

ECONOMETRIC FILTERS



D.S.G. POLLOCK, University of Leicester

ECONOMETRIC FILTERS

By D.S.G. POLLOCK*

University of Leicester

Email: stephen_pollock@sigmapl.u-net.com

A variety of filters that are commonly employed by econometricians are analysed with a view to determining their effectiveness in extracting well-defined components of economic data sequences. These components can be defined in terms of their spectral structures—i.e. their frequency content—and it is argued that the process of econometric signal extraction should be guided by a careful appraisal of the periodogram of the detrended data sequence.

Whereas it is true that many annual and quarterly economic data sequences are amenable to relatively unsophisticated filtering techniques, it is often the case that monthly data that exhibit strong seasonal fluctuations require a far more delicate approach. In such cases, it may be appropriate to use filters that work directly in the frequency domain by selecting or modifying the spectral ordinates of a Fourier decomposition of data that have been subject to a preliminary detrending.

Keywords: Time series, Spectral analysis, Business cycles, Turning points, Seasonality.

1. Introduction

In many cases, the macroeconomic data sequences that are subject to filtering are tolerant of poorly designed filters of a sort that would not be acceptable in other applications, such as in audio-acoustic engineering. For that reason, and for other reasons, such as a lack of knowledge on the part of economists and econometricians of the essential frequency-domain analysis, the development of appropriate filtering methods has been somewhat retarded.

The purpose of this paper is to describe some of the filters that are commonly employed by economists and to define the limits of their applicability. The outcome should be some clear suggestions for how one should approach the matter of econometric signal extraction in various circumstances. Whenever the detailed derivation of a filter is omitted, an appropriate reference will be provided.

It is appropriate to begin the analysis of econometric filters by considering a macroeconomic data sequence that delivers similar results from a wide variety of filters. Then, other sequences that are much less tractable can be considered. The tractable sequence in question is that of the quarterly data on aggregate consumption in the U.K. from 1955Q1 to 1994Q4. The logarithms of the data are plotted in Figure 1, which also displays a linear trend that has been fitted to the data by a least-squares regression.

* This paper was written in support of a lecture delivered in the Department of Econometrics and Business Statistics of Monash University. I wish to acknowledge the outstanding hospitality of the Department.

ECONOMETRIC FILTERS




Figure 1. The quarterly sequence of the logarithms of household consumption expenditure in the U.K. for the years 1955 to 1994, with an interpolated linear trend.




Figure 2. The periodogram of the logarithmic consumption data.




Figure 3. The periodogram of the residual sequence from the linear detrending of the logarithmic consumption data.

2. The Periodogram

To understand the data from the point of view of filtering theory, it is necessary to look at their periodogram. The periodogram depicts the squared amplitudes $\rho_j^2; j = 0, 1, \dots, [T/2]$ of the phase-displaced cosine functions into which the data sequence $\{y_t; t = 0, 1, \dots, T-1\}$ can be decomposed. (Here, $[T/2]$ denotes the integer quotient from the division of T by 2.) The elements of the data sequence can be reconstituted by summing these cosine functions. Thus

$$\begin{aligned} y_t &= \sum_{j=0}^{[T/2]} \rho_j \cos(\omega_j t + \theta_j) \\ &= \sum_{j=0}^{[T/2]} \{\alpha_j \cos(\omega_j t) + \beta_j \sin(\omega_j t)\}, \end{aligned} \tag{1}$$

where $\omega_j = 2\pi j/T; j = 0, 1, \dots, [T/2]$ are the so-called Fourier frequencies, which are evenly distributed in the interval $[0, \pi]$. The second expression resolves each of the displaced cosine functions into the sum of a sine function and a cosine function, weighted by the appropriate coefficients α_j and β_j . The two expressions are related via the identities $\rho_j^2 = \alpha_j^2 + \beta_j^2$ and $\theta_j = \tan^{-1}(\beta_j/\alpha_j)$.

The sines and cosines are perpetual functions of constant amplitude that are defined on the entire set of positive and negative integers. Equivalently, they can be envisaged as functions defined on the perimeter of a circle. In projecting a finite data sequence onto these functions, we are constrained to adopt the fiction that the sequence represents a single cycle of a periodic function that would be obtained by a perpetual replication of the sequence.

This perpetual replication of the data over all preceding and subsequent time is described as their periodic extension. The data may also be regarded, equivalently, as forming a circular sequence, described as the circular wrapping of the data.

When it is replicated perpetually, a finite trended sequence will give rise, not to a continuously increasing function, but, instead, to a saw tooth function. This function has a one-over- f periodogram, resembling a rectangular hyperbola, in which the low-frequency component will far outweigh the other elements of the Fourier transform. The periodogram of the trending consumption data, which is shown in Figure 2, has this feature.

In order to assess the remaining components of the data, one should examine the periodogram of the detrended data. In the case of the logarithmic consumption data, it is appropriate to examine the periodogram of the residual sequence from a linear detrending. The vector of the ordinates of the linear function interpolated into the data sequence by an ordinary least-squares regression is given by

$$\begin{aligned} x &= y - Q(Q'Q)^{-1}Q'y \\ &= y - e, \end{aligned} \tag{2}$$

where e is the vector of the residual sequence, and where

$$Q' = \begin{bmatrix} 1 & -2 & 1 & \dots & 0 & 0 \\ 0 & 1 & -2 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & \dots & -2 & 1 \end{bmatrix} \quad (3)$$

is the matrix version of the twofold difference operator.

Figure 3 shows the periodogram of the residual vector e . There is a low-frequency structure that extends no further than the frequency value of $\pi/8$ radians or 22.5 degrees, which corresponds to the minimum cyclical duration of four years. This is followed by a wide dead space that extends to a point somewhat short of the frequency value of $\pi/2$, where there is a tall spike. This is the fundamental seasonal frequency, which, given that the data are quarterly, corresponds to one cycle per annum. The spike is followed by another dead space that extends almost to the Nyquist frequency value of π , which is the harmonic of the seasonal frequency, and which is marked by another spike.

