Economic Quarterly—Volume 93, Number 2—Spring 2007—Pages 143–161

How Accurate Are Real-Time Estimates of Output Trends and Gaps?

Mark W. Watson

Trends and gaps play an important role in macroeconomic discussions. For example, the output gap (the deviation of output from its trend, or potential value) and the unemployment gap (the deviation of the unemployment rate from its trend, or "NAIRU") are standard business cycle indicators and key ingredients for Phillips curve forecasts of inflation, and likewise the trend, or long-run level of inflation is a central concern of central banks.

Trends and gaps (deviations of series from trends) are inherently two-sided concepts. By this I mean that the value of the trend in real GDP in 1987, for example, depends on how the observed value of GDP in 1987 compares to its past values (in 1986, 1985, and so forth) and to future values (in 1988, 1989, etc.). For historical analysis, the need for past and future values of the series does not pose a problem. Looking at a plot of the postwar values of real GDP, it is fairly easy to estimate its trend value by drawing a smooth curve through the plot, and various statistical formulae have been developed to mimic this freehand trend estimation procedure. However, because trends and gaps require both past and future data, it is much more difficult to estimate their values at the beginning of the sample (where there is no past data) and at the end of the sample (where there is no future data). The end-of-sample uncertainty

I have benefited from discussions with Robert Hall, Robert King, Athanasios Orphanides, and James Stock, and comments from colleagues at the Federal Reserve Bank of Richmond. Data and replication files for this research can be found at http://www.princeton.edu/~mwatson.
Mark W. Watson is a professor in the Department of Economics and the Woodrow Wilson School at Princeton University, Research Associate at the National Bureau of Economic Research, and Visiting Scholar with the Federal Reserve Bank of Richmond. The views expressed in this article are those of the author and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

in the trend is particularly problematic because these are the observations most relevant for real-time policy analysis.

The accuracy of real-time estimates of trends and gaps depends on the series under study. For example, if a series shows essentially random fluctuations around a linear trend, then the value of the trend can be accurately estimated from past observations. On the other hand, when a series shows serially correlated fluctuations around a slowly evolving trend, then future values of the series are critical to accurately separate the trend from the fluctuations. This article studies four economic indicators: industrial production, unemployment rate, employment, and real GDP to quantify the accuracy of real-time or one-sided estimates of output trends, gaps, and business cycle components.[1]

The analysis must begin with a definition of a trend and several reasonable definitions suggest themselves. Low-order polynomials in time are natural candidates, but these methods can yield unrealistic estimates of estimation errors at the ends of the sample. Martingales (or "random walks") and integrated martingales (processes for which first differences are random walks) can approximate smooth sample paths and are used to represent trends in unobserved component models (see Harvey 1989 for a detailed discussion). However, these models imply that the trend value cannot be estimated with certainty, even using an infinite amount of past and future data. This feature may or may not be reasonable, but it often leads to the conclusion that the estimated trend is inaccurate.

This article defines trends, gaps, and business cycle components using band-pass filters. These filters are moving averages of the data designed to isolate variation at specific frequencies; they are analogous to filters on an audio system that allow a user to eliminate specific frequency bands (for example, the sound from a low-frequency bass guitar or a high-frequency piccolo). In this article, the trend is defined as the cyclical movements in the time series with periods longer than the business cycle (that is, longer than 8 years); the gap includes components with periods shorter than 8 years; and the business cycle component includes components with periods between $1\frac{1}{2}$ and 8 years.

The advantage of this definition is twofold. First, it produces reasonable-looking and flexible estimates of trends, gaps, and business cycle components

[1] There are two distinct problems using real-time data to estimate trends and gaps. First, data published in real time are often subsequently revised, and these revisions can be large. Second, for the purpose of estimating trends and gaps, future values of the series are needed, so that estimates of a trend at time t will change as data becomes available for time $t+1, t+2$, etc., even if the data at time t is not revised. This article is concerned with the second problem. In particular, all analysis in this article is carried out using a 2006-vintage dataset, and to avoid confusion with actual real-time estimates, I will refer to estimates constructed using current and past values of a series as "one-sided" estimates. Orphanides (2003a) studies many of the same problems studied here and also includes analysis of data revisions.

(see the discussion and examples in Baxter and King 1999 and Stock and Watson 1999), and second, it means that historical values of these components can be estimated precisely allowing a sharp distinction between historical and one-sided analysis. Importantly, for interpreting the results shown in this article, uncertainty about the correct definition of trends, gaps, and business cycle components will only increase the real-time uncertainty of the estimates.

This is not the first article to look at this issue. For example, Staiger, Stock, and Watson (1997, 2002) quantify the uncertainty in estimates of the NAIRU; Orphanides and van Norden (2002) and Orphanides (2003a) discuss uncertainty in estimates of the output gap; Orphanides (2003b) and Orphanides and Williams (2002) discuss the effects of output gap uncertainty on monetary policy; and Hall (2005) contains a thoughtful critique of the usefulness of decomposing series in smooth trend and gap components.

The following section provides a brief review (or primer) on band-pass filtering and the Appendix contains some additional details. Section 2 presents benchmark results for one-sided estimates of the gaps based on the index of industrial production, the unemployment rate, payroll employment, and real GDP. As it turns out, the one-sided gap estimates are quite imprecise and capture only 50 percent of the variability in the gap as determined by two-sided estimates. Section 3 discusses improving the precision by using multivariate methods, but these produce only marginal improvements in the precision of the one-sided estimates. This section also shows that the reduction in volatility associated with the "Great Moderation" has greatly increased the (absolute) precision of one-sided estimates. The final section contains a brief summary and some concluding remarks.

