



DEPARTMENT OF ECONOMICS

**THE REALISATION OF FINITE-SAMPLE
FREQUENCY-SELECTIVE FILTERS**

D.S.G. Pollock, University of Leicester, UK

**Working Paper No. 08/13
April 2008**

THE REALISATION OF FINITE-SAMPLE FREQUENCY-SELECTIVE FILTERS

By D.S.G. POLLOCK

University of Leicester

Email: stephen_pollock@sigmapl.u-net.com

This paper shows how a frequency-selective filter that is applicable to short trended data sequences can be implemented via a frequency-domain approach. A filtered sequence can be obtained by multiplying the Fourier ordinates of the data by the ordinates of the frequency response of the filter and by applying the inverse Fourier transform to carry the product back into the time domain. Using this technique, it is possible, within the constraints of a finite sample, to design an ideal frequency-selective filter that will preserve all elements within a specified range of frequencies and that will remove all elements outside it.

Approximations to ideal filters that are implemented in the time domain are commonly based on truncated versions of the infinite sequences of coefficients derived from the Fourier transforms of rectangular frequency response functions. An alternative to truncating an infinite sequence of coefficients is to wrap it around a circle of a circumference equal in length to the data sequence and to add the overlying coefficients. The coefficients of the wrapped filter can also be obtained by applying a discrete Fourier transform to a set of ordinates sampled from the frequency response function. Applying the coefficients to the data via circular convolution produces results that are identical to those obtained by a multiplication in the frequency domain, which constitutes a more efficient approach.

Key words: Linear filtering, Frequency-domain analysis

JEL Classification: C22

Date: April, 2008

1. Introduction: The Problem of the Ideal Filter

Recently, business cycle analysts have become interested in extracting, from macroeconomic indices, data components that fall within specified intervals of the frequency spectrum. Examples are to be found in the papers of Baxter and King (1999), Christiano and Fitzgerald (2003) and Iacobucci and Noullez (2005). In particular, Baxter and King have proposed that, according to the definition of Burns and Mitchell (1946), the business cycles should comprise cyclical elements with durations of no less than 18 months and of no more than 8 years.

It is commonly believed that, in the case of finite-length samples, it is impossible to design a filter that will preserve completely all elements within a specified range of frequencies and that will remove all elements outside it. A filter that would achieve such an objective is described as an ideal filter.

This belief is based on the fact that, when a (classical) Fourier transform is applied to a periodic square wave or boxcar function, representing the ideal

frequency response of the filter, the result is a symmetric doubly-infinite sequence of filter coefficients. To obtain a practical filter, it seems that one must truncate the sequence, retaining only a limited number of its central elements (Figure 1).

The truncation gives rise to a filter of which the frequency response has certain undesirable characteristics. In particular, there is a ripple effect whereby the gain of the filter fluctuates within the pass band, where it should be constant with a unit value, and within the stop band, where it should be zero-valued. Within the stop band, there is a corresponding problem of leakage whereby the truncated filter transmits elements that ought to be blocked (Figure 2).

The classical approach to these problems, which has been pursued by electrical engineers, has been to modulate the truncated filter sequence with a so-called window sequence, which applies a gradual taper to the higher-order filter coefficients. (A full account of this has been given by Pollock, 1999.) The effect is to suppress the leakage that would otherwise occur in regions of the stop band that are remote from the regions where the transitions occur between stop band and pass band. The detriment of this approach is that it exacerbates the extent of the leakage within the transition regions (Figure 3).

The purpose of this paper is to show that none of the above-mentioned problems need afflict the filtering of finite data sequences. It shows that an alternative to truncating the filter is to wrap the infinite sequence of coefficients around a circle of a circumference T equal to the length of the data sequence. The overlying coefficients are added to give the coefficients of the wrapped filter.

It is impractical to perform the operation of filter wrapping in the time domain by summing the infinite sequences of the overlying coefficients directly. Instead, one may resort to the equivalent operation of sampling the frequency response function of the filter at T equally-spaced points. The coefficients of the wrapped filter may be obtained by applying a discrete Fourier transform to this frequency-domain sample to carry its effects into the time domain.

The wrapped filter can be applied, via an ordinary convolution, to a periodic extension of the data sequence. Alternatively, it can be applied, via circular convolution, to the ordinary data sequence, with the same results. However, such a convolution can be realised most effectively via an equivalent modulation in the frequency domain of the Fourier transform of the data, followed by an inverse Fourier transform to carry the results back to the time domain.

This implementation of an ideal filter is but one instance of a general approach to the problem of finite-sample filter design, which we shall expound in this paper, that recognises the finite nature of the data sample at the outset. Other approaches begin, in effect, with the assumption of a doubly-infinite sample; and then they makes amends for the fact that the sample is finite by resorting to a variety of ingenious adaptations.

In order to pursue the circular approach successfully, it is necessary to detrend the data and to ensure that there are no radical disjunctions in the periodic extension of the differenced data where the end of one replication of the sequence meets the beginning of the next replication.

The data can be detrended by differencing. Once the relevant components have been extracted from the differenced data, the corresponding components of the trended data can be recovered by a process of anti-differencing, or cumulation,




Figure 1. The central coefficients of the Fourier transform of the frequency response of an ideal lowpass filter with a cut-off point at $\omega = \pi/2$. The sequence of coefficients extends indefinitely in both directions.




Figure 2. The frequency response of a filter obtained by applying a 17-point rectangular window to the coefficients of an ideal lowpass filter with a cut-off point at $\omega = \pi/2$, superimposed upon the frequency response of the ideal filter.




Figure 3. The frequency response of a filter obtained by applying a 17-point Blackman window to the coefficients of an ideal lowpass filter with a cut-off point at $\omega = \pi/2$.

that requires some initial conditions. These are readily available. In the case of highpass or bandpass filtering, the cumulation process can be avoided. If the cumulation operator is cancelled with the differencing operator that is embodied by the filter, then a reduced filter is derived that will deliver the required product directly.

This paper has a frequency-domain orientation. Reference to the frequency domain is becoming increasingly common amongst statisticians and econometricians. Thus, for example, Haywood and Tunnicliffe Wilson (1997) and Proietti (2005) have recently devised modified lowpass filters in reference to their effects in the frequency domain.

There is clear evidence that the central statistical agencies, which are responsible for producing seasonally adjusted data series and for estimating the trends in official statistics, are placing increasing emphasis in the frequency domain. Examples from the U.S. Census Bureau are provided by the recent papers of Findley and Martin (2003) and of Bell and Martin Bell (2004).

Amongst Europeans, the SEATS–TRAMO program for the canonical analysis of unobserved components in time series has been influential in fostering a growing awareness of the frequency domain (see Caporello and Maravall 2004).

2. Approximations to the Ideal Filter

The theory underlying the spectral analysis of statistical time series deals preponderantly with sequences that are defined over the entire set of positive and negative integers. Such a sequence, which may be denoted by $y(t) = \{y_t; t = 0, \pm 1, \pm 2, \dots\}$, can be described in the frequency domain as a linear combination of trigonometrical functions of which the frequencies, denominated in radians per sampling interval, range from zero to the limiting Nyquist value of π .

An infinite sequence generated by a stationary stochastic process is liable to be expressed as a weighted integral of a non denumerable set of sines and cosines indexed by a frequency value ω that varies continuously within the interval $[0, \pi]$. Since

$$\cos(\omega t) = \frac{1}{2} (e^{i\omega t} + e^{-i\omega t}) \quad \text{and} \quad \sin(\omega t) = \frac{i}{2} (e^{i\omega t} - e^{-i\omega t}), \quad (1)$$

the value generated at time t can also be expressed as a weighted integral over the interval $[-\pi, \pi]$ of a complex exponential function $\exp\{i\omega t\}$:

$$y_t = \int_{-\pi}^{\pi} e^{i\omega t} dZ(\omega). \quad (2)$$

Here, the complex element $dZ(\omega)$, which constitutes the stochastic weighting function, represents the infinitesimal increments of a cumulative function $Z(\omega)$ that is everywhere continuous but nowhere differentiable. The expectation of the squared modulus of $dZ(\omega)$ constitutes an increment of the cumulative spectrum: $dF(\omega) = E\{dZ(\omega)dZ^*(\omega)\}$. In the case of a purely stochastic process, the cumulative spectrum $F(\omega)$ is an analytic function of which the derivative is $f(\omega)$ is described as the spectral density function or the “spectrum”.