There are various things that one can do to the consumption data with a linear filter. One can effect the seasonal adjustment of the data by removing the spikes at $\pi/2$ and π . It matters little if one removes everything that resides within wide vicinities of these values, since the spectral dead spaces contribute very little to the data.

One might also choose to isolate the low-frequency structure that falls in the interval $[0, \pi/8)$, which can be regarded as the spectral signature of the business cycle. In this case, it matters little if what is isolated in pursuit of the business cycle comprises elements that fall in an interval that runs almost to $\pi/2$, since these elements are virtually insignificant.

(The elements that lie within the dead-space interval $[\pi/8, \pi/2)$, which excludes the fundamental seasonal frequency at $\pi/2$, contribute to variance of the detrended logarithmic consumption data barely 15 percent of what is contributed by the ordinates in the interval $[0, \pi/8)$, which are identified with the business cycle.)

3. Local Polynomials: The Henderson Filters

The so-called trend-cycle component of the consumption data can be successfully estimated by using one of the time-honoured Henderson filters. In this case, the filter can be applied to the trended data rather than to the linearly detrended data. If it were applied to the latter, then one should add the filtered sequence to the linear trend in order to obtain the trend-cycle function.

The Henderson filters are derived by pursuing a concept of local polynomial regression. A polynomial is fitted to the points that fall within a window, spanning $2m + 1$ data points, that advances step-by-step through the data. At each step, the smoothed value that replaces the corresponding data value is the central ordinate of the fitted polynomial. The outcome of the regression is a set of moving-average coefficients $\psi_j; j = 0, \pm 1, \dots, \pm m$ that are disposed

ECONOMETRIC FILTERS




Figure 4. The coefficients of the symmetric Henderson filter of 23 points.




Figure 5. A trend, determined by a Henderson filter with 23 coefficients, interpolated through the 160 points of the logarithmic consumption data.




Figure 6. The frequency response function of the Henderson moving-average filter of 23 terms.

ECONOMETRIC FILTERS

symmetrically around the central value ψ_0 , with $\psi_{-j} = \psi_j$. These are applied throughout the sample except at the beginning and the end.

In the case of the Henderson filters, a cubic polynomial is fitted to the windowed data points. The consequence is that the resulting filters will transmit, without alteration, the ordinates of any polynomial of degree three or less to which it might be applied. The polynomial regression in question employs a generalised least-squares criterion that minimises the sum of squares of the third differences of the polynomial ordinates.

The filters, which were derived by Robert Henderson (1916, 1924), are used within the X-11 family of seasonal adjustment programs, where they are applied to data that have already been seasonally adjusted. However, they can be applied directly to seasonal data in pursuit of an estimate of the trend-cycle component. Detailed accounts of the filters have been provided by Kenney and Durbin (1982) and by Pollock (2009a), and the X-11 program has been described in detail by Ladiray and Quenneville (2001).

Figure 4 displays the coefficients of the 23-point Henderson filter and Figure 5 shows the effect of applying the Henderson filter directly to the logarithmic consumption data. The filter appears to do a reasonable job of estimating the trend-cycle function. Notice also that the filter runs to the ends of the sample, whereas one might expect it to fall short, leaving $m = 11$ points unprocessed at either end.

This feat is achieved by virtue of some cunning modifications of the filter that adapts its coefficients as it nears the ends. The adaptations are equivalent to an extrapolation of the data beyond the ends of the sample, sufficient to support the filter coefficients. The resulting asymmetric filters were proposed originally by Musgrave (1964a, b) in two unpublished notes, and their rationale has been described more fully by Doherty (2001).

The gain of the Henderson filter is depicted in Figure 6. This function indicates the extent to which the filter will alter the amplitudes of the trigonometrical functions of frequencies $\omega \in [0, \pi]$, which are the elements of the spectral decomposition of a stationary stochastic process. The gain effect is one aspect of the frequency response of the filter. The other aspect is the phase effect, by which the elements are displaced in time.

The two effects can be revealed by mapping the complex exponential sequence $x(t) = \cos(\omega t) + i\sin(\omega t) = \exp\{i\omega t\}$ through the filter defined by the coefficients $\{\psi_j\}$ to give

$$y(t) = \sum_j \psi_j e^{i\omega(t-j)} = \left\{ \sum_j \psi_j e^{-i\omega j} \right\} e^{i\omega t} = \psi(\omega) e^{i\omega t}. \quad (4)$$

The effects are summarised by the complex function

$$\psi(\omega) = |\psi(\omega)| e^{i\theta(\omega)}. \quad (5)$$

On the RHS, there is the gain effect $|\psi(\omega)|$, which corresponds to the modulus of the complex function, and the phase effect $\theta(\omega)$, which corresponds to its argument.




Figure 7. The central coefficients of the ideal bandpass filter defined on the frequency interval $[\pi/16, \pi/3]$.

In the case of a symmetric Henderson filter, where $\psi_j = \psi_{-j}$, the associated complex exponential functions combine to form $\cos(\omega_j) = \{\exp(-i\omega_j) + \exp(i\omega_j)\}/2$. Therefore, the frequency response function is real-valued and there is no phase effect.

The frequency response of the filter allows the business-cycle component of the consumption data, which is represented by the spectral structure in Figure 3 on the interval $[0, \pi/8)$, to be transmitted with only a minor attenuation of the elements in the vicinity of the upper limit. There are no other significant elements of the data that fall within the pass band of the filter. Therefore, it serves the purpose of extracting the trend-cycle function well enough.

It will be observed that the frequency response function of the Henderson filter shows a very gradual transition from the pass band, where the elements of the Fourier decomposition are fully preserved, to the stop band, where they should be wholly nullified. There are circumstances where one would wish to have a more rapid transition.

