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

Marianne Baxter
Robert G. King

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

ABSTRACT

This paper develops a set of approximate band-pass filters designed for use in a wide range of economic applications. In particular, we design and implement a specific band-pass filter which isolates business-cycle fluctuations in macroeconomic time series. This filter was designed to isolate fluctuations in the data which persist for periods of two through eight years. This filter also "detrends" the data, in the sense that it will render stationary time series that are integrated of order two or less, or that contain deterministic time trends. We apply our filter to several of the key macroeconomic time series, and describe the picture of the U.S. postwar business cycle that emerges from our analysis. We also provide detailed comparisons with several alternative detrending methods.

Marianne Baxter
Department of Economics
Rouss Hall
University of Virginia
Charlottesville, VA 22903
and NBER

Robert G. King
Department of Economics
Rouss Hall
University of Virginia
Charlottesville, VA 22903
and NBER

1 Introduction

The study of business cycles necessarily begins with the *measurement* of business cycles. The seminal contribution of Burns and Mitchell [1946] was influential 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 did Burns and Mitchell 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 macroeconomists 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 only requires that the detrending procedure produce a stationary business-cycle component, and does not otherwise specify the statistical characteristics of business cycles. Examples of techniques in common use are (i) application of two-sided moving averages; (ii) first-differencing; (iii) removal of linear or quadratic time trends; and (iv) application of the Hodrick-Prescott [1980] filter. Many recent studies using 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 [1946] viewed as central: the *definition* of a business cycle. In this paper, we develop methods for measuring business cycles which require that the researcher begin by specifying characteristics of these cyclical components. Our procedures then isolate business cycle components by simply applying moving averages to macroeconomic data. Technically, we develop approximate band-pass filters that are constrained to produce stationary outcomes when applied to growing time series.[2]

---

[1] However, it is possible to implement a judgment-free version of the Burns and 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] In recent work, Englund, Persson, and Svensson [1992] and Hassler, Lundvik, Persson, and Soderlind [1992] 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. Canova [1993] also uses high-pass and band-pass filters in the frequency domain in his extensive analysis of detrending and business-cycle facts. We discuss these methods in more detail later in the paper.

For the empirical applications in this paper, we adopt the definition of the business cycle suggested by the procedures and findings of NBER researchers like Burns and Mitchell [1946]. We apply our method to a several quarterly post-war U.S. time series. 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 thirty two quarters (eight years). We adopt these limits as our definition of the business cycle.

Specifying 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 which passes through components of the time series with fluctuations between six and thirty-two quarters, while removing components at higher and lower frequencies. However, the resulting moving average is of infinite order, so an approximation to this filter is necessary for it to be applicable to finite time series. 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.[3]

In approaching this problem of filter design, we require that our method meet six objectives.[4] 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 specific loss function for discrepancies between the exact and approximate filter. Fourth, we require that the application of an approximate band-pass result in a stationary time series even when applied to trending data. Given recent empirical which suggests the presence of stochastic trends in economic time series, we design our filters so that they will make a 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 zero at the zero frequency). This requirement also means that our band-pass filters will eliminate quadratic 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. 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

---

[3]Like many of the *ad hoc* filters listed above, our approximate filters are moving averages which can readily be applied to time series. However, our filters have the important advantage that the researcher can specify the objective of the data transformation explicitly, which in business-cycle research is presumably to isolate the component of a time series with particular periodicities.

[4]These requirements are very similar to those that Prescott [1986] discusses in justifying use of the Hodrick-Prescott [1980] filter.

3

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 about the tradeoffs involved. We recommend that researchers use moving averages based on six years of data for both quarterly and annual time series.

The organization of the paper is as follows. Section 2 describes the construction of approximate band-pass filters. In section 3, we define our business cycle filter and apply it to post-war 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 post-war 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 4, we contrast our business cycle filter to the results of other commonly used procedures. In section 5, 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 6 concludes the paper with a brief review of the goals and findings of the paper. Three appendices provide information on the stochastic and deterministic trend reduction implications of a class of moving average filters that contains our approximate filters and many others; details of the derivation of the optimal approximate band-pass filter; and printouts of the MATLAB programs used to implement the filters developed in this paper.

2 Band-pass filters for economic time series

This section describes the construction of moving averages that isolate the periodic components of an economic time series which lie in a specific band of frequencies. 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, defined as a linear filter which 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 (for a recent presentation, see Harvey [1981, chapter 3]). If the time series y_t is stationary, then we can use frequency domain methods to study the implications of applying moving averages to y_t. In this paper, we employ 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, readers who are simply

4

interested in the practical results of our filtering methods may skip ahead to section 3.

2.1 Applying moving averages to time series

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

$$y_t^* = \sum a_k y_{t-k} \ . \tag{1}$$

For convenience, 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 $a_k = a_{-k}$ for $k = 1, ...K$.

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., $\sum_{k=-K}^{K} a_h = 0$, then we show in appendix A that it has trend reduction properties. That is, if the weights sum to zero, we can always factor $a(L)$ as:

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

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 which 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]).

The Cramer representation of the stationary time series y_t is:

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

That is, the time series can be expressed as the integral of random periodic components, the $\xi(\omega)$, which 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

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

where $\alpha(\omega) = \sum_{h=-K}^{K} a_h e^{-i\omega h}$ is the frequency response function of the linear filter. That is, $\alpha(\omega)$ 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:

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

5

where $|\alpha(\omega)|^2$ is the squared gain or transfer function of the linear filter at frequency ω and $f_y(\omega) = var(\xi(\omega))$ is the spectral density of the series y at frequency ω. The squared gain thus indicates the extent to which a moving average raises or lowers the contribution to variance in the filtered series from the level in the original series.

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 reduction 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 $\alpha(0) = 0$.

2.2 The low-pass filter

A basic building block in filter design is the low-pass filter, by which we mean a filter which retains only slow-moving components of the data. An ideal low-pass filter, which passes only frequencies $-\underline{\omega} \leq \omega \leq \underline{\omega}$, is illustrated in Panel A of Figure 1.[5] The ideal low-pass filter we will study thus has a frequency response function given by $\beta(\omega) = 1$ for $|\omega| \leq \underline{\omega}$, and $\beta(\omega) = 0$ for $|\omega| > \underline{\omega}$. Notice that it is symmetric, $\beta(\omega) = \beta(-\omega)$.

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:

$$b_h = \int_{-\pi}^{\pi} \beta(\omega)e^{i\omega h}d\omega \tag{6}$$

Evaluating the integral above (see Appendix B for the details), the filter weights b_h for the ideal filter are $b_0 = \underline{\omega}/\pi$, and $b_h = \sin(h\underline{\omega})/h\pi$ for $h = 1,2,\dots$. 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 $\alpha_K(\omega) = \sum_{h=-K}^{K} a_h e^{-i\omega h}$.

---

[5] In this Figure, as in others below, we measure frequency ω as a fraction of π, so that the horizontal axis ranges from -1 to 1. For the figures, this means that periodicity of the frequency component is simply $p = 2/\omega$, so that the most rapid oscillations shown in Figure 1 have period two. The "cutoff frequency" for the low 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, $\underline{\omega} = 2/32 = 1/16 \cong .07$. However, for the analytical results below, we use the more conventional definition that the frequency ω has as its domain the interval $-\pi \leq \omega \leq \pi$.

6

2.3 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:

$$Q = \int_{-\pi}^{\pi} |\delta(\omega)|^2 \, d\omega, \tag{7}$$

where $\delta(\omega) \equiv \beta(\omega) - \alpha_K(\omega)$ is the discrepancy arising from approximation at frequency ω. This loss function thus 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 a_h at lag K. This result reflects the fact that each of the truncated terms in a symmetric linear filter is orthogonal to the included terms. Thus the optimal approximate low-pass filter sets $a_h = b_h$ for $h = 0, 1, ..., K$, and $a_h = 0$ for $h \geq K + 1$, where the weights b_h are those given in section 2.2 above.[6]

2.4 Construction of high-pass and band-pass filters

High-pass and band-pass filters are easily constructed from low-pass filters. Before defining these additional filters, we establish some notation which we use throughout the rest of the paper. Since it is more natural for us to think in terms of periodicity of cycles than frequencies, we let $LP_K(p)$ denote the approximate low-pass filter which is truncated at lag K and which 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, as illustrated in panel B of Figure 1. If the weights of the low-pass filter in Figure 1 panel A are b_h for $h = 0$ and $h = \pm 1, 2, ...$, then the weights of the high-pass filter are $1 - b_0$ at $h = 0$ and $-b_h$ at $h = \pm 1, 2, ...$ Correspondingly, the optimal approximate high-pass filter, $HP_K(p)$ is simply constructed by truncating the weights of $HP_\infty(p) = 1 - LP_K(p)$.[7]

The ideal band-pass filter passes only frequencies in the ranges $\underline{\omega} \leq |\omega| \leq \overline{\omega}$. It is therefore constructed from the two low-pass filters with cutoff frequencies $\underline{\omega}$ and $\overline{\omega}$: we denote the frequency response of these filters as $\overline{\beta}(\omega)$ and $\underline{\beta}(\omega)$. Then, to get the desired frequency response, we form the band-pass filter's frequency response as

---

[6] A classic reference on the approximation of linear filters using Kth order linear filters is Koopmans [1974].

[7] This is implied by the result discussed in section 2.3: that approximation of the ideal low-pass filter simply involves truncation of the ideal filter's weights at lag K.

7

$\bar{\beta}(\omega) - \underline{\beta}(\omega)$ 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 $\bar{b}_h - \underline{b}_h$. Panel C of Figure 1 plots an ideal band-pass filter which passes through cycles of length between 6 and 32 quarters, which corresponds to the Burns and Mitchell [1946] 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 approximation band-pass filter which 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 1-C, $p = 6$ and $q = 32$). We construct $BP_K(p,q)$ by truncating the ideal band-pass filter.

2.5 Constraints on specific points

The minimization problem described above may be reformulated to recognize that certain points are of particular concern to the researcher. In our context, we want to design a low-pass filter that places unit weight at the zero frequency ($\alpha_K(\omega) = 1$ at $\omega = 0$). If we construct a low-pass filter 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.

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 C, 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,

$$a_h = b_h + \theta, \tag{8}$$

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 place unit weight at the zero frequency results in a relatively simply 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 which 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 cut-off 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 adjusted as follows:

$$(\overline{b}_h - \underline{b}_h) + (\overline{\theta} - \underline{\theta}) \tag{9}$$

where $\overline{\theta}$ is the adjustment coefficient associated with the upper-cutoff filter and $\underline{\theta}$ is the adjustment coefficient associated with the lower cut-off filter (see Appendix C for additional discussion of this point). That is, the constrained optimal Kth order band-pass filter is simply the difference between two constrained optimal Kth 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 which passes cycles between p and q periods.

2.6   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, so that the best choice of K in a particular instance will depend on the length of the data period, and the necessity to obtain a good 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 2 illustrates the effect of truncation on the shape of the low-pass filter which has been constrained to have unit weight at the zero frequency. The ideal filter, illustrated by the dotted line in each panel, passes frequencies ω which corresponds to cycles in the quarterly data of length greater than or equal to 32 quarters. This Figure shows that there are important effects on the shape of the approximate low-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 above the cutoff frequency $\underline{\omega} = \pi/16$. 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. Correspondingly, the approximating filter has less than unit frequency response on the range $|\omega| \leq \pi/16$, which we define as "compression." As the value of K increases, the truncated filter more closely approximates the true filter. With $K = 8$, the problems of leakage and compression have been substantially reduced relative to the $K = 4$ case. Further reductions in leakage and compression are obtained with $K = 16$ and $K = 32$.

9

Figure 3 displays the frequency response function for approximate band-pass filters. As with the approximate low-pass filters, there is substantial "leakage" and "compression" for small values of K. However, it is an empirical question whether improvement in approximating the ideal filter (by use of larger values of K) lead to important changes in moments computed from the filtered time series. In the next section we explore the effects of changes in K on the behavior of filtered macroeconomic time series.

2.7 Why filter in the time domain?

One common approach to band-pass filtering is the frequency domain method used by Hassler, et al. [1992], Canova [1993], and Li, et al. [1994]. 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, $\{\tilde{y}_1, , ...\tilde{y}_T\}$. We see two major drawbacks with this explicitly frequency domain procedure, relative to our time domain method. First, since there are likely to be stochastic trends (unit root components) in most economic time series, it is necessary to first detrend the series prior to taking the Fourier transform. That is: in order to accomplish band-pass filtering, one must first choose a detrending method. Working with annual data, Hassler et. al. [1992] use the Hodrick-Prescott filter with $\lambda = 10$ for this initial detrending step. Working with quarterly data, Li, et. al. [1994] 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" component of a time series at a particular date t, denoted $\tilde{y}_t$, obtained from a study of economic data in a study of length T_1. When the sample length is extended to T_2, the discrete Fourier transform of $\{y_1, y_2, ...y_T\}$ 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 filtered series, i.e., the cyclical observations, $\{\tilde{y}_1, , ...\tilde{y}_T\}$. Thus, the cyclic component of output at a particular date will change when the sample period changes. This time variation violates the fifth requirement that we discussed in section 1 above, which is also one that we share with Prescott [1986].

