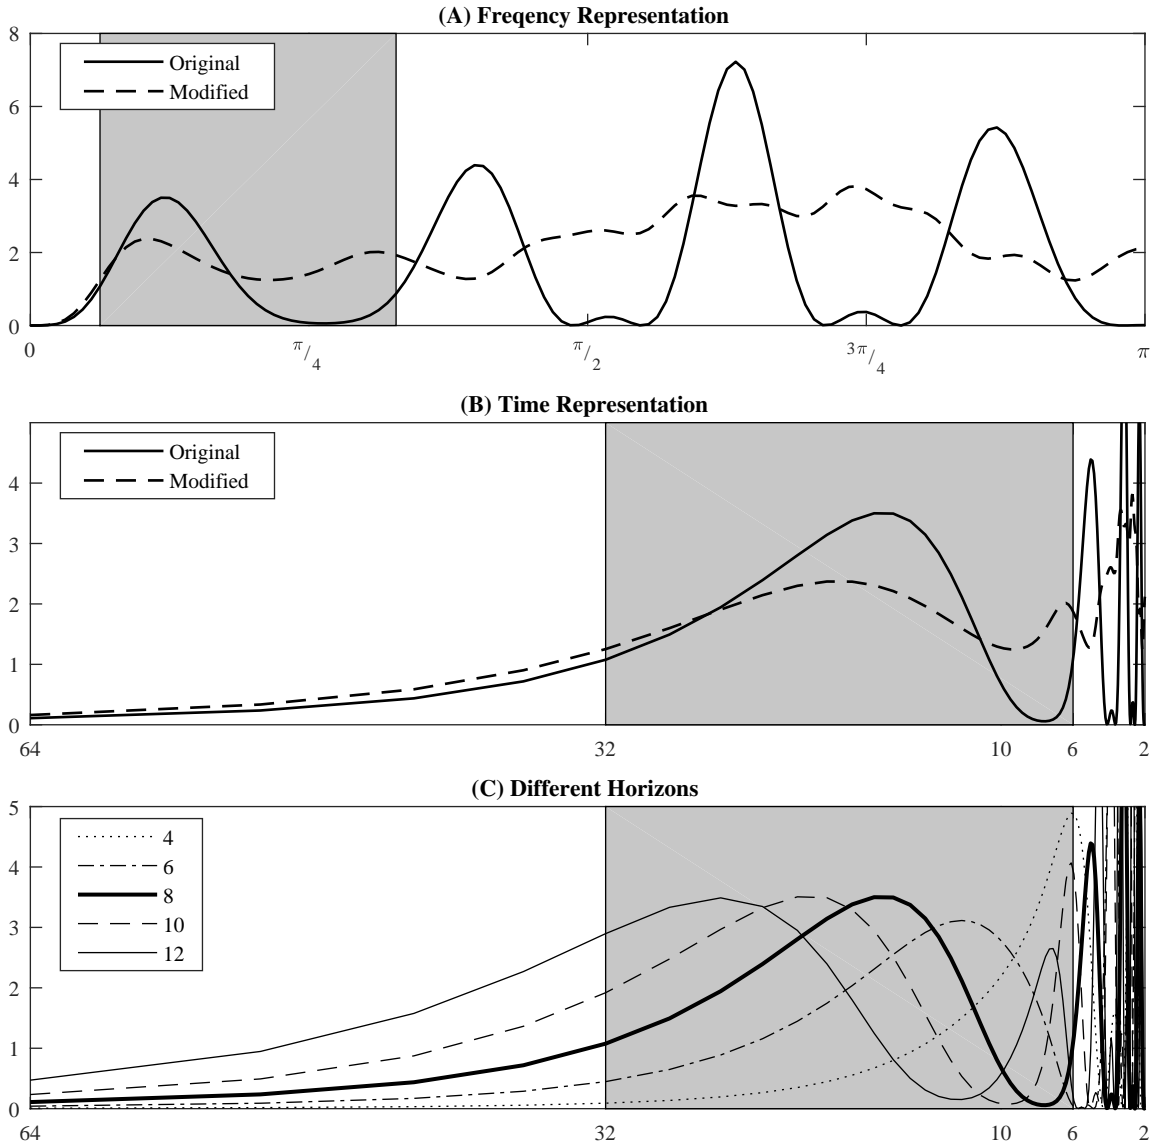


Figure 1: Power Transfer Functions in Frequency and Time Representation

It is apparent that the original Hamilton filter suppresses and dampens business cycles ranging from 6 to 10 quarters. Hence, short business cycle frequencies are not present in a Hamilton filtered output gap. On the other hand, cycles with a duration of 10 to 20 quarters are amplified substantially compared to short cycles and long cycles of 20 to 32 quarters. Panel C shows PTFs of the Hamilton filter for different forecast horizons. Short horizons emphasize short business cycles and mute long business cycles in the output gap, while it is the other way around for long horizons. Hence, taking an average over different horizons leads to a more even coverage of business cycle frequencies as shown by our proposed modified filter in panels A and B. Ideally, only frequencies between 6 and 32 quarters would be present in the output gap. Yet, only the two-sided Bandpass filter is able to get close to such an ideal filter (see Christiano and Fitzgerald, 2003, for PTFs of the approximate Bandpass filter). However, such is only obtainable in retrospect using two-sided filtering, but not in real-time applications.

The proposed modified filter achieves a very good real-time reliability due to its one-sided nature, but covers business cycle frequencies more evenly than the original approach by Hamilton (2018b). Yet, it is very simple to compute.

Further, the smoothing of different PTF peaks and troughs improves substantially the treatment of

high frequency noise in the obtained output gap and trend measures. While the gap and the trend include erratic movements each quarter for the original Hamilton filter, those from the modified Hamilton filter are very smooth as shown in Figure 2 (Panel A and B). This difference is most clearly visible for the estimated trend growth rate (Panel C) that is as volatile as the growth rate of GDP itself for the original Hamilton filter. Due to the structure of equation (2) the period t trend growth rate of GDP and the period $t - 8$ GDP growth rate are very similar (correlation coefficient: 0.93). Hence, the trend measure extracted via the original Hamilton filter can hardly be interpreted as potential GDP, which in turn makes the economic interpretation of the resulting output gap difficult. By contrast, the extracted trend based on the modified Hamilton filter is very smooth and its growth rate varies slowly over time like those of other often used trend measures as, for example, based on the HP or the BP filter. Further, the smooth trend of the modified Hamilton filter leads to a smoother output gap compared to the original Hamilton filter.

