

An Investigation of the Cycle Extraction Properties of Several Bandpass Filters Used to Identify Business Cycles

By

Melvin. J. Hinich

Applied Research Laboratories, University of Texas at Austin,
Austin, TX 78712-1087

Phone: +1 512 232 7270

Email: hinich@mail.la.utexas.edu

John Foster

School of Economics, University of Queensland,
St Lucia, QLD, 4072, Australia

Phone: +61 7 3365 6780

Email: j.foster@economics.uq.edu.au

and

Phillip Wild*

School of Economics and ACCS, University of Queensland,
St Lucia, QLD, 4072, Australia

Phone: +61 7 3346 9258

Email: p.wild@economics.uq.edu.au

* Address all correspondence to Phillip Wild, School of Economics, University of Queensland, St Lucia, QLD, 4072, Australia; Email: p.wild@economics.uq.edu.au, Phone: +61 7 3346 9258.

ABSTRACT

The purpose of this article is to investigate the ability of bandpass filters commonly used in economics to extract a known periodicity. The specific bandpass filters investigated include a Discrete Fourier Transform (DFT) filter, together with those proposed by Hodrick and Prescott (1997) and Baxter and King (1999). Our focus on the cycle extraction properties of these filters reflects the lack of attention that has been given to this issue in the literature, when compared, for example, to studies of the trend removal properties of some of these filters. The artificial data series we use are designed so that one periodicity deliberately falls within the passband while another falls outside. The objective of a filter is to admit the 'bandpass' periodicity while excluding the periodicity that falls outside the passband range. We find that the DFT filter has the best extraction properties. The filtered data series produced by both the Hodrick-Prescott and Baxter-King filters are found to admit low frequency components that should have been excluded.

Keywords: business cycles, bandpass filter, cycle extraction, Discrete Fourier Transform (DFT).

1 INTRODUCTION

There has been considerable attention paid to the design of various filters to extract business cycle components from macroeconomic time series. Two of the best-known are the Hodrick-Prescott (HP) and Baxter-King (BK) filters – see Hodrick and Prescott (1997) and Baxter and King (1999), respectively. These filters share a common property that they are implemented in the time domain. This contrasts with the approach commonly found in mathematical statistics and signal processing where cycle extraction algorithms are implemented in the frequency domain using Finite Fourier Transforms (FFT) methods.

While much attention has been given in the economics literature to establishing the properties of the two above-mentioned time domain filters - particularly with regard to their ability to make the time series data stationary but less attention has been given to their efficacy in successfully extracting cyclical components.

In this article, we investigate the ability of selected bandpass filters commonly used in economics to extract a known deterministic periodicity while, at the same time, requiring the filter to pass over another periodicity, deliberately designed to fall outside the passband. This is a simple task that we would reasonably expect any filtering algorithm to be able to successfully accomplish.

We investigate this capability both when the periodicity is purely deterministic and when it is embedded in stationary mixing noise. Here, we ignore the issues and implications arising when the periodicity is embedded in a trend process. If the filters cannot easily remove the appropriate components when the process is purely deterministic, or embedded in stationary Gaussian noise, then it is not likely to be successful when embedded in a trend process, irrespective of the ability of the filter to de-trend the underlying data.

2 THE DISCRETE FOURIER TRANSFORM

A linear filter, a convolution of the filter's impulse response and input signal, is a product of the complex gain of the filter and the Fourier transform of the input series $x(t)$. For a finite sample of a discrete-time series, the appropriate Fourier transform is the Discrete Fourier Transform (DFT).

Suppose that we have a finite length discrete time series $\{x(t_n)\}$ where $t_n = n\tau$, τ is the sampling interval and the record length is $T = N\tau$. In the discussion below we set $\tau = 1$ to suppress the time unit, use N and T interchangeably and set the first observation index to zero. The DFT can be viewed as mapping a sequence of N (or T) data points $\{x(0), x(1), \dots, x(N-1)\}$ in the time domain to a set of N equally spaced ordinates in the frequency domain at points $f_k = \frac{k}{N}$ termed Fourier

frequencies. The (time to frequency) DFT is the vector of complex variables $\{A_x(0), A_x(1), \dots, A_x(N/2)\}$ defined by

$$A_x(k) = \sum_{n=0}^{N-1} x(n) \exp\left(-i2\pi \frac{kn}{N}\right), \quad (1)$$

where $\exp\left(i2\pi \frac{kn}{N}\right) = \cos\left(2\pi \frac{kn}{N}\right) + i \sin\left(2\pi \frac{kn}{N}\right)$ and $i = \sqrt{-1}$. Since the observations are real variables and $\exp(i2\pi n) = 1$ for all n , $A(N-k) = A(-k) = A^*(k)$ where the asterisk denotes the complex conjugate of $A(k)$. Thus the highest Fourier frequency index is $[N/2] = N/2$ if N is even and $[N/2] = (N-1)/2$ if N is odd.

The complex representation of a sinusoidal periodicity whose period is N is of the form $x(n) = c \exp\left(i2\pi \left(\frac{k_o n}{N} + \theta\right)\right)$, where c is the amplitude of the sinusoid and θ is its phase shift. It follows from the algebra of these complex variables that the Fourier transform of this complex sinusoid is $A_x(k) = c \exp(i2\pi\theta)$ if $k = k_o$ and $A_x(k) = 0$ otherwise. The unit amplitude vector $\left\{\exp\left(i2\pi \frac{k_o n}{N}\right)\right\}$ for a given integer k_o is called a phasor of frequency $f_o = k_o / N$ (Bracewell (1978, p. 21)). The DFT values are efficiently computed using the mixed radix Fast Fourier Transform (FFT) as long as N is not a large prime number.

The inverse (frequency to time) DFT, producing the discrete-time filtered series $y(n)$ is

$$y(n) = \frac{1}{N} \sum_{k=0}^{N-1} A(k) \exp\left(i2\pi \frac{kn}{N}\right). \quad (2)$$

Thus the data vector $y(n)$ is a dot product of the Fourier transform's complex amplitudes and the vector of phasors for time index n . The difference between adjacent Fourier frequencies is determined by the fundamental frequency $f_1 = \frac{1}{N}$, which is the upper limit for resolving cycles in the finite segment of the time series. There is no way to extract more information about cycles from the sample of time length T given the mathematical fact of equation (2).

