University of Leicester

DEPARTMENT OF ECONOMICS

Alternative Methods of Seasonal Adjustment

Stephen Pollock, University of Leicester, UK

Emi Mise, University of Leicester, UK

Working Paper No. 11/12

January 2011

ALTERNATIVE METHODS OF SEASONAL ADJUSTMENT

D.S.G. Pollock and Emi Mise
University of Leicester
Email: stephen_pollock@sigmapi.u-net.com

Abstract

Alternative methods for the seasonal adjustment of economic data are described that operate in the time domain and in the frequency domain. The time-domain method, which employs a classical comb filter, mimics the effects of the model-based procedures of the SEATS–TRAMO and STAMP programs. The frequency-domain method eliminates the sinusoidal elements of which, in the judgement of the user, the seasonal component is composed.

It is proposed that, in some circumstances, seasonal adjustment is best achieved by eliminating all elements in excess of the frequency that marks the upper limit of the trend-cycle component of the data. It is argued that the choice of the method seasonal adjustment is liable to affect the determination of the turning points of the business cycle.

Keywords: Wiener–Kolmogorov Filtering, Frequency-Domain Methods, The Trend-Cycle Component

1 Introduction

For a great many years, and until recently, there has been a broad consensus amongst central statistical agencies in the matter of how to perform the seasonal adjustment of economic data. The prevalent techniques have been those that were developed by the U.S. Bureau of the Census and which have been encapsulated in the X-11 computer program. The program was the culmination of the pioneering work undertaken by Julius Shiskin in the 1950's and the 1960's. (See Shiskin *et al.* 1967.)

The X-11 program has undergone numerous improvements and modifications, leading to the X-11-ARIMA software packages of 1975 and 1988. (See Dagum 1980, 1988.) Its latest incarnation is in the X-12-ARIMA package. (See Findlay *et al.* 1998.) Much of the relevant information on the method has been provided in a monograph of Ladiray and Quenneville (2001).

Recently, some alternative methods of seasonal adjustment have been making headway amongst central statistical agencies. Foremost amongst these is the ARIMA-model-based method of the TRAMO–SEATS package, which has been adopted by Eurostat, the central statistical agency of the European Union (See Caporello and Maravall 2004.) The STAMP program is also capable of seasonal adjustment. (See Koopman *et al.* 2000.) There are also indications that the Census Bureau itself is becoming more eclectic in its approach to seasonal adjustment. (See Monsell, Aston and Koopman 2003.)

In view of such developments, it seems that the time is ripe for a re-examination of the concepts and methods of seasonal adjustment. In this paper, we examine three methods that can be used for removing from an economic data sequence the elements associated with the annual seasonal cycle and its harmonics.

The first method, which operates in the time domain, is derived from a simple ARIMA model of a seasonal fluctuation superimposed upon a white-noise process. The Wiener–Kolmogorov methodology is used in deriving a filter that is appropriate for removing the seasonal component.

When there is a trend in the data, this can be removed by differencing before the filter is applied. Thereafter, the filtered data can be reinflated by anti-differencing or summation, which requires the estimation of some initial conditions. Alternatively, the trend can be represented by an interpolated polynomial, which is subtracted from the data prior to filtering. Thereafter, the polynomial sequence can be added back to the filtered sequence.

Unlike the models that underlie the STAMP and SEATS–TRAMO programs, which are intended to be realistic representations of the processes generating the data, our model is an heuristic device that is intended only for the purpose of deriving the seasonal-adjustment filter. Nevertheless, the filter has much the same effects as those of the aforementioned programs.

The second method that we shall examine is a more flexible one, which operates in the frequency domain. By inspecting the periodogram of the residuals from a polynomial detrending of the data, one can determine the width of the frequency bands that contain the elements of the seasonal component that is to be eliminated.

In addition to the elements at the seasonal frequencies and its harmonics, these bands are liable to comprise several adjacent elements, which may also be removed from the data. Thereafter, the seasonally-adjusted data can be synthesised from the remaining spectral elements. (Equivalently, the seasonal component may be synthesised from the seasonal elements and subtracted from the original data.) The manner of dealing with trends in the case of the frequency-domain method is the same as in the case of the time domain-method.

The frequency-domain method envisages a circular sequence, which is the result of mapping the finite data sequence onto the circumference of a circle of the same length. The circular sequence gives rise to an infinite sequence, which is the periodic extension of the finite sequence. It is important that there should be no disjunctions in the periodic extension at the points where the end of one replication of the data sequence joins the beginning of the next.

Various devices are proposed for this purpose, which include extrapolating the data sequence and tapering it.

The spectral analysis of an econometric data sequence often reveals a fundamental component that falls within a low-frequency band running from the zero frequency up to a well-defined limit. The seasonal component of the data is liable to be separated from the fundamental component by a wide dead space containing nothing but the spectral traces of minor elements of noise.

Such a discovery may suggest that, instead of removing only the seasonal component from the data, one might aim to isolate the fundamental component. The process of isolating this band-limited component can be described as one of cleansing the data. This represents the third method of seasonal adjustment.

A band-limited data sequence cannot be modelled by an ARMA process, since this is supported on the entire frequency range running from zero up to the Nyquist frequency π radians per sample period. To make the band-limited process amenable to ARMA modelling, its frequency band must be mapped onto the interval $[0, \pi]$. This can be achieved by synthesising a continuous trajectory from the relevant Fourier ordinates, which can be re-sampled at the appropriate rate, which must be less than the original sampling rate.

A question arises concerning the effects of the alternative methods of seasonal adjustment. The choice of the method is bound to affect the parameters of any model that is build upon the data. However, a full investigation of this would carry us too far field. Instead, we choose to examine the extent to which the alternative methods give rise to differences in the dating of the turing points of a typical business cycle.

The methods that are described in this paper have been implemented in a computer program IDEOLOG that is available, together with its code in Pascal, at the address

`http://www.le.ac.uk/users/dsgp1/`

2 The Time-Domain Method

The method that we shall adopt for seasonal adjustment in the time domain has some affinities with the ARMA-model-based methods of seasonal adjustment that are represented most prominently by the TRAMO–SEATS and STAMP programs.