1. A REVIEW OF BAND-PASS FILTERING

Let Y_t denote a stationary scalar stochastic process. The Spectral Representation (sometimes called the Cramér Representation) of Y is given by

$$Y_t = \int_0^{\pi} \cos(\omega t)d\alpha(\omega) + \int_0^{\pi} \sin(\omega t)d\delta(\omega), \tag{1.1}$$

where $d\alpha(\omega)$ and $d\delta(\omega)$ are zero-mean random variables that are mutually uncorrelated, are uncorrelated across frequency, and have variances that depend on frequency. The representation decomposes Y_t into a set of heteroskedastic, mutually uncorrelated, strictly periodic components. The business cycle component of Y can be defined as $Y_t^{BC} = \int_{\omega_1}^{\omega_2} \cos(\omega t)d\alpha(\omega) + \int_{\omega_1}^{\omega_2} \sin(\omega t)d\delta(\omega)$, where ω_1 and ω_2 demarcate business cycle frequencies, for example, frequencies with periods between $1\frac{1}{2}$ and 8 years. Similarly, the trend component of Y can be defined as the lower-than-business-cycle components of Y, $Y_t^{Trend} = \int_0^{\omega_1} \cos(\omega t)d\alpha(\omega) + \int_0^{\omega_1} \sin(\omega t)d\delta(\omega)$, and the gap is $Y_t^{Gap} = Y_t - Y_t^{Trend}$.

Band-pass filtering uses moving averages of the data to estimate frequency components of Y over specific frequency bands. To see how a band-pass filter

Figure 1 Band-Pass Filter Weights (Periods Less Than Eight Years—Monthly Data)

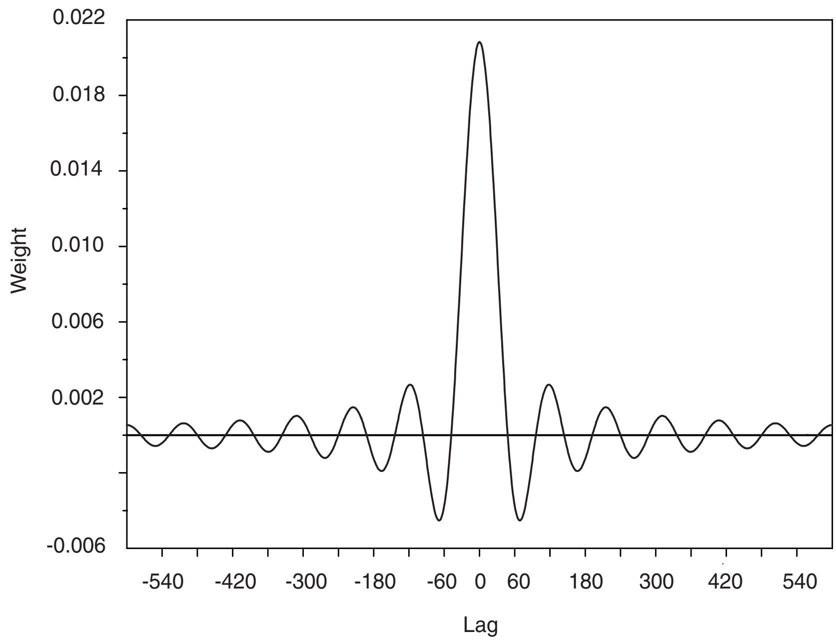

Notes: Let $c(L) = \sum_{j=-\infty}^{\infty} c_{|j|} L^{-j}$ denote the band-pass filter for monthly data frequencies with periods greater than 96 months. This figure shows the first 600 values of c_j, corresponding to weights that are used in a symmetric 100-year moving average of monthly data.

works, let X_t denote a moving average of Y_t with moving average weights c_j

$$X_t = \sum_{j=-r}^{s} c_j Y_{t-j} = c(L)Y_t, \tag{1.2}$$

where $c(L) = \sum_{j=-r}^{s} c_j L^j$ is a polynomial in the lag-operator L with coefficients c_j. As shown in the Appendix, the ω^{th} cyclical component of X_t is the ω^{th} component of Y_t transformed in two distinct ways: (1) it is shifted backward or forward in time, and (2) it is amplified or attenuated. Specifically, letting $X(t,\omega)$ and $Y(t,\omega)$ denote the ω^{th} components, $X(t,\omega) = g(\omega)Y(t-\rho(\omega),\omega)$, so that $\rho(\omega)$ denotes the time shift and $g(\omega)$ denotes the amplification factor. A calculation presented in the Appendix shows that $g(\omega) = |c(e^{-i\omega})|$ and $\rho(\omega) = \omega^{-1} \times \tan^{-1}\{Im[c(e^{-i\omega})/Re[c(e^{-i\omega})]\}$, where $i = \sqrt{-1}$ is a complex number and $c(e^{-i\omega}) = \sum_{j=-r}^{s} c_j e^{-ji\omega}$, with imaginary and real parts given by $Im[c(e^{-i\omega})]$ and $Re[c(e^{-i\omega})]$. Because the moving average operation modifies the cyclical components, $c(L)$ is called a *filter.*

A band-pass filter chooses the coefficients c_j to isolate (or "pass") a specific range (or "band") of cyclical components. To be specific, a band-pass filter that isolates frequencies between ω_{Lower} and ω_{Upper} chooses the moving average weights c_j so that that $g(\omega)$ and $\rho(\omega)$ satisfy two properties:

$$\rho(\omega) = 0 \text{ and} \tag{1.3}$$

$$g(\omega) = \left\{ \begin{array}{c} 1 \text{ for } \omega_{Lower} \leq \omega \leq \omega_{Upper} \\ 0 \text{ otherwise} \end{array} \right\}. \tag{1.4}$$

The restriction (1.3) means that the series is not shifted in time. This constraint can be satisfied by making the filter symmetric, that is, by choosing $c_j = c_{-j}$ for all j. (This makes $Im[c(e^{-i\omega})] = 0$, so that $\rho(\omega) = 0$.) The restriction (1.4) is more complicated. The Appendix shows that this constraint is satisfied by choosing

$$c_j = \left\{ \begin{array}{c} \frac{1}{j\pi}[\sin(j\omega_{Upper}) - \sin(j\omega_{Lower})] \text{ for } j \neq 0 \\ \frac{1}{\pi}[\omega_{Upper} - \omega_{Lower}] \text{ for } j = 0 \end{array} \right\}. \tag{1.5}$$

Figure 1 plots these weights for the monthly trend band-pass filter with $\omega_{Upper} = \frac{2\pi}{96}$ and $\omega_{Lower} = 0$, which passes components with periods greater than 8 years (= 96 months). The weights die out slowly. The figure plots the weights for the first 600 values of c_j, corresponding to a symmetric 100-year moving average of the data. Evidently, the weights are nonnegligible even outside this 100-year window.