The use of the Fourier representation for bandpass filtering a time series is shown in the next section. Our approach is similar to the convolved bandpass filter method of Iacobucci and Noullez (2004, 2005) in that we also operate directly in the frequency domain.

3 BANDPASS FILTERING A TIME SERIES

Assume that we wish to analyze the periodic nature of the time series in the passband (f_{k_1}, f_{k_2}) . We therefore want to filter out the Fourier frequencies whose indices are less than k_1 and greater than k_2 . A simple way to bandpass the input using the FFT is to zero out all the complex

FFT values outside the passband. This is accomplished by applying the following ideal frequency response at the discrete set of Fourier Frequency ordinates

$$\psi(k) = \begin{cases} 1, & \text{if } k \in (k_1, k_2) \\ 0, & \text{otherwise} \end{cases}. \quad (3)$$

Applying (3) to (2), the filtered time series is

$$\tilde{y}(t_n) = \frac{1}{N} \left[\sum_{k=k_1}^{k_2} A(f_k) \exp\left(i2\pi \frac{kn}{N}\right) + \sum_{k=N-k_2}^{N-k_1} A(f_k) \exp\left(i2\pi \frac{kn}{N}\right) \right]. \quad (4)$$

The complex amplitudes of the time series in the passband (f_{k_1}, f_{k_2}) and their complex conjugates are exactly the same as for the original time series.

The filtered data series $\tilde{y}(t_n)$ in (4) will be generated by the following sequence of operations:

- 1) apply the time to frequency DFT using (1) to the source data series $x(t_n)$;
- 2) apply the ideal frequency response (3) to each Fourier Frequency ordinate of the frequency response function obtained from the first step; and
- 3) generate the filtered series $\tilde{y}(t_n)$ by applying the inverse (frequency to time) DFT using (2) to the output from the previous step.

The ideal frequency response function of the simple DFT bandpass filter yields zero values only at the Fourier frequencies in the stopband and have values equal to unity at all Fourier frequencies falling within the passband. This result can be formally demonstrated by applying the above DFT filter algorithm to a unit impulse sequence $\{x(0), x(1), x(2), \dots, x(N-1)\} = \{1, 0, 0, \dots, 0\}$. We adopt a sample size of 120 observations corresponding, for example, to 30 years of quarterly data and a passband of (6, 1.5) years or equivalently (24, 6) quarters. In terms of frequency (or inverted period), the passband is given by $(1/24, 1/6) = (0.042, 0.167)$.

The frequency response is depicted in [Figure 1](#) and is derived by applying (1) to the unit impulse series and applying the ideal frequency response in (3) at the Fourier frequency ordinates (defined by the ‘dot’ points in Figure 1). This operation is implemented in the frequency domain and it is evident from inspection of Figure 1 that the ideal frequency response outlined in (3) is synthesized at the Fourier frequency ordinates. It should be noted that while the curve in Figure 1 appears to be continuous, the frequency response function is only strictly defined at the discrete Fourier frequency ordinates themselves.

Figure 1 about here.

The application of the inverse DFT (2) to the discrete ideal frequency response outlined in Figure 1 is documented in [Figure 2](#). This figure

contains a plot of the resulting Dirichlet function, which is symmetrical about $N/2$, i.e. data point 60 in Figure 2. Once again, this function is strictly defined only at the discrete data points represented by the “dots” in Figure 2. It is not a continuous function in time.

Figure 2 about here.

The theory underpinning Fourier analysis is based on the concept of stationary time series. However, most economic time series have strong trends. Therefore, it is crucial to establish what mechanism (deterministic or stochastic) is generating the trend and propose actions to remove the trend prior to performing any DFT filtering operation. We assume that the data has been rendered stationary. It should also be noted that in the context of the artificial data simulated in later sections of this article, this is the case by construction.

The two other filters considered in this paper are those of Baxter-King (BK) and Hodrick-Prescott (HP). Because these filters are well known, we refer interested readers to the seminal articles by Baxter and King (1999) and Hodrick and Prescott (1997). Other useful references include King and Rebelo (1993), Cogley and Nason (1995), Pedersen (2001) and Iacobucci and Noullez (2004,2005). It should be noted that we employ the band-pass version of the Hodrick-Prescott filter proposed in Iacobucci and Noullez (2004, 2005). Specifically, we choose a value of λ ($=215.32$) which passes cycles of 6 years or less followed by an application in which a value of λ ($=1$) is chosen that will pass cycles of

1.5 years or less assuming quarterly data. To extract cycles of 6 to 1.5 years, we then subtract the second HP filtered series from the first HP filtered series - see Iacobucci and Noullez (2004, p. 18) for further details.¹

4 SIMULATION MODEL

We wrote a FORTRAN 95 program to conduct the reported simulations. In general terms, the artificial data model can be viewed as a periodic process that can be deterministic or embedded in stationary Gaussian noise. 'Complete' periodicity is defined as the sum of two orthogonal periodicities. The first is a low frequency periodicity that is deliberately designed to fall outside the chosen passband, while the second periodicity is designed to fall within the passband. We adopt the same parameter settings that were outlined in relation to Figure 1 in the previous section. Specifically, the sample size is 120 and the passband is defined to be (0.042, 0.167).

