

MEASURING BUSINESS CYCLES: APPROXIMATE BAND-PASS FILTERS FOR ECONOMIC TIME SERIES

Marianne Baxter and Robert G. King*

Abstract—Band-pass filters are useful in a wide range of economic contexts. This paper develops a set of approximate band-pass filters and illustrates their application to measuring the business-cycle component of macroeconomic activity. Detailed comparisons are made with several alternative filters commonly used for extracting business-cycle components.

I. Introduction

THE STUDY of business cycles necessarily begins with the measurement of business cycles. The seminal contribution of Burns and Mitchell (1946) was influential partly because it provided a comprehensive catalogue of the empirical features of the business cycles of developed countries, notably the United States. However, their work was also important because it developed methods for measuring business cycles that could be used by other researchers working with other countries or other sample periods.

Contemporary students of the business cycle still face the same basic issue as Burns and Mitchell did fifty years ago: How should one isolate the cyclical component of an economic time series? In particular, how should one separate business-cycle elements from slowly evolving secular trends and rapidly varying seasonal or irregular components? The decomposition used by Burns and Mitchell is no longer in common use, due both to its complexity and its central element of judgment.¹ In its place, modern empirical macro-economists employ a variety of detrending and smoothing techniques to carry out trend-cycle decompositions. These decompositions are frequently ad hoc in the sense that the researcher requires only that the detrending procedure produce a stationary business-cycle component, but does not otherwise explicitly specify the statistical characteristics of business cycles. Examples of techniques in common use are application of two-sided moving averages, first-differencing, removal of linear or quadratic time trends, and application of the Hodrick-Prescott (1980) filter. Many recent studies use a battery of such methods to measure business cycles.

In our view, this proliferation of techniques for measuring business cycles has resulted from a lack of attention to an issue which Burns and Mitchell viewed as central: the definition of a business cycle. In this paper, we develop methods for measuring business cycles that require that the

researcher begin by specifying characteristics of these cyclical components. Our procedures isolate business-cycle components in a straightforward way, transforming the macroeconomic data by applying particular moving averages that are implied by these defining characteristics. Technically, we develop approximate band-pass filters that are constrained to produce stationary outcomes when applied to growing time series.²

For the empirical applications in this paper, we adopt the definition of *business cycles* suggested by the procedures and findings of NBER researchers like Burns and Mitchell. Burns and Mitchell specified that business cycles were cyclical components of no less than six quarters (eighteen months) in duration, and they found that U.S. business cycles typically last fewer than 32 quarters (eight years). We adopt these limits as our definition of the business cycle. We apply our method to several major quarterly postwar U.S. time series, including output and inflation.

Defining the business cycle as fluctuations with a specified range of periodicities results in a particular two-sided moving average (a linear filter). In the particular case of the NBER definition of the business cycle, the desired filter is a band-pass filter, i.e., a filter that passes through components of the time series with periodic fluctuations between six and 32 quarters, while removing components at higher and lower frequencies. However, the exact band-pass filter is a moving average of infinite order, so an approximation is necessary for practical applications. Thus, a central problem addressed by this paper is how to construct a good approximation to the optimal filter (i.e., the filter that accomplishes the business-cycle decomposition specified by the researcher).

In approaching this problem of filter design, we require that our method meet six objectives.³ First, as suggested above, the filter should extract a specified range of periodicities and otherwise leave the properties of this extracted component unaffected. Second, we require that the ideal band-pass filter should not introduce phase shift, i.e., that it not alter the timing relationships between series at any frequency. These two objectives define an ideal moving average of the data with symmetric weights on leads and lags. Third, we require that our method be an optimal approximation to the ideal band-pass filter; we specify a particular quadratic loss function for discrepancies between

Received for publication March, 1999. Revision accepted for publication June, 1999.

* University of Virginia

We thank Bennett McCallum, Alexei Onatski, Jim Stock, Mark Watson, and seminar participants at the Federal Reserve Bank of Richmond for helpful comments and suggestions. Filtering programs and other replication materials are available from the authors.

¹ However, it is possible to implement a judgment-free version of the Burns-Mitchell procedure, using the business-cycle dating algorithm of Bry and Boschan (1981). Two recent examples are King and Plosser (1994) and Watson (1994).

² Englund et al. (1992) and Hassler et al. (1994) proceed as we do by first defining a business cycle and then developing methods to extract business-cycle components from time series. They employ a two-step procedure in which they first detrend the time series using the Hodrick-Prescott (1980) filter, and then extract business-cycle components by band-pass filtering in the frequency domain. We discuss their method in more detail later in the paper.

³ These requirements are very similar to those that Prescott (1986) discusses in justifying use of the Hodrick-Prescott (1980) filter.

the exact and approximate filter. Fourth, we require that the application of an approximate band-pass must result in a stationary time series even when applied to trending data. Given the large body of empirical work that suggests the presence of stochastic trends in economic time series, we design our filters so that they will make the filtered time series stationary if the underlying time series is integrated of order one or two. (Equivalently, we impose the requirement that the approximate filter's frequency response is exactly zero at the zero frequency). This requirement also means that our band-pass filters will eliminate quadratic deterministic trends from a time series. Fifth, we require that the method yield business-cycle components that are unrelated to the length of the sample period. Technically, this means that the moving averages we construct are time invariant, in that the coefficients do not depend on the point in the sample. Sixth, and finally, we require that our method be operational. In the general filter-approximation problem, there is an important tradeoff involved: The ideal band-pass filter can be better approximated with the longer moving averages, but adding more leads and lags also means that observations must be dropped at the beginning and end of the sample, thus leaving fewer for analysis. We therefore experiment extensively with the application of our filter to macroeconomic time series and provide some guidance to readers about the tradeoffs involved. We recommend that researchers use moving averages based on three years of past data and three years of future data, as well as the current observation, when working with both quarterly and annual time series.

The organization of the paper is as follows. Section II describes the construction of approximate band-pass filters. In section III, we define our business-cycle filter and apply it to postwar U.S. data. Further, we investigate the implication of changing the number of leads and lags used to construct the approximate filter for certain summary statistics, using both postwar U.S. data and a specified stochastic data-generating process (for which we can compute the influence of the length of the moving average on population moments). In section IV, we contrast our business-cycle filter to the results of other commonly used procedures. In section V, we provide a detailed comparison of two "HP" filters: the cyclical filter of Hodrick and Prescott (1980) and a high-pass filter constructed using our methods. Particular attention is directed to two practical problems that researchers encounter using the Hodrick-Prescott method: unusual behavior of cyclical components near the end of the sample and the choice of the smoothing parameter for data sampled at other than the quarterly frequency. Section VI concludes the paper with a brief review of the goals and findings of the paper.

II. Filter Design

This section describes the construction of moving averages that isolate the periodic components of an economic time series that lie in a specific band of frequencies. That is, in the jargon of time-series analysis, we are interested in constructing band-pass linear filters. We are particularly

interested in designing a business-cycle filter, i.e., a linear filter that eliminates very slow-moving ("trend") components and very high-frequency ("irregular") components while retaining intermediate ("business-cycle") components.

It has long been understood that moving averages alter the relative importance of the periodic components in a time series. (See, for example, Harvey (1981, ch. 3).) If the time series y_t is stationary, then we can use frequency-domain methods to consider these implications of applying moving averages. In this paper, we employ the frequency-domain analysis to consider the design of linear filters, but we ultimately will undertake our filtering entirely in the time domain (i.e., we will simply apply moving averages to macroeconomic data). Thus, for readers who are simply interested in the practical results of our filtering methods, the current section may be skimmed or skipped.

A. Applying Moving Averages to Time Series

Applying a moving average to a time series, y_t , produces a new time series y_t^* , with

$$y_t^* = \sum_{k=-K}^K a_k y_{t-k}. \quad (1)$$

For convenience in the discussion below, we will write the moving average as a polynomial in the lag operator L , $a(L) = \sum_{k=-K}^K a_k L^k$, with L defined so that $L^k x_t = x_{t-k}$ for positive and negative values of k . We will further specialize our attention to symmetric moving averages, i.e., those for which the weights are such that $a_k = a_{-k}$ for $k = 1, \dots, K$.

One traditional use of moving averages has been to isolate or to eliminate trends in economic time series. If a symmetric moving average has weights that sum to zero, i.e., $\sum_{k=-K}^K a_k = 0$, then we show in appendix A that it has trend elimination properties. That is, if the weights sum to zero, we can always factor $a(L)$ as

$$a(L) = (1 - L)(1 - L^{-1})\psi(L) \quad (2)$$

where $\psi(L)$ is a symmetric moving average with $K - 1$ leads and lags. Symmetric moving averages with weights that sum to zero will thus render stationary series that contain quadratic deterministic trends (i.e., components of the form $\tau_t = \gamma_0 + \gamma_1 t + \gamma_2 t^2$). Further, these moving averages can also make stationary the stochastic trends that arise when a time series is a realization of an integrated stochastic process (of the I(1) or I(2) type in the lexicon of Engle and Granger (1987)).

Turning to analyzing the effect of filtering from a frequency-domain perspective, the Cramer representation of a zero-mean stationary time series y_t is

$$y_t = \int_{-\pi}^{\pi} \xi(\omega) d\omega. \quad (3)$$

That is, the time series is expressed as the integral of random periodic components, the $\xi(\omega)$, that are mutually orthogonal

$(E\xi(\omega_1)\xi(\omega_2))' = 0$ for $\omega_1 \neq \omega_2$). In turn, the filtered time series can be expressed as

$$y_t^* = \int_{-\pi}^{\pi} \alpha(\omega) \xi(\omega) d\omega, \quad (4)$$

where $\alpha(\omega) = \sum_{h=-K}^K a_h e^{-i\omega h}$ is the frequency-response function of the linear filter. (The frequency response $\alpha(\omega)$ indicates the extent to which y_t^* responds to y_t at frequency ω , in the sense that $\alpha(\omega)$ is the weight attached to the periodic component $\xi(\omega)$.) Since the periodic components $\xi(\omega)$ are orthogonal, it follows that we can write the variance of the filtered series as

$$\text{var}(y_t^*) = \int_{-\pi}^{\pi} |\alpha(\omega)|^2 f_y(\omega) d\omega. \quad (5)$$

where $|\alpha(\omega)|^2$ is the squared gain or transfer function of the linear filter at frequency ω and $f_y(\omega) = \text{var}(\xi(\omega))$ is the spectral density of the series y at frequency ω . At a given frequency, the squared gain thus indicates the extent to which a moving average raises or lowers the variance of the filtered series relative to that of the original series. The gain $|\alpha(\omega)|$ is similarly the effect on the standard deviation at a particular frequency: We thus use it in various figures below as a measure of the consequences of filtering.

In terms of our discussion below, it is important to note that the frequency-response function $\alpha(\omega)$ takes on a value of zero at frequency zero if and only if we require that the sum of the filter weights is zero ($\alpha(0) = \sum_{h=-K}^K a_h e^{-i0h} = 0$ if and only if $\sum_{h=-K}^K a_h = 0$).

We turn next to the problem of designing filters to isolate specific frequencies in the data. Our method is to use frequency-domain logic to design a moving average that emphasizes specified frequency bands, but we also require that our business-cycle filter have the trend-elimination properties discussed in this section, so that it can be meaningfully applied to economic time series which are nonstationary. We thus require that our business-cycle filter has a frequency response function with $\alpha(0) = 0$.