3 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

10

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.

3.1 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 filters. We examine the effect of variation in K on the autocovariances of the following first-order autoregression:

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

with $\sigma_\varepsilon = 1$. 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.[8] Looking first at the variance of x_t, (the autocovariance at lag 0), we see that when K is small, so that the moving average covers only a few observations, the approximate filter produces a series whose variance is much smaller than the true or "exact" variance of 1.38. The approximation error for the filtered variance becomes quite small once $K \geq 12$. This phenomenon 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 3). 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 problem 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 $K \geq 12$.

3.2 Empirical effects of variation in K

This sub-section explores the effect of the length of the moving average on summary statistics for several post-war 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. post-war quarterly macroeconomic time

---

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

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 [1980] filter, but we defer discussion of these results until Section 4.

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 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 $K \geq 12$.

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 $K = 12$.

3.3 Inspecting the results for GNP

Figure 4 displays the results of applying five filters to the natural logarithm of gross national product. Throughout the four graphs, we use the band-pass business-cycle filter with $K = 12$ as our reference point: it the dashed line which is present in all of the graphs. The common sample period for these graphs is 1947-1993, 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 4 shows the quarterly growth rate of real GNP vs. 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 4 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 displays a high-pass filter constructed using our procedures which 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 out of irregular components has little effect on the overall volatility.

The Deviation from 5 year Moving Average Filter: Finally, Panel D 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 quite close, with the moving average filter being somewhat more volatile.

3.4   Inspecting the results for inflation

In Figure 5, we present the results of applying the same five filters to the inflation rate. As before, the solid line in each panel is the $BP_K(6,32)$ business-cycle filter.

The First Difference Filter: Panel A of Figure 5 shows the quarterly growth rate of inflation vs. the band-pass filter. As before, the first difference filter produces a highly volatile time series which bears little resemblance to the band-pass filter.

The Hodrick-Prescott Filter: Panel B of Figure 5 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 which are passed by the Hodrick-Prescott filter, but which 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 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 (0.48 for the high-pass filter, versus 0.32 for the band-pass filter over this sample period).

The Deviation from 5 year Moving Average Filter: Finally, Panel D displays deviations from an equally-weighted 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.

4 Detailed Comparison with Other Filters

This section compares the properties of our proposed business cycle filter with other commonly used filters. We evaluate each filter in terms of its ability to achieve the following characteristics which we have argued are necessary for a "good" business-cycle filter: (i) ability to remove unit roots; (ii) absence of phase shift; (iii) ability to isolate business cycle frequencies without re-weighting components at the desired 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.

4.1 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 which 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 re-weight frequencies, the failure to remove unit root components from the data means that linear detrending is undesirable for most macroeconomic time series.

4.2 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 phase shift for this filter). Second, this filter involves a dramatic re-weighting of frequencies. Figure 6-A plots the frequency response function for this filter; the first-difference filter re-weights strongly toward the higher frequencies, while down-weighting lower frequencies. If the goal of a business cycle filter is to isolate fluctuations in the data which occur between specific periodicities, without special emphasis on any particular frequency, the first-difference filter is a poor choice.

14

4.3 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:

$$y_t^c = \left(\frac{\lambda(1 - L)^2(1 - L^{-1})^2}{1 + \lambda(1 - L)^2(1 - L^{-1})^2} \right) y_t \tag{10}$$

where λ is a parameter which 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 (10) 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):

$$y_t^g = \frac{\theta_1 \theta_2}{\lambda} \left[\sum_{j=0}^{\infty} \left(A_1 \theta_1^j + A_2 \theta_2^j \right) y_{t-j} + \sum_{j=0}^{\infty} \left(A_1 \theta_1^j + A_2 \theta_2^j \right) y_{t+j} \right] \tag{11}$$

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.[9]

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

$$\tilde{C}(\omega) = \frac{4\lambda \left(1 - \cos(\omega)\right)^2}{1 + 4\lambda \left(1 - \cos(\omega)\right)^2} \tag{12}$$

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 6-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 $\underline{\omega} = \pi/16$.

In terms of the objectives that we specified for our filter design problem, the Hodrick-Prescott cyclical filter has several desirable features. First, it is a symmetric

---

[9]Equation (11) 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 5 below; the discussion here focuses on the exact Hodrick-Prescott filter.

15

filter so that no phase shift is introduced. Second, it has trend reduction 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 rise 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 (10) 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 5 below.

4.4 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

$$y_t^g = \frac{1}{2K+1} \sum_{j=-K}^{K} y_{t-j}. \tag{13}$$

Thus the cyclic component of y_t is generated as $y_t^c = a(L)y_t$ with $a_0 = 1 - \frac{1}{2K+1}$, and $a_j = a_{-j} = \frac{1}{2K+1}$ for $j = 1, 2, ..., K$. This filter places zero weight at the zero frequency since $\sum a_k = 0$, and is symmetric. Figure 6-C plots the gain for the centered moving average filter for several values of K. The general shape of this filter is very similar to that of the approximate high-pass filter, plotted in Figure 6-D, although the "side-lobes" are more exaggerated for the moving average filter.

4.5 A high-pass filter

We have defined a high-pass business-cycle filter, $HP_K(32)$, as a filter which passes components of the data with periodicity less than or equal to 32 quarters. Figure 6-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., $K \geq 12$).

4.6 Comparisons across filters

Table 2 shows how application of these alternative filters affects moments computed from several postwar U.S. time series. We focus on three set 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 retain some high-frequency volatility which 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 re-weighting 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 3.4 above, inflation contains sizable high-frequency components—components which 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. 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 which 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).

5 Comparing HP's

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 two important dimensions: its ease of application to data sampled at frequencies other than quarterly, and its appropriate treatment of observations near the endpoints of the sample.

5.1 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 6, discussed in section 3.2 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.

5.2 The Hodrick-Prescott filter in finite samples

Many individuals currently use the Hodrick-Prescott filter with $\lambda = 1600$ to define cyclical components of quarterly economic time series. One main rationale for this, suggested 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). The results presented above suggested that there is indeed a close correspondence between alternative HP filters.

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 $t = 1, ...T$, the Hodrick-Prescott procedure produces estimates of the cyclical component, y_t^c, for $t = 1, ...T$.

To understand this result, 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^τ 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 $\Delta^2 y_t^\tau = \eta_t$ and $y_t^c = \epsilon_t$. If one knows the relative magnitude of σ_η^2 and σ_ϵ^2, then it is possible to extract estimates of y_t^τ and y_t^c at each date of a finite sample $t = 1, ...T$. 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 computing the Hodrick-Prescott filter makes it easy to recover the coefficients 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 reduction properties at every date.[10]

To begin our more detailed look at the time-varying filter, Figure 7 plots the gain of the linear filter $d_t(L) = \sum_{h=1}^{T} d_{ht} L^{(h-t)}$ for a range of dates t = {1,2,3}, {4,6,8},

---

[10]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.

19

{12,16,24}, {32,48,60}. These choices are motivated by the idea that we are studying a quarterly sample period of post-war 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 7.

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 observation 12, it would seem natural to drop 12 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. Further, our HP filter embeds a mechanical rule for handling (i.e., dropping) endpoints.

5.3 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 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 filter 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.

20

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, figure 8 plots annual GNP filtered with our $BP_3(2,8)$ filter together with data filtered with the Hodrick-Prescott filter, for several values of λ.[11] Examining the top two panels of figure 8 shows that the commonly-used values of $\lambda = 400$ and $\lambda = 100$ do not produce a filtered time series for GNP that closely resembles that produced by the band-pass filter. However, setting $\lambda = 10$, as in the third panel, produces a much better correspondence between the Hodrick-Prescott and band-pass filters. The bottom panel of this figure shows that little improvement is made when the length of the moving average is increased from $K = 3$ to $K = 6$. Figure 9 plots the gain for the Hodrick-Prescott filter for the three values of λ against the ideal filter. This figure shows 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 9 and 16 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."[12]

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. Figure 10 plots the gain for the finite-sample version of the Hodrick-Prescott filter for $\lambda = 10$, by observation number, in a manner comparable with figure 7 presented earlier. As in the prior case, the finite-sample version of the filter produces serious departures from the ideal filter for the first three observations, but improves dramatically from observation 4 onward.

Overall, we find that our approximate band-pass filters are more straightforward to apply, as they do not require the researcher to choose a new "smoothing parameter" when changing data frequencies. Second, we find that the commonly-used values of $\lambda = 100$ and $\lambda = 400$ for annual versions of the Hodrick-Prescott filter produce very poor approximations to a business-cycle filter. Third, we find that it is important to

---

[11]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.

[12]Hassler et al. [1992] also argue that $\lambda = 10$ is the appropriate value for the smoothing parameter when applying the Hodrick-Prescott filter to annual data.

drop at least three data points from each end of the sample when using the Hodrick-Prescott filter on annual data, even if one chooses $\lambda = 10$.

6   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, deviations from an equally-weighted moving-average 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, while producing a better approximation to the ideal filter. While the main 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. These filters are illustrated in Figure 11 and the weights are given in Table 4.

For quarterly macroeconomic data, we recommend the "Burns and 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 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. Figures 11-A and 11-B provide plots of these weights, which are also given 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 between two and eight years, with $K = 3$. The filter weights are illustrated in Figure 11-C and given in table 4.

We have applied the filters constructed in this paper in some recent work, 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 [1993] 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 eighteen months to ninety-six 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 of economic time series that is fast, flexible, and easy 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

[1] Backus, David and Patrick Kehoe, "International Evidence on the Historical Properties of Business Cycles," *American Economic Review*, vol. 82, no.4, (September 1992), 864-888.

[2] Baxter, Marianne, "Real Exchange Rates and Real Interest Differentials: Have We Missed the Business-Cycle Relationship?" *Journal of Monetary Economics*, vol. 33, no. 1 (February 1994), 5-37.

[3] Bry, Gerhard, and Charlotte Boschan, *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, New York, N.Y.: National Bureau of Economic Research, New York, 1981.

[4] Burns, Arthur M. and Wesley C. Mitchell, *Measuring Business Cycles*, New York, N.Y.: National Bureau of Economic Research, 1946.

[5] Canova, Fabio, "Detrending and business cycle facts," CEPR discussion paper #792, June 1993.

[6] Engle, Robert and Clive Granger, "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica*, vol. 55, no. 2 (March 1987), 251-276.

[7] Englund, Peter; Torsten Persson and Lars E.O. Svensson, "Swedish Business Cycles: 1861-1988," *Journal of Monetary Economics*, vol. 30, no. 3 (December 1992), 343-371.

[8] Granger, Clive, "The Typical Spectral Shape of An Economic Variable," *Econometrica*, vol 34, no. 1 (January 1966), 150-161.

[9] Harvey, Andrew C., *Time Series Models*, New York, N.Y.: John Wiley and Sons, 1981.

[10] Hassler, John; Petter Lundvik, Torsten Persson, and Paul Soderlind, "The Swedish Business Cycle: Stylized Facts over 130 Years," Monograph Series No. 22, Institute for International Economic Studies, Stockholm University, 1992.

[11] Hodrick, Robert J. and Edward Prescott, "Post-war U.S. Business Cycles: An Empirical Investigation," working paper, Carnegie-Mellon University, 1980.

[12] King, Robert G. and Sergio T. Rebelo, "Low Frequency Filtering and Real Business Cycles," *Journal of Economic Dynamics and Control*, vol. 17, no. 1 (January 1993), 207-231.

[13] King, Robert G. and Charles I. Plosser, "Real Business Cycles and the Test of the Adelmans," *Journal of Monetary Economics*, forthcoming,1994.

[14] King, Robert G. and Mark W. Watson, "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," Carnegie Rochester Conference on Public Policy, November 1994.

[15] Koopmans, Leonid, *The Spectral Analysis of Time Series*, New York, N.Y.:Academic Press, 1974.

[16] Li, Yikang, Mark Rush and Li Zhu, "Filtering methodology and fit in real business cycle models," manuscript, University of Florida, September 1994.

[17] Prescott, Edward C., "Theory Ahead of Business Cycle Measurement," *Carnegie-Rochester Conference Series on Public Policy*, 25, Fall 1986, 11-66.