A time-invariant linear filter forms a weighted combination of adjacent elements of the sequence $y(t)$. The filter is defined by the sequence of these weights or

filter coefficients, which is the impulse response of the filter. Its effect can also be represented by the manner in which it alters the sinusoidal elements of which $y(t)$ is composed.

Mapping a (doubly-infinite) cosine sequence $y(t) = \cos(\omega t)$, of a given frequency ω , through a filter defined by the coefficients $\{\phi_k\}$ produces the output

$$\begin{aligned} x(t) &= \sum_k \phi_k \cos(\omega[t - k]) \\ &= \sum_k \phi_k \cos(\omega k) \cos(\omega t) + \sum_k \phi_k \sin(\omega k) \sin(\omega t) \\ &= \alpha \cos(\omega t) + \beta \sin(\omega t) = \lambda \cos(\omega t - \theta), \end{aligned} \tag{3}$$

where $\alpha = \sum_k \phi_k \cos(\omega k)$, $\beta = \sum_k \phi_k \sin(\omega k)$, $\lambda^2 = \alpha^2 + \beta^2$ and $\theta = \tan^{-1}(\beta/\alpha)$. These results follow in view of the trigonometrical identity $\cos(A - B) = \cos(A) \cos(B) + \sin(A) \sin(B)$.

The effect of the filter is to alter the amplitude of the cosine via the gain factor λ and to induce a delay that corresponds to the phase angle θ . It is apparent that, if the filter is symmetric about the central coefficient ϕ_0 , with $\phi_{-k} = \phi_k$, then $\beta = \sum_k \phi_k \sin(\omega k) = 0$ and, therefore, $\theta = 0$. That is to say, a symmetric filter that looks equally forward and backwards in time has no phase effect.

The z -transform of the sequence of filter coefficients is the polynomial

$$\phi(z) = \sum_k \phi_k z^k, \tag{4}$$

wherein z stands for a complex number. Setting $z = \exp\{-i\omega\} = \cos(\omega) - i \sin(\omega)$ constrains this number to lie on the unit circle in the complex plane. The resulting function

$$\begin{aligned} \phi(\exp\{-i\omega\}) &= \sum_k \phi_k \cos(\omega k) - i \sum_k \phi_k \sin(\omega k) \\ &= \alpha(\omega) - i\beta(\omega) \end{aligned} \tag{5}$$

is the frequency response function, which is, in general, a periodic complex-valued function of ω with a period of 2π . In the case of a symmetric filter, it becomes a real-valued and even function, which is symmetric about $\omega = 0$. When the frequency response function is defined over the interval $[-\pi, \pi)$, or equally over the interval $[0, 2\pi)$, it conveys all of the information concerning the gain and the phase effects of the filter.

For a more concise notation, we may write $\phi(\omega)$ in place of $\phi(\exp\{-i\omega\})$. This allows us to denote the frequency response by

$$\phi(\omega) = |\phi(\omega)|e^{-i\theta(\omega)}, \quad \text{where} \quad |\phi(\omega)| = \sqrt{\alpha^2(\omega) + \beta^2(\omega)}. \tag{6}$$

Here, $|\phi(\omega)|$ denotes the gain the filter at the frequency ω , whereas $\theta(\omega)$ indicates the phase effect. In the case of a symmetric filter, there is $\theta(\omega) = 0$ and $\phi(\omega) = \sum_k \phi_k \cos(\omega k)$; and $\exp\{-i\theta(\omega)\}$ is evaluated as either $+1$ or -1 , according to the sign of $\phi(\omega)$.

An ideal frequency-selective filter has the effect of nullifying all trigonometric sequences of which the frequencies fall within the stop band and of preserving, without alteration, all those of which the frequencies fall within the pass band.

The ideal phase-neutral lowpass filter with a cut-off at frequency $\omega = \alpha$ has the following frequency response over the interval $[-\pi, \pi]$:

$$\phi(\omega) = \begin{cases} 1, & \text{if } |\omega| \in (0, \alpha), \\ 1/2, & \text{if } \omega = \pm\alpha, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Here, the gain of the filter coincides with its frequency response. The coefficients of a filter may be obtained via the (inverse) Fourier transform of $\phi(\omega)$. In the case of the ideal filter, they are given by the sampled ordinates of a sinc function:

$$\phi_k = \frac{1}{2\pi} \int_{-\alpha}^{\alpha} e^{i\omega k} d\omega = \begin{cases} \alpha, & \text{if } k = 0; \\ \frac{\sin(\alpha k)}{\pi k}, & \text{if } k \neq 0. \end{cases} \quad (8)$$

The coefficients constitute a doubly-infinite sequence. Figure 1 shows the central coefficients of the ideal lowpass filter with a cut-off frequency of $\alpha = \pi/2$.

The coefficients of a bandpass filter with a gain of unity within the interval $[\alpha, \beta]$ and a gain of zero outside the interval are given by $\phi_k = \{\sin(\beta k) - \sin(\alpha k)\} / \pi k$ when $k \neq 0$ together with $\phi_0 = \beta - \alpha$. This is just the difference of two lowpass filters. The sum of the coefficients of a bandpass filter is zero.

In practice, all data sequences are finite, and it is impossible to apply a filter that has an infinite number of coefficients. However, a practical filter may be obtained by selecting a limited number of the central coefficients of an ideal infinite-sample filter. In the case of a truncated filter based on $2q + 1$ central coefficients, the elements of the filtered sequence are given by

$$\begin{aligned} x_t = & \phi_q y_{t-q} + \phi_{q-1} y_{t-q+1} + \cdots + \phi_1 y_{t-1} + \phi_0 y_t \\ & + \phi_1 y_{t+1} + \cdots + \phi_{q-1} y_{t+q-1} + \phi_q y_{t+q}. \end{aligned} \quad (9)$$

Given a sample y_0, y_1, \dots, y_{T-1} of T data points, only $T - 2q$ processed values $x_q, x_{q+1}, \dots, x_{T-q-1}$ are available, since the filter cannot reach the ends of the sample, unless it is extrapolated.

If the coefficients of the truncated bandpass or highpass filter are adjusted so that they sum to zero, then the z -transform polynomial $\phi(z)$ of the coefficient sequence will contain two roots of unit value. The adjustments may be made by subtracting $\sum_k \phi_k / (2q + 1)$ from each coefficient. The sum of the adjusted coefficients is $\phi(1) = 0$, from which it follows that $1 - z$ is a factor of $\phi(z)$. The condition of symmetry, which is that $\phi(z) = \phi(z^{-1})$, implies that $1 - z^{-1}$ is also a factor. Thus, the polynomial contains the factor

$$(1 - z)(1 - z^{-1}) = -z^{-1}(1 - z)^2, \quad (10)$$

within which $\nabla^2(z) = (1 - z)^2$ corresponds to the square of the difference operator.

Since it incorporates the factor $\nabla^2(z)$, the effect of applying the filter to a data sequence with a linear trend will be to produce an untrended sequence with a zero mean. The effect of applying it to a sequence with a quadratic trend will be to produce an untrended sequence with a nonzero mean. Such filters have been used by Baxter and King (1999) in extracting the business cycle from strongly trended aggregate economic indices.

It is possible to remove $\nabla^2(z)$ from $\phi(z) = \psi(z)\nabla^2(z)$. Then, the corresponding differencing operator can be applied to the data with the aim of reducing it to stationarity before applying a reduced filter, of which $\psi(z)$ is the z -transform. However, when the computations are wholly within the time domain, such an approach has no practical advantage over the approach that applies the symmetric filter $\phi(z)$ directly to the undifferenced data.