The airline passenger model

The TRAMO–SEATS program employs the airline passenger model of Box and Jenkins (1976) as its default model. This is represented by the equation

$$y(z) = \frac{N(z)}{P(z)}\varepsilon(z) = \left\{ \frac{(1-\rho z)(1-\theta z^s)}{(1-z)(1-\ z^s)} \right\} \varepsilon(z), \tag{1}$$

where $N(z)$ and $P(z)$ are polynomial operators and $y(z)$ and $\varepsilon(z)$ are, respectively, the z-transforms of the output sequence $y(t) = \{y_t; t = 0, \pm 1, \pm 2, \ldots\}$ and of the input sequence $\varepsilon(t) = \{\varepsilon_t; t = 0, \pm 1, \pm 2, \ldots\}$ of unobservable white-noise disturbances. The integer s stands for the number of periods in the year, which are $s = 4$ for quarterly data and $s = 12$ for monthly data. Without loss of generality as far as the derivation of the filters is concerned, the variance of the input sequence can be set to unity.

Given the identity $1 - z^s = (1 - z)\Sigma(z)$, where $\Sigma(z) = 1 + z + \cdots + z^{s-1}$ is the seasonal summation operator, it follows that

$$P(z) = (1 - z)(1 - z^s) = \nabla^2(z)\Sigma(z), \tag{2}$$

where $\nabla(z) = 1 - z$ is the backward difference operator. The polynomial $\Sigma(z)$ has zeros at the points $\exp\{\mathrm{i}(2\pi/s)j\}; j = 1, 2, \ldots, s - 1$, which are located on the circumference of the unit circle in the complex plane at angles from the horizontal that correspond to the fundamental seasonal frequency $\omega_s = 2\pi/s$ and its harmonics.

The TRAMO–SEATS program effects a decomposition of the data into a seasonal component and a non-seasonal component that are described by statistically independent processes driven by separate white-noise forcing functions. It espouses the principal of canonical decompositions that has been expounded by Hillmer and Tiao (1982).

The first step in this decomposition entails the following partial-fraction decomposition of the generating function of the autocovariances of $y(t)$:

$$\frac{N(z^{-1})N(z)}{P(z^{-1})P(z)} = \frac{U^*(z^{-1})U^*(z)}{\nabla^2(z^{-1})\nabla^2(z)} + \frac{V^*(z^{-1})V^*(z)}{\Sigma(z^{-1})\Sigma(z)} + \rho\theta. \tag{3}$$

Here, $\rho\theta$ is the quotient of the division of $N(z^{-1})N(z)$ by $P(z^{-1})P(z)$, which must occur before the remainder, which will be a proper fraction, can be decomposed.

In the preliminary decomposition of (3), the first term on the RHS corresponds to the trend component, the second term corresponds to the seasonal component and the third term corresponds to the irregular component. Hillmer and Tiao have provided expressions for the numerators of the RHS, which are somewhat complicated, albeit that the numerators can also be found by numerical means.

When $z = e^{\mathrm{i}\omega}$, equation (3) provides the spectral ordinates of the process and of its components at the frequency value of ω. The corresponding spectral density functions are obtained by letting ω run from 0 to π. The quotient $\rho\theta$ corresponds to the spectrum of a white-noise process, which is constant over the frequency range.

The principal of canonical decomposition proposes that the estimates of the trend and of the seasonal component should be devoid of any elements of white noise. Therefore, their spectra must be zero-valued at some point in the interval $[0, \pi]$. Let q_T and q_S be the minima of the spectral density functions associated with the trend and the seasonal components respectively.

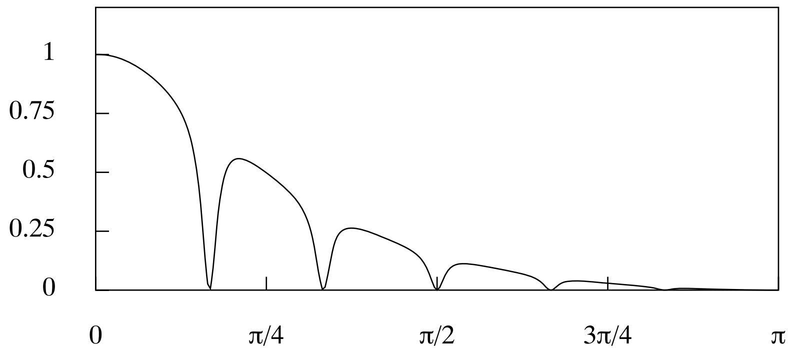

Figure 1: The gain of the canonical trend extraction filter associated with the airline monthly passenger model.

By subtracting these values from their respective components, a revised decomposition is obtained that fulfils the canonical principal. This is

$$\frac{N(z^{-1})N(z)}{P(z^{-1})P(z)} = \frac{U(z^{-1})U(z)}{\nabla^2(z^{-1})\nabla^2(z)} + \frac{V(z^{-1})V(z)}{\Sigma(z^{-1})\Sigma(z)} + q, \tag{4}$$

where $q = \rho\theta + q_T + q_S$.

The Wiener–Kolmogorov principle of signal extraction indicates that the filter that serves to extract the trend from the data sequence $y(t)$ should take the form of

$$\begin{aligned}\beta_T(z) &= \frac{U(z^{-1})U(z)}{\nabla^2(z^{-1})\nabla^2(z)} \times \frac{P(z^{-1})P(z)}{N(z^{-1})N(z)} \\ &= \frac{U(z^{-1})U(z)}{N(z^{-1})N(z)} \times \Sigma(z^{-1})\Sigma(z).\end{aligned} \tag{5}$$

This is the ratio of the autocovariance generating function of the trend component to that of the process as a whole. This filter nullifies the seasonal component in the process of extracting a trend that is relatively free of high-frequency elements. The nullification of the seasonal component is due to the factor $\Sigma(z)$.

The gain of the trend-extraction filter is depicted in Figure 1. Here, $s = 12$ and the values of $\rho = 0.4$ and $\theta = 0.6$ that determine the polynomial $N(z)$ are the estimates of Box and Jenkins (1976). The filter is an appropriate device for seasonal adjustment if the high-frequency elements that it serves to attenuate are liable to be regarded as a noisy contamination of no economic significance. In that case, as we shall propose later, it might be best to be remove them completely from the data.