B. The Low-Pass Filter

A basic building block in filter design is the low-pass filter—a filter that retains only slow-moving components of the data. An ideal symmetric low-pass filter, which passes only frequencies $-\underline{\omega} \leq \omega \leq \underline{\omega}$, has a frequency-response function given by $\beta(\omega) = 1$ for $|\omega| \leq \underline{\omega}$, and $\beta(\omega) = 0$ for $|\omega| > \underline{\omega}$. The frequency-domain implication of symmetry in the weights is that $\beta(\omega) = \beta(-\omega)$.

Let $b(L) = \sum_{h=-\infty}^{\infty} b_h L^h$ denote the time-domain representation of this ideal low-pass filter. The filter weights b_h may be found by the inverse Fourier transform of the frequency response function

$$b_h = \frac{1}{2\pi} \int_{-\pi}^{\pi} \beta(\omega) e^{i\omega h} d\omega. \quad (6)$$

Evaluating the integral above (see appendix B for the details), the filter weights b_h for the ideal filter are

$$b_0 = \underline{\omega}/\pi \quad \text{and} \quad b_h = \sin(h\underline{\omega})/h\pi \quad \text{for } h = 1, 2, \dots \quad (7)$$

While the weights tend to zero as h becomes large, notice that an infinite-order moving average is necessary to construct the ideal filter. Hence, we are led to consider approximation of the ideal filter with a finite moving average $a(L) = \sum_{h=-K}^K a_h L^h$; this approximating filter has a frequency-response function $\alpha_K(\omega) = \sum_{h=-K}^K a_h e^{-i\omega h}$.

C. Approximation of Symmetric Filters

If one is considering the general problem of choosing an approximate filter, $\alpha_K(\omega)$, to approximate a specific filter $\beta(\omega)$, then a natural approximation strategy is to choose the approximating filter's weights a_h to minimize

$$Q = \frac{1}{2\pi} \int_{-\pi}^{\pi} |\delta(\omega)|^2 d\omega, \quad (8)$$

where $\delta(\omega) = \beta(\omega) - \alpha_K(\omega)$ is the discrepancy arising from approximation at frequency ω . This loss function attaches equal weight to the squared approximation errors at different frequencies.

There is a remarkable, general result for this class of optimization problems: The optimal approximating filter for given maximum lag length, K , is constructed by simply truncating the ideal filter's weights b_h at lag K . Thus, the optimal approximate low-pass filter sets $a_h = b_h$ for $h = 0, 1, \dots, K$, and $a_h = 0$ for $h \geq K + 1$, where the weights b_h are those given in equation (7) above.

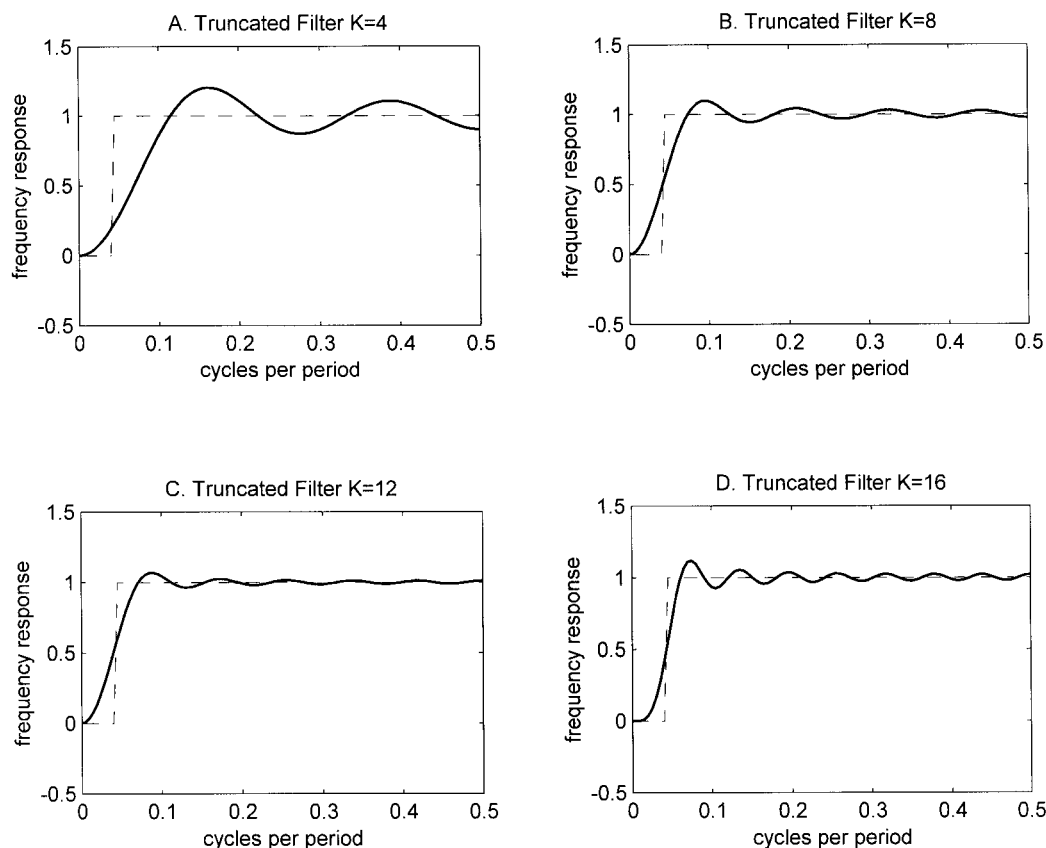
D. Construction of High-Pass and Band-Pass Filters

High-pass and band-pass filters are easily constructed from low-pass filters. Before precisely defining these additional filters, we establish some notation that we use throughout the rest of the paper. It is more natural for us to work empirically using terms of periodicity of cycles than frequencies (periodicity is related to frequency via $p = 2\pi/\omega$). Thus, we let $LP_K(p)$ denote the approximate low-pass filter that is truncated at lag K and that passes components of the data with periodicity greater than or equal to p . Since the ideal filter involves $K = \infty$, the ideal low-pass filter is denoted $LP_{\infty}(p)$.

The ideal high-pass filter $HP_{\infty}(p)$ passes components of the data with periodicity less than or equal to $p = 32$ (illustrated by the dashed line in figure 1).⁴ A low-pass filter removes high-frequency cycles while retaining low ones, the

⁴ In this figure, as in others below, the horizontal axis is labeled “cycles per period” and runs from 0 to 1/2. More traditionally, figures like these run from 0 to π , but we use our normalization since it makes it easy to calculate the periodicity by taking the reciprocal of the value on the axis. For example, the “cutoff frequency” for the high-pass filter corresponds to a period of $p = 32$ time units (presumed to be quarters of a year in view of empirical work below), and, hence, $\underline{\omega} = 1/32 \approx 0.03$. However, for the analytical results below, we use the conventional definition that the frequency ω has as its domain the interval $-\pi \leq \omega \leq \pi$.

FIGURE 1.—CONSTRAINED APPROXIMATE BAND-PASS FILTERS



high-pass filter does the reverse task, and the original time series is just the sum of its low-frequency and high-frequency components. Thus, the high-pass filter weights are $1 - b_0$ at $h = 0$ and $-b_h$ at $h = \pm 1, 2, \dots$. Correspondingly, the optimal approximate high-pass filter, $HP_K(p)$, is simply constructed by truncating the weights of $HP_\infty(p) = 1 - LP_K(p)$.⁵

The ideal band-pass filter passes only frequencies in the ranges $\underline{\omega} \leq |\omega| \leq \bar{\omega}$. It is therefore constructed from two low-pass filters with cutoff frequencies $\underline{\omega}$ and $\bar{\omega}$: We denote the frequency responses of these filters as $\beta(\omega)$ and $\underline{\beta}(\omega)$. Then, to get the desired band-pass frequency response, we form the band-pass filter's frequency response as $\beta(\omega) - \underline{\beta}(\omega)$ since this will give unit frequency response on the frequency bands $\underline{\omega} \leq |\omega| \leq \bar{\omega}$ and zero elsewhere. It is then easy to derive the filter weights for a band-pass filter. If we let \underline{b}_h and \bar{b}_h be the filter weights for the low-pass filters with cutoffs $\underline{\omega}$ and $\bar{\omega}$, then the band-pass filter has weights $\bar{b}_h - \underline{b}_h$. The dashed line in figure 2 plots an ideal band-pass filter that passes through cycles of length between 6 and 32 quarters, which corresponds to the Burns-Mitchell definition of business-cycle frequencies.

We use a similar notation for the approximate band-pass filters to that developed above for the high- and low-pass filters: $BP_K(p, q)$ denotes our approximate band-pass filter

that passes cycles between p and q periods in length, for given truncation point K , where p denotes the shortest cycle length passed by the band-pass filter and q denote the longest cycle length (in figure 2, $p = 6$ and $q = 32$). We construct $BP_K(p, q)$ by truncating the ideal band-pass filter.

E. Constraints on Specific Points

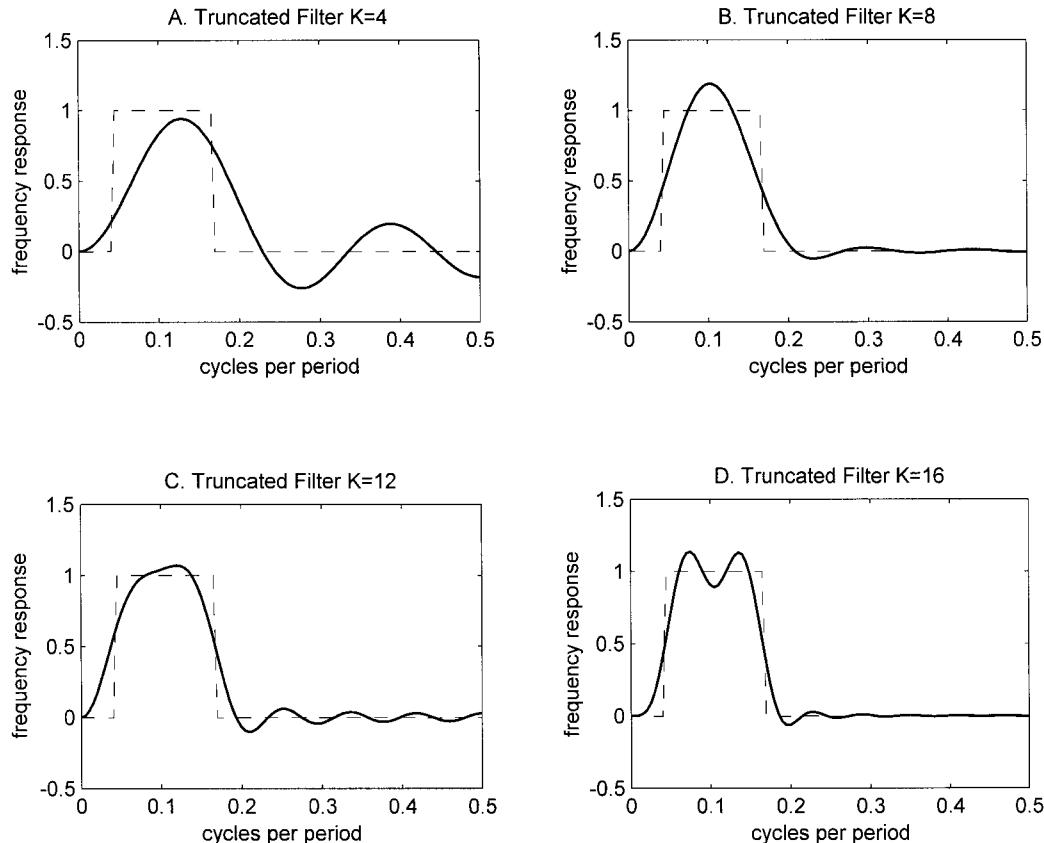
The minimization problem described above may be reformulated to recognize that certain points are of particular concern to the researcher. This approach to filter design has been advanced by Craddock (1957) in the statistics literature and discussed in the context of designing filters to eliminate trend by Granger and Hatanaka (1964, section 8.4), but does not appear to have been much followed up on in applied work in economics.