Recall that the objective of the bandpass filtering exercise is to admit any periodicities falling within the passband while passing over any

¹ The FORTRAN code for both the HP and BK Filters was kindly provided by Houston H Stokes, Department of Economics, University of Illinois at Chicago.

periodicities falling outside the passband range. Formally, we define the low frequency periodicity as

$$x_l(t) = amp_1 * \sin(2\pi f_1 * (t + 10)), \quad (5)$$

and the 'bandpass' periodicity is defined as

$$x_b(t) = amp_2 * \cos(2\pi f_2 * (t - 4)), \quad (6)$$

where amp_1 and amp_2 are amplitude parameters while f_1 and f_2 are frequency parameters. The complete periodicity can be represented by

$$x(t) = x_l(t) + x_b(t) + \varepsilon(t), \quad (7)$$

where $\varepsilon(t)$ is a stationary random process. In this article $\varepsilon(t)$ can take the following forms:

$$\varepsilon(t) = 0, \quad (8a)$$

where $x_l(t)$, $x_b(t)$ and $x(t)$ are deterministic periodic processes; or

$$\varepsilon(t) = u(t), \quad (8b)$$

where $u(t)$ and $\varepsilon(t)$ are stationary Gaussian noise processes, $x(t)$ is a stochastic periodic process while both $x_l(t)$ and $x_b(t)$ are deterministic periodic processes.

Figures 3-4 show plots of the artificial data series corresponding to the deterministic model (8)a for the true low frequency and bandpass periodicities obtained from applying equations (5) and (6) with two

particular parameter settings for (6). The particular parameter settings associated with both equations are listed in each figure's caption. It should be noted that the true low frequency periodicity given by $f_1 = 0.025$ is perfectly synchronized with the Fourier frequency ordinate 0.025 in the low frequency region of the stopband – see Figure 3. As such, the true low frequency periodicity is both a sub-multiple of the series length and a harmonic of the fundamental frequency.

Figure 3 about here.

The different parameter settings adopted in (6) for the true 'bandpass' periodicity reflect our desire to examine the implications on filtering of two specific circumstances. The first corresponds to the situation when the true 'bandpass' periodicity is perfectly synchronized with a Fourier frequency ordinate in the passband. In this case, the true 'bandpass' periodicity is also a harmonic of the fundamental frequency. The second circumstance is when the true 'bandpass' periodicity lies between two adjacent Fourier frequency ordinates in the passband. In this particular case, the true periodicity is not a sub-multiple of the series length or a harmonic of the fundamental frequency.

Data series corresponding to the perfectly synchronized and unsynchronized cases are listed in Figure 4. In the synchronized case, the f_2 parameter is set to 0.0667, which corresponds to a specific Fourier frequency ordinate in the passband. In the unsynchronized

case, a parameter setting for f_2 of 0.0625 was adopted that, by design, falls half way between the two adjacent Fourier frequency ordinates of 0.0583 and 0.0667. In Figure 4, the unsynchronized data series is denoted as $[xb(t): f2 = 0.0625]$ while the synchronized data series is denoted as $[xb(t): f2=0.0667]$, respectively.

Figure 4 about here.

Plots of the two complete deterministic data series, corresponding to (7)-(8)a, are documented in [Figure 5](#). The complete deterministic periodicity for the unsynchronized case is presented [i.e. the $x(t): f2=0.0625$ data series] together with the synchronized case [i.e. the $x(t): f2=0.0667$ data series].

Figure 5 about here.

The time to frequency DFT in (1) is applied to the respective ‘input’ data series $x(t_n)$ outlined in Figure 5. This is in accordance with the first stage of the DFT filter algorithm, outlined in the previous section. This data also serves as the source input data series to which both the BK and bandpass version of the HP filter are applied. The data series presented in Figure 4 are the respective targets of each filter in the next two sections. The goal of the bandpass filtering operations applied to data generated by model [(7)-(8)a] (i.e. the data presented in Figure 5) is to produce a filtered data series which either coincides exactly

(synchronized case) or closely approximates (unsynchronized case) the data series depicted in Figure 4.

The ‘input’ data associated with the stochastic model [(7)-(8)b] is presented in Figure 6 for particular choices of parameter f_2 in (6), corresponding to the synchronized and unsynchronized cases, respectively. Specifically, the complete stochastic periodicity (7) for the unsynchronized case is presented in Figure 6 [i.e. the $x(t)$: $f_2=0.0625$ series] together with the synchronized case [i.e. the $x(t)$: $f_2=0.0667$ series]. This data is the source input data to the various bandpass filters investigated. In the case of model [Eq (7)-(8)b], the target ‘bandpass’ periodicity is still the deterministic data series documented in Figure 4. However, it should be noted that the filtered data series no longer reproduces the target data series exactly because the input data is no longer purely deterministic. We expect the filtered data series to broadly track the periodic structure of the target data series documented in Figure 4.

Figure 6 about here.

5 RESULTS ASSOCIATED WITH THE SYNCHRONIZED

DETERMINISTIC PERIODIC MODEL (7)-(8)A: EQ (6): $F_2 = 0.0667$

In this section, we investigate the ability of the bandpass filters to extract the deterministic ‘bandpass’ cycle that is perfectly synchronized with a Fourier frequency ordinate in the passband. The testing

procedure involves the application of the bandpass filters to the artificial ‘synchronized’ data series presented in Figure 5. Figure 7 contains a plot of the results from the application of the DFT filter. The artificial data series associated with the true ‘bandpass’ periodicity (6) and the bandpass filtered data series from the DFT filter are displayed together. Because of the perfect synchronization of the true ‘bandpass’ periodicity and the Fourier frequency ordinate 0.0667, then, if the bandpass filtering operation has been successful, the two data series should coincide exactly. In fact, it is apparent from inspection of Figure 7 that this occurs – the filtering algorithm has successfully extracted the deterministic cycle corresponding to the perfectly synchronized ‘bandpass’ periodicity.

Figure 7 about here.

In order to confirm that the DFT filter operation has ignored the low frequency periodicity contained in (7) [subsequently generated by (5)], we make use of the periodogram of the bandpass filtered data series. This is calculated as the squared modulus (or square of the absolute value) of the complex variable $A_x(k)$ determined from (1) for each Fourier frequency k , divided by the number of sample points N . If the low frequency cycle has been expunged from the filtered data series, then there should be no ‘power’ (i.e. no non-zero value) evident at the low frequency ordinate (0.025) in the periodogram of the filtered series. The

periodogram of the DFT filtered data series, together with the periodograms of the other bandpass filtered data series are presented in [Figure 8](#). The periodogram results for the DFT filter is depicted by the ‘DFT’ series.

Figure 8 about here.