4. Approximate Bandpass Filters

One case, where a rapid transition is desired, concerns a definition of the business cycle that is due to Arthur Burns and Wesley Mitchell (1946), who were working at the U.S. National Bureau of Economic Research throughout the 1930's and the 1940's. According to their definition, the business cycle comprises all the elements of the data that have cyclical durations of no less than one and a half years and of no more than eight years.

Baxter and King (1999) have sought to implement an appropriate filter in the time domain by taking the inverse Fourier transform of the rectangle, defined on the interval $[\alpha, \beta]$ within the frequency range $[0, \pi]$, that constitutes the ideal frequency response. For quarterly data, the values in radians that correspond to the definition of Burns and Mitchell of the business cycle are $\alpha = \pi/16$ (11.25°) and $\beta = \pi/3$ (60°).

ECONOMETRIC FILTERS

A difficulty arises from the fact that the Fourier transform of a frequency-domain rectangle gives rise to a doubly-infinite sequence of filter coefficients, of which the central values are displayed in Figure 7. The coefficients are provided by the sampled ordinates of the function

$$\begin{aligned}\psi(k) &= \frac{1}{\pi k} \{\sin(\beta k) - \sin(\alpha k)\} = \frac{2}{\pi k} \cos\{(\alpha + \beta)k/2\} \sin\{(\beta - \alpha)k/2\} \\ &= \frac{2}{\pi k} \cos(\gamma k) \sin(\delta k),\end{aligned}\tag{6}$$

described as a displaced sinc function, where $k \in \{0, \pm 1, \pm 2, \dots\}$. Here, γ , which is the displacement parameter, represents the centre of the pass band, whereas δ is half its width.

This sequence of coefficients must be drastically truncated if it is to become a moving average that can be applied to a finite data sequence. Figure 8 shows the frequency response of a truncated bandpass filter of 25 coefficients, and it compares this with the rectangle of the ideal frequency response.

The truncated filter allows elements within the stop bands to be transmitted to a significant extent. This so-called problem of leakage greatly subverts the original intentions. However, we shall have reason to doubt whether the definition of Burns and Mitchell is an appropriate one in any case. Figure 9 shows the effect of applying the filter to the logarithmic data sequence, and it also shows how the filter fails to reach the ends of the sample. The filter automatically detrends the data.

The end-of-sample problem has been tackled by Christiano and Fitzgerald (2003) who proposed a simple way of extrapolating the ends of the data. They proposed that it is reasonable to imagine that the data have been generated by a random-walk process. In that case, the optimal forecasts and backcasts are obtained simply by horizontal extrapolations of the values at the ends of the sample.

When a branch of the infinite sequence of coefficients extends beyond the end of the sample, the unsupported coefficients are summed and multiplied by the data value at the end. Then, the product is added to the sum of the products of the within-sample coefficients and the sample values. Thus, the filtered value at time t may be denoted by

$$\begin{aligned}x_t &= Ay_0 + \psi_t y_0 + \dots + \psi_1 y_{t-1} + \psi_0 y_t \\ &\quad + \psi_1 y_{t+1} + \dots + \psi_{T-1-t} y_{T-1} + By_{T-1},\end{aligned}\tag{7}$$

where A and B are the sums of the extra-sample coefficients at either end.

The method of Christiano and Fitzgerald is not appropriate to a sequence that shows a clear upward trend. Such a sequence should be subject to some form of detrending that will deliver a mean-reverting residual sequence with a mean of zero. In that case, the extra-sample values of the detrended sequence may be represented by zeros, which might stand for their unconditional expectations.

ECONOMETRIC FILTERS




Figure 8. The rectangular frequency response of the ideal bandpass filter defined on the interval $[\pi/16, \pi/3]$, together with the frequency response of the truncated filter of 25 coefficients.




Figure 9. The effect of applying the truncated bandpass filter of 25 coefficients to the quarterly logarithmic data on U.K. consumption.




Figure 10. The effect of applying the filter of Christiano and Fitzgerald to the linearly detrended quarterly logarithmic data on U.K. consumption.

ECONOMETRIC FILTERS

In the case of a data sequence that is free of trend and that can be regarded as the product of a mean-zero stationary stochastic process, there is a straightforward way of overcoming the end-of-sample problem. The role of the data can be interchanged with that of the filter. Instead of attempting to fit a truncated filter within the confines of a finite data sequence, one can run the data sequence along the central part of the infinite sequence of filter coefficients. That is to say, the data can be treated as the moving average and the filter coefficients can be treated as the data.

This is what has been done in creating Figure 10. A little thought will serve to show that this is equivalent to setting the required extra-sample values to zero. An alternative interpretation of this procedure is derived by considering a banded Toeplitz matrix Ψ of the same order as the sample.

The elements of the principal diagonal of this matrix have the value ψ_0 , given by the function of (6) when $t = 0$, and the elements of the t th subdiagonal and supradiagonal bands have the value of $\psi_t = \psi(t)$. The form of the Toeplitz matrix is adequately represented by the case of $T = 4$:

$$\Psi = \begin{bmatrix} \psi_0 & \psi_1 & \psi_2 & \psi_3 \\ \psi_1 & \psi_0 & \psi_1 & \psi_2 \\ \psi_2 & \psi_1 & \psi_0 & \psi_1 \\ \psi_3 & \psi_2 & \psi_1 & \psi_0 \end{bmatrix}. \quad (8)$$

The vector $x = \Psi d$ of the filtered values is obtained by premultiplying the vector d of the detrended data by this matrix.

Figure 10 plots the filtered sequence against the backdrop of the sequence from which it is derived, which has been obtained by a linear detrending of the logarithmic consumption data. The filtered sequence fails to follow the underlying trajectory of the detrended data. The fault lies in the partial exclusion of the low-frequency elements that are in the range $[0, \pi/16)$. An appropriate recourse would be to set $\alpha = 0$ to create a lowpass filter in place of the approximate bandpass filter.