As an example of our approach, suppose that we want to design a low-pass filter that places unit weight at the zero frequency ($\alpha_K(\omega) = 1$ at $\omega = 0$). This is equivalent to requiring that the filter weights sum to unity (since $1 = \alpha_K(0) = \sum_{k=-K}^K a_k e^{0} = \sum_{k=-K}^K a_k$). If we construct approximating low-pass filters in this way, then the corresponding high-pass and band-pass filters will place zero weight at the zero frequency, and, as we have seen above, this will mean that they give rise to stationary time series when applied to a range of nonstationary time series.⁶

⁵ This is implied by the result discussed in section II, B, that approximation of the ideal low-pass filter simply involves truncation of the ideal filter's weights at lag K .

⁶ Equivalently, one can consider the problem of approximating a desired band-pass or high-pass filter subject to the constraint that the weights sum

FIGURE 2.—CONSTRAINED APPROXIMATE HIGH-PASS FILTERS



The constraint that $\alpha_K(0) = 1$ may be incorporated as a side condition to the minimization problem discussed above. Using the results of appendix B, we find the following modification of the optimal approximate filter weights, a_h , as functions of the weights of the ideal low-pass filter, b_h ,

$$a_h = b_h + \theta, \quad (9)$$

where θ is a constant that depends on the specified maximum lag length, K . That is, since we require that the filter weights sum to one, ($\sum_{h=-K}^K a_h = 1$), the normalizing constant is $\theta = (1 - \sum_{h=-K}^K b_h)/(2K + 1)$. Thus, the constraint that the low-pass filter places unit weight at the zero frequency results in a relatively simple adjustment of the filter weights.

Similar adjustments are necessary when constructing optimal truncated high-pass and band-pass filters subject to constraints on the frequency-zero value of the frequency-response function. As discussed above, the unconstrained band-pass filter has weights that are the difference between two low-pass filters; i.e., the weights are $\bar{b}_h - \underline{b}_h$ where \bar{b}_h is the filter weight at lag/lead h for the upper-cutoff filter and \underline{b}_h is the weight for the lower-cutoff filter. The constrained band-pass filter involves the requirement that the sum of its weights must be zero. Hence, the weights in the constrained

optimal band-pass filter are

$$(\bar{b}_h - \underline{b}_h) + (\bar{\theta} - \underline{\theta}), \quad (10)$$

where $\bar{\theta}$ is the adjustment coefficient associated with the upper-cutoff filter and $\underline{\theta}$ is the adjustment coefficient associated with the lower-cutoff filter. (See appendix B for additional discussion of this point.) That is, the constrained optimal K th-order band-pass filter is simply the difference between two constrained optimal K th-order low-pass filters. Throughout the remainder of the paper, we consider only band-pass filters with this zero-frequency constraint imposed. We use the notation defined above, $BP_K(p, q)$, to denote our approximation to the ideal band-pass filter that passes cycles between p and q periods.

F. The Effects of Truncation

This section explores the effect of changes in the maximum lag length, K , on the shape of the constrained low-pass and high-pass filters. If we choose an approximating moving average with maximum lag length K , implementing the filter means that we lose $2K$ observations (i.e., K leads and K lags). There is no “best” value of K ; increasing K leads to a better approximation to the ideal filter, but results in more lost observations. Thus, the researcher will have to balance these opposing factors: The best choice of K in a particular

to zero (that the frequency response is zero at the zero frequency). Accordingly, in appendix B, we study constrained approximation problems with the generic constraint $\phi = \alpha_K(0)$.

instance will depend on the length of the data period and the importance attached to obtaining an accurate approximation to the ideal filter. The next section will explore this trade-off in the context of postwar U.S. macroeconomic time series. In this section, however, we are simply concerned with describing the effect of variations in K on the shape of the approximating filters.

Figure 1 illustrates the effect of truncation on the shape of the high-pass filter that has been constrained to have unit weight at the zero frequency. The ideal filter is illustrated by the dashed line in each panel; it passes frequencies ω that correspond to cycles of length less than or equal to 32 quarters, assuming that the underlying data is measured quarterly. This figure shows that there are important effects on the shape of the approximate high-pass filter of changes in K . When $K = 4$ (so that the moving average covers only the preceding and subsequent four quarters), there is a major departure from the ideal filter. In particular, the approximate filter admits substantial components from the range of frequencies just below the cutoff frequency ($\omega = \pi/16$ or $\omega/s\pi = \frac{1}{32} = 0.03$ cycles per period). This phenomenon is conventionally called *leakage*: This term captures the notion that the filter has passed through frequencies that the filter was designed to suppress, including them with those the filter was designed to retain. The approximating filter has less than unit frequency response on the range just above the cutoff frequency, which we can similarly define as *compression*. Finally, when the index of cycles per period lies roughly between 0.14 and 0.22, then there is a frequency response of more than one-for-one, which we can define as *exacerbation*. As the value of K increases, the truncated filter more closely approximates the true filter. With $K = 8$, the problems of leakage, compression, and exacerbation have been substantially reduced relative to the $K = 4$ case. Further reductions in these departures from the exact filter are obtained with $K = 16$ and $K = 32$. These oscillatory departures of the approximating filter from the exact filter arise even when we do not impose the constraint that $\alpha(0) = 0$ and have been extensively studied in this context. They are typically referred to as the Gibbs phenomenon, after the researcher who initially stressed their importance.⁷

Figure 2 displays the frequency-response function for approximate band-pass filters. As with the approximate high-pass filters, there is substantial leakage, compression, and exacerbation for smallest values of K . The frequency responses oscillate around zero above the higher-cutoff frequency. The fact that there are some small negative weights in the frequency response in these approximate

filters means that they do not exactly display the “no phase shift” requirement that we impose on the ideal filter, but we regard these departures as minor.⁸ The deviations from the exact filter are attenuated with increases in K , so that these again appear small by $K = 12$. However, it is an empirical question whether improvement in approximating the ideal filter (by use of larger values of K) leads to important changes in a filtered time series or moments computed from it. In section III below, we explore the effects of changes in K on the behavior of filtered macroeconomic time series.

G. Why Filter in the Time Domain?

One common approach to band-pass filtering is the frequency-domain method used by Hassler et al. (1994) and Rush et al. (1997). This method works as follows. First, one takes a discrete Fourier transform of the economic data, computing the periodic components associated with a finite number of “harmonic” frequencies. Second, one “zeros out” the frequencies that lie outside of the band of interest. Third, one computes the inverse Fourier transform to get the time-domain filtered series, $\{\tilde{y}_1 \dots \tilde{y}_T\}$. We see two major drawbacks with this explicitly frequency-domain procedure, relative to our time-domain method. First, since there are likely “stochastic trends” in most economic time series, arising from unit root components, it is necessary to first detrend the series prior to taking the Fourier transform: In order to accomplish band-pass filtering, one must therefore make a choice of detrending method. Working with annual data, Hassler et al. use the Hodrick-Prescott filter with $\lambda = 10$ for this initial detrending step. Working with quarterly data, Rush et al. argue for a much larger value ($\lambda = 10,000$) in the initial detrending step so as to avoid distorting business-cycle outcomes. Second, the results of the frequency-domain method at all dates are dependent on the sample length T . Consider, for example, the “business cycle” outcome \tilde{y}_t obtained from a study of quarterly economic data in a study of length T_1 , e.g., the observation on cyclical output in 1970:2, obtained using data through 1985. When the sample length is extended to T_2 , the discrete Fourier transform of $\{y_1, y_2, \dots, y_T\}$ must be recomputed and each of its elements will change. Consequently, so too will each of the elements of the inverse Fourier transform of the

⁸ Our reasoning is as follows. The frequency-response function $\alpha(\omega)$ is an imaginary number at each frequency, which can be written in polar form as $g(\omega)e^{-i\phi(\omega)}$, where $g(\omega) = |\alpha(\omega)|$ is the gain and $\phi(\omega)$ is the phase shift. Our approximate filter’s frequency-response function is always real-valued since it is symmetric. But, to represent a negative value of the frequency response, since gain is positive, the implied phase shift must be $\pm\pi$, which makes $e^{-i\phi(\omega)} = -1$. (This phase shift is similar to that which arises when one considers the series $-y_t$: In order to represent it in terms of the standard definitions of gain and phase, it must be viewed as having a gain of 1 and a phase shift of $\pm\pi$.) We view these departures as small for two reasons. First, the negative values of the frequency response are numerically small. Second, it appears to be an artifact of the mathematical convention that gain is defined positive, rather than actually reflecting a translation of the series through time. If gain had alternatively been defined to admit negative values, then there would be no phase shift implied by a negative frequency response. (If this alternative definition had been employed, then there would also be no phase shift of $-y_t$ relative to y_t .)

⁷ A formal analysis of the Gibbs phenomenon as it derives from truncation for the unconstrained filter proceeds as follows. (See, for example, Koopmans (1974, ch. 6).) First, the truncation procedure at K leads and lags is viewed as a filter, with Fourier transform $\psi_K(\omega) = \sum_{k=-K}^K e^{-i\omega k} = \sin((2K+1)\omega/2)/\sin(\omega/2)$. Second, frequency-response function of the truncated filter is $\beta_K(\omega) = \int_{-\pi}^{\pi} \beta(\mu)\psi_K(\omega - \mu) d\mu$, using the fact that the Fourier transform of a product is the convolution of the Fourier transforms. Thus, the Gibbs phenomenon arises as a consequence of the oscillatory nature of the truncation “window” $\psi_K(\omega)$.

filtered series, i.e., the cyclical observations, $\{\tilde{y}_1, \dots, \tilde{y}_T\}$. Thus, the outcome for cyclical output in 1970:2 will necessarily be different when data is added from 1986 to 1994. This time variation violates the fifth requirement that we discussed in section I above—a requirement that we share with Prescott (1986).

III. Measuring Business Cycles

This section explores several empirical issues raised by the foregoing discussion of approximate band-pass filters. As discussed earlier, an ideal business-cycle filter is defined to be the $BP_\infty(6, 32)$ filter, and its optimal approximation is the $BP_K(6, 32)$ filter for $0 < K < \infty$. First, we describe the effect of changes in the truncation point K on moments computed from a specified data-generating process. Second, we explore the effect of variation in K on moments computed from several macroeconomic time series.

