



BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers

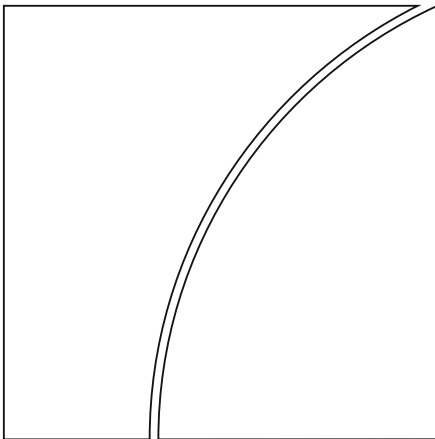
No 442

A parsimonious approach to incorporating economic information in measures of potential output

by Claudio Borio, Piti Disyatat and Mikael Juselius

Monetary and Economic Department

February 2014



JEL classification: E10, E40, E44, E47, E52, E60.

Keywords: Potential output, output gap, Phillips curve, financial cycle.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2014. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)

ISSN 1682-7678 (online)

A parsimonious approach to incorporating economic information in measures of potential output*

Claudio Borio, Piti Disyatat and Mikael Juselius[†]

Abstract

A popular strategy for estimating output gaps is to anchor them to structural economic relationships. The resulting output gaps, however, are often highly sensitive to numerous auxiliary assumptions inherent in the approach. This complicates their use in policymaking. We illustrate the point using the Phillips curve, arguably the most popular structural relationship in this context. Depending on the specification, we show that conditioning on this relationship either introduces a trend in the output gap – which is conceptually unappealing – or has little effect on it – which defeats the purpose of the exercise. Moreover, the estimated gaps perform poorly in real time, with large ex-post revisions. The opaqueness of the approach, which increases greatly with the dimension of the estimated system, can mask these problems. In order to address these limitations, we propose a more parsimonious and transparent approach to embedding economic information that is less vulnerable to misspecification. As an illustration, we apply the corresponding parsimonious multivariate filter to US data. We find that proxies for the financial cycle, notably credit growth, but also unemployment contain significant information and help generate robust real-time output gap estimates.

JEL classification: E10, E40, E44, E47, E52, E60.

Keywords: Potential output, output gap, Phillips curve, financial cycle.

* This paper is a technical companion to, and generalises, Borio et al (2013). We would like to thank Katrin Assenmacher, Éric Dubois, Stan Fischer, Adrian Pagan, Andrzej Sławiński, Frank Smets, Birger Vikøren and Marco Lombardi for comments on an earlier draft. Magdalena Erdem and Anamaria Illes provided excellent statistical assistance. All remaining errors are ours. The views expressed are those of the authors and do not necessarily represent those of the Bank for International Settlements or the Bank of Thailand.

[†] Borio: Monetary and Economic Department, Bank for International Settlements, claudio.borio@bis.org; Disyatat: Research Department, Bank of Thailand, pitid@bot.or.th; Juselius: Monetary and Economic Department, Bank for International Settlements, mikael.juselius@bis.org

Table of Contents

Introduction.....	1
I. Univariate Kalman-filter estimates of potential output.....	3
The HP filter benchmark output gap: static version.....	4
The HP filter benchmark output gap: dynamic version	5
II. Multivariate estimates: Adding structural economic relationships.....	7
Output gap estimates based on the Phillips curve.....	8
III. An alternative approach to embedding economic information	13
The parsimonious multivariate filter: simplicity, transparency and robustness...14	
Potentially informative variables	15
Results.....	17
IV. Real-time performance	20
Conclusion.....	22
References.....	24
Appendix A: Data definitions and sources	26
Appendix B: The Kalman filter	27
Graphs.....	29

Introduction

Potential output and its corresponding deviations from actual output (“output gaps”) are not directly observable: they must be estimated from the data. A common strategy is to go beyond purely statistical approaches that rely exclusively on the path of output itself and to incorporate additional economic information.¹ Incorporating such information is very appealing because it holds out the promise of improving the estimates in at least two respects. First, it can yield *more economically meaningful* measures. Relying on theoretically plausible relationships between the output gap and other economic variables can tie the estimates of potential output more tightly to the concept, ie the level of sustainable output that maximises the use of available resources. Second, in principle it can *improve the statistical properties* of the estimates. In particular, it can alleviate the well-known end-point problem associated with the purely statistical approaches and yield estimates that perform better in real time, ie that are less subject to revisions as the future unfolds. When revisions occur, history, as it were, gets continuously rewritten. This is especially troublesome in policymaking.

The conventional, most popular approach to embedding economic information is adding structural economic relationships to a system of equations. The system may be a blend of a purely statistical representation of the data and reduced-form economic relationships, as in the case of some multivariate filter methods, or an explicit theoretical model, as in the case of DSGE-based estimates. The most prominent example is the Phillips curve, in which the output gap is an explanatory variable for inflation.

In this paper, we do two things. First, we argue that the conventional approach is in fact opaque and highly vulnerable to specification errors that have not been sufficiently appreciated. For example, even if the specific assumptions for the individual economic relationships in the system are clearly spelled out, their interaction can nevertheless produce unpredictable outcomes. This is particularly disappointing since, as is well known, the approach has generally failed to improve real-time performance significantly. Second, we propose an alternative approach that is less vulnerable to misspecification, can yield much more robust real-time estimates and is more transparent. By “more transparent” we mean that one can clearly link outcomes to specification assumptions and more easily assess the contribution of conditioning variables.

We illustrate the vulnerability of the conventional approach with the performance of the Phillips curve on US data. As is well known, the US inflation rate has trended down since its peak in the late 1970s. This trend is difficult to reconcile with standard formulations of the Phillips curve, which posit a close link between inflation and the output gap – a variable not expected to trend over long horizons. Conditioning output gap estimates on this relationship gives rise to a catch-22-like dilemma. If the estimation assigns a sufficiently high weight to the relationship, the estimated output gap will mimic the inflation rate and inherit its low-frequency trend. As a result, the corresponding output gaps do not accord well with general economic intuition. If the weight on the relationship is freely estimated, the output gap will look “sensible”, largely coinciding with that obtained from purely statistical

¹ See eg Canova (1998) and Cogley and Nason (1995) for early critiques of univariate methods.

filters, but the Phillips curve will have little influence on the estimates. Moreover, the dilemma cannot be avoided by filtering out the low-frequency component from the inflation rate: this reintroduces the end-point problem through the back door and undermines real-time performance.²

Probably even more troublesome is the opaqueness of the conventional approach. For one, it may yield Phillips curve estimates with plausible and significant coefficients even if the relationship is effectively redundant in the estimation, thereby providing a false sense of comfort. The reason is that the relevance of the estimated output gap for inflation is determined through various auxiliary assumptions inherent in the estimation, such as the prior confidence placed in the Phillips curve, priors on the parameters, and the assumed dynamics. Hence, for example, one should generally resist the temptation to interpret the t-value on the filtered output gap term in the Phillips curve as a sign of modelling success. More generally, the complexity and opaqueness of the approach become rapidly overwhelming as the number of conditioning economic relationships (equations) increases. This makes it harder to spot possible specification problems and to assess what is driving the results.

We argue that the success of any model that adds economic information to output gap estimates should be evaluated on three criteria: (i) the economic plausibility of the estimated gap; (ii) the contribution of the economic variables or relationships to the estimated gaps; and (iii) the size of the revisions between real-time and ex-post estimates.

Our alternative approach seeks to score highly on these three criteria. Rather than adding structural equations that embody priors on economic relationships, it evaluates directly the ability of plausible *observable* economic variables to explain cyclical output fluctuations *at a specific frequency*.³ Our approach can be seen as a restricted version of the more general system-based approach used in the literature. But rather than adding additional equations as the number of conditioning variables expands, we keep the dimension of the system small by including the variables directly in an equation where the output gap is the dependent variable. As we discuss below, keeping the system small turns out to be critical for good real-time performance.

This parsimonious multivariate filter approach has several advantages. First and foremost, compared to the conventional approach, it calls for fewer a priori judgements about the relevance of the additional economic relationships or variables. One simply needs to choose the desired frequency at which to explain cyclical fluctuations and then see whether the candidate (conditioning) economic variables actually do so. Thus, the method can easily accommodate many conditioning variables while keeping the dimensionality of the system low. Second, the size and statistical significance of the coefficients on the conditioning variables

² While, for simplicity, we focus on the Phillips curve and limit ourselves to US data, the problem is more general. For instance, as inflation has similar statistical properties across countries, the widely documented weak explanatory power of output gaps for inflation, or alternatively the poor empirical performance of the Phillips curve (eg Nason and Smith (2008), ECB (2011), IMF (2013)), may be at least in part a reflection of the same problem. Likewise, our critique would also apply, for instance, to the inversion of Okun's Law, à la Benes et al (2010).

³ Here we simply adopt the HP-filter frequency, the most widely used in the analysis of the business cycle. That said, the "right" frequency depends on the specific question. We leave this for future work.

offers a cleaner, albeit still imperfect, benchmark by which to judge their relative importance. Third, the simplicity and transparency of the approach make it easy to interpret the results and judge the economic plausibility of the output gap estimates. Finally, the approach can improve real-time performance substantially. As we show, statistically good real-time performance requires that the output gap estimates be anchored to observable variables with stable means (see also Borio et al (2013)).

We illustrate the usefulness of the approach by applying it to US data and considering a range of plausible explanatory variables. We find that proxies for the financial cycle, notably credit growth, possibly combined with property prices, as well as the *raw* unemployment rate, are particularly promising. They help generate economically plausible output gaps with good real-time performance. The results for financial cycle proxies confirm those in Borio et al (2013), which considers a broader set of countries.⁴ For unemployment, they highlight the benefits of relying on the raw data as opposed to *deviations* from a natural rate. We leave for further research the question of whether the unemployment rate is equally valuable for other countries; that said, preliminary evidence suggests that it is unlikely to be.

The rest of the paper is organised as follows. Section I describes our starting point and benchmark for adding economic information, namely the popular Hodrick-Prescott (HP) filter couched in a state-space (Kalman filter) framework. This relies exclusively on information about output. Section II illustrates the key pitfalls that may arise when incorporating additional economic information in the traditional approach, based on the performance of variants of the Phillips curve. Section III lays out our alternative approach. Section IV compares the real-time performance of the different approaches.

I. Univariate Kalman-filter estimates of potential output

The basic state-space Kalman filter framework underlies most attempts to incorporate economic information in potential output estimates. We start from the univariate case, which relies only on observations of real GDP and accommodates the popular HP filter. This is the most familiar benchmark for output gaps at conventional business cycle frequencies.⁵ We illustrate how different specifications of the state-space system of equations and of the relative weights attached to each equation impact on potential output estimates. Understanding the univariate case is critical before assessing the value added of embedding additional economic information in the filter.

⁴ Our preliminary estimates for a much broader set of countries tend to confirm the general nature of the information content of these variables.

⁵ Appendix B provides a short introduction to the Kalman filter. The filter is a general way of estimating potential output from a system of linear equations. See, for instance, Hamilton (1994) for an in-depth treatment and Kuttner (1994) for an early application.

The HP filter benchmark output gap: static version

Given T observations on actual output (in logs), y_t , the HP filter chooses potential output, y_t^* , to minimise the loss function

$$\sum_{t=1}^T \left(\frac{1}{\sigma_1^2} (y_t - y_t^*)^2 + \frac{1}{\sigma_0^2} (\Delta y_{t+1}^* - \Delta y_t^*)^2 \right) \quad (1)$$

where σ_1^2 is the variance of the output gap, σ_0^2 is the variance of the change in potential output, and Δ is the difference operator. A closed-form solution to the minimisation problem exists for a fixed value of the scaling parameter $\lambda_1 = \sigma_1^2 / \sigma_0^2$ (eg Cogley and Nason (1995)). As formula (1) indicates, the scaling parameter determines the relative weight attached to deviations of potential output from actual output and to the smoothness of the potential output series itself. The standard value for λ_1 in quarterly data is 1600, which limits the maximum length of the business cycle to approximately 8 years.

Alternatively, one can minimise the loss function in (1) using the Kalman filter. Specifically, note that one can interpret the second term in (1) as the squared residuals from the following *state equation*

$$\Delta y_{t+1}^* = \Delta y_t^* + \varepsilon_{0,t+1} \quad (2)$$

where $\varepsilon_{0,t}$ is assumed to be normally distributed with mean zero with variance σ_0^2 . Similarly, one can interpret the first term in (1) as the squared residuals from the *observation equation*

$$y_t = y_t^* + \varepsilon_{1,t} \quad (3)$$

where $\varepsilon_{1,t}$ has mean zero and variance σ_1^2 . The errors $\varepsilon_{0,t}$ and $\varepsilon_{1,t}$ are assumed to be uncorrelated at all lags and leads. Since the Kalman filter is an algorithm for calculating the linear *least squares* forecasts for the variables of the system, it jointly minimises the squared residuals in (2) and (3). As a result, the solution for y_t^* will coincide with that of the HP-filter provided one restricts the ratio between the residual variances to be λ_1 . The Kalman filter framework also highlights the implicit assumption in the standard HP filter that output gaps are not serially correlated. We refer to this standard form of the filter as the "*static*" version, in order to distinguish it from its "*dynamic*" counterpart, which allows for serial correlation and which we introduce below.