Examination of Figure 8 indicates that, in the case of the DFT filtered data, the low frequency component has been successfully expunged from the bandpass filtered data series – there is no power corresponding to Fourier frequency ordinate 0.025. In fact, the only power corresponds to the spike at Fourier frequency ordinate 0.0667 reflecting the perfect synchronization with the true ‘bandpass’ periodicity generated by (6).

The results of applying the bandpass HP filter to the data series are documented in [Figure 9](#). Recall that we adopt the same values for λ as cited in Iacobucci and Noullez (2004, p.18)) – namely, $\lambda_{low} = 215.3225$ and $\lambda_{high} = 1$. It should also be noted that, in this Figure, we have lost 16 data points from the start and end of the data series because the results from the BK filter are also represented in this Figure – the data series is now defined from data points 17 to 104 instead of 1 to 120. Once again, the artificial (target) data series associated with the true ‘bandpass’ periodicity (6) is included as a point of reference.

Figure 9 about here.

It is evident from inspection of Figure 9 that the bandpass HP filtered data series, [i.e. the 'BPHP' series], does not reproduce the exact pattern of the target 'bandpass' periodicity generated from (6). As such, the filtering operation has induced distortions and cannot be viewed as successfully extracting the true 'bandpass' periodicity. This result contrasts with the outcome depicted in Figure 7 in relation to the DFT filter.

The cause of this distortion can be found by inspecting the periodogram of the HP filtered data series presented in Figure 8 – [i.e. the 'BPHP' series]. Some leakage from the low frequency periodicity is apparent in the periodogram of the HP filtered data series. The power associated with Fourier frequency ordinate 0.025 clearly falls outside the passband range of (0.042,0.167). The other key feature is the noticeable spike at Fourier frequency ordinate 0.0667. However, the power at this ordinate is smaller than that associated with the DFT filter, reflecting the leakage associated with the presence of the low frequency periodicity in the data.

The results obtained from application of the BK filter to the artificial data series are also outlined in Figure 9 – [i.e. the 'BK(16)' series]. We followed the recommendations in Iacobucci and Noullez (2005, p87) concerning filter resolution and chose a number of MA lags corresponding to $K=16$ (in terms of Iacobucci and Noullez's notation) thus ensuring that the length of the filter is greater than the longest cycle we wish to extract (of 24 quarters).

It is apparent from examination of Figure 9 that the filtered data series obtained from the BK filter does not reproduce the exact pattern of the target bandpass periodicity generated by (6). In common with the bandpass version of the HP filter, the BK filter has also generated some distortions although, overall, these are smaller than those associated with the bandpass HP filter.

The cause of the distortion can, once again, be found by inspecting the periodogram of the BK filtered data series documented in Figure 8 – [i.e. the ‘BK’ series]. The extent of the leakage appears to be marginally less compared with the bandpass HP filter. However, the power at frequency 0.0667 seems to be comparable. To summarize, the DFT algorithm is superior in extracting the true ‘bandpass’ periodicity. The DFT filter is the only one that can completely extract the target (bandpass) periodicity. The results reported in this section were obtained for a deterministic periodic data series in which the true ‘bandpass’ periodicity was constructed to be perfectly synchronized with a Fourier frequency ordinate in the passband. However, it is possible that this is not the case in practice.

6 RESULTS ASSOCIATED WITH THE UNSYNCHRONIZED DETERMINISTIC PERIODIC MODEL (7)-(8)A: EQ (6): F2 = 0.0625

In this section, we investigate the ability of the various bandpass filters to extract a true deterministic ‘bandpass’ cycle which is not synchronized

with any Fourier frequency ordinate in the passband. The test procedure, once again, is to apply the various bandpass filters to the ‘unsynchronized’ artificial data series presented in Figure 5 and examine the extent to which the filtered data series coincides with the ‘unsynchronized’ (target) ‘bandpass’ data in Figure 4. Recall from Section 4 that, in constructing the true ‘bandpass’ periodicity using (6) for the unsynchronized case, we adopted a value for parameter f_2 of 0.0625, which was designed to fall half way between the adjacent Fourier frequency ordinates 0.0583 and 0.0667. As such, the true ‘bandpass’ periodicity falls half way between the two Fourier frequency ordinates cited above. However, the DFT for the input data series is only defined at the Fourier frequency ordinates themselves. Therefore, the true periodicity must be ‘smeared’ between the Fourier frequency ordinates that border the true periodicity.

In general, if the true periodicity falls between two Fourier frequency ordinates but is closer to one ordinate than the other, more of the contribution of the true periodicity will go to the closer Fourier frequency ordinate. If the true periodicity is approximately half way between the two Fourier frequency ordinates, then around half of the true periodicities contribution will go to each ordinate. Hence, the situation we deal with here, where the true periodicity falls half way between two, could be viewed as a worst case scenario in terms of possible deviations from the results cited in the previous section.

Figure 10 contains plots of the periodograms of the filtered data series obtained from application of the filters to the ‘unsynchronized’ data series documented in Figure 5.

Figure 10 about here.

The first thing to note in relation to the periodogram of DFT filtered data, (i.e. the “DFT” series), is that the spike associated with the synchronized case outlined in Figure 8 at frequency 0.0667 has disappeared because the true periodicity is no longer synchronized with Fourier frequency ordinate 0.0667. Instead, in Figure 10 the true periodicity has been spread over frequency ordinates 0.0583 and 0.0667. Moreover, there is also some power spread to other neighboring Fourier frequency ordinates adjacent to ordinates 0.0583 and 0.0667 in the passband region, although at a diminishing rate. Importantly, the low frequency component has been successfully expunged from the bandpass filtered data series – there is no power corresponding to Fourier frequency ordinate 0.025. The impact of ‘smearing’ distortions can be discerned from Figure 11 which contain plots of the DFT filtered data series against the unsynchronized true ‘bandpass’ periodicity.

Figure 11 about here.

It is apparent that, apart from some noticeable minor amplitude deviations at both endpoints, the filtered data series seems to track the true ‘bandpass’ periodicity remarkably well. Upon closer inspection, there

is some slight differences between the two series – for example, the ‘dot points’ for both series do not coincide exactly although they are typically very close to each other, except at the endpoints.

