Department of Economics

Econometric Methods of Signal Extraction

Stephen Pollock

Working Paper No. 530 May 2005 ISSN 1473-0278

Queen Mary
University of London

ECONOMETRIC METHODS OF SIGNAL EXTRACTION

by D.S.G. POLLOCK

Queen Mary, University of London

Email: stephen_pollock@sigmapi.u-net.com

The Wiener–Kolmogorov signal extraction filters, which are widely used in econometric analysis, are constructed on the basis of statistical models of the processes generating the data. The models may be heuristic devices that can be specified in whichever ways are appropriate to ensure that the filters have the desired characteristics. The digital Butterworth filters, which are described and illustrated in the paper, are specified in this way. The components of an econometric time series often give rise to spectral structures that fall within well-defined frequency bands that are isolated from each other by spectral dead spaces. In such cases, we are inclined to use a Fourier-based method that operates in the frequency domain. This new method can be assimilated to a finite-sample Wiener–Kolmogorov framework

JEL Classification: C22

Keywords: Signal extraction, Linear Filtering, Frequency-domain analysis, Trend estimation

1. Introduction

The econometric methods of signal extraction that are based on linear filters have attained a high level of sophistication. They have now coalesced into two distinct categories.

In the first category are the methods that are favoured by the central statistical agencies of many of the OECD counties. These were developed in the North American agencies, notably in Statistics Canada (Dagum 1980) and in the U.S. Bureau of the Census (Findley *et al.* 1998), and they have been supported and refined in other agencies throughout Europe and in the Antipodes. The methods are based upon a variety of moving-average smoothing filters, and they are used, principally, in trend extraction and in deseasonalising data series.

The success of the methods of the central statistical agencies has entailed a widespread acceptance of a set of common conventions and definitions. Thus, it is commonly agreed that a deseasonalised data series can be defined, fairly and simply, as the product of the relevant methods of the central statistical office.

However, the acceptance of such conventions, whilst necessary to ensure comparability of statistics across nations, inhibits scientific innovation and discovery. Therefore, academic interest has been focused mainly on the so-called model-based procedures, which constitute the second category of econometric signal-extraction methodology.

The model-based approaches are derived from the idea that the components of an econometric times series, which are its trend, its secular cycles, its seasonal cycles and its irregular component, can all be modelled by autoregressive integrated moving-average (ARIMA) processes of low orders. There are several exponents of this genre who have taken slightly different approaches.

In the structural approach, which originated with Harrison and Stevens (1976) and which has been developed by Harvey (1989) and others, such models are fitted jointly to the data in a manner that renders their parameters readily accessible. (These methods have been implemented in the STAMP program of Koopman et al. 2000.)

In the alternative canonical approach, which has been advocated by Hillmer and Tiao (1982) and by Maravall and Pierce (1987), amongst others, the models of the individual components must be disentangled from a fitted ARIMA model that represents their joint effects. (An accessible implementation is in the SEATS–TRAMO program of Gómez and Maravall, 1994, 1996—the DOS version—and of Caporello and A. Maravall 2004—the Windows version.)

Another approach, which combines aspects of both the structural approach and the canonical approach, has been pursued in the CAPTAIN MatLab Toolbox program of Young et al. (2004). An account of some of the features of this program and of its uses has been provided by Young, Pedregal and Tych (1999).

The model-based procedures have the seeming advantage that they subsist within a framework that facilitates conventional statistical inference. Thus, for example, confidence intervals are easily generated that can surround the estimated data components. However, the validity of such inferences depends crucially upon the cogency of the linear time-invariant ARIMA models that are applied to the components. The ability of such models to reflect the underlying data structures is limited. In particular, the models are liable to be subverted whenever the structures show any significant tendency to evolve through time.

In this paper, we shall also use statistical models in deriving the filters that are used to isolate the data components, but the models will be treated mainly as heuristic devices that will allow us to exploit the mathematical formalisms of the Wiener–Kolmogorov theory of signal extraction. This theory indicates that the optimal estimates of the data components are provided by their conditional expectations that are formed in the light of the observed data and of the models that are presumed to have generated them.

When the models themselves are to be fitted to the data, it is important that they should be realistic—otherwise they will provide an insecure basis for forming the supposedly optimal filters. However, in practice, they rarely achieve much realism. When the models are used merely as heuristic devices, they can be specified in whichever ways are appropriate to ensure that the filters have the desired characteristics. Moreover, the desirable characteristics will be unaffected, in the main, by the evolutions of the data components that the resulting filters are designed to isolate.

The original Wiener–Kolmogorov theory was developed under the fictional assumption that the data are generated by a stationary stochastic process and that they form a doubly infinite sequence. In econometrics, one has to contend with short nonstationary sequences; and, to cope with these, it is common to resort to the Kalman filter.

The Kalman filter and the associated smoothing algorithms are complicated and powerful devices, of which the workings can often seem obscure. The difficulty can be attributed to the all-encompassing nature of the algorithms. In this paper, we shall pursue a simpler approach that fulfils the same objective of obtaining quasi minimum-mean-square-error estimates, but which deals directly with the specific features of the problem at hand. We shall use the finite-sample version of the bidirectional Wiener–Kolmogorov filter that has been expounded in previous papers of the present author (see Pollock, 1997, 2000, 2001, and 2002).

One of the contentions of this paper is that the components of an econometric time series often give rise to spectral structures that fall within well-defined frequency bands that are isolated from each other by spectral dead spaces, wherein the are no Fourier ordinates of any significant magnitude. This leads us to consider the nature of band-limited stochastic processes, which are characterised by singular dispersion matrices. We find that the finite-sample Wiener–Kolmogorov formulation lends itself readily to a specialisation that is appropriate for dealing with band-limited components.

2. Filtering Short Stationary Sequences

We begin by considering the problem of estimating the signal component $\xi(t)$ and the noise component $\eta(t)$ of a data sequence

$$y(t) = \xi(t) + \eta(t), \tag{1}$$

where $t \in \{0, \pm 1, \pm 2, \ldots\}$ is the index of the discrete-time observations. According to the classical assumptions, which we shall later amend in various ways, the signal and the noise are generated by stationary stochastic processes that are mutually independent. It follows that the autocovariance generating function of the data, defined by

$$\gamma(z) = \gamma_0 + \sum_{\tau=1}^{\infty} \gamma_\tau (z^\tau + z^{-\tau}), \tag{2}$$

is the sum of the autocovariance generating functions of its components. Thus

$$\gamma(z) = \gamma_\xi(z) + \gamma_\eta(z). \tag{3}$$

In practice, the available data will form a finite sequence, which constitutes a vector $y = [y_0, y_1, \ldots, y_{T-1}]'$ with a signal component ξ and a noise component η such that

$$y = \xi + \eta. \tag{4}$$

The data might owe their stationarity to a prior differencing operation or to an operation that has involved the extraction of a trend from the original data and the retention of the residue. For these vectors, the moment matrices are

$$\begin{aligned} E(\xi) = 0, \qquad & D(\xi) = \Omega_\xi, \\ E(\eta) = 0, \qquad & D(\eta) = \Omega_\eta, \\ \text{and} \quad & C(\xi, \eta) = 0. \end{aligned} \tag{5}$$

The independence of ξ and η implies that $D(y) = \Omega = \Omega_\xi + \Omega_\eta$.

Here, the various variance–covariance or dispersion matrices, which have a Toeplitz structure, may be obtained by replacing the argument z within the relevant generating function by the matrix

$$L_T = [e_1, \ldots, e_{T-1}, 0], \tag{6}$$

which is obtained from the identity matrix $I_T = [e_0, e_1, \ldots, e_{T-1}]$ by deleting the leading column and appending a column of zeros to the end of the array. The matrix L_T, which has units on the first subdiagonal and zeros elsewhere, is the finite-sample version of the lag operator. Using it in place of z in $\gamma(z)$ gives

$$D(y) = \Omega = \gamma_0 I + \sum_{\tau=1}^{T-1} \gamma_\tau (L_T^\tau + F_T^\tau), \tag{7}$$

where $F_T = L_T'$ is in place of z^{-1}. Since L_T and F_T are nilpotent of degree T, such that $L_T^q, F_T^q = 0$ when $q \geq T$, the index of summation has an upper limit of $T - 1$.

The optimal predictors of the components ξ and η are their conditional expectations, denoted by x and h, respectively, in (8) and (9):

$$\begin{aligned} E(\xi|y) &= E(\xi) + C(\xi, y) D^{-1}(y)\{y - E(y)\} \\ &= \Omega_\xi(\Omega_\xi + \Omega_\eta)^{-1} y = x, \end{aligned} \tag{8}$$

$$\begin{aligned} E(\eta|y) &= E(\eta) + C(\eta, y) D^{-1}(y)\{y - E(y)\} \\ &= \Omega_\eta(\Omega_\xi + \Omega_\eta)^{-1} y = h. \end{aligned} \tag{9}$$

These are the finite-sample versions of the so-called Wiener–Kolmogorov estimates. They constitute the minimum-mean-square-error estimates of the components, under the assumption that the specifications in (5) are correct. The assumptions provide a set of ordinary positive-definite variance–covariance matrices that pertain to conventional linear stochastic processes, which have spectra that extend across the frequency range.

We should observe that adding the estimates gives

$$y = x + h, \tag{10}$$

which is to say that the estimated components add up to the data vector y, as do the true components ξ and η in equation (4).