A. Effect of Variation in K on an AR(1) Process

A useful way to explore the approximation error induced by application of the approximate band-pass filter is to compute moments for a known stochastic process using both the ideal and approximate versions of our business-cycle filter, $BP_K(6, 32)$. We examine the effect of variation in K on the autocovariances of the following first-order autoregression:

$$x_t = 0.95x_{t-1} + \epsilon_t$$

with σ_ϵ set so that the variance of x_t is 100 as a convenient normalization (i.e., $\sigma_\epsilon = 100 * (1 - 0.95^2)$). Table 1 gives the autocovariances of x_t for the ideal business-cycle filter and for several approximations to this filter (i.e., several values of K).⁹ The first point to be made is that the band-pass filter (exact or approximate) substantially lowers the variance of the series (the autocovariance at lag 0), from a base of 100 to at most about 13.5. Looking next at how the variance of filtered x_t varies with the details of the approximation, we see that, when K is small (so that the moving average covers only a few observations), the approximate filter produces a filtered series whose variance is much smaller than the true or “exact” variance of 13.5. The approximation error for the filtered series’ variance becomes quite small once $K \geq 12$. These findings can be understood by recalling that the $K = 4$ approximation to the ideal filter involved both leakage and compression near the cutoff frequency. (See figure 2.) For variables possessing Granger’s (1966) typical spectral shape, such as this highly persistent AR(1) process, the effect of the compression is to filter out large components of frequencies for which there is substantial power in the original time series. As K rises and the accuracy of the approximate filter improves, this prob-

TABLE 1.—EFFECT OF K ON MOMENTS OF AN AR(1) PROCESS

K	Autocovariance at Lag:				
	0	1	2	4	8
2	0.23	0.07	−0.10	0.00	0.00
3	1.43	0.89	−0.05	−0.64	0.00
4	4.07	3.11	1.00	−2.01	0.01
6	8.45	7.23	4.09	−2.66	−1.69
8	9.14	7.91	4.75	−2.30	−2.32
12	13.08	11.78	8.43	0.79	−3.41
16	12.58	11.28	7.91	0.33	−3.59
20	12.10	10.77	7.37	−0.30	−4.42
24	12.19	10.86	7.44	−0.28	−4.60
32	13.01	11.67	8.22	0.42	−4.23
48	13.08	11.72	8.25	0.38	−4.48
60	13.00	11.64	8.15	0.26	−4.68
90	13.10	11.74	8.23	0.31	−4.73
exact	13.51	12.14	8.60	0.59	−4.74
no filter	100.00	95.00	90.25	81.45	66.34

lem becomes smaller. Interestingly, the variance computed from the approximate filter does not converge monotonically to the true variance as K rises. However, the departures from the true value are small for large values of K . A similar picture emerges for the other autocovariances: Small values of K generally produce autocovariances smaller, in absolute value, than those produced by the ideal filter. Throughout, the approximation error is small for $K \geq 12$.

B. Empirical Effects of Variation in K

This subsection explores the effect of the length of the moving average on summary statistics for several postwar U.S. time series. To provide some information about how one’s view of the macroeconomic “facts” might depend on K , we have computed a set of summary statistics for several U.S. postwar quarterly macroeconomic time series using a range of values for K . Table 2 presents statistics on standard deviations, serial correlation coefficients, and contemporaneous correlations with GNP for $K = \{4, 8, 12, 16, 20\}$. Throughout the table, moments are computed for the time period associated with the shortest filtered time series (i.e., the $K = 20$ filter), so differences in moments are not due to differences in the sample period. Summary statistics are also presented for three other filters: a centered moving average, the first-difference filter, and the Hodrick-Prescott filter—but we defer discussion of these results until section IV.

Table 2-A shows that one commonly used measure of volatility—the standard deviation—is sensitive to the choice of K . Specifically, the measured volatility of every time series studied is about half as large for the lowest value of K ($K = 4$) compared with the value generated by the largest value of K ($K = 32$). This table shows that there is little effect of increases in K on the standard deviations of the filtered time series for $K \geq 12$. These results are consistent with the results obtained above for the AR(1): Small values of K yielded low variances, while a good approximation was obtained for $K \geq 12$.

⁹ These autocovariances were not generated from Monte Carlo experiments. They are population moments and were computed by applying the approximate band-pass filter’s transfer function, $|\alpha_K(\omega)|^2$, to the spectral density of the first-order autoregression and then numerically integrating the result.

TABLE 2.—EFFECT OF FILTERING ON MOMENTS: QUARTERLY DATA, 1947:1–1997:2

A. STANDARD DEVIATIONS

Variable	K: Truncation Point for Band-Pass Filter					Moving Average	Hodrick- Prescott	First Difference
	4	8	12	16	20			
GNP	0.95	1.49	1.75	1.71	1.71	2.07	1.82	1.08
Cons: durables	2.61	4.09	4.97	4.85	4.83	6.12	5.47	3.81
Cons: nondurables	0.60	0.99	1.15	1.13	1.09	1.37	1.23	0.76
Cons: durables	0.33	0.50	0.65	0.63	0.60	0.81	0.71	0.49
Investment	2.40	4.04	5.14	4.99	5.08	6.18	5.42	2.68
Hours per person	0.24	0.37	0.39	0.39	0.38	0.45	0.41	0.28
Employment	0.67	1.15	1.43	1.40	1.38	1.69	1.47	0.68
Exports	2.42	4.05	4.97	4.90	4.94	6.31	5.54	4.43
Imports	2.50	3.91	4.60	4.48	4.42	5.71	5.22	4.05
Net exports*	6.17	11.11	16.94	16.09	14.60	21.52	18.96	10.03
Gov't purchases	1.00	2.00	3.19	3.04	2.86	4.01	3.27	1.25
GNP deflator	0.29	0.59	0.95	0.90	0.80	1.20	0.91	0.64
Inflation*	0.58	0.82	1.01	0.99	1.01	1.45	1.32	1.50

Notes: Application of these filters involves loss of data points at both ends of the sample. For consistency, the moments reported are for the truncated sample 1952:1–1992:2 (the longest period available for the $K = 20$ band-pass filter). The sample period for the hours variable is 1947:1–1996:3. Except for starred variables, natural logs were taken before filtering. See the Data Appendix for a description of data sources.

B. FIRST-ORDER AUTOCORRELATION

Variable	K: Truncation Point for Band-Pass Filter					Moving Average	Hodrick- Prescott	First Difference
	4	8	12	16	20			
GNP	0.80	0.87	0.91	0.90	0.90	0.87	0.84	0.35
CONS: DURABLES	0.79	0.87	0.92	0.91	0.91	0.81	0.77	−0.02
CONS: NONDURABLES	0.82	0.88	0.92	0.91	0.91	0.86	0.83	0.27
CONS: DURABLES	0.78	0.87	0.92	0.91	0.90	0.84	0.81	0.24
INVESTMENT	0.83	0.89	0.93	0.93	0.92	0.91	0.89	0.44
HOURS PER PERSON	0.80	0.86	0.88	0.88	0.88	0.81	0.78	0.26
EMPLOYMENT	0.84	0.89	0.93	0.92	0.92	0.92	0.91	0.70
EXPORTS	0.79	0.89	0.92	0.91	0.91	0.75	0.69	−0.19
IMPORTS	0.78	0.87	0.91	0.90	0.90	0.76	0.72	−0.09
NET EXPORTS*	0.83	0.92	0.96	0.95	0.94	0.91	0.89	0.22
GOV'T PURCHASES	0.84	0.95	0.97	0.97	0.96	0.95	0.94	0.33
GNP DEFLATOR	0.89	0.94	0.96	0.96	0.95	0.95	0.94	0.84
INFLATION*	0.64	0.87	0.89	0.89	0.89	0.49	0.40	−0.37

C. CONTEMPORANEOUS CORRELATION WITH GNP

Variable	K: Truncation Point for Band-Pass Filter					Moving Average	Hodrick- Prescott	First Difference
	4	8	12	16	20			
GNP	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
CONS: DURABLES	0.85	0.83	0.78	0.79	0.78	0.74	0.75	0.65
CONS: NONDURABLES	0.73	0.81	0.82	0.82	0.83	0.79	0.78	0.50
CONS: DURABLES	0.53	0.70	0.76	0.76	0.77	0.75	0.72	0.37
INVESTMENT	0.90	0.90	0.87	0.87	0.89	0.85	0.85	0.74
HOURS PER PERSON	0.85	0.85	0.83	0.83	0.84	0.80	0.80	0.69
EMPLOYMENT	0.81	0.81	0.83	0.83	0.82	0.83	0.82	0.72
EXPORTS	0.26	0.24	0.28	0.29	0.32	0.29	0.27	0.20
IMPORTS	0.74	0.80	0.77	0.77	0.79	0.72	0.71	0.35
NET EXPORTS*	−0.44	−0.45	−0.41	−0.41	−0.45	−0.38	−0.39	−0.18
GOV'T PURCHASES	0.20	0.12	0.18	0.17	0.11	0.19	0.17	0.24
GNP DEFLATOR	−0.39	−0.46	−0.49	−0.49	−0.51	−0.49	−0.55	−0.26
INFLATION*	−0.06	0.05	0.14	0.12	0.18	0.09	0.05	−0.17

Table 2-B presents serial correlation coefficients. As with the standard deviations, the serial correlations of the filtered time series depend on K . In particular, this measure of persistence is uniformly lower for the smallest value of K , compared with the largest. The reason, once again, can be traced to the effects of leakage and compression for small K

on the filtered time series. Since the most-persistent components of economic time series occur at the lower frequencies, the effect of compression in particular is to reduce the measured persistence of the filtered time series. As with standard deviations, the problem is most severe for $K = 4$, and there is little change for $K \geq 12$.

Table 2-C presents results for the contemporaneous correlation of various aggregates with GNP, which is one commonly used measure of the comovement of a variable with the business cycle. This table shows that there is a tendency for a variable's correlation with GNP to increase as K increases, although this is not uniformly true. As before, there is a tendency for the estimated moments not to change much for $K \geq 12$. Overall, our results suggest that summary statistics computed from the key macroeconomic time series are largely invariant to further improvements in the approximate business-cycle filter beyond $K = 12$.

C. Inspecting the Results for GNP

Figure 3 displays the results of applying five filters to the natural logarithm of gross national product.¹⁰ Throughout the four graphs, we use the band-pass business-cycle filter with $K = 12$ as our reference point: It is the dark line which is present in all of the graphs. The common sample period for these graphs is 1947.1–1997.1, but, since we use $K = 12$, we lose three years of data at each end of the plots for the band-pass and high-pass filters.

The First-Difference Filter: Panel A of figure 3 shows the quarterly growth rate of real GNP versus the band-pass filter. The first-difference filter's heavy weight on high-frequency components of the data lead to the very jagged appearance of the filtered time series. There is little correspondence between the time series produced by the first-difference and the band-pass filters.

The Hodrick-Prescott Filter: Panel B of figure 3 plots Hodrick-Prescott filtered real GNP. There is a very close correspondence between the cycles isolated by this filter and those generated by the band-pass filter, although the Hodrick-Prescott filtered series is somewhat less smooth.

The High-Pass Filter ($HP_K(32)$): Panel C of figure 3 displays a high-pass filter constructed using our procedures that isolates periodic components of 32 quarters (eight years). We have chosen the same K value for this filter as for the reference band-pass filter, so that the panel simply illustrates the effect of the smoothing of high-frequency components introduced by our band-pass filter. For GNP, the panel makes clear that this smoothing of irregular components has little effect on the overall volatility.