The results of applying the bandpass HP filter to the data series are documented in [Figure 12](#). We adopt the same values for λ as in the previous section, that is, $\lambda_{low} = 215.3225$ and $\lambda_{high} = 1$. Also observe that, because we have included the results from the BK filter in this Figure, we once again lose 16 data points from the beginning and end of the filtered data series.

Figure 12 about here.

It is evident from inspection of Figure 12 that the bandpass HP filtered data series (i.e. the ‘BPHP’ series) does not reproduce the exact pattern of the true target ‘bandpass’ periodicity generated from (6). As was the case in the previous section, the HP filtering operation has continued to generate distortions, producing significant amplitude based deviations in particular.

The cause of these distortions can be found by inspecting the periodogram of the HP filtered data series documented in Figure 10 – i.e. the ‘BPHP’ series. Some leakage from the low frequency periodicity is apparent in the periodogram of the HP filtered data series. This is particularly evident with the power associated with Fourier frequency ordinate 0.025. The other key feature is the smearing effect identified in

Figure 10 in the passband around the ordinates 0.0583 and 0.0667. In this case, however, there are qualitative differences when compared to the pattern observed for the DFT filter. Specifically, in the case of HP filter, Fourier frequency ordinate 0.0667 has significantly more relative power when compared with ordinate 0.0583.

The results obtained from application of the BK filter to the artificial data series are also presented in Figure 12 – i.e. the ‘BK(16)’ series. It is evident that the filtered data series obtained from the BK filter does not reproduce the exact pattern of the target bandpass periodicity generated by (6). In common with the bandpass HP filter, the BK filter also generates some distortions.

The cause of these can be found by inspecting the periodogram of the BK filtered data series in Figure 10 – i.e. the ‘BK’ series. Some leakage from the low frequency periodicity is evident. However, comparison of results for HP and BK filters indicate that the extent of the leakage is marginally less in the latter. This finding is consistent with the comparable results cited in the previous section. The smearing effect is evident in the bandpass region. The observed pattern around the 0.0583 and 0.0667 Fourier frequency ordinates more closely matches the qualitative pattern observed with the DFT filter compared to the bandpass HP filter.

To summarize, the substantive conclusions in the previous section continue to hold. First, the DFT algorithm is markedly superior in extracting the true ‘bandpass’ periodicity. It is the only filter that comes close to extracting the true target ‘bandpass’ periodicity, provided that we ignore some minor ‘smearing’ distortions at the endpoints of the filtered data. For the most part (i.e. away from the endpoints), these distortions are very small in magnitude. In contrast, the other two filters could not competently extract the ‘bandpass’ periodicity with continued contamination arising from the low frequency periodicity.

7 RESULTS ASSOCIATED WITH THE STOCHASTIC STATIONARY GAUSSIAN PERIODIC MODEL (7)-(8)B

In this section, we investigate the ability of the filters to extract the ‘bandpass’ periodicity when the deterministic periodicities (5)-(6) are embedded in a stationary Gaussian noise process according to the stochastic model [(7)-(8)b]. The test procedure employed is, once again, to apply the various bandpass filters to the artificial data series outlined in Figure 6 and assess the extent to which the properties of the deterministic ‘bandpass’ periodic data, outlined in Figure 4, appear to be reflected in the respective filtered data series. The embedding of (7) in stationary Gaussian noise is expected to generate some deviations in the waveform of the filtered data series when compared with the deterministic periodicity generated by (6). As such, we are now in a

world where the filtered data series can only provide an approximation to the true deterministic ‘bandpass’ periodicity. Examination of filter performance is now based on standard summary measures of goodness of fit in order to ascertain how ‘close’ the filtered series approximates the target ‘bandpass’ periodicities. Again a key consideration is whether the filtering operations generate data series in which the low frequency periodicity associated with (5) is successfully expunged.

In order to give some perspective on tracking performance, Figures 13 and 14 contain plots of the results of applying the DFT filter to the synchronized and unsynchronized input data respectively. In both figures, the data associated with the true ‘bandpass’ periodicity (6) and the bandpass filtered data from the DFT filter are displayed together.

Figure 13 about here.

Figure 14 about here

It is apparent from inspection of Figures 13 and 14 that the two filtered data series broadly track the true periodicities, although the Gaussian noise does generate some deviations in terms of both amplitude and phase variation of the DFT filtered data series, compared with the deterministic ‘bandpass’ periodicities. The tracking observed in Figures 13 and 14 is comparable to that found in earlier sections. So more plots will not add any meaningful information about comparative filter performance. Instead, conventional goodness-of-fit measures can help.

In order to assess whether or not the bandpass filter operation has ignored the low frequency periodicity generated by (5), we inspect the periodograms of the filtered data series documented in [Figures 15](#) and [16](#) for the unsynchronized and synchronized cases, respectively.

Figure 15 about here.

Figure 16 about here

In relation to the DFT filtered series, (i.e. the ‘DFT’ series), the low frequency component has been successfully expunged from both filtered data series – there is no power corresponding to Fourier frequency ordinate 0.025. For the synchronized case (Figure 16), the main power in the filtered series occurs at Fourier frequency ordinates 0.0667 and 0.1583 respectively. The first ordinate is the Fourier frequency ordinate that is synchronized with the f_2 value of 0.0667 used in (6). For the unsynchronized case, corresponding to $f_2 = 0.0625$ (Figure 15), the spike associated with the synchronized case has disappeared. Instead, the true periodicity of 0.0625 has been spread over frequency ordinates 0.0583 and 0.0667. The power identified at Fourier frequency 0.1583 is also evident in the unsynchronized case as well.

The power at the second Fourier frequency ordinate 0.1583 is most probably an artifact of the noise process which, when combined with the deterministic periodicity, produces a harmonic effect. This was not observed in Sections 5 or 6 and must, therefore, be an artifact of the

noise process and not of the underlying deterministic process. These periodicities are all clearly in the passband range. It should also be noted that, for both cases, the noise process, more generally, produces small amounts of power at all Fourier frequencies in the passband. This contrasts with the lack of ‘power’ outside of the passband. In fact, it is clear from inspection of both figures that the DFT filtered data series has no discernible power associated with any Fourier frequencies falling outside of the passband region, thus clearly encapsulating the desired ideal response in the stopband region.