[18] Watson, Mark W. "Business Cycle Duration and Postwar Stabilization of the U.S. Economy," *American Economic Review*, vol. 84, no. 1 (March 1994), 24-46.

A   Trend-reduction properties of symmetric moving average filters

In this appendix, we consider how symmetric moving average filters reduce series with deterministic and stochastic trends to series that are stationary. In particular, we consider the filter:

$$a(L) = \sum_{k=-K}^{K} a_k L^k \tag{14}$$

where L is the lag operator. We impose two conditions, which are that the filter's coefficients sum to zero and that the filter is symmetric:

$$a(1) = \sum_{k=-K}^{K} a_k = 0 \tag{15}$$

$$a_k = a_{-k}. \tag{16}$$

A.1   Deterministic trends

Consider a quadratic trend specification,

$$\tau_t = \gamma_0 + \gamma_1 t + \gamma_2 t^2.$$

We are interested in the effects of applying the two sided moving average to this trend, i.e., creating a new variable

$$s_t = a(L)\tau_t = \sum_{k=-K}^{K} a_k \tau_{t-k} = \gamma_0 \sum_{k=-K}^{K} a_k + \gamma_1 \sum_{k=-K}^{K} a_k(t-k) + \gamma_2 \sum_{k=-K}^{K} a_k(t-k)^2.$$

Writing out $(t - k)^2$ as $t^2 - 2tk + k^2$ and consolidating terms we find that

$$s_t = \{\gamma_0 \sum_{k=-K}^{K} a_k - \gamma_1 \sum_{k=-K}^{K} a_k k + \gamma_2 \sum_{k=-K}^{K} a_k k^2\} + \{\gamma_1 \sum_{k=-K}^{K} a_k - \gamma_2 \sum_{k=-K}^{K} a_k k\}t + \{\gamma_2 \sum_{k=-K}^{K} a_k\}t^2$$

It follows that the general conditions for trend reduction—elimination of t for all values of γ_1, γ_2—are as follows:

$$\sum_{k=-K}^{K} a_k = 0 \tag{17}$$

26

$$\sum_{k=-K}^{K} a_k k = 0 \tag{18}$$

The first of these conditions is imposed as eq. (15) above. Further, any symmetric moving average with $\sum_{k=-K}^{K} a_k = 0$ implies that there is trend reduction, since

$$\sum_{k=-K}^{K} a_k k = \sum_{k=1}^{K} (a_k - a_{-k}) k = 0$$

directly from the symmetry condition ($a_k = a_{-k}$). Thus, eq. (18) is satisfied under our assumptions. Hence, the filters defined by eqs. (14)-(16) reduce series containing quadratic deterministic trends series to ones with no influence of time.[13]

An Example: The simplest example arises if $K = 1$. Then, there is a single free parameter of the specification $a(L)$. In particular, symmetry implies that $a_1 = a_{-1} = -\theta$. The $a(1) = 0$ condition then implies that $a_0 = 2\theta$. Applying this filter directly to eq. (14), we find that:

$$s_t = 2\,\theta\,\gamma_2.$$

Thus, the influence of time is eliminated but there is not necessarily a zero-mean series as a result of the application of this filter.

Note that this filter may alternatively be written as:

$$a(L) = \theta[(-L^{-1} + 2 - L)] = \theta[(1 - L^{-1})(1 - L)],$$

which indicates that the basic trend reduction filter contains two differencing operations, a forward and backward difference. It is this property which permits it to remove the influence of time from eq. (14). We will see next that this "double difference" property holds for all members of the class.

A.2   Stochastic trends

We now consider the "differencing operations" implicit in more general trend reduction schemes. We write:

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

The first equality follows from assumption (15), $\sum_{k=-K}^{K} a_k = 0$. The second follows from the symmetry assumption (16), $a_k = a_{-k}$.

Now, consider the individual terms in the preceding sum. We can write:

---

[13]If we also wish to require that there is a zero mean for the series s_t, then we must also require that $\sum_{k=-K}^{K} a_k k^2 = 0$. However, we do not impose this condition in our analysis; means can always be removed after filtering the series of interest.

$$(L^k + L^{-k} - 2) = -(1 - L^k)(1 - L^{-k}).$$

We know that $(1 - L^k) = (1 - L)[1 + L + L^2 + ... \; L^{k-1}]$. Further, with a little bit of algebra, we can show

$$[1 + L^1 + L^2 + ... + L^{k-1}][1 + L^{-1} + L^{-2} + ... + L^{-(k-1)}] = \sum_{h=-(k-1)}^{(k-1)} (k - |h|)L^h.$$

Hence, we can write:

$$a(L) = \sum_{k=-K}^{K} a_k L^k = \sum_{k=-K}^{K} a_k L^k - a_k = -\sum_{k=1}^{K} a_k [(1-L^k)(1-L^{-k})] = -(1-L)(1-L^{-1})\psi_K(L)$$

where $\psi_K(L) = \{\sum_{k=1}^{K} a_k \sum_{h=-(k-1)}^{(k-1)} (k - |h|)L^h]\}$ and is a symmetric moving average with $K - 1$ leads and lags.

That is, our general moving average filter $a(L)$ contains (at least) two differences, i.e., that it has the ability to render stationary I(2) stochastic processes.[14] This accords with the finding on deterministic trend specifications in the prior section.

B Weights for the ideal low-pass filter

The inverse Fourier transform of the ideal low-pass filter implies that

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

where the second line derives from the fact the $\beta(\omega) = 1$ for $|\omega| \le \underline{\omega}$ and $\beta(\omega) = 1$ for $|\omega| > \omega$.

Hence, it follows that

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

and that

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

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

---

[14]Since $\psi_K(L)$ is a finite-term moving average, it does not alter the stationarity properties of the series to which it is applied.

C Optimal approximation of symmetric linear filters

In this section, we consider the optimal approximation of an ideal symmetric linear filter by a K^{th} order symmetric moving average. (This filter might be a low-pass filter, but it is useful to consider the more general case, for reasons that will become clear below). The problems is to minimize

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

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 $\alpha(0) = \phi$.

To solve this constrained maximization problem, we form the Lagrangian,

$$\mathcal{L} = Q + \lambda[\phi - \alpha(0)].$$

The first order conditions are that:

$$\frac{\partial \mathcal{L}}{\partial a_0} = \frac{\partial Q}{\partial a_0} - \lambda = 0$$

$$\frac{\partial \mathcal{L}}{\partial a_h} = \frac{\partial Q}{\partial a_h} - 2\lambda = 0 \text{ for } h = 1, ...K$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \phi - \alpha(0) = 0.$$

To evaluate $\partial Q/\partial a_h$, it is desirable to proceed as follows. First, we compute $\partial|\delta(\omega)|/\partial a_h$. Second, we compute the relevant integrals.

C.1 Computation of the partial derivative

The partial derivative of interest is found as follows. We begin by noting that:

$$\frac{\partial}{\partial a_h}[\delta(\omega)\delta(-\omega)] = \frac{\partial\delta(\omega)}{\partial a_h}\delta(-\omega) + \delta(\omega)\frac{\partial\delta(-\omega)}{\partial a_h} = 2\,\delta(\omega)\frac{\partial\delta(\omega)}{\partial a_h},$$

where the second equality follows from the symmetry of the discrepancy measures, $\delta(\omega) = \beta(\omega) - \alpha(\omega) = \beta(-\omega) - \alpha(-\omega) = \delta(-\omega)$.

Further, since the frequency response function of the approximating filter is $\alpha(\omega) = \sum_{h=-K}^{K} a_h e^{-i\omega h}$, it follows that:

$$\partial\delta(\omega)/\partial a_0 = -1$$

29

and that

$$\partial \delta(\omega)/\partial a_h = -(e^{i\omega h} + e^{-i\omega h}) \text{ for } h = 1, 2, ...K$$

.

These results imply that $\partial Q/\partial a_h = \int_{-\pi}^{\pi} \partial |\delta(\omega)|/\partial a_h d\omega$ takes the form

$$\partial Q/\partial a_h = -\int_{-\pi}^{\pi} 2 \, \delta(\omega) d\omega \text{ for } h = 0$$

$$\partial Q/\partial a_h = -\int_{-\pi}^{\pi} 2 \, \delta(\omega)(e^{i\omega h} + e^{-i\omega h}) d\omega \text{ for } h = 1, 2, ...$$

Our next task is to evaluate these integrals, which it turns out is best done by deriving a set of intermediate results.

C.2 Intermediate results

If we have a symmetric linear filter $g(L) = \sum_{h=-\infty}^{\infty} g_h L^h$ with frequency response function $\gamma(\omega) = \sum_{h=-\infty}^{\infty} g_h e^{-i\omega h}$, then it follows that there are simple expressions for integrals of the form $\int_{-\pi}^{\pi} \gamma(\omega) d\omega$ and $\int_{-\pi}^{\pi} \gamma(\omega)(e^{i\omega j} + e^{-i\omega j}) d\omega$ for integer j. To evaluate these integrals, we make repeated use of the facts that $\int_{-\pi}^{\pi} e^{i\omega n} e^{-i\omega m} d\omega = 0$ for $n \neq m$ and $\int_{-\pi}^{\pi} e^{i\omega n} e^{-i\omega m} d\omega = 2\pi$ for $n = m$, where n and m are integers.

Using this pair of results in the first integral, we find that:

$$\int_{-\pi}^{\pi} \gamma(\omega) d\omega = \int_{-\pi}^{\pi} \sum_{h=-\infty}^{\infty} g_h e^{-i\omega h} d\omega = 2\pi \, g_0 \, .$$

since all terms integrate to zero for $h \neq 0$. Using them in the second integral, we find that:

$$\int_{-\pi}^{\pi} \gamma(\omega)(e^{i\omega j} + e^{-i\omega j}) d\omega = \int_{-\pi}^{\pi} \sum_{h=-\infty}^{\infty} g_h e^{-i\omega h}(e^{i\omega j} + e^{-i\omega j}) d\omega = 2\pi(g_j + g_{-j}) = 4 \pi \, g_j.$$

C.3 Evaluation of the first order conditions

The $K + 1$ first order conditions, $\partial Q/\partial \alpha_h - \lambda = 0$, then can be expressed as:

$$-4\pi[b_0 - a_0] + \lambda = 0$$

and as

$$-8 \pi[b_h - a_h] + 2\lambda = 0 \text{ for } h = 1, 2, ...K.$$

Hence, if there is no constraint on $\alpha(0)$, i.e., $\lambda = 0$, then it follows that the optimal approximate filter is simply derived by truncation of the ideal filter's weights. Notice that this result is quite general in that it applies for the approximation of any

symmetric ideal filter, *i.e.*, it does not make use of any properties of $\beta(\omega)$ except for symmetry.

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:

$$a_0 = b_0 + \theta \text{ and } a_h = b_h + \theta, \text{ for } h = 1, 2, ...K,$$

where $\theta = \lambda/(8\pi)$. Then, requiring that $\alpha(0) = \sum_{h=-K}^{K} a_h = \phi$, we find that the required value adjustment is

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

Again, the result is a general one. It implies that construction of the optimal approximating filter contains two steps: first, truncation of the ideal filter's weights and, second, addition of correction term θ which depends on the extent to which the truncation disturbs the desired zero frequency behavior.

Further, the form of this correction process makes clear the origins of some of the observations made in the main text. For example, it is easy to see 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.

31

D MATLAB programs

Following are the programs used to implement the approximate band-pass filters developed in this paper. The main program is called BPF.M; this program calls a second program called FILTK.M. A replication diskette is available from the authors which contains these and other programs used to produce this paper.