The Deviation from Five-Year Moving Average Filter: Finally, Panel D of figure 3 displays deviations from a centered equally weighted moving average, which is a detrending method long used by business-cycle researchers. As with the Hodrick-Prescott filter and the high-pass filter, the

correspondence with the band-pass filter is quite close, with the moving average filter being somewhat more volatile.

D. Inspecting the Results for Inflation

In figure 4, we present the results of applying the same five filters to the inflation rate. As before, the dark line in each panel is the $BP_K(6, 32)$ business-cycle filter.

The First-Difference Filter: Panel A of figure 4 shows the quarterly growth rate of inflation versus the band-pass filter. As before, the first-difference filter produces a highly volatile time series that bears little resemblance to the band-pass filter.

The Hodrick-Prescott Filter: Panel B of figure 4 plots Hodrick-Prescott filtered real GNP. In contrast to the results for GNP, there is a notable difference between the Hodrick-Prescott filter and the band-pass filter. The reason is that inflation contains important high-frequency components that are passed by the Hodrick-Prescott filter, but that are removed by the band-pass filter. GNP, by contrast, does not have important variation at high frequencies.

The High-Pass Filter ($HP_K(32)$): Panel C of figure 4 displays results for the $HP_{12}(32)$ filter. Like the Hodrick-Prescott filter, this filter passes the high-frequency components of inflation, leading to a more volatile filtered time series compared with that produced by the band-pass filter.

The Deviation from Five-Year Moving Average Filter: Finally, Panel D of figure 4 displays deviations from a moving average. As with the Hodrick-Prescott filter and the high-pass filter, the correspondence with the band-pass filter is weaker when we consider inflation compared with GNP. Once again, the reason is that high-frequency variation is much more important as a source of overall variation in inflation, compared with GNP.


IV. Comparison with Other Filters

This section compares the properties of our proposed business-cycle filter with other commonly used filters.¹¹ We evaluate each filter in terms of its ability to achieve the following characteristics that we have argued are necessary for a "good" business-cycle filter: ability to remove unit roots, absence of phase shift, and ability to isolate business-cycle frequencies without reweighting the passed frequencies. Further, since model evaluation involves comparison of model moments with moments computed from the data, it is desirable that a business-cycle filter be easily (and consistently) applied both to the data and to economic models.

¹⁰ We used excerpts from the database provided with Stock and Watson's (1999) extensive cataloging of U.S. business-cycle facts. Exact definitions of variables are contained in replication materials available from the authors.

¹¹ Our comparison is motivated, in part, by the fact that previous studies have shown that business-cycle statistics are quite sensitive to the detrending procedure. (See, for example, Baxter (1991) and Kydland and Prescott (1990).)

FIGURE 3.—THE EFFECTS OF ALTERNATIVE FILTERS ON GDP




A. Removal of Linear Trends

Although the removal of linear (or log-linear) trends historically was a standard method for separating trends from cycles, a large and growing body of evidence suggests that many macroeconomic time series contain unit root (stochastic trend) components that would not be removed by

this procedure. Primarily for this reason, this approach to detrending has fallen out of favor in empirical macroeconomic investigations. Although this procedure does not induce phase shift (nor does it reweight frequencies), the failure to remove unit root components from the data means that linear detrending is undesirable for most macroeconomic time series.

FIGURE 4.—THE EFFECTS OF ALTERNATIVE FILTERS ON INFLATION

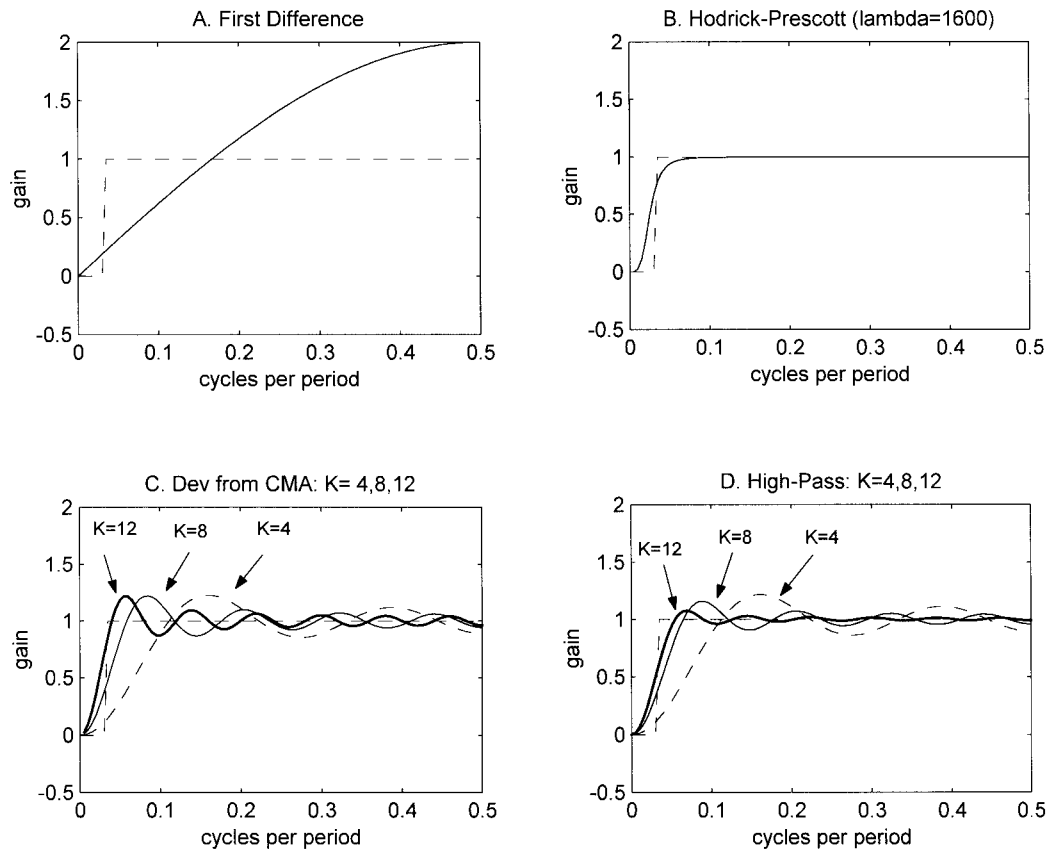


B. The First-Difference Filter

The first-difference filter extracts the cyclic component y_t^c from a time series y_t as follows: $y_t^c = (1 - L)y_t$. It is evident that this filter removes unit root components from the data; for this reason, use of the first-difference filter has been popular in recent years. However, there are several problems with this filter with respect to the criteria listed above. First,

because this filter is not symmetric, it alters timing relationships between variables (i.e., there is substantial phase shift for this filter). Second, this filter involves a dramatic reweighting of frequencies. Figure 5 panel A plots the gain function for this filter; the first-difference filter reweights strongly toward the higher frequencies, while down-weighting lower frequencies. If the goal of a business-cycle

FIGURE 5.—GAIN OF VARIOUS FILTERS RELATIVE TO HIGH-PASS FILTER (CUTOFF AT 32 QUARTERS)



filter is to isolate fluctuations in the data that occur between specific periodicities, without special emphasis on any particular frequency, the first-difference filter is a poor choice.

C. The Hodrick-Prescott Filter

Use of the business-cycle filter proposed by Hodrick and Prescott (1980) has grown dramatically in recent years, especially in investigations involving the quantitative-equilibrium approach to constructing aggregative models. The properties of this filter were previously studied by King and Rebelo (1993), and the following discussion borrows heavily from their analysis.

The infinite-sample version of the Hodrick-Prescott filter defines the cyclic component of a time series y_t as follows:

$$y_t^c = \left(\frac{\lambda(1-L)^2(1-L^{-1})^2}{1 + \lambda(1-L)^2(1-L^{-1})^2} \right) y_t \quad (11)$$

where λ is a parameter that penalizes variation in the growth component. (For quarterly data, Hodrick and Prescott recommend a value of $\lambda = 1600$.) From this equation, we see that the Hodrick-Prescott filter removes unit root components from the data. (In fact, it will remove nonstationary components that are integrated of order four or less.) Further, the

filter is symmetric so there is no phase shift. Expanding equation (11) gives the following time-domain representation of the growth component extracted by the Hodrick-Prescott filter. (See appendix A to King and Rebelo (1989) for the derivation.):

$$y_t^g = \frac{\theta_1 \theta_2}{\lambda} \left(\sum_{j=0}^{\infty} (A_1 \theta_1^j + A_2 \theta_2^j) y_{t-j} + \sum_{j=0}^{\infty} (A_1 \theta_1^j + A_2 \theta_2^j) y_{t+j} \right) \quad (12)$$

where A_1 and A_2 depend on θ_1 and θ_2 ; the coefficient $A_1 \theta_1^j + A_2 \theta_2^j$ is a real number for each j ; and A_1 and A_2 are complex conjugates.¹²

As noted by King and Rebelo, the Fourier transform of the cyclical component of the Hodrick-Prescott filter has a

¹² Equation (12) makes it clear that the Hodrick-Prescott filter is a two-sided moving average, as are several of the filters we consider. This equation also shows that the moving average is of infinite order, so that in empirical applications some approximation to this filter is required. We discuss the issue of approximation of the Hodrick-Prescott filter in section V below; the discussion here focuses on the exact Hodrick-Prescott filter.

particularly simple form:

$$\tilde{C}(\omega) = \frac{4\lambda(1 - \cos(\omega))^2}{1 + 4\lambda(1 - \cos(\omega))^2}. \quad (13)$$

Thus, the cyclical component of the Hodrick-Prescott filter places zero weight on the zero frequency ($\tilde{C}(0) = 0$), and close to unit weight on high frequencies ($\tilde{C}(\pi) = 16\lambda/(1 + 16\lambda)$). Figure 5-B plots the frequency-response function of the Hodrick-Prescott filter for $\lambda = 1600$. Visually, this filter looks remarkably like an approximate high-pass filter with cutoff frequency $\underline{\omega} = \pi/16$ or 32 cycles per period.

In terms of the objectives that we specified for our filter design problem, the Hodrick-Prescott cyclical filter has several desirable features. First, no phase shift is introduced. Second, it has trend-elimination properties: It places zero weight at the zero frequency or, equivalently, contains multiple differencing operations. Third, with $\lambda = 1600$, it approximates the high-pass filter $HP_\infty(32)$ reasonably well since its gain rises sharply from near zero to near unit in the vicinity of the cutoff frequency $\omega = \pi/16$. However, since the Hodrick-Prescott filter of equation (4.2) is an infinite order moving average, some modification is necessary in order to apply it to data. We return to discussion of this topic in section V below.

D. Moving Averages

Another widely used method of detrending economic time series is to define the growth or trend component as a two-sided or centered moving average, with the cyclic component defined in the usual way as the deviation of a particular observation from the trend line. That is, the growth or trend component is formed as

$$y_t^g = \frac{1}{2K + 1} \sum_{j=-K}^K y_{t-j}. \quad (14)$$

Thus, the cyclic component of y_t is generated as $y_t^c = a(L)y_t$ with $a_0 = 1 - 1/2K + 1$, and $a_j = a_{-j} = 1/2K + 1$ for $j = 1, 2, \dots, K$. This filter places zero weight at the zero frequency since $\sum a_k = 0$, and is symmetric. Figure 5-C plots the gain for the centered moving-average filter for several values of K .¹³

E. A High-Pass Filter

We have defined a high-pass business-cycle filter, $HP_K(32)$, as a filter that passes components of the data with periodicity less than or equal to 32 quarters. Figure 5-D plots the gain

for this filter for several values of K . As with the moving-average filter, this filter yields a good approximation to an ideal high-pass filter for sufficiently large values of K (i.e., $K \geq 12$).