In order to assess whether the bandpass HP filter has expunged the low frequency periodicity from the filtered data series, we inspect the periodograms associated with the unsynchronized and synchronized cases documented in Figures 15 and 16.

It is apparent that some leakage from the low frequency periodicity is evident in the periodograms of the HP filtered data series – i.e. the ‘BPHP’ series. There is power associated with Fourier frequency ordinate 0.025, which corresponds to the “low” frequency f_1 periodicity in (5) and which clearly falls outside the passband range. In common with the result cited for the DFT filter, there is also power in the filtered series at Fourier frequency ordinates 0.067 and 0.1583, which both fall within the passband region for the synchronized case (Figure 16). For the unsynchronized case (Figure 15), the true periodicity is once again

smearred around the Fourier frequency ordinates of 0.0583 and 0.0667. However, as found in the previous section, the Fourier frequency ordinate 0.0667 has significantly more relative power when compared with ordinate 0.0583, thus producing a different qualitative pattern to that found for the unsynchronized case. Finally, there is also power at ordinate 0.1583 for the unsynchronized case.

The results obtained from applying the BK filter to the unsynchronized and synchronized input data series are also outlined in Figures 15 and 16 - .i.e. the 'BK(16)' series. It is evident from inspection of these figures that some leakage from the low frequency periodicity is apparent in the periodogram of the BK filtered data series especially at Fourier frequency ordinate 0.025. This result was also found in both figures for the bandpass HP filter. In addition, and, in common with the evidence cited for the other filters, there is also power in the filtered series at Fourier frequency ordinates 0.067 and 0.1583 for the synchronized case. In the case of the unsynchronized case, the true periodicity is once again smearred over the ordinates 0.0583 and 0.0667 with more relative weight being given to ordinate 0.0583. Note that this pattern differs qualitatively from the patterns observed for both DFT and bandpass HP filters. In particular, the pattern does not closely follow that observed for the DFT filter, as was the case in the previous section.

In general, it was difficult to assess the relative performance of each filter by simply 'eye-balling' plots of the filtered data series relative to the

respective target periodicity. Because of the Gaussian noise in (7), all filtered data series display some amplitude and phase variation when compared with the target periodicities. However, some substantive points can be drawn, particularly from the periodogram analysis of the filtered data from the various filters. First, the DFT filter most clearly encapsulates the ideal response desired in a bandpass filter - there is no power associated with components outside the passband. This contrasts with the HP and BK filters which have some power at frequency 0.025, reflecting contamination of the filtered data series by the low frequency periodicity generated by (5). Second, the DFT filter seems to generate the most pronounced power in the passband at frequency 0.0667 for the synchronized case. This suggests that the DFT filter continues to pass on the deterministic “passband” component more completely to the filtered data series when compared to the other filters. However, the DFT filter also generates the most relative power at frequency 0.1583, although this “harmonic” effect also holds in relative terms for the other filters as well. As such, the DFT filter has the most sensitive reaction to the heightened volatility in the input data generated by the Gaussian noise variates. Third, assessment of the power of the periodogram of the BK filter, against that of the HP filter, indicates that it seems to pass less of the low frequency periodicity to the filtered data series.

Because of the apparent amplitude and phase variation in the filtered data generated by all filters considered, we use conventional goodness-

of-fit measures to get some idea of the comparative performance of the different filters. In this approach, the target periodicity is viewed heuristically as the dependent variable in a regression model while the filtered data series from the various filters is treated as the predicted value from the regression model itself. In this context, the filter that produces the best fit according to the standard goodness-of-fit criteria would be deemed to be best.

The goodness-of-fit results for the synchronized case are outlined in Table 7.1. It should be noted that the first two rows of the Table relating to the mean and standard deviation are applied to a calculated residual series, determined as the difference between the filtered data series and the true 'bandpass' periodicity. All the other statistics are calculated directly from the two respective data series.² It is apparent from inspection of Table 7.1 that the DFT based filter produces the best fit. The DFT filter (column 2) has the smallest standard deviation, largest correlation coefficient, highest coefficient of determination (R Square) and smallest standard error.

Comparison of columns 2 and 3 of Table 7.1 clearly demonstrates that the DFT filter is superior on every statistical count to the bandpass HP

² All statistics are calculated using the following worksheet functions in Excel – "AVERAGE", "STDEV", "CORREL", "RSQ" and "STEYX".

filter. We cannot strictly compare the results for the BK filter (column 4) with the results for the DFT and HP filters listed in the earlier columns in Table 7.1 because of the differences in length of data series. However, in column 5 we include the set of statistics calculated for the DFT filter with endpoint adjustments so that the dimension of the data series matches that for the BK filter. It is evident from inspection of the last two columns that the BK filter does marginally better than the DFT filter on every count except for the standard deviation. However, this result is at the expense of the loss of data points and comparison of column 2 and 5 in Table 7.1 indicate that the additional data points seem to improve the fit associated with the DFT filter on all statistical counts – in this context, the end points are not redundant.

The results for the unsynchronized data are listed in Table 7.2. It is evident that our conclusions concerning the DFT filter continue to hold on all statistical counts, as do those relating to the bandpass HP filter. In relation to the BK and adjusted DFT data listed in the last two columns of Table 7.2, it is apparent that the DFT filter now achieves the best results on all statistical counts. Moreover, examination of columns 2 and 5 of Table 7.2 indicates that the endpoint data contribute to the overall goodness-of-fit. The endpoints are not redundant data.

8 CONCLUSIONS

In this article, we have investigated the ability of bandpass filters, widely used in economics, to extract a known periodicity while passing over another periodicity that lies outside the passband. The bandpass filters employed in this article included a DFT filter together with the popular filters proposed in Hodrick and Prescott (1997) and Baxter-King (1999).