3 Real-time Reliability

We analyze the real-time reliability of different output gaps in the spirit of Orphanides and van Norden (2002) by computing output gaps based on real-time data vintages and comparing them to those based on revised data. We obtain real-time data on quarterly real GDP from the Federal Reserve Bank of Philadelphia's real-time data set. The first data vintage is from 1965Q4 and the last data vintage is from 2018Q1. Data in all vintages starts in 1947Q1 and ends one quarter before the publication date of the data vintage, i.e. in 1965Q3 for the first and in 2017Q4 for the last data vintage. We take logs of real GDP and apply the various filtering techniques to the real GDP vintages to get real-time output gap vintages. We then collect the end-point estimates of these vintages to get a real-time output gap series for each filtering method.

We use the last data vintage from 2018Q1 as our measure of final revised data and define the output gap revision as the difference between the final and the real-time estimate. To make sure that a comparison of real-time and revised data is not biased by the last data vintages in which real-time and revised data are very close to each other, we follow Orphanides and van Norden (2002) and discard the last two years of observations. Hence, our results for the revised gaps are based on the 2018Q1 vintage with gaps being estimated until 2015Q4. This proceeding is reasonable since most revisions take place in the first two years after the initial publication of GDP (Orphanides and van Norden, 2002; Edge and Rudd, 2016). Hence, the sample of output gaps that we study starts in 1965Q3 and ends in 2015Q4.

To distinguish between data revisions and revisions due to the filter-induced structure, we compute a quasi real-time series. That is, we estimate output gaps recursively as in the real-time estimation, but based on the final revised data. By that we are able to isolate the impact of pure data revisions—defined as the quasi real-time output gap minus the real-time output gap—as the estimates of the real-time and the quasi real-time series are based on data covering the exact same time periods.

Figure 3 shows that revisions to output gaps computed with the original and the modified Hamilton filter are small relative to the amplitude of the output gap and are mainly caused by data revisions. By contrast final and real-time output gap estimates of the HP and the BP filter differ markedly from each other. Revisions for these are of the same sizes as the output gaps themselves confirming the results by Orphanides and van Norden (2002) for an updated sample. The revisions for the HP and BP filter are to a large extend filter-induced revisions, while data revisions play only a minor role.