There are limits to the power of time-domain FIR filters to resolve the data into components within well-defined bands. A superior performance can be obtained from filters that employ feedback. Time-invariant feedback filters are represented by rational polynomial transfer functions. Since the series expansion of a rational function is, in general, a power series with an infinite numbers of terms, such filters are also described as infinite impulse response or IIR filters.

5. Wiener–Kolmogorov Filters

Leading examples of IIR filters are the Wiener–Kolmogorov or W-K filters. On the one hand, there are the time-invariant filters that are derived on the assumption that the data sequence is doubly infinite. These can be represented by ratios of polynomials in the lag operator. In that case, there is no treatment of the end-of-sample problem. On the other hand, there are W-K filters that are adapted to finite samples of specific lengths. These have coefficients that

ECONOMETRIC FILTERS




Figure 11. The frequency response function of the Hodrick–Prescott lowpass smoothing filter—or Leser filter—for various values of the smoothing parameter.




Figure 12. A trend function obtained by applying a Hodrick–Prescott filter with a smoothing parameter of $\lambda = 1,600$ to 160 points of the logarithmic consumption data.




Figure 13. The frequency response function of the Butterworth filters of orders $n = 6$ and $n = 12$ with a nominal cut-off point of $\pi/6$ radians (30°).

vary as the filters move through the sample, and they must be represented by matrix transformations that are applied to the vector of the sample elements.

The W-K filters have the virtue that every lowpass filter is accompanied by a complementary highpass filter. Therefore, the data sequence can be reconstituted by adding together the products of the two filters. By the same token, the output of the lowpass filter can be obtained by subtracting the output of the highpass filter from the original data sequence.

For such filters, there is a need to take steps to cater to trended data sequences. There are several ways of doing so. Perhaps the most straightforward way is to remove a linear trend or a polynomial trend of higher degree from the data, and, thereafter, to filter the residual sequence. The low-frequency filtered sequence can be added back to the trend, if it is required to represent the trend-cycle.

An alternative to removing a linear trend is to apply a twofold difference operator to the data. The differenced data can be filtered and, thereafter, they can be reinflated by a double summation, which represents the inverse of the differencing operation. Such a summation requires the provision of some initial conditions. An exposition of this method has been provided by Pollock (2007).

The requirement for an explicit estimation of the initial conditions can be avoided if attention is concentrated on the highpass filter. The initial conditions that are required for the inflation of a differenced sequence that has been subjected to a highpass filter are nothing other than zero values. The complementary lowpass sequence can be obtained by subtracting the inflated product of the highpass filter from the data. This subtraction procedure is already manifest in equation (2), where the residual vector e represents the highpass component.

It is perhaps remarkable that, given the appropriate conditions, all three methods of dealing with the problems of a trended sequence are algebraically equivalent. The equivalence of the subtraction procedure and the procedure of polynomial detrending will be illustrated hereafter.

The W-K filter that is most familiar to econometricians is undoubtedly the so-called Hodrick–Prescott filter, (described in Hodrick and Prescott, 1980, 1997 and properly attributable to Conrad Leser 1961.) This is a simple filter comprising a single adjustable parameter λ , which is the smoothing parameter. The equation of the lowpass time-varying filter is

$$\begin{aligned} x &= y - Q(\lambda^{-1}I + Q'Q)^{-1}Q'y \\ &= y - h, \end{aligned} \tag{9}$$

where h represents the highpass component. It will be observed that, as $\lambda \rightarrow \infty$, the equation converges on that of the linear detrending regression, represented by equation (2).

Given that the matrix transformation of (9) has an order that is equal to the size of the sample, care must be taken to economise on the use of the memory of the computer. This can be done by exploiting the fact that the component matrices have a limited number of adjacent diagonal bands.

ECONOMETRIC FILTERS

First, the differenced vector $d = Q'y$ and the matrix $W = \lambda^{-1}I + Q'Q$ of five diagonal bands may be formed. Then, the equation $d = Wb$ is solved for $b = (\lambda^{-1}I + Q'Q)^{-1}d$. This is achieved via a Cholesky factorisation that sets $W = GG'$, where G is a lower triangular matrix of three nonzero bands. The equation $GG'b = d$ may be cast in the form of $Gp = d$ and solved recursively for p . Then, $G'b = p$ can be solved for b by backsubstitution. It is then straightforward to calculate $x = y - Qb$.

Figure 11 shows the frequency response functions of time-invariant versions of the lowpass Hodrick–Prescott filter. Proceeding from the innermost curve, the corresponding values of the smoothing parameter λ are 14,400, 1,600 and 100, which are the values commonly prescribed for monthly, quarterly and annual data, respectively. In all cases, there is only a gradual transition from the pass band to the stop band. The consequence is that the filter is unable clearly to isolate spectral structures that lie within strictly limited frequency bands.

In the case of quarterly data, such as the consumption data, the recommended value of the smoothing parameter is 1,600. As is evident in Figure 12, this value is too great for the purpose of extracting the trend-cycle from the consumption data, since it results in a function that is too inflexible. This is confirmed by comparing Figure 12 with Figure 5. A more appropriate value for the parameter would be 100, which is the value recommended for annual data. This works adequately with the tractable consumption data.

What is often required in place of the H-P filter is a filter that gives a clearer demarcation between the stop band and the pass band. The point at which the transition occurs should be freely specified by the user in the light of the spectral structure of the detrended data, which is revealed by the periodogram.

A Wiener–Kolmogorov filter that goes some way towards achieving this is the Butterworth filter. This filter, which was conceived, originally, by the British physicist Stephen Butterworth (1930) as an analogue filter, is common in electrical engineering. The digital version has been described in an econometric context by Pollock (2000).