The usual effect of the truncation will be to cause a considerable spectral leakage. Thus, if the filter is applied to trended data, then it is liable to transmit some powerful low-frequency elements that will give rise to cycles of high amplitudes within the filtered output.

An alternative filter that is designed to reach the ends of the sample has been proposed by Christiano and Fitzgerald, (2003). The filter is described by the equation

$$x_t = Ay_0 + \phi_t y_0 + \cdots + \phi_1 y_{t-1} + \phi_0 y_t + \phi_1 y_{t+1} + \cdots + \phi_{T-1-t} y_{T-1} + B y_{T-1}. \quad (11)$$

This equation comprises the entire data sequence y_0, \dots, y_{T-1} ; and the value of t determines which of the coefficients of the infinite-sample filter are involved in producing the current output. Thus, the value of x_0 is generated by looking forwards to the end of the sample, whereas the value of x_{T-1} is generated by looking backwards to the beginning of the sample.

If the process generating the data is stationary, then it is appropriate to set $A = B = 0$, which is tantamount to approximating the extra-sample elements by zeros. In the case of a data sequence that appears to follow a first-order random walk, it has been proposed to set A and B to the values of the sums of the coefficients that lie beyond the span of the data on either side. Since the filter coefficients must sum to zero, it follows that

$$A = -\left(\frac{1}{2}\phi_0 + \phi_1 + \cdots + \phi_t\right) \quad \text{and} \quad B = -\left(\frac{1}{2}\phi_0 + \phi_1 + \cdots + \phi_{T-t-1}\right). \quad (12)$$

The effect is tantamount to extending the sample at either end by constant sequences comprising the first and the last sample values respectively. For data that have the appearance of having been generated by a first-order random walk with a constant drift, it is appropriate to extract a linear trend before filtering the residual sequence. In fact, this has proved to be the usual practice in most circumstances.

Christiano and Fitzgerald (1993) have also proposed another time-varying filter that is intended to be a superior approximation to the ideal bandpass filter. The coefficients of this filter, which may be denoted by $\phi_j^{(t)}$ and which are determined for each value of t , are comprised by the equation

$$x_t = \phi_t^{(t)} y_0 + \cdots + \phi_1^{(t)} y_{t-1} + \phi_0^{(t)} y_t + \phi_1^{(t)} y_{t+1} + \cdots + \phi_{T-1-t}^{(t)} y_{T-1}. \quad (13)$$

Let $\phi(z)$ be the z -transform of the coefficients of the ideal infinite-sample filter and let $\phi^{(t)}(z)$ be the z -transform of the finite-sample filter for time t . Setting $z = \exp\{-i\omega\}$ gives the frequency response functions of the filters. It is proposed that the coefficients of the finite-sample filter should be the values that jointly minimise the function

$$\int_{-\pi}^{\pi} |\phi(e^{-i\omega}) - \phi^{(t)}(e^{-i\omega})|^2 f(\omega) d\omega, \quad (14)$$

where $f(\omega)$ is the spectral density function of the process generating the data.

The intention of this criterion is to minimise the discrepancy between the finite-sample filter and the ideal filter in those regions of the frequency domain, indicated by the values of $f(\omega)$, where it matters most. However, given that a data sequence of T elements is represented in the frequency domain by a set of complex exponential functions defined on T frequency values, described as the Fourier frequencies, there is no cause for assessing the discrepancy at every frequency in the interval $[\pi, \pi]$.

Moreover, it is possible to devise an ideal finite-sample filter that eliminates the discrepancy completely at the T Fourier frequencies. To demonstrate this point, it is necessary to consider the discrete Fourier transform of the finite data sequence.

3. The Discrete Fourier Transform

The discrete Fourier transform is a one-to-one mapping from a set of T data points to a set of T coefficients associated with a set of harmonically related trigonometric functions. The vectors of the ordinates sampled from the trigonometric functions constitute an orthogonal basis of the T -dimensional space that contains the data vector.

The inverse Fourier transform, which is a mapping from the coefficients to the data, gives rise to the following equation, which describes the Fourier synthesis of the data:

$$y_t = \sum_{j=0}^{[T/2]} \{\alpha_j \cos(\omega_j t) + \beta_j \sin(\omega_j t)\}; \quad t = 0, 1, \dots, T-1. \quad (15)$$

Here, $[T/2]$ denotes the integer quotient of the division of T by 2. The harmonically related Fourier frequencies $\omega_j = 2\pi j/T$; $j = 0, \dots, [T/2]$, which are equally spaced in the interval $[0, \pi]$, are integer multiples of the fundamental frequency $\omega_1 = 2\pi/T$, which relates to a sinusoidal function that completes a single cycle in the time spanned by the sample. A stochastic nature is imparted to y_t by the coefficients α_j, β_j , which are to be regarded as random variables

The temporal index t of the above equation ranges from 0 to $T-1$. However, strictly for analytic purposes, we may regard the data sequence as a single cycle of a periodic function defined over the entire set of positive and negative integers, which is described as the periodic extension the data. Considering the periodic extension does not entail making any assumption that the data have been generated by an underlying periodic process.

For mathematical convenience, we may express the trigonometric functions of (15) in terms of complex exponential functions:

$$\cos(\omega_j t) = \frac{1}{2} (e^{i\omega_j t} + e^{-i\omega_j t}), \quad \sin(\omega_j t) = \frac{i}{2} (e^{i\omega_j t} - e^{-i\omega_j t}). \quad (16)$$

Then, on defining

$$\zeta_j = \alpha_j + i\beta_j \quad \text{and} \quad \zeta_{T-j} = \alpha_j - i\beta_j, \quad (17)$$

equation (15) can be written as

$$y_t = \sum_{j=0}^{T-1} \zeta_j e^{i\omega_j t} = \sum_{j=0}^{T-1} \zeta_j W^{jt}; \quad t = 0, 1, \dots, T-1, \quad (18)$$

where $W^{jt} = \exp\{2\pi jt/T\}$. Here, W^q is a T -period function of q such that $W \uparrow q = W \uparrow (q \bmod T)$, where the upward arrow signifies exponentiation. The complex values $W^0 = 1, W, W^2, \dots, W^{T-1}$, which describe one cycle of the function, are known as the T roots of unity. Figure 4 shows them inscribed on the circumference of the unit circle in the case of $T = 8$.

Using the exponential notation, the Fourier transform and its inverse can be denoted by

$$\zeta_j = \frac{1}{T} \sum_{t=0}^{T-1} y_t e^{-i\omega_j t} dt \quad \longleftrightarrow \quad y_t = \sum_{j=0}^{T-1} \zeta_j e^{i\omega_j t}. \quad (19)$$

For a matrix representation of these transforms, one may define

$$\begin{aligned} U &= T^{-1/2} [\exp\{-i2\pi tj/T\}; t, j = 0, \dots, T-1], \\ \bar{U} &= T^{-1/2} [\exp\{i2\pi tj/T\}; t, j = 0, \dots, T-1], \end{aligned} \quad (20)$$

which are unitary complex matrices such that $U\bar{U} = \bar{U}U = I_T$. Then,

$$\zeta = T^{-1/2} U y \quad \longleftrightarrow \quad y = T^{1/2} \bar{U} \zeta. \quad (21)$$

where $y = [y_0, y_1, \dots, y_{T-1}]'$ and $\zeta = [\zeta_0, \zeta_1, \dots, \zeta_{T-1}]'$ are the vectors of the data and of their spectral ordinates, respectively.

Observe that, under the assumption that the data are generated by a stationary stochastic process, the limiting form of equation (18), as $T \rightarrow \infty$, is equation (2).