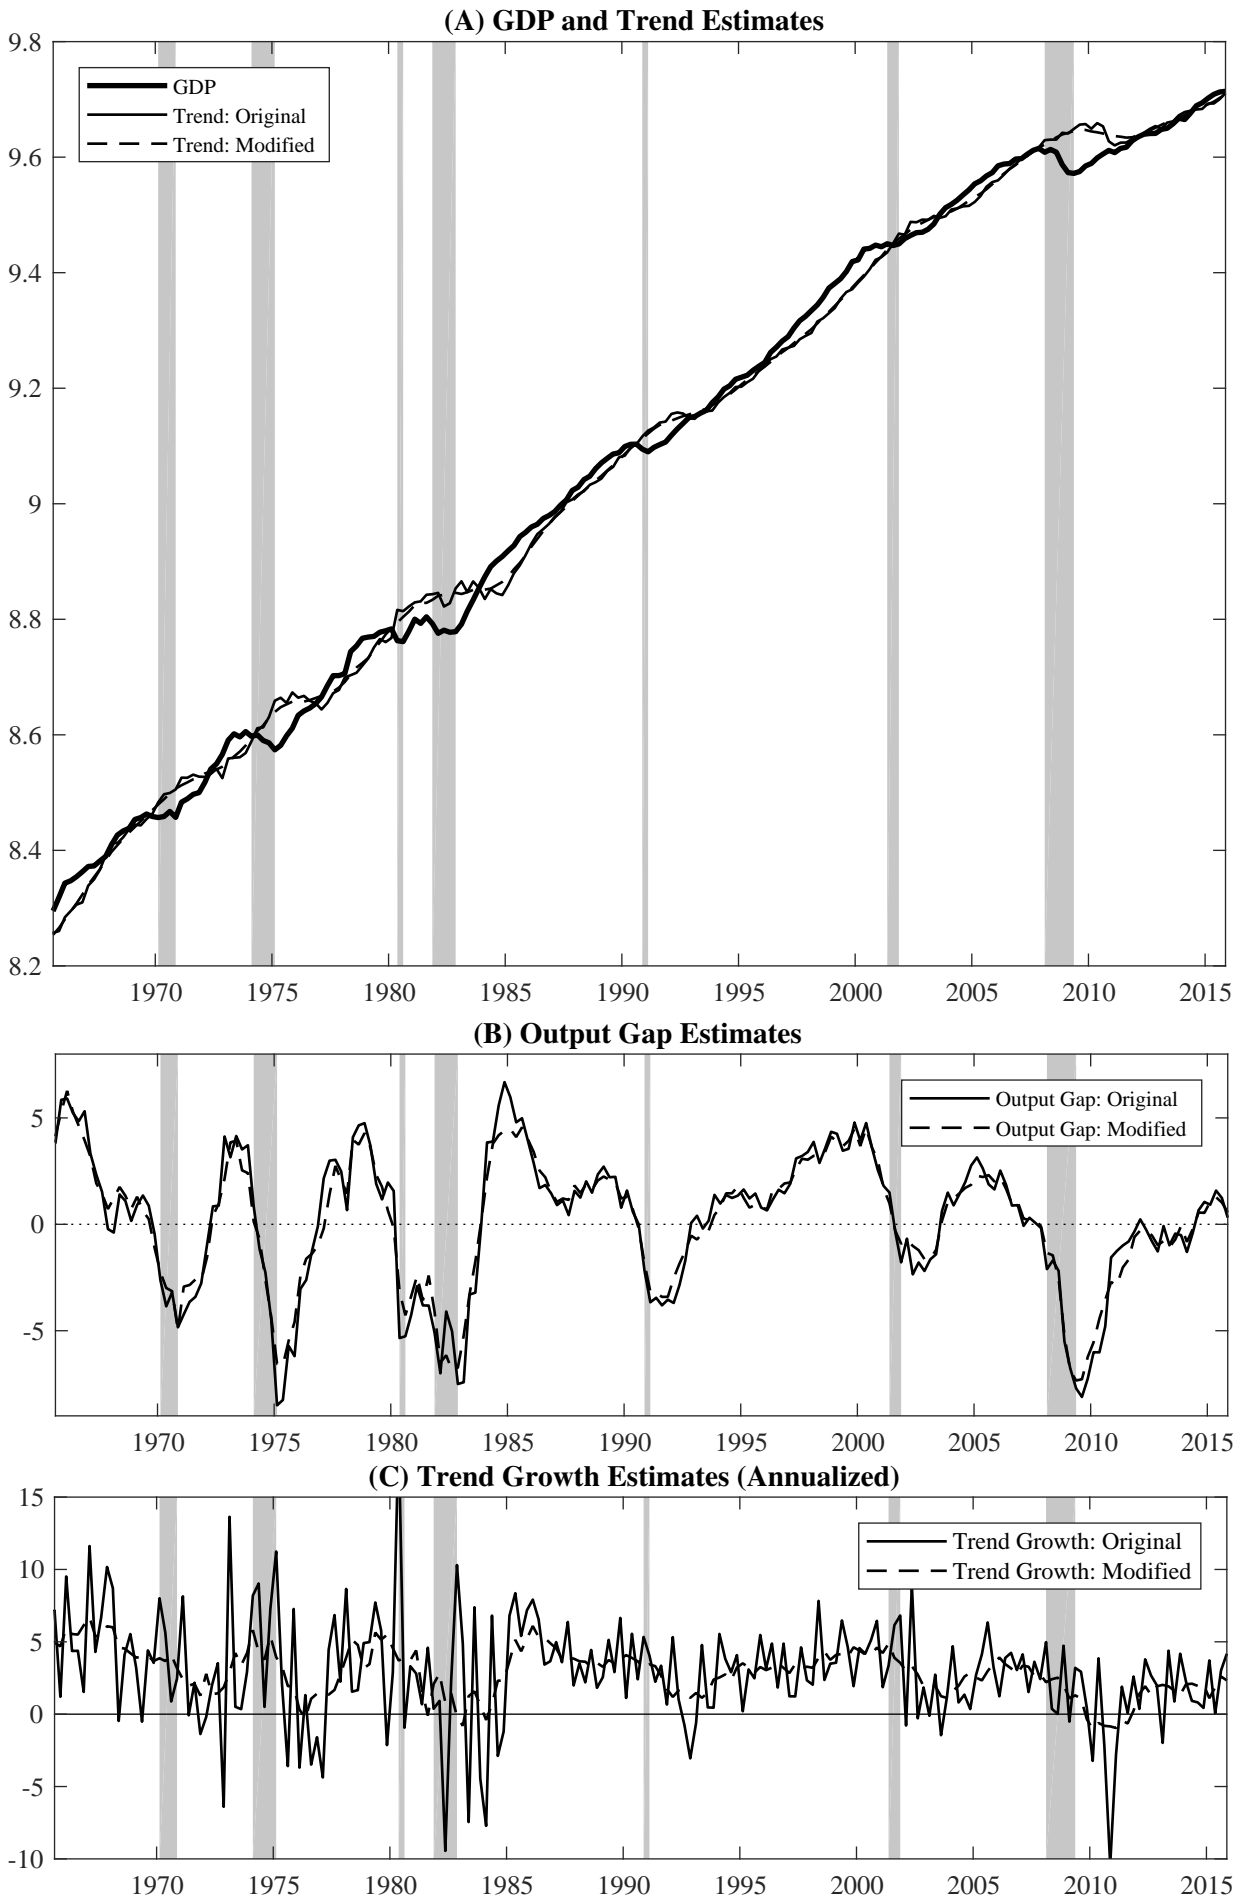


Figure 2: Trend, Output Gap and Trend Growth: Original and Modified Hamilton Filter