Program name: BPF.M

```
function yf=bpf(y,up,dn,K);

% bpf.m
% Program to compute band-pass filtered series
% Inputs are
% y:   data (rows = observations, columns=series)
% up:  period corresponding to highest frequency (e.g., 6)
% dn:  period corresponding to lowest frequency (e.g., 32)
% K:   number of terms in approximating moving average
% [calls filtk.m (filter with symmetric weights) as subroutine]
x=[up dn];
disp(' ')
disp('bpf(y,up,dn,K): band-pass filtering of series y with symmetric MA(2K+1)')
disp(' ')
disp('             for additional information see: ')
disp(' ')
disp('             M. Baxter and R.G. King      ')
disp(' ')
disp('             Measuring Business Cycles:   ')
disp('             Approximate Band-Pass Filters')
disp('             for Macroeconomic Time Series')
disp(' ')
disp('Filter extracts components between periods of: ')
disp('  up   dn')
disp(x)
% pause(2)

if (up>dn)
 disp('Periods reversed: switching indices up & dn')
 disp(' ')
 dn=x(1); up=x(2);
end
```

```
if (up<2)
 up=2;
 disp('Higher periodicity > max: Setting up=2')
 disp(' ')
end

% convert to column vector
[r c]=size(y);
if (r<c)
 y=y';
 disp('There are more columns than rows: Transposing data matrix')
 disp(' ')
end

% Implied Frequencies
omubar=2*pi/up;
omlbar=2*pi/dn;

% An approximate low pass filter, with a cutoff frequency of "ombar",
% has a frequency response function
%
%  alpha(om) = a0 + 2*a1 cos(om) + ... 2*aK cos(K om)
%
% and the ak's are given by:
%
% a0 = ombar/(pi)          ak = sin(k ombar)/(k pi)
%
% where ombar is the cutoff frequency.

% A band-pass filter is the difference between two
% low-pass filters,
%  bp(L) = bu(L) - bl(L)
% with bu(L) being the filter with the high cutoff point and bl(L) being
% that with the low cutoff point.  Thus, the weights are differences
% of weights for two low-pass filters.

% Construct filter weights for bandpass filter (a(0)....a(K)).

akvec=zeros(1,1:K+1);

akvec(1)=(omubar-omlbar)/(pi);  % weight at k=0
```

```
for k=1:K;
  akvec(k+1)=(sin(k*omubar)-sin(k*omlbar))/(k*pi); % weights at k=1,2,...K
end

% Impose constraint on frequency response at om = 0
% (If high pass filter, this amounts to requiring that weights sum to zero).
% (If low pass filter, this amounts to requiring that weights sum to one).

if (dn>1000)
 disp('dn > 1000: assuming low pass filter')
 phi=1;
else
 phi=0;
end

% sum of weights without constraint
theta=akvec(1)+2*sum(akvec(2:K+1));
% amount to add to each nonzero lag/lead to get sum = phi
theta=phi-(theta/(2*K+1));
% adjustment of weights
akvec=akvec+theta;

% filter the time series
yf=filtk(y,akvec);

if (r<c)
 yf=yf';
end
→
```

Program name: FILTK.M

```
function  yf=filtk(y,a);

% Filter data with a filter with symmetric filter with weights
% data is organized (rows=obs, columns=series)
% a=[a0 a1 ... aK];

K=max(size(a))-1;  % max lag;

T=max(size(y));    % number of observations;
```

```matlab
% Set vector of weights

avec=zeros(1,2*K+1);
avec(K+1)=a(1);
for i=1:K;
avec(K+1-i)=a(i+1);
avec(K+1+i)=a(i+1);
end
yf=zeros(y);
for t=K+1:1:T-K
 yf(t,:)=avec*y(t-K:t+K,:);
end
```

TABLE 1

Effect of K on moments of an AR(1) process

K	Autocovariance at lag:				
	0	1	2	4	8
2	0.02	0.01	-0.01	0.00	0.00
3	0.15	0.09	-0.01	-0.07	0.00
4	0.42	0.32	0.10	-0.21	0.00
6	0.87	0.74	0.42	-0.27	-0.17
8	0.94	0.81	0.49	-0.24	-0.24
12	1.34	1.21	0.86	0.08	-0.35
16	1.29	1.16	0.81	0.03	-0.37
20	1.24	1.10	0.76	-0.03	-0.45
24	1.25	1.11	0.76	-0.03	-0.47
32	1.33	1.20	0.84	0.04	-0.43
48	1.34	1.20	0.85	0.04	-0.46
60	1.33	1.19	0.84	0.03	-0.48
90	1.34	1.20	0.84	0.03	-0.48
exact	1.38	1.23	0.87	0.05	-0.49
no filter	10.26	9.74	9.26	8.35	6.80

TABLE 2

Effect of filtering on moments:
quarterly data, 1947:1 - 1993:1

A. Standard deviations

Variable	K: truncation point for band-pass filter					Moving average	Hodrick-Prescott	First difference
	4	8	12	16	20			
GNP	0.84	1.31	1.63	1.59	1.59	1.97	1.71	1.00
Cons: durables	2.44	3.91	4.86	4.75	4.80	6.06	5.38	3.81
Cons: nondurables	0.59	0.96	1.10	1.09	1.05	1.32	1.18	0.75
Cons: durables	0.32	0.47	0.60	0.58	0.56	0.75	0.66	0.50
Investment	2.42	4.10	5.25	5.11	5.25	6.33	5.57	2.72
Hours per person	0.25	0.39	0.40	0.40	0.40	0.46	0.42	0.29
Employment	0.71	1.20	1.46	1.43	1.42	1.72	1.50	0.72
Exports	2.46	4.04	4.98	4.91	4.96	6.37	5.58	4.57
Imports	2.52	3.97	4.67	4.56	4.53	5.78	5.22	4.10
Net exports*	5.38	9.89	15.41	14.62	13.44	19.51	16.77	8.30
Gov't purchases	0.98	2.03	3.32	3.16	2.96	4.19	3.31	1.21
GNP deflator	0.34	0.63	0.95	0.90	0.77	1.18	0.90	0.70
Inflation*	0.18	0.25	0.28	0.27	0.28	0.42	0.39	0.50

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 - 1988:1 (the longest period available for the K=20 band-pass filter). Except for starred variables, natural logs were taken before filtering.

TABLE 2, cont'd.

B. First-order autocorrelation

Variable	K: truncation point for band-pass filter					Moving average	Hodrick-Prescott	First difference
	4	8	12	16	20			
GNP	0.79	0.87	0.91	0.91	0.91	0.87	0.84	0.31
Cons: durables	0.78	0.88	0.92	0.92	0.91	0.80	0.76	-0.05
Cons: nondurables	0.81	0.88	0.91	0.91	0.91	0.85	0.82	0.24
Cons: durables	0.75	0.86	0.91	0.90	0.90	0.79	0.75	0.10
Investment	0.83	0.90	0.93	0.93	0.93	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	0.92	0.92	0.92	0.91	0.90	0.61
Exports	0.78	0.89	0.92	0.91	0.91	0.74	0.67	-0.19
Imports	0.78	0.87	0.91	0.90	0.90	0.76	0.72	-0.11
Net exports*	0.84	0.93	0.96	0.96	0.94	0.92	0.90	0.31
Gov't purchases	0.85	0.95	0.97	0.97	0.96	0.96	0.94	0.32
GNP deflator	0.86	0.92	0.96	0.95	0.94	0.94	0.90	0.75
Inflation*	0.65	0.82	0.86	0.86	0.85	0.28	0.20	-0.47

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 - 1988:1 (the longest period available for the K=20 band-pass filter). Except for starred variables, natural logs were taken before filtering.

TABLE 2, cont'd.

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	0.77	0.75	0.68	0.68	0.67	0.64	0.65	0.56
Cons: nondurables	0.72	0.80	0.81	0.81	0.82	0.77	0.77	0.47
Cons: durables	0.50	0.69	0.73	0.72	0.75	0.70	0.70	0.40
Investment	0.88	0.87	0.83	0.83	0.85	0.81	0.82	0.74
Hours per person	0.85	0.84	0.82	0.82	0.83	0.80	0.80	0.68
Employment	0.83	0.85	0.87	0.86	0.87	0.86	0.85	0.71
Exports	0.33	0.28	0.28	0.29	0.33	0.28	0.27	0.21
Imports	0.68	0.76	0.74	0.75	0.78	0.70	0.69	0.29
Net exports*	-0.36	-0.41	-0.37	-0.37	-0.42	-0.35	-0.36	-0.13
Gov't purchases	0.15	0.10	0.25	0.24	0.17	0.28	0.23	0.24
GNP deflator	0.09	-0.11	-0.34	-0.32	-0.32	-0.38	-0.38	-0.11
Inflation*	0.32	0.34	0.31	0.30	0.35	0.22	0.23	0.07

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 - 1988:1 (the longest period available for the K=20 band-pass filter). Except for starred variables, natural logs were taken before filtering.

TABLE 3

Effect of Hodrick-Prescott filter with time-varying weights

observation	variance
1	1.7468
2	1.2320
3	1.0222
4	0.9970
6	1.1831
8	1.4053
12	1.6038
16	1.6167
24	1.6614
32	1.6967
48	1.6985
60	1.6990
90	1.6990