4. The Ideal Finite-Sample Filter

In terms of the frequency domain, the process of filtering a finite data sequence consists of altering the values of the spectral ordinates within the vector $\zeta = [\zeta_0, \zeta_1, \dots, \zeta_{T-1}]'$. These ordinates, which correspond to the Fourier frequencies

$$\omega_j = 2\pi j/T; \quad j = 0, 1, 2, \dots, T-1, \quad (22)$$

may be envisaged as a set of spikes erected on the circumference of the unit circle in the complex plane at locations that are indicated by the T roots of unity

$$W^j = \exp\{-i\omega_j\} = \cos(\omega_j) - i \sin(\omega_j); \quad j = 0, 1, 2, \dots, T-1. \quad (23)$$




Figure 4. The 8th roots of unity. The diagram on the left represents the frequency response of the ideal half-band lowpass filter at the Fourier frequencies corresponding to the angles $\omega_j = 2\pi j/8; j = 0, 1, \dots, 7$. Unit responses are represented by black dots, zero responses by circles and the responses at the transition points by encircled dots. The diagram on the right shows the values of the cosine function at those angles.

The frequency response at ω_j , determined in accordance with the ideal specification of (7), is

$$\lambda_j = \phi(\omega_j) = \sum_{k=-\infty}^{\infty} \phi_k W^{jk}. \quad (24)$$

Since W^q is a T -periodic function, it follows that

$$\begin{aligned} \lambda_j &= \left\{ \sum_{q=-\infty}^{\infty} \phi_{qT} \right\} + \left\{ \sum_{q=-\infty}^{\infty} \phi_{qT+1} \right\} W^j + \dots + \left\{ \sum_{q=-\infty}^{\infty} \phi_{qT+T-1} \right\} W^{j(T-1)} \\ &= \phi_0^\circ + \phi_1^\circ W^j + \dots + \phi_{T-1}^\circ W^{j(T-1)}, \quad \text{for } j = 0, 1, 2, \dots, T-1. \end{aligned} \quad (25)$$

These equations serve to determine the circular filter coefficients $\phi_0^\circ, \phi_1^\circ, \dots, \phi_{T-1}^\circ$.

Let $\phi^\circ = [\phi_0^\circ, \phi_1^\circ, \dots, \phi_{T-1}^\circ]'$ be the vector of the coefficients of a circular filter and let $\lambda = [\lambda_0, \lambda_1, \dots, \lambda_{T-1}]'$ be the vector of the values of the frequency response at the Fourier frequencies. Then, in terms of the matrices of (20), the mapping from ϕ° to λ and the corresponding inverse mapping can be represented by

$$\lambda = T^{-1/2} U \phi^\circ \quad \longleftrightarrow \quad \phi^\circ = T^{1/2} \bar{U} \lambda. \quad (26)$$

The filtering operation can be performed by multiplying the spectral ordinates within the vector ζ by the weights within λ . The resulting vector may be transformed from the frequency domain to the time domain to produce the filtered output.

Let $\Lambda = \text{diag}\{\lambda_0, \lambda_1, \dots, \lambda_{T-1}\}$ be the diagonal matrix of the weights. Then, with reference to (21), it can be seen that the weighted values of the spectral ordinates

are given by the vector

$$\Lambda\zeta = T^{-1/2}\Lambda Uy. \quad (27)$$

Subjecting this vector to the inverse Fourier transform gives the filtered output

$$x = T^{1/2}\bar{U}\Lambda\zeta = \{\bar{U}\Lambda U\}y = \Phi^\circ y, \quad (28)$$

where $\Phi^\circ = \bar{U}\Lambda U$ is the matrix of the filtering operation in the time domain. The next section is devoted to revealing the nature of this matrix and of the associated time-domain filtering operation.

5. Filtering via Circular Convolution

Consider the following matrix equation:

$$\begin{aligned} & \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & W & 0 & 0 \\ 0 & 0 & W^2 & 0 \\ 0 & 0 & 0 & W^3 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & W & W^2 & W^3 \\ 1 & W^2 & W^4 & W^6 \\ 1 & W^3 & W^6 & W^9 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & W & W^2 & W^3 \\ 1 & W^2 & W^4 & W^6 \\ 1 & W^3 & W^6 & W^9 \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ W & W^2 & W^3 & 1 \\ W^2 & W^4 & W^6 & 1 \\ W^3 & W^6 & W^9 & 1 \end{bmatrix} \end{aligned} \quad (29)$$

The first equality can be represented in summary notation by

$$DU = UK. \quad (30)$$

This example is readily generalised to encompass matrices of any order.

In general, $K = [e_1, e_2, \dots, e_{T-1}, e_0]$ is a circulant matrix operator that is formed from the identity matrix $I = [e_0, e_1, e_2, \dots, e_{T-1}]$ by moving the leading vector to the back of the array, whereas $D = \text{diag}\{1, W, W^2, \dots, W^{T-1}\}$ is a diagonal matrix containing the roots of unity. The remaining matrix U is in accordance with the definitions of (20).

The following conditions hold for the circulant operator:

$$\begin{aligned} & \text{(i)} \quad K^{-q} = K^{T-q}, \\ & \text{(ii)} \quad K^0 = K^T = I, \\ & \text{(iii)} \quad K' = K^{T-1} = K^{-1}. \end{aligned} \quad (31)$$

From (30), it follows that

$$\begin{aligned} & \text{(i)} \quad K = \bar{U}DU = U\bar{D}\bar{U}, \\ & \text{(ii)} \quad K' = K^{-1} = UD\bar{U} = \bar{U}\bar{D}U, \\ & \text{(iii)} \quad D = UK\bar{U}, \\ & \text{(iv)} \quad \bar{D} = \bar{U}KU, \end{aligned} \quad (32)$$

where $\bar{D} = \text{diag}\{1, W^{T-1}, W^{T-2}, \dots, W\}$ is both the complex conjugate and the inverse of D . More extensive accounts of this algebra have been given by Pollock (1999), (2002).

Recall that $\Lambda = \text{diag}\{\lambda_0, \lambda_1, \dots, \lambda_{T-1}\}$ is a diagonal matrix containing values sampled at equal intervals from the frequency response function of a linear filter. Let $\iota = [1, 1, \dots, 1]'$ be the summation vector of order T , and observe that $\Lambda \iota = \lambda = [\lambda_0, \lambda_1, \dots, \lambda_{T-1}]'$ and that $T^{-1/2} \iota = U e_0$, where $e_0 = [1, 0, \dots, 0]'$. In terms of this notation, the second equation of (26) can be written as

$$\phi^\circ = \bar{U} \Lambda U e_0 = \Phi^\circ e_0. \quad (33)$$

Premultiplying by the circulant operator $K = \bar{U} D U$ gives

$$\begin{aligned} K \phi^\circ &= \bar{U} D (U \bar{U}) \Lambda U e_0 & [U \bar{U} &= I] \\ &= \bar{U} (D \Lambda) U e_0 & [D \Lambda &= \Lambda D] \\ &= \bar{U} \Lambda (D U) e_0 & [D U &= U K] \\ &= \bar{U} \Lambda U (K e_0) & [K e_0 &= e_1] \\ &= \bar{U} \Lambda U e_1. \end{aligned} \quad (34)$$

More generally, there is $K^q \phi^\circ = \bar{U} \Lambda U e_q$. Letting $q = 0, 1, \dots, T-1$ and gathering the resulting vectors in a matrix array creates the following circulant matrix:

$$\Phi^\circ = [\phi^\circ, K \phi^\circ, K^2 \phi^\circ, \dots, K^{T-1} \phi^\circ] = \bar{U} \Lambda U. \quad (35)$$

This is the symmetric real-valued circulant matrix of which the general form is adequately represented by the case where $T = 4$:

$$\Phi^\circ = \begin{bmatrix} \phi_0^\circ & \phi_1^\circ & \phi_2^\circ & \phi_1^\circ \\ \phi_1^\circ & \phi_0^\circ & \phi_1^\circ & \phi_2^\circ \\ \phi_2^\circ & \phi_1^\circ & \phi_0^\circ & \phi_1^\circ \\ \phi_1^\circ & \phi_2^\circ & \phi_1^\circ & \phi_0^\circ \end{bmatrix}. \quad (36)$$