To investigate the mapping from y to x or, equally, the mapping from y to h, we must take account of the various symmetries manifested by the Toeplitz matrices Ω_η and Ω_ξ. The generic Toeplitz matrix Ω is symmetric about the principal (northwest-southeast) diagonal, which is ordinary symmetry. It is symmetric about the secondary (northeast-southwest) diagonal, which is persymmetry. It is invariant with respect to rotations of 180° around the central point at the intersection of its two diagonals, which is centrosymmetry.

Let $H = [e_{T-1}, \ldots, e_1, e_0]$ be the counter-identity matrix, which has units on the secondary diagonal and zeros elsewhere, and let $\Omega^\#$ be the counter transpose, which is the reflection of Ω about the secondary diagonal. Then, the symmetries of Ω may be recorded as follows:

$$\begin{array}{lll} \text{(i)} & \text{Symmetry:} & \Omega = \Omega', \\ \text{(ii)} & \text{Persymmetry:} & \Omega = \Omega^\#, \quad \text{equivalently} \quad H\Omega H = \Omega', \\ \text{(iii)} & \text{Centrosymmetry:} & \Omega = (\Omega^\#)' = \Omega^{\mathrm{r}}, \quad \text{equivalently} \quad H\Omega H = \Omega. \end{array} \tag{11}$$

The matrix of $x = Zy$, which is the estimating equation of the signal ξ, is determined by the equation $\Omega_\xi = Z\Omega$, wherein both Ω_ξ and $\Omega = \Omega_\xi + \Omega_\eta$ are Toeplitz matrices. Therefore, since $H\Omega_\xi H = \Omega_\xi$, $H\Omega H = \Omega$ and $HH = I$, it follows that

$$\Omega_\xi = H\Omega_\xi H = \{HZH\}\{H\Omega H\} = \{HZH\}\Omega = Z\Omega. \tag{12}$$

In view of the nonsingularity of the factors, we conclude from this that $HZH = Z$, which is to say that $Z = \Omega_\xi(\Omega_\xi + \Omega_\eta)^{-1}$ is a centrosymmetric matrix, albeit that it is not a Toeplitz matrix.

Let y^{r} and x^{r} be y and x in reverse. Then, the centrosymmetric property of Z ensures that both $x = Zy$ and $x^{\mathrm{r}} = Zy^{\mathrm{r}}$. This feature is in accordance with the fact that the direction of time can be reversed without affecting the statistical properties of a stationary process.

The filter weights that are provided by the rows of the matrices Z vary as the filter progresses through the sample. As the sample size increases, the weights in the central row of Z, when it has an odd number of rows, will tend to the set of constant coefficients that would be derived under the assumption of a doubly-infinite data sequence. These coefficients are symmetric about a central value.

The weights of the final row of Z correspond to the coefficients of a one-sided causal filter, whereas those in the first row correspond to the same filter working in reversed time. As the sample size increases, the weights of the final row tend to the coefficients of the expansion of the rational feedback filter that represents the real-time component of the classical bidirectional Wiener–Kolmogorov filter that is appropriate to infinite samples.

In calculating the estimates, we should avoid inverting any of the matrices directly, since they are of the order of the sample size T, which is liable to be large. Therefore, the centrosymmetric filtering matrix Z is to be regarded as a theoretical entity rather than a practical one. A simple procedure begins by solving the equation

$$(\Omega_\xi + \Omega_\eta)b = y \tag{13}$$

for the value of b. Thereafter, we can find

$$x = \Omega_\xi b \qquad \text{and} \qquad h = \Omega_\eta b. \tag{14}$$

As will be shown, it may be presumed, without loss of generality, that Ω_ξ and Ω_η correspond to the dispersion matrices of moving-average processes. That is to say, they are sparse matrices with a limited number of central nonzero bands and with zeros elsewhere. Therefore, the solution to equation (13) may be found via an efficient Cholesky factorisation that sets $\Omega_\xi + \Omega_\eta = GG'$, where G is a lower-triangular matrix with a limited number of nonzero bands. The system $GG'b = y$ may be cast in the form of $Gp = y$ and solved for p. Then $G'b = p$ can be solved for b.

The solution via the Cholesky decomposition constitutes a recursive bidirectional filtering process that generates the vector b via two passes running in opposite directions through the data. The vector p is the product of a pass that runs forwards in time, and the vector b is generated from p in a reverse-time pass. Then b is subjected to further non-recursive filtering operations, described by (14), which produce x and h.

An alternative way of calculating the estimates is to use the Kalman filter in combination with a fixed-interval smoothing algorithm. The Kalman filter provides a sequence of the minimum-mean-square-error estimates, indexed by t, in the form of $E(\xi_t|\mathcal{I}_t)$ and $E(\eta_t|\mathcal{I}_t)$, which make use only of the information in $\mathcal{I}_t = [y_0, y_1, \ldots, y_t]$ that is available at time t. To make full use of all of the sample information, the estimates must be subjected, via the smoothing algorithm, to a retrospective enhancement using the information that has transpired after time t up to the end of the sample.

A detailed account of the Kalman filter and of a variety of fixed-interval smoothing algorithms has been provided by Pollock (2003). Also included is an account of method of Ansley and Kohn (1985), which appears to be the definitive way of handling the problem finding the initial conditions for the Kalman filter algorithm when the data are assumed to be the generated by a nonstationary linear process.

The Wiener–Kolmogorov principle of signal extraction is the foundation of the model-based methods of unobserved components analysis that are nowadays in widespread use. The parameters of the filters are determined in the process of fitting ARIMA models to the data components. However, the principle also supports a variety of heuristic filters of which the parameters are determined by rule of thumb or in view of the desired characteristics of their frequency response functions.

Amongst such heuristic filters is the digital version of the Butterworth filter, which has been advocated by Pollock (1997, 1999, 2000, 2001a, 2001b). The frequency response of the filter is maximally flat in the vicinity of the zero frequency and it has a transition band, centred on a chosen cut-off frequency, that can be narrowed by increasing the filter order. (Gómez 2001 has also advocated the Butterworth filter, but he has widened its definition to include filters, such as the filter of Hodrick and Prescott 1997, that do not share the property of maximal flatness.)

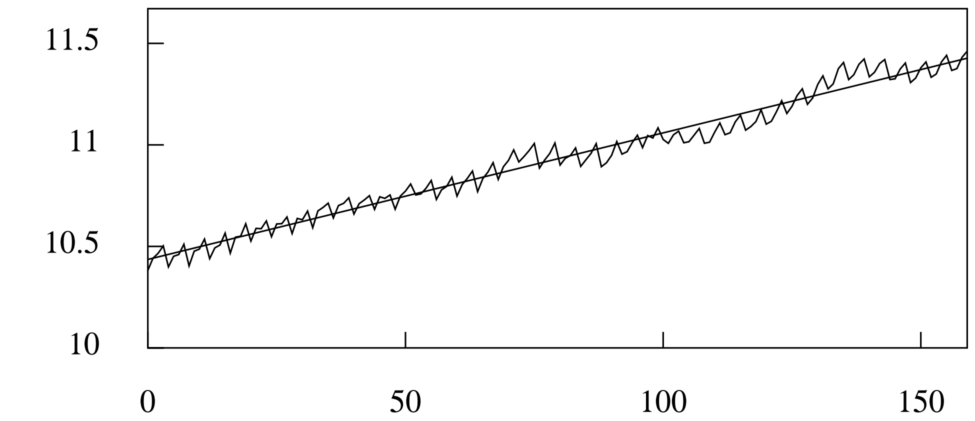

Figure 1. The quarterly series of the logarithms of consumption in the U.K., for the years 1955 to 1994, with a linear function interpolated by least-squares regression.

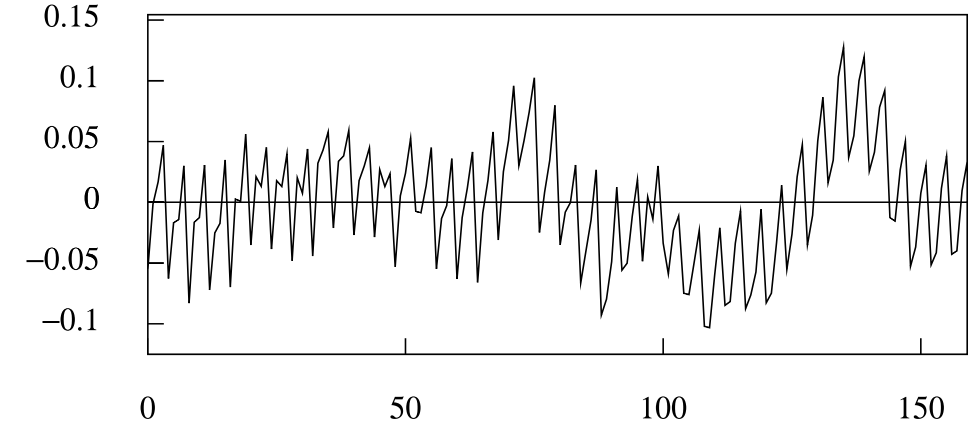

Figure 2. The residuals obtained by fitting a linear trend through the logarithmic consumption data of Figure 1.