The state and observation equations play different roles. Observation equation (3) anchors potential to actual output by positing that their difference is white noise. State equation (2) expresses potential output growth as a random walk, allowing it to evolve freely without bounds as a non-stationary process. In doing so, the equation imparts persistence, or inertia, to potential output growth, implying smoothness in the potential output path and permitting protracted one-sided deviations from actual output. Thus, it is the relative weight, λ_1 , between these two equations that determines how strictly potential output is anchored to actual output and, therefore, also *which output frequencies belong to the business cycle*.

There are two extreme cases. If one attaches a very high relative penalty to potential output growth in (1), ($\lambda_1 \rightarrow \infty$), the minimising solution amounts to setting $\varepsilon_{0,t}$ to zero and fitting a straight line (linear trend) to actual output. Potential output is then only weakly anchored to actual output since the error term, $\varepsilon_{1,t}$, in (3) can be very large. At the other extreme, when one attaches a very high relative penalty to

deviations of potential output from actual output in (1), ($\lambda_1 = 0$), the solution actually equates the two at all points in time.⁶

It is possible to estimate the scaling factor, λ_1 , in the Kalman filter rather than fixing it at the outset, but this will generally involve a substantial upward bias. In other words, it will tend to smooth potential output too much and reduce correspondingly the degree of mean reversion in the output gap: the business cycle will appear more persistent or longer than it really is. This is known as the “pile-up problem”.⁷ Put differently, it is difficult to infer the “right” value of the noise-to-signal ratio from the data. Below, we will be very explicit about how the scaling factor is set and report its value for each estimated model.⁸

The HP filter benchmark output gap: dynamic version

The first extension we consider is to include dynamic terms in (3) to account for the pronounced serial correlation that is clearly visible in the HP-filtered output gap (the red line in Graph 1, left panel). Ostensibly, the implicit assumption that the HP-filtered output gaps are white noise is a poor representation of the data. The simplest possible specification of the dynamic HP filter is to assume an AR(1) process for the output gap

$$y_t = y_t^* + \beta(y_{t-1} - y_{t-1}^*) + \varepsilon_{2,t} \quad (4)$$

where $\varepsilon_{2,t}$ is assumed to be normally distributed with mean zero and variance σ_2^2 . In this case the Kalman filter seeks solutions to β and y_t^* that minimise the loss function

$$\sum_{t=1}^T \left(\frac{1}{\sigma_2^2} (\varepsilon_{2,t})^2 + \frac{1}{\sigma_0^2} (\Delta y_{t+1}^* - \Delta y_t^*)^2 \right) \quad (5)$$

In this alternative specification the variability of the output gap no longer depends solely on the corresponding scaling factor, $\lambda_2 = \sigma_2^2 / \sigma_0^2$; it is also a function of the auto-regressive parameter, β . For example, if $\beta = 1$, then the output gap becomes a random walk and its cumulative variance can go to infinity even if λ_2 is a fixed number. However, provided that β is sufficiently smaller than 1, so that the

⁶ One can see equation (3) as establishing a one-to-one long-run correspondence between potential and actual output. This ensures that the estimated output gap is mean reverting. One may view this as a necessary, but by no means sufficient, property for an output gap estimate to make economic sense. It would be hard to reconcile output gaps that trended up or down persistently in the long-run with existing theoretical notions of potential output. As shown below, while weakening the restrictions that enforce this long-run one-to-one correspondence has only a minor impact on the estimates in the univariate case, it can have a much larger one in the multivariate case.

⁷ The technical reason is that in the presence of a non-stationary state variable, maximum likelihood estimates of the noise-to-signal ratio, defined by $1/\lambda_1$, has point mass at zero when its true value is small (Shephard (1993) and Stock and Watson (1998)). The pile-up problem reflects a discontinuity in the distribution of the state variable. For example, let u_t be a non-stationary state variable assumed to follow a random walk without drift. If the variance of its innovation is equal to zero, u_t becomes stationary (in fact a constant) and this accounts for the discontinuity in the distribution function.

⁸ Note that, by itself, the exact functional form assumed for the evolution of potential output, equation (2), does not have a large impact on the results, provided that it allows for stochastic trending and the scaling factor is set to cut out comparable frequencies for the business cycle. For example, the results based on a random walk with drift are similar to those reported below.

estimated output gap is mean-reverting, it is possible to generate a measure that has comparable cyclicity to the HP-filtered gap. Note that (4) implies $\text{var}(y_t - y_t^*) = (1 - \beta^2)^{-1} \sigma_2^2$ under the assumptions above. Hence, to ensure that $\text{var}(y_t - y_t^*)/\sigma_0^2 = 1600$ we should set $\lambda_2 = \sigma_2^2/\sigma_0^2 = 1600(1 - \beta^2)$. If β is estimated, one can do this by iterating until the condition holds.

That said, there is a small sample problem. When the true representation of output gap is (4), the residual that results from using (3) in the filter will be auto-correlated. This implies that the *empirical* HP-filter scaling factor will generally be substantially larger than 1600 in small samples, converging only slowly towards this number as the sample size increases. Hence, setting $\lambda_2 = 1600(1 - \beta^2)$ in the dynamic HP filter will put too little weight on reducing potential output variability compared to the HP-filter.

To overcome this problem, we iterate over different values of λ_2 until

$$\text{var}(y_t - y_{HP,t}^*)/\text{var}(\Delta y_{HP,t}^* - \Delta y_{HP,t-1}^*) = \text{var}(y_t - y_{ALT,t}^*)/\text{var}(\Delta y_{ALT,t}^* - \Delta y_{ALT,t-1}^*) \quad (6)$$

holds, where $y_{HP,t}^*$ and $y_{ALT,t}^*$ are the potential output estimates from the standard HP-filter and any model that uses equation (2) together with (4), respectively. This basically restricts the empirical scaling factor of the HP filter and its alternative specifications to be the same, ensuring comparability.⁹

While adding dynamic terms reveals a high degree of persistence in the output gap, *in the univariate case* the effect on the point estimates is surprisingly small. For example, estimating (2) and (4) on US quarterly real GDP data from 1980q1-2011q4 yields a value for β of 0.96, which is statistically difficult to distinguish from unity. However, as can be seen from the left panel of Graph 1, the estimated output gap is *virtually identical* to the one obtained from the standard static HP filter.

Why is this so? The main reason is that *without additional information that can drive a wedge between them*, the estimation procedure tends to anchor potential output to actual output. To see this, note that the error term $\varepsilon_{2,t}$ in equation (4) is minimised when potential output equals actual output. Minimising (5) then implies that potential output growth should accelerate (ie. shocks to the state equation (2) should always be positive) when the output gap is positive: this reduces the error in (4), unless β is exactly equal to one. The estimation procedure ensures this by raising potential output when it is below actual output and vice versa. Thus, there is a mechanical tendency for potential to be anchored to actual output and for output gaps to self-correct even when β is close to unity. This, however, *no longer holds once one adds economic information* to the filter, as we discuss in more detail below. Hence, in the multivariate case, one should carefully restrict dynamic terms to avoid trending (non-stationary) output gaps.

While dynamic terms do not change the point estimates of potential output to any significant degree, they do change drastically our perception of their precision: the standard errors in the static version of the filter are biased downwards. The right-hand panel of Graph 1 illustrates this, by comparing confidence bands from the static and dynamic HP-filtered gaps. Clearly, the precision of the HP gap is much

⁹ An alternative, more technical way, would be to express the estimate of potential output that would result from the Kalman-filter using a doubly-infinite sample as a linear, time-invariant, two-sided filter and to apply similar techniques as in eg Baxter and King (1999) to extract the relevant frequency band.

lower once the high degree of serial correlation is taken into account. Here we have only considered an AR(1) extension of (3), but we could equally well have tried a more general ARMA specification. As pointed out by Harvey and Jaeger (1993), by fixing the noise-to-signal ratio, the HP filter estimate of $\varepsilon_{1,t}$ is no longer a white noise process. However, the standard errors obtained under the Kalman filter still assume that it is. This leads to estimates of the errors that are deceptively precise.

II. Multivariate estimates: Adding structural economic relationships

Consider next adding economic information, which is generally motivated by the poor real-time performance of the univariate estimates and by concerns that they may fail to reflect “true” business cycle movements. Arguably, the most common approach in the literature is to add structural economic relationships as observation equations to the Kalman filter (eg Kuttner (1994) and Benes et al (2010)).

The features of these structural relationships can have a major impact on the resulting output gap estimate. Typically, the output gap is an explanatory variable in (at least one) of these relationships, which is assumed to explain the behaviour of some observable variable. The most prominent example is the Phillips curve, which aims to exploit an assumed link between economic slack, as proxied by the output gap, and inflation. We next explore how this general approach affects output gap estimates, using the Phillips curve as illustration. We highlight, in particular, its sensitivity to the various specification assumptions involved.

Additional observation equations

Suppose that in addition to (2) and (3), we add a structural economic relationship of the form

$$z_t = \gamma_0 + \gamma_2' x_t + \gamma_3(L)(y_t - y_t^*) + \varepsilon_{3,t} \quad (7)$$

where z_t is the economic variable to be explained (“conditioning variable”), x_t is a vector of economic variables, possibly including lags or leads of z_t itself, $\gamma_3(L)$ is a lag polynomial, and $\varepsilon_{3,t}$ is normally distributed with mean zero and variance σ_3^2 . The corresponding loss function is

$$\sum_{t=1}^T \left(\frac{1}{\sigma_1^2} (y_t - y_t^*)^2 + \frac{1}{\sigma_0^2} (\Delta y_{t+1}^* - \Delta y_t^*)^2 + \frac{1}{\sigma_3^2} (\varepsilon_{3,t})^2 \right) \quad (8)$$

This is identical to the loss function associated with the multivariate HP filter introduced by Laxton and Tetlow (1992). Compared to (1) and (5), not just one but two scaling factors $\lambda_1 = \sigma_1^2/\sigma_0^2$ and $\lambda_3 = \sigma_1^2/\sigma_3^2$ determine the relative weights of the three terms in (8).

Importantly, both scaling factors, λ_1 and λ_3 , in (8) *jointly* influence key aspects of the outcome, such as the cyclicity of the estimated gap and the degree to which the conditioning economic relationships contribute to it. For example, the first scaling factor, λ_1 , seems identical to the one from the standard HP filter. Hence, it is tempting to fix its value at 1600 (indeed, we do this in several examples below) in the hope of generating comparable cyclicity as that which results from the standard HP-filter. This, however, will not necessarily be the case: the third term in

(8) also matters. The relative weight accorded to this term is controlled by λ_3 , which reflects the importance of relationship (7), as either determined by the data, when the parameter is estimated, or as decided a priori by the researcher, when the parameter is imposed. If this scaling factor is sufficiently large compared to λ_1 , minimising (8) becomes almost equivalent to minimising the residual in (7). This can result in an output gap that mimics z_t regardless of its cyclical nature. For example, if z_t has a trend, the output gap will inherit it.

Complicating matters further, all the terms and parameters on the right-hand side of (7) change the size of the residual $\varepsilon_{3,t}$ and therefore also matter in (8). Different dynamic terms, in particular, can have a large impact. This makes it difficult to provide clear guidelines on how to appropriately set the scaling factors, for example, in order to ensure comparable cyclical nature. In principle, one could try to estimate both scaling factors but this would again involve a similar pile-up problem to that discussed above. If one is willing to fix the value of λ_1 , the scaling factor that ties down the variability of potential output, the pile-up problem should be less of a concern when estimating the other scaling factor.

These difficulties highlight the main pitfall of the standard approach: the estimated output gap is heavily dependent on a complex interaction between the scaling factors and specification choices for the key equations. Even small and seemingly innocuous changes can have large effects on the estimated gaps. Nor is there a clear basis to guide those choices. And the problem quickly becomes unwieldy, if not overwhelming, as the number of conditioning relationships increases.

Output gap estimates based on the Phillips curve

We illustrate this point with the help of the Phillips curve applied to US data. A cornerstone of prevailing output gap estimates is the reliance on inflation as the main indicator of output deviations from potential. This has deep economic roots. From at least Okun (1962) on, it is the behaviour of inflation that is assumed to signal whether output is above or below potential.

Before proceeding, it is worth recalling some key properties of the inflation process. As shown in the top right-hand panel of Graph 3, inflation in the United States, as in most other countries, has trended downwards from quite a high level in 1980 to a moderate and steady one from mid-1990s onwards. Moreover, statistically, inflation in the United States is highly persistent (Fuhrer and Moore (1995), Stock and Watson (2007)). Bai and Ng (2004) and Henry and Shields (2004), among many others, find evidence of a unit root in post-war inflation. Not surprisingly, trend inflation is often modelled as a drift-less random walk (eg Cogley and Sargent (2005), Ireland 2007; Stock and Watson 2007). As will become apparent below, the high persistence and downward trend in inflation makes it treacherous as a conditioning variable for output gap estimates.

We now specify (7) as a Phillips curve and investigate how this relationship affects our output gap estimates. The specification we consider is

$$\pi_t = \gamma_0 + \gamma_1 \pi_{t-1} + \gamma_2 (y_t - y_t^*) + \varepsilon_{3,t} \quad (9)$$

where π_t is the CPI inflation rate.¹⁰ The data are quarterly and the estimation period is 1980q1-2012q4.