F. Moment Implications

Table 2 shows how application of these alternative filters affects moments computed from several postwar U.S. time series. We focus on three sets of moments of particular interest to business-cycle analysis: volatility, persistence, and correlation with output.

Volatility. Table 2-A presents volatility statistics. As discussed earlier, the band-pass filter with $K \geq 12$ yields a very good approximation to the ideal band-pass filter. For this reason, we regard the statistics computed with the $K = 20$ band-pass filter as the best measure of business-cycle volatility, and then compare the other filters to this benchmark. Except for inflation (which we discuss separately below), a clear pattern emerges. The Hodrick-Prescott filter produces volatility statistics that exceed those of the ideal band-pass filter, although in many cases not by a large amount. The moving-average filter produces volatility statistics that are larger still, although again the changes are not dramatic. The first-difference filter, by contrast, produces volatility statistics that are smaller (in many cases, much smaller) than those produced by the band-pass filter. Having studied the gain functions of these filters, these results are easy to understand. The Hodrick-Prescott and moving-average filters are rough approximations to a high-pass filter, which means that they retain some high-frequency volatility that is removed by the band-pass filter. These macroeconomic time series do not have a great deal of power at high frequencies, so including these components leads to only small increases in the volatility of the filtered time series. The first-difference filter produces smaller measures of volatility because it removes more of the low-frequency components of the time series than the band-pass filter, while reweighting the frequencies to emphasize the higher frequencies. For all the variables studied except inflation, most of the power is at the lower frequencies.

The pattern described above is reversed for inflation: Here, the first-difference filter produces the highest measure of cyclic volatility. As discussed in section III.D. above, inflation contains sizable high-frequency components—components that are emphasized by the first-difference filter. This also explains why the moving-average and Hodrick-Prescott filters produce significantly higher volatility measures compared with the band-pass filter: The band-pass filter removes the high-frequency components, while these alternative filters do not.

Persistence. Table 2-B presents statistics on the first-order autocorrelation of filtered macroeconomic time series.

¹³ The general shape of this filter is very similar to that of the approximate high-pass filter, plotted in figure 5-D, although the “side lobes” are more exaggerated for the moving-average filter.

As before, we take the band-pass filter (for $K \geq 12$) as our benchmark. Compared with this benchmark, each of the other filters produces a lower measure of persistence. Excepting, once again, the inflation series, the differences are relatively small for the moving-average and Hodrick-Prescott filters. However, the first-difference filter produces dramatically smaller measures of persistence compared with the other filters. Once again, this is due to the fact that the first-difference filter removes more of the highly persistent, low-frequency components, and emphasizes the much-less-persistent, high-frequency components. As before, the inflation series behaves differently than the other time series, because of its important high-frequency components. (With the emphasis on these components provided by the first-difference filter, the measured persistence of inflation is actually negative.)

Correlation with GNP. Finally, Table 2-C provides statistics on the correlation between various macro variables and GNP. Once again, we find that the moving-average and Hodrick-Prescott filters produce statistics that are roughly similar to those computed using the band-pass filter. The first-difference filter produces correlations that are, in many cases, significantly smaller (in absolute value). Overall, researchers using the band-pass filter, the moving-average filter, or the Hodrick-Prescott filter on quarterly postwar U.S. time series are likely to obtain a similar impression of the nature of business cycles. However, use of the first-difference filter will yield a markedly different view of the central business cycle “facts.”

In general, the first-difference procedure produces filtered time series with lower volatility than those generated by the band-pass filters or the Hodrick-Prescott filter. This is a direct consequence of the fact that the first-difference filter downweights the lower frequencies relative to the alternative filters. For the same reason, the first-difference filter produces time series that exhibit much lower persistence than those produced by other filters (see table 2-B) and whose correlation with GNP is also much lower (table 2-C).

V. Comparing HPs

In this section, we undertake a detailed comparison of the Hodrick-Prescott filter with high-pass filters constructed using our approach. For the purposes of many users of the Hodrick-Prescott filter, we shall conclude that our high-pass filter is better in one important dimension: its ease of application to data sampled at other-than-quarterly frequencies.

A. The Quarterly HP Filters Can Be Very Close

The first observation is that our $HP_{12}(32)$ filter and the conventional Hodrick-Prescott filter give essentially similar results for quarterly GNP, thus reinforcing the idea—discussed in the previous section—that the Hodrick-Prescott filter is a reasonable approximation to the band-pass filter.

This result is suggested by comparison of panels C and D of figure 5, discussed in section III.B, above: The two series look very much like each other. In fact, the correlation of the Hodrick-Prescott cyclical component and the $HP_{12}(32)$ cyclical component is 0.994 over the common sample period.

B. The Hodrick-Prescott Filter in Finite Samples

Many individuals currently use the Hodrick-Prescott filter with $\lambda = 1600$ for defining cyclical components of quarterly economic time series. One main rationale for this, given by Prescott (1986), is that the filter is approximately a band-pass filter that passes cyclical components of periodicity greater than eight years (32 quarters).

To apply the Hodrick-Prescott cyclical filter to data, one strategy would be to truncate its weights at some fixed lag K , which would be analogous to our approximation of the ideal band-pass filter. However, in actual practice, an alternative procedure is typically used. This procedure has the apparently attractive feature that there is no loss of data from filtering. That is, for a time series y_t for $t = 1, \dots, T$, the Hodrick-Prescott procedure produces estimates of the cyclical component, y_t^c for $t = 1, \dots, T$.¹⁴

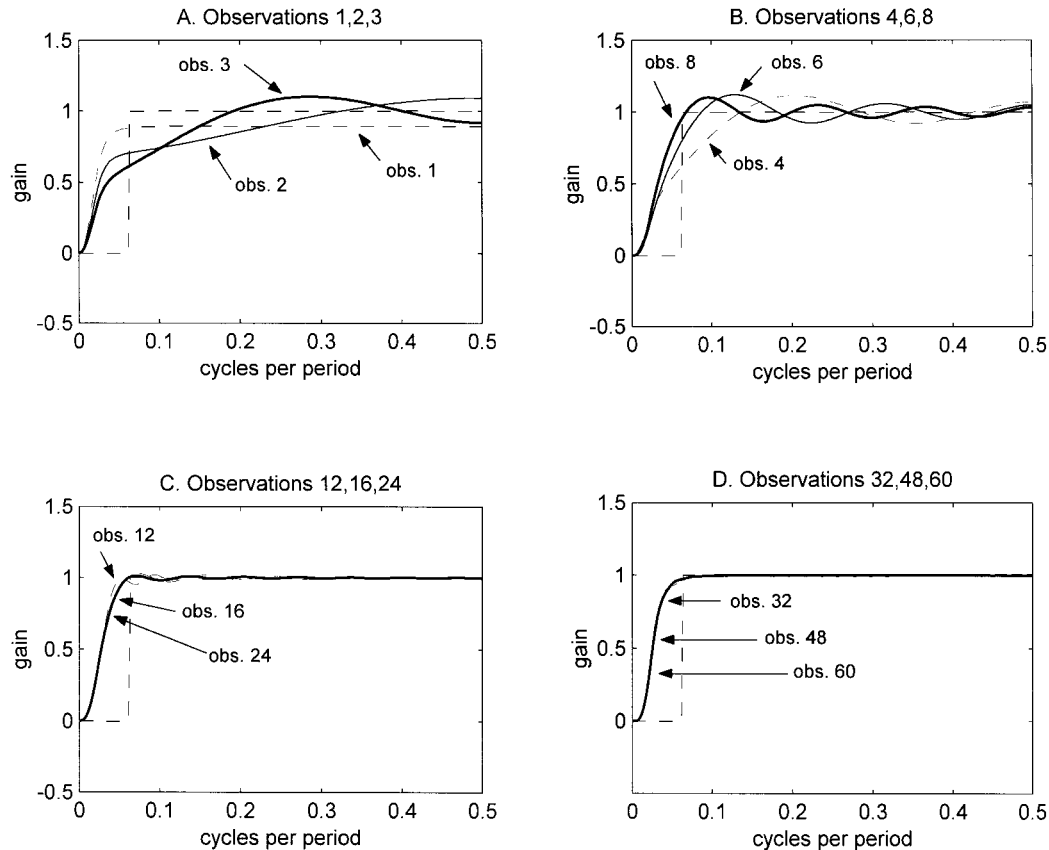
To understand this outcome, it is useful to return to the original derivation of the Hodrick-Prescott filter as the solution to a specific econometric problem, which is essentially to find the optimal estimates of trend and cycle corresponding to a particular known probability model. If we let y_t^t denote the trend component and continue to let y_t^c denote the cyclical component, this probability model is that trend and cycle are driven by independent white noises (η_t and ϵ_t respectively) and that their dynamics are $\Delta^2 \tau_t = \eta_t$ and $c_t = \epsilon_t$. If one knows the relative magnitude of σ_η^2 and σ_ϵ^2 , then it is possible to extract estimates of y_t^t and y_t^c at each date of a finite sample $t = 1, \dots, T$. Further, these estimates are simply weighted averages of the original data, so that the cyclical component at date t is

$$y_t^c = \sum_{h=1}^T d_{ht} y_h.$$

While this derivation makes the date t cyclical component a moving average of the data, the linear filter is not time invariant: The weights depend on the date t as well as the lead/lag index h . However, the algorithm that we use for

¹⁴ We implement the finite-sample Hodrick-Prescott filter as follows. First, we stack the data into a column vector Y . Second, we define a matrix Γ that links the corresponding column vector of “growth components,” Y^G , to the data: $Y = \Gamma Y^G$. Third, we compute the vector of “cyclical components” as $Y^C = Y - Y^G = (I - \Gamma^{-1})Y$. The matrix Γ is implied by the equations that link the growth components to the data. The general equation is $y_t = \lambda y_{t+2}^g - 4\lambda y_{t+1}^g + (1 + 6\lambda)y_t^g - 4\lambda y_{t-1}^g + \lambda y_{t-2}^g$, but this expression must be modified near the endpoints. For example, at the beginning of the sample, we use $y_1 = (1 + \lambda)y_1^g + (-2\lambda)y_2^g + (1 + \lambda)y_3^g$ and $y_2 = (-2\lambda)y_1^g + (1 + 5\lambda)y_2^g + (-4\lambda)y_3^g + \lambda y_4^g$, and comparable modifications must be made near the end of the sample.

FIGURE 6.—GAIN OF THE HODRICK-PRESCOTT FILTER IN FINITE SAMPLES



computing the Hodrick-Prescott filter makes it easy to recover the coefficient d_{ht} so that we can study their properties. One feature that emerges is that, for each date t , $\sum_{h=1}^T d_{ht} = 0$ so that, in this fashion, the time-varying linear filter displays trend-elimination properties at every date.