The Butterworth filter that is appropriate to short trended sequences can be represented by the equation

$$x = y - \Sigma Q(\lambda^{-1}M + Q'\Sigma Q)^{-1}Q'y. \quad (10)$$

Here, the matrices are

$$\Sigma = \{2I_T - (L_T + L'_T)\}^{n-2} \quad \text{and} \quad M = \{2I_T + (L_T + L'_T)\}^n, \quad (11)$$

where L_T is the lag-operator matrix of order T , which has units on the first subdiagonal and zeros elsewhere. It can be verified that

$$Q'\Sigma Q = \{2I_T - (L_T + L'_T)\}^n. \quad (12)$$

This filter has two parameters that can be chosen at will. The first parameter to be chosen is the order n of the filter. The higher is the order of the

filter, the more rapid is the transition from the pass band to the stop band. The second parameter is the nominal cut-off point ω_c , which is the midpoint in the transition. The cut-off point is mapped to the smoothing parameter via the function $\lambda = \{1/\tan(\omega_c/2)\}^n$.

Figure 13 shows the frequency responses of the Butterworth filters of orders 6 and 12 with a nominal cut-off point of $\pi/6$ radians or 30 degrees. These are both quite adequate for isolating the low-frequency structure that is evident in Figure 3 and which has been identified with the business cycle.

The spectral structure in question extends no further in frequency than $\pi/8$ radians, which is 22.5 degrees. The Butterworth filter with a nominal cut off point at a slightly higher value and with a reasonably rapid transition serves the purpose well enough. In consequence of the succeeding dead space, it does not contaminate the business cycle estimate with any significant extraneous elements.

It should now be observed that, in deploying the finite-sample Wiener–Kolmogorov filters, it makes no difference to the calculation of the highpass component h whether it is the original data vector y or the residual vector $e = Py$ from a linear detrending that is subject to the filtering.

The matrix, within equation (2) that maps from y to $e = Py$ is $P = Q(Q'Q)^{-1}Q'$. The matrix of the highpass Hodrick–Prescott filter is $H = Q(\lambda^{-1}I + Q'Q)^{-1}Q'$. It can be seen that $HP = H$ and, therefore, that $h = Hy = HPy = He$.

Moreover,

$$(I - H)y = (I - H)Py + (I - P)y, \quad (13)$$

which is to say that the Hodrick–Prescott trend (or trend-cycle) can be calculated by adding the filtered residuals to the linear trend. An analogous identity arises when the matrix H is replaced by the matrix $B = \Sigma Q(\lambda^{-1}M + Q'\Sigma Q)^{-1}Q'$ of the Butterworth filter.

6. Frequency-Domain Filters

Components of the data that have well-defined spectral structures can be isolated by synthesising them from their spectral ordinates. Thus, in implementing a bandpass filter that is intended to capture a component that lies within a specific range of frequencies, the spectral elements that fall within the corresponding pass band should be preserved and those elements that lie within the stop band should be nullified, or replaced by zeros.

This is not the only thing that can be achieved by operating directly in the frequency domain. Any required frequency response can be realised, simply by multiplying the spectral elements by the appropriate factors that are indicated by the response function.

An example is provided by the linearly detrended logarithmic consumption data. The objective is to isolate the business cycle, which has the spectral structure that is displayed in Figure 3 that falls in the frequency interval $[0, \pi/8]$. In terms of equation (1), the business cycle component is synthesised by run-

ning the summation up to the index q for which the associated frequency value $\omega_q = 2\pi q/T = q \times \omega_1$ is closest to $\pi/8 = \beta$.

The frequency-domain filter has a time-domain representation that may be compared with the filter of Christiano and Fitzgerald, specialised to the case where $\alpha = 0$ and $\beta = \pi/8$. In that case, the elements of the symmetric Toeplitz matrix Ψ of the mapping $x = \Psi d$ from the detrended data vector d to the filtered vector x are provided by the sinc function:

$$\psi_k = \begin{cases} \beta, & \text{if } k = 0, \\ \frac{\sin(\beta k)}{\pi k}, & \text{if } k \neq 0, \end{cases} \quad (14)$$

where k is the index of the supra-diagonal and sub-diagonal bands.

The time-domain representation of the frequency-domain filter entails a circulant matrix Ψ° in place of the Toeplitz matrix Ψ . Its diagonal elements are provided by the Dirichlet kernel:

$$\psi_k^\circ = \begin{cases} (2q + 1)/T, & \text{if } k = 0, \\ \frac{\sin([q + 1/2]\omega_1 k)}{T \sin(\omega_1 k/2)}, & \text{if } k \neq 0. \end{cases} \quad (15)$$

The kernel, which is also described as an aliased sinc function, represents the Fourier transform of a set of values sampled from the frequency-domain rectangle defined on the interval $[-\beta, \beta]$. The effect of the sampling is to wrap the sinc function around a circle of circumference T and to add its overlying ordinates. A derivation has been provided by Pollock (2009b).

The filtered values would be obtained by the circular convolution of the data with the coefficients of the filter; and the effect would be the same as that of applying the sinc function to an indefinite periodic extension of the data sequence by a linear convolution. The circular convolution can be represented by the matrix equation $x = \Psi^\circ y$. Here, the structure of the symmetric circulant matrix may be illustrated adequately by the case where $T = 4$:

$$\Psi^\circ = \begin{bmatrix} \psi_0^\circ & \psi_1^\circ & \psi_2^\circ & \psi_1^\circ \\ \psi_1^\circ & \psi_0^\circ & \psi_1^\circ & \psi_2^\circ \\ \psi_2^\circ & \psi_1^\circ & \psi_0^\circ & \psi_1^\circ \\ \psi_1^\circ & \psi_2^\circ & \psi_1^\circ & \psi_0^\circ \end{bmatrix}. \quad (16)$$

In the process of a circular convolution, the data are treated as a circular sequence, with the effect that the filtered values towards the end of the sequence are liable to be formed partly from data values at the beginning of the sequence—and vice versa for the filtered values at the beginning the sequence.

There can be problems if the beginning and the end of the data do not join seamlessly, as they appear to do in the case of the detrended logarithmic consumption data of Figure 14, through which the business cycle is interpolated. Therefore, in the next section, we shall outline a recourse that is effective in overcoming such problems.