To estimate the parameters in (9) we adopt a conventional Bayesian approach. We use the Kalman filter to form the likelihood of the system, specify prior distributions for the parameters, and maximise the posterior density function with respect to the parameters.¹¹ As prior distribution, we assume the gamma distribution with standard deviation of 0.3 for all the parameters. The inflation persistence parameter, γ_1 , is restricted to lie between 0 and 1 with a prior mean of 0.70. The constant, γ_0 , and the parameter on the output gap, γ_2 , are unrestricted in \mathbb{R}^+ with prior means equal to $3.5 \cdot 0.3$ (approximately the average US inflation rate times one minus the prior on γ_1) and 0.30, respectively. For simplicity, we set the scaling factor, λ_1 , associated with (3) to 1600 in line with the standard static HP filter.¹² We also try a dynamic specification, in which we replace (3) with (4). In this case, we restrict β to lie between 0 and 1, with a prior mean of 0.70 and set the scaling factor, λ_2 , such that (6) holds.¹³ For the scaling factor λ_3 , we try both freely estimated and imposed values.

As noted at the outset, we need to use an appropriate method to evaluate the marginal information content of the conditioning relationship. This is because a seemingly economically and statistically good performance of the Phillips curve can be misleading. We rely on the decomposition analysis based on the results in Koopman and Harvey (2003).

We next illustrate the sensitivity of the results to the various specification assumptions, in particular with respect to how the Phillips curve relationship is incorporated in the analysis. We highlight the "catch-22" dilemma. If one adopts specifications that ultimately assigns a sufficiently high weight to the Phillips curve, the trend in inflation will infect the output gap estimate and the resulting output gap measure will not look economically plausible. Alternatively, if the specifications adopted deliver intuitively sensible output gaps, the Phillips curve will be largely irrelevant as a conditioning relationship.

The first specification, Model 1, combines state equation (2), common to all models, with the static equation for the output gap, (3), and the Phillips curve relationship, (9). The scaling factor, λ_3 , is estimated freely. As the table indicates, all parameter values look reasonable: the process for inflation is clearly stationary, although somewhat persistent, and the term on the output gap is significant. On this basis, the Phillips curve appears to help condition the potential output estimate.

¹⁰ We also tried versions with additional lags and leads of inflation in the equation. This affects the standard errors but not the main conclusions of our analysis.

¹¹ We use the IRIS toolbox add-on to Matlab to perform these calculations. Instead of using Bayesian estimation, we could alternatively obtain maximum likelihood estimates of the parameter directly from the Kalman filter likelihood function.

¹² Again the empirical scaling factor will be larger due to the serial correlation in the estimated output gap.

¹³ We do not consider the added complexity of estimating λ_2 . However, if this parameter is sufficiently high, potential output becomes a line and the other equations will have relatively little impact on the results. This can effectively mask specification errors in (9).

Phillips curve estimates

t-statistics in parenthesis

Table 1

Model†	γ_0	γ_1	γ_2	γ_3	λ_1 (or λ_2)	λ_3	β
Model 1	1.27 (8.22)	0.61 (13.38)	0.31 (2.42)	—	1600*	0.43	—
Model 2	1.36 (3.86)	0.06 (0.94)	0.70 (6.26)	—	1600*	100*	—
Model 3	0.97 (7.24)	0.25 (3.70)	0.64 (3.43)	—	49**	0.04	0.99 (155.81)
Model 4	0.47 (5.37)	0.81 (11.85)	0.07 (0.60)	—	1600*	0.36	—
Model 5	0.95 (0.24)	0.19 (0.53)	0.53 (3.55)	—	1600*	0.83	—
Model 6	0.66 (7.10)	0.12 (3.00)	0.55 (3.68)	0.66 (13.25)	1600*	0.55	—

† All models are estimated on quarterly US data over the sample 1980q1–2012q4. Models 1 and 2: system contains equations (2), (3), and (9); Model 3: system contains equations (2), (4), and (9); Model 4 contains equations (2), (3), and (9) with $\Delta\pi$ in place of π ; Model 5: system contains equations (2), (3), and (9) in “gap” form; Model 6: system contains equations (2), (3), and (10). * λ_1 value imposed. ** λ_2 Set iteratively so that (6) holds.

But appearances are highly misleading: it turns out that the Phillips curve in Model 1 is virtually irrelevant for the estimate of potential output. As shown in the top left panel of Graph 2, the output gap estimate is hardly distinguishable from the HP-filtered one and inflation only accounts for a tiny fraction of its variation. In other words, the value of $\lambda_3 = 0.43$ as determined freely by the data, in effect implies that the Phillips curve has little weight. Conversely, the bulk of the variation in inflation is due to the residual term in the regression, with the output gap accounting for at most 2 percentage points of this variation at any given time. This can be seen from the upper right panel of Graph 2, which shows the contribution of the output gap to the variation in inflation.

Forcing the weight of the Phillips curve to be higher, rather than estimating it freely, produces the opposite problem. Now the relationship matters, but the output gap does not look economically plausible. Specifically, in Model 2 we set λ_3 equal to 100, compared with its freely estimated value of 0.43. This naturally forces the estimated output gap to mimic inflation increasingly closely: the coefficient estimates confirm the close correspondence between the output gap and inflation (Table 1);¹⁴ likewise, inflation accounts for a substantial fraction of the estimated output gap and vice versa (second row panels in Graph 2). But this occurs at the expense of the relevance of equation (3), the static output gap specification, which, critically, anchors potential output to actual output in the long run. Hence, the downward trend in inflation results in an estimated output gap which is large and positive at the beginning of the sample and then trends downwards (left second row panel in Graph 2). Such trending output gaps are hard to reconcile with prevailing conceptual notions of potential output.

We obtain a similar result if we replace the static output gap specification (3), with its dynamic counterpart, (4), and let the data speak (Model 3). In this case, we estimate λ_3 freely and set λ_2 so that (6) holds. The results are reported in Table 1

¹⁴ In particular, the auto-regressive coefficient for inflation, γ_1 , is now statistically insignificant, while the coefficient for the output gap, γ_2 , is strongly significant.

and the third row of panels in Graph 2. Now the coefficient estimate for β is very close to unity, indicating a very high persistence (quasi-unit root behaviour) in the output gap. The estimate of λ_3 is 0.04, which roughly translates into a weight on inflation of $(1 - \beta^2)^{-1} \sigma_1^2 / \sigma_3^2 = (1 - \beta^2)^{-1} \lambda_3 \approx 50 * 0.04 = 2$ compared with that on the output gap. Again, the coefficient for the output gap, γ_2 , is large and statistically significant and the associated graphs are very similar to the ones corresponding to Model 2. The resulting output gap is not economically plausible.¹⁵

Imposing a “natural rate of output” hypothesis, ie precluding a long-run relationship between inflation and the output gap, does not change the results much. To see this, we express the Phillips curve as a relationship between the level of the output gap and the change in inflation (Model 4).¹⁶ Again, the Phillips curve is virtually irrelevant for the estimated output gap: the coefficient on the output gap is small and insignificant coefficient (Table 1), and the left fourth-row panel of Graph 2 shows the limited information content.¹⁷

In all this, the fundamental problem with the Philips curve is the presence of a low-frequency trend in inflation. Whatever the true reason for this trend, a conventional mean-reverting output gap appears unable to fully account for it.¹⁸ In other words, the relationship in (9) is likely to be misspecified. A natural question is whether allowing for a trend in the inflation process improves matters.

The answer, unfortunately, is “no”. Model 5 seeks to do this, by rewriting equation (9) in “gap” form and specifying the unobserved inflation trend as a random walk.¹⁹ The estimated coefficients bear close resemblance to those reported in Model 3 and the Phillips curve is again largely irrelevant for the estimated output gap (fifth-row panels in Graph 2). Thus while the modification does help to mitigate the misspecification, it does so at the cost of virtually eliminating the role of inflation as a conditioning variable. Moreover, the inflation trend introduces an

¹⁵ On the other hand, if we restrict the auto-regressive coefficient β in (4) to be firmly less than unity (say between 0 and 0.9), so that the output gap is mean-reverting, the results become similar to those reported in Model 1. This highlights the high sensitivity to specification assumptions under this approach.

¹⁶ We do this by adding $(1 - \gamma_1) \pi_{t-2}$ to the right hand side of (9). This is equivalent to replacing inflation, π , in (9) by its first difference, $\Delta\pi$.

¹⁷ The main difference is that the output gap now contributes a lot to inflation (right fourth-row panel). Why is this so? By imposing the restriction on the auto-regressive coefficients, inflation becomes a random walk. As a result, it is the *sum* of the current and past output gaps that affects it. This makes it technically easier to replicate the trend in inflation even as the output gap is mean-reverting. But the link between the output gap and inflation is rather weak, as the estimates in Table 1 indicate. Therefore, the error term in (9) – which also affects the inflation rate cumulatively – is forced to exhibit a substantial degree of persistence in order to account for the mismatch in the relationship between the output gap and inflation.

¹⁸ One possible explanation for the trend is central banks’ long-run target inflation, which has arguably declined over time. For example, Woodford (2008) and Goodfriend and King (2009) derive a forward-looking New Keynesian Phillips curve that allows inflation to have a time-varying trend. In empirical applications, Milani (2006) and Ireland (2007) incorporate trend inflation as a drift-less random walk in a DSGE setting to analyse changes in the Federal Reserve’s implicit target for post-war US inflation.

¹⁹ The inflation equation is now $\hat{\pi}_t = \gamma_0 + \gamma_1 \hat{\pi}_{t-1} + \gamma_2 (y_t - y_t^*) + \varepsilon_{3,t}$ where $\hat{\pi}_t = \pi_t - \bar{\pi}_t$ with trend inflation, $\bar{\pi}_t$, assumed to follow a random walk. The inclusion of an equation for unobserved trend inflation entails another scaling factor between it and the “gap” Phillips curve. This is estimated freely.

additional end-point problem, which is bound to reduce the real-time performance of the output gap (see below).

The same is true if, instead of treating trend inflation as an unobservable state variable, we anchor it to observable inflation expectations provided by professional forecasters. This is so despite the fact that this specification has been shown to improve the forecasting performance of DSGE models substantially in a number of papers (eg Del Negro and Eusepi (2011) and Del Negro and Schorfheide (2012)). The argument is that inflation forecasts contain information about future changes in inflation and thereby implicitly about changes in the inflation target.

To show that conditioning on inflation expectations is not helpful in the current context, we rewrite the Phillips curve as follows:

$$\pi_t = \gamma_0 + \gamma_1\pi_{t-1} + \gamma_2(y_t - y_t^*) + \gamma_3\pi_t^e + \varepsilon_{3,t} \quad (10)$$

where π_t^e is a forecast of future inflation formed at date t , which we treat as exogenous in the filter.²⁰ As a proxy for π_t^e we use one-year ahead forecasts obtained from the Blue Chip Economic Indicators and the Survey of Professional Forecasters (SPF).²¹ We set the prior means of γ_1 and γ_3 to be 0.3 and 0.7, respectively, and keep the remaining priors as before. Model 6 in Table 1 reports the estimates.

At first glance, the results look impressive: the coefficient on the output gap is significant, the coefficient on the inflation forecast is large and highly significant, and the lagged inflation term is small albeit significant. However, the estimated output gap is again hardly affected by the inclusion of (10) and the output gap only accounts for a modest share of changes in inflation (last row panels of Graph 2). A possible reason is that the one-year ahead professional inflation forecasts are statistically very close to forecasts based on simple backward-looking rules.²² Hence, conditioning on these forecasts is not so different from de-trending inflation by an exponential moving average of past inflation rates.

From a technical viewpoint, the results in Table 1 underscore two general messages.

First, adding to the Kalman filter structural relationships of the form (7), where the output gap is an explanatory variable, makes output gap estimates vulnerable to specification errors. The outcome is quite sensitive to a number of specification assumptions, such as the scaling parameters, the priors on the coefficients and the dynamic terms. This complexity would increase greatly with the number of

²⁰ An alternative way, consistent with theory, would be to specify an additional observation equation that relates the forecasts to unobserved inflation expectations derived from the filter directly (eg Del Negro and Schorfheide (2012)).

²¹ The data are available from FRB Philadelphia's Real-Time Data Research Center (<http://www.phil.frb.org/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/inflation-forecasts.cfm>). We used linear interpolations to fill in missing observations at the beginning of the sample, for which the data are available only bi-annually. We also tried ten-year-ahead forecasts but the results were virtually identical.

²² For example, the correlation between the professional inflation forecasts and the forecast, π_t^f , based on the backward-looking rule $\pi_t^f = 0.9\pi_{t-1}^f + 0.1\pi_t$ is 0.93. Both series also display very similar correlation patterns with past and future inflation.

relationships and unobservable trends added to the model. And, importantly, there are no clear criteria that can help choose these auxiliary assumptions.²³

Second, and for much the same reason, it is generally difficult to interpret the significance of the estimated coefficients in a conventional way. Importantly, a high and statistically significant coefficient on the output gap in a conditioning relationship need not reflect that relationship's influence on the output gap. High significance levels can, for example, easily be obtained simply through the choice of a priori weights. A more comprehensive assessment of the contribution of conditioning variables requires evaluating their marginal impact on the output gap as well as the extent to which the estimated output gaps account for movements in the conditioning variable through procedures such as those adopted here.

III. An alternative approach to embedding economic information

The results above show that specification problems in the standard approach may easily lead to output gap estimates that are either unaffected by the additional economic information or not economically sensible. In this section, we propose an alternative way of embedding economic information that is both transparent and more robust to such specification errors.