To begin our more detailed look at the time-varying filter, we compute the gain of the linear filter $d_t(L) = \sum_{h=1}^T d_{ht} L^{(h-t)}$ in figure 6 for a range of dates $t = (1, 2, 3), (4, 6, 8), (12, 16, 24), (32, 48, 60)$. These choices are motivated by the idea that we are studying a quarterly sample period of postwar size, so that there are about 180 observations, and we want to explore the effects of time variation near the endpoints and in the middle of the sample. (It is sufficient to look at the initial values because there is a symmetry property to the weights: $d_{1T} = d_{T1}$, etc.) These figures show that the d_{ht} coefficients at the beginning of the sample period are such that the $d_t(L)$ has very different properties than an exact high-pass filter: The gain functions differ sharply from each other for $t = 1, 2, 3$ and from the gain of the exact high-pass filter. (There is also phase shift near the endpoints, since $d_t(L)$ is not close to being a symmetric linear filter for t close to 1 or T .) But, as we move toward the middle of the sample period, the gain of the filter differs less sharply from one observation to the next, and the overall filter looks closer to the ideal band-pass filter.

Another perspective on the extent of time variation in the filter weights is afforded by considering the effect of $d(L)$ if

it is applied to a specific data-generating process. While it is feasible to undertake this for standard macroeconomic models, we opted for the simpler procedure of evaluating the effects of the filter on population variance of a first-order autoregression, $y_t = \rho y_{t-1} + e_t$ with $\sigma_e^2 = 1$ and $\rho = 0.95$. Table 3 gives the variance by observation with the time-varying weight version of the Hodrick-Prescott filter. (This variance should be viewed as calculated across many realizations of the time series generated by this first-order autoregressive process.) Although each observation has the same variance before filtering, time-variation in the filter applied to the process leads to different variances across observations. In fact, the change in the variance is not even monotonic, as suggested by the gain patterns in Figure 6.

This investigation thus suggests that the Hodrick-Prescott filter does not really generate as many useful estimates of the cyclical component as there are data points. Since the filter weights settle down after about the twelfth observation, it would seem natural to drop twelve observations from the beginning and end of the sample period. But, then, there would be little reason to prefer the Hodrick-Prescott filter to our high-pass filter for quarterly data.

C. HP Filters at Other Data Frequencies

Is the Hodrick-Prescott filter an adequate approximation to a high-pass filter when used with data sampled at other

TABLE 3.—EFFECT OF
HODRICK-PRESCOTT FILTER
WITH TIME-VARYING WEIGHTS

Observation	Variance
1	17.50
2	12.01
3	9.97
4	9.72
6	11.54
8	13.70
12	15.64
16	15.76
24	15.89
32	16.54
48	16.56
60	16.56
90	16.56

frequencies? The answer to this question is important to researchers concerned with international and public finance questions: Very often, the data used by these researchers are available only at the annual frequency. For our procedures, it is clear how to move between different data frequencies. For example, if we are considering results from the high-pass filter $HP_{12}(32)$ with data at the quarterly frequency, then the natural first filter to consider for annual data is $HP_3(8)$: We isolate the same frequencies (periodicities of eight years and higher), and we lose the same number of years of data at the ends of the sample.

However, it is much less clear how to proceed with the Hodrick-Prescott method. The difficulty is that the Hodrick-Prescott filter requires the researcher to specify the “smoothing parameter,” λ . For quarterly data, we found that $\lambda = 1600$ produces a reasonable approximation to a high-pass filter. For annual data, current empirical practice is to use $\lambda = 400$ or $\lambda = 100$. (For example, Backus and Kehoe (1992) use $\lambda = 100$ in their study of international business cycles.) To investigate whether these values of λ yield a good approximation to a band-pass filter for annual data, we applied our $BP_3(2, 8)$ filter and the HP filter for several values of λ to U.S. annual GNP.¹⁵ The commonly used values of $\lambda = 400$ and $\lambda = 100$ did not produce a filtered time series for GNP that closely resembled that produced by the band-pass filter.¹⁶ However, setting $\lambda = 10$ produced a much better correspondence between the Hodrick-Prescott and band-pass filters. Figure 7 plots the gain for the Hodrick-Prescott filter for the three values of λ against the ideal filter. This figure reveals why $\lambda = 100$ and $\lambda = 400$ produce such different pictures for filtered GNP compared with the optimal approximate band-pass filter: For these values of λ , the Hodrick-Prescott filter is a poor approximation to the ideal filter. In particular, these filters contain a great deal of leakage from low frequencies. (That is, the $\lambda = 100$ and $\lambda = 400$ filters pass through nearly all of the

components of the data with cycles between nine and sixteen years—components that most researchers would not identify as business-cycle components.) The approximation to the ideal band-pass filter is significantly better for $\lambda = 10$. However, even the $\lambda = 10$ filter contains significant leakage as well as significant compression.¹⁷

The foregoing discussion concerned the properties of the exact Hodrick-Prescott filter. In practice, however, a finite-moving-average approximation to this exact filter must be used. Looking at a figure similar to figure 6 but designed for annual data, we found that the finite-sample version of the filter produces serious departures from the ideal filter for the first three observations, but improves dramatically after the fourth observation. We thus recommend dropping at least three data points from each end of the sample when using the Hodrick-Prescott filter on annual data, even if one chooses $\lambda = 10$, which is the same number of data points dropped by our business-cycle filter.

VI. Summary and Conclusions

This paper develops a set of approximate band-pass filters designed for use in a wide range of economic applications. The empirical focus of the paper is on isolating cyclic fluctuations in economic time series, defined as cycles in the data between specified frequency bands. We make detailed comparisons of our band-pass business-cycle filter with other commonly used filters, and evaluate these alternative filters in terms of their ability to isolate business-cycle fluctuations in the data. We found that linear detrending and first-differencing the data are not desirable business-cycle filters. On the other hand, moving-average analysis and Hodrick-Prescott filtering can, in some cases, produce reasonable approximations to an ideal business-cycle filter. However, the optimal approximate band-pass filter that we develop in this paper is more flexible and easier to implement than these filters and produces a better approximation to the ideal filter. While the main motivation for and focus of our investigation is on construction of a business-cycle filter, the results should be of more-general interest since the defining periodicities may be readily specified by a researcher and applied to data at any observation frequency. Based on the results of this paper, we recommend three filters for use with quarterly and annual macroeconomic data.


For quarterly macroeconomic data, we recommend the Burns-Mitchell band-pass filter, which admits frequency components between 6 and 32 quarters, with $K = 12$. This filter removes low-frequency trend variation and smooths high-frequency irregular variation, while retaining the major features of business cycles. Some macroeconomists, particularly those who have extensively used the Hodrick-Prescott filter, may prefer to employ the high-pass filter, which

¹⁵ Since the shortest detectable cycle in a time series is one that lasts two periods, the annual business-cycle filter passes components with cycle length between two and eight years. Note that, in this case, the band-pass filter is equivalent to a high-pass filter.

¹⁶ This finding was not altered by increasing the K parameter from 3 to 6.

¹⁷ Hassler et al. (1994) also argue that $\lambda = 10$ is the appropriate value for the smoothing parameter when applying the Hodrick-Prescott filter to annual data.

FIGURE 7.—ALTERNATIVE ANNUAL HODRICK-PRESCOTT FILTERS



admits frequency components between 2 and 32 quarters with $K = 12$. Essentially, this filter removes the trend variation without removing the higher-frequency irregular variation in the series. Relative to the Hodrick-Prescott method, this filter does involve dropping three years of data at the beginning and end of the sample; we have seen, however, that this loss is more apparent than real because the weights in the Hodrick-Prescott filter are rapidly changing near the ends of the sample, resulting in substantial distortions of these cyclical observations. The filter weights are provided in the first two columns of table 4.

For annual macroeconomic data, band-pass and high-pass business-cycle filters are equivalent. We accordingly recommend a single filter that admits periodic components be-

tween two and eight years, with $K = 3$. The filter weights are given in the last column of table 4.

We have applied the filters constructed in this paper in various research contexts, which provides an additional demonstration of their flexibility and usefulness. For example, Baxter (1994) uses the methods of this paper to study the relationship between real exchange-rate differentials and real interest rates at low frequencies (trend components), medium frequencies (business-cycle components) and high frequencies (irregular components). She concludes that prior studies have missed interesting relationships between these variables because a concern for producing stationary data led researchers to use the first-difference filter. This procedure emphasized irregular (high-frequency) components where little relationship exists at the expense of the business-cycle components where a striking, positive relationship emerges. In another application, King and Watson (1994) show that the “Phillips correlations” (defined as a negative correlation of inflation and unemployment) appear strong at the business-cycle frequencies even though they are hard to see in the original inflation and unemployment time series. This latter investigation uses monthly data and thus defines the business-cycle periodicities as 18 months to 96 months. It thus highlights one important strength of our approach: It is easy to alter the filter construction when the sampling frequency changes.

In conclusion, the primary goal of this paper was to “build a better mousetrap”—that is, to develop an approach to filtering economic time series that is fast, flexible, and easy

TABLE 4.—MOVING AVERAGE WEIGHTS FOR BUSINESS-CYCLE FILTERS

Lag	BP(6,32)	BP(2,32)	BP(2,8)
0	0.2777	0.9425	0.7741
1	0.2204	−0.0571	−0.2010
2	0.0838	−0.0559	−0.1351
3	−0.0521	−0.0539	−0.0510
4	−0.1184	−0.0513	
5	−0.1012	−0.0479	
6	−0.0422	−0.0440	
7	0.0016	−0.0396	
8	0.0015	−0.0348	
9	−0.0279	−0.0297	
10	−0.0501	−0.0244	
11	−0.0423	−0.0190	
12	−0.0119	−0.0137	

to implement. Our goal in this undertaking is to encourage empirical researchers to adopt a common approach to filtering, which will greatly aid in replication and comparison of results across researchers.