To summarise these results, we observe that the factorisation $K = \bar{U} D U$ of the circulant operator indicates that

$$\Phi^\circ = \phi^\circ(K) = \bar{U} \phi^\circ(D) U = \bar{U} \Lambda U, \quad (37)$$

where $\Phi^\circ = \phi^\circ(K)$ is obtained by replacing z by K in the z -transform $\phi^\circ(z) = \phi_0^\circ + \phi_1^\circ z + \phi_2^\circ z^2 + \dots + \phi_{T-1}^\circ z^{T-1}$ and $\Lambda = \phi^\circ(D)$ is obtained by replacing z by D . Moreover, the j th diagonal element of $\phi^\circ(D)$ is just $\phi^\circ(W^j) = \phi(W^j)$; which is to say that it is an ordinate of the frequency response sampled at the j th Fourier frequency $\omega_j = 2\pi j/T$. Thus, the coefficients of the wrapped filter are obtained by carrying into the time domain, via the inverse Fourier transform, a sample of T ordinates of the frequency response function.

Applying the filter matrix Φ° to the data vector $y = [y_0, y_1, \dots, y_{T-1}]'$ gives

$$x = \Phi^\circ y = \bar{U} \Lambda U y. \quad (38)$$

This equation indicates that there are two ways of forming the filtered vector x . The first way is via the circular convolution in the time domain of the vector ϕ° of the filter coefficients and the vector y of the data. To elucidate this operation, define the circulant data matrix $Y = [y, Ky, K^2y, \dots, K^{T-1}y]$ and observe that circulant matrices commute in multiplication. It follows that

$$\Phi^\circ y = \Phi^\circ Y e_0 = Y \Phi^\circ e_0 = Y \phi^\circ. \quad (39)$$

This expression manifests a symmetry that puts the data vector y and the filter vector ϕ° on an equal footing.

The second way of obtaining the filtered output is via the Fourier transform and its inverse. First, the discrete Fourier transform is applied to the data vector to carry it into the frequency domain. Then a differential weighting is applied to the elements of the resulting vector $\zeta = Uy$ to give $\Lambda\zeta = \Lambda Uy$. (In the case of the ideal frequency-selective filter of (7), the weights, which are the diagonal elements of Λ , are units in the pass band and zeros in the stop band; and there are ordinates with a value one half on the points of transition.) Finally, the filtered vector $x = \bar{U}\Lambda Uy$ is obtained by applying the inverse Fourier transform.

6. The Finite-Sample Frequency Response

It should be emphasised that $\phi(z)$ and $\phi^\circ(z)$ are distinct functions of which, in general, the values coincide only at the roots of unity. Strictly speaking, these are the only points for which the latter function is defined. Nevertheless, there may be some interest in discovering the values that $\phi^\circ(z)$ would deliver if its argument were free to travel around the unit circle.

Figure 5 is designed to satisfy such curiosity. Here, it can be seen that, within the stop band, the function $\phi^\circ(\omega)$ is zero-valued only at the Fourier frequencies. Since the data are compounded from sinusoidal functions with Fourier frequencies, this is enough to eliminate from the sample all elements that fall within the stop band and to preserve all that fall within the pass band.

This explains the seeming paradox whereby we are able to achieve a perfect frequency selection via a finite filter. In theory, we do have a doubly infinite sequence at our disposal in the form of the periodic extension of the data. Applying a wrapped filter to a finite data sequence by circular convolution is equivalent to applying an infinite, unwrapped filter to its periodic extension by ordinary linear convolution.

In Figure 5, the transition between the pass band and the stop band occurs at a Fourier frequency. This feature is also appropriate to other filters that one can design, which have a more gradual transition with a mid point that also falls on a Fourier frequency. Figure 6 shows that it is possible, nevertheless, to make an abrupt transition within the space between two adjacent frequencies. The formulae for the filter coefficients in both cases are given in the appendix.

The method of filter design that we are pursuing in this paper allows considerable flexibility in specifying the form of the frequency response function. Often, there is an advantage in departing from the ideal specification so as to allow a more gradual transition between the pass band and the stop band. However, for




Figure 5. The frequency response of the 16-point wrapped filter defined over the interval $[-\pi, \pi)$. The values at the Fourier frequencies are marked by circles and dots. (Note that, when the horizontal axis is wrapped around the circumference of the unit circle, the points at π and $-\pi$ coincide. Therefore, only one of them is included in the interval.)




Figure 6. The frequency response of the 17-point wrapped filter defined over the interval $[-\pi, \pi)$. The values at the Fourier frequencies are marked by circles.

freely specified responses, it may be difficult to find analytic expressions for the corresponding filter coefficients.

In the case of the ideal function, the coefficients of the wrapped filter are readily available. In the appendix, the coefficients are found of the lowpass filter that is obtained by sampling the following periodic frequency response function at T Fourier points $\omega_j = 2\pi j/T$ that lie in the interval $[-\pi, \pi)$:

$$\phi(\omega) = \begin{cases} 1, & \text{if } \omega \in (-\omega_d, \omega_d), \\ 1/2, & \text{if } \omega = \pm\omega_d, \\ 0, & \text{for } \omega \text{ elsewhere in } [-\pi, \pi). \end{cases} \quad (40)$$

Here, $\pm\omega_d = \pm d\omega_1 = \pm 2\pi d/T$ are the points of discontinuity. The filter coefficients

are given by

$$\phi_d^\circ(k) = \begin{cases} \frac{2d}{T}, & \text{if } k = 0 \\ \frac{\cos(\omega_1 k/2) \sin(d\omega_1 k)}{T \sin(\omega_1 k/2)}, & \text{for } k = 1, \dots, [T/2], \end{cases} \quad (41)$$

where $\omega_1 = 2\pi/T$ and where $[T/2]$ is the integral part of $T/2$.

One might wish to construct a wrapped filter according to the more general bandpass specification:

$$\phi(\omega) = \begin{cases} 1, & \text{if } |\omega| \in (-\omega_a, \omega_b), \\ 1/2, & \text{if } \omega = \pm\omega_a, \pm\omega_b, \\ 0, & \text{for } \omega \text{ elsewhere in } [-\pi, \pi), \end{cases} \quad (42)$$

where $\omega_a = a\omega_1$ and $\omega_b = b\omega_1$. For the ideal filter, this can be achieved by subtracting one filter from another to create

$$\begin{aligned} \phi_{[a,b]}^\circ(t) &= \phi_b^\circ(t) - \phi_a^\circ(t) \\ &= \frac{\cos(\omega_1 t/2) \{\sin(b\omega_1 t) - \sin(a\omega_1 t)\}}{T \sin(\omega_1 t/2)} \\ &= 2 \cos(g\omega_1 t) \frac{\cos(\omega_1 t/2) \sin(d\omega_1 t)}{T \sin(\omega_1 t/2)}. \end{aligned} \quad (43)$$

Here, $2d = b - a$ is the width of the pass band (measured in terms of a number of sampled points) and $g = (a + b)/2$ is the index of its centre. The final expression follows from the identity $\sin(A + B) - \sin(A - B) = 2 \cos A \sin B$. The expression can be interpreted as the result of shifting a lowpass filter with a cut-off frequency at the point d so that its centre is moved from 0 to the point g . The technique of frequency shifting is not confined to the ideal frequency response. It can be applied to any frequency response function.

7. Filtering Trended Sequences

The problems of filtering a trended data sequence may be overcome by reducing it to stationarity by differencing. The differenced sequence can be filtered and, if necessary, it can be reinflated thereafter to obtain an estimate of a trended data component. If one is seeking to estimate a stationary component of a nonstationary sequence, then the reinflation can be avoided.