The basic idea is to limit the number of specification assumptions and the complexity of their interactions with the data by keeping the dimensionality of the system small. Rather than adding additional structural relationships, we augment directly equation (4), the dynamic version of the univariate output gap process, with additional observable economic variables that could explain the evolution of the output gap. As such, our approach can be seen as a restricted version of the more general system-based approach used in the literature that possibly includes as many equations as the number of conditioning variables and estimates many trends jointly. Reflecting this, we will refer to our approach as the "parsimonious multivariate filter". We start from the dynamic version because, as noted earlier, its static counterpart is clearly inconsistent with the data.

Our parsimonious multivariate filter approach has several advantages. First, it can accommodate information from many economic variables while keeping the dimensionality of the system low. There is no need to introduce additional scaling factors that interact in complex ways and can have a large impact on the outcome. All that is required is to choose the desired frequency at which one wishes to explain cyclical fluctuations and then see whether candidate (conditioning) economic variables actually do so. Technically, the setup involves setting only one scalar parameter that determines the relevant frequency. This can be set, for instance, to ensure the same cyclicity as that of the traditional HP-filter. Second, the size and statistical significance of the coefficients on the conditioning variable

²³ The underlying problem is also of broader relevance for the estimation of structural DSGE models. In these settings, prior assumptions imposed on the statistical properties of the structural shocks influence strongly how the model interprets history. There is a substantial risk that misspecification may considerably contaminate historical shock decomposition. This could, for example, result in misleading conclusions about which factors drive the economy.

offers a cleaner, albeit still imperfect, benchmark by which to judge their relative importance. Third, the simplicity and transparency of the method makes it easy interpret the results, spot potential specification problems, and judge the economic plausibility of the resulting output gap estimates. Finally, the method can lead to good real-time performance, as long as the candidate variables have the right properties.

We next outline the approach and discuss its properties. We then suggest a set of natural candidate conditioning variables. Finally, we report the results from applying our approach.

The parsimonious multivariate filter: simplicity, transparency and robustness

Suppose that we add a vector z_t of economic variables, to the right hand side of (4), producing

$$y_t = y_t^* + \beta(y_{t-1} - y_{t-1}^*) + \gamma' z_t + \varepsilon_{4,t} \quad (11)$$

where $\varepsilon_{4,t}$ has mean zero and variance σ_4^2 . Together with (2), equation (11) corresponds to the loss function

$$\sum_{t=1}^T \left(\frac{1}{\sigma_4^2} (\varepsilon_{4,t})^2 + \frac{1}{\sigma_0^2} (\Delta y_{t+1}^* - \Delta y_t^*)^2 \right) \quad (12)$$

This loss function is similar to (1) with the output gap term swapped for $\varepsilon_{4,t}$, which now depends on the dynamic terms and economic variables. Since there are only two terms in the loss function, only one scaling factor, $\lambda_4 = \sigma_4^2 / \sigma_0^2$, determines their relative weight.

The scaling factor, λ_4 , can be set to preserve the same cyclicalty as that assumed by the standard HP filter, provided that $|\beta|$ is statistically strictly less than one. We follow a similar line of reasoning as in the case of the dynamic HP filter above and iterate over different values of λ_4 until (6) holds, where $y_{ALT,t}^*$ now refers to the estimate from (2) and (11).

The approach has a highly desirable property. Minimising (12) subject to (6) implies that only variables that are directly relevant in explaining output fluctuations at the chosen frequencies will receive a non-zero weight in (11). That is, a variable will enter (11) with a statistically significant coefficient *if and only if* its presence reduces the unexplained part of the estimated output gap relative to (4).

When will this happen exactly? First, for a variable to survive in (11) it is not sufficient that it be correlated with output, it must also be correlated at the frequencies implicitly set by the choice of λ_4 . In addition, the conditioning variable should have a stable mean.²⁴ For example, adding a trending variable in z_t will cause $\text{var}(y_t - y_t^*)$ to increase with t . To ensure that (6) holds in the presence of such a variable, one would have to increase the variability of the change in potential output

²⁴ Of course, stationary linear combinations of non-stationary variables may also work. Note that our setup does not require that the conditioning variables have mean zero. That said, if one of the conditioning variables has a non-zero mean, it will show up in a non-zero mean for the estimated output gap. If one wishes to exclude this possibility, the relevant variables should be de-meanned.

growth at the same rate. This, however, would also increase the loss function, (12), implying that the minimising solution is to set the coefficient on the trending variable to zero when t becomes large enough.²⁵

There are, however, at least two potential difficulties *in small samples*.²⁶

First, in such samples variables that display a high degree of cyclicity – and are therefore often informative in the present context – will tend to have pro-cyclical means. This would tend to understate booms and busts in real time and hence reduce the information content of the output gap estimates. One can mitigate the problem by using Cesàro means, as we discuss below.

Second, and conversely, in small samples even a trending (or non-stationary) variable can receive a positive weight in (11) if its stationary variation dominates its low-frequency trend. In this case, there is the risk that the low frequency trend will be passed onto the output gap estimate. This would distort the estimate of the output gap over time and undermine its real-time performance. We illustrate this possibility with an example below.

All this highlights the importance of checking the stability of the means of the candidate variables prior to the analysis. One should pay particular attention to variables whose mean exhibits persistent one-sided trends over time, because of the insidious distortions they can generate.

Potentially informative variables

There are several variables, in addition to inflation, that may contain information about the state of the business cycle.

Most naturally, the unemployment rate, u_t , and the rate of capacity utilisation, κ_t , spring to mind. Each of these is a reasonable candidate for z_t in (11). Similarly, depending on the reaction function of the central bank, the short-term real interest rate, r_t , may well move countercyclically.

Following the recent financial crisis, there has been widespread recognition that financial sector developments can have real effects (eg Aikman et al (2011), Claessens et al (2011), Schularick and Taylor (2011), Borio (2012)). To allow for this possibility, as done more comprehensively in Borio et al (2013), we add two variables that capture these factors: the growth in real total non-financial private sector credit, Δcr_t , and real residential property price growth, Δrp_t .²⁷ For instance, these variables can reflect the interaction between financing constraints, collateral

²⁵ The stationary component of a trending variable may be dominant in a small sample, but will eventually be surpassed by the trend component.

²⁶ Note also that, technically, it is not possible to identify and, hence, estimate a constant in (11). The reason is that the initial value of y_t^* already plays this role. As a result; any additional constant, included either directly or indirectly through conditioning variables with non-zero means, simply increase potential output uniformly by that amount, without any other effects on the results. For this reason, in our estimation below, we demean all conditioning variables.

²⁷ We focus on growth rates rather than levels for two reasons. First, growth rates display less persistence and have more stable means. Hence, they are more suitable for our framework. Second, using levels either requires that the variables are de-trended individually prior to the analysis or de-trended jointly with output in the Kalman-filter. Both alternatives exacerbate the end-point problem and thus worsen the real-time performance (see below).

values and wealth effects (eg Kiyotaki and Moore (1997)). In other work, they have been found to be the best proxies for the “financial cycle” (Drehmann et al (2012)).

In all, therefore, we consider six variables: inflation, the unemployment rate, capacity utilisation, the real interest rate, real credit growth, and real property price growth. Again, the estimation sample is 1980q1-2012q4. As noted above, before embarking on estimation we investigate the stability of the means for the different variables. In addition, we check for large outliers, as they can have a disproportionate impact on the estimates.

The stability of the mean of the six variables differs. This is shown in Graph 3, which plots each variable together with a sequence of its estimated mean. We obtain the sequence by extending the sample successively by one observation starting from the initial period 1970q1-1979q4. This initial period is long enough to get a meaningful starting estimate of the underlying mean provided that this is well-defined (constant).

The means of the inflation rate and, to a somewhat lesser extent, capacity utilisation, show persistent downward trends over the sample. The low-frequency trend seems to dominate the cyclical variation in the inflation rate, whereas the opposite seems to hold for capacity utilisation. Hence, capacity utilisation may be more problematic for our approach, as the gradual trend in the mean may be passed onto output gap estimates.

The means of the remaining variables show less obvious trends. That said, the means of unemployment and the real interest rate display sizeable and persistent changes over the sample. This suggests that these variables may contain low-frequency variation that reduces the reliability of their information content for output gaps. By contrast, the means of the financial variables appear comparatively more stable.²⁸

We attempt to mitigate the problem of cyclical means by using the Cesàro mean.²⁹ Its key property is that it converges faster to the population mean whenever this mean exists. The sample Cesàro mean is simply the mean of the sequence of means that we created earlier. To demonstrate its convergence properties, we also plot the sequence of Cesàro means in Graph 3 (black lines). As can be seen, these means generally fluctuate less as the sample size grows. However, if the mean is not well-defined in the sample, as seems to be the case at least for inflation and capacity utilisation, applying a Cesàro mean will generally result in a greater divergence between the mean and the actual data than if a simple mean is employed. The reason is that the Cesàro mean adjusts more slowly, so that it will lag behind.

In what follows, we apply the Cesàro mean to de-mean all conditioning variables in (11). This is consistent with the logic of the approach, which *by construction* works best when the conditioning variables have stable means.

²⁸ Ideally, one would like to test directly for mean-stability. Unfortunately, we have been unable to find a test that works well in our setup with small samples. Unit root tests are not suitable here because they test for stationarity whereas, in fact, we are only interested in mean-stability, which is a weaker condition. Stationarity tests also have been shown to have low power.

²⁹ Named after Ernesto Cesàro, a 19th century mathematician, who proved that if a sequence of numbers converges to a constant, then the sequence of arithmetic means taken over the n first elements also converge to the same constant.

However, this may also compound the shortcomings of using inflation, capacity utilisation, and possibly also the real interest and unemployment rates, which exhibit varying degrees of mean instability.

Results

We proceed to estimate equation (11) by successively including each of the six variables separately and in various combinations; we do so based on their statistical significance when considered in isolation. As before, we adopt a Bayesian approach with gamma distributed priors on the parameters. In particular, we assume that the auto-regressive parameter has a prior mean of 0.70 and standard deviation of 0.3 and that it lies in the interval between 0 and 0.95.³⁰ The remaining coefficients have prior means and standard deviations equal to 0.3. In each case, we multiply the elements of γ with the expected signs of the corresponding conditioning variables but otherwise allow them to be unrestricted in \mathbb{R}^+ . Moreover, we permit each variable to enter with a lag of up to four periods. We choose the specific lag length to maximise the contribution of the conditioning variable to the estimated output gap.³¹

Table 2, Graph 4 (individual variables) and Graph 5 (combinations) summarise the results.

The estimates for inflation (Model 1) confirm that this variable contains virtually no information about the business cycle, in line with our previous results. For example, the auto-regressive parameter, β , reaches the upper boundary of 0.95 and the coefficient on inflation is practically zero.³² The upper left panel of Graph 4 plots the corresponding output gap, which is practically identical to the one obtained from the standard univariate HP filter. The panel also shows the low value added of inflation; almost all of the variation in the estimated output gap is due to the residual in (11).

A similar result holds for the real interest rate (Model 2 and the top right-hand panel). This may be partly a reflection of the influence of inflation on central bank policies over the sample.

The results for unemployment are more encouraging (Model 3). The coefficient on unemployment is large and significant, while the autoregressive parameter is now much lower, indicating substantial explanatory power. More interestingly, the corresponding output gap differs substantially from the HP filter benchmark (second row, left panel in Graph 4). In particular, the recession in the beginning of the 1980s and the recent booms are amplified. Moreover, the unemployment rate explains the bulk of the variation in the estimated output gap.³³

³⁰ The upper boundary is rather high for a quarterly sample and very close to a unit-root. A more prudent value would be 0.85, for example.

³¹ An interesting alternative approach would be to allow for phase shifts in the spirit of Rünstler (2004).

³² The coefficient on the inflation rate is significant if the huge outlier in 2008q4 is left in the data. The graph of the associated output gap is then very similar to those corresponding to models 2 and 3 in Graph 2.

³³ This result is not likely to hold across countries, as unemployment rates in several economies have undergone structural shifts. Preliminary estimates confirm this.

The problem of conditioning on variables with subtle low-frequency trends but dominant cyclical fluctuations becomes apparent in the results for capacity utilisation (Model 4). The slight downward trend in this variable, most strongly visible in the more recent period, infects the estimated output gap (second row, right panel in Graph 4). For example, the gap is below or equal to zero *for the entire period* following the burst of the dot.com bubble in 2000. In this case, the short-run variability in capacity utilisation explains a significant fraction of the cyclical movement in output and dominates the effects of the gradual downward trend on its mean. The dominance of the cyclical component implies that capacity utilisation is assigned a significant role in the small sample despite the downward trend, which is passed onto the output gap estimates. We return to this issue in the next section, when we discuss real-time performance.

In line with the results in Borio et al (2013), we find that the proxy variables for the financial cycle, credit and property price growth, yield statistically significant results and deliver sensible output gap estimates. The auto-regressive parameter in the model for credit growth (Model 5) is statistically squarely in the stationary region and the coefficient on credit growth is highly significant. The associated output gap (third row, left panel in Graph 4) differs substantially from the HP filter benchmark, capturing the outside boom prior to the early 1990s crisis, the bust of the dot.com bubble, and the recent financial crisis. As can be seen in the graph, most of this variation is explained by credit growth. Compared with the HP filter, the approach also delivers deeper and longer recessions in the early 1990s and recently. The growth rate in the residential property price index is also statistically significant (Model 6), but explains a smaller portion of cyclical output variability overall. The associated output gap (third row, right panel) shows large deviations from the HP filter benchmark during the recent housing boom and the subsequent recession.