The weights shown in Figure 1 produce the estimated trend in a series. The deviation of the series from the trend is the gap: $Y_t^{Gap} = Y_t - Y_t^{Trend} = [1 - c^{Trend}(L)]Y_t$, so that the band-pass filter for the gap is $1 - c^{Trend}(L)$. Thus, the weights used to construct the band-pass estimates of the gap will also decay very slowly.

Evidently, accurate estimates of the trend or gap in a series require a long two-sided moving average. This leads to problems for estimating the trend or gap for dates near the beginning of the sample period (when long lags of the series are not available) and near the end of the sample (when long leads are not available).

Two results suggest how these problems are best handled. First, Baxter and King (1999) consider the problem of constructing a finite order filter $\hat{c}(L) = \sum_{j=-s}^{s} c_{|j|}L^j$ that provides the best L^2 (or "least squares") approximation to the ideal $g(\omega)$ given in (1.4). They show that the best approximation is simply the truncated version of the infeasible infinite order filter. Second, Geweke (1978) makes the following general observation about constructing optimal estimates of filtered series: Let $X_t = \sum_{j=-r}^{s} c_j Y_{t-j}$, and suppose that data are available on a vector of random variables Z_τ from $1 \leq \tau \leq T$. Then the best

Figure 2 Two-Sided Band-Pass Estimates: Logarithm of the Index of Industrial Production

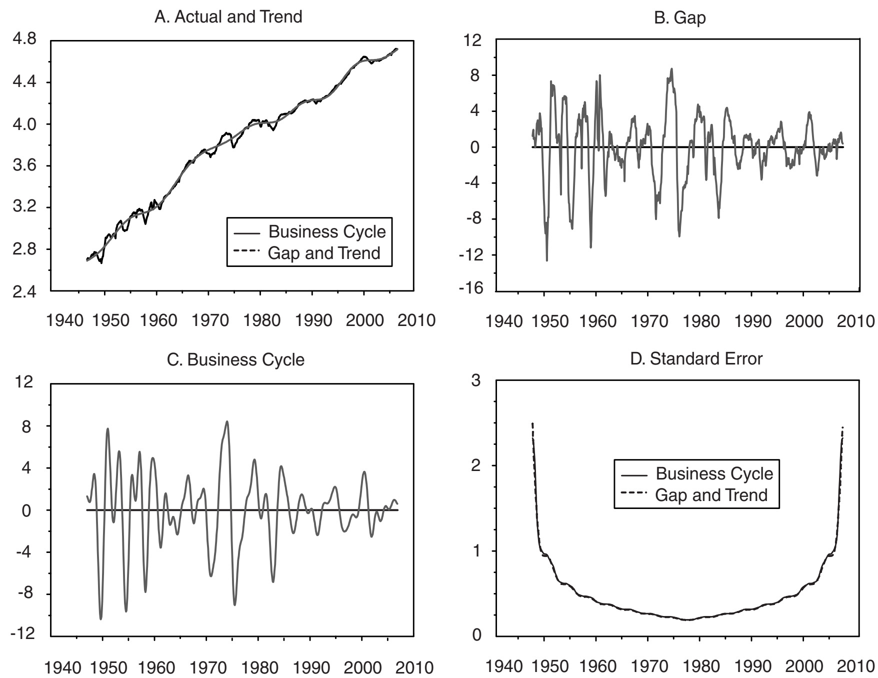

Notes: These panels show that estimated values of the band-pass estimates of the trend (periods > 96 months), the gap (periods < 96 months), and the business cycle (periods between 18 and 96 months). Panel D shows the standard errors of the estimates relative to values constructed using a symmetric 100-year moving average. Values shown in Panel A correspond to logarithms, while values shown in Panels B-D are percentage points.

(minimum mean square error) estimator of X_t is given by $E(X_t \mid \{Z_\tau\}_{\tau=1}^T) = \sum_{j=-r}^{s} c_j E(Y_{t-j} \mid \{Z_\tau\}_{\tau=1}^T)$.

Taken together, the Baxter and King (1999) and Geweke (1978) results suggest the following procedure for constructing band-pass estimates of the trend and gap. First, approximate the ideal filter using $\hat{c}(L) = \sum_{j=-s}^{s} c_{|j|} L^j$ with filter weights given by (1.4) and s chosen sufficiently large ($s = 600$). Second, letting $\{Z_\tau\}_{\tau=1}^T$ denote the sample observations on Y, construct $Y_{t|T}^{Trend} = \sum_{j=-s}^{s} c_{|j|} Y_{t/\tau}$, where $Y_{t/T} = E(Y_t \mid \{Y_\tau\}_{\tau=1}^T)$. That is, $Y_{t|T}^{Trend}$ is constructed using the ideal filter, truncated after a large number of terms and applied to the Y_t series padded into the future and past using forecasts and backcasts of the series.[2] Truncating the filter using a small value of s (an approach used by some applied researchers) is not necessary when the series is padded with forecast values of the series, and as Geweke's (1978) analysis implies, this produces a more accurate estimate of the ideal band-pass filtered series. Readers familiar with seasonal adjustment will recognize that this essentially is the procedure used in the Census X-12-ARIMA seasonal adjustment procedure (see Findley et al. 1998), and Christiano and Fitzgerald (2003) propose a one-sided band-pass filtered estimator using this procedure implemented with random-walk forecasts of Y_t.

The error in $Y_{t|T}^{Trend}$ is

$$Y_{t|T}^{Trend} - Y_t^{Trend} = \sum_{j=-s}^{s} c_{|j|}(Y_{t-j|T} - Y_{t-j}) + \sum_{|j|>s} c_{|j|} Y_{t-j}. \quad (1.6)$$

With s chosen sufficiently large, the second term is negligible and the variance of the first term can be computed from the autocovariances of the forecast/backcast errors of the Y process. (Details are provided in the Appendix.) Standard errors based on this variance formula will be used in the next section, which studies estimates of the trend, gap, and business cycle component of several economic time series.[3]

2. EMPIRICAL RESULTS

Figure 2 shows the results for computing the estimated trend (Panel A), gap (Panel B), and business cycle component (Panel C) of the logarithm of the index of industrial production (*IP*) using data from 1947:2–2006:11. These estimates are computed using a 600-term approximation to the band-pass filters and forecasts and backcasts constructed from an AR(6) model for $\Delta ln(IP_t)$. Panel D of the figure shows the standard error of the estimated components, where the standard error is computed by estimating the standard deviation of the first term on the right-hand side of (1.6) using the estimated parameters of the AR model. The estimated trend and gap components have the same standard error (because the gap and trend add to the observed series), while the estimated business cycle component is slightly smaller.