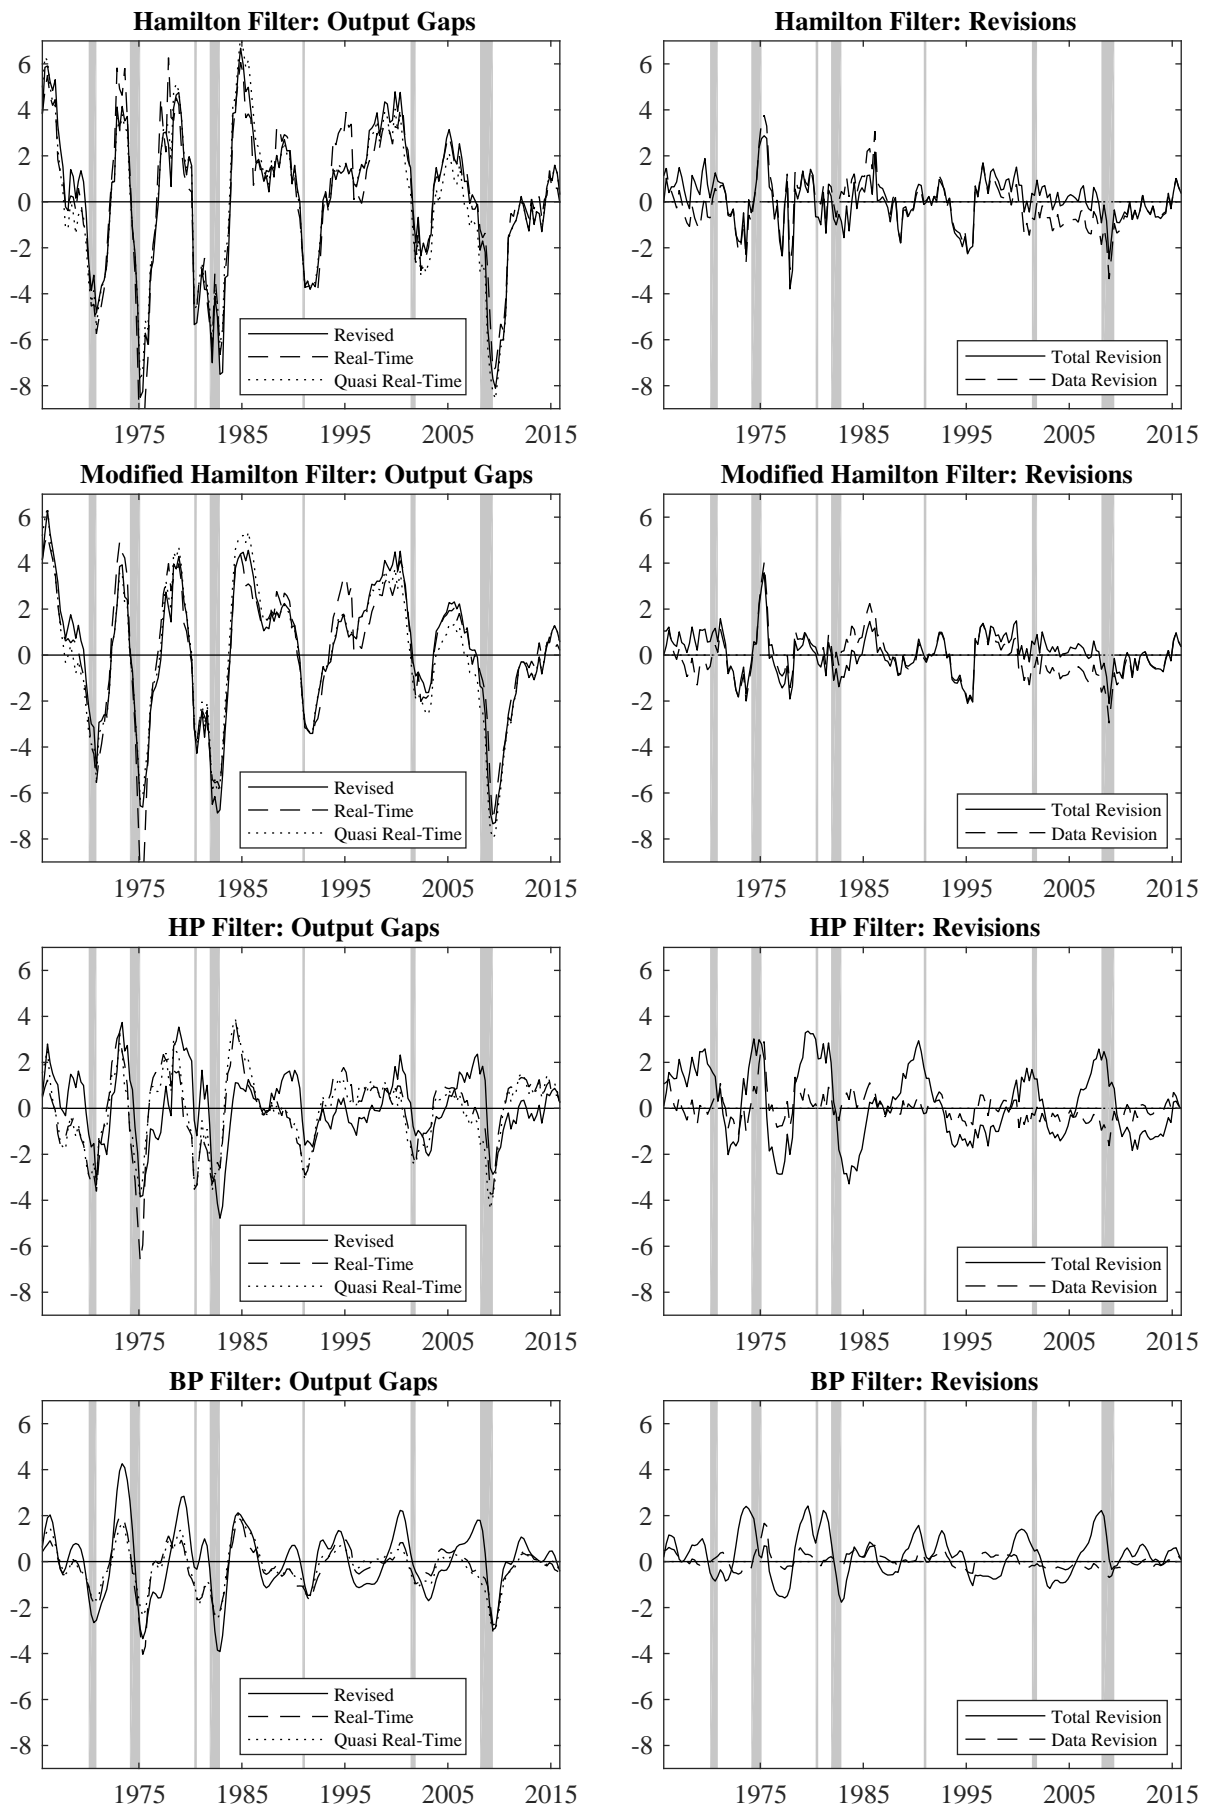


Figure 3: Output Gaps and Revisions

Table 1 presents summary statistics for the output gap revisions. The upper part of the table shows statistics on total revisions. The Hamilton filtered output gap has the smallest mean error with a value of 0.04 and the BP filter the largest with a value of 0.29, which is still not very large given that the BP filtered output gap fluctuates between ± 4 . The standard deviation is smallest for the modified Hamilton filter taking a value of 0.86 and almost twice as large for the HP filter with the original Hamilton and the BP filter being in between. However, this value is not very informative without comparing it to the standard deviation of the revised estimate of the output gap, because the amplitudes of the two versions of the Hamilton filtered output gap are quite a bit larger than the amplitudes of the other two output gaps. Such a comparison is given by the noise-to-signal ratios. The first measure of the noise-to-signal ratio compares the standard deviation of the revision to the standard deviation of the final revised output gap estimate. This value is only 0.3 for the two versions of the Hamilton filtered output gap, while it is with values of 0.6 and almost 1.0 much larger for the BP and HP filtered output gaps. In principle, a noise-to-signal ratio based on dividing the RMSE of the gap revision by the standard deviation of the final revised output gap is more informative, because it reflects biases in the real-time output gaps. However, the noise-to-signal ratios based on both measures are almost the same, because the biases of all four real-time output gap estimates are small. Further, we report the differences between the noise-to-signal ratios relative to the results obtained for the modified measure in columns 6 and 8. To analyze whether they are statistically significant, we proceed similarly to Edge and Rudd (2016) and compute empirical distributions of the noise-to-signal ratios based on a naive block bootstrapping procedure with replacement (Kilian and Lütkepohl (2017)).⁶ Compared to the HP and BP based ratios, the total revision series based on the modified measure is significantly less noisy. There are small absolute, yet insignificant, gains compared to the original Hamilton filter. Finally, we also report the fraction of observations in which the final revised and the real-time output gap estimates have opposite signs. This is the case for only 3% and 6% of observations for the modified and the original Hamilton filter, respectively, but for 19% and 31% of observations for the BP and HP filtered output gaps, respectively.

Table 1: Statistics of Total Revisions and Data Revisions of Output Gaps

	Mean	SD	RMSE	Noise-Signal Ratios				Opposite Sign
				SD	Diff	RMSE	Diff	
Total Revisions								
Hamilton	0.04	1.00	1.00	0.30	0.01	0.29	0.01	0.06
Modified	0.06	0.86	0.86	0.29	-	0.29	-	0.03
HP	0.17	1.58	1.59	0.97	0.68***	0.97	0.68***	0.31
BP	0.29	0.92	0.96	0.59	0.30***	0.62	0.33***	0.19
Data Revisions								
Hamilton	-0.12	1.10	1.10	0.32	0.00	0.33	0.00	0.07
Modified	-0.11	0.96	0.96	0.32	-	0.32	-	0.06
HP	0.01	0.60	0.60	0.37	0.05	0.37	0.05	0.04
BP	0.04	0.34	0.34	0.22	-0.10**	0.22	-0.10**	0.08

Notes: *, **, and *** denote significance on the 10, 5, and 1 % significance level.

The lower part of the table shows statistics on the share of the revisions that is caused by data revisions. This share is computed by subtracting the real-time output gap estimates from the quasi real-time estimates. For the two versions of the Hamilton filter, the standard deviations and the noise-to-signal ratios based on data revisions and total revisions are very similar. This reflects that almost all revisions

⁶We use 5000 bootstrap replications and a block size of 4 to compute the empirical distributions of the noise-to-signal ratios.

of these output gap measures are due to data revisions. For the HP and BP filter, the standard deviations, RMSEs and the two noise-to-signal ratio measures are much smaller than for total revisions, reflecting that data revisions are relatively unimportant. As the differences in the noise-to-signal ratios associated with the data revisions are insignificant or even in favor of the BP based measure, the significantly higher presence of noise in the total revision series is merely due to the choice of the filtering procedure. Further, there are relatively few sign switches between the real-time and the quasi real-time output gap estimates for all four methods.