The results so far suggest that the unemployment rate, the capacity utilisation rate, real credit growth, and real property price growth all carry important information about the business cycle. A key strength of our approach is that it facilitates the evaluation of combinations of these variables.

Since the two variables that jointly proxy the financial cycle – credit and property prices – also appear to have stable means over the sample, and thus better satisfy the prerequisites of our approach, we first try a specification with these variables only. This combination yields a “finance-neutral” output gap discussed extensively in Borio et al (2013), shown in the top left panel of Graph 5 and as Model 7 in Table 2. Both credit and property price growth are statistically significant, with coefficients of a similar size to those found when included individually. In relative terms, credit growth explains a substantially larger portion of the estimated output gap than property price growth.

Next, we extend the “finance neutral” specification by including the unemployment rate. While the mean of this variable is somewhat unstable over the sample (Graph 3) it does not seem to trend in any specific direction. Hence, this variable – *at least for the United States* – might satisfy the prerequisites of our approach. The results are reported in Table 2 (Model 8) and the top right panel of Graph 5.

The results cast further light on the statistical properties of the approach and the impact of seasonal adjustments. The coefficients on all three variables are significant, but the larger influence, in terms of both statistical significance and contribution to the estimated output gap, comes from the unemployment rate. The joint contribution of the two financial cycle proxies amounts up to 1 percentage

point of the output gap at each point of time. However, the relative strength of unemployment rate is partly due to the presence of *high-frequency components* in the financial cycle proxies, which are not seasonally adjusted (Graph 3). The filter tries to offset these components in order to match the relatively smooth seasonally-adjusted GDP series and therefore puts more weight on the unemployment rate. If the high-frequency components are removed, the contribution of financial variables to the output gap roughly doubles and they become more significant (Model 9 in Table 2 and Graph 5 bottom left panel).³⁴

Moreover, regardless of the relative statistical contribution, the financial cycle proxies carry useful additional *economic* information. While the three variables signal similar information in the run-up to the 2007 crisis, the rapid boom in credit from mid-1985 to 1990 and the subsequent sharp credit contraction in the early 1990s add complementary information about the sustainability of output trajectories during these periods.³⁵ Conditioning on these developments yields significantly higher output gaps in first period (1985-1990) and materially lower ones in the latter (1990-1995). The resulting output gap estimates better reflect the extent to which economic activity at those times was sustainable. In other words, the financial cycle proxies appear to contain additional information in these periods and similar information in the rest of the sample compared to the unemployment rate. As a result, if one were to use them as conditioning variables, they would arguably be the most informative overall. And if the choice was restricted to only one, on the same basis credit growth would arguably be the single best one.

For completeness, we try all of the four statistically significant variables simultaneously as a final specification. This highlights the potential pitfall of relying exclusively on the relative strength of the contribution of a variable as the criterion for choosing conditioning variables. The results indicate that the capacity utilisation rate clearly dominates the other variables, both in terms of significance and value-added (Model 10 in Table 2 and bottom right-hand panel of Graph 5). That said, the slight downward trend in capacity utilisation is clearly transmitted to the estimated output gap and distorts the measure: the output gap is negative for almost the entire period since 2000. This confirms the difficulty of conditioning on variables with persistent one-sided trending means despite their apparently high informational content.

Together, the results indicate that there may be large gains from combining variables. Nevertheless, this has to be done with care, going beyond purely statistical criteria and including economic ones. Moreover, a full evaluation requires also an analysis of the real-time performance of the estimates, to which we now turn.

³⁴ We use a simple HP-filter with $\lambda = 1$ to remove the high-frequency components from the financial growth rates in Model 9. This has the advantage of roughly preserving the timing of the variables, but may introduce an end-point problem if the high-frequency swings are large enough. A more stable alternative would, for example, be to use exponential moving averages, but this would lead to phase shifts in the variables.

³⁵ In addition, one should be careful when interpreting the results of finance-neutral gaps, especially during recessions associated with financial busts. The estimates used in this paper pertain to purely linear specifications. Non-linear ones, as those also employed in Borio et al (2013), should better capture the forces at work, resulting in output gaps that are larger in the boom and smaller in the bust. Moreover, we would conjecture that capturing the non-linearities during busts requires further work.

Estimated parameters of equation (11)								Table 2
Model [†]	β	π_t	r_{t-2}	u_t	κ_t	Δcr_t	Δrp_{t-4}	λ_4
Model 1	0.95 (24.65)	0.00 (0.27)	—	—	—	—	—	72.25
Model 2	0.95 (23.46)	—	-0.00 (0.20)	—	—	—	—	71.40
Model 3	0.36 (7.40)	—	—	-1.00 (10.72)	—	—	—	25.00
Model 4	0.43 (9.17)	—	—	—	0.29 (13.37)	—	—	27.04
Model 5	0.86 (17.22)	—	—	—	—	0.43 (4.44)	—	25.50
Model 6	0.91 (19.92)	—	—	—	—	—	0.10 (2.97)	43.56
Model 7	0.83 (15.91)	—	—	—	—	0.43 (9.48)	0.07 (3.06)	20.70
Model 8	0.44 (8.79)	—	—	-0.73 (-7.65)	—	0.27 (2.19)	0.06 (1.99)	19.80
Model 9	0.36 (7.34)	—	—	-0.74 (-7.17)	—	0.61 (2.95)	0.13 (2.07)	16.81
Model 10	0.37 (7.85)	—	—	-0.19 (-2.22)	0.26 (10.83)	0.14 (1.43)	0.02 (0.76)	24.01

[†] All models estimate a system containing equations (2) and (11) on quarterly US data over the sample 1980q1-2012q4. The scaling parameter, λ_4 , is set so that (6) holds.

IV. Real-time performance

A critical factor that determines the usefulness of output gap measures, especially in policy use, is their real-time performance. It is well known that output gaps based on purely statistical filters are often subject to large ex-post revisions, undermining their value in policymaking.³⁶ This is illustrated in the upper left-hand panel of Graph 6, which compares real-time and ex-post output gap estimates for the standard HP filter. The real-time estimates completely miss, for instance, the large boom ahead of the recent financial crisis, which becomes visible only in the ex-post estimates. The average absolute deviation between the real-time and ex-post gaps is 0.65 percentage points per standard deviation in the (ex-post) gap. In a policy context, this problem is discussed in more detail and across a number of countries in Borio et al (2013).

The reason for the large ex-post revisions is simple: purely statistical filters put a disproportionate weight on the last observations in the sample, ie they are generally subject to an “end-point problem” (eg Mise et al (2005)). The HP-filtered potential output will tend to “catch-up” faster with the most recent realised values of actual output compared to past ones. Since booms are typically slow to build up but quick to end, only busts will be visible in real time.

³⁶ Various solutions to the end-point problem have been suggested in the literature. These range from using GDP forecasts in the estimations to statistical modifications to the filter, but none has been entirely successful (eg Mise et al (2005) and Garratt et al (2008)).

Given the poor performance of purely statistical filters, a natural question is whether adding economic information yields better results. The answer is “no” when the Phillips curve is considered in the standard approach (Section II), regardless of its additional information content. If the estimation assigns little weight to the Phillips curve, the real-time performance of the gap is almost identical to that of the standard HP-filter. This is illustrated in the upper middle panel of Graph 6, which depicts the real-time performance for Model 1 in Table 1. But even if the estimation does assign a high weight, the problem persists. This is illustrated for Model 2 in Table 1 in the right upper panel of Graph 6. The corresponding average absolute deviation is 0.63 percentage points per standard deviation in the gap, almost identical to that of the pure HP filter.

Can our approach deliver better real-time performance? At least three properties are important here: (i) the explanatory power of the conditioning variables; (ii) the stability of their means; and (iii) the stability of the coefficients in equation (11). Obviously, the larger the share of output fluctuations (at the desired frequencies) that can be attributed to some mean-stable variable, the smaller will be the discrepancy between real-time and ex-post estimates. Similarly, the more stable the estimated relationship, the more robust will it be to the passage of time. Here, the stability of the long-run impact of the conditioning variables on the output gap, $\gamma(1 - \beta)^{-1}$, is particularly relevant.

From this perspective, the output gaps associated with either the inflation rate or the real interest rate will obviously not perform well: the explanatory power of these conditioning variables is too low and trends in their means too pronounced. In fact, their performance is almost identical to that obtained for the HP-filter (not shown).³⁷ Hence, we turn to the variables that were found to be significant in (11) and which offer more promise.

The real-time performance of capacity utilisation is rather good, but the trend in the variable raises its ugly head again. The average absolute deviation is 0.22 percentage points per standard deviation in the gap (left second-row panel of Graph 6). But the downward trend results in a marked trend decline in the output gap, which is negative over the entire sample and there is no satisfactory way to correct for it. Applying a simple, rather than Cesàro, mean to attenuate this bias almost doubles the size of the revisions, to 0.40 percentage points (middle second row panel of Graph 6), as the mean becomes more sensitive to the passage of time. Naturally, these revisions grow larger further back in time. Alternatively, correcting for the trending mean by treating it as an unobserved latent variable re-introduces the end-point problem (right second row panel of Graph 6): the average revisions are even bigger, amounting to 0.57 percentage points. On the whole, the real-time performance of capacity utilisation deteriorates rapidly once we seek to correct for the downward trend. And if we do not, the economic plausibility of the output gap estimate is questionable.

For the remaining variables, which are generally trendless, the deviations between real-time and ex-post estimates are fairly small. The unemployment rate generates the smallest overall, with an average absolute deviation of 0.10 percentage points (left third row panel of Graph 6). If, however, unemployment is expressed as deviations from an unobserved trend in line with a natural rate

³⁷ Allowing for unobservable trends in inflation does not improve their contribution much and results in real-time performance close to that of the pure HP filter.

hypothesis, the real-time performance deteriorates to an average absolute deviation of 0.55 percentage points per standard deviation in the gap (middle third row panel of Graph 6). This reflects the additional end-point problem that is introduced by estimating a trend for unemployment.

The real-time performance of the output gap measure conditioned on credit growth comes second (third row, right-hand panel of Graph 6), with an average deviation of 0.18 percentage points, whereas property price growth (bottom left-hand panel of Graph 6) yields somewhat bigger ones (0.34 percentage points).

The combined specifications perform quite well, generally producing revisions which are almost identical to those for the best-performing variable included in each specification. The “finance-neutral” output gap, which combines credit and property prices, yields an average absolute deviation of 0.19 percentage points (bottom middle panel of Graph 6), compared with 0.18 percentage points for credit growth alone. And the absolute deviation of the measure conditioned on unemployment, credit growth and property prices is 0.12 percentage points (bottom right-hand panel of Graph 6), compared with 0.10 for unemployment alone.

Overall, conditioning on economic information can indeed alleviate the end-point problem associated with the HP filter. Our method does this flexibly, with a minimum set of assumptions, and is relatively robust to potential misspecifications. That said, it is worth reiterating that real-time performance is not the only criterion that matters. Economic plausibility of the output gap estimates matters just as much.

A key property of conditioning variables, both for real-time performance and economically sensible output gaps, is that of mean-stability. This means that trending variables should either be avoided or de-trended in a way that does not entail an end-point problem. Existing approaches typically de-trend the conditioning variables jointly in the system along with the estimate of potential output. Unfortunately, this introduces additional end-point problems that hamper real-time performance. Our approach side-steps this problem by focusing on conditioning variables that have a stable mean in the first place or, if needed, are de-trended by taking growth rates. That said, without reliable tests for mean instability, this still poses potential problems and one must be mindful of the pitfalls involved.

Conclusion

In this paper we proposed a parsimonious multivariate filter approach to incorporating economic information in estimates of potential output in order to overcome several of the drawbacks of prevailing methodologies. Rather than adding structural equations in which the output gap is an explanatory variable, we evaluate directly the information content of observable conditioning variables for the evolution of the output gap at particular frequencies. The approach is simple, transparent and more robust to misspecification. When the right variables are chosen, it can deliver estimates of output gaps that are economically plausible and perform well in real time.

A good conditioning variable has two key properties. First, it is correlated with output at the chosen frequencies. This endows it with explanatory power. Second, it

has a stable mean. This ensures that the corresponding output gap is trendless and mean-reverting – a weak condition for economic plausibility. Together, these two properties provide the basis for a good real-time performance. Admittedly, the approach is vulnerable to possible misspecification in small samples, especially in the presence of low-frequency trends. However, the simplicity of the methodology makes it easier to detect misspecification problems.

When illustrating the parsimonious approach on US data, we find that two sets of variables are especially useful. The first are proxies for the “financial cycle” (Drehmann et al (2012), namely credit and, to a lesser extent, property prices. When considered jointly, these yield the so-called “finance-neutral” output gap (Borio et al (2013)). The second is the raw unemployment rate – *not* its deviation from an unobservable natural rate. Such a deviation would generate all the problems to which the standard approach is vulnerable, not least poor real-time performance. The combination of the two sets of variables performs best on strictly statistical criteria. But the final choice depends also on economic criteria and the policy issue under examination.