The seasonal-adjustment filter, which nullifies the seasonal component without further attenuating the high-frequency elements of the data, is marginally more complicated. Define

$$\frac{W(z^{-1})W(z)}{\nabla^2(z^{-1})\nabla^2(z)} = \frac{U(z^{-1})U(z)}{\nabla^2(z^{-1})\nabla^2(z)} + q. \tag{6}$$

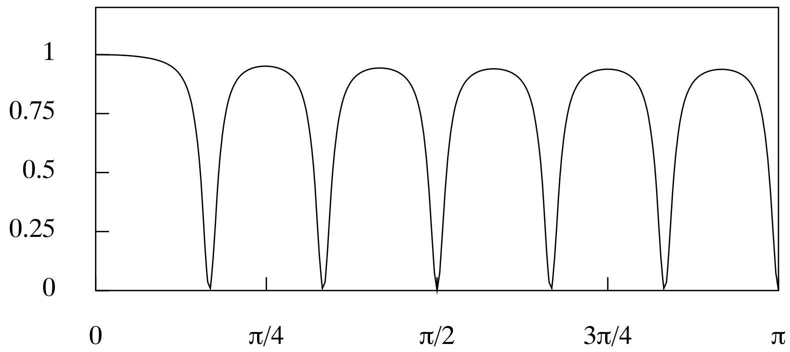

Figure 2: The gain of the seasonal adjustment filter associated with the airline passenger model.

Then, the seasonal adjustment filter is

$$\beta_A(z) = \frac{W(z^{-1})W(z)}{N(z^{-1})N(z)} \times \Sigma(z^{-1})\Sigma(z). \tag{7}$$

The gain of this filter is shown in Figure 2. Further examples of its gain, for various values of the parameters ρ and θ, are provided in a paper of Findlay (2005).

An alternative model

The filter that we shall propose is derived from a model that combines a white-noise component $\eta(t)$ with a seasonal component obtained by passing an independent white noise $\nu(t)$ through a rational filter with poles located on the unit circle at angles corresponding to the seasonal frequencies, and with zeros at the same angles, which lie inside the circle. The z-transform of the sequence $g(t)$, which is the data sequence or some transformation thereof that is free of trend, such as the differenced data or the residuals from the interpolation of a polynomial trend, is given by

$$g(z) = \frac{R(z)}{\Sigma(z)}\nu(z) + \eta(z), \tag{8}$$

where

$$R(z) = 1 + \rho z + \rho^2 z^2 + \cdots + \rho^{s-1} z^{s-1} \tag{9}$$

with $\rho < 1$, and where $\Sigma(z) = 1 + z + z^2 + \cdots + z^{s-1}$ is the seasonal summation operator, as previously defined.

This equation can be compared with that of the basic structural time series model

$$y(z) = \frac{z\zeta(z)}{\nabla^2(z)} + \frac{\xi(z)}{\nabla(z)} + \frac{\nu(z)}{\Sigma(z)} + \eta(z), \tag{10}$$

wherein $\zeta(z), \xi(z), \nu(z)$ and $\eta(z)$ are the z-transforms of independently distributed white-noise sequences. Equation (8) lacks the terms in $\nabla(z)$ and

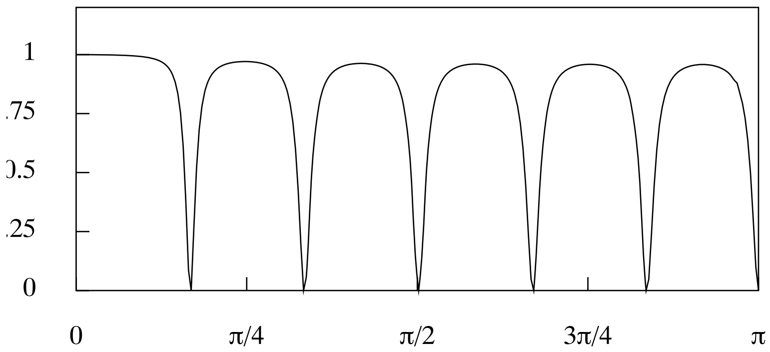

Figure 3: The gain of the seasonal adjustment filter associated with equation (8).

$\nabla^2(z)$, which give rise to a stochastic trend; and it includes the polynomial $R(z)$ in the numerator of the term in $\nu(z)$. The basic structural time series model is the default model of the STAMP program.

The seasonal-adjustment filter that extracts $\eta(z)$ from $g(z)$ of equation (8) is also the filter the extracts $\Sigma(z)\eta(z)$ from

$$\Sigma(z)g(z) = R(z)\nu(z) + \Sigma(z)\eta(z), \tag{11}$$

The z-transform of the filter, which is derived according to the Wiener–Kolmogorov principle, is

$$\beta_C(z) = \frac{\sigma_\eta^2\Sigma(z)\Sigma(z^{-1})}{\sigma_\eta^2\Sigma(z)\Sigma(z^{-1}) + \sigma_\nu^2 R(z)R(z^{-1})}. \tag{12}$$

Setting $z = \exp\{-\mathrm{i}\omega\}$ and letting ω run from 0 to π generates the frequency response of the filter, of which the modulus or gain is plotted in Figure 3 for the case where $s = 12$, $\rho = 0.6$ and $\lambda = \sigma_\eta^2/\sigma_\nu^2 = 0.125$.

There is hardly a difference between Figure 3 and Figure 2, which relates to the seasonal-adjustment filter associated with the airline passenger model. This is notwithstanding the fact that the airline passenger model incorporates a stochastic trend, whereas there is no trend term in equation (8). Also, the airline filter $\beta_A(z)$ has a numerator of degree $2s$ and a denominator of degree $2(s+1)$, whereas both the numerator and the denominator of $\beta_C(z)$ are of degrees $2(s-1)$.

In fact, $\beta_C(z)$ is a classic comb filter in which the zeros of the numerator polynomial $\Sigma(z)$, which are located on the perimeter of the unit circle, are balanced by poles of the denominator that fall on the same radii but which are located within the circle at a short distance from the perimeter. The poles counteract the effects of the zeros, except in the neighbourhoods of the seasonal frequencies and its harmonics, where the zeros account for the notches in the gain function.

The formulae for the seasonal-adjustment filters are not fully adapted to the circumstances of a finite data sequence with an underlying trend. In the following section, we shall show how the comb filter can be implemented.

The method depends upon eliminating the trend in the first instance before applying the filter. After the filtering, the trend can be restored. We shall deal first with the means of adapting the filter to a finite sample.