Overall, we find that the output gaps based on the modified and the original Hamilton filter are by far the most reliable real-time output gap measures, with the former having an even slightly lower noise-to-signal ratio than the latter. It is also worth noting that the BP filter performs better than the popular HP filter which shows the least reliable performance among the four output gap measures.

The reason for the high reliability of the two Hamilton filtered output gaps with respect to revisions is that they rely on one-sided filters. Revisions can occur only for two reasons: data revisions and changes in the estimated parameters of the AR(4) processes. Data revisions lead only to small revisions of the output gap estimates as discussed above. Regarding the estimated AR(4) parameters, Figure 4 shows for the original Hamilton filter how they change over time via the recursive estimation procedure and how they converge to the full sample estimates. While there are some changes, overall the parameters are rather stable. When looking at the sum of the AR-coefficients there are almost no changes at all. GDP follows a unit-root or near unit-root process for all subsamples. Regarding the estimated constant, there are some changes until the subsample ending in 1985. Afterwards, the estimates decrease somewhat reflecting the growth slowdown in the 1980s.

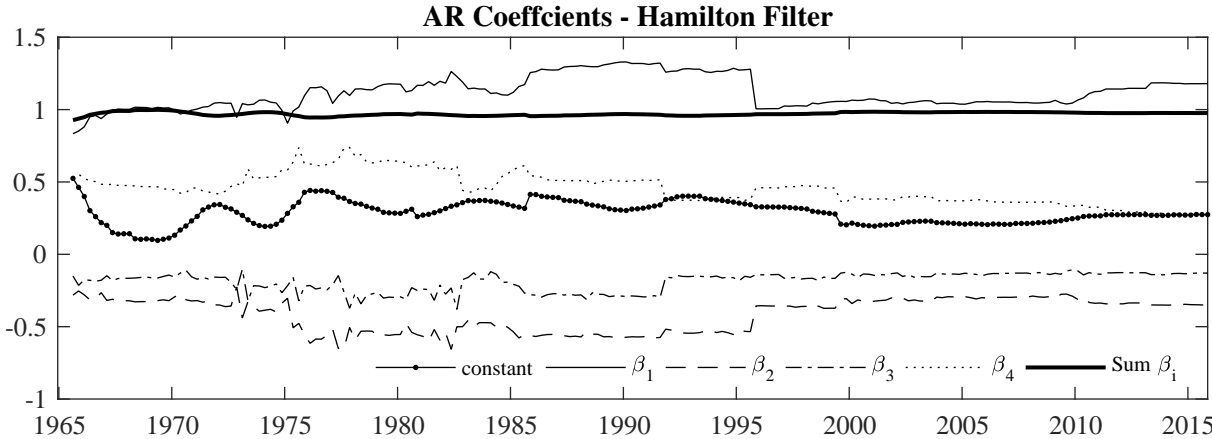


Figure 4: Recursive estimates of the AR(4) parameters of the original Hamilton filter

Determining the current output gap is an important element in the policy process at policy institutions like, for example, central banks. To understand to which extent the revisions of output gaps matter in the policy process of a central bank one can think about the implied interest rate distortions. For example, when using the standard Taylor rule, the interest rate reaction to the output gap is 0.5. Hence, the interest rate distortions through output gap revisions are given by multiplying the output gap revisions with 0.5. The usage of real-time compared to revised output gap estimates would imply interest rates that are on average 0.08 percentage points too low for the HP filter, 0.14 percentage points too low for the BP filter, but only 0.02-0.03 percentage points too low for the two Hamilton filter based output gaps. More interesting might be the maximal mistake when using real-time compared to revised output gap estimates. The numbers (1.90, 1.79, 1.20, 1.68 and percentage points for the original Hamilton, the

modified Hamilton, the BP and the HP filter, respectively) are however not very meaningful. Due to the higher amplitude of the two versions of the Hamilton filter, the output gap coefficients in a policy rule would be chosen smaller. The standard deviations of the revised Hamilton filtered output gaps are approximately twice as large (3.26 and 2.92) for the original and the modified version) compared to the HP filtered (1.52) and the BP filtered (1.45) output gaps. Hence, we adjust the numbers for the BP and the two Hamilton filter to make them comparable to the HP filter by dividing by the standard deviation of the respective output gap and multiplying the result with the standard deviation of the HP filtered output gap. This results in maximal interest rate setting mistakes of 1.68 percentage points for the HP filtered (no adjustment), 1.26 for the BP filtered output gap ($1.20/1.45 \cdot 1.52$) and only 0.89 ($1.90/3.26 \cdot 1.52$) for the original and 0.93 ($1.79/2.92 \cdot 1.52$) for the modified Hamilton filtered output gap. Thus, the maximal interest rate mistake implied by output gap revisions would be almost twice as large for the HP filtered compared to using one of the two Hamilton filter based output gaps.

4 Economic Meaningfulness of Output Gap Estimates

While there is evidence for a superior real-time performance of the two versions of the Hamilton filtered output gap compared to the HP and BP filtered ones, it is yet unclear whether the Hamilton filtered gaps are economically meaningful. A first basic requirement of the output gap being stationary is fulfilled as long as the original time series is integrated of order 4 or less (Hamilton, 2018b).⁷ This is clearly the case for real GDP. Further, at least the modified Hamilton filter yields a smooth trend estimate and an output gap that covers the relevant business cycle frequencies relatively evenly.

A second requirement for a meaningful output gap is that the method used is able to successfully disentangle trend and cycle. Hamilton’s method is able to do so to the extent that no large trend changes occur between periods $t - 8$ and t . Hamilton (2018b) justifies the exclusion of trend changes over a two-year horizon arguing that the primary reason for forecast errors at this horizon are transitory factors such as whether a recession occurs and the timing of recoveries. Other output gap estimation methods are based on a similar a-priori belief that the cyclical volatility is much larger than the volatility of the trend. An example is the standard smoothing parameter of 1600 of the HP filter. A statistical formalization of the choice of the HP smoothing parameter would result in a much lower cyclical relative to trend volatility.⁸ Similarly, when using the Beveridge-Nelson decomposition, only a version with a strong prior belief that the volatility of the cyclical component is much higher than the one of the trend component leads to sensible output gap estimates, while without imposing such a prior belief, changes in the trend component dominate (Kamber et al., 2018). Hence, the a-priori choice that over a two-year horizon trend changes are much less important than cyclical fluctuations is in line with other output gap estimation methods and the modified version of the Hamilton filter accounts for trend changes to a certain extent already after one rather than two years.