The matrix that takes the p -th (backward) difference of a vector of order T is given by

$$\nabla_T^p = (I - L_T)^p, \quad (44)$$

where $L_T = [e_1, e_2, \dots, e_{T-1}, 0]$ is the matrix lag operator that is formed from the identity matrix $I_T = [e_0, e_1, e_2, \dots, e_{T-1}]$ by deleting the leading vector and appending a zero vector to the end of the array.

The differencing matrix may be partitioned such that $\nabla_T^p = [Q_*, Q']'$, where Q_* has p rows. The inverse matrix is partitioned conformably to give $\nabla_T^{-p} = [S_*, S]$. It follows that

$$[S_* \quad S] \begin{bmatrix} Q_*' \\ Q' \end{bmatrix} = S_* Q_*' + S Q' = I_T, \quad (45)$$

and that

$$\begin{bmatrix} Q_*' \\ Q' \end{bmatrix} [S_* \quad S] = \begin{bmatrix} Q_*' S_* & Q_*' S \\ Q' S_* & Q' S \end{bmatrix} = \begin{bmatrix} I_d & 0 \\ 0 & I_{T-d} \end{bmatrix}. \quad (46)$$

When the difference operator is applied to the data vector y , the first p elements of the product, which are in g_* , are not true differences and they are liable to be discarded:

$$\nabla_T^p y = \begin{bmatrix} Q_*' \\ Q' \end{bmatrix} y = \begin{bmatrix} g_* \\ g \end{bmatrix}. \quad (47)$$

However, if the elements of g_* are available, then the vector y can be recovered from $g = Q' y$ via the equation

$$y = S_* g_* + S g. \quad (48)$$

The columns of the matrix S_* provide a basis for the set of polynomials of degree $p - 1$ defined over the integer values $t = 0, 1, \dots, T - 1$. Therefore, $S_* g_*$ is a vector of polynomial ordinates, whilst g_* can be regarded as a vector of p polynomial parameters.

Taking differences via the backward-looking operator ∇^p induces a phase lag. This is reflected in the indexing of the elements of transformed data vector when it is written with $Q_*' y = g_* = [g_0, \dots, g_{p-1}]'$ and $Q' y = g = [g_p, \dots, g_{T-1}]'$. The summation operator, on the other hand, reverses the phase lag by inserting the elements of $S_* g_* = [y_0, \dots, y_{p-1}]'$ at the head of the recovered data vector and by moving the other elements forwards in time.

For an example of the differencing operator, we may consider the case where the degree of differencing is $p = 2$ and the length of the data sequence is $T = 5$. Then, there is

$$\nabla_5^2 = \begin{bmatrix} Q_*' \\ Q' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ -2 & 1 & 0 & 0 & 0 \\ \hline 1 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & 1 & -2 & 1 \end{bmatrix}. \quad (49)$$

A trended data sequence may be filtered as follows. First, the data is reduced to stationarity by differencing it an appropriate number of times. (We rarely need to difference the data more than twice.) Next, the relevant filters are applied to the differenced data to isolate its components. Finally, the components of the differenced data may be integrated, with an appropriate choice of initial conditions, to provide estimates of the components of the original trended sequence.

The initial conditions will be determined according to a criterion that assumes different forms depending on whether the component to be extracted is trended or non-trended. A trended component will contain the zero-frequency Fourier element,

which is a constant vector, together with other elements of adjacent frequencies in the vicinity of the zero. We presume that there can be only one trended component. The remaining components will comprise fluctuations around mean values of zero.

In accordance with this categorisation, we can represent the generic decomposition of the data vector as

$$y = x + h, \quad (50)$$

where x is the trend component and h is the complementary detrended component, which might be subject to further decompositions. The differenced data would be

$$\begin{aligned} Q'y &= Q'x + Q'h \\ &= d + k = g. \end{aligned} \quad (51)$$

The vectors d and k require to be cumulated to form

$$x = S_*d_* + Sd \quad \text{and} \quad h = S_*k_* + Sk. \quad (52)$$

However, given the adding-up constraint that is posed by (50), the initial conditions within d_* and k_* must be equivalent.

The initial conditions should be chosen so as to ensure that the trend is aligned with the data as closely as possible or, equivalently, that the deviations of the trend from the data are minimised. This entails minimising the quadratic norm $h'h = (y - x)'(y - x)$. The criterion for finding k_* is, therefore,

$$\text{Minimise } (S_*k_* + Sk)'(S_*k_* + Sk) \quad \text{with respect to } k_*. \quad (53)$$

The solution for the starting values is

$$k_* = -(S'_*S_*)^{-1}S'_*Sk. \quad (54)$$

The equivalent criterion for finding d_* is

$$\text{Minimise } (y - S_*d_* - Sd)'(y - S_*d_* - Sd) \quad \text{with respect to } d_*. \quad (55)$$

The solution for the starting values is

$$d_* = (S'_*S_*)^{-1}S'_*(y - Sd). \quad (56)$$

In terms of the notation

$$P_* = S_*(S'_*S_*)^{-1}S'_*, \quad (57)$$

the equations of (52) can be written as

$$x = P_*y + (I - P_*)Sd, \quad \text{and} \quad h = (I - P_*)Sk. \quad (58)$$

Since $(I - P_*)\iota = 0$, where $\iota = [1, 1, \dots, 1]'$ is the summation vector, it follows that $\iota'h = 0$, which is to say that the detrended data has a mean of zero. Then, given

that $S(d+k) = Sg$ and that $(I - P_*)Sg = (I - P_*)(Sg + S_*g) = (I - P_*)y$, since $(I - P_*)S_* = 0$, it follows that

$$\begin{aligned} x + h &= (I - P_*)S(d+k) + P_*y \\ &= (I - P_*)y + P_*y = y, \end{aligned} \tag{59}$$

which is to say that the sum of the estimated components is the original data vector, in accordance with (50). In practice, it is redundant to compute both x and h . Only one of them is required, since the other can be found by subtracting from y .

In extracting the stationary component h contained within a trended data sequence y , the business of reinflating the filtered sequence can be avoided, thereby dispensing with the initial conditions. The highpass filter that would serve to extract $k = Q'h$ from $g = Q'y$ will contain an implicit differencing operator, which serves to nullify the low-frequency elements of the data. If the filter is symmetric, then it will embody at least a twofold differencing operator. The need for reinflation can be avoided by cancelling the inflating summation operator with the differencing factors within the filter.

We may begin by considering the symmetric version of the twofold differencing operator, which is to be applied to the data at the outset. This is

$$\begin{aligned} N(z) &= z^{-1} - 2 + z = z^{-1}(1 - z)^2 \\ &= z^{-1}\nabla^2(z). \end{aligned} \tag{60}$$

The matrix version of the operator is obtained by setting $z = L_T$ and $z^{-1} = L'_T$, which gives

$$N(L_T) = N_T = L_T - 2I_T + L'_T. \tag{61}$$

The first and the final rows of this matrix do not deliver true differences. Therefore, they are liable to be deleted, with the effect that the two end points are lost from the twice-differenced data. Deleting the rows e'_0N_T and $e'_{T-1}N_T$ from N_T gives the matrix Q' , which can also be obtained by from $\nabla_T^2 = (I_T - L_T)^2$ by deleting the matrix Q'_* , which comprises the first two rows $e'_0\nabla_T^2$ and $e'_1\nabla_T^2$. In the case of $T = 5$ there is

$$N_5 = \begin{bmatrix} Q'_{-1} \\ Q' \\ Q_{+1} \end{bmatrix} = \begin{bmatrix} -2 & 1 & 0 & 0 & 0 \\ \hline 1 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & 1 & -2 & 1 \\ \hline 0 & 0 & 0 & 1 & -2 \end{bmatrix}. \tag{62}$$

On deleting the first and last elements of the vector N_Ty , which are $Q'_{-1}y = e'_1\nabla_T^2y$ and $Q_{+1}y$, respectively, we get $Q'y = [q_1, \dots, q_{T-2}]'$.