In order to investigate the cycle extraction properties of the various filters, we conducted simulations involving artificial data generated from a model of a periodic process that is deterministic or embedded in stationary Gaussian noise. The ‘complete’ periodicity was defined as the sum of two orthogonal periodicities. The first was a low frequency periodicity that was deliberately designed to fall outside the passband while the second periodicity was designed to fall within the passband.

We also distinguished between cases where the true periodicity was synchronized with a Fourier frequency ordinate in the passband. This is the case when the ‘true’ periodicity is a harmonic of the fundamental frequency or equivalently, when the true cycle is a sub-multiple of the length of the underlying data series. It was established that, under this particular circumstance, the DFT filter works optimally. The second case dealt with is when the true periodicity is not synchronized with a Fourier frequency ordinate. In this case, it was demonstrated that the true periodicity is smeared between the Fourier frequency ordinates bordering it and this effect is capable of producing distortions between the true periodicity and the filtered data series, especially at the start and end

points of the filtered data series. However, apart from the endpoints, the distortions appeared to be very small in magnitude.

The major conclusion is that the DFT algorithm is superior in extracting the true 'bandpass' periodicity, when compared with the other two filters considered. This finding is supported by both visual inspection of charts (for the purely deterministic periodicities outlined in Sections 5 and 6) and from goodness-of-fit measures for the case when the deterministic periodicities were embedded in Gaussian noise (Section 7). This general conclusion is also supported by the periodograms of the various filtered data series. In particular, inspection of the periodograms of the bandpass filtered data series, derived from the bandpass HP and BK filters, indicated contamination from low frequency periodicities that should have been excluded. It was only in the DFT case that the low frequency periodicity was completely expunged from the bandpass filtered data.

Table 7.1. Summary Statistics of Goodness of Fit of Synchronized Deterministic Target Bandpass Periodicity [Eq (8)-(9)b: $f_2 = 0.0667$] and Filtered Data Obtained From Various Bandpass Filters

	DFT	Bp HP	BK(16)	DFT(16)
Mean	0.0000	0.0000	-0.0404	-0.0123
Std Dev	0.6182	0.6778	0.6660	0.6332
Correlation	0.7276	0.6216	0.6872	0.6784
R Squared	0.5294	0.3863	0.4722	0.4602
Std Error	0.4891	0.5586	0.5240	0.5299

Table 7.2. Summary Statistics of Goodness of Fit of Unsynchronized Deterministic Target Bandpass Periodicity [Eq (8)-(9)b: $f_2 = 0.0625$] and Filtered Data Obtained From Various Bandpass Filters

	DFT	Bp HP	BK(16)	DFT(16)
Mean	-0.0419	-0.0419	-0.0392	-0.0013
Std Dev	0.6346	0.7154	0.6699	0.6333
Correlation	0.6496	0.5328	0.5895	0.6247
R Squared	0.4220	0.2839	0.3476	0.3902
Std Error	0.5412	0.6024	0.5759	0.5567

REFERENCES

- Baxter, M. and R. G. King (1999) Measuring business cycles: approximate band-pass filters for economic time series. *The Review of Economics and Statistics* 81, 575-593.
- Bracewell, R. N. (1978) *The Fourier Transform and its Applications. Second Edition*. New York: McGraw-Hill.
- Cogley, T. and J.M. Nason (1995) Effects of the Hodrick-Prescott filter on trend and difference stationary time series. Implications for business cycle research. *Journal of Economic Dynamics and Control* 19, 253-278.
- Hodrick, R. J. and E. C. Prescott (1997) Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit, and Banking* 29, 1-16.
- Iacobucci, A. and A. Noullez (2004) A frequency selective filter for short-length time series. *OFCE Working paper N 2004-5*. (Available at: <http://repec.org/sce2004/up.4113.1077726997.pdf>).
- Iacobucci, A. and A. Noullez (2005) A frequency selective filter for short-length time series. *Computational Economics* 25, 75-102.
- King, R.G. and S.T. Rebelo (1993) Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17, 207-231.
- Pedersen, T.M. (2001) The Hodrick-Prescott filter, the Slutsky effect, and the distortionary effects of filters. *Journal of Economic Dynamics and Control* 25, 1081-1101.

Figure 1. Plot of Frequency Response of Bandpass Filter for Unit Impulse - Sample Size = 120, Passband = (0.042,0.167)




Figure 2. Plot of Inverse DFT (Dirichlet Function) of Bandpass Filter for Unit Impulse - Sample Size=120, Passband = (0.042,0.167)




Figure 3. Plot of Deterministic Low Frequency Periodicity [Eq (5) and (8)a] - Sample Size = 120, amp1 = 5.0, $f_1 = 0.025$




Figure 4. Plot of Deterministic BandPass Periodicity [Eq (6) and (8)a] - Sample Size = 120, amp2=1.0, $f_2 = 0.0625$ (unsynchronized case) and $f_2=0.0667$ (synchronized case), respectively




Figure 5. Plot of Artificial Deterministic Periodicity [Eq (7) and (8)a] - Sample Size = 120, Various Eq (6): $f_2 = 0.0625$ (unsynchronized case) and $f_2 = 0.0667$ (synchronized case), respectively




Figure 6. Plot of Artificial Stochastic Periodicity [Eq (7)-(8)b] - Sample Size = 120, Various Eq (6): $f_2 = 0.0625$ (unsynchronized case) and $f_2 = 0.0667$ (synchronized case), respectively




Figure 7. Comparison of DFT BandPass Filtered Data Series From Deterministic Model [Eq (7) and Eq (8)a] and Actual (Target) Synchronized Bandpass Periodicity Data [Eq (6)] - Sample Size = 120




Figure 8. Plots of Periodograms of Bandpass Filtered Data Series Derived From Synchronized Deterministic Model [Eq (7) and Eq (8)a]: DFT, HP and BK Filters - Sample Size=120, Passband = (0.042,0.167)




Figure 9. Comparison of HP and BK BandPass Filtered Data Series From Deterministic Model [Eq (7) and Eq (8)a] and Actual (Target) Synchronized Bandpass Periodicity Data [Eq (6)] - Sample Size = 120