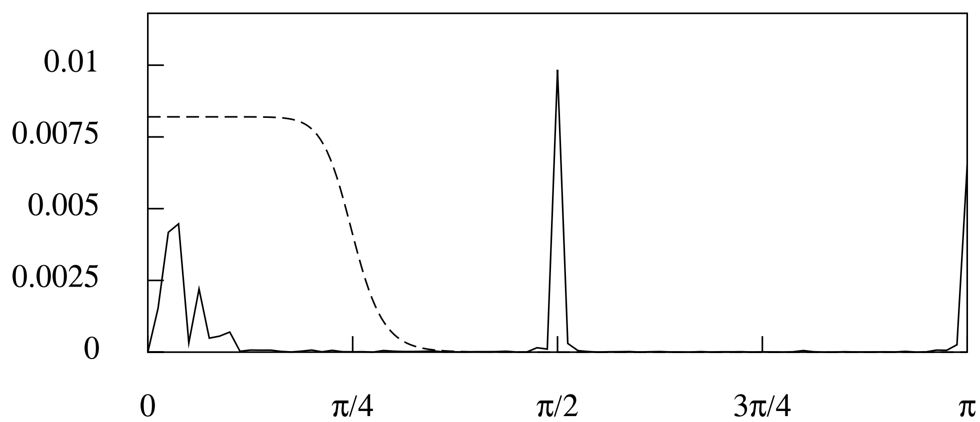

Figure 3. The periodogram of the residuals obtained by fitting a linear trend through the logarithmic consumption data of Figure 1. The gain of the lowpass Butterworth filter of order $n = 6$ and with a cut-off frequency of $\pi/4$ is represented by the dotted line. (The gain is unity at zero frequency.)

The Butterworth filter, which is of the lowpass Wiener–Kolmogorov variety, originates in a function that is the ratio of two quasi autocovariance generating functions:

$$\begin{aligned}\psi(z) &= \frac{\sigma_\xi^2\gamma_\xi(z)}{\sigma_\xi^2\gamma_\xi(z)+\sigma_\eta^2\gamma_\eta(z)} \\ &= \frac{(1+z)^n(1+z^{-1})^n}{(1+z)^n(1+z^{-1})^n+\lambda(1-z)^n(1-z^{-1})^n}.\end{aligned} \tag{15}$$

(Here, the normalised autocovariance functions $\gamma_\xi(z)$ and $\gamma_\eta(z)$, which are autocorrelation functions in other words, need to be scaled by the factors σ_ξ^2 and σ_η^2 respectively, which stand for the variances of the white-noise processes from which the signal and the noise components are supposedly derived by linear filtering.)

The autocovariance generating functions relate to an heuristic statistical model as opposed to a realistic one. The denominator function $\gamma(z)=\sigma_\xi^2\gamma_\xi(z)+\sigma_\eta^2\gamma_\eta(z)$ stands in place of that of the data process and the numerator function $\sigma_\xi^2\gamma_\xi(z)$ corresponds to that of the signal. Here, $\lambda=\sigma_\eta^2/\sigma_\xi^2=\{1/\tan(\omega_C/2)\}^{2n}$ incorporates the nominal cut-off frequency of ω_C. By setting $z=L_T$ and $z^{-1}=F_T$ in the numerator and the denominator of $\psi(z)$, we derive the matrices Ω_ξ and $\Omega_\xi+\Omega_\eta$, respectively, which can be entered into equations (8) and (9).

We should observe that $\psi(z)$ will continue to be expressed a rational function, regardless of whether the vectors ξ and η are generated by are moving-average processes, as we have supposed so far, of by more general autoregressive moving–average processes. Whichever is the case, the structures of the estimating equations under (8) and (9), which entail sparse Toeplitz matrices, will not be affected.

To derive the two unidirectional filters, the rational function is factorised as $\psi(z)=\beta(z)\beta(z^{-1})$, where $\beta(z)$, which relates to the direct-time filter, contains the poles that lie outside the unit circle, and $\beta(z^{-1})$, which relates to the reverse-time filter, contains the poles that lie inside the circle. This factorisation is described as the Cramér–Wold decomposition. In the case of the Butterworth filter, analytic expressions for the roots of both the denominator and the numerator, i.e. the poles and the zeros of the filter, are available. The roots of the denominator have been given by Pollock (2000).

For most other Wiener–Kolmogorov filters specified in the manner of the Butterworth filter, it is necessary to use an iterative procedure for finding the Cramér–Wold decomposition. (See for, example Pollock 2003b.) The algorithm of Wilson (1969), which is based on the Newton–Raphson procedure, is an effective way of achieving the factorisation; and versions which are coded in *C* and in *Pascal* have been provided by Pollock (1999). (See, also, Laurie 1980, 1982.)

There can be a reasonable objection to the assumption that the data components are generated by ordinary linear stochastic processes that comprise the full range of frequencies from zero up to the limiting Nyquist frequency of π radians per period. (In discretely sampled systems, the frequencies in excess of the Nyquist value will be aliased by frequencies within the interval $[0, \pi]$.) We shall illustrate the grounds for questioning the assumption via an analysis of a leading economic index.

Example 1. Figure 1 show the logarithms of the quarterly consumption data for the U.K. for the years 1955–1994, through which a linear trend has been interpolated by least-squares regression. When a quadratic polynomial trend was fitted, it was discovered that the coefficient associated with t^2 was not significantly different from zero.

This implies that, over the years in question, the underlying growth of the economy was at a constant exponential rate. Moreover, an exponential trend represents a benchmark that can be used in defining the business cycle. (Flexible methods of trend extraction that employ linear filters are described later, beginning in section 4.) The residual deviations from the trend, which are shown in Figure 2, represent a variable multiplicative factor by which the underlying trend is modulated; and the residuals reveal both secular and seasonal variations in consumption.

The periodogram of the residuals is shown in Figure 3. This has a low-frequency spectral structure, which extends no further than the frequency value of $\pi/8$. The remainder of the periodogram shows a dead space that is punctuated by tall spikes in the vicinities of the frequencies of $\pi/2$ and π. The first of these spikes corresponds to the fundamental frequency of the seasonal fluctuations that play on the back of the more gentle variations that surround the ascending line in Figure 1. The spike at π is corresponds to the first harmonic of the seasonal frequency.

The low-frequency structure of Figure 3, which occupies the frequency interval $[0, \pi/8]$, can be isolated successfully by any of a wide variety of filters. All that is required of such a filter is that its transition from pass band to stop band occurs within the spectral dead space that stretches from $\pi/8$ to the vicinity of $\pi/2$, where the spectral structure of the seasonal fluctuations is first encountered. The Butterworth filter of order $n = 6$ with a cut off frequency of $\pi/4$ fulfils this requirement. Its frequency response function is superimposed on Figure 3. The effect of applying this filter to the data of Figure 2 is shown in Figures 4 and 5.

A more exacting task is the extraction of the low-frequency components from data that is observed at monthly intervals. In that case, the fundamental seasonal frequency is at $\pi/6$ and the transition of the filter must occur within a correspondingly reduced interval.

The sharpening of the transition can be achieved by raising the order n of the filter. However, a sharp transition in the low frequency range can be achieved with a recursive Wiener–Kolmogorov filter only at the cost of bringing the poles of the filter into close proximity with the perimeter of the unit circle. This can lead to problems of filter instability, which include the propagation of numerical rounding errors and the prolongation of the transient effects of ill-chosen start-up conditions.

These problems have been addressed within the context of the Wiener–Kolmogorov specification by Pollock (2003b). Alternative specifications for recursive filters have been investigated in Pollock (2003a). In the next section, we shall also deal with the problems of "sharp filtering" within the context of the Wiener–Kolmogorov theory; but we shall forsake the method of recursive filtering in favour of a method based on Fourier analysis.

3. Filtering via Circulant Matrices

A finite-sample analogue of a stationary stochastic process is a circular or periodic process $y(t) = \{y_t; t = 0, \pm 1, \pm 2, \ldots\}$ that is completely specified by its values at T consecutive points such that $y_t = y_{t \bmod T}$. For such processes, the lag operator is replaced by the circulant matrix

$$K_T = [e_1, \ldots, e_{T-1}, e_0], \tag{16}$$

which is formed from the identity matrix I_T by moving the leading vector to the back of the array.

This operator effects the cyclic permutation of the elements of any (column) vector of order T. The matrix is T-periodic such that $K_T^{q+T} = K_T^q$. Whereas $L_T y = [0, y_0, \ldots, y_{T-2}]'$ is obtained from $y = [y_0, y_1, \ldots, y_{T-1}]'$ by deleting the final element and placing a zero in the leading position, the vector $K_T y = [y_{T-1}, y_0, \ldots, y_{T-2}]'$ is obtained from y by moving the final element to the leading position.