Figure 14. The residual sequence from fitting a linear trend to the logarithmic consumption data with an interpolated line representing the business cycle, obtained by the frequency-domain method.




Figure 15. The turning points of the business cycle marked on the horizontal axis by black dots. The solid line is the business cycle of Figure 14. The broken line is the derivative function.

The business cycle that is represented in Figure 14 is formed from a Fourier synthesis based on trigonometrical functions with frequencies $\omega_j = 2\pi j/T$, with $T = 160$ and $j = 0, 1, \dots, 9$, that fall in the interval $[0, \pi/8)$ and with squared amplitudes ρ_j^2 that are provided by the periodogram of the data. However, in contrast to the Fourier synthesis of equation (1), wherein the temporal index t takes integer values, t is now a continuous variable,

As a sum of continuous trigonometrical functions, the business-cycle trajectory is an analytic function of which the derivatives exist of all orders. This implies that it is straightforward to find the maxima and minima of the function, and hence the turning points of the business cycle, simply by identifying the points where the first derivative is zero-valued. The simplicity of this procedure contrasts markedly with the complexity of some other well-known procedures for locating the turning points of the business cycle, such as that of Bry and

Boschan (1971).

Figure 15 shows the function that is obtained by differentiating the business cycle function of Figure 14. The turning points of the business cycle are marked by dots on the horizontal axis. Also plotted on the diagram is a line that is parallel to the horizontal axis, depressed by a distance that corresponds to the slope of the log-linear trend line of Figure 1, which represents the underlying rate of growth of U.K. consumption.

The intersection of the derivative function with this line indicates the turning points of the trend-cycle function that is obtained by adding the trajectory of the business cycle to the linear trend. It will be observed that the majority of the business-cycle turning points are absent from the trend-cycle function, wherein they have become points of inflection. Compared with those of the business cycle, the downturns of the trend-cycle are postponed and its upturns come sooner.

The function of Figure 14 has flat tops and broad valleys. Although the method of derivatives attributes exact dates to the turning points, this suggests that one should not expect great precision in such dates. There is undoubtedly a common desire to determine precise dates for the turning points and to identify them with economic events. However, this may be an unreasonable demand, since the index of consumption is an aggregate of numerous trajectories that turn at different times.

In practice, the turning points are too often located by the conventional dating methods at the pits and the pinnacles in the aggregate index that are the result of its inadequate seasonal adjustment. Nevertheless, the dates that are displayed on the horizontal axis of Figure 15 should enable the reader to interpret it in the light of a knowledge of the recent economic history of the U.K.

7. Monthly Seasonal Data

A more exacting exercise, for which the time-domain filters are barely adequate, concerns the extraction of the trend from a monthly sequence of the logarithms of the U.S. money supply. Figure 16 shows the logarithmic money supply data, through which a quadratic trend has been interpolated via a least-squares regression.

The data are affected by a marked pattern of seasonal variation that entails elements in the vicinity of the fundamental seasonal frequency of $\pi/6$ radians or 30 degrees and in the vicinities of the various harmonic frequencies of $\pi/3$, $\pi/2$, $2\pi/3$, $5\pi/6$ and π .

Figure 17 displays the periodogram of the quadratically detrended data sequence. There is evidence here of a low-frequency component that extends almost to the seasonal frequency. The shaded area covers this component. If the low-frequency component is to be isolated, then a filter is required of which the transition from pass band to stop band occurs at a point. For this, a frequency-domain filter is required.

The end-of-sample problem, as previously described, does not arise with circular wrapping or, equivalently, with the periodic extension of the data that

ECONOMETRIC FILTERS




Figure 16. The plot of 132 monthly observations on the logarithms of the U.S. money supply, beginning in January 1960. A quadratic function has been interpolated through the data.




Figure 17. The periodogram of the residuals from the quadratic detrending of the logarithmic money-supply data.




Figure 18. The residuals from a linear detrending of the money-supply data, with an interpolation of four years length inserted between the end and the beginning of the circularised sequence, marked by the shaded band.

is entailed in a Fourier analysis. However, as we have already indicated, a problem can arise with trended data where the end of one replication of the sample, where the values are at a maximum, joins the start of the succeeding replication, where the values are at a minimum. The resulting disjunctions give rise to a saw tooth function.

The problem of the disjunction, which is acute when the data are trended, can arise even when the trend has been removed, since the beginning and the end of the sample may not meet at the same level. The most common recourse for overcoming this problem is to taper the ends of the data so that they are both reduced to zero. However, this tends to falsify the data.

An alternative recourse is to interpolate a section pseudo data between the end and the beginning that will effect a smooth transition. At an appropriate stage, the pseudo data can be discarded. In the case of data that show a strong seasonal variation, which may have evolved over the course of the sample, it is appropriate to construct a segment of pseudo data by morphing the pattern of seasonal variation so that it changes from the pattern at the end of the data sequence to the pattern at the start.

Each point of the pseudo data will be a convex combination of a point within the final pattern and a point within the initial pattern. The weights of the combinations will vary between unity and zero. The weight on the points in the final pattern will be close to unity near the start of the segment of pseudo data, and they will become close to zero near the end. Their trajectory is governed by a half cycle of a raised cosine function: $\{\cos(\omega) + 1\}/2$, with $\omega \in [0, \pi]$.

Figure 18 shows the segment of pseudo data that has been interpolated into the circularised sequence of the residuals from a quadratic detrending of the logarithmic money supply data. From this augmented data sequence, a low-frequency cycle is estimated by the frequency-domain method. This is added to the quadratic trend to create the trend-cycle function that is plotted in Figure 19. Figure 20 shows the deviations of the data from this function.

8. Interrupted Trends

There have been wide differences of opinion in the econometrics literature on how a trend should be defined and on how it should be extracted from the data. It seems appropriate to approach this matter with an open mind. The definition of the trend may be influenced by the characteristics of the data, by the objectives of the analysis and by the methodological and aesthetic preferences of the analyst.