[2] As a practical consideration, it is useful to follow a suggestion by Baxter and King (1999) and modify the coefficients in the truncated trend filter so that they sum to unity. This produces an $I(0)$ estimate of the gap when the filter is applied to an $I(1)$ process and assures a bounded mean square error for the one-sided band-pass filtered estimates. The empirical analysis presented in the next section uses this modification.

[3] Harvey and Trimbur (2003) suggest an alternative procedure for approximate band-pass filtering based on an unobserved components model. An attractive feature of their proposal is that the end-of-the-sample problem is easily handled by the Kalman filter.

Figure 3 Two-Sided (Solid) and One-Sided (Dashed) Band-Pass Estimates: Index of Industrial Production

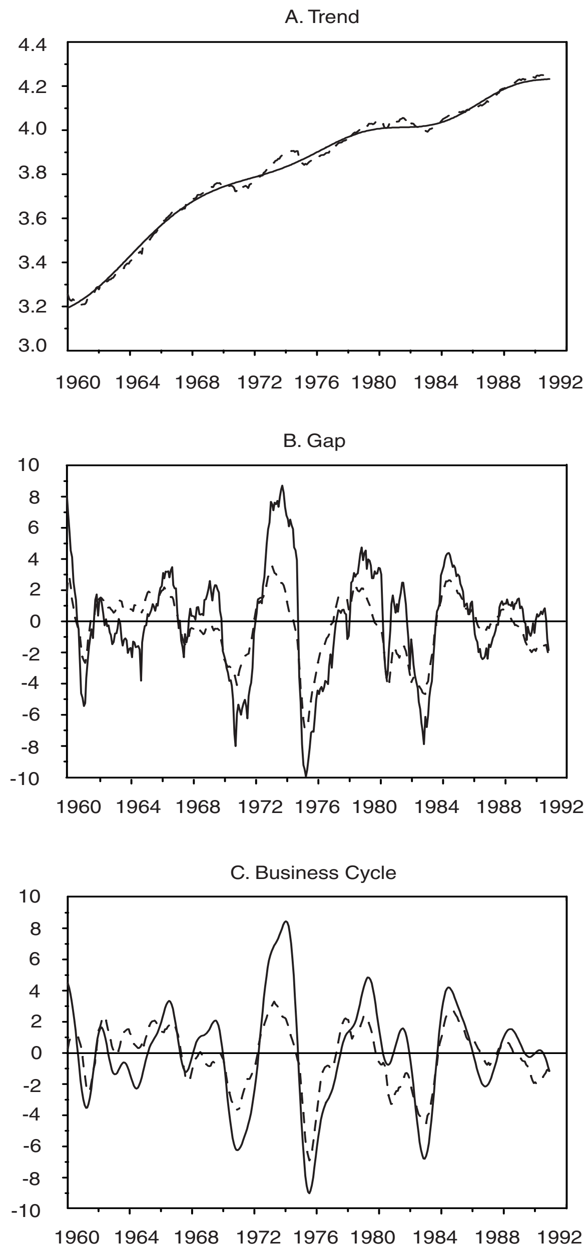

Notes: The solid lines are the two-sided estimates shown in Figure 2. The dashed lines are one-sided estimates that do not use data after the date shown on the horizontal axis.

Table 1 Joint Frequency Distribution of the Sign of $Y_{2-sided}^{BusinessCycle}$ and $Y_{1-sided}^{BusinessCycle}$—Industrial Production: 1960–1990

	$Y_{2-sided}^{BusinessCycle} > 0$	$Y_{2-sided}^{BusinessCycle} < 0$
$Y_{1-sided}^{BusinessCycle} > 0$	0.36	0.12
$Y_{1-sided}^{BusinessCycle} < 0$	0.15	0.37

Notes: This table shows the relative frequency of the events $Y_{2-sided}^{BusinessCycle} > 0$, $Y_{2-sided}^{BusinessCycle} < 0$, $Y_{1-sided}^{BusinessCycle} > 0$, and $Y_{1-sided}^{BusinessCycle} < 0$ for 1960–1990. $Y_{2-sided}^{BusinessCycle}$ is computed using the logarithm of the index of industrial production over 1947:2–2006:11, while $Y_{1-sided}^{BusinessCycle}$ uses a one-sided sample from 1947 through the date of the index.

Panel D shows that there is substantial uncertainty associated with the estimated value of the trend, gap, or business cycle components near the beginning and ends of the sample. For example, the business cycle component has a standard deviation of 2.3 percentage points at the end of the sample, which corresponds to an R^2 of only slightly greater than 50 percent. The uncertainty falls as data accumulates: when there are 15 years of data after the endpoint, the standard error falls to less than 0.4 percentage points, which corresponds to an R^2 of 99 percent.

Figure 3 shows the full-sample estimates of the components over the period 1960–1990, together with the one-sided estimates of the components. The one-sided estimates are computed using "pseudo-real-time" methods; that is, the results shown for date t are constructed using data from the beginning of the sample through time period t. Thus, for example, to compute the one-sided estimate for 1969:12, data from 1947:2–1969:12 are used to estimate an AR(6) model. This model, in turn, is used to forecast and backcast 300 observations, and the 600-term band-pass filter is applied to the resulting series.

Figure 3 shows that the one-sided estimates are considerably different than the historical estimates, consistent with the standard error results shown in Figure 2. The one-sided estimates of the gap and business cycle components are less variable than their two-sided counterparts. This "attenuation" is a property of optimal estimates: the difference between the two-sided and one-sided estimates reflects unforecastable shocks that are uncorrelated with the one-sided estimates. The figure shows the underestimation of the output gap in the late 1960s and early 1970s as highlighted in Orphanides's (2003a) discussion of the "Great Inflation."