The powers of K_T form the basis for the set of circulant matrices. In particular, we may define a matrix of circular autocovariances via the formula

$$\begin{aligned} D^\circ(y) = \Omega^\circ &= \gamma(K) \\ &= \gamma_0 I + \sum_{\tau=1}^{\infty} \gamma_\tau (K_T^\tau + K_T^{-\tau}) \\ &= \gamma_0^\circ I + \sum_{\tau=1}^{T-1} \gamma_\tau^\circ (K_T^\tau + K_T^{-\tau}). \end{aligned} \tag{17}$$

Here, $\gamma_\tau^\circ; \tau = 0, \ldots, T-1$ are the circular autocovariances defined by

$$\gamma_\tau^\circ = \sum_{j=0}^{\infty} \gamma_{(jT+\tau)}. \tag{18}$$

The matrix operator K_T has a spectral factorisation that is particularly useful in analysing the properties of the discrete Fourier transform. The basis of this factorisation is the so-called Fourier matrix. This is a symmetric matrix

$$U = T^{-1/2}[W^{jt}; t, j = 0, \ldots, T-1], \tag{19}$$

of which the generic element in the jth row and tth column is

$$\begin{aligned} W^{jt} &= \exp(-\mathrm{i}2\pi tj/T) = \cos(\omega_j t) - \mathrm{i}\sin(\omega_j t), \\ &\text{where} \quad \omega_j = 2\pi j/T. \end{aligned} \tag{20}$$

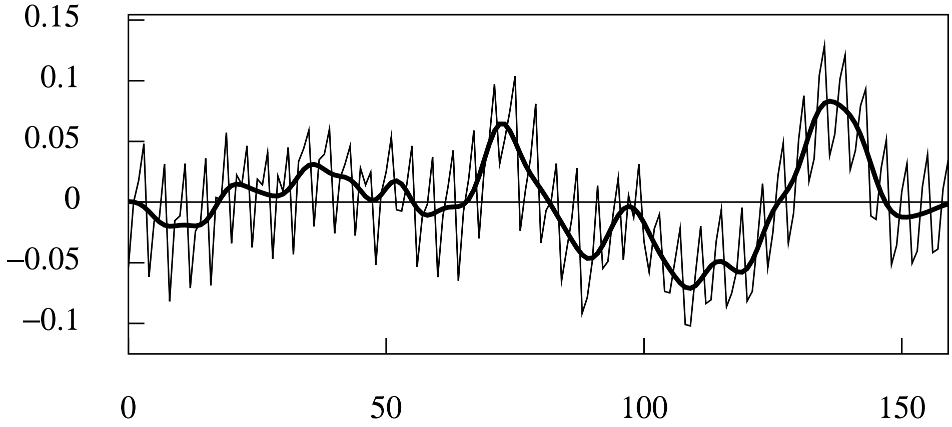

Figure 4. The low-frequency component of the consumption residuals of Figure 2. The component has been extracted by applying a lowpass Butterworth filter of order $n = 6$ with a cut off point at $\omega_c = \pi/4$.

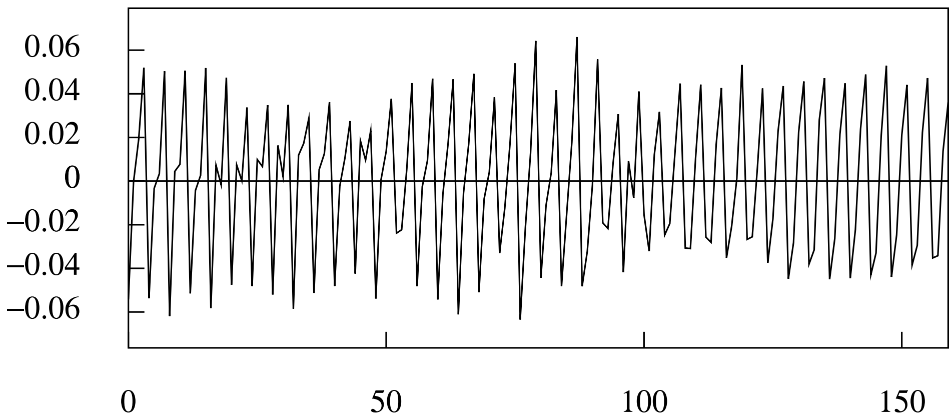

Figure 5. The component extracted from the consumption residuals by applying a highpass Butterworth filter of order 6 with a cut off point at $\omega_c = \pi/4$.

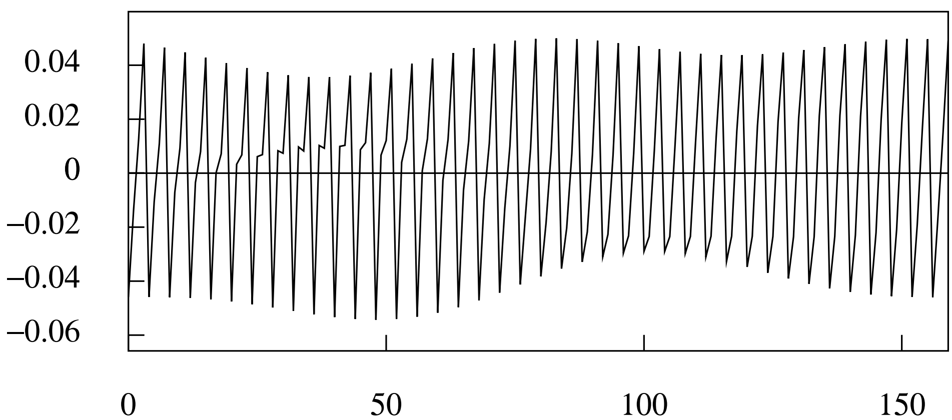

Figure 6. The seasonal component of the consumption residuals, synthesised from the Fourier ordinates in the vicinities of $\pi/2$ and π.

The matrix U is a unitary, which is to say that it fulfils the condition

$$\bar{U}U = U\bar{U} = I, \tag{21}$$

where $\bar{U} = T^{-1/2}[W^{-jt}; t, j = 0, \ldots, T-1]$ denotes the conjugate matrix.

The operator K_T can be factorised as

$$K = \bar{U}DU = U\bar{D}\bar{U}, \tag{22}$$

where

$$D = \mathrm{diag}\{1, W, W^2, \ldots, W^{T-1}\}, \tag{23}$$

with $W = \exp\{2\pi/T\}$, is a diagonal matrix whose elements are the T roots of unity, which are found on the circumference of the unit circle in the complex plane. Observe also that D is T-periodic, such that $D^{q+T} = D^q$, and that $K_T^q = \bar{U}D^qU = U\bar{D}^q\bar{U}$ for any integer q.

The spectral factorisation of the circulant autocovariance matrix gives

$$\Omega^\circ = \gamma(K_T) = \bar{U}\gamma(D)U. \tag{24}$$

Here, the jth element of the diagonal matrix $\gamma(D) = \Lambda$ is

$$\gamma(\exp\{\mathrm{i}\omega_j\}) = \gamma_0 + 2\sum_{\tau=1}^{\infty} \gamma_\tau \cos(\omega_j\tau). \tag{25}$$

This represents the cosine Fourier transform of the sequence of the ordinary autocovariances; and it corresponds to an ordinate (scaled by 2π) sampled at the point ω_j from the spectral density function of the linear (i.e. non-circular) stationary stochastic process. (An account of the algebra of circulant matrices has been provided by Pollock 2002. See, also, Gray 2002.)

The circulant autocovariance matrices that are the counterparts of the ordinary autocovariance matrices defined in (5) are

$$\begin{aligned} &\Omega_\xi^\circ = \bar{U}\Lambda_\xi U, \qquad \Omega_\eta^\circ = \bar{U}\Lambda_\eta U,\\ &\Omega^\circ = \bar{U}\Lambda U = \bar{U}(\Lambda_\xi + \Lambda_\eta)U, \end{aligned} \tag{26}$$

where Λ_η and Λ_ξ are diagonal matrices of spectral ordinates. Any of these autocovariance matrices may be singular in consequence of the presence of zero elements on the diagonals. Using the circulant matrices instead of the ordinary autocovariance matrices in the Wiener–Kolmogorov formulae of (8) and (9) gives

$$x = \bar{U}\Lambda_\xi\{\Lambda_\xi + \Lambda_\eta\}^+Uy = \bar{U}J_\xi Uy, \tag{27}$$

$$h = \bar{U}\Lambda_\eta\{\Lambda_\xi + \Lambda_\eta\}^+Uy = \bar{U}J_\eta Uy. \tag{28}$$

To accommodate the possibility that $\Lambda_\xi + \Lambda_\eta$ is singular, a generalised inverse has been applied to it instead of an ordinary inverse. A generalised inverse can obtained by replacing the zero-valued diagonal elements, which correspond to spectral ordinates falling within dead spaces, by nonzero values and, thereafter, by inverting the matrix that has acquired the full rank. Observe that, if Λ_ξ and Λ_η are disjoint such that $\Lambda_\xi \Lambda_\eta = 0$, then $J_\xi = \Lambda_\xi\{\Lambda_\xi + \Lambda_\eta\}^+$ is a matrix with units on the diagonal wherever Λ_ξ has nonzero elements and with zeros elsewhere. Analogous conditions apply to $J_\eta = \Lambda_\eta\{\Lambda_\xi + \Lambda_\eta\}^+$.

The formulae of (27) and (28) have a simple interpretation. First, the discrete Fourier transform is applied to the data vector y to translate it into the frequency domain. Then, a differential weighting, which might entail setting some values to zero, is applied to the spectral ordinates of the transformed vector via the diagonal matrices J_ξ or J_η. Finally, to produce the estimate of the component, the inverse Fourier transform is applied.

Implicit in the use of the discrete Fourier transform is the assumption that the data sequence represents a single cycle of a periodic function. In the periodic extension of the data, the values from the interval $[0, T)$ are reproduced in successive segments of length T that precede and follow the data.

In one sense, there is no start-up problem affecting a Fourier-based filtering procedure, since the periodic extension constitutes a doubly-infinite sequence. However, there may be radical disjunctions at the points where one replication of the data ends and another begins.

Such features are liable to be reflected in the periodogram in a way that can obscure the underlying data structures. Thus, the ordinates of the Fourier transform may be affected by a slew of values which serve the purpose only of synthesising the end-of-sample disjunctions. One recourse is to taper both ends of the sample so that they arrive the same level. Another recourse is to join the sample to it mirror-image reflection and to use this combination in place of the original data.