The vector $Q'y$ of differenced data is to be used both in the procedure that reinflates the filtered sequence and in the present procedure that avoids doing so. However, in the absence of reinflation, the missing elements are not restored to the filtered vector, nor is there any accompanying phase alteration.

The loss of the two elements from either end of the (centrally) twice-differenced data can be overcome by supplementing the original data vector y with two extrapolated end points y_{-1} and y_T . Alternatively, the differenced data may be supplemented by attributing appropriate values to q_0 and q_{T-1} . These could be zeros or some combination of the adjacent values. In either case, we will obtain a vector of order T denoted by $q = [q_0, q_1, \dots, q_{T-1}]'$.

In describing the method for implementing a highpass filter, let Λ be the matrix which selects the appropriate ordinates of the Fourier transform $\gamma = Uq$ of the twice differenced data. These ordinates must be reinflated to compensate for the differencing operation, which has the frequency response

$$f(\omega) = 2 - 2 \cos(\omega). \quad (63)$$

The response of the anti-differencing operation is $1/f(\omega)$; and γ is reinflated by pre-multiplying by the diagonal matrix

$$V = \text{diag}\{v_0, v_1, \dots, v_{T-1}\}, \quad (64)$$

comprising the values $v_j = 1/f(\omega_j)$; $j = 0, \dots, T-1$, where $\omega_j = 2\pi j/T$.

Let $H = V\Lambda$ be the matrix that is applied to $\gamma = Uq$ to generate the Fourier ordinates of the filtered vector. The resulting vector is transformed to the time domain to give

$$h = \bar{U}H\gamma = \bar{U}HUq. \quad (65)$$

It will be seen that $f(\omega)$ is zero-valued when $\omega = 0$ and that $1/f(\omega)$ is unbounded in the neighbourhood of $\omega = 0$. Therefore, a frequency-domain reinflation is available only when there are no nonzero Fourier ordinates in this neighbourhood. That is to say, it can work only in conjunction with highpass or bandpass filtering. However, it is straightforward to construct a lowpass filter that complements the highpass filter. The low-frequency trend component that is complementary to h is

$$x = y - h = y - \bar{U}HUq. \quad (66)$$

8. Filtering in Practice

The previous sections of this paper have provided some alternative frequency-domain methods for filtering trended and non-trended sequences. The fine details of how these methods should be applied in practice will depend upon the precise characteristics of the data sequences. The choice of an appropriate method is also liable to reflect the underlying purpose of the analysis.

It might be the intention to isolate a composite trend-cycle component in order to depict the underlying growth trajectory of the economy. Alternatively, a pure business cycle component might be sought that is devoid of any trend. Another purpose that can be served by these methods is to provide a seasonally-adjusted sequence by removing from the data the sinusoidal elements that have frequencies in the vicinity of the fundamental seasonal frequency and its harmonics.

A program has been written in conjunction with this paper that is capable of serving all of these purposes. The program is designed to extract components of




Figure 7. The logarithms of the monthly data on retail sales in U.S.A. for the years 1954 to 1964, together with an interpolated trend.




Figure 8. The periodogram of residuals obtained by fitting a linear trend through the logarithmic sales data of Figure 7.




Figure 9. The seasonal fluctuations of the data of Figure 7, which are virtually identical to the deviations of the data from the interpolated trend.

the data that lie within well-defined frequency bands and to do so without altering the relative amplitudes of the sinusoidal elements of which they are composed. However, it should be emphasised that the methods of frequency-domain filtering can be applied more flexibly to give differential weightings to these elements. By such means, a frequency-domain method is capable of mimicking the effects of any time-domain filter that has a well-defined frequency response function.

When the objective is that of extracting a trend-cycle component, there are three alternative approaches, which are liable to generate results that are virtually identical. Two of these methods have already been expounded in the previous section. They differ in the manner in which they effect the cumulation of a filtered version of a differenced sequence, obtained by applying a differencing operator to the original trended data sequence.

The third method depends upon removing the trend from the data by interpolating a polynomial function and by applying the filter to the residual sequence. A lowpass filter can be applied to this sequence to extract the cycles, whereafter the result can be added to the polynomial function to create the trend-cycle component.

Each of the foregoing methods depends for its success upon an adequate detrending of the data. It is necessary to avoid any radical disjunctions in the periodic extension of the data sequence where the end of one replication of the sample joins the beginning of the next. The same disjunction will arise when the data sequence is mapped onto the circumference of a circle, at the point where the head of the sequence joins the tail.

The traditional means of avoiding disjunctions has been by tapering the ends of the mean-adjusted data sequence so that they both decline to zero. (see Bloomfield 1976, for example.) The disadvantage of this recourse is that it tends to falsify the data at the ends of the sequence. This is particularly inconvenient if, as is often the case in economics, attention is liable to be focussed of the most recent data.

The difficulty can be overcome by extrapolating the data at both ends via an interpolated polynomial, of which the degree should be equal to the order of the differencing to which the data will be subjected subsequently. The polynomial can be fitted to the data by a weighted least-squares regression that gives large weights to the points close to the ends in order to ensure that it passes through their midst.

Tapered versions of the residual sequence that have been reflected around the endpoints of the sample can be added to the extrapolated branches of the polynomial. Alternatively, if the data show strong seasonal fluctuations, then a tapered sequence based on successive repetitions of the ultimate seasonal cycle can be added to the upper branch, and a similar sequence based on the first cycle can be added to the lower branch. After the augmented data have been filtered, the extrapolations can be discarded.

The method of extrapolation will prevent the end of the sample from being joined directly to its beginning. When the data are supplemented by extrapolations, the circularity of the filter will effect only the furthest points the extrapolation, which will usually be discarded after the filtering has taken place.

In many cases, such extrapolations have proved to be unnecessary. However, due care must be exercised in cases where the data are subject to significant seasonal fluctuations, which can be of an amplitude that equals or exceeds that of the secular or business-cycle fluctuations that one might wish to uncover. In that case, it may

necessary either to deseasonalise the data in advance or else to take steps to ensure that it contains an integral number of seasonal cycles.

Consider a seasonally fluctuating sequence that spans an integral number of years. When this is subjected to the centralised twofold differencing operator, it will lose an observation from each end. Then, the data will no longer span those years.

There are two ways of overcoming the detriment. Either the sequence can be further shortened so that there is once more an integral number of years, or else the two lost points can be replaced by appropriate estimates. It is usually sufficient to use as estimates the points from the corresponding seasons in the adjacent years.

Example. The techniques expounded in this paper can be explored within the above-mentioned computer program that has been devised in order to test and to illustrate them. A data sequence that demands to be handled with care is provided by the 144 monthly observations on retail sales in the US from 1954 to 1964, which were recorded in the paper of Shiskin, Young and Musgrave (1967) that presented the X-11 program of the U.S. Census Bureau. This program, which was based on the Henderson (1916) filter, is intended for seasonal adjustment and trend estimation.

Figure 7 shows the logarithms of the data and Figure 8 shows the periodogram of the residuals from fitting a linear trend. The periodogram has a prominent spike at the seasonal frequency of $\pi/6$ and at the harmonic frequencies of $\pi/3$, $\pi/2$, $2\pi/3$, $5\pi/6$ and π . With the exception of the interval $(2\pi/3, 5\pi/6)$, which does contain one significant ordinate, the interstices between these seasonal frequencies are virtual dead spaces. The presence of some nonzero ordinates in the interval $[0, \pi/6)$, which covers the trend frequencies, indicates that the log-linear detrending is inadequate.

The trend that is portrayed in Figure 7 is based on the Fourier ordinates of the twice-differenced data that lie in the interval $[0, \pi/6)$. To avoid some distortionary end effects, the sample has been extrapolated for a short distance in the manner described above.

Figure 9 shows a sequence of seasonal fluctuations that has been synthesised from a selection of the Fourier ordinates of the differenced data. The ordinates correspond to the seasonal frequency and its harmonics and to the two frequency points immediately above $2\pi/3$. The periodogram of Figure 8 has provided the necessary guidance in selecting these frequencies.