While there is substantial error in the level of the business cycle gap, the sign of the one-sided estimate of the output gap is a useful indicator of the sign of the two-sided gap. Table 1 summarizes the joint distribution of the signs of the one-sided and two-sided estimates of the business cycle component of industrial production. During the 1960–1990 sample, $\hat{P}(Y_{2-sided}^{BusinessCycle} > 0) = 0.51$, while $\hat{P}(Y_{2-sided}^{BusinessCycle} > 0) \mid Y_{1-sided}^{BusinessCycle} > 0) = 0.71$, where $\hat{P}$ denotes the relative frequency in the sample. Similarly $\hat{P}(Y_{2-sided}^{BusinessCycle} < 0) = 0.49$, while $\hat{P}(Y_{2-sided}^{BusinessCycle} < 0 \mid Y_{1-sided}^{BusinessCycle} < 0) = 0.76$. Thus, at least over this sample period, positive and negative realizations of $Y_{1-sided}^{BusinessCycle}$ served as reasonably reliable indicators of the sign of $Y_{2-sided}^{BusinessCycle}$.

The index of industrial production is one of several cyclical indicators. Figure 4 summarizes results for three other indicators: the civilian unemployment rate and the logarithm of employment, both available monthly, and the logarithm of real GDP, a quarterly time series. The figure compares the two-sided and one-sided estimates of the trend, gap, and business cycle component for each of these series over 1960–1990. Table 2 summarizes uncertainty in the one-sided estimates by showing the estimated standard error associated with the one-sided band-pass filter estimate, the corresponding R^2, and the values of $\hat{P}(Y_{2-sided}^{BusinessCycle} > 0 | Y_{1-sided}^{BusinessCycle} > 0)$ and $\hat{P}(Y_{2-sided}^{BusinessCycle} < 0 | Y_{1-sided}^{BusinessCycle} < 0)$. The results for these series are similar to those obtained from the index of industrial production. There is significant error in the end-of-sample estimates with R^2 values of approximately 50 percent. That said, the sign of the filtered estimates predicts the sign of the two-sided estimates with a probability of approximately 70 percent.

3. IMPROVING THE ACCURACY OF ONE-SIDED BAND-PASS ESTIMATES

The error in one-sided band-pass estimates arises from the use of forecasts of future values of Y_t in place of true values. The resulting forecast errors lead to errors in the one-sided band-pass estimates. More accurate forecasts have smaller forecast errors, and, therefore, result in more accurate one-sided band-pass estimates. Forecasts may become more accurate through the use of improved forecasting methods or because of good luck associated with smaller shocks. This section quantifies the effect of both of these sources of increased accuracy for one-sided band-pass estimates of the output gaps.

The forecasts constructed in the last section were based on univariate information sets; that is, future values of Y_t were forecast using current and lagged values of Y_t. Several authors have noted that multiple indicators can, in principle, be used to increase the accuracy of output gaps. For example, Basistha and Startz (2005), Kuttner (1994), and Orphanides and van Norden

Figure 4 Two-Sided (Solid) and One-Sided (Dashed) Band-Pass Estimates: Unemployment Rate, Employment, and Real GDP

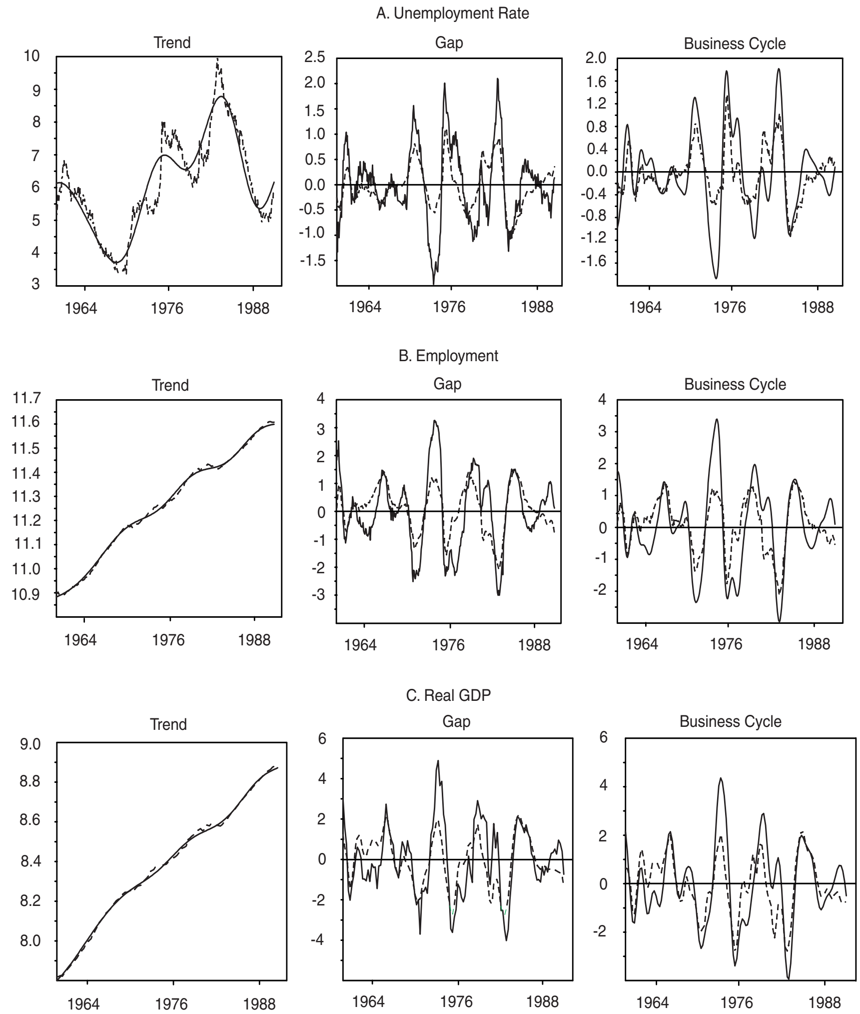

Notes: See the notes in Figure 3 for a description of the series plotted. Panel A uses monthly data on the civilian unemployment rate beginning in 1948; Panel B uses monthly data on the logarithm of nonfarm payroll employment beginning in 1947; and Panel C uses data on the logarithm of real GDP (chained $2,000) beginning in 1947.

(2002) discuss the issue in the context of Kalman filter estimates in unobserved components models, and Altissimo et al. (2006) and Valle e Azevedo (2006) discuss the issue in the context of one-sided band-pass filtered estimates.