The problems of an end-of-sample disjunction are particularly acute in the case of nonstationary data sequences that follow rising or falling trends; and the trends have to be eliminated before the filters are applied. So far, we have succeeded in eliminating the trend by fitting a polynomial function to the data. An alternative recourse, which we shall pursue in the next section, is to make use of differencing.

Example 2. Consider the task of extracting the seasonal component from the residuals that have been obtained by fitting a linear function to the logarithmic consumption data. The periodogram of Figure 3 suggests that the seasonal sequence should be synthesised from a small number of Fourier ordinates that are in the vicinity of the seasonal frequency and its harmonic. Apart from those at either end of the frequency range, the ordinates come in conjugate complex pairs. In addition to the pair of ordinates at $\pi/2$, we may take two ordinates below and one above. Also, we may take ordinate at π, which is real-valued, and the pair of ordinates immediately below.

The seasonal sequence, which is plotted in of Figure 6, is equally a component of the sequence of Figure 4, which represents the residuals from the linear detrending of the logarithmic consumption data, and a component of the sequence of Figure 5, which has been derived by applying a highpass Butterworth filter to remove a further low-frequency component—represented by the thick line in Figure 4—that is unrelated to the seasons. In terms of the variances, the seasonal sequence represents 47 percent of the Figure 4 sequence and 94 percent of the Figure 5 sequence.

4. Filtering Nonstationary Sequences

The problems of a trended data sequence may be overcome by differencing. The matrix that takes the d-th difference of a vector of order T is given by

$$\nabla_T^d = (I - L_T)^d. \tag{29}$$

We may partition the matrix so that $\nabla_T^d = [Q_*, Q]'$, where Q_*' has d rows. The inverse matrix is partitioned conformably to give $\nabla_T^{-d} = [S_*, S]$. We may observe that

$$[\,S_* \quad S\,]\begin{bmatrix} Q_*' \\ Q' \end{bmatrix} = S_*Q_*' + SQ' = I_T, \tag{30}$$

and that

$$\begin{bmatrix} Q_*' \\ Q' \end{bmatrix}[\,S_* \quad S\,] = \begin{bmatrix} Q_*'S_* & Q_*'S \\ Q'S_* & Q'S \end{bmatrix} = \begin{bmatrix} I_d & 0 \\ 0 & I_{T-d} \end{bmatrix}. \tag{31}$$

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

$$\nabla_T^d y = \begin{bmatrix} Q_*' \\ Q' \end{bmatrix} y = \begin{bmatrix} g_* \\ g \end{bmatrix}. \tag{32}$$

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

$$y = S_*g_* + Sg. \tag{33}$$

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

We may approach the filtering of a trended data sequence in the following manner. First, we reduce the data to stationarity by differencing it an appropriate number of times. (We rarely need to difference the data more than twice.) From the differenced data, viewed in an appropriate manner, we may discern the nature and the frequency ranges of the various data structures that we wish to isolate.

Next, the components of the differenced data that correspond to these structures may be extracted, either by a recursive filtering process—using, for example, a Butterworth filter—or via the Fourier method described in the preceding section.

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.

An apparent problem with this procedure is that the act of differencing is liable to attenuate the components of the low-frequency data structure to such an extent that they become invisible in the periodogram of the differenced data. The problem is illustrated in Figure 8, which shows the periodogram of $g = Q'y$ in the case of the the once-differenced consumption data.

The problem vanishes when we recognise that we can discern the low-frequency structure via the periodogram of the residual sequence

$$\begin{aligned} y - p &= y - S_*(S_*'S_*)^{-1}S_*'y \\ &= Q(Q'Q)^{-1}Q'y, \end{aligned} \tag{34}$$

obtained by fitting to the data, by least-squares, a polynomial of degree $d-1$. The identity $Q(Q'Q)^{-1}Q' = I - S_*(S_*'S_*)^{-1}S_*'$ follows from the fact that Q and S_* are complementary matrices with $\text{Rank}[Q, S_*] = T$ and $Q'S_* = 0$. It will be recognised that the residuals contain the same information as does the differenced data $Q'y$. Their periodogram, in the case of the consumption data, has been displayed already in Figure 3.

To elucidate the procedures for extracting the components of a trended data sequence, let us consider the case of the data vector $y = \xi + \eta$, where η, which has $E(\eta) = 0$ and $D(\eta) = \Omega_\eta$, is from a stationary stochastic process and where ξ is from a process that requires a d-fold differencing in order to reduce it to a vector $\zeta = Q'\xi$ with a stationary distribution. Then we shall have

$$\begin{aligned} Q'y &= Q'\xi + Q'\eta, \\ &= \zeta + \kappa = g, \end{aligned} \tag{35}$$

and we may assume, by analogy with (5), that ζ and κ are characterised by their first and second moments, which are

$$\begin{aligned} &E(\zeta) = 0, \qquad D(\zeta) = \Omega_\zeta = Q'\Omega_\xi Q, \\ &E(\kappa) = 0, \qquad D(\kappa) = \Omega_\kappa = Q'\Omega_\eta Q, \\ &\quad \text{and} \quad C(\zeta, \kappa) = 0. \end{aligned} \tag{36}$$

Here, the derived dispersion matrices Ω_ζ and Ω_κ retain the Toeplitz structure that is a feature of Ω_ξ and Ω_η.

Let the estimates of ζ and κ be denoted by z and k. If x and h are the estimates of ξ and η respectively, then it is reasonable to require that $Q'x = z$ and $Q'h = k$ so that

$$\begin{aligned} Q'y &= Q'x + Q'h \\ &= z + k = g. \end{aligned} \tag{37}$$

The estimates z and k must be integrated to give

$$x = S_*z_* + Sz \qquad \text{and} \qquad h = S_*k_* + Sk. \tag{38}$$

The criterion for finding the starting value z_* is

$$\text{Minimise} \quad (y - x)'\Omega_\eta^{-1}(y - x) = (y - S_*z_* - Sz)'\Omega_\eta^{-1}(y - S_*z_* - Sz). \tag{39}$$

This requires that the estimated trend x should adhere as closely as possible to the data. The minimising value is

$$z_* = (S_*'\Omega_\eta^{-1}S_*)^{-1}S_*'\Omega_\eta^{-1}(y - Sz). \tag{40}$$

Since $y - x = h$, an equivalent criterion is

$$\text{Minimise} \quad h'\Omega_\eta^{-1}h = (S_*k_* + Sz)'\Omega_\eta^{-1}(S_*k_* + Sk), \tag{41}$$

for which the minimising value is

$$k_* = -(S_*'\Omega_\eta^{-1}S_*)^{-1}S_*'\Omega_\eta^{-1}Sk. \tag{42}$$

Using

$$P_* = S_*(S_*'\Omega_\eta^{-1}S_*)^{-1}S_*'\Omega_\eta^{-1}, \tag{43}$$

we get, from (38), the following values:

$$x = P_*y + (I - P_*)Sz, \qquad \text{and} \qquad h = (I - P_*)Sk. \tag{44}$$

The disadvantage in using these formulae directly is that the inverse matrix Ω_η^{-1}, which is of order T, is liable to have nonzero elements in every location. (This will be so whenever Ω_η has the form of an autocovariance matrix of a moving-average process.)

The appropriate recourse is to use the identity

$$\begin{aligned} I - P_* &= I - S_*(S_*'\Omega_\eta^{-1}S_*)^{-1}S_*'\Omega_\eta^{-1} \\ &= \Omega_\eta Q(Q'\Omega_\eta Q)^{-1}Q' \end{aligned} \tag{45}$$

to provide an alternative expression for the projection matrix $I - P_*$ that incorporates the band-limited matrix Ω_η instead of its inverse. The equality follows from the fact that, if $\text{Rank}[R, S_*] = T$ and if $S_*'\Omega_\eta^{-1}R = 0$, then

$$I - S_*(S_*'\Omega_\eta^{-1}S_*)^{-1}S_*'\Omega_\eta^{-1} = R(R'\Omega_\eta^{-1}R)^{-1}R'\Omega_\eta^{-1}. \tag{46}$$

Setting $R = \Omega_\eta Q$ gives the result. Given that $x = y - h$, it follows that we can write

$$\begin{aligned} x &= y - (I - P_*)Sk \\ &= y - \Omega_\eta Q(Q'\Omega_\eta Q)^{-1}k, \end{aligned} \tag{47}$$

where the second equality depends upon $Q'S = I$.

So far, we have not specified the method by which the estimates z and k of the differenced components have been obtained. They may be obtained equally via a recursive filtering method or via the Fourier that has been outlined in the preceding section. In case we have used the Fourier method, we might be inclined to use the circulant version of the dispersion matrix Ω_η within the foregoing formulae.

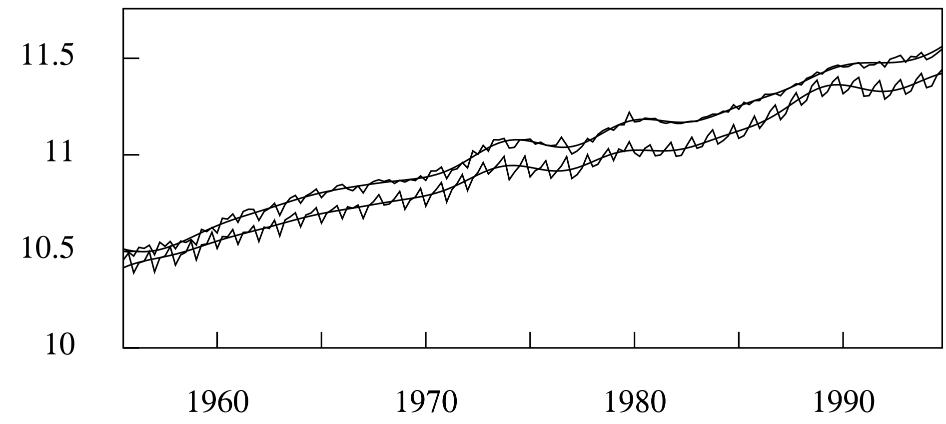

Figure 7. The quarterly series of the logarithms of income (upper) and consumption (lower) in the U.K. for the years 1955 to 1994 together with their interpolated trends.