REFERENCES

- Backus, D., and P. Kehoe, "International Evidence on the Historical Properties of Business Cycles," *American Economic Review* 82 (Sep., 1992), 864–888.
- Baxter, M., "Business Cycles, Stylized Facts, and the Exchange Rate Regime: Evidence from the United States," *Journal of International Money and Finance* 10 (Jan., 1991), 71–88.
- , "Real Exchange Rates and Real Interest Differentials: Have We Missed the Business-Cycle Relationship?" *Journal of Monetary Economics* 33 (Feb., 1994), 5–37.
- Bry, G., and C. Boschan, *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs* (New York: National Bureau of Economic Research, 1981).
- Burns, A. M., and W. C. Mitchell, *Measuring Business Cycles* (New York: National Bureau of Economic Research, 1946).
- Craddock, J. M., "An Analysis of the Slower Temper Variations at Kew Observatory by Means of Mutually Exclusive Bandpass Filters," *Journal of the Royal Statistical Society*, 120 (1957), 387–397.
- Engle, R., and C. Granger, "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica* 55 (March, 1987), 251–276.
- Englund, P., T. Persson, and L. E. O. Svensson, "Swedish Business Cycles: 1861–1988," *Journal of Monetary Economics* 30 (Dec., 1992), 343–371.
- Granger, C., "The Typical Spectral Shape of An Economic Variable," *Econometrica* 34 (Jan., 1966), 150–161.
- Granger, C. W. J., and M., *Spectral Analysis of Economic Time Series* (Princeton, N.J.: Princeton University Press, 1964).
- Harvey, A. C., *Time Series Models* (New York: John Wiley and Sons, 1981).
- Hassler, J., P. Lundvik, T. Persson, and P. Soderlind, "The Swedish Business Cycle: Stylized Facts over 130 Years," in V. Bergstrom, A. Vredin (eds.), *Measuring and interpreting business cycles*. (Oxford and New York: Oxford University Press, Clarendon Press, 1994), 9–108.
- Hodrick, R. J., and E. C. Prescott, "Post-war U.S. Business Cycles: An Empirical Investigation," working paper, Carnegie-Mellon University (1980).
- King, R. G., and S. T. Rebelo, "Low Frequency Filtering and Real Business Cycles," Rochester Center for Economic Research working paper No. 205, University of Rochester (Oct., 1989).
- , "Low Frequency Filtering and Real Business Cycles," *Journal of Economic Dynamics and Control* 17 (Jan., 1993), 207–231.
- King, R. G., and C. I. Plosser, "Real Business Cycles and the Test of the Adelmans," *Journal of Monetary Economics* 33 (April, 1994), 405–438.
- King, R. G., and M. W. Watson, "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," Carnegie Rochester Conference Series on Public Policy 41 (Fall, 1994), 157–219.
- Koopmans, L., *The Spectral Analysis of Time Series* (New York: Academic Press, 1974).
- Kydland, F., and E. C. Prescott, "Real Facts and A Monetary Myth," Federal Reserve Bank of Minneapolis, *Quarterly Review* 14 (Spring, 1990), 3–18.
- Prescott, E. C., "Theory Ahead of Business Cycle Measurement," *Carnegie-Rochester Conference Series on Public Policy* 25 (Fall, 1986), 11–66.
- Stock, J. H., and M. W. Watson, "Business Cycle Fluctuations in U.S. Macroeconomic Time Series, in J. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Amsterdam: Elsevier Science Publishers, 1999).
- Rush, M., Y. Li, and L. Zhu, "Filtering Methodology and Fit in Dynamic Business Cycle Models," University of Florida working paper (1997).
- Watson, M. W., "Business Cycle Duration and Postwar Stabilization of the U.S. Economy," *American Economic Review* 84 (March, 1994), 24–46.

APPENDIX

A. Trend-Elimination Properties

In this appendix, we demonstrate that symmetric moving-average filters can render stationary economic time series with deterministic and stochastic trends. In particular, we consider the filter $a(L) = \sum_{k=-K}^K a_k L^k$ on which we impose two conditions: that the coefficients sum to zero, $a(1) = \sum_{k=-K}^K a_k = 0$; and that the filter is symmetric, $a_k = a_{-k}$. Using these conditions, we rewrite the moving average as

$$a(L) = \sum_{k=-K}^K a_k L^k = \sum_{k=-K}^K a_k L^k - a_k = \sum_{k=1}^K a_k (L^k + L^{-k} - 2),$$

where the first equality follows from the sum of coefficients requirement and the second follows from the symmetry assumption. The individual terms $(L^k + L^{-k} - 2)$ can be rewritten as $-(1 - L^k)(1 - L^{-k})$. Using the additional fact that $(1 - L^k) = (1 - L)[1 + L + L^2 + \dots + L^{k-1}]$, a little bit of algebra demonstrates that $[1 + L + L^2 + \dots + L^{k-1}][1 + L^{-1} + \dots + L^{-(k-1)}]$ is equal to $\sum_{h=-(k-1)}^{(k-1)} (k - |h|) L^h$. Hence, we conclude that

$$\begin{aligned} a(L) &= - \sum_{k=1}^K a_k [(1 - L^k)(1 - L^{-k})] \\ &= -(1 - L)(1 - L^{-1}) \psi_K(L), \end{aligned}$$

with $\psi_K(L)$ being a symmetric moving average with $K - 1$ leads and lags, which is defined by $\psi_K(L) = [\sum_{k=1}^K a_k \sum_{h=-(k-1)}^{(k-1)} (k - |h|) L^h]$. Since $\psi_K(L)$ is a finite-term moving average, it does not alter the stationarity properties of any series to which it is applied.

We have shown that any symmetric moving filter $a(L)$ whose weights sum to zero contains a backward difference $(1 - L)$ and a forward difference $(1 - L^{-1})$. Consequently, $a(L)$ has the ability to render stationary I(2) stochastic processes and quadratic deterministic trends.

APPENDIX

B. Optimal Approximation

In this appendix, we consider the optimal approximation of an ideal symmetric linear filter by a K th-order symmetric moving average. We pose the filter design problem in the frequency domain.

Preliminaries

Consider a filter $g(L) = \sum_{h=-\infty}^{\infty} g_h L^h$ with square-summable weights. The filter and its frequency-response function $\gamma(\omega)$ are a Fourier transform pair. Operationally, this means that the frequency-response function can be obtained from the filter weights by the Fourier sum,

$$\gamma(\omega) = \sum_{h=-\infty}^{\infty} g_h e^{-i\omega h}.$$

The filter weights can be obtained from the frequency-response function by the Fourier integral,

$$g_h = \frac{1}{2\pi} \int_{-\pi}^{\pi} \gamma(\omega) e^{i\omega h} d\omega.$$

To accomplish this integration in particular contexts, we employ the facts that

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-i\omega(j-k)} d\omega = 1 \quad \text{for } j = k$$

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-i\omega(j-k)} d\omega = 0 \quad \text{for } j \neq k.$$

(For example, these facts imply that $g_h = (1/2\pi) \int_{-\pi}^{\pi} \gamma(\omega) e^{i\omega h} d\omega = (1/2\pi) \int_{-\pi}^{\pi} [\sum_{h=-\infty}^{\infty} g_h e^{-i\omega h}] e^{i\omega h} d\omega = g_h$, which is a “reality check” of sorts.) We use these facts and the Fourier integral repeatedly in our analysis below, given that we design the optimal filter in the frequency domain and must derive the filter weights.

Application to Deriving Weights for the Ideal Low-Pass Filter

The Fourier integral of the ideal low-pass filter $\beta(\omega)$ implies that the filter coefficients satisfy

$$b_h = \frac{1}{2\pi} \int_{-\pi}^{\pi} \beta(\omega) e^{i\omega h} d\omega = \frac{1}{2\pi} \int_{-\bar{\omega}}^{\bar{\omega}} e^{i\omega h} d\omega,$$

where the second line derives from the fact that $\beta(\omega) = 1$ for $|\omega| \leq \bar{\omega}$ and $\beta(\omega) = 0$ for $|\omega| > \bar{\omega}$. Hence, it follows that

$$b_0 = \frac{1}{2\pi} \int_{-\bar{\omega}}^{\bar{\omega}} d\omega = \frac{\bar{\omega}}{\pi}$$

$$b_h = \frac{1}{2\pi} \int_{-\pi}^{\pi} \beta(\omega) e^{i\omega h} d\omega = \frac{1}{2\pi} \left[\frac{1}{ih} e^{i\omega h} \right]_{-\bar{\omega}}^{\bar{\omega}} = \frac{1}{\pi h} \sin(\bar{\omega} h)$$

where the last equality follows from $2i \sin(x) = e^{ix} - e^{-ix}$.

The Filter Design Problem in the Frequency Domain

The problem is to minimize $Q = 1/2\pi \int_{-\pi}^{\pi} |\delta(\omega)|^2 d\omega$, with $\delta(\omega)$ being the discrepancy between the exact and approximating filters at frequency ω , $\delta(\omega) = \beta(\omega) - \alpha(\omega)$. Some versions of the problem discussed in the text require that the approximating filter take on a specified value at the zero frequency, which we represent as $\alpha(0) = \phi$. (Equivalently, since $e^0 = 1$, this restriction is $\sum_{k=-K}^K a_k = \phi$.) To solve this as a constrained-maximization problem, we form the Lagrangian, $\mathcal{L} = -Q + \lambda[\alpha(0) - \phi]$, which may be expressed alternatively as

$$L = -\frac{1}{2\pi} \int_{-\pi}^{\pi} \left[\beta(\omega) - \sum_{k=-K}^K a_k e^{-i\omega k} \right] \left[\beta(\omega) - \sum_{k=-K}^K a_k e^{-i\omega k} \right]' d\omega + \lambda \left[\sum_{k=-K}^K a_k - \phi \right].$$

The first-order conditions are

$$\begin{aligned} a_j : 0 &= \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-i\omega j} \left[\beta(\omega) - \sum_{k=-K}^K a_k e^{-i\omega k} \right]' d\omega \\ &+ \frac{1}{2\pi} \int_{-\pi}^{\pi} \left[\beta(\omega) - \sum_{k=-K}^K a_k e^{-i\omega k} \right] e^{i\omega j} d\omega + \lambda \\ \lambda : 0 &= \sum_{k=-K}^K a_k - \phi \end{aligned}$$

where $-K \leq j \leq K$.

Restrictions on the Filter Weights From the First-Order Conditions

Repeatedly using the facts that $(1/2\pi) \int_{-\pi}^{\pi} e^{-i\omega(j-k)} d\omega = 1$ for $j = k$ and $(1/2\pi) \int_{-\pi}^{\pi} e^{-i\omega(j-k)} d\omega = 0$ for $j \neq k$, the $2K + 1$ first-order conditions with respect to a_j can be expressed as

$$0 = 2(b_j - a_j) + \lambda.$$

(For an example of this process, the term $(1/2\pi) \int_{-\pi}^{\pi} [\sum_{k=-K}^K a_k e^{-i\omega k}] e^{i\omega j} d\omega$ is equal to $(1/2\pi) \int_{-\pi}^{\pi} a_j d\omega = a_j$.)

Thus, if there is no constraint on $\alpha(0)$ so that $\lambda = 0$, then it follows that the optimal approximate filter simply involves truncation of the ideal filter's weights.

If there is a constraint on $\alpha(0)$, then λ must be chosen so that the constraint is satisfied. For this purpose, it is useful to write the FOCs as $a_h = b_h + \theta$, where $\theta = \lambda/2$. Then, requiring that $\alpha(0) = \sum_{h=-K}^K a_h = \phi$, we find that the required adjustment is

$$\theta = \frac{\phi - \sum_{h=-K}^K b_h}{2K + 1}.$$

Conclusions and Extensions

We have derived the general result discussed in the text. Construction of the optimal approximating filter contains two steps: truncation of the ideal filter's weights and addition of the correction term θ . Further, the form of this correction process makes clear the origins of some of the observations made in the main text which are not explicitly derived here. For example, the same logic implies that the constrained K th-order approximate band-pass filter is the difference between two constrained K th-order approximate low-pass filters. Since the ideal band-pass filter weights are simply differences between the weights of two low-pass filters, $\bar{b}_h - \underline{b}_h$ it follows that the weights for an optimal truncated band-pass filter are $(\bar{b}_h - \underline{b}_h) - [\sum_{h=-K}^K (\bar{b}_h - \underline{b}_h)]/[2K + 1]$. As this may be rearranged as $[\bar{b}_h + [1 - \sum_{h=-K}^K \bar{b}_h]/[2K + 1]] - [\underline{b}_h + [1 - \sum_{h=-K}^K \underline{b}_h]/[2K + 1]]$, it follows that the weights of the optimal, constrained approximate band-pass filter are simply the difference in the weights of the two constrained K th-order low-pass filters.