The preference expressed in this paper has been for a trend function that represents a firm benchmark against which the cyclical fluctuations of the economy may be measured. In periods of sustained economic growth, the trend can be represented by a polynomial function.

The periodogram of the detrended data often shows a clear spectral signature of the business cycle that can guide its extraction. By adding the business cycle to the polynomial trend, a trend-cycle component can be estimated that

ECONOMETRIC FILTERS




Figure 19. The plot of the logarithms of 132 monthly observations on the U.S. money supply, beginning in January 1960. A trend-cycle, estimated by the Fourier method, has been interpolated through the data.




Figure 20. The sequence of residual deviations of the logarithmic money supply data from the estimated trend-cycle function.




Figure 21. The logarithms of annual U.K. real GDP from 1873 to 2001 with an interpolated trend. The trend is estimated via a filter with a variable smoothing parameter.

can provide a benchmark against which to measure the seasonal fluctuations of the data.

Sometimes, there are major interruptions that halt the steady progress of the economy and which can give rise to wide deviations from an interpolated polynomial trend. If such interruptions are deemed to have an enduring effect on the underlying trajectory of the economy, then it may be appropriate to describe them as structural breaks and to absorb them into the trend.

A device that will serve this purpose is a form of the Hodrick–Prescott filter in which the smoothing parameter can take different values in different localities. In the vicinity of the break, the smoothing parameter can be set to a sufficiently low value to allow the function to absorb the break. Elsewhere, it should be set to a high value to make it sufficiently stiff to prevent it from absorbing the cyclical fluctuations of the data.

Figure 21 shows the logarithms of the annual real GDP of the UK from 1873 to 2001. The value of the smoothing parameter has been reduced radically within the highlighted regions in order to absorb the effects of the economic recessions that followed the two world wars. Elsewhere, the parameter has been given a high value to generate a stiff curve. In particular, no attempt has been made to accommodate the downturn of the recession of 1929. An alternative purpose would be to show the full extent of these three interruptions. For that purpose, one might fit a polynomial function of degree four or more to the data.

9. Filtering and the Autocovariances

The act of filtering invariably alters the spectral structure and the autocovariance structure of the data. Wallis (1974) noted this fact in a paper concerning the effect of seasonal adjustment on the relationships between variables.

Subsequent authors have investigated the effects of seasonal adjustment upon the outcomes of statistical tests, such as the tests for the presence of unit roots in ARIMA models (see Diebold 1993). Others have shown the effects of seasonal adjustment on the estimation of dynamic models (see, for example, del Barrio Castro and Osborn 2004).

There are clear advantages in employing filters with ideal or rectangular frequency responses, such as the frequency-domain filters that have been advocated in this paper. Such filters do not alter the frequency components of the data that fall within their pass bands. Moreover, the covariance relationships between the elements of two sequences that have been subjected to such ideal filters are unaltered.

However, the filtering is bound to affect the usual methods for estimating dynamic models. Indeed, given that frequency-limited sequences are associated with rank-deficient autocovariance matrices, an essential assumption of the ARMA processes is falsified.

To understand the appropriate recourse in such cases, we may consider the effect of applying an ideal lowpass filter to a white-noise sequence. This will give rise to a sequence of which the autocovariance function contains the sampled ordinates of a sinc function. Figure 22 shows the autocovariance function of a




Figure 22. The central coefficients of the Fourier transform of the frequency response of an ideal lowpass filter with a cut-off point at $\omega = \pi/2$, which are sampled from the sinc function $\sin(\pi t/2)/\pi t$.




Figure 23. The sinc function $\sin(\pi t)/\pi t$ comprising frequencies in the interval $[0, \pi]$.

white-noise sequence that has been subjected to an ideal lowpass filter with a cut-off point at $\omega = \pi/2$.

(The ideal frequency response, which is a rectangle on the interval $[-\pi/2, \pi/2]$, is a symmetric and idempotent function, and the spectrum of the filtered process is the same rectangle scaled by the variance of the white noise. Therefore, the autocovariance function, which is the (inverse) Fourier transform of the spectrum, is a sampled sinc function scaled by the white-noise variance.)

In order to transform the filtered data into a sequence that has the covariance properties of white noise, it is necessary to expand the spectrum of the filtered data, which is a rectangle on the interval $[-\pi/2, \pi/2]$, into a rectangle on the Nyquist interval of $[-\pi, \pi]$. This is readily accomplished by subsampling the data by discarding alternate data points.

The autocovariance function of the resulting sequence is shown in Figure 23. This function also comprises ordinates sampled from a sinc function. How-

ever, in this case, the ordinates, apart from the central one, coincide with the zeros of the sinc function, which is to say that they correspond to the autocovariances of a white-noise process.

The simplicity of this procedure is due to the fact that the cut-off frequency of $\omega_c = \pi/2$ divides the Nyquist frequency of π an integral number of times. In general, a more complicated procedure of subsampling will be required. This leads one to consider a continuous version of the data trajectory created by the method of Fourier synthesis. The continuous function will need to be subsampled by taking points separated by time intervals of π/ω_c . In general, these points will not coincide with any of the original data points.

The subsampling procedure can be applied to any isolated spectral structure. In the case of the structure on the interval $[0, \pi/8)$ in Figure 3, one in every 8 data points should be taken; and these will allow an ARMA model to be fitted. A truth that some have found hard to accept is that this process of subsampling entails no loss of the information. In fact, the filtered data can be recovered readily from the subsampled data.

10. Summary and Conclusions

This paper has presented several alternative methods for filtering economic data that may be used in estimating of their trends. These methods have been implemented in a computer program called IDEOLOG that is freely available at the following address:

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

The code of the program, which is in Pascal, has also been provided. This should assist the implementation of the methods within alternative environments.

The purpose of the program has been both to demonstrate the favoured methods and to highlight the problems that may arise with the other methods. For the favoured methods, a set of documents, described as log files, have been produced that guide the user through the operations that are involved in processing an accompanying collection of data sets.