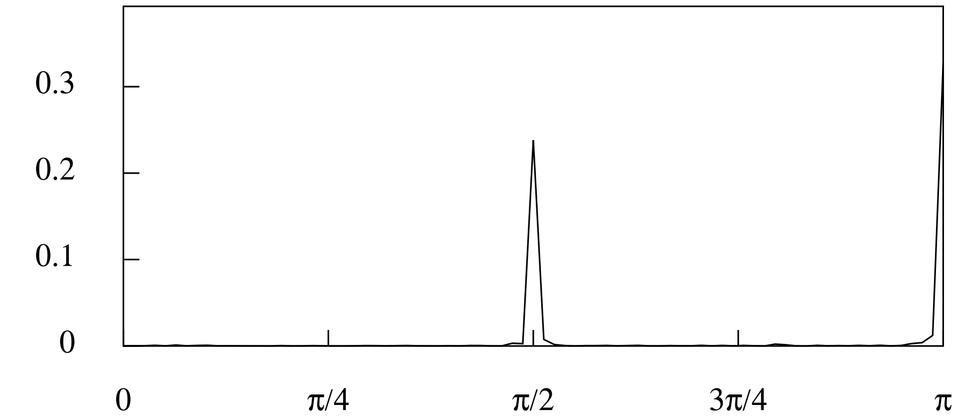

Figure 8. The periodogram of the first differences of the logarithmic consumption data.

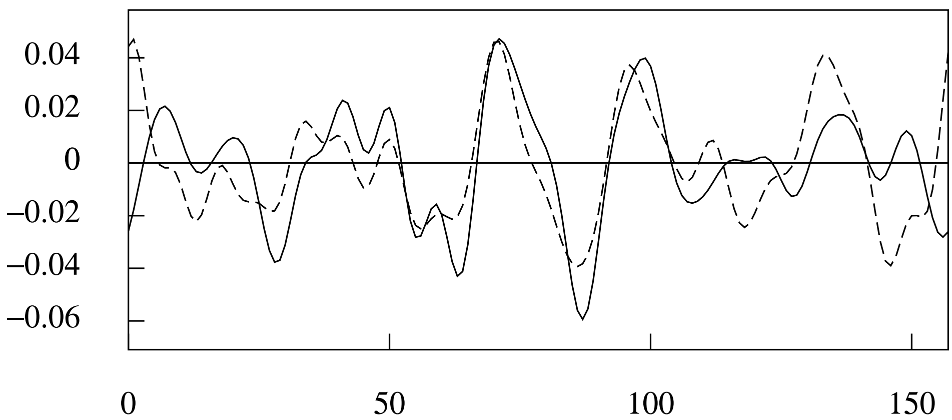

Figure 9. The bandpass estimates of the fluctuations, within the range of the business-cycle frequencies, of the logarithmic income series (solid line) and of the logarithmic consumption series (broken line).

Let us consider, instead, the possibility of obtaining the estimate k via recursive filtering. Then, with reference to equation (9), we can see that the assumptions of (36) imply that the estimate should take the form of

$$k = Q'\Omega_\eta Q(\Omega_\zeta + Q'\Omega_\eta Q)^{-1}Q'y, \tag{48}$$

On substituting this in the equation of (47), we get

$$x = y - \Omega_\eta Q(\Omega_\zeta + Q'\Omega_\eta Q)^{-1}Q'y. \tag{49}$$

In the case of Butterworth filter, we take the quasi-autocorrelation functions of the nonstationary signal sequence $\xi(t)$ and of the stationary noise sequence $\eta(t)$ to be

$$\gamma_\xi(z) = \sigma_\xi^2 \frac{(1+z)^n(1+z^{-1})^n}{(1-z)^d(1-z^{-1})^d} \quad \text{and} \quad \gamma_\eta(z) = \sigma_\eta^2(1-z)^{n-d}(1-z^{-1})^{n-d} \tag{50}$$

respectively, which become the elements of (15) in the case where $d = 0$. We may also define

$$\begin{aligned}\gamma_\zeta(z) = \sigma_\xi^2(1+z)^n(1+z^{-1})^n \qquad \text{and} \qquad \gamma_\kappa(z) &= (1-z)^d\gamma_\eta(z)(1-z^{-1})^d\\ &= \sigma_\eta^2(1-z)^n(1-z^{-1})^n,\end{aligned} \tag{51}$$

which is a matter of renaming the elements of (15) when $d > 0$. The matrices Ω_ζ and $\Omega_\kappa = Q'\Omega_\eta Q$ are generated by setting $z = L_{T-d}$ and $z^{-1} = L'_{T-d} = F_{T-d}$ in $\gamma_\zeta(z)$ and $\gamma_\kappa(z)$ respectively and by scaling the resulting matrices by the appropriate variances. Observe that the generating functions of (51) are not affected by the order d of the differencing operator. Therefore, for the Butterworth filter, only the dimension of the matrix $\Omega_\zeta + Q'\Omega_\eta Q$ changes when d varies. Its essential structure remains the same.

The computational procedure that has been described in section 2 can also be applied when $d > 0$. That is to say, the solution of the equation $(\Omega_\zeta + Q'\Omega_\eta Q)b = g$, where $g = Q'y$, is found via the Cholesky factorisation of $\Omega_\zeta + Q'\Omega_\eta Q = GG'$. Thereafter, $h = \Omega_\eta b$ and $x = y - h$ are found.

Example 3. Figure 7 shows the quarterly sequences of the logarithms of income (upper) and consumption (lower) in the U.K. for the years 1955 to 1994 together with their interpolated trends. We can afford to treat the income sequence in the same manner as we treat the consumption sequence; and, in what follows, we shall concentrate on the latter.

The periodogram of Figure 3, suggests that both the trend component and the seasonal component of the consumption data are generated by band-limited processes. The trend component is confined to the frequency interval $[0, \pi/8]$ and the seasonal component comprises a handful of nonzero Fourier ordinates in the vicinities of $\pi/2$ and π. The remainder of the periodogram consists of virtual dead spaces. When equation (4) is applied to these circumstances, ξ, which is estimated by x, becomes the trend component and η, which is estimated by h, becomes the seasonal component.

The trend that interpolates the consumption data has been constructed by extracting from the untrended, twice-differenced data sequence $g = Q'y$ the Fourier elements that lie in the frequency interval $[0, \pi/8]$. The sequence z that is synthesised from these elements has then been integrated to create the trend $x = S_* z_* + Sz$. In seeking the starting value z_* with which to initiate the process of integration, we may consider minimising a criterion function in the form of $h'(\Omega_\eta^\circ)^+ h = h'\bar{U}\Lambda_\eta^+ Uh$, where $(\Omega_\eta^\circ)^+ = \bar{U}\Lambda_\eta^+ U$ represents the generalised inverse of the singular circulant autocovariance matrix $D^\circ(\eta) = \Omega_\eta^\circ$.

The elements of Λ_η^+ that correspond to zero-valued elements of Uh, which lie in spectral dead spaces, can take arbitrary values. These values will have no effect upon the value of the criterion function. Therefore, the generalised inverse can be formed by replacing the nonzero elements of Λ_η by their inverses and by placing arbitrary values elsewhere on the diagonal.

For want of a better assumption, we may assume that the Fourier ordinates of the seasonal process are distributed uniformly within their designated bands. In that case, the corresponding elements of Λ_η should all have same value, and so, likewise, should the corresponding elements of Λ_η^+.

The remaining elements of Λ_η^+, which correspond to zero-valued Fourier ordinates and which can take arbitrary values, may be set to the same values as the elements corresponding to the seasonal ordinates. Thus Λ_η^+, which needs to be determined only up to a scalar factor, becomes an arbitrary multiple of the identity matrix—and it may as well become the identity matrix itself. In that case, we should have $\Omega_\eta^\circ = \bar{U}U = I$ and $(\Omega_\eta^\circ)^+ = I$.

This simplification allows us to specialise equation (40) to give

$$z_* = (S_*' S_*)^{-1} S_*'(y - Sz). \tag{52}$$

In the case where the data is differenced twice, there is

$$S_*' = \begin{bmatrix} 1 & 2 & \ldots & T-1 & T \\ 0 & 1 & \ldots & T-2 & T-1 \end{bmatrix} \tag{53}$$

The elements of the matrix $S_*' S_*$ can be found via the formulae

$$\begin{aligned} \sum_{t=1}^{T} t^2 &= \frac{1}{6}T(T+1)(2T+1) \quad \text{and} \\ \sum_{t=1}^{T} t(t-1) &= \frac{1}{6}T(T+1)(2T+1) - \frac{1}{2}T(T+1). \end{aligned} \tag{54}$$

(A compendium of such results has been provided by Jolly 1961, and proofs of the present results were given by Hall and Knight 1899.) The matrix is somewhat ill-conditioned. Moreover, when the order of differencing exceeds two or three, it is necessary, in calculating the polynomial ordinates of $p = S_* z_*$, to use an orthogonal basis in place of the monomial basis that is provided by the columns of S_*. However, this case is rare.