The finite-sample filter

To derive the finite-sample version of a Wiener–Kolmogorov filter, we may consider a data vector $g = [g_0, g_1, \ldots, g_{T-1},]'$ that has a seasonal component ξ and a noise component η:

$$g = \xi + \eta. \tag{13}$$

To cast the equations into a form analogous to (11), it is necessary to define the matrix counterparts of the polynomial operators $R(z)$ and $\Sigma(z)$.

Let $L_T = [e_1, e_2, \ldots, e_{T-1}, 0]$ be the matrix lag operator of order T, which is obtained from the identity matrix $I_T = [e_0, e_1, \ldots, e_{T-1}]$ by deleting the leading column and by appending a column of zeros to the end of the array. Then, by replacing z by L_T in the polynomial operators, we get

$$R(L_T) = \begin{bmatrix} Q'_{*R} \\ Q'_R \end{bmatrix} \quad \text{and} \quad \Sigma(L_T) = \begin{bmatrix} Q'_{*\Sigma} \\ Q'_\Sigma \end{bmatrix}. \tag{14}$$

These are banded lower-triangular Toeplitz matrices with units on the diagonals. Since the matrices Q'_{*R} and $Q'_{*\Sigma}$ suffer from end effects, they are liable to be discarded leaving Q'_R and Q'_Σ, which are of order $(T - s + 1) \times T$.

With these matrices, the following equation is formed, which is the matrix analogue of (11):

$$Q'_\Sigma g = Q'_R \nu + Q'_\Sigma \eta. \tag{15}$$

Here, ν is a vector of order T of white noise elements from the sequence $\nu(t)$. A demonstration by Pollock (2007) serves to show that the minimum-mean-square-error estimate of the vector η is given by

$$h = Q_\Sigma (Q'_\Sigma Q_\Sigma + \lambda^{-1} Q'_R Q_R)^{-1} Q'_\Sigma g. \tag{16}$$

The matrix of the transformation mapping from g to h is seen to be the analogue of the filter function $\beta_C(z)$ of (12).

A simple procedure for calculating the estimates of h is to solve the following equations in succession:

$$(Q'_\Sigma Q_\Sigma + \lambda^{-1} Q'_R Q_R) b = Q'_\Sigma g \quad \text{and} \quad h = Q_\Sigma b. \tag{17}$$

Since $Q'_\Sigma Q_\Sigma$ and $Q'_R Q_R$ correspond to the narrow-band dispersion matrices of moving-average processes, the solution to the first equation of (17) may be found via a Cholesky factorisation that sets $Q'_\Sigma Q_\Sigma + \lambda^{-1} Q'_R Q_R = GG'$, where G is a lower-triangular matrix with a limited number of nonzero bands. The system $GG'b = Q'_\Sigma g$ may be cast in the form of $Gp = Q'_\Sigma g$ and solved for p. Then, $G'b = p$ can be solved for b.

In dealing with trended economic data, there are two approaches that can be taken. The first approach depends upon eliminating the trend by applying a twofold differencing operator to the data. The difference operator is obtained by replacing the polynomial argument z of $\nabla^2(z) = 1 - 2z + z^2$ by the matrix lag operator L_T. The result is the matrix

$$\nabla^2(L_T) = \begin{bmatrix} Q'_* \\ Q' \end{bmatrix}. \tag{18}$$

Here, the sub-matrix Q'_* of order $2 \times T$ is liable to be discarded. The inverse of $\nabla^2(L_T)$ is the summation operator

$$\nabla^{-2}(L_T) = [\, S_* \quad S \,]. \tag{19}$$

We observe that, if $g_* = Q'_* y$ and $g = Q'y$ are available, then y can be recovered via the equation

$$y = S_* g_* + S g. \tag{20}$$

The two columns of the matrix S_* provide a basis for the set of all linear functions defined over the set of integers $t = 0, 1, \ldots, T-1$. Therefore, $f = S_* g_*$ is the vector of the ordinates of a linear trend, whilst the elements of g_* may be regarded as the parameters of the trend. If the elements of the vector $\delta' = [g'_*, g']$ have a nonzero mean value $\bar{\delta}$, such that $\delta = \bar{\delta}\iota_T + (\delta - \bar{\delta}\iota_T)$, where $\iota_T = [1, 1, \ldots, 1]'$ is the summation vector of order T, then the twofold summation $\nabla^2(I_T)\iota_T\bar{\delta}$ will give rise to a quadratic sequence.

If g within equation (16) represents the vector of differenced data and if h is the seasonally-adjusted version, then h requires to be reinflated in order to provide a seasonally-adjusted version of the original data vector. This can be denoted by $x = S_* h_* + Sh$, where h_* contains the appropriate initial conditions or constants of integration.

The way to determine h_* is to find the value that minimises the function

$$(y - x)'(y - x) = (y - S_* h_* - Sh)'(y - S_* h_* - Sh). \tag{21}$$

The effect should be to make the seasonally-adjusted data adhere as closely as possible to the original data. The solution is

$$h_* = (S'_* S_*)^{-1} S'_*(y - Sg). \tag{22}$$

The alternative procedure depends upon estimating a trend function, which can be removed from the data to create a set of residual values that are to be subjected to the process of seasonal adjustment. We shall denote the vector of the residual sequence also by g.

Often it is appropriate to represent the trend by an ordinary polynomial function of time. In the case of the U.K. consumption, which provides our first example, it is sufficient to interpolate a straight line through the logarithms of the data. Pollock (2007) has shown that an appropriate formula for the vector of residuals is

$$g = Q(Q'Q)^{-1}Q'y. \tag{23}$$

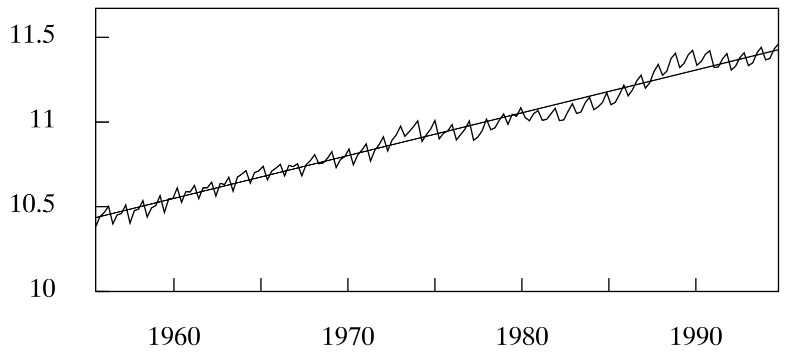

Figure 4: The quarterly sequence of the logarithms of consumption in the U.K., for the years 1955 to 1994, together with a linear trend interpolated by least-squares regression.