Table 2 Summary of Results for Four Cyclical Indicators

	$Y^{Gap}_{1-sided}$		$Y^{BusinesCycle}_{1-sided}$		$\hat{P}[Y^{BusinessCycle}_{2-sided} > 0 \mid Y^{BusinessCycle}_{1-sided} > 0]$	$\hat{P}[Y^{BusinessCycle}_{2-sided} < 0 \mid Y^{BusinessCycle}_{1-sided} < 0]$
Series	**SE**	$\mathbf{R^2}$	**SE**	$\mathbf{R^2}$		
Industrial Production	2.49	0.50	2.32	0.53	0.76	0.71
Unemployment Rate	0.56	0.48	0.53	0.52	0.63	0.66
Employment	0.99	0.50	0.94	0.53	0.69	0.68
Real GDP	1.12	0.55	1.05	0.58	0.71	0.71

Notes: This table summarizes results for the four series shown in the first column. SE indicates the standard error of the end-of-sample band-pass estimate of Y^{Gap} (column 2) or $Y^{BusinessCycle}$ (column 4), and R^2 is the corresponding R^2 of the one-sided estimate. $\hat{P}[Y^{BusinessCycle}_{2-sided} > 0 \mid Y^{BusinessCycle}_{1-sided} > 0]$ shows the relative frequency of $Y^{BusinessCycle}_{2-sided} > 0$ conditional on $\mid Y^{BusinessCycle}_{1-sided} > 0$ over the 1960–1990 sample period, and similarly for $\hat{P}[Y^{BusinessCycle}_{2-sided} < 0 \mid Y^{BusinessCycle}_{1-sided} < 0]$. Estimates were constructed using data beginning in 1947:2, except those involving the unemployment rate began in 1948:2.

Table 3 Standard Errors of One-Sided Band-Pass Estimates: AR (Univariate) and VAR (Multivariate) Forecasts

	$Y_{1-sided}^{Gap}$		$Y_{1-sided}^{BusinessCycle}$	
Series	**AR**	**VAR**	**AR**	**VAR**
Industrial Production	2.01	1.88	1.88	1.80
Unemployment Rate	0.46	0.41	0.43	0.40
Employment	0.78	0.77	0.75	0.75
Real GDP	1.03	0.86	0.95	0.83

Notes: This table summarizes results for the four series shown in the first column. The entries under "AR" are the standard errors of one-sided band-pass estimates constructed using forecasts constructed by univariate AR models with six lags. The entries under "VAR" are the standard errors of one-sided band-pass estimates constructed using forecasts constructed by VAR models with six lags for monthly models and four lags for quarterly models. The VAR models included the series of interest and first difference of inflation, the term spread, and building permits. Monthly models were estimated over 1960:9–2006:11, and quarterly models were estimated over 1961:III–2006:III.

Table 3 shows results from constructing one-sided estimates using forecasts from univariate time series models with forecasts constructed from Vector Autoregressive (VAR) models. The VAR models include the first difference of inflation (for personal consumption expenditures [all items] deflator), the term spread (the difference between ten-year Treasury bond yields and three-month Treasury bill rates), and housing starts (new permits). Inflation is included because it is often used as an indicator for the output gap, and the other variables are standard leading indicators of economic activity. VARs for the monthly series (industrial production, unemployment rate, and employment) use six lags of each of the variables and the quarterly VAR for real GDP uses four lags. Results are shown for the VAR estimated from 1960:9–2006:11 for the monthly series and 1961:III–2006:III for real GDP. The autocovariances of the forecast errors, which together with the band-pass filter weights, determine the standard error of the one-sided band-pass estimates, and were computed from the estimated parameters of the VAR.

The univariate standard errors are in the columns labeled "AR" in Table 3, and the multivariate standard errors are in the columns labeled "VAR".[4] There is a small but nonnegligible increase in precision associated with the VAR forecasts. For example, the standard error of $Y^{BusinessCycle}$ falls by approximately 5 percent (from 1.88 to 1.80) for industrial production and by

[4] The univariate standard errors shown in Table 3 are slightly smaller than the values shown in Table 2 because the standard errors in Table 2 included observations from the late 1940s and 1950s, which were somewhat more volatile than those in the 1960–2006 sample period used in Table 3.

Table 4 Standard Errors of One-Sided Band-Pass Estimates

	$Y^{Gap}_{1-sided}$		$Y^{BusinessCycle}_{1-sided}$	
Series	**1960–1983**	**1984–2006**	**1960–1983**	**1984–2006**
Industrial Production	2.27	1.47	2.12	1.37
Unemployment Rate	0.54	0.39	0.51	0.37
Employment	0.84	0.37	0.80	0.35
Real GDP	1.28	0.58	1.20	0.54

Notes: This table summarizes results for the four series shown in the first column. The standard errors for the one-sided band-pass estimates are computed using the same AR models as Table 2, but the standard deviation of the AR residual is computed over the sample period shown in the column headings.

over 10 percent (from 0.95 to 0.83) for real GDP. That said, the standard errors of the one-sided estimates remain large.

The standard errors for the one-sided band-pass estimates shown in Table 2 were based on autoregressive models estimated using data from the late 1940s through 2006, and those in Table 3 used estimates from 1960 through 2006. But, as is now widely appreciated, the volatility of real economic activity over the past 20 or so years has been much lower than the volatility in the preceding 30 years.[5] This Great Moderation is evident in Figures 2–4. For real variables, such as those considered here, the reduction in volatility is well characterized as a reduction in the volatility in the "shocks" to the AR model, rather than a change in the AR coefficients. (See Ahmed, Levin, and Wilson 2004, Blanchard and Simon 2001, and Stock and Watson 2002.) This implies that AR forecasting formulae have been relatively constant over the postwar period, but that the variance of forecast errors has fallen. This, in turn, implies that the standard error of one-sided band-pass estimates has fallen.

Table 4 presents estimates of the standard errors for one-sided band-pass estimates of Y^{Gap} and $Y^{BusinessCycle}$ over the 1960–1983 and 1984–2006 sample periods. These estimates are based on the same full-sample estimated AR models used in Table 2, but with error standard deviations that are allowed to be different in the two sample periods. The standard errors shown in Table 4 for 1960–1983 are computed using the AR error standard deviation estimated over 1960–1983, and the results for 1984–2006 use standard deviations estimated over 1984–2006. The reduction in volatility has been large: the standard deviation of the AR errors has fallen by approximately 50 percent, and this reduction is reflected in an increase in the precision of the one-sided band-pass estimates. For example, these results suggest that the one-sided

[5] For example, see Blanchard and Simon (2001), Kim and Nelson (1999), McConnell and Perez-Quiros (2000), and Stock and Watson (2002).

estimate of the GDP output gap was 1.3 percentage points during 1960–1983, but fell to 0.6 percentage points in the post-1984 period.