In this paper we do not assess the robustness of the choice of variables across countries: the objective is simply to explain the methodology in detail. That said, other work we have done suggests that finance-neutral output gaps, and especially credit growth, work quite well across countries (see eg Borio et al (2013) for an example). The unemployment rate may be less robust: some preliminary analysis confirms that it exhibits structural shifts in several economies. In addition, other variables may perform quite well. Capacity utilisation is an obvious candidate. Its inferior performance in the case of the United States reflects the presence of a pronounced low-frequency downward trend, not low correlation with the output gap at the chosen frequencies. This trend may well be absent in other countries.

The approach can be extended in several directions. Some are purely statistical. One may, for instance, abandon Bayesian estimation and adopt classical maximum likelihood methods. This has the advantage of removing additional elements of judgment, such as those imposed through various priors on the parameters, and makes standard inference more applicable. In addition, it could be useful to adopt frequency-based methods, which would permit a more precise choice of desired frequencies. Other directions are more conceptual. In this paper, we have simply adopted the usual business cycle frequencies, effectively up to 8 years, as reflected in the typical Hodrick-Prescott filter. While this choice has become standard – indeed, part of the macroeconomics furniture, as it were – its optimality is far from obvious. The answer surely depends partly on the question, the dynamics of output and the posited “model”. In addition, while we have applied the approach to the output gap, in principle one could apply it to *any* “gap” with mean-reverting properties and to the corresponding unobservable variable. Examples include the natural rate of unemployment and natural rate of interest. We leave these issues to future research.

References

- Aikman, D, A Haldane and B Nelson (2011): "Curbing the credit cycle," Speech delivered at Columbia University Center on Capitalism and Society Annual Conference, New York.
- Bai, J and S Ng (2004): "A PANIC attack on unit roots and cointegration," *Econometrica*, 72(4), pp. 1127-1177.
- Baxter, M and R King (1999): "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 81: pp. 575–593.
- Benes, J, K Clinton, R Garcia-Saltos, M Johnson, D Laxton, P Manchevand T Matheson (2010): "Estimating potential output with a multivariate filter," *IMF Working Paper* WP/10/285.
- Borio, C (2012): "The financial cycle and macroeconomics: what have we learnt?," BIS Working Papers, no 395, December. Forthcoming in the *Journal of Banking & Finance*.
- Borio, C, P Disyatat and M Juselius (2013): "Rethinking potential output: Embedding information about the financial cycle," *BIS Working Papers*, no 404, February.
- Canova, F (1998): "Detrending and business cycle facts," *Journal of Monetary Economics*, 41, pp. 475-540.
- Claessens, S, M Kose and M Terrones (2011): "How do business and financial cycles interact?" *IMF Working Paper*, WP/11/88.
- Cogley, T and J Nason (1995): "Output dynamics in real business-cycle models," *American Economic Review*, 85(3), pp. 492-511.
- Cogley, T and T Sargent (2005): "Drifts and volatilities: Monetary policies and outcomes in the post-WWII U.S.," *Review of Economic Dynamics*, 8(2): pp. 262–302.
- Del Negro, M and Eusepi, S (2011): "Fitting observed inflation expectations," *Journal of Economic Dynamics and Control*, vol 35(12), pp 2105-2131.
- Del Negro, M and F Schorfheide (2012): "DSGE model-based forecasting," *Staff Reports*, no 554, Federal Reserve Bank of New York.
- Drehmann, M, C Borio, and K Tsatsaronis (2012): "Characterising the financial cycle: don't lose sight of the medium-term!" *BIS Working Papers*, no 380.
- European Central Bank (2011): "Trends in potential output," *Monthly Bulletin*, pp. 73-85.
- Fuhrer, J and G Moore (1995): "Inflation persistence," *Quarterly Journal of Economics* 110(1), pp. 127-159.
- Garratt A, K Lee, E Mise, and K Shields (2008): "Real-time representations of the output gap," *The Review of Economics and Statistics*, 90: pp. 792 – 804.
- Goodfriend, M and R King (2009): "The great inflation drift," *NBER Working Paper*, no. 14862.
- Hamilton, J (1994): *Time series analysis*, Princeton University Press, New Jersey.
- Harvey, A and A Jaeger (1993): "De-trending, stylized facts and the business cycle." *Journal of Applied Econometrics*, 8, pp. 231–247.

- Henry, O and K Shields (2004): "Is there a unit root in inflation?" *Journal of Macroeconomics* 26(3), pp. 481-500.
- International Monetary Fund (2013): "The dog that didn't bark: Has inflation been muzzled or was it just sleeping?" *World Economic Outlook*, April.
- Ireland, P (2007): "Changes in the Federal Reserve's inflation target: Causes and consequences," *Journal of Money, Credit, and Banking*, 39(8), pp. 1851-82.
- Kiyotaki, N and J Moore (1997): "Credit cycles," *Journal of Political Economy*, 105(2), pp. 211-48.,
- Koopman, S J and A Harvey (2003): "Computing observation weights for signal extraction and filtering," *Journal of Economic Dynamics and Control*, 27: pp. 1317-1333.
- Kuttner, K (1994): "Estimating potential output as a latent variable," *Journal of Business & Economic Statistics*, 12(3), pp. 361-68.
- Laxton, D and R Tetlow (1992): "A simple multivariate filter for the measurement of potential output," *Bank of Canada Technical Report*, no. 59.
- Milani, F (2006): "A Bayesian DSGE model with infinite horizon learning: Do mechanical sources of persistence become superfluous?" *International Journal of Central Banking* 2(6).
- Mise E, T Kim and P Newbold (2005): "On suboptimality of the Hodrick-Prescott filter at time series endpoints," *Journal of Macroeconomics*, 27, pp. 53-67.
- Nason, J and G Smith (2008): "The New Keynesian Phillips Curve: Lessons from single-equation econometric estimation," *Federal Reserve Bank of Richmond Economic Quarterly*, 94 (4), pp. 361-395.
- Okun, A (1962): "Potential GNP, its measurement and significance," Cowles Foundation, Yale University.
- Rünstler, G (2004): "Modelling Phase Shifts Among Stochastic Cycles," *Econometrics Journal*, 7: pp. 232-248.
- Schularick, M and A Taylor (2011): "Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008," *NBER Working Paper*, no 15512.
- Shephard, N (1993): "Maximum likelihood estimation of regression models with stochastic trend components," *Journal of American Statistical Association*, 88, pp. 590-595.
- Stock, J and M Watson (2007): "Why has U.S. inflation become harder to forecast?" *Journal of Money, Credit, and Banking*, 39(s1), pp. 3-33.
- (1998): "Median unbiased estimation of coefficient variance in a time-varying parameter model," *Journal of the American Statistical Association*, 93(441), pp. 349-358.
- Woodford, M (2008): "How important is money in the conduct of monetary policy?" *Journal of Monetary Economics* 44, pp. 195-222.

Appendix A: Data definitions and sources

- y_t = log real seasonally adjusted GDP. Source: *OECD Economic Outlook*.
- i_t = nominal three-month money market rate. Source: National data.
- p_t = log consumer price index. Source: *OECD Economic Outlook*.
- cr_t = log real credit to the non-financial private sector. Source: national data.
- ph_t = log real residential property price index deflated by the CPI. Source: national data.

Appendix B: The Kalman filter

A large class of dynamic models for the vector variable, y_t , can be written in so called *state-space* form³⁸

$$z_{t+1} = Fz_t + \omega_{t+1} \quad (\text{A1})$$

$$y_t = Ax_t + Hz_t + \vartheta_t \quad (\text{A2})$$

where z_t is a *state vector* that consists of possibly unobserved variables, x_t is a vector of predetermined variables, and the residual vectors ω_t and ϑ_t are assumed to be white noise processes with means zero and covariance matrices Ω and Φ , respectively. The residual vectors are further assumed to be mutually uncorrelated at all lags and leads. Equation (A1) is known as the *state equation* and equation (A2) as the *observation equation*.

The main benefits from expressing a dynamic system in the state-space form is that it can be used to derive a recursive algorithm – the Kalman filter – for formulating the least squares forecast $\hat{z}_{t+1|t}$ of z_{t+1} based on current and past values of y_t and x_t . The filter starts from initial guesses for $\hat{z}_{1|0}$, as well as the mean square error (MSE) associated with this forecast, $P_{1|0}$, and then iterates on

$$\hat{z}_{t+1|t} = F\hat{z}_{t|t-1} + FP_{t|t-1}H(H'P_{t|t-1}H + \Phi)^{-1}(y_t - Ax_t - H\hat{z}_{t|t-1}) \quad (\text{A3})$$

$$P_{t+1|t} = F(P_{t|t-1} - P_{t|t-1}H(H'P_{t|t-1}H + \Phi)^{-1}H'P_{t|t-1})F' + \Omega \quad (\text{A4})$$

The initial guesses are usually the unconditional mean and variance of z_1 . Given a sequence of estimates from (A3) and (A4), it is also possible to formulate estimates $\hat{z}_{t|T}$ – called smoothed estimates – which are based on the full sample information. This is convenient when z_t has a structural interpretation and is of interest on its own, as in the present context where it represents potential output.

To understand how the Kalman filter works, consider first the model in (A1) and (A2) as a data generating process for z_t and y_t : in each period, new realisations of ω_t , ϑ_t , and x_t determine the values for these two variables. The Kalman filter can then be seen as a way of backward engineering the shocks ω_t and ϑ_t from observations on x_t and y_t assuming that the model structure, given by (A3) and (A4), is true. For instance, the Kalman filter guess of the shock ω_{t+1} , based on the information available at t , is given by the second right hand term in (A3). Clearly this term depends on x_t and y_t through the assumed relationship in (A2) and the parameters of the model.

An additional advantage of the Kalman filter compared to other filtering techniques is that it automatically produces estimates of the MSE associated with $\hat{z}_{t|t}$, as is evident from (A4). In the present context, this implies that we can readily formulate confidence intervals for our output gap estimates. As always, the validity of such inference critically rests on statistical fit of the model given by (A1) and (A2). If this model produces a poor fit for the data on x_t and y_t , the confidence intervals will generally not be meaningful.

Finally, it is also possible to decompose the estimated state vector, z_t , into the contribution from the observed variables, x_t and y_t . This decomposition is based in the results in Koopman and Harvey (2003), who present algorithms for computing

³⁸ This section draws heavily on the exposition in chapter 13 of Hamilton (1994).

the weights implicitly assigned to unobserved components that are estimated from a model in state space form. This allows us to assess the value-added of each auxiliary variable and associated economic relationship for the estimated output gaps.

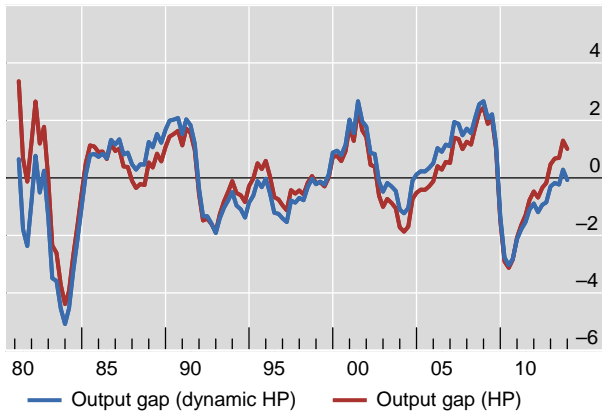
Graphs

Static and dynamic Hodrick-Prescott filter: United States

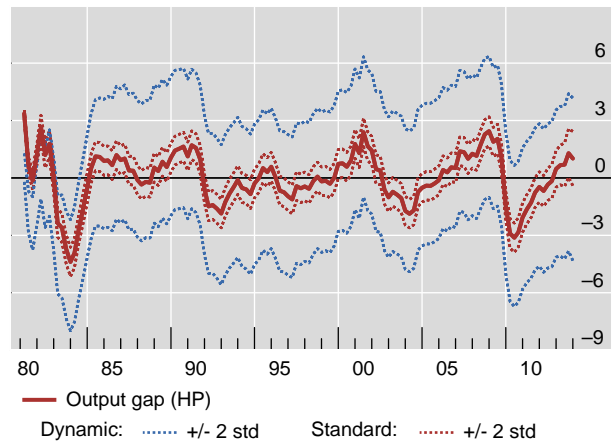
In per cent

Graph 1

Output gaps



Output gaps: confidence bands



Source: Authors' calculations.