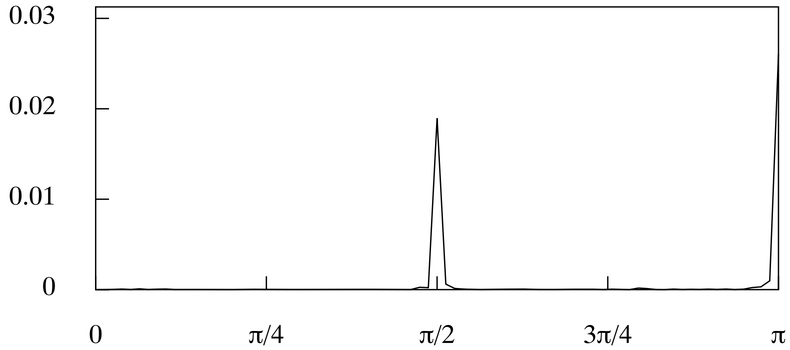

Figure 5: The periodogram of the first differences of the U.K. logarithmic consumption data.

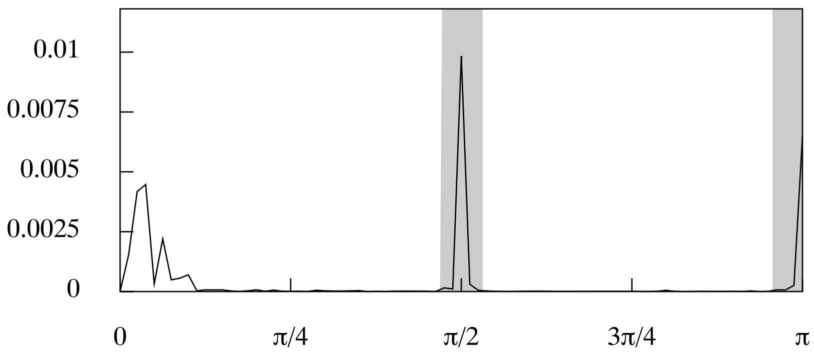

Figure 6: The periodogram of the residual sequence obtained from the linear detrending of the logarithmic consumption data. The shaded bands contain the elements of the seasonal component.

It is notable that this vector contains exactly the same information as does the vector $Q'y$ of the second differences of the data. The difference operator has the effect of nullifying the element of zero frequency and of attenuating radically the adjacent low-frequency elements. This is true even of the first-difference operator. Therefore, the low-frequency spectral structures of the data are not perceptible in the periodogram of the differenced sequence. Figure 5 provides evidence of this.

However, the periodogram of the residuals of the polynomial regression can be relied upon to reveal the spectral structures at all frequencies. Figure 6 shows that the low-frequency structure of the U.K. consumption data is fully evident in the periodogram of the residuals from fitting a linear trend to the logarithmic data.

Once the residual sequence has been seasonally adjusted, it can be added back to the interpolated polynomial. The effects of this procedure are indistinguishable from those of the seasonal-adjustment procedure that employs the operations of differencing to eliminate the trend and anti-differencing or summation to restore it.

Figure 7 shows the plot of the seasonally-adjusted data that would be obtained from either of the two procedures. Figure 8 shows the seasonal component that is extracted in the process of seasonal adjustment, and Figure 9 shows the periodogram of the residual sequence of the linearly detrended data after it has been subjected to seasonal adjustment. The removal of seasonal spikes from the periodogram is evident from the comparison of Figure 9 with of Figure 6.

The seasonal-adjustment filter eliminates from the data the elements at the seasonal frequencies, and it takes little else from the data. It is for this reason that the estimated seasonal component has the regular appearance that is see in Figure 8, for it is synthesised, in the main, from a strictly limited number of sinusoidal elements.

However, the periodograms of Figures of 6 and 9 suggest that the true seasonal component of the data comprises a wider range of elements, including some at frequencies that are the adjacent to the seasonal frequencies. Contributions from such elements can be included in the estimated seasonal component, to some extent, by widening the notches of the comb file. This can be done by reducing the value of the parameter ρ; but this is not a wholly satisfactory recourse.

The alternative method of seasonal adjustment that we propose in the next section requires the elements that constitute the seasonal component to be identified by inspecting the periodogram of the detrended data. Once these elements have been identified, the seasonal component can be synthesised and subtracted thereafter from the original data.

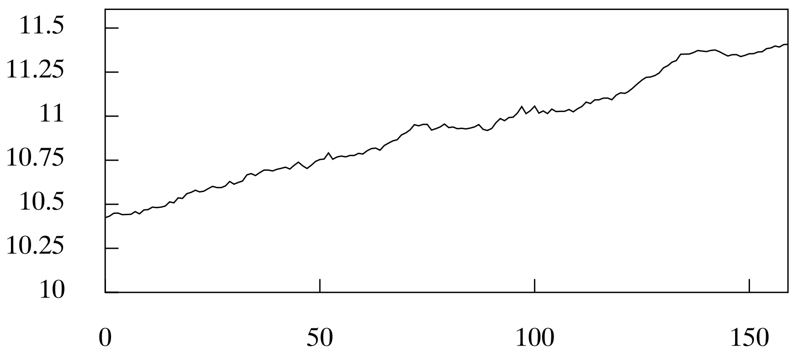

Figure 7: The plot of a seasonally-adjusted version of the logarithmic consumption data of Figure 4.

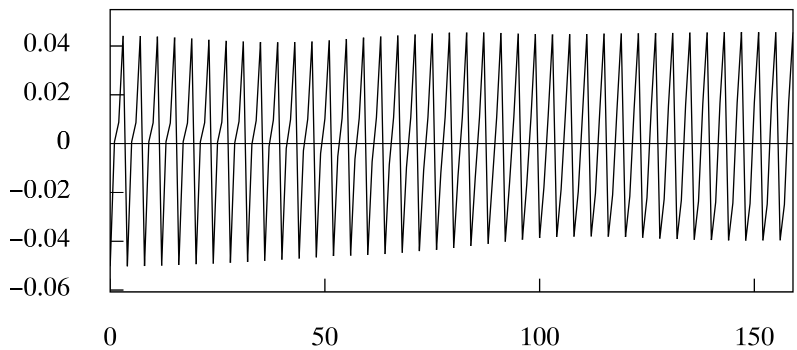

Figure 8: The seasonal component extracted from the logarithmic consumption data.