5. Bandpass Filtering

Econometricians often characterise the business cycle in terms of a sinusoid that fluctuates around a slow-moving trend. According to the definitions of Burns and Mitchell (1946), the effects of the business cycle within an economic index correspond to the sinusoidal elements therein that have periods of no less than one-and-a-half years and of no more than eight years. A duration of one-and-a-half years seems too short, and we prefer to set the shortest duration at 2 years—and this seems to be a common preference (see, for example, Christiano and Fitzgerald 1998).

The business cycle, defined in this manner, is unlikely to correspond to any self-contained spectral structure that might be discerned by inspecting the relevant periodogram. In the case of quarterly data, the business cycle frequencies range from $\pi/16$ radians per period to $\pi/4$ radians per period (corresponding to a duration of 2 years.) Neither of these values corresponds to a natural break in the periodogram of the consumption residuals of Figure 3.

The business-cycle frequencies may be extracted from the data using a bandpass filter with nominal cut-off points at the designated frequencies. For this purpose, economists have tended to use finite-impulse-response (FIR) or moving-average filters that are derived by truncating the doubly-infinite sequence of filter coefficients associated with the unrealisable ideal bandpass filter. (See, for example, Baxter and King, 1999.) The effect of the truncation is to create ripples in the stopbands of the frequency response function, which entail considerable spectral leakage.

A superior bandpass filter can be realised using the Butterworth formulation. One way of creating a bandpass filter is to apply the so-called Constantinides (1970) transformation to a prototype lowpass filter with a nominal cut-off point at $\pi/2$. The method is also described by Pollock (1999). In the current application of business cycle analysis, this transformation will result in a filter with a frequency response that has a far wider transition band at the upper cut-off frequency than at the lower cut-off frequency.

A better way of creating a bandpass filter for the current application is to apply two filters in succession. The first filter is a lowpass filter that is intended to remove the components of frequencies in excess of $\pi/4$. The second is a highpass filter that preserves the remaining components of frequencies in excess of $\pi/16$ and eliminates those of lesser frequencies. The order of the first filter should exceed that of the second filter so as to enhance the rate of transition at the upper cut-off frequency. (Harvey and Trimbur 2003 have shown how this sort of bandpass filter can be derived from an heuristic statistical model.)

Figure 10 shows the pole–zero diagrams of the 12th order lowpass and the 6th order highpass filters; while Figure 11 shows the frequency response functions of the two filters superimposed on the same diagram. It can be seen that some of the poles of the highpass filter come very close to the circumference of the unit circle. This feature can lead to problems of numerical instability.

One way of overcoming the problems of numerical instability is to subsample the data that has resulted from applying the first filter. Since there is no information in this data remaining in the interval $[\pi/2, \pi]$, we can afford

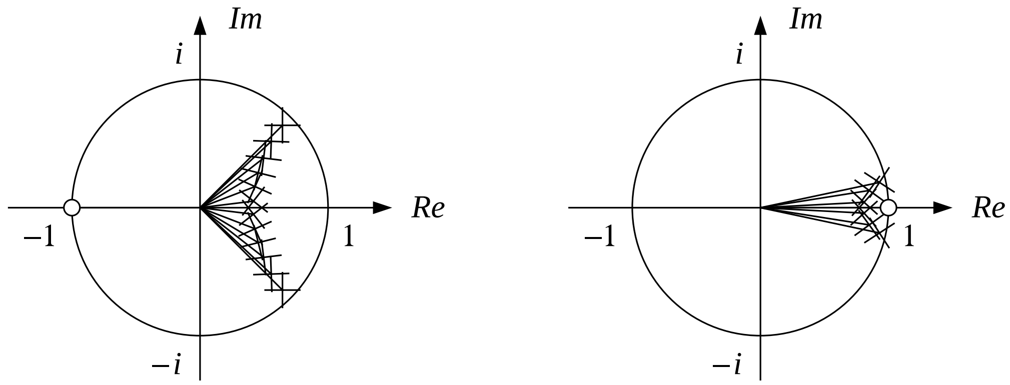

Figure 10. The pole–zero diagrams (left) of the lowpass Butterworth filter of order $n = 12$ with a cut-off frequency of $\omega_U = \pi/4$ and (right) of the highpass filter of order $n = 6$ with a cut-off frequency of $\omega_L = \pi/16$.

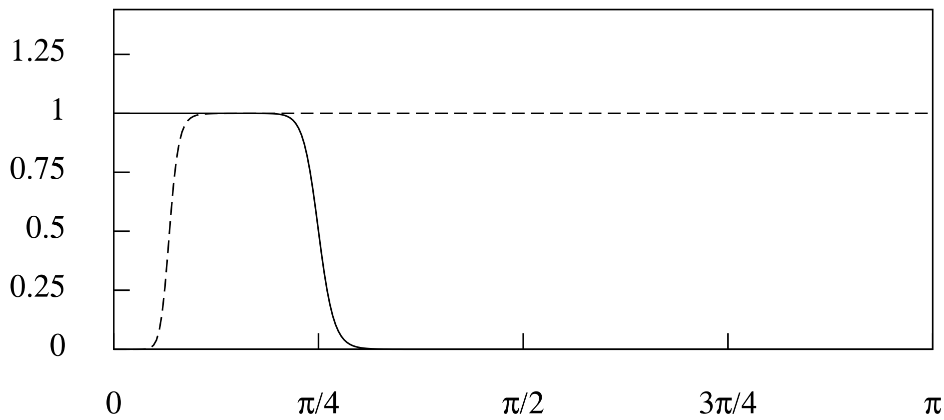

Figure 11. The gain of the 12th order digital Butterworth lowpass filter with cut-off frequency of $\omega_U = \pi/4$ and of 6th order highpass filter with cut-off frequency of $\omega_L = \pi/16$, superimposed on the same diagram.

to omit alternate points so as to create a semi-annual sequence. The effect is that the contents of the original data that lie in the frequency interval $[0, \pi/2]$ are mapped into the wider interval $[0, \pi]$. In the process, the lower cut-off frequency moves from $\pi/16$ to $\pi/8$. The poles of the 6th-order Butterworth filter with this cut-off point are no longer so close to the perimeter of the unit circle, which implies a greater numerical stability. (More general methods of sample-rate conversion have been described by Vaidyanathan 1993, amongst others.)

An alternative recourse is to base the estimate of the business cycle component on the Fourier ordinates of the data that fall within the specified frequency range. In principal, the method entails no spectral leakage so long as it is applied to data that have been detrended in a manner that will ensure that there are no disjunctions in the periodic extension where the end of one data segment joins the beginning of another. This can be achieved by a process of differencing followed by a judicious tapering of the ends of the data segment.

Since the business cycle is an artificial construct, it is difficult to relate the method of extraction to an underlying statistical model. However, under certain assumptions, it becomes appropriate to treat this component in the same manner as the noise component η within the trended vector $y = \xi + \eta$, which has been the subject of the previous section.

Now the component vector ξ becomes the repository of the Fourier elements with frequencies that are less than the value of the lower cut-off frequency of the pass band. The components of frequencies in excess of the upper cut-off frequency can be assigned to a third component, which is eliminated via the first lowpass filtering operation. The vector y can be taken to represent the product of this operation.

Under these constructions, there are no spectral overlaps amongst the various components; and the appropriate statistical model is one that comprises separable band-limited processes. It follows that the appropriate method for extracting the business-cycle component is, indeed, the Fourier-based method of Section 3. This is well-adapted to dealing with band-limited processes. The relevant Fourier components of the business cycle, contained in the vector k, must be extracted from a data vector $g = z + k$ that has been detrended by differencing. The estimate

$$h = S_* k_* + Sk \tag{55}$$

of the business cycle component is obtained by a process of summation that reverses the differencing.

In the absence of prior knowledge of the distribution of the spectral ordinates, we may set $\Omega_\eta = I$. In that case, the starting values are provided by the simplified formula

$$k_* = (S_*' S_*)^{-1} S_*' Sk. \tag{56}$$

which is derived from equation (42) by setting $\Omega_\eta^{-1} = I$. The simplification extends to the identity of (45), which becomes

$$\begin{aligned} P_* &= S_*(S_*' S_*)^{-1} S_*' \\ &= I - Q(Q'Q)^{-1}Q' = I - P_Q. \end{aligned} \tag{57}$$

Therefore, the estimate of the business cycle component is also provided by

$$h = (I - P_*)Sk = Q(Q'Q)^{-1}k, \tag{58}$$

wherein the condition $Q'S = I$ has been effective in simplifying the final expression.

Example 4. Figure 9 shows the business cycle fluctuations that have been extracted from the quarterly logarithmic income and consumption data for the U.K. over the period 1955 to 1994. In both cases, a Fourier bandpass filter, of the sort describe in section 3, has been applied to data that have been reduced to stationarity by second differencing.

A lower cut-off point at $\pi/16$ radians per period (corresponding to a cycle of 8 years duration) and an upper cut-off point of $\pi/4$ radians per period (corresponding to a cycle of 2 years duration) have be chosen. The stationary sequences, denoted by z in equation (37), which have been synthesised from the selected Fourier ordinates, have been re-inflated in accordance with the formula for the estimated trend x given by (38), wherein the constants of integration of z_* are given by (52).

There is evidence here that the fluctuations in consumption precede those in income. This contradicts the common supposition that the business cycle is driven by variations in "autonomous expenditures", which do not include consumption, and in the rate of investment.