Output gaps: United States

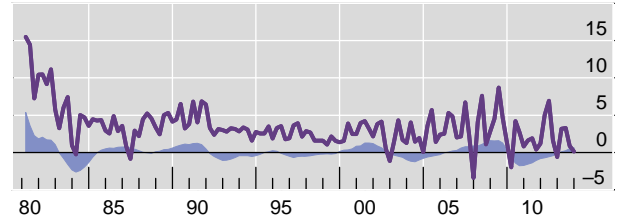
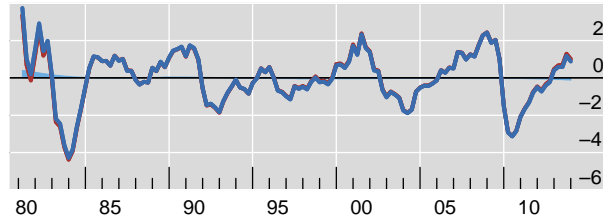
In per cent

Graph 2

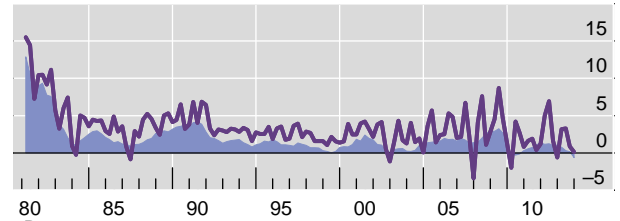
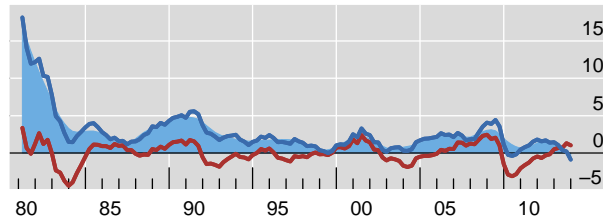
Contribution of inflation to the output gap

Contribution of the output gap to inflation

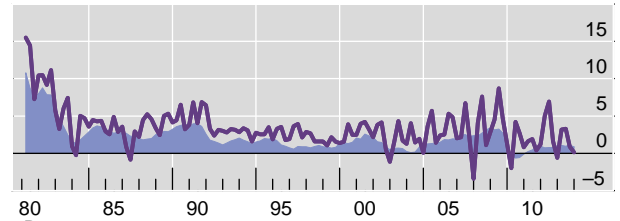
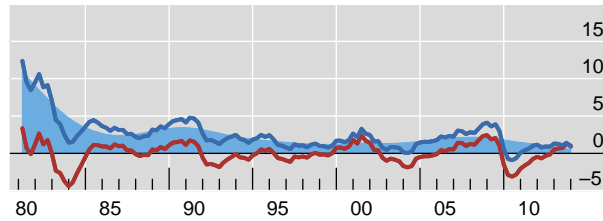
Model 1



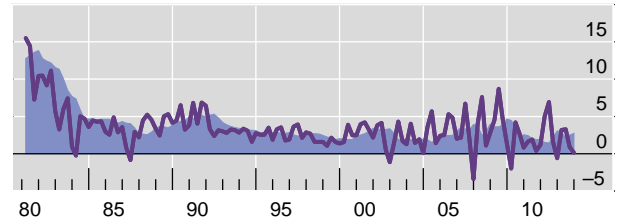
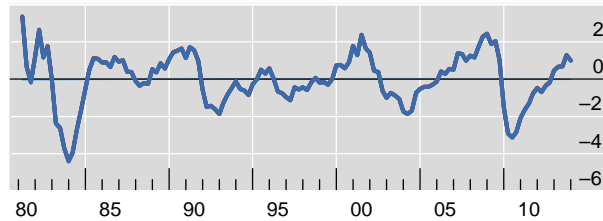
Model 2



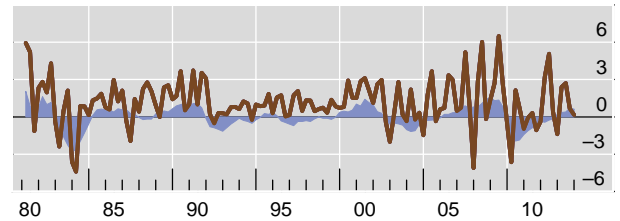
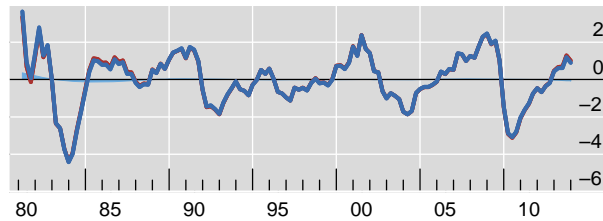
Model 3



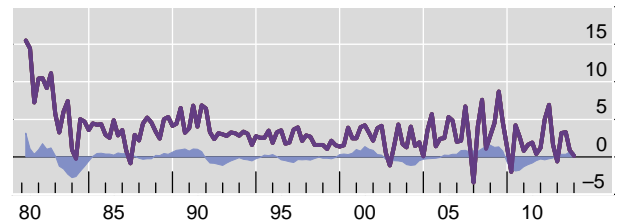
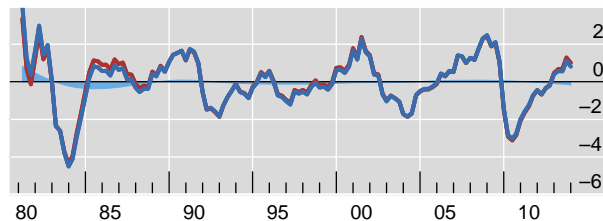
Model 4



Model 5



Model 6



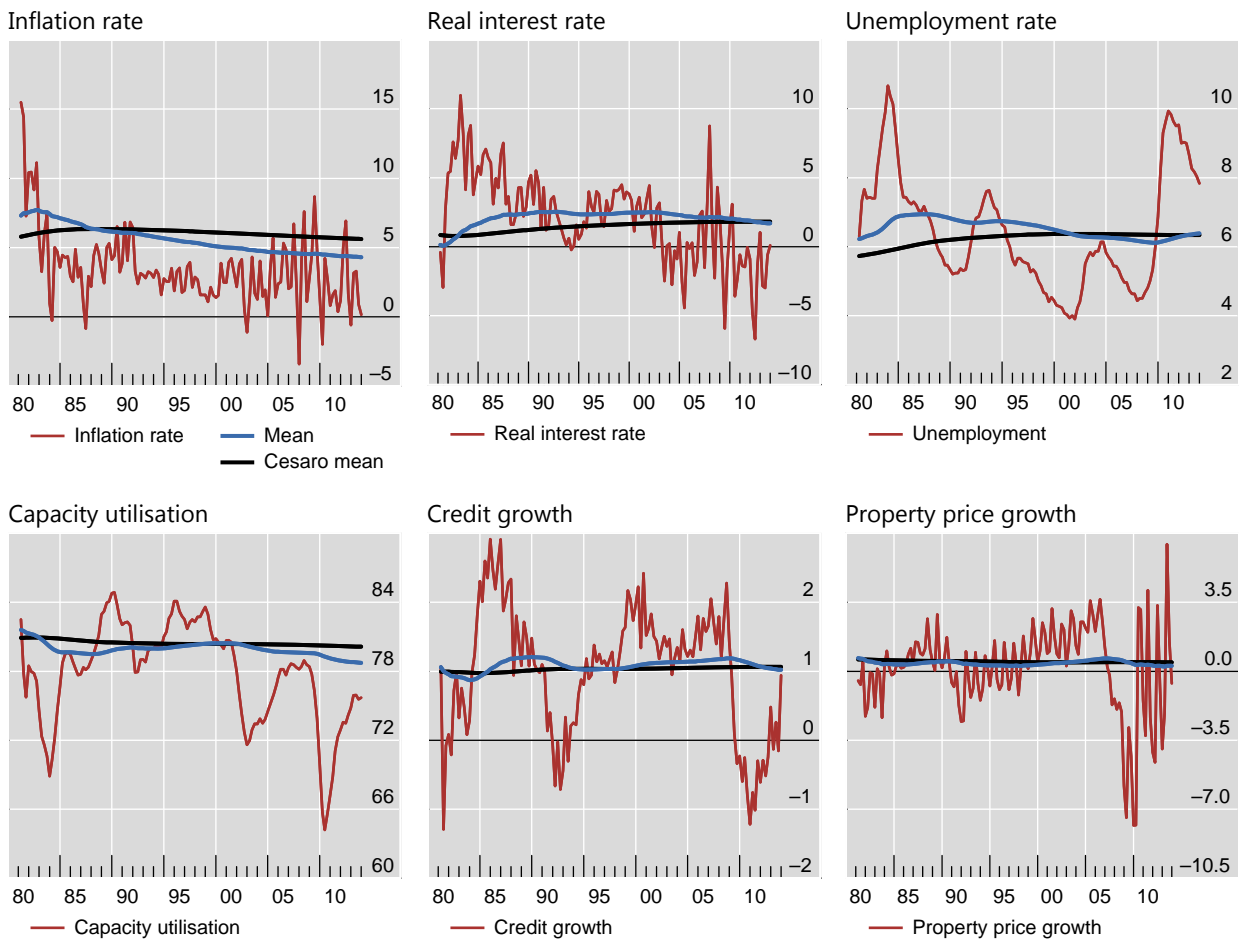
— Output gap — Output gap (HP) Due to: Inflation — Actual inflation — Inflation gap Due to: Output gap

Source: Authors' calculations.

Assessing the behaviour of the means

In per cent

Graph 3



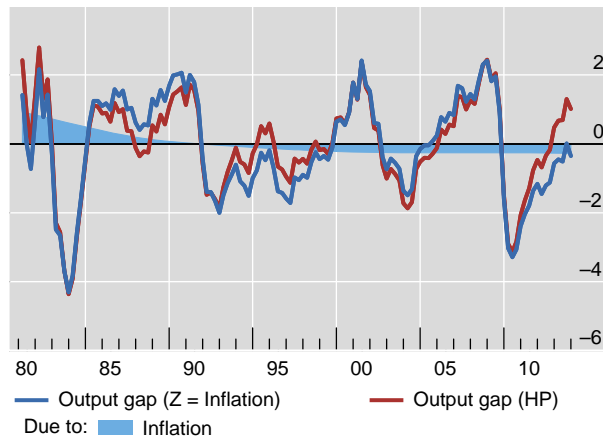
Source: Authors' calculations.

US output gaps based on individual conditioning variables: decomposition

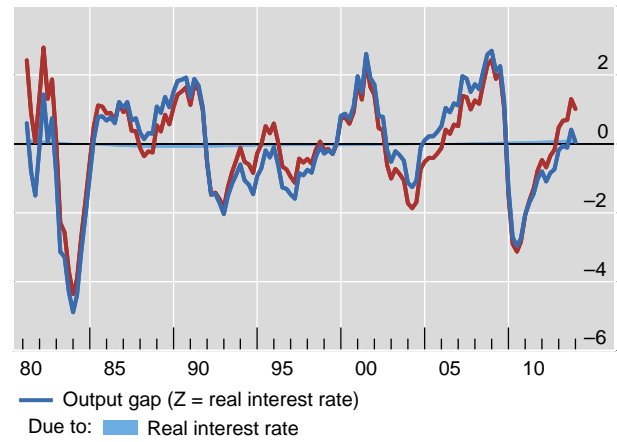
In per cent

Graph 4

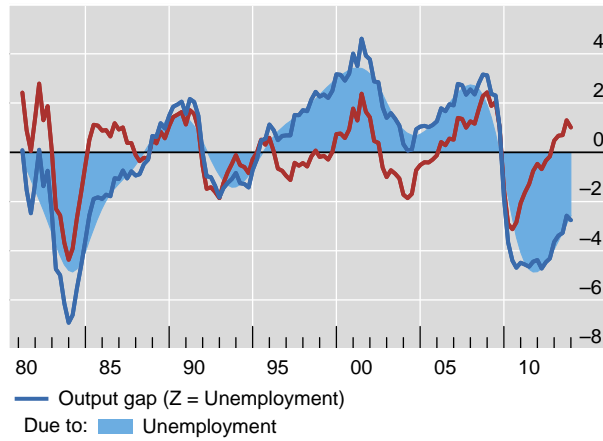
Inflation



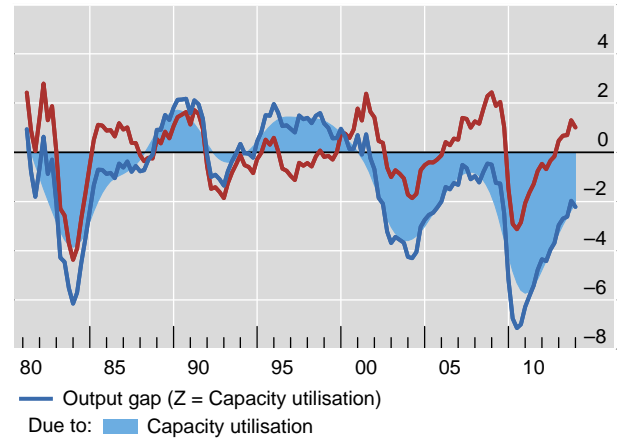
Real interest rate



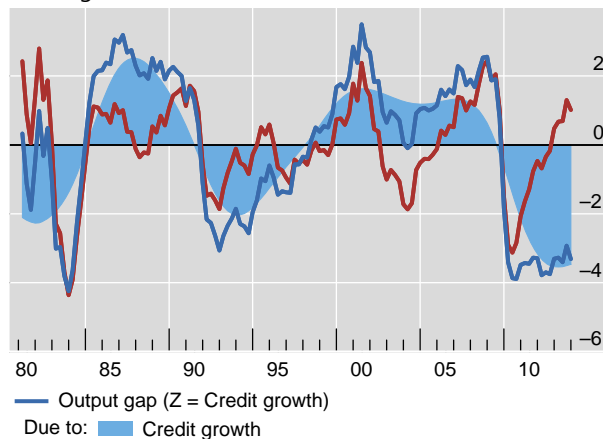
Unemployment



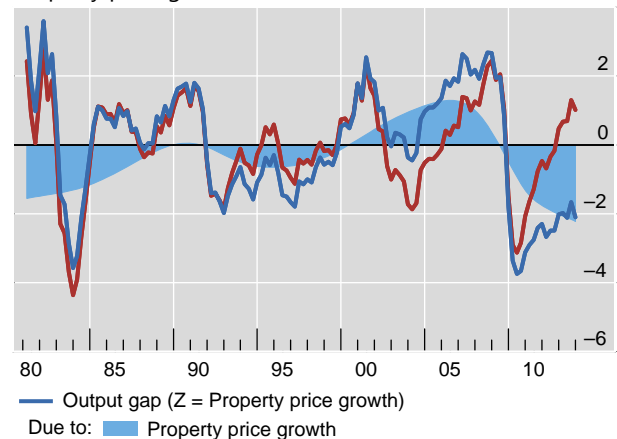
Capacity utilisation



Credit growth



Property price growth



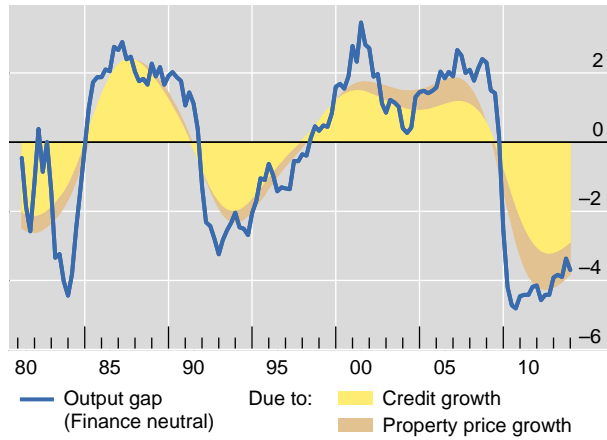
Source: Authors' calculations.