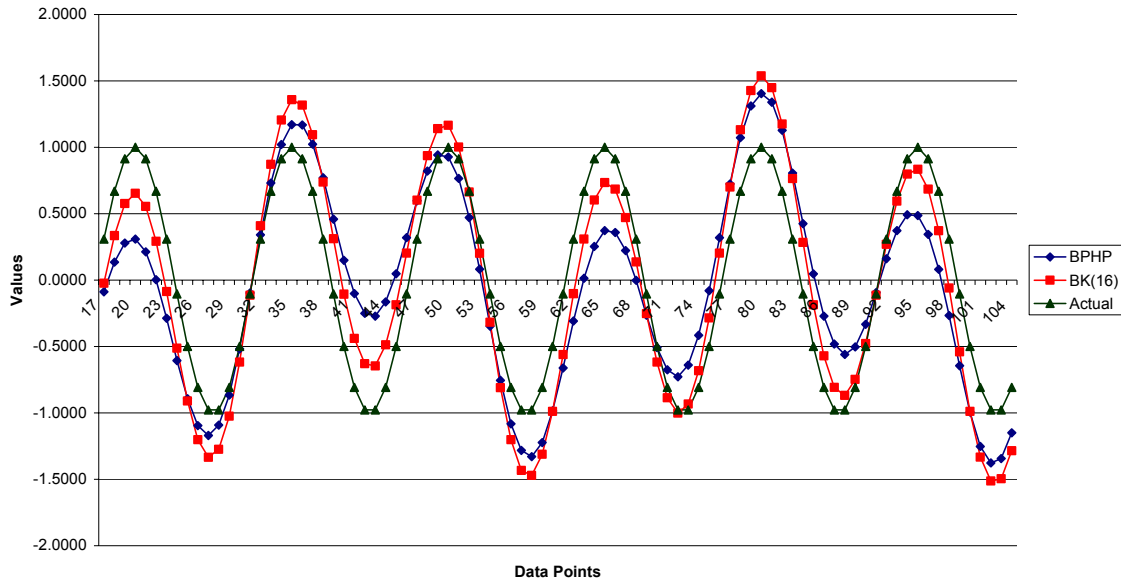


Figure 10. Plots of Periodograms of Bandpass Filtered Data Series Derived From Unsynchronized Deterministic Model [Eq (7) and Eq (8)a]: DFT, HP and BK Filters - Sample Size=120, Passband = (0.042,0.167)




Figure 11. Comparison of DFT BandPass Filtered Data Series From Deterministic Model [Eq (7) and Eq (8)a] and Actual (Target) Unsynchronized Bandpass Periodicity Data [Eq (6)] - Sample Size = 120




Figure 12. Comparison of HP and BK BandPass Filtered Data Series From Deterministic Model [Eq (7) and Eq (8)a] and Actual (Target) Unsynchronized Bandpass Periodicity Data [Eq (6)] - Sample Size = 120

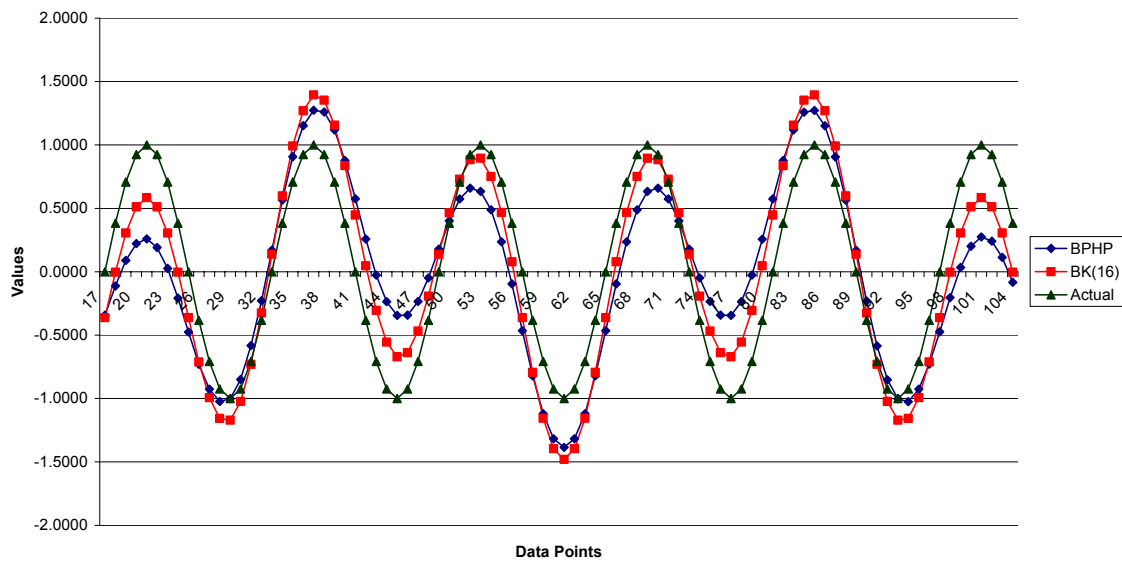


Figure 13. Comparison of DFT BandPass Filtered Data Series From Stochastic Model [Eq (7)-(8)b) and Synchronized Deterministic "Bandpass" Data [Eq (6); $f=0.0667$] - Sample Size=120




Figure 14. Comparison of DFT BandPass Filtered Data Series From Stochastic Model [Eq (7) and Eq (8)b) and Unsynchronized Deterministic "Bandpass" Periodicity [Eq (6); $f_2=0.0625$] - Sample Size=120




Figure 15. Plots of Periodograms of Unsynchronized Bandpass Filtered Data Series Derived From Stochastic Model [Eq (7) and Eq (8)b]: DFT, HP and BK Filters - Sample Size = 120, Passband = (0.042,0.167)




Figure 16. Plots of Periodograms of Synchronized Bandpass Filtered Data Series Derived From Stochastic Model [Eq (7) and Eq (8)b]: DFT, HP and BK Filters - Sample Size = 120, Passband = (0.042,0.167)