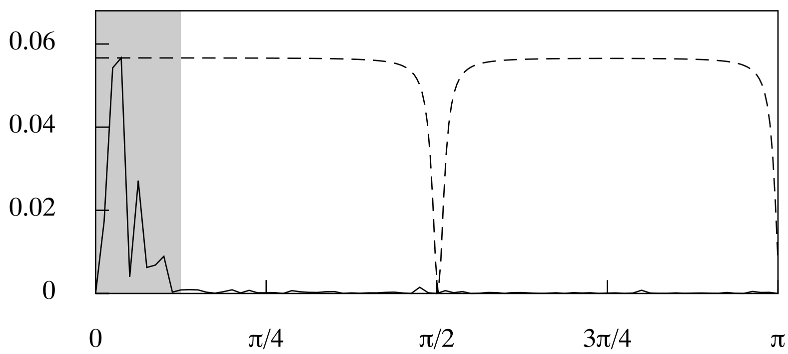

Figure 9: The periodogram of the residuals from a linear detrending the seasonally-adjusted logarithmic consumption data, with the frequency response function of the filter superimposed. The shaded band contains the elements of the business cycle.

3 Seasonal Adjustment in the Frequency Domain

In a Fourier analysis, an arbitrary function is resolved into a weighted combination of sine and cosine functions or, alternatively, of complex exponential functions. The domain of such periodic functions is either the perimeter of a circle, or else it is the entire real line.

A data sequence that is subject to a Fourier analysis must be regarded as a single cycle of a periodic function. The periodicity is achieved either by mapping the sequence onto the circumference of a circle, or else by extending it indefinitely by successive replications.

If a data sequence contains a significant trend, then there will be a sharp disjunction at the point on the circle where the beginning of the sequence joins its end. Alternatively, there will be successive disjunctions in the periodic extension of the sequence at points where the end of one replication of the data joins the start of the next replication. Then, the effect will be to create the appearance of the serrated edge of a saw blade.

The spectrum of a saw tooth function has a *one-over-f* profile in the form of a rectangular hyperbola that extends from a high point adjacent to the zero frequency to a low point at the limiting Nyquist frequency of π radians per sample interval. Such a profile is liable to mask the spectral information that is of genuine interest. Therefore, for a spectral analysis to be successful, the trend must be eliminated from the data at first.

The data can be detrended by differencing. Alternatively, a trend function can be interpolated into the data and the residual deviations can be subjected to the analysis. In the context of the seasonal adjustment of the data, the latter approach is preferable, since the pattern of seasonal fluctuations can be obscured by taking differences. For the present, we shall assume that the detrended data are available in the vector g. Later, we shall describe various devices that can accompany the process of polynomial detrending, which are intended to minimise the problem of the disjunctions.

It is more convenient to work with complex Fourier coefficients and with complex exponential functions in place of sines and cosines. In these terms, the Fourier transform and its inverse are given by

$$\gamma_j = \frac{1}{T}\sum_{t=0}^{T-1} g_t e^{-\mathrm{i}\omega_j t} \quad \longleftrightarrow g_t = \sum_{j=0}^{T-1} \gamma_j e^{\mathrm{i}\omega_j t}, \tag{24}$$

where $\omega_j = 2\pi j/T$ is the jth Fourier frequency, which, in the case of $j < T/2$, relates to a sinusoidal element that completes j cycles in the period spanned by the data. The conjugate frequencies ω_{T-j} are to be found within $\cos(\omega_j) = \{\exp(\omega_j)+\exp(\omega_{T-j})\}/2$ and $\cos(\omega_j) = \mathrm{i}\{\exp(\omega_j)-\exp(\omega_{T-j})\}/2)$, wherein $\exp(\omega_{T-j}) = \exp(\omega_{-j})$.

For a matrix representation of these transforms, one may define

$$\begin{aligned} U &= T^{-1/2}[\exp\{-\mathrm{i}2\pi tj/T\}; t,j = 0,\ldots,T-1], \\ \bar{U} &= T^{-1/2}[\exp\{\mathrm{i}2\pi tj/T\}; t,j = 0,\ldots,T-1], \end{aligned} \tag{25}$$

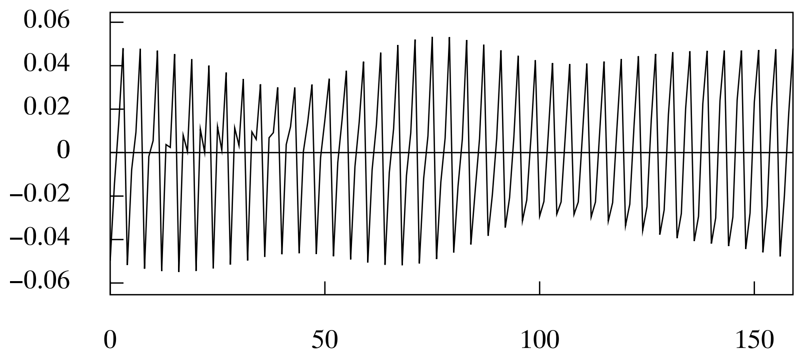

Figure 10: The seasonal component extracted from the logarithmic consumption data by the frequency-domain method.

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

$$\gamma = T^{-1/2}Ug \quad \longleftrightarrow \quad g = T^{1/2}\bar{U}\gamma, \tag{26}$$

where $g = [g_0, g_1, \ldots g_{T-1}]'$ and $\gamma = [\gamma_0, \gamma_1, \ldots \gamma_{T-1}]'$ are the vectors of the data and of their spectral ordinates, respectively.