After all, the appropriateness of a filtering technique depends on the researcher’s objective. For output gap estimations the objective is often to match important historical business cycle episodes. Computing correlations of the quarter-on-quarter change of the different real-time output gap measures with

⁷Some well-known output gap estimation methods fail to do so. Besides deterministic detrending procedures also DSGE model-based output gaps can be nonstationary (Wolters, 2018).

⁸The standard HP smoothing parameter of 1600 is based on the assumption that the volatility of the cyclical component is $\sigma_c = 5$ and the volatility of the trend component is $\sigma_t = 0.125$ resulting in $\lambda = \sigma_c^2 / \sigma_t^2 = 1600$. A statistical formalization of the choice of λ leads to a value of 0.245 for U.S. quarterly GDP (Hamilton, 2018b).

a dummy variable that takes the value of 1 during NBER defined expansions and 0 during NBER defined recessions, yields a correlation of 0.46 for the original Hamilton, 0.54 for the modified Hamilton, 0.40 for the HP and 0.41 for the BP filtered output gaps.⁹ Overall, the analysis yields first indications that the modified Hamilton filtered output gap is economically meaningful. More systematic evaluations are provided below.

4.1 Correlation with output gaps from policy institutions

First, we compare real-time output gap estimates to revised US output gap estimates of the Federal Reserve, the Congressional Budget Office, the IMF and the OECD. Output gap estimates of policy institutions should be useful benchmarks for three reasons. First, these output gaps should entail economic considerations regarding past courses of the US business cycle that stem from economic expertise rather than being based only on statistical models. Second, their assessment should capture output gap dynamics that are deemed important from a practitioner's perspective. Third, recent papers show that output gaps from policy institutions have been more reliable than those based on statistical methods over the last 20 years (see Edge and Rudd, 2016, for the Fed's output gap and Champagne et al., 2018, for the Bank of Canada's output gap).

There are some limitations regarding available samples and frequencies of output gap estimates from policy institutions. For the Fed and the CBO quarterly data is available, while for the IMF and the OECD only annual data is available. Further, for the Fed data is only available until 2012Q3. The data is based on the Greenbook which is made available to the public with a 5 year lag. Data for the Fed and the CBO is available from the start of our real-time output gap sample in 1965Q3, while data for the IMF and the OECD starts in 1985.

Table 2 shows correlations for the quarterly output gap series covering the period 1965Q3 until 2012Q3 on the left and for the annual series for data from 1985 to 2012 on the right.¹⁰ It gets apparent that the two Hamilton filtered real-time output gaps are highly correlated with all policy institutions' output gaps with correlation coefficients ranging from 0.55 to 0.86. The output gap based on the modified Hamilton filter has an even higher correlation with all four expert output gap estimates than the one based on the original filter. The HP and BP filtered real-time output gap correlations with the policy institutions' gaps are much lower ranging from 0.16 to 0.59, with those of the HP filter being particularly low. Both Hamilton filtered output gaps have a significantly higher correlation with the expert output gaps compared to the output gaps based on the HP or the BP filter. Interestingly, for all four statistical output gap measures the correlations are higher with the output gaps of the Fed and the CBO—the two U.S. institutions—than with the IMF and OECD—the two international institutions. Overall, these results imply that the real-time Hamilton filtered output gaps—and in particular the version based on the modified Hamilton filter—are able to reflect the ex post expert evaluation of the U.S. business cycle to a considerable degree in real time, while the BP and in particular the HP filter do this to a much smaller

⁹One would expect that output falls below potential during a recession implying a decreasing output gap, while one would expect it to rise above potential during a boom implying an increasing output gap. Based on this, we compute correlations with NBER recession dates based on the changes in the output gap rather than the output gap itself.

¹⁰The sample end is chosen to make the results for the Fed's output gap comparable to the other ones. Extending the sample for the other output gap series yields very similar results. The sample for the quarterly correlations includes 189 observations. For simplicity, we base the annualized series for the statistical real-time gap measures on yearly averages of the real-time series. When we account for within year revisions for quarters one to three of a given year to mimic the information set of a policy maker we obtain very similar results, both regarding numerical as well as p-values. Note that the annual correlations are based on a sample of 28 observations.

extent.

Table 2: Evaluation based on expert output gap estimates: Correlations

	Quarterly Data (1965Q3-2012Q3)		Annual Data (1985-2012)			
	FED	CBO	FED	CBO	IMF	OECD
Hamilton	0.59	0.68	0.78	0.82	0.69	0.55
Modified	0.64	0.73	0.83	0.86	0.73	0.60
HP	0.25***	0.36***	0.33***	0.43***	0.29**	0.16**
BP	0.49**	0.59**	0.51**	0.57**	0.49	0.37

Notes: *, **, and *** denote significance on the 10, 5, and 1 % significance level relative to the modified measure.

4.2 Forecasting Performance: Output Growth

The previous analysis shows that the assessment of the meaningfulness of output gap estimates is to some extent subjective depending on the researcher’s objective regarding what the output gap should measure. However, with the assessment of the predictive content of output gap estimates also a more objective criterion is available.

Nelson (2008) proposed the evaluation of competing output gap measures via their output growth forecasting performance. If an output gap is negative, one would expect above average output growth rates in the future, so that output reverts back to trend. Conversely, if the output gap is positive, output growth should be below average sometimes in the future. We use a standard forecast equation in which output growth h periods ahead is predicted using the real-time output gap:

$$y_{t+h} - y_t = \alpha + \beta \hat{c}_t + \varepsilon_{t+h|t}, \quad (9)$$

where y denotes log real GDP, \hat{c} the estimated real-time output gap and $\varepsilon_{t+h|t}$ the forecast error. The equation is estimated with OLS. The first subsample goes from 1965Q3 to 1975Q2 covering 10 years of data. The sample is then recursively extended quarter-by-quarter. Forecasts are computed for horizons 1 to 12. Based on the intuition developed above, we expect $\beta < 0$ at some horizon, essentially indicating the ability of the output gap to predict trend-reverting tendencies of the output growth series.

Table 3 shows RMSEs and the estimated slope coefficient β averaged over all subsamples. The RMSEs are reported for the Hamilton, modified Hamilton and BP filtered output gaps relative to the HP filtered one. A number smaller than 1 indicates a superior forecasting performance of the Hamilton, modified Hamilton or BP filtered output gap, while a number larger than 1 indicates that the HP output gap dominates. The RMSEs of output growth forecasts based on the four different real-time output gap measures are almost identical. The small differences are insignificant based on the Dieboldt-Mariano test. The signs of the slope coefficients show that all four output gap measures predict trend-reverting output growth rates. The slope coefficients are negative and significant from horizon 6 onwards for the two Hamilton and the BP filtered output gap and from horizon 11 onwards for the HP filtered output gap. The size of the slope coefficients cannot be directly compared, because of the different amplitude of the three output gap measures.