4. SUMMARY AND CONCLUSIONS

This article has discussed the problem of estimating output trends, gaps, and business cycle components using the "one-sided" data samples that are available in real time. The results indicate that one-sided estimates necessary for real-time policy analysis are substantially less accurate than the two-sided estimates used for historical analysis. The quantitative results suggest that one-sided estimates of gaps and business cycle components have an R^2 of approximately 0.50; that is, they forecast only 50 percent of the variability in historically measured gaps and business cycle components. Thus, the answer to the question posed in the title of this article, "How Accurate are Real-Time Estimates of Trends and Gaps?" is "not very." Small improvements can be achieved using leading indicators to help forecast future values of the output series used in the construction of the one-sided estimates. The Great Moderation has led to an increase in the accuracy of forecasts of real economic variables and this accuracy, in turn, has led to an increase in the precision of one-sided output trend, gap, and business cycle estimates.

The analysis in this article was based on one-sided estimates constructed using band-pass filters, but the conclusion coincides with the conclusion reached by other authors using different methods (see, for example, Staiger, Stock, and Watson 1997 for an analysis of the unemployment rate gap using spline methods and unobserved component models and Orphanides and van Norden 2002 for an analysis of output gaps using a wide variety of methods).

APPENDIX: LINEAR FILTERS

This Appendix reviews some key results on linear filters. Let $X_t = c(L)Y_t$, where $c(L) = c_{-r}L^{-r} + \ldots + c_{-1}L^{-1} + c_0L^0 + c_1L + \ldots + c_sL^s$ is a time-invariant linear filter. From (1.1), the ω^{th} component of Y_t, $Y(t,\omega)$ is a weighted average of $cos(\omega t)$ and $sin(\omega t)$. For notational simplicity, suppose that $Y(t,\omega) = 2cos(\omega t) = e^{i\omega t} + e^{-i\omega t}$. In this case, $X(t,\omega)$ has a simple representation:

$$
\begin{aligned}
X(t,\omega) &= \sum_{j=-r}^{s} c_j Y(t-j,\omega) \\
&= \sum_{j=-r}^{s} c_j [e^{i\omega(t-j)} + e^{-i\omega(t-j)}] \\
&= e^{i\omega t} \sum_{j=-r}^{s} c_j e^{-i\omega j} + e^{-i\omega t} \sum_{j=-r}^{s} c_j e^{i\omega j} \\
&= e^{i\omega t} c(e^{-i\omega}) + e^{-i\omega t} c(e^{i\omega}).
\end{aligned}
$$

To simplify this expression further, write the complex number $c(e^{i\omega})$ in polar form, as $c(e^{i\omega}) = a + ib$, where $a = Re[c(e^{i\omega})]$ and $b = Im[c(e^{i\omega})]$. Then $c(e^{i\omega}) = (a^2 + b^2)^{\frac{1}{2}}[\cos(\theta) + i\sin(\theta)] = ge^{i\theta}$ where $g = (a^2 + b^2)^{\frac{1}{2}} = [c(e^{i\omega})c(e^{-i\omega})]^{\frac{1}{2}}$and $\theta = \tan^{-1}[\frac{b}{a}] = \tan^{-1}[\frac{Im[c(e^{i\omega})]}{Re[c(e^{i\omega})]}]$. This means that $X(t,\omega)$ can be written as

$$
\begin{aligned}
X(t,\omega) &= e^{i\omega t} g e^{-i\theta} + e^{-i\omega t} g e^{i\theta} \\
&= g[e^{i\omega[t-\frac{\theta}{\omega}]} + e^{-i\omega[t-\frac{\theta}{\omega}]}] \\
&= 2g\cos(\omega(t-\omega^{-1}\theta)) \\
&= gY(t-\omega^{-1}\theta,\omega).
\end{aligned}
$$

This expression shows that the filter $c(L)$ "amplifies" $Y(t,\omega)$ by the factor g and shifts $Y(t,\omega)$ back in time by $\omega^{-1}\theta$ time units.

Note that g and θ depend on ω, and so it makes sense to write them as $g(\omega)$ and $\theta(\omega)$. $g(\omega)$ is called the filter *gain* (or sometimes the *amplitude gain*), $g(\omega)^2 = [c(e^{i\omega})c(e^{-i\omega})]$ is called the *power transfer function* of the filter, and $\theta(\omega)$ is called the filter *phase*. In the expression below equation (1.2), $\rho(\omega) = \omega^{-1}\theta(\omega)$.

To derive the band-pass filter, first consider the problem of constructing the low-pass filter with frequency cutoff $\underline{\omega}$. Then, the gain of the band-pass filter is given by

$$
gain(c(L)) = \mid c(e^{i\omega}) \mid = c(e^{i\omega}) = \left\{ \begin{array}{c} 1 \text{ for } -\underline{\omega} \leq \omega \leq \underline{\omega} \\ 0 \text{ elsewhere} \end{array} \right\},
$$

where the second equality follows because $c(e^{i\omega})$ is real, ($c(L)$ is symmetric). Since $c(e^{-i\omega}) = \sum_{j=-\infty}^{\infty} c_j e^{-i\omega j}$, then $c_j = (2\pi)^{-1}\int_{-\pi}^{\pi} e^{i\omega j} c(e^{-i\omega})d\omega$ follows generally from $\int_{-\pi}^{\pi} e^{i\omega k} d\omega = \left\{ \begin{array}{c} 2\pi \text{ for } k = 0 \\ 0 \text{ for } k \neq 0 \end{array} \right\}$. Setting the gain equal to unity over the desired frequencies and carrying out the integration yields

$$
c_j = (2\pi)^{-1}\frac{1}{ij} e^{i\omega j} \mid_{-\underline{\omega}}^{\underline{\omega}} = \left\{ \begin{array}{c} \frac{1}{j\pi}\sin(\underline{\omega} j) \text{ for } j \neq 0 \\ \frac{\underline{\omega}}{\pi} \text{ for } j = 0 \end{array} \right\}.
$$

The difference between low-pass filters with cutoffs ω_{Lower} and ω_{Upper} is a band-pass filter that passes frequencies between ω_{Lower} and ω_{Upper}, and this difference yields the filter weights given in (1.5).