This notation can be used to advantage for representing the process of estimating the seasonal component. Let J be a diagonal selection matrix of order T of zeros and units, wherein the units correspond to the frequencies of the pass band and the zeros to those of the stop band. Then, the selected Fourier ordinates are the nonzero elements of the vector $J\gamma$. By an application of the inverse Fourier transform, the selected elements are carried back to the time domain to form the filtered sequence, which is the estimated component w. Thus, there is

$$w = \bar{U}JUg = \Psi g. \tag{27}$$

Here, $\bar{U}JU = \Psi = [\psi^{\circ}_{|i-j|}; i, j = 0, \ldots, T-1]$ is a circulant matrix of the filter coefficients that would result from wrapping the infinite sequence of the coefficients of the ideal bandpass filter around a circle of circumference T and adding the overlying elements. Thus

$$\psi^{\circ}_k = \sum_{q=-\infty}^{\infty} \psi_{qT+k}. \tag{28}$$

Applying the wrapped filter to the finite data sequence via a circular convolution is equivalent to applying the original filter to an infinite periodic extension of the data sequence. In practice, the wrapped coefficients of the time-domain filter matrix Ψ would be obtained from the Fourier transform of the vector of the diagonal elements of the matrix J. However, it is more efficient to perform the filtering by operating upon the Fourier ordinates in the frequency domain, which is how the program operates.

The periodogram of Figure 6 shows evidence of some minor spectral elements at frequencies adjacent to the seasonal frequencies. The seasonal component of Figure 10 has been synthesised from the five elements in the interval $[\pi/2 - 4\pi/T, \pi/2 + 4\pi/T]$ and from the four elements in the interval $[\pi - 6\pi/T, \pi]$. These are covered by the highlighted bands. It shows more variability than the component of Figure 8, which has been extracted by the comb filter that operates in the time domain.

The method of frequency-domain filtering can be used to mimic the effects of any linear time-invariant filter, operating upon stationary data in the time domain, that has a well-defined frequency-response function. All that is required is to replace the selection matrix J of equation (27), consisting of zeros and units, by a diagonal matrix containing the ordinates of the desired frequency response, sampled at points corresponding to the Fourier frequencies.

Tapering and extrapolations

Various devices are available for ensuring that, when the data sequence is wrapped around the circumference of a circle, there is no disjunction at the point where its head joints its tail.

The conventional means of avoiding such disjunctions is to taper a detrended and mean-adjusted sequence so that both ends decay to zero. (See Bloomfield 1976, for example.) The disadvantage of this recourse is that it falsifies the data at the ends of the sequence, which is particularly inconvenient if, as is often the case in economics, attention is focussed on the most recent data. To avoid this difficulty, the tapering can be applied to some extrapolations, which can be added to the data, after it has been subject to a preliminary detrending.

The preliminary detrending can be achieved by interpolating a polynomial function of time of by using the Leser or Hodrick–Prescott filter. (This filter, which is commonly attributed by economists to Hodrick and Prescott (1980, 1997), was expounded by Leser (1961) in an earlier publication.) The interpolated function should be a stiff one containing only periodic elements of the lowest frequencies. It is also desirable that the function should pass through the midst of the scatters of points at either end of the data sequence. For this purpose, a method of weighted least-squares polynomial regression can be used that allows extra weight to be placed upon the initial and the final runs of observations.

The method of weighted least-squares can also be used in the context of the Leser filter. Here, an additional flexibility is available by allowing the value of the smoothing parameter to vary. By attributing a low value to the parameter within the appropriate locality, a sharp turn in the data or an evident structural break can be absorbed by the trend, thereby allowing the residual sequence to maintain its normal behaviour.

A tapered sequence, based on successive repetitions of the ultimate seasonal cycle, can be added to the end of the data sequence, and a similar sequence, based on the first cycle, can be added before the start. However, such extrapolations tend to misrepresent the seasonal fluctuations by impos-

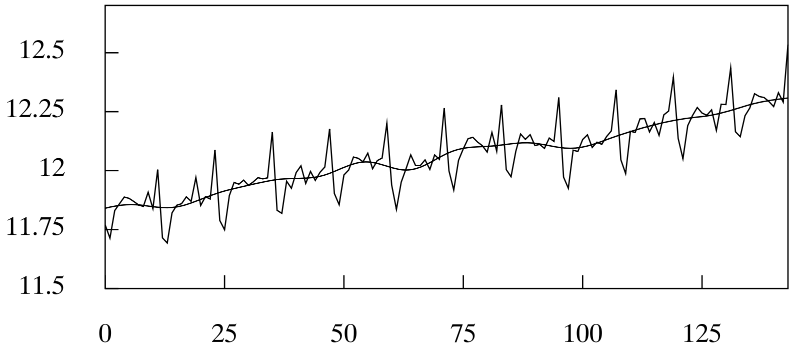

Figure 11: The logarithms of U.S. total retail sales from January 1953 to December 1964 with an interpolated trend-cycle function.

ing gradual reductions in their amplitude. The amplitudes can be preserved by inserting a segment into the circular data sequence in which the seasonal pattern in the final year is transformed gradually into the pattern of the first year. This process can be described as one of morphing the data, which is an allusion to a popular technique in computer graphics.

Let s be the number of months or seasons in the year, and let the data be supplemented by a sequence of points, indexed by $j = 0, 1, \ldots, Ns - 1$ that correspond to an integral number of years. To avoid the use of subscripted indices, let the (detrended) sample points be denoted by $g[t]; t = 0, 1, \ldots, T - 1$. Then, the first year and the final year are replicated N times in sequences in which the elements are defined, respectively, by

$$g^S[j] = g[j \bmod s] \qquad \text{and} \qquad g^F[j] = g[T - s + (j \bmod s)]. \tag{29}$$

A convex combination of these sequences with varying weights is given by

$$\begin{aligned} g^E[j] &= \lambda_j g^F[j] + (1 - \lambda) g^S[j], \quad \text{with} \\ \lambda_j &= \frac{1}{2}\{\cos(\theta_j) + 1\}, \quad \text{where} \quad \theta_j = \frac{\pi j}{Ns}. \end{aligned} \tag{30}$$

The weights λ_j are described by a half-cycle of a raised cosine function. The resulting sequence can be added to the end of the linear data sequence, which means that it will be interpolated between the finish and the start of the circular sequence.

The various devices described in this section are illustrated in a sequence of graphs. Figure 11 shows the logarithms of 144 monthly observations on retail sales in U.S. for the period from January 1953 to December 1964. The data have been taken from the monograph of Siskin *et al.* (1967) that described the X-11 program for seasonal adjustment. Interpolated through these data is a trend line, which is one of the eventual products of the analysis.