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	

FIGURE 1: Ideal filters

A. Ideal low-pass filter: cutoff frequency = 32 quarters

B. Ideal high-pass filter: cutoff frequency = 32 quarters

C. Ideal band-pass filter: 6-32 quarters
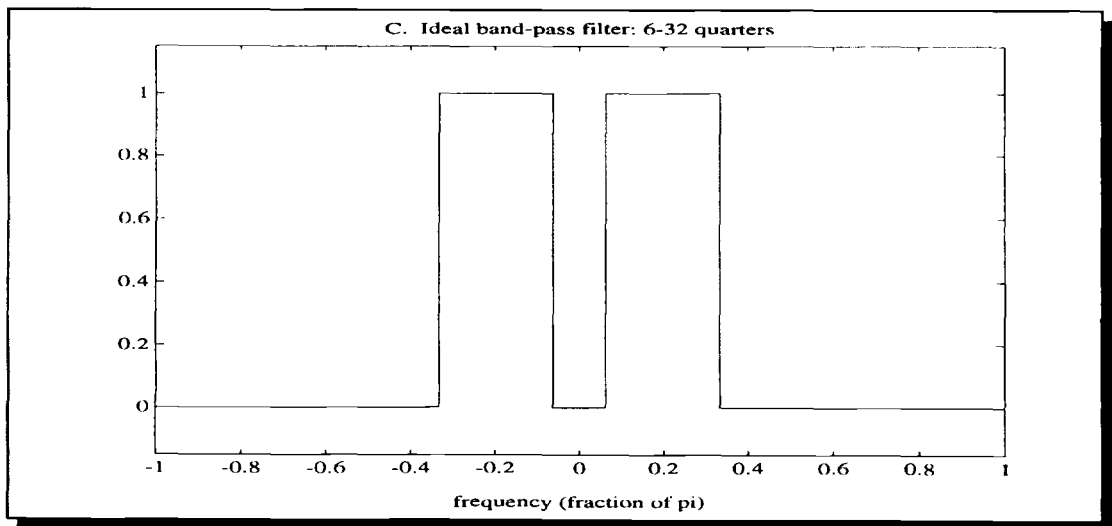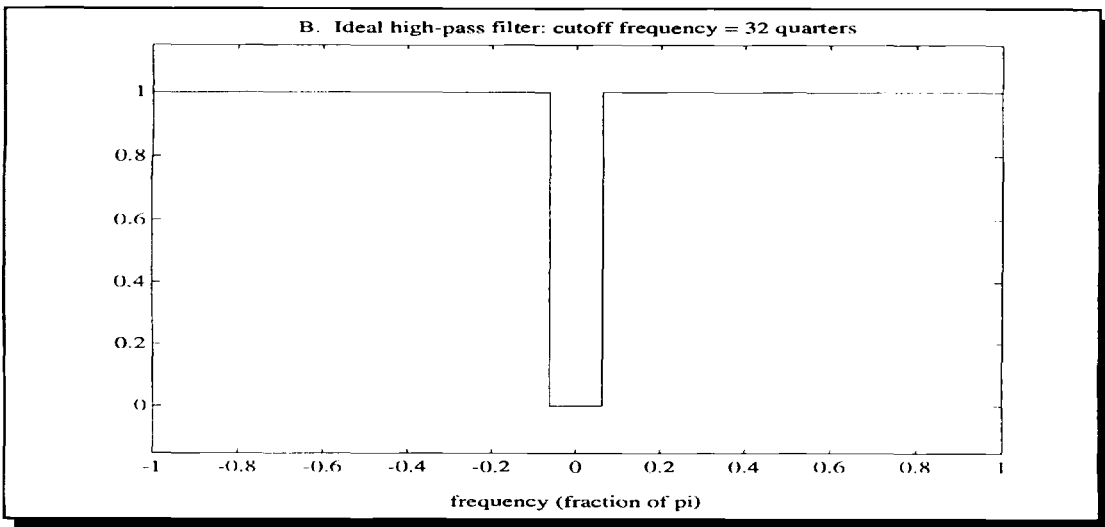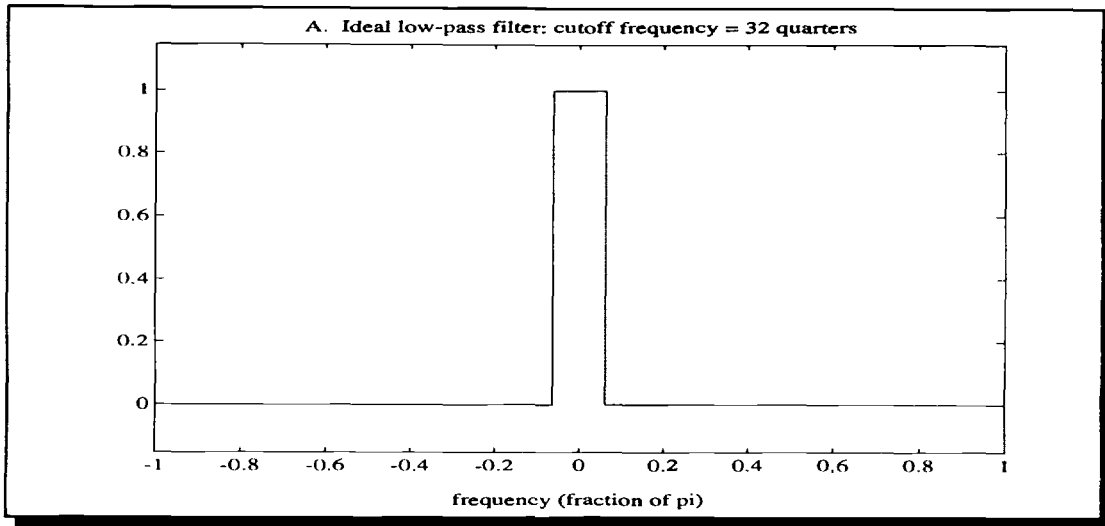
FIGURE 2: Constrained approximate low-pass filters
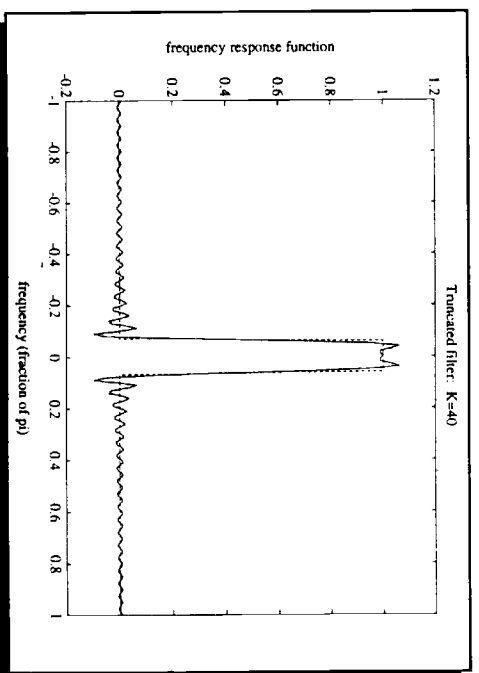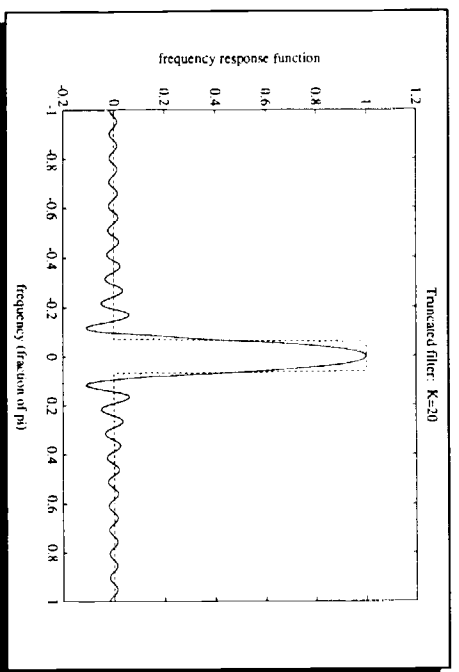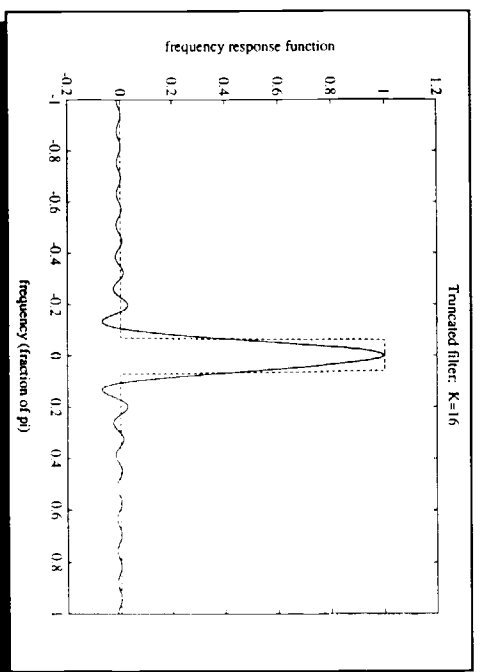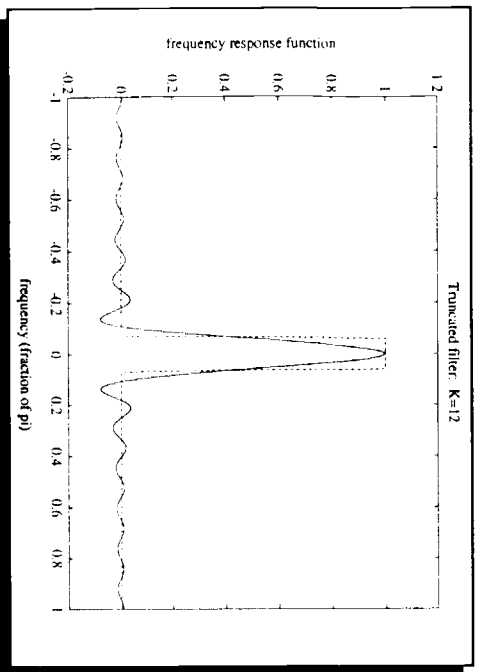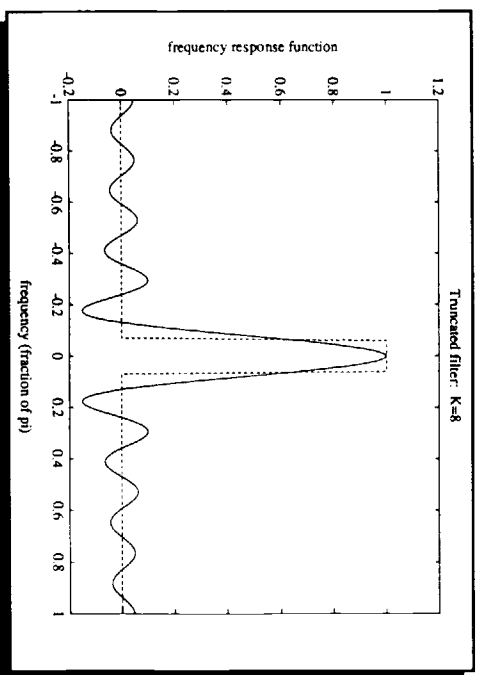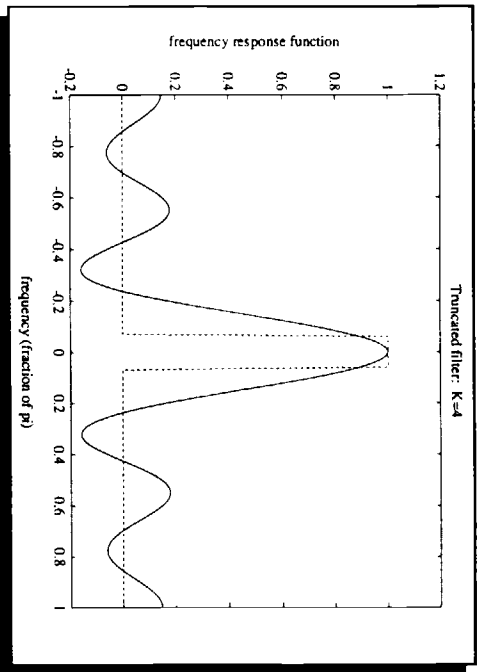
FIGURE 3: Constrained approximate band-pass filters (6-32 quarters)
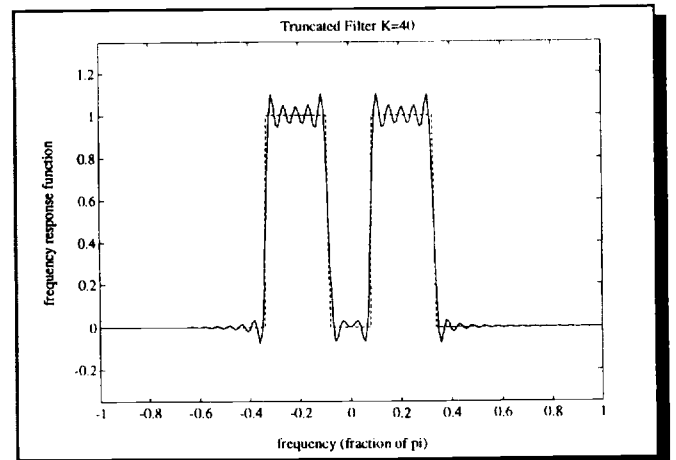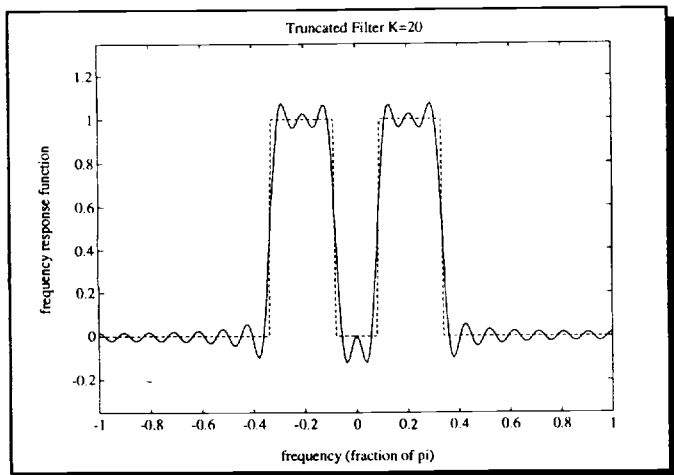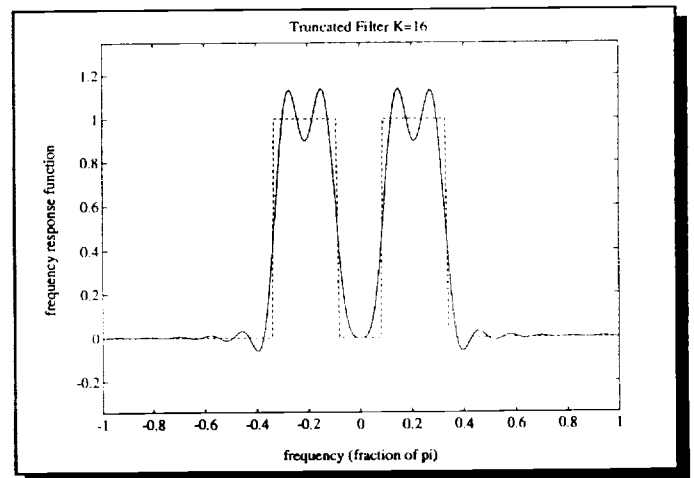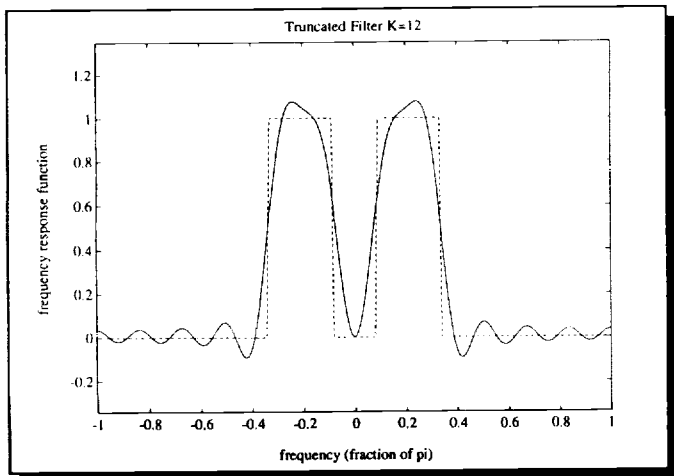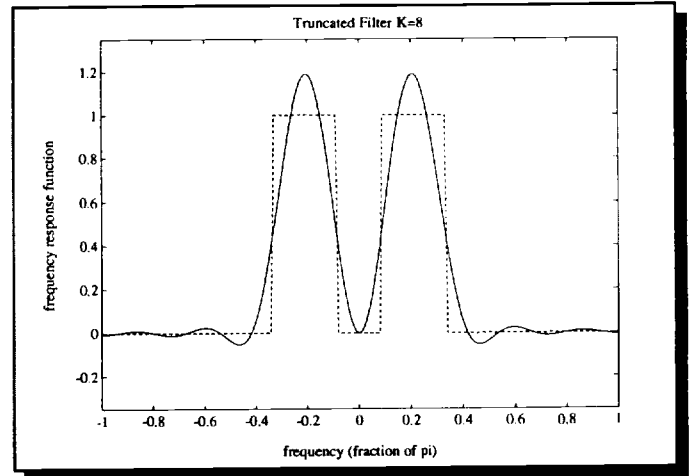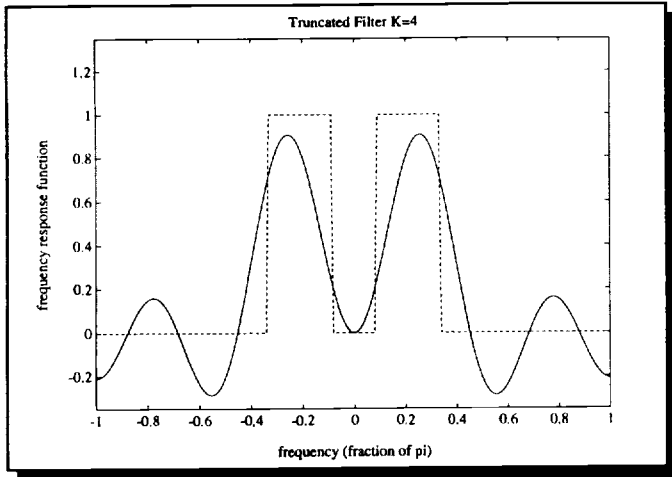
FIGURE 4: Effects of alternative filters on GNP