To compute the standard error of the one-sided band-pass filtered estimate, suppose initially that Y_t is $I(0)$ with moving average representation $Y_t = \theta(L)\varepsilon_t$. For any date t, $Y_t^{BP}$is a function of values of Y_j with $j \le T$ and values of Y_j for $j > T$, where T represents the final date in the sample. Write these two components as $Y_t^{BP} = w(L)Y_T + v(L^{-1})Y_T$, where $w(L)$ is a polynomial in nonnegative powers of L, and $v(L^{-1})$ is a polynomial in negative powers of L. The term $w(L)Y_T$ represents the part of Y_t^{BP}determined by values of Y_j with $j \le T$, and the term $v(L^{-1})Y_T$ represents the part of Y_t^{BP}determined by Y_j with $j > T$. The variance of the one-sided estimate of $Y_t^{BP}$is then the variance of $\{v(L^{-1})Y_T - E[v(L^{-1})Y_T \mid Y_j, j \le T]\}$. Write $v(L^{-1})Y_T = v(L^{-1})\theta(L)\varepsilon_T$, so that $v(L^{-1})Y_T - E[v(L^{-1})Y_T \mid Y_j, j \le T] = d(L^{-1})\varepsilon_T$, where $d(L^{-1}) = [v(L^{-1})\theta(L)]_-$, and the polynomial operator $[.]_-$ retains terms involving negative powers of L. The variance of $v(L^{-1})Y_T - E[v(L^{-1})Y_T \mid Y_j, j \le T]$ is then $\sigma_\varepsilon^2 \sum_j d_j^2$. Because the autocovariance generating function is symmetric, the variance associated with pre-sample values of Y_j can be computed using the same formula after time reversing the stochastic process. Finally, the same type of calculations can be used for $I(1)$ processes by computing the variance of $(v(L^{-1})-1)Y_T - E[(v(L^{-1})-1)Y_T \mid Y_j, j \le T]$.

REFERENCES

Ahmed, S., A. Levin, and B. A. Wilson. 2004. "Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?" *The Review of Economics and Statistics* 86 (3): 824–32.

Altissimo, F., R. Cristadora, M. Forni, M. Lippi, and G. Veronese. 2006. "New Eurocoin: Tracking Growth in Real Time." CEPR Discussion Paper No. 5633.

Basistha, A., and R. Startz. 2005. "Measuring the NAIRU with Reduced Uncertainty: A Multiple Indicator–Common Component Approach." Working Paper No. 04-07, Department of Economics, University of Washington.

Baxter, M., and R. G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *The Review of Economics and Statistics* 81 (4): 575–93.

Blanchard, O., and J. Simon. 2001. "The Long and Large Decline in U.S. Output Volatility." *Brookings Papers on Economic Activity* 2001 (1): 135–64.

Christiano, Lawrence J., and Terry J. Fitzgerald. 2003. "The Band Pass Filter." *International Economic Review* 44 (2): 435–65.

Findley, D. F., B. C. Monsell, W. R. Bell, M. C. Otto, and B. C. Chen. 1998. "New Capabilities and Methods of the X-12-ARIMA Seasonal Adjustment Program." *Journal of Business and Economic Statistics* 16 (2): 127–52.

Geweke, J. 1978. "The Revision of Seasonally Adjusted Time Series." Proceedings of the Business and Economics Statistics Section, American Statistical Association: 320–25.

Hall, R. E. 2005. "Separating the Business Cycle from Other Economic Fluctuations." Paper presented at the Federal Reserve Bank of Kansas City Economic Symposium, "The Greenspan Era: Lessons for the Future." Jackson Hole, WY.

Harvey, A. C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press.

Harvey, Andrew C., and Thomas M. Trimbur. 2003. "General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series." *The Review of Economics and Statistics* 85 (2): 244–55.

Kim, C. J., and C. R. Nelson. 1999. "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle." *The Review of Economics and Statistics* 81 (4): 608–16.

Kuttner, K. N. 1994. "Estimating Potential Output as a Latent Variable." *Journal of Business and Economic Statistics* 12 (3): 361–68.

McConnell, M. M., and G. Perez-Quiros. 2000. "Output Fluctuations in the United States: What has Changed Since the Early 1980's?" *American Economic Review* 90 (5): 1,464–76.

Orphanides, A. 2003a. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics* 50 (3): 605–31.

Orphanides, A. 2003b. "The Quest for Prosperity Without Inflation." *Journal of Monetary Economics* 50 (3): 633–63.

Orphanides, A., and S. van Norden. 2002. "The Unreliability of Output-Gap Estimates in Real Time." *The Review of Economics and Statistics* 84 (4): 569–83.

Orphanides, A., and J. C. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity* 2002 (2): 63–145.

Staiger, D., J. H. Stock, and M. W. Watson. 1997. "How Precise are Estimates of the Natural Rate of Unemployment?" In *Reducing Inflation: Motivation and Strategy*, eds. Christine Romer and David Romer. Chicago: University of Chicago Press.

Staiger, D., J. H. Stock, and M. W. Watson. 2002. "Prices, Wages and the U.S. NAIRU in the 1990s." In *The Roaring Nineties: Can Full Employment Be Sustained?* eds. Alan B. Krueger and Robert Solow. New York: Russell Sage Foundation.

Stock, J. H., and M. W. Watson. 1999. "Business Cycle Fluctuations in U.S. Macroeconomic Time Series." In *Handbook of Macroeconomics,* vol. 1, eds. J. Taylor and M. Woodford. Amsterdam: North Holland.

Stock, J. H., and M. W. Watson. 2002. "Has the Business Cycle Changed and Why?" In *NBER Macroeconomics Annual 2002*, eds. M. Gertler and K. Rogoff. Cambridge, MA: MIT Press.

Valle e Azevedo, J. 2006. "A Multivariate Band-Pass Filter." Job Market Paper, Department of Economics, Stanford University.