The results are robust to changes in the sample except for leaving out the period of the Great Recession and its aftermath. When ending the forecast evaluation in 2007Q4 the RMSEs of the two Hamilton filtered output gaps relative to the HP filtered one drop to values around 0.9 for horizons 8 to 12 and the difference is significant on the 10% level. The RMSEs relative to the BP filter drop to values around 0.97

Table 3: Output growth forecast evaluation

Horizon	Relative RMSE			Slope Coefficient (p-value)			
	Hamilton/HP	Modified/HP	BP/HP	Hamilton	Modified	HP	BP
1	1.02	1.02	1.02	0.03 (0.24)	0.03 (0.22)	0.13 (0.02)	0.08 (0.34)
2	1.01	1.01	1.01	0.02 (0.53)	0.03 (0.54)	0.18 (0.05)	0.04 (0.60)
3	1.01	1.01	1.01	-0.02 (0.67)	-0.02 (0.69)	0.17 (0.14)	-0.08 (0.70)
4	1.02	1.01	1.02	-0.06 (0.37)	-0.07 (0.34)	0.13 (0.27)	-0.26 (0.26)
5	1.01	1.00	1.02	-0.12 (0.13)	-0.15 (0.10)	0.05 (0.45)	-0.47 (0.07)
6	1.01	0.99	1.01	-0.18 (0.04)	-0.22 (0.03)	-0.05 (0.63)	-0.69 (0.02)
7	0.99	0.98	1.00	-0.24 (0.01)	-0.29 (0.01)	-0.15 (0.57)	-0.89 (0.00)
8	0.98	0.97	0.99	-0.30 (0.00)	-0.36 (0.00)	-0.27 (0.29)	-1.07 (0.00)
9	0.97	0.97	0.98	-0.35 (0.00)	-0.41 (0.00)	-0.36 (0.15)	-1.21 (0.00)
10	0.97	0.97	0.97	-0.39 (0.00)	-0.45 (0.00)	-0.44 (0.07)	-1.33 (0.00)
11	0.97	0.97	0.97	-0.42 (0.00)	-0.49 (0.00)	-0.51 (0.03)	-1.45 (0.00)
12	0.97	0.97	0.97	-0.44 (0.00)	-0.53 (0.00)	-0.60 (0.01)	-1.54 (0.00)

for horizons 8 to 12, but differences remain insignificant. One reason why the Hamilton filtered output gap's relative performance improves when leaving out the post 2007 period is that it adjusts to structural breaks in the level of GDP late as discussed above. The Hamilton filter needed about three years to adjust to the level shift of GDP after the Great Recession. This is why the output gap of the Hamilton filter is highly negative until the end of 2010. The other HP and BP filter based output gap measures are more flexible with respect to trend adjustments if there are very large business cycle fluctuations, so that the output gap goes back to zero more quickly after the crisis. The recovery after the crisis was slow, so that there are no above average output growth rates after 2010 as one would expect from the prolonged highly negative output gap of the Hamilton filter until then end of 2010.

Overall, the results from this exercise are inconclusive in judging which output gap measure has favorable properties. On the one hand the two Hamilton filtered output gaps perform somewhat better than the HP filtered one, but on the other hand a drawback of the Hamilton filter is its inflexibility to adjust to structural breaks early on. Differences to the BP filter with respect to the prediction of trend-reverting output growth rates are small.¹¹

4.3 Forecasting Performance: Inflation

Since theory predicts that output gap estimates should be useful predictors for forecasting inflation, we also consider a Phillips curve type forecasting inflation to evaluate the competing output gap measures. We follow Stock and Watson (1999b), Clark and McCracken (2006), Stock and Watson (2008) and Kamber et al. (2018) in specifying an autoregressive distributed lag (ADL) Phillips Curve forecasting equation:

$$\pi_{t+h} - \pi_t = \alpha + \sum_{i=0}^p \beta_i \Delta \pi_{t-i} + \sum_{i=0}^q \gamma_i \hat{c}_{t-i} + \varepsilon_{t+h|t}, \quad (10)$$

¹¹When conducting the same exercise with final revised output gaps, the RMSEs of the Hamilton filtered output gaps are significantly larger than the ones of the HP and BP filter. However, this is not a suitable criterion to conclude that for historical analyses revised HP and BP filtered output gaps are preferable. The low RMSEs merely reflect the two-sided nature of the HP and BP filter. When using future information to measure the output gap, then it is not surprising that the predictive content regarding these future realizations is higher compared to using a one-sided filter like the Hamilton filter.

where π_t denotes US PCE inflation, \hat{c}_t the estimated real-time output gap and $\varepsilon_{t+h|t}$ the forecast errors.¹² We use final revised data for inflation. The lag lengths p for inflation and q for the output gap are determined based on the entire sample using the SIC. We consider $p \in [0, 12]$ and $q \in [0, 12]$. When doing the recursive forecast evaluation we then assume that the optimal lag order is known a priori. As above the initial sample runs from 1965Q3 to 1975Q2, the sample is recursively extended quarter-by-quarter and forecasts are computed for horizons 1 to 12. For comparison we also compute results for a model that omits the output gap, but is otherwise identical to equation (10).

Table 4: Inflation forecast evaluation

Horizon	Relative RMSE					
	Hamilton/HP	Modified/HP	Hamilton/BP	Modified/BP	Hamilton/No Gap	Modified/No Gap
1	1.02	1.01	1.02	1.02	1.10***	1.10***
2	1.01	0.98	0.99	0.97	1.01	0.99
3	1.00	0.99	1.00	0.99	1.04	1.03
4	1.00	1.01	1.00	1.01	1.01	1.02
5	1.01	1.00	0.99	0.98	0.99	0.97
6	1.01	1.00	0.97	0.96	0.96	0.95
7	0.97	0.97	0.97	0.97	0.93	0.92
8	1.01	1.01	1.01	1.01	0.96	0.95
9	1.01	1.01	1.00	1.00	0.97	0.96
10	1.01	1.01	1.00	1.00	0.97	0.96
11	1.01	1.01	1.02**	1.02**	0.96	0.95
12	1.01	1.01	1.01	1.00	0.96	0.95

Notes: *, **, and *** denote significance on the 10, 5, and 1 % significance level based on a two-sided Diebold-Mariano test.

Table 4 shows root mean squared inflation forecast errors based on the two Hamilton filtered output gaps relative to those based on the HP and the BP filter and based on the specification without output gap. The differences between the forecasting models are marginal and insignificant. Further, models that condition on an output gap measure do not significantly improve upon a univariate inflation forecast. While these results unfortunately do not help in evaluating competing output gap measures, they are fully in line with the literature. Among others, Stock and Watson (2007), Stock and Watson (2008), Edge and Rudd (2016), and Kamber et al. (2018) find that it is generally difficult to beat univariate inflation forecast models through conditioning on output gaps. We find very similar results, when conditioning on final revised output gaps. Hence, these results are due to the general decline in the forecastability of inflation in recent decades (Stock and Watson, 2007, 2008), rather than to real-time output gap measurement problems.