The differenced seasonal vector, synthesised from these elements, may be denoted by k . The vector h of the seasonal fluctuations is obtained by cumulating k via the formula $h = S_* k_* + Sk$ of (52), wherein k_* is calculated according to (54). It transpires that this synthesised seasonal vector, which is represented by Figure 9, is virtually indistinguishable from the vector of the residuals obtained by subtracting the trend vector x of Figure 7 from the corresponding data vector y .

The Programs

The program that has been written for the purposes of this paper is available on request from the author. There are versions both for the Macintosh computer and for Microsoft Windows. The Windows version of the program is also available at the following web address:

<http://www.le.ac.uk/users/dsgp1/>

Appendix: The Wrapped Coefficients of the Ideal Lowpass Filter

Let the sample size be T , and consider a set of T frequency-domain ordinates sampled, at the Fourier frequencies $\omega_j = 2\pi j/T$ that fall within the interval $[-\pi, \pi)$, from a boxcar function, centred on $\omega_0 = 0$. If the cut-off points are at $\pm\omega_d = \pm d\omega_1 = \pm 2\pi d/T$, then the ordinates of the sample will be

$$(A.1) \quad \lambda_j = \begin{cases} 1, & \text{if } j \in \{1-d, \dots, d-1\}, \\ 1/2, & \text{if } j = \pm d, \\ 0, & \text{otherwise.} \end{cases}$$

Their (discrete) Fourier transform is the sequence of the coefficients

$$(A.2) \quad \phi_k^\circ = \frac{1}{T} \sum_j \lambda_j e^{i\omega_j k},$$

defined for $k = 0, 1, \dots, T-1$.

The ordinates $\lambda_d = \lambda_{-d} = 1/2$ cause some inconvenience in evaluating this transform. To overcome this, we may begin by evaluating the function

$$(A.3) \quad S^+(z) = z^{1-d} + \dots + z^{-1} + 1 + z + \dots + z^d,$$

where $z = e^{i\omega}$, together with the function $S^-(z) = z^{-1}S^+(z)$. Then we may form the symmetric function $\phi^\circ(z) = \{S^-(z) + S^+(z)\}/(2T)$, wherafter we may set $\omega = \omega_k = 2\pi k/T = k\omega_1$ to obtain the k th coefficient.

First, consider

$$(A.4) \quad \begin{aligned} S^+(z) &= z^{1-d}(1 + z + \dots + z^{2d-1}) \\ &= z^{1-d} \frac{(1 - z^{2d})}{1 - z}. \end{aligned}$$

Multiplying top and bottom by $z^{-1/2}$ gives

$$(A.5) \quad \begin{aligned} S^+(z) &= z^{(1/2)-d} \frac{(1 - z^{2d})}{z^{-1/2} - z^{1/2}} \\ &= z^{1/2} \frac{(z^{-d} - z^d)}{z^{-1/2} - z^{1/2}}. \end{aligned}$$

Then, by setting $z = e^{i\omega}$, we get

$$(A.6) \quad S^+(e^{i\omega}) = e^{i\omega/2} \frac{(e^{i\omega d} - e^{-i\omega d})}{e^{i\omega/2} - e^{-i\omega/2}} = e^{i\omega/2} \frac{\sin(\omega d)}{\sin(\omega/2)}.$$

When $\omega = \omega_k = k\omega_1$, this becomes

$$(A.7) \quad S^+(k) = e^{i\omega_1 k/2} \frac{\sin(d\omega_1 k)}{\sin(\omega_1 k/2)},$$

which is the Dirichlet function multiplied by a complex exponential that owes its presence to the non-symmetric nature of $S^+(z)$. There is also

$$(A.8) \quad S^-(k) = e^{-i\omega_1 k/2} \frac{\sin(d\omega_1 k)}{\sin(\omega_1 k/2)},$$

Therefore, for $k \neq 0$, there is

$$(A.9) \quad \phi^\circ(k) = \frac{1}{2T} \{S^-(k) + S^+(k)\} = \frac{\cos(\omega_1 k/2) \sin(d\omega_1 k)}{T \sin(\omega_1 k/2)},$$

whereas for $k = 0$ there is $\phi_0^\circ = 2d/T$, which comes from setting $z = e^0 = 1$ in the expression for $S^+(z)$ of (A.3) and in the analogous expression for $S^-(z)$.

Setting $d = T/4$ in (A.9) gives the wrapped version of the lowpass half band filter that is the subject of Section 2 of the paper.

In an alternative specification of the ideal filter, the cut-off points fall between Fourier frequencies. Then, the ordinates sampled from the frequency response function are

$$(A.10) \quad \lambda_j = \begin{cases} 1, & \text{if } j \in \{1-d, \dots, d-1\}, \\ 0, & \text{otherwise.} \end{cases}$$

In place of $\{S^-(z) + S^+(z)\}/2$, there is

$$(A.11) \quad \begin{aligned} S(z) &= z^{1-d} + \dots + z^{-1} + 1 + z + \dots + z^{d-1} \\ &= \frac{z^{(1/2)-d} - z^{d-1/2}}{z^{-1/2} - z^{1/2}}. \end{aligned}$$

Then, for $k = 0$, there is $\phi_0^\circ = (2d-1)/T$, whereas, for $k \neq 0$, the formula is

$$(A.12) \quad \phi^\circ(k) = \frac{\sin([d-1/2]\omega_1 k)}{T \sin(\omega_1 k/2)}.$$

References

Baxter, M., and R.G. King, (1999), Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series, *Review of Economics and Statistics*, 81, 575–593.

Bell, W.M., and D.E.K. Martin, (2004), Computation of Asymmetric Signal Extraction Filters and Mean Squared Error for ARIMA Components Models, *Journal of Time Series Analysis*, 25, 603–623.

Bloomfield, P., (1976), *Fourier Analysis of Time Series: An Introduction*, John Wiley and Sons, New York.

Burns, A.M., and W.C. Mitchell, (1946), *Measuring Business Cycles*, New York: National Bureau of Economic Research.

Caporello, G., and A. Maravall (2004), *Program TSW: Revised Reference Manual*, Working Paper, Research Department, Banco de España.

Christiano, L.J., and T.J. Fitzgerald, (2003), The Band-pass Filter, *International Economic Review*, 44, 435–465

Findley, D.F., and D.E.K. Martin, (2003), *Frequency Domain Analyses of SEATS and X-11/12-ARIMA Seasonal Adjustment Filters for Short and Moderate-Length Time Series*, Statistical Research Division U.S. Bureau of the Census Washington D.C. 20233.

Haywood, J., and G. Tunnicliffe Wilson, (1997), Fitting Time Series Models by Minimising Multistep-Ahead Errors: A Frequency Domain Approach, *Journal of the Royal Statistical Society, Series B (Methodological)*, 59, 237–254.

Henderson, R., (1916), Note on Graduation by Adjusted Averages, *Transactions of the Actuarial Society of America*, 17, 43–48.

Iacobucci, Alessandra, and A. Noullez, (2005), A Frequency Selective Filter for Short-Length Time Series, *Computational Economics*, 25, 75–102.

Pollock, D.S.G., (1999), *A Handbook of Time-Series Analysis, Signal Processing and Dynamics*, Academic Press, London.

Pollock, D.S.G., (2002), Circulant Matrices and Time-Series Analysis, *The International Journal of Mathematical Education in Science and Technology*, 33, 213–230.

Proietti, T., (2005), Forecasting and Signal Extraction with Misspecified Models, *Journal of Forecasting*, 24, 539–556.

Shiskin, J., A.H. Young and J.C Musgrave, (1967), *The X-11 Variant of the Census Method II Seasonal Adjustment Program*, Technical Paper no. 15, U.S. Dept. of Commerce, Bureau of the Census.

www.census.gov/srd/www/historicpapers.html