A. First Difference vs. Band-pass(6,32) for K=12
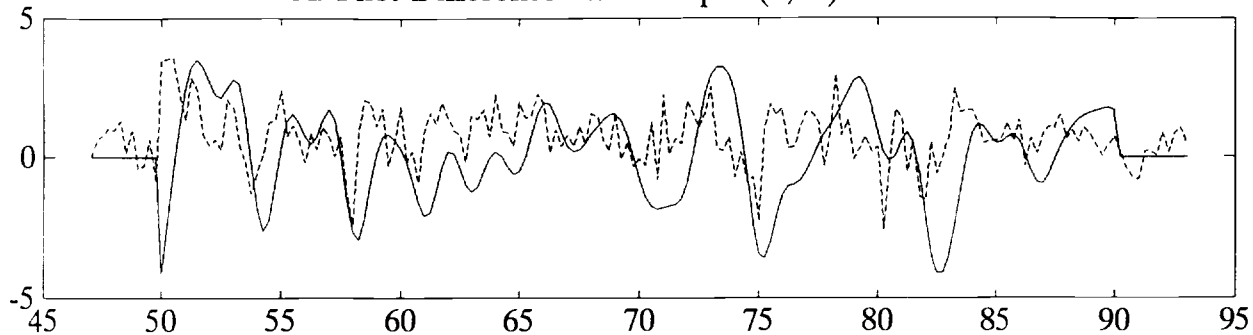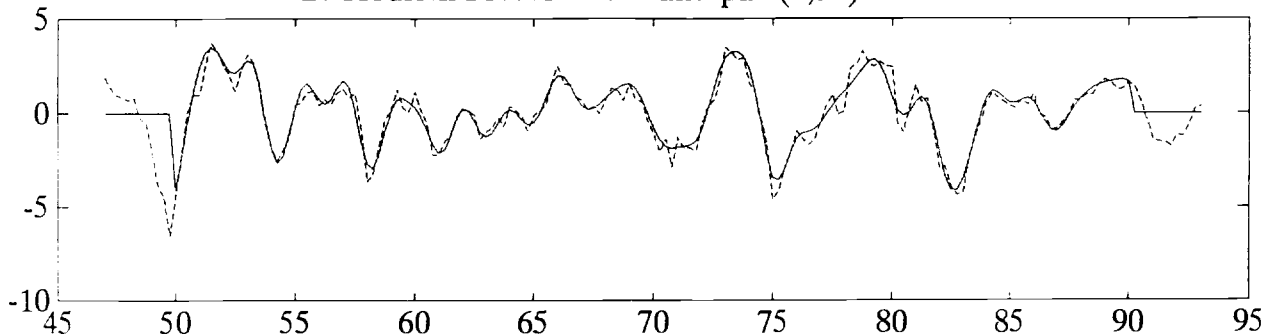
B. Hodrick-Prescott vs. Band-pass(6,32) for K=12

C. High-Pass (2,32) vs. Band-pass(6,32) for K=12
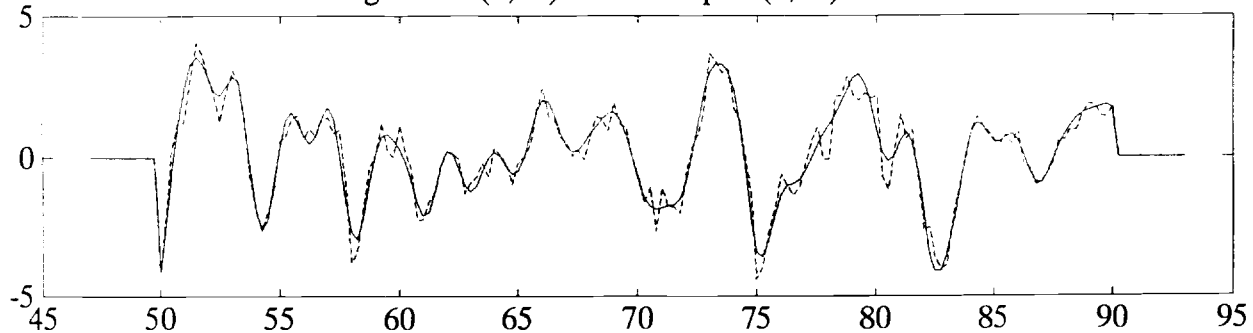
D. Deviation from Moving Average vs. Band-pass(6,32) for K=12
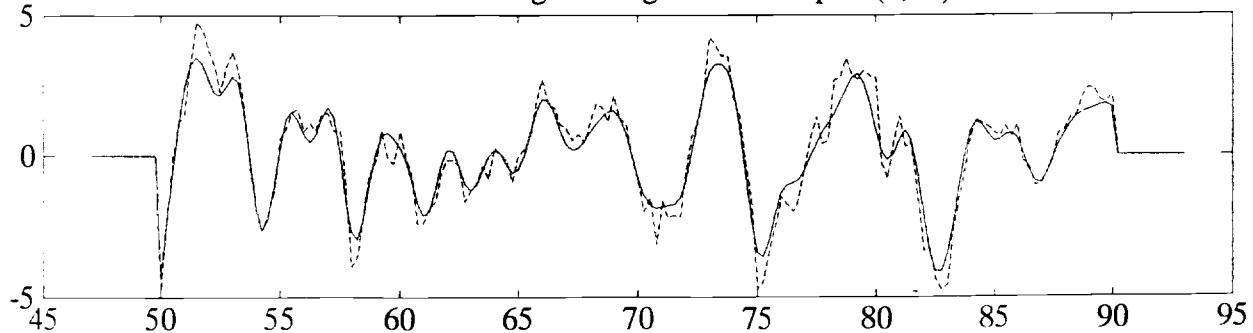
FIGURE 5: Effects of alternative filters on inflation

A. First Difference vs. Band-pass(6,32) for K=12
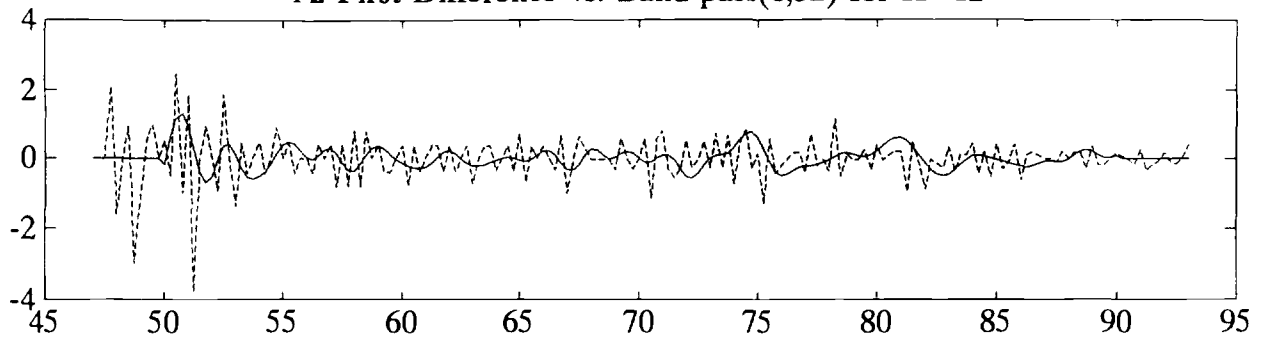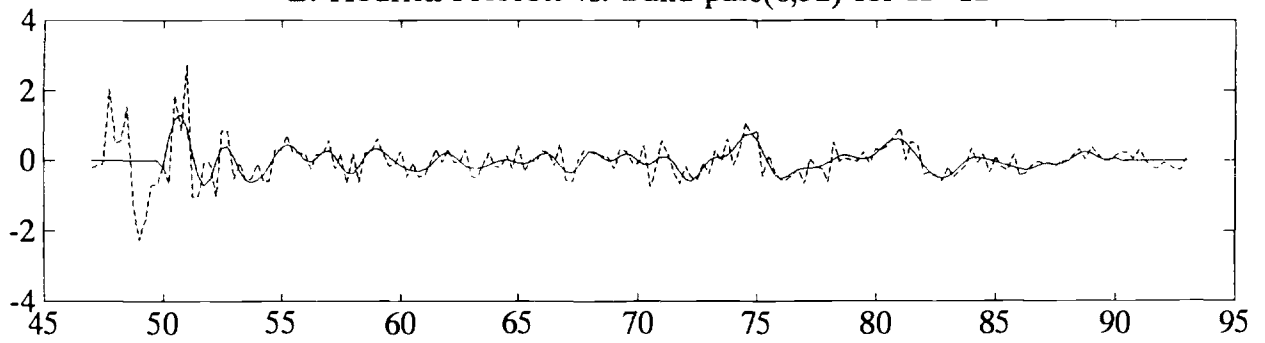
B. Hodrick-Prescott vs. Band-pass(6,32) for K=12

C. High-Pass (2,32) vs. Band-pass(6,32) for K=12
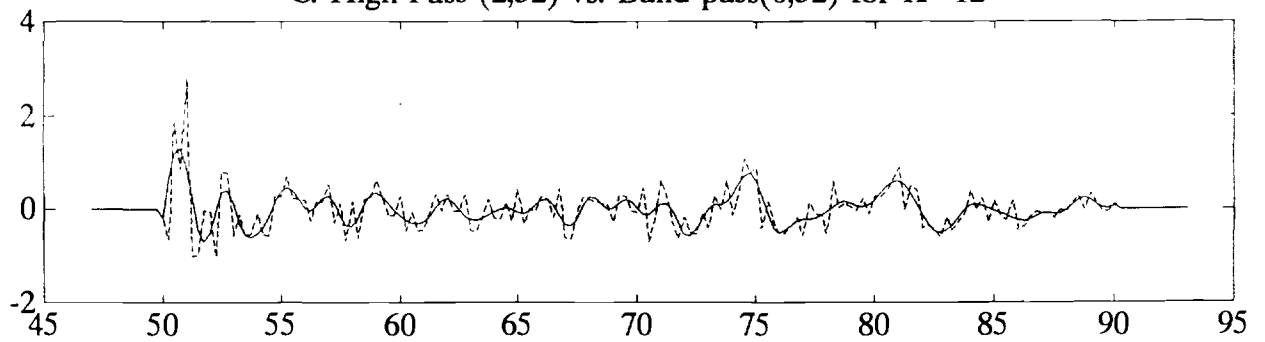
D. Deviation from Moving Average vs. Band-pass(6,32) for K=12
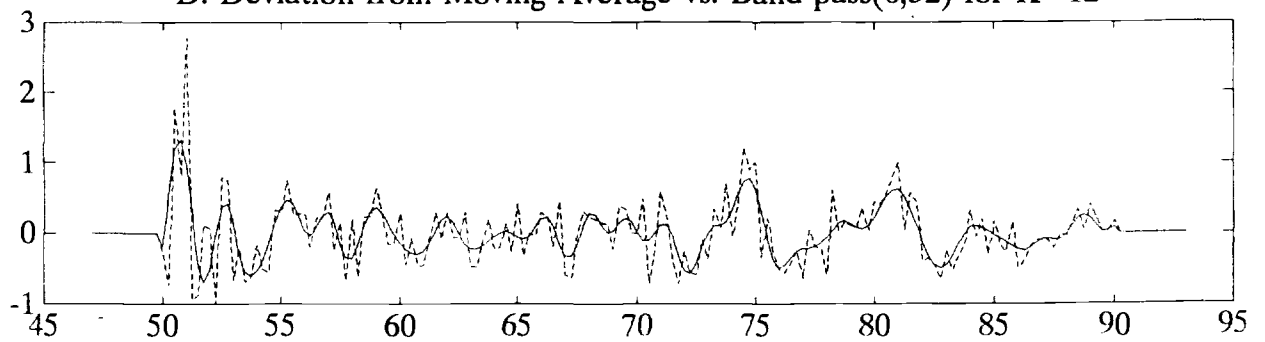
FIGURE 6: Alternative filters vs. an ideal high-pass filter (cutoff=32 quarters)