One might be doubtful of the comparisons at the beginning and the end of the sample, where the interpolated functions are not tied down by preceding or succeeding data points and where they appear to be heading in opposite directions. The problem could be overcome by adding a few extrapolated points at either end of the sample that would serve to tie down the functions.

6. Multiple Components

The problems of econometric signal extraction have been handled, so far, within the context of a model, described by equation (4), that has only a signal component and a noise component. Allowance has been made for a non stationary signal component. However, it might be required to partition the data amongst more than two components. Thus, in a classical econometric time-series analysis, at least four components are identified. These are the trend, the business cycle, the seasonal cycle and an irregular component.

The two-component model can also serve the purpose of extracting several components, for the reason that its components are readily amenable, if necessary, to further decompositions. Thus, for example, an initial decomposition of the data sequence into a trend/cycle component and a residue can be followed by decomposition of the residue into a seasonal cycle an irregular cycle. If the data are stationary, it is unnecessary to perform such a multiple decomposition sequentially—each component can be extracted separately.

If the data are nonstationary and if there are more than one nonstationary component, then a sequential decomposition might be called for. A typical model of an econometric time series, described by the equation $y = \xi + \eta = (\mu + \rho) + \eta$, comprises both a trend/cycle component μ and a seasonal component ρ that are described by ARIMA models with real and complex unit roots respectively.

To reduce the data to stationarity, an operator is used that is the product of the d-fold difference operator $\nabla_T^d = (I - L_T)^d$ and a deseasonalising operator $\Sigma_T = (I - L_T^s)(I - L_T)^{-1}$. (The operator Σ is used instead of $(I - L_T^s)$ because it can be assumed, without loss of generality, that the seasonal deviations from the trend have zero mean.) Let the product of the two operators be denoted by $M_T = \Sigma_T \nabla_T^d = [Q_*, Q]'$, where Q_*' contains the first $d + s - 1$ rows of the matrix, and let the inverse operator $M_T^{-1} = [S_*, S]$ be partitioned conformably such that S_* contains the first $d + s - 1$ columns. The factors of M_T^{-1} are further partitioned as $\Sigma_T^{-1} = [S_{\Sigma *}, S_\Sigma]$ and $\nabla_T^{-d} = [S_{\nabla *}, S_\nabla]$.

Let the components of the differenced data be denoted by $Q'\xi = \zeta$, $Q'\mu = \zeta_\mu$ and $Q'\rho = \zeta_\rho$. Then there is

$$\begin{aligned} Q'y &= Q'\xi + Q'\eta \\ &= Q'(\mu + \rho) + \kappa = (\zeta_\mu + \zeta_\rho) + \kappa. \end{aligned} \tag{59}$$

Also, let the estimates of μ and ρ be denoted by m and r and those of ζ_μ and ζ_ρ by z_m and z_r. Then, in parallel with equation (59), there is

$$\begin{aligned} Q'y &= Q'x + Q'h \\ &= Q'(m + r) + k = (z_m + z_r) + k. \end{aligned} \tag{60}$$

The estimates z_m, z_r and k may be obtained from the differenced data $g = Q'y$ by a process of linear filtering. It is then required to form m, r and h from these elements. First, consider

$$\begin{aligned} x = (m + r) &= S_* z_* + Sz \\ &= S_* z_* + S(z_m + z_r). \end{aligned} \tag{61}$$

Here, z_* is computed according the formula of (40). Given x, an estimate $h = y - x$ of the irregular component can be formed. Next, there is an equation

$$S_* z_* = [\, S_{\nabla *} \quad S_{\Sigma *} \,] \begin{bmatrix} z_{*m} \\ z_{*r} \end{bmatrix}. \tag{62}$$

This may be solved uniquely for z_{*m} and z_{*r}; and, for this purpose, only the first $s + d - 1$ rows of the system are required. Thereafter, the estimates of μ and ρ are given by

$$m = S_{\nabla *} z_{*m} + S z_m \quad \text{and} \quad r = S_{\Sigma *} z_{*r} + S z_r. \tag{63}$$

References

Ansley, C.F., and R. Kohn, (1985), Estimation, Filtering and Smoothing in State Space Models with Incompletely Specified Initial Conditions, *The Annals of Statistics,* 13, 1286–1316.

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

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, (1998), The Business Cycle, It's Still a Puzzle, *Federal Reserve Bank of Chicago, Economic Perspectives,* 22, 56–83.

Constantinides, A.G., (1970), Spectral Transformations for Digital Filters, *Proceedings of the IEE,* 117, 1585–1590.

Dagum, E.B., (1980), The X-11-ARIMA Seasonal Adjustment Method. Catalogue No. 12-564E, Statistics Canada, Ottawa.

Findley, D.F., B.C. Monsell, W.R. Bell, M.C. Otto and B.-C. Chen, (1998), New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program, *Journal of Business and Economic Statistics,* 16, 127–177.

Gómez, V., (2001), The Use of Butterworth Filters for Trend and Cycle Estimation in Economic Time Series, *Journal of Business and Economic Statistics,* 19, 365–373.

Gómez, V., and A. Maravall, (1994), Program TRAMO: Time Series Regression with ARIMA Noise, Missing Observations, and Outliers—Instructions for the User, EUI Working Paper Eco No. 94/31, Department of Economics, European University Institute.

Gómez, V., and A. Maravall, (1996), Programs TRAMO and SEATS. Instructions for the User, (with some updates). Working Paper 9628, Servicio de Estudios, Banco de España.

Gray, R.M., (2002), *Toeplitz and Circulant Matrices: A Review,* Information Systems Laboratory, Department of Electrical Engineering, Stanford University, California, `http://ee.stanford.edu/gray/~toeplitz.pdf`.

Hall, H.S., and S.R. Knight, (1899), *Higher Algebra,* Macmillan and Co., London.

Harrison, P.J., and C.F. Stevens, (1976), Bayesian Forecasting (With a Discussion), *Journal of the Royal Statistical Society, Series B,* 38, 205–247.

Harvey, A.C., (1989), *Forecasting, Structural Time Series Models and the Kalman Filter,* Cambridge University Press, Cambridge.

Harvey, A.C., andT.M. Timbur, (2003), General Model Based Filters for Extracting Cycles and Trends in Economic Time Series, *The Review of Economics and Statistics*, 85, 244–255.

Hillmer, S.C. and G.C. Tiao, (1982), An ARIMA-Model Based Approach to Seasonal Adjustment, *Journal of the American Statistical Association,* 77, 63–70.

Hodrick, J.R., and E.C. Prescott, (1997), Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit and Banking,* 29, 1–16.

Jolly, L.B.W., (1961), *Summation of Series: Second Revised Edition,* Dover Publications: New York.

Koopman, S.J., A.C. Harvey, J.A. Doornik and N. Shephard, (2000), STAMP: Structural Time Series Analyser, Modeller and Predictor, Timberlake Consultants Press, London.

Laurie, D.P., (1980), Efficient Implementation of Wilson's Algorithm for Factorising a Self-Reciprocal Polynomial, *BIT,* 20, 257–259.

Laurie, D.P., (1982), Cramér-Wold Factorisation, Algorithm AS 175, *Applied Statistics,* 31, 86–90.

Maravall, A., and D.A. Pierce (1987), A Prototypical Seasonal Adjustment Model, *Journal of Time Series Analysis,* 8, 177–193. Reprinted in Hylleberg, S. (ed.), *Modelling Seasonality,* Oxford University Press, 1992.

Pollock, D.S.G., (1997), Data Transformations and De-trending in Econometrics, Chapter 11 (pps. 327–362) in Christian Heij et al. (eds.), *System Dynamics in Economic and Financial Models,* John Wiley and Sons.

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

Pollock, D.S.G., (2000), Trend Estimation and Detrending via Rational Square Wave Filters, *Journal of Econometrics,* 99, 317–334.

Pollock, D.S.G., (2001a), Methodology for Trend Estimation, *Economic Modelling,* 18, 75–96.

Pollock, D.S.G., (2001b), Filters for Short Nonstationary Sequences, *Journal of Forecasting,* 20, 341–355.

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

Pollock, D.S.G., (2003a), Sharp Filters for Short Sequences, *Journal of Statistical Inference and Planning,* 113, 663–683.

Pollock, D.S.G., (2003b), Improved Frequency-Selective Filters, *Journal of Computational Statistics and Data Analysis,* 42, 279–297.

Pollock, D.S.G., (2003c), Recursive Estimation in Econometrics, *Journal of Computational Statistics and Data Analysis,* 44, 37–75.

Vaidyanathan, P.P., (1993), *Multirate Systems and Filter Banks,* Prentice-Hall, Englewood Cliffs, New Jersey.

Wilson, G.T., (1969), Factorisation of the Covariance Generating Function of a Pure Moving Average Process, *SIAM Journal of Numerical Analysis,* 6, 1–7.

Young, P. C., Pedregal, D., and Tych, W. (1999), Dynamic harmonic regression. *Journal of Forecasting,* 18, 369–394.

Young, P.C., C.J. Taylor, W. Tych, D.J. Pedregal, and C.J. McKenna, (2004), *The Captain Toolbox,* Centre for Research on Environmental Systems and Statistics, Lancaster University. (www.es.lancs.ac.uk/cres/captain).

Queen Mary
University of London

This working paper has been produced by
the Department of Economics at
Queen Mary, University of London

Copyright © 2005 Stephen Pollock
All rights reserved

Department of Economics
Queen Mary, University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 5096
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/papers/wp.htm