The Wiener–Kolmogorov Butterworth square-wave filter, which was described by Pollock (2000), is amongst the favoured methods. The computer code for this filter has also been translated into the C language and incorporated in the *Gretl* program of Cottrell and Lucchetti (2015).

The two parameters of the Butterworth filter, which are the filter order and the nominal cut-off frequency, can be used to specify a version of the filter that has a rapid transition from the pass band to the stop band in the vicinity of an appropriate cut-off point. A combination of two Butterworth filters, applied in succession as a lowpass filter and a highpass filter should serve to create an effective bandpass filter.

The IDEOLOG program also implements the Hodrick–Prescott (H-P) filter or Leser filter, which is a Wiener–Kolmogorov filter. The misuse of this filter has been criticised by Pollock (2000), who has shown that it has insufficient flexibility to enable it to isolate components with well-defined spectral

structures within bounded frequency ranges.

Nevertheless, the H-P filter finds a use in the IDEOLOG program, which allows the smoothing parameter to take different values in different locations. This facility enables the filter to generate a stiff trend function that incorporates sharp bends in places where there appear to be a structural breaks in the processes generating the data.

The frequency-domain filters are also amongst the favoured methods, An essential requirement for their successful application is that the data should form a circular sequence in which the head and the tail are joined seamlessly. This can be achieved by ensuring that the preliminary trend function interpolates the data at both ends.

For this purpose, the program provides a method of polynomial regression that places arbitrarily large weights on the end points and on the adjacent points. However, the method that has been illustrated in the paper, which has proved to be invariably successful, interpolates a segment of pseudo data between the beginning and the end of the circularly wrapped data.

A further innovation associated with the frequency-domain filters is a new methodology for determining the turning points of a continuous frequency-limited function based on a Fourier synthesis. The analytic nature of such a function makes it amenable to the methods of differential calculus.

It may be claimed that the frequency-domain methods, together with their supporting devices, have a greater flexibility than any other methods of filtering. This is because they allow an arbitrary choice of weights to be applied to the Fourier ordinates of the data to produce whatever frequency response is required. In the frequency-domain filters of IDEOLOG, the weights are restricted to be zeros of units, which gives rise to rectangular pass bands.

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.
- Bry, G., and C. Boschan, (1971), *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, New York, National Bureau of Economic Research.
- Burns, A.M., and W.C. Mitchell, (1946), *Measuring Business Cycles*, New York, National Bureau of Economic Research. Available in electronic form at the address <http://papers.nber.org/books/burn46-1>
- Butterworth, S., (1930), On the Theory of Filter Amplifiers, *Experimental Wireless and the Wireless Engineer*, 7, 536–541.
- Christiano, L.J., and T.J. Fitzgerald, (2003), The Band-pass Filter, *International Economic Review*, 44, 435–465.
- Cottrell, A., and R. Lucchetti, (2015), *Gretl: Gnu Regression, Econometrics and Time-series Library*. Freely available for the Linux, Windows and Macintosh OS X platforms at the address <http://gretl.sourceforge.net/>

ECONOMETRIC FILTERS

- del Barrio Castro, T., and Denise R. Osborn, (2004), The Consequences of Seasonal Adjustment for Periodic Autoregressive Processes, *The Econometrics Journal*, 7, 307–321.
- Diebold, F.X., (1993), The Effect of Seasonal Adjustment Filters on Tests for a Unit Root, *Journal of Econometrics*. — 55 99–103.
- Doherty, M., (2001), The Surrogate Henderson Filters in X-11, *The Australian and New Zealand Journal of Statistics*, 43, 385–392.
- Henderson, R., (1916), Note on Graduation by Adjusted Average, *Transactions of the Actuarial Society of America*, 17, 43–48.
- Henderson, R., (1924), A New Method of Graduation, *Transactions of the Actuarial Society of America* 25, 29–40.
- Hodrick, R.J., and E.C. Prescott, (1980), *Postwar U.S. Business Cycles: An Empirical Investigation*, Working Paper, Carnegie–Mellon University, Pittsburgh, Pennsylvania.
- 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.
- Kenny, P.B., and J. Durbin, (1982), Local Trend Estimation and Seasonal Adjustment of Economic and Social Time Series, *Journal of the Royal Statistical Society, Series A* , 145, 1–41.
- Ladiray, D., and B. Quenneville, (2001), *Seasonal Adjustment with the X-11 Method*, New York, Springer-Verlag.
- Leser, C.E.V., (1961), A Simple Method of Trend Construction, *Journal of the Royal Statistical Society, Series B*, 23, 91–107.
- Musgrave, J.C., (1964a), A Set of End Weights to End all End Weights. *Unpublished Working Paper of the U.S. Bureau of Commerce*.
- Musgrave, J.C., (1964b), Alternative Sets of Weights Proposed for X-11 Seasonal Factor Curve Moving Averages, *Unpublished Working Paper of the U.S. Bureau of Commerce*.
- Pollock, D.S.G., (2000), Trend Estimation and De-Trending via Rational Square Wave Filters, *Journal of Econometrics*, 99, 317–334.
- Pollock, D.S.G., (2007), Wiener–Kolmogorov Filtering, Frequency Selective Filtering and Polynomial Regression, *Econometric Theory*, 23, 71–83.
- Pollock, D.S.G., (2009a), Statistical Signal Extraction: A Partial Survey, Chapter 9 in D.A. Belsley and E.J. Kontoghiorghes (eds.) *Handbook of Computational Econometrics*, John Wiley and Sons, Chichester.
- Pollock, D.S.G., (2009b), Realisations of Finite-sample Frequency-selective Filters, *Journal of Statistical Planning and Inference*, 139, 1541–1558.
- Wallis, K.F., (1974), Seasonal Adjustment and Relations Between Variables, *Journal of the American Statistical Association*, 69, 18–31.