5 Conclusions

We have proposed a modified version of the Hamilton filter (Hamilton, 2018b) for the estimation of reliable and economically meaningful real-time output gaps. It shares the favorable real-time properties of the Hamilton filter and is similarly easy to compute, but has a much better coverage of typical business cycle frequencies and yields a smooth estimated trend, while the original approach does not. While the original approach is very useful for detrending, the modified version allows a meaningful economic interpretation of the cyclical component as an output gap and of the trend as potential GDP. This is particularly important, because the existing papers applying the Hamilton filter use it not merely for detrending, but

¹²This specification is based on changes in inflation, i.e. it imposes a unit root in inflation. We also estimated models in terms of inflation levels and specifications that include a relative import price inflation term. The forecasting results are very similar.

attach an economic interpretation to the filtered time series. Compared to other simple statistical trend-cycle decomposition techniques such as the HP or the BP filter, the real-time output gap based on the modified Hamilton filter shows a much higher correlation with ex post assessments of output gaps from important policy institutions. Hence, the method yields a real-time output gap that captures business cycle dynamics that are deemed important from a practitioner's perspective.

References

- Ahn, H. J. and J. D. Hamilton (2016). Heterogeneity and unemployment dynamics. *National Bureau of Economic Research* (No. w22451).
- Atkeson, A. and L. E. Ohanian (2001). Are Phillips curves useful for forecasting inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25(1), 2–11.
- Baxter, M. (1991). Business cycles, stylized facts, and the exchange rate regime: evidence from the United States. *Journal of International Money and Finance* 10(1), 71–88.
- Baxter, M. and A. C. Stockman (1989). Business cycles and exchange rate regime: some international evidence. *Journal of Monetary Economics* 23(3), 377–400.
- Bernhardsen, T., Ø. Eitrheim, A. S. Jore, and Ø. Røisland (2005). Real-time data for Norway: Challenges for monetary policy. *North American Journal of Economics and Finance* 16(3), 333–349.
- Bordo, M. D. and P. L. Siklos (2017). Central banks: Evolution and innovation in historical perspective. *National Bureau of Economic Research* (No. w23847).
- Burns, A. F. and W. C. Mitchell (1946). *Measuring business cycles*. Nber Books.
- Canova, F. (1998). Detrending and business cycle facts. *Journal of Monetary Economics* 41(3), 475–512.
- Canova, F. (2011). *Methods for Applied Macroeconomic Research*. Princeton University Press.
- Cayen, J.-P. and S. van Norden (2005). The reliability of Canadian output-gap estimates. *North American Journal of Economics and Finance* 16(3), 373–393.
- Champagne, J., G. Poulin-Bellisle, and R. Sekkel (2018). The real-time properties of the Bank of Canada’s staff output gap estimates. *Journal of Money, Credit and Banking* 50(6), 1167–1188.
- Christiano, L. J. and T. J. Fitzgerald (2003). The band pass filter. *International Economic Review* 44(2), 435–465.
- Clark, T. E. and M. W. McCracken (2006). The predictive content of the output gap for inflation: resolving in-sample and out-of-sample evidence. *Journal of Money, Credit and Banking* 38(5), 1127–1148.
- Cogley, T. and J. 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.
- Danielsson, J., M. Valenzuela, and I. Zer (2018). Learning from history: volatility and financial crises. *The Review of Financial Studies* 31(7), 2774–2805.
- Edge, R. M. and J. B. Rudd (2016). Real-time properties of the Federal Reserve’s output gap. *Review of Economics and Statistics* 98(4), 785–791.
- Fisher, J. D., C. T. Liu, and R. Zhou (2002). When can we forecast inflation? *Federal Reserve Bank of Chicago Economic Perspectives* 1Q/2002, 30–42.

- Gruen, D., T. Robinson, and A. Stone (2005). Output gaps in real time: How reliable are they? *Economic Record* 81(252), 6–18.
- Hamilton, J. D. (2018a). Measuring global economic activity. *University of California at San Diego*.
- Hamilton, J. D. (2018b). Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics* 100(5), 831–843.
- Harvey, A. C. and A. Jaeger (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8(3), 231–247.
- Kamber, G., J. Morley, and B. Wong (2018). Intuitive and reliable estimates of the output gap from a Beveridge-Nelson filter. *Review of Economics and Statistics* 100(3), 550–566.
- Kilian, L. and H. Lütkepohl (2017). *Structural vector autoregressive analysis*. Cambridge University Press.
- King, R. G. and S. T. Rebelo (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17(1-2), 207–231.
- López-Salido, D., J. C. Stein, and E. Zakrajšek. (2017). Credit-market sentiment and the business cycle. *The Quarterly Journal of Economics* 132(3), 1373–1426.
- Marcellino, M. and A. Musso (2011). The reliability of real-time estimates of the euro area output gap. *Economic Modelling* 28(4), 1842–1856.
- Nelson, C. R. (2008). The beveridge-nelson decomposition in retrospect and prospect. *Journal of Econometrics* 146(2), 202–206.
- Orphanides, A. and S. van Norden (2002). The unreliability of output-gap estimates in real time. *Review of Economics and Statistics* 84(4), 569–583.
- Orphanides, A. and S. van Norden (2005). The reliability of inflation forecast based on output gap estimates in real time. *Journal of Money, Credit and Banking* 37, 583–600.
- Richter, B., M. Schularick, and I. Shim (2018). The costs of macroprudential policy. *National Bureau of Economic Research* (No. w24989).
- Richter, B., M. Schularick, and P. Wachtel (2017). When to lean against the wind. *Centre for Economic Policy Research* (DP12188).
- Schüler, Y. S. (2018). On the cyclical properties of Hamiltons regression filter. *Deutsche Bundesbank Discussion Paper* (No. 03/2018), 1–18.
- Stock, J. H. and M. W. Watson (1999a). Business cycle fluctuations in us macroeconomic time series. *Handbook of macroeconomics* 1, 3–64.
- Stock, J. H. and M. W. Watson (1999b). Forecasting inflation. *Journal of Monetary Economics* 44(2), 293–335.
- Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* 39(s1), 3–33.

- Stock, J. H. and M. W. Watson (2008). Phillips curve inflation forecasts. In J. Fuhrer, Y. K. Kodrzycki, J. S. Little, and G. P. Olivei (Eds.), *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, pp. 101–186. Cambridge, MA: MIT Press.
- Van Zandweghe, W. (2017). The changing cyclicalness of labor force participation. *Economic Review* (Q II), 1–30.
- Wolters, M. H. (2018). How the baby boomers' retirement wave distorts model-based output gap estimates. *Journal of Applied Econometrics* 33(5), 680–689.