US output gaps based on combinations of variables: contributions

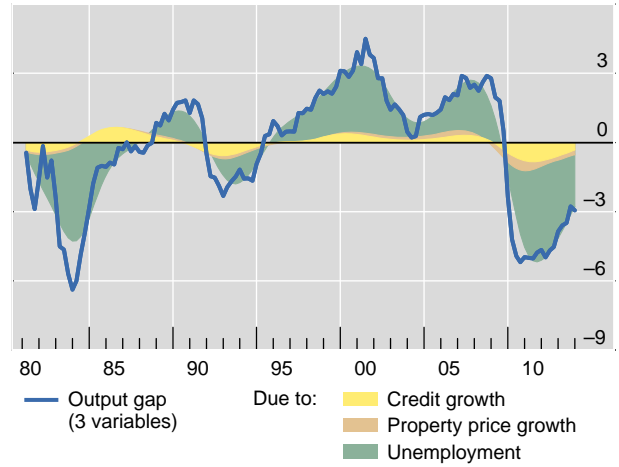
In per cent

Graph 5

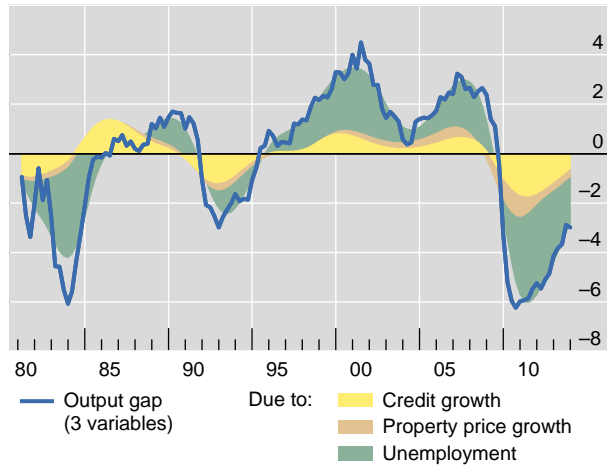
Credit and property prices: finance-neutral output gaps



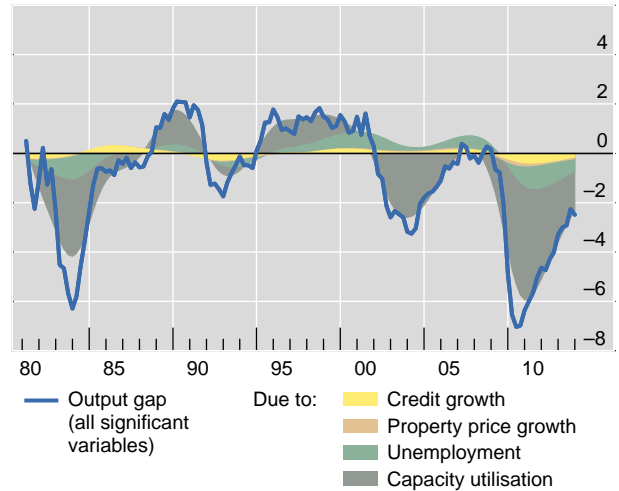
Credit, property prices and unemployment: unsmoothed



Credit, property prices and unemployment: smoothed



Four variables

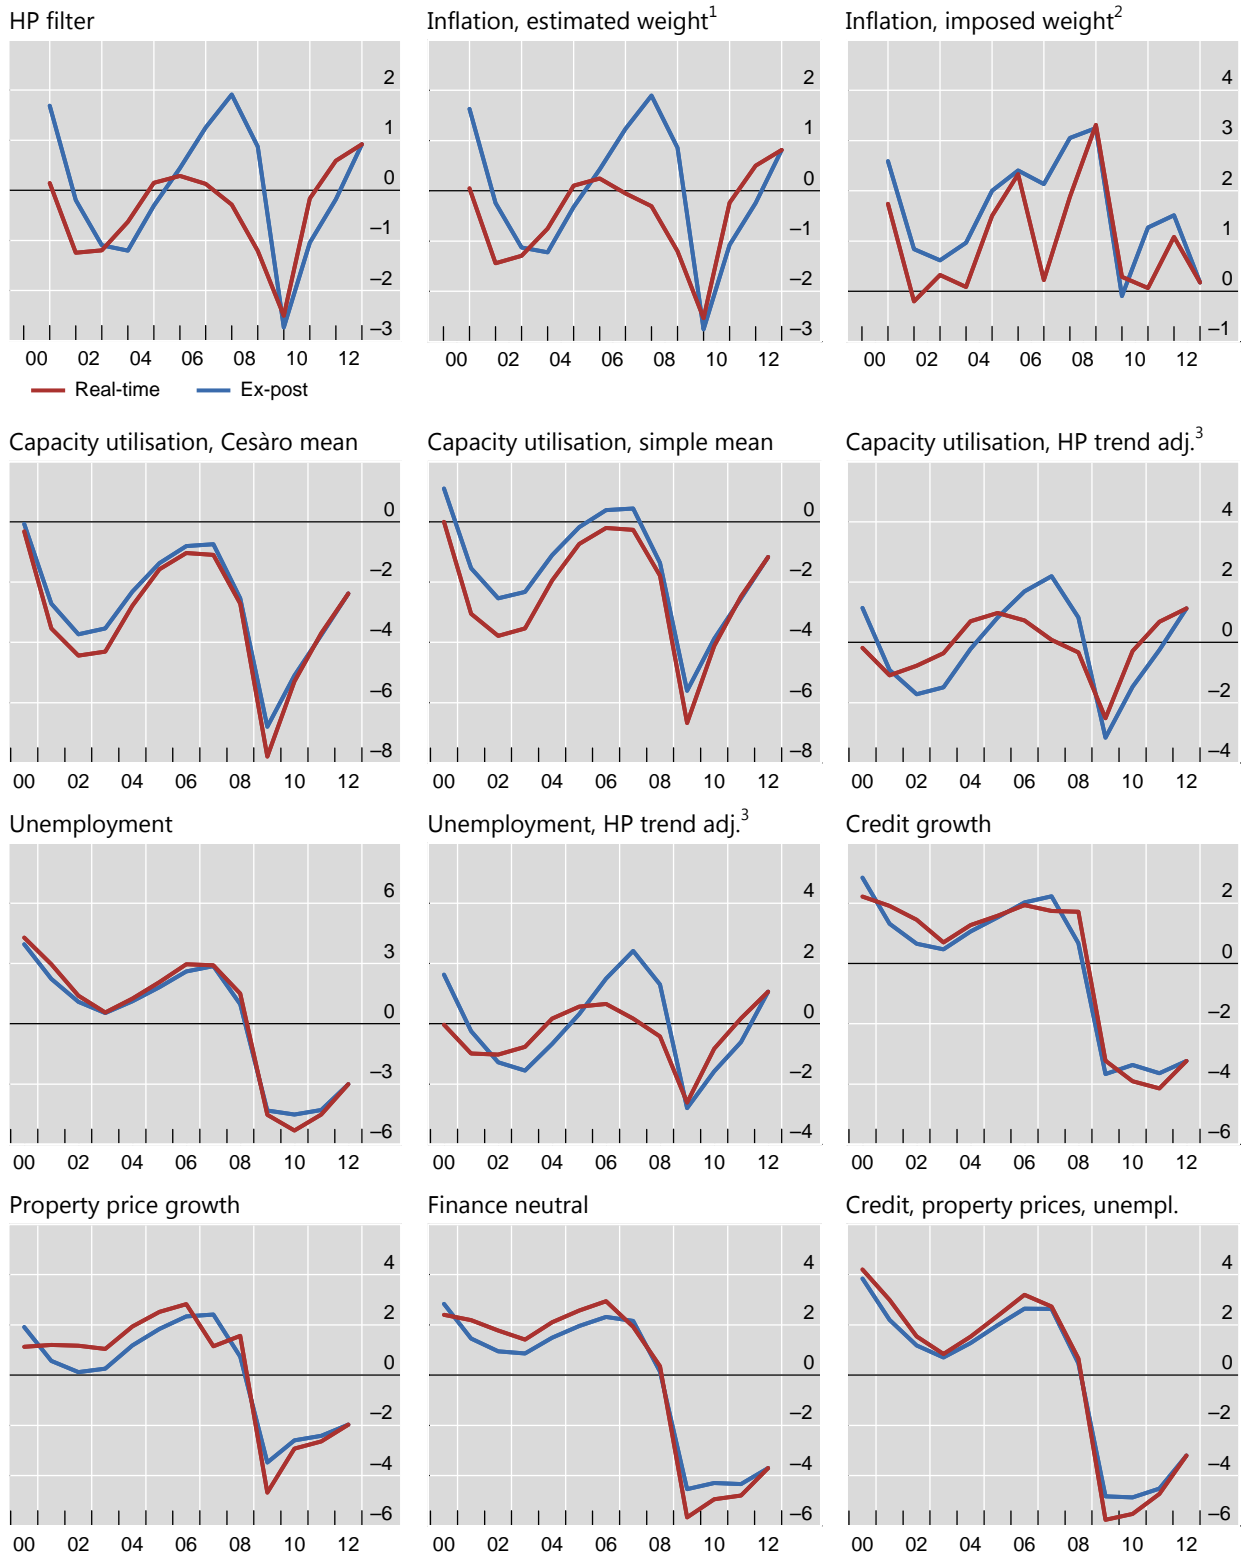


Source: Authors' calculations.

US output gaps: real-time performance

In per cent

Graph 6



¹ Value of the relative weight on the Phillips curve estimated freely (Model 1, Table 1). ² High relative weight on the Phillips curve imposed (Model 2, Table 2). ³ The conditioning variable is adjusted in real-time by a HP-filtered trend (lambda 1600).

Source: Authors' calculations.

Previous volumes in this series

No	Title	Author
441 February 2014	The global long-term interest rate, financial risks and policy choices in EMEs	Philip Turner
440 January 2014	Monetary policy and financial stability: What role in prevention and recovery?	Claudio Borio
439 January 2014	On the economics of committed liquidity facilities	Morten L Bech and Todd Keister
438 December 2013	Asia's decoupling: fact, forecast or fiction?	Lillie Lam and James Yetman
437 December 2013	International monetary policy coordination: past, present and future	John B Taylor
436 December 2013	Global spillovers and domestic monetary policy	Menzie D Chinn
435 December 2013	Is monetary policy overburdened?	Athanasios Orphanides
434 December 2013	Cyclical macroeconomic policy, financial regulation and economic growth	Philippe Aghion and Enisse Kharroubi
433 November 2013	Can non-interest rate policies stabilise housing markets? Evidence from a panel of 57 economies	Kenneth N Kuttner and Ilhyock Shim
432 October 2013	Liquidity regulation and the implementation of monetary policy	Morten L Bech and Todd Keister
431 October 2013	Transmitting global liquidity to East Asia: policy rates, bond yields, currencies and dollar credit	Dong He and Robert N McCauley
430 September 2013	Asymmetric effects of FOREX intervention using intraday data: evidence from Peru	Erick Lahura and Marco Vega
429 September 2013	On central bank interventions in the Mexican peso/dollar foreign exchange market	Santiago García-Verdú and Miguel Zerecero
428 September 2013	The impact of pre-announced day-to-day interventions on the Colombian exchange rate	Juan José Echavarría, Luis Fernando Melo, Santiago Téllez and Mauricio Villamizar
427 September 2013	Interventions and inflation expectations in an inflation targeting economy	Pablo Pincheira
426 September 2013	Order Flow and the Real: Indirect Evidence of the Effectiveness of Sterilized Interventions	Emanuel Kohlscheen
425 September 2013	The Response of Tail Risk Perceptions to Unconventional Monetary Policy	Masazumi Hattori, Andreas Schrimpf and Vladyslav Sushko
424 September 2013	Global and euro imbalances: China and Germany	Guonan Ma and Robert N McCauley

All volumes are available on our website www.bis.org.