A. First difference

B. Hodrick-Prescott filter (lambda = 1600)

C. Centered moving average: K = maximum lag/lead (qtrs)

D. Approximate high-pass filter: K = truncation point
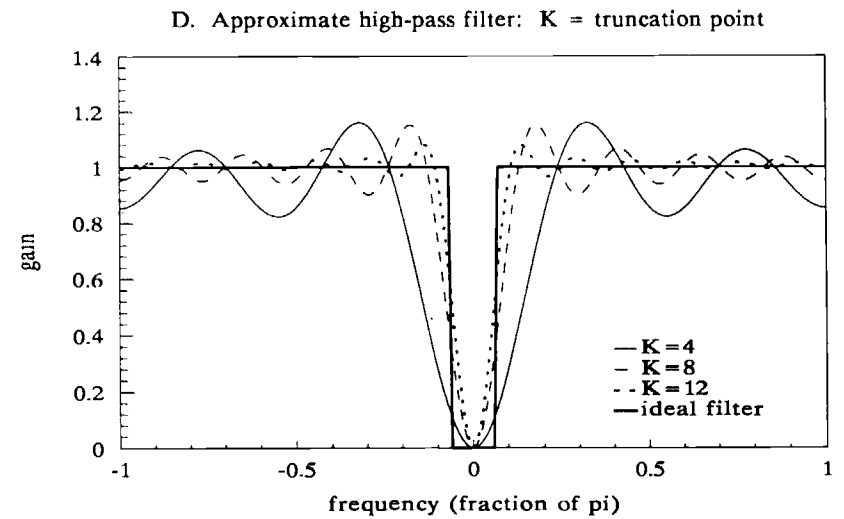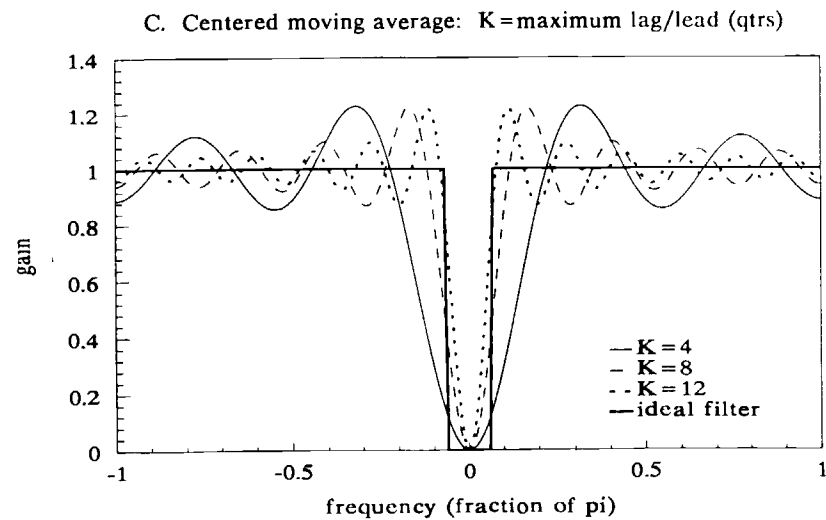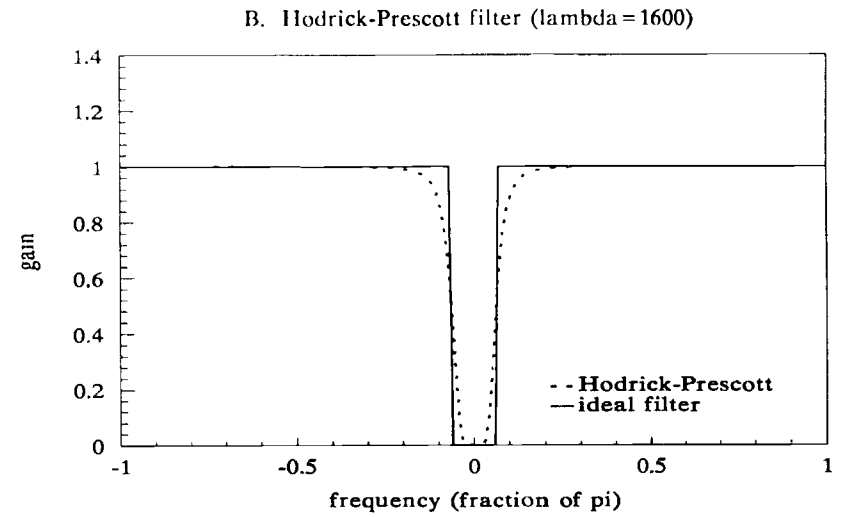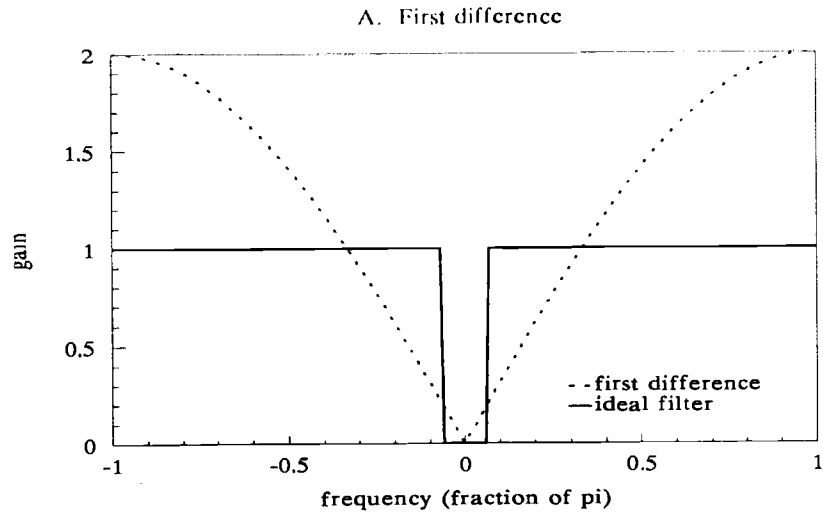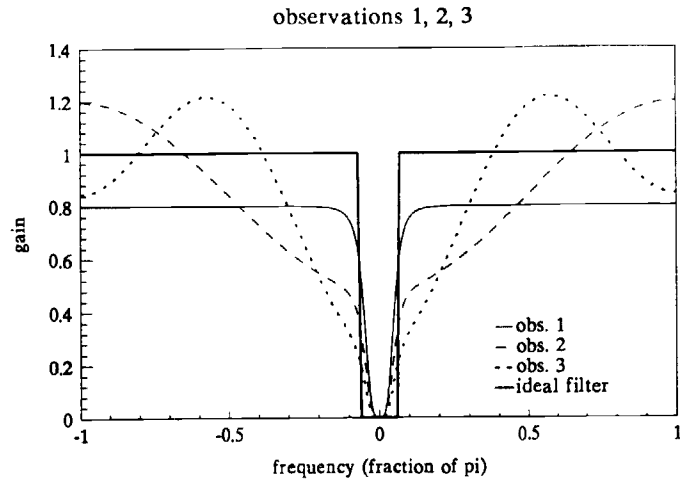
FIGURE 7: The Hodrick-Prescott filter in finite samples

observations 1, 2, 3

observations 4, 6, 8

observations 12, 16, 24

observations 32, 48, 60
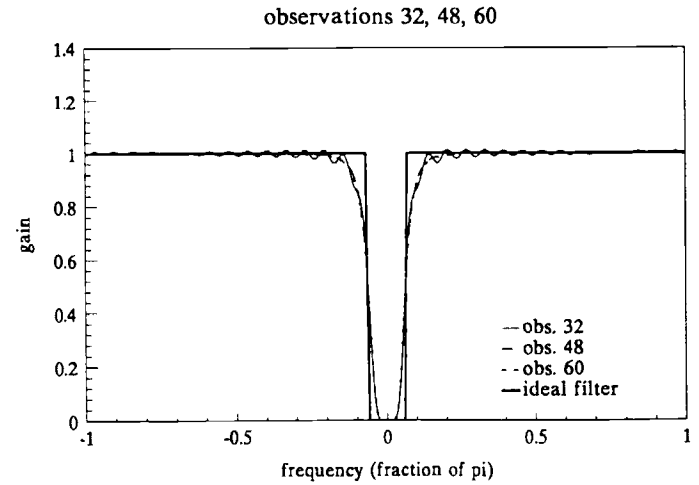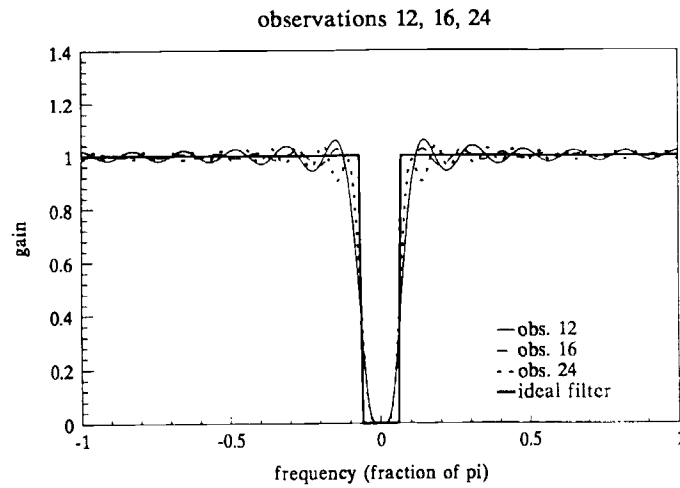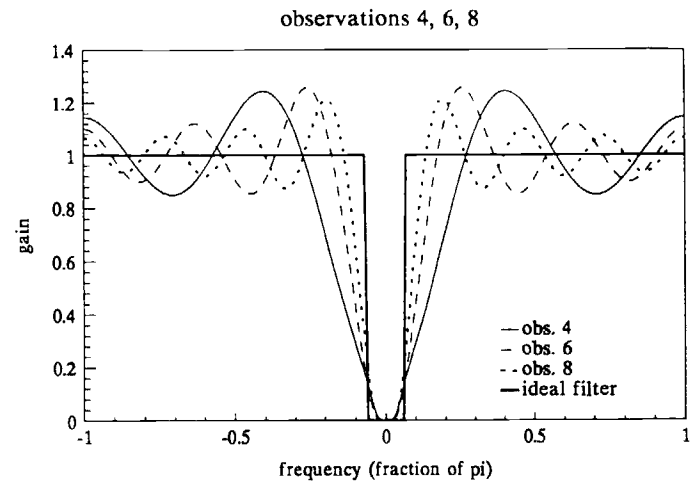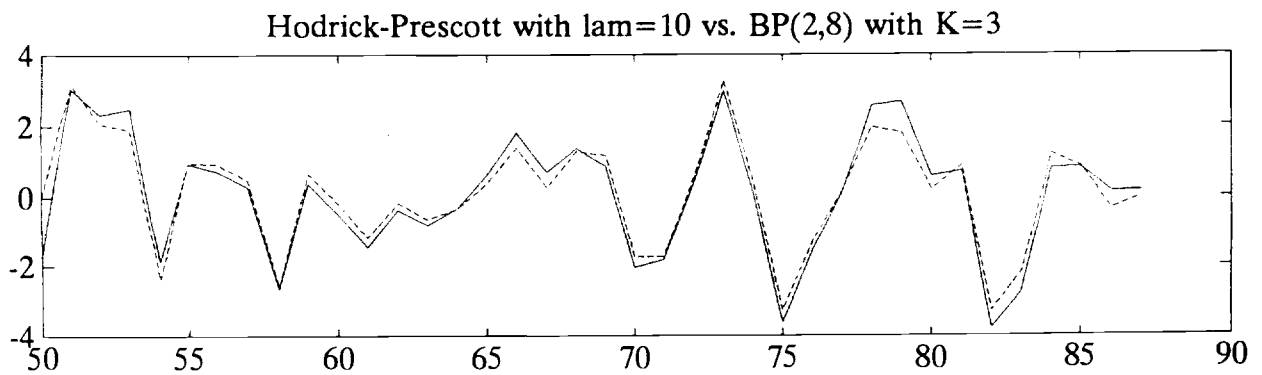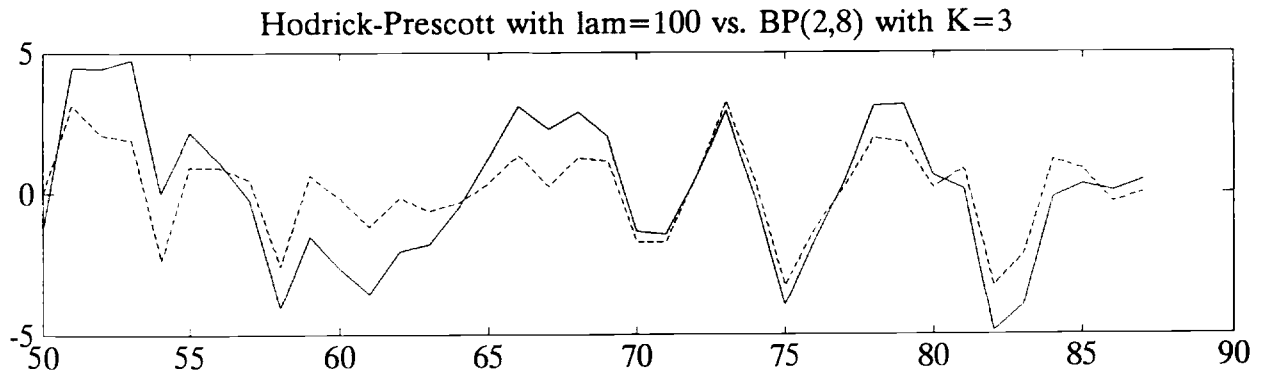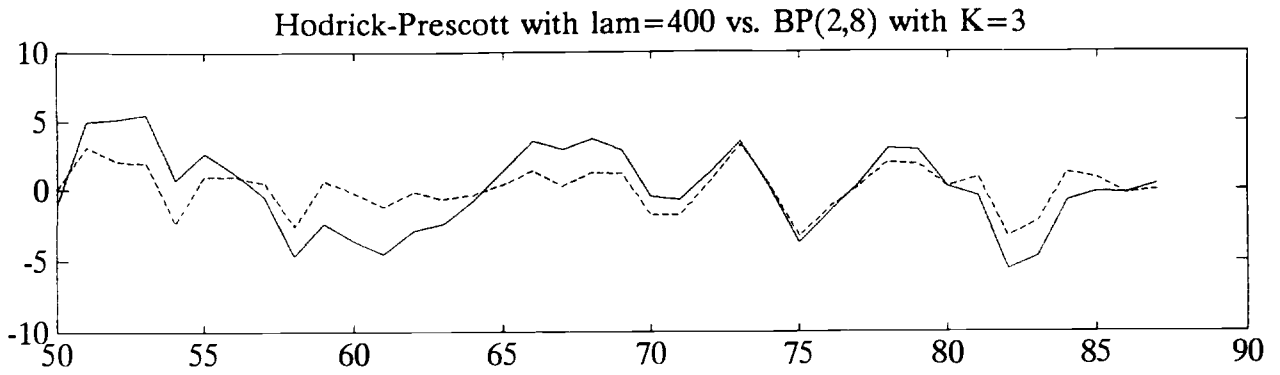
FIGURE 8: Effects of alternative Hodrick-Prescott filters on GNP

Hodrick-Prescott with lam=400 vs. BP(2,8) with K=3

Hodrick-Prescott with lam=100 vs. BP(2,8) with K=3

Hodrick-Prescott with lam=10 vs. BP(2,8) with K=3

BP(2,8) with K=6 vs. BP(2,8) with K=3
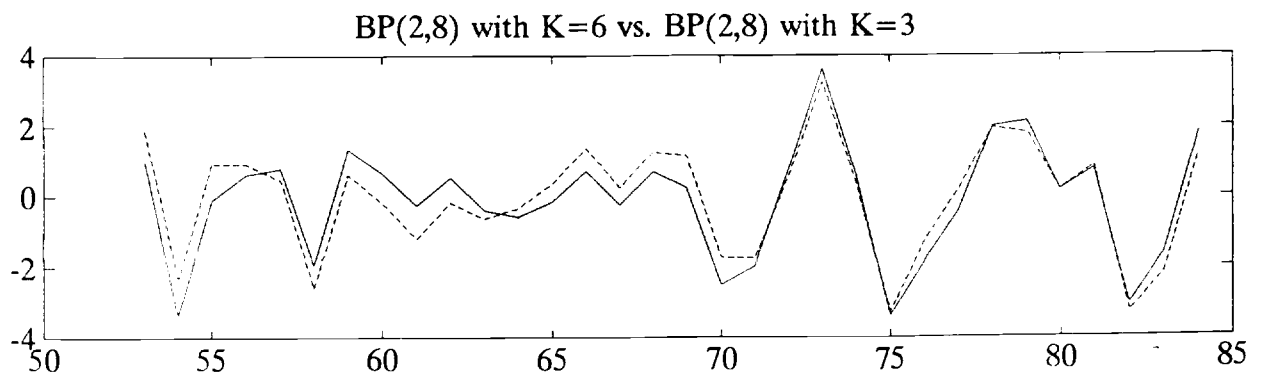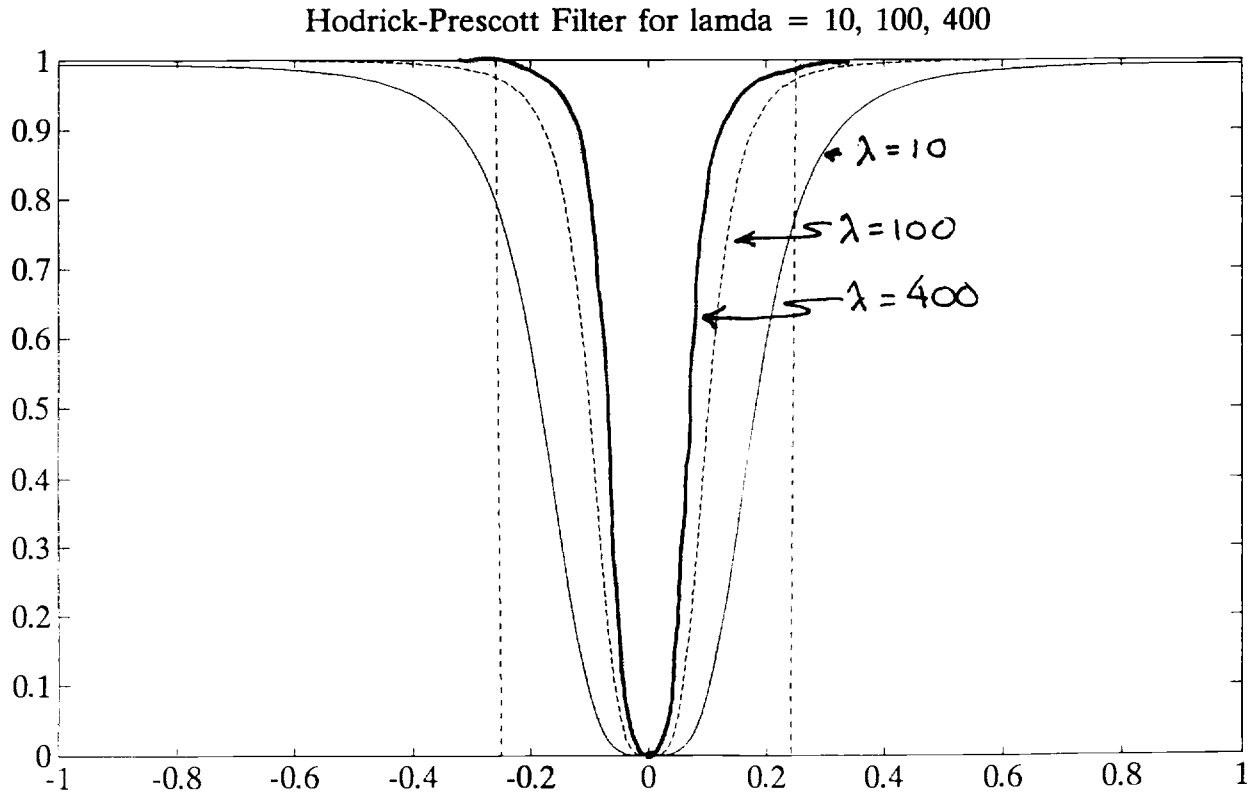
FIGURE 9: Alternative annual Hodrick-Prescott filters

Hodrick-Prescott Filter for lamda = 10, 100, 400

FIGURE 10: Annual Hodrick-Prescott filters in finite samples

A. Observations 1,2,3
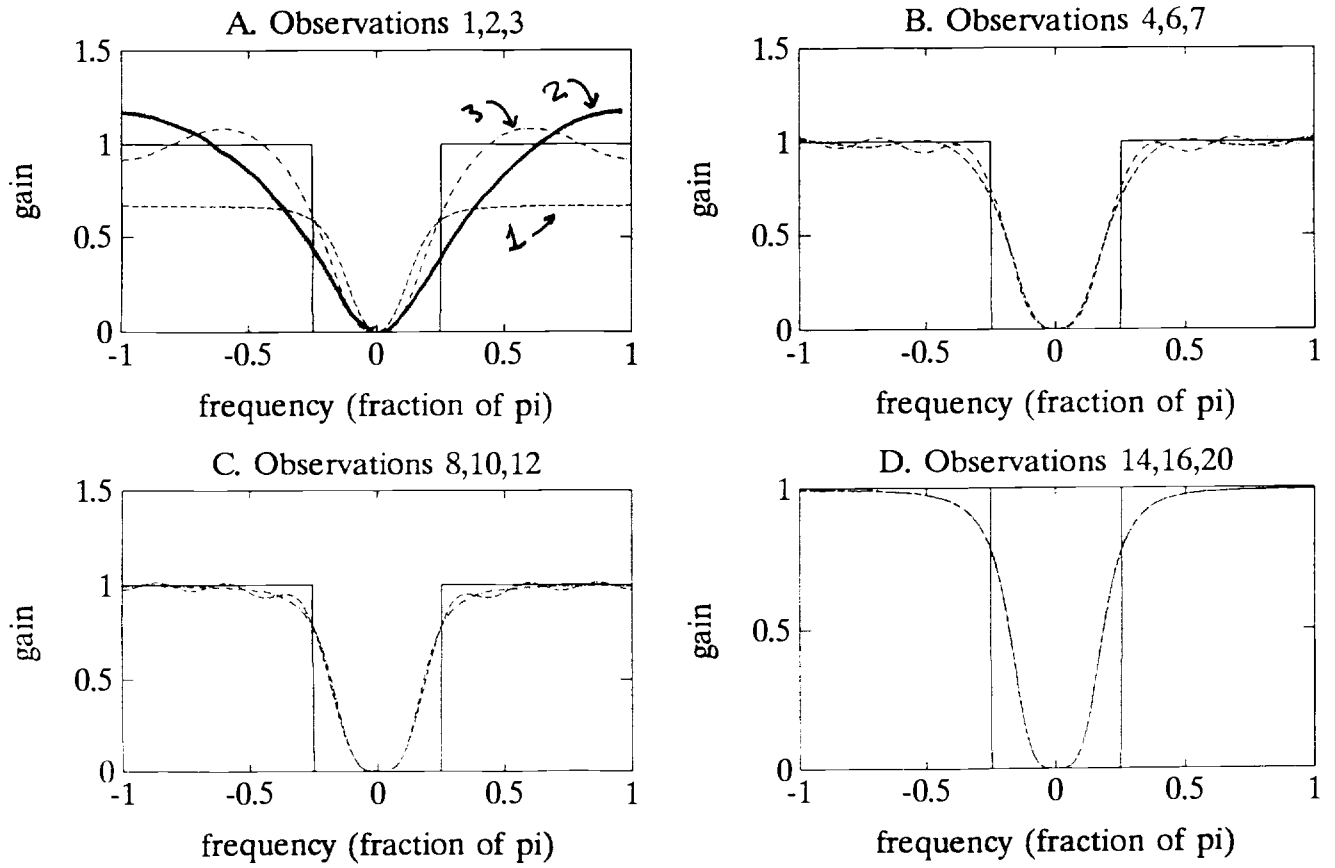
B. Observations 4,6,7

C. Observations 8,10,12

D. Observations 14,16,20

FIGURE 11: Lag weights for business-cycle filters

A. Lag weights for the BP(6,32) filter with K=12

B. Lag weights for the BP(2,32) filter with K=12

C. Lag weights for the BP(2,8) filter with K=3