Figure 12 contains the residual deviations obtained by interpolating a linear trend through the data of Figure 11. The residuals from the mid point of the sample to its end are displayed on the left of the figure. They

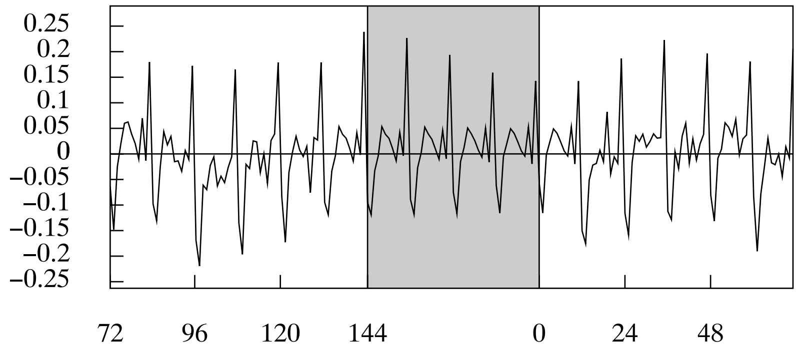

Figure 12: The residuals from a linear detrending of the sales data, with an interpolation of four years length inserted between the end and the beginning of the circularised sequence, marked by the shaded band.

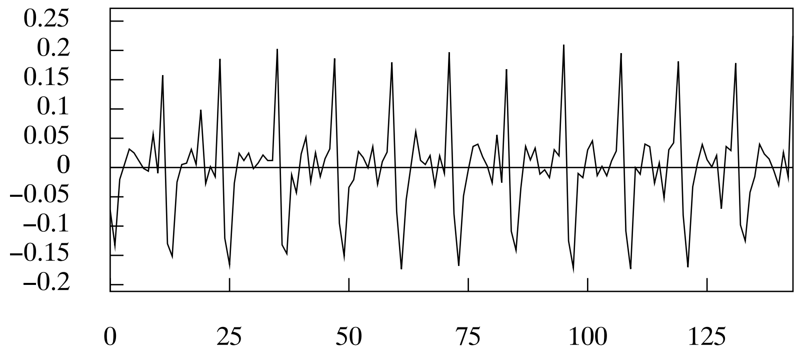

Figure 13: The sequence of residual deviations of the sales data from their trend, which may be regarded as the seasonal component.

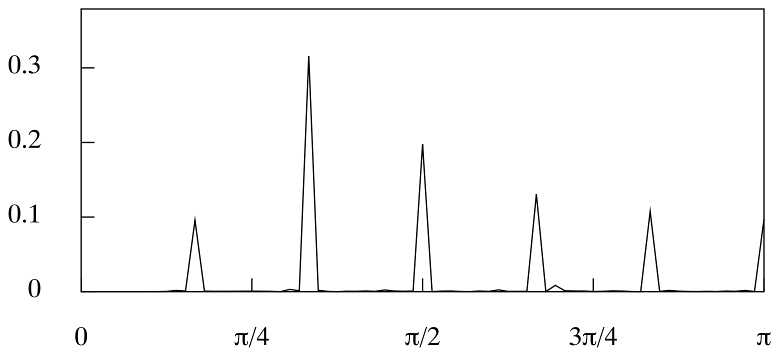

Figure 14: The periodogram of the residual deviations of the sales data from their trend.

are followed by a segment of artificial data, of four years duration, in which the pattern of the fluctuations of the final year of the data is gradually transformed into the pattern of the first year. This is followed by the residuals from the beginning of the sample to its mid point.

An inspection of the periodogram of the residuals of Figure 12 would show that they posses a low-frequency component falling in the interval $[0, \pi/8]$. By removing this component, discarding the artificial elements and adding the remainder to the linear trend, the revised trend of Figure 11 is generated, which is apt to be described as the trend-cycle component. The deviations of the data from the revised trend are shown in Figure 13 and their periodogram is displayed in Figure 14.

The periodogram of Figure 14 contain virtually nothing that cannot be attributed to the seasonal fluctuations. Therefore, the complement, within the original data, of the sequence of Figure 13, which is the trend line of Figure 11, can be regarded as the seasonally-adjusted data. It should be recognised that these seasonally-adjusted data has been obtained by cleansing the original data of all components of frequencies in excess of $\pi/8$.

4 Resampling Band-limited Data

We now consider the means of generating a continuous function to represent the trajectory of the business cycle that underlies the sampled data. We may assume that the data in the vector $g = [g_0, g_1, \ldots, g_{T-1}]'$ is the residual sequence from fitting a polynomial trend to the original data.

Since its frequency content resides within a sub interval of the Nyquist range, the business cycle trajectory has a minimal discrete-time representation comprising a set of points whose number is a fraction of the original sample size. These points are obtained by sampling the continuous trajectory at regular intervals. The estimate of a discrete-time ARMA model, intended to represent the business cycle, should be based on the minimal discrete-time representation.

When the value of t is allowed to vary continuously in the interval $[0, T)$, the formula for the Fourier synthesis of a finite sequence will generate a continuous periodic function that interpolates the ordinates of the sequence. The formula in question, which is from (24), can be written as

$$g(t) = \sum_{j=0}^{T} \gamma_j e^{\mathrm{i}\omega_j t} = \alpha_0 + \sum_{j=1}^{[T/2]} \{\alpha_j \cos(\omega_j t) + \beta_j \cos(\omega_j t)\}, \qquad (31)$$

where $\alpha_0 = \gamma_0$ and $\alpha_j = \gamma_j + \gamma_{T-j}$ and $\beta_j = \mathrm{i}(\gamma_j - \gamma_{T-j})$ for $j > 0$, and where $[T/2]$ is the integral part of $T/2$. It simplifies matters to assume that $T/2 = n$ is a whole number, which we shall do hereafter. (Observe that the segment of $g(t)$ that falls in the interval $(T-1, T)$ bridges the gap in the circular sequence between the final sample point g_{T-1} and the initial point $g_T = g_0$.)

Imagine that the maximum frequency within the business cycle component is $\omega_d = 2\pi d/T < \pi$. In that case, there will be $N = 2d$ Fourier coefficients, and the continuous function that is limited to the frequency interval $[0, \omega_d]$ will be given by

$$\begin{aligned} h(t) &= \alpha_0 + \sum_{j=1}^{d}\left\{\alpha_j \cos\left(\frac{2\pi jt}{T}\right)\right\} + \sum_{j=1}^{d-1}\left\{\beta_j \sin\left(\frac{2\pi jt}{T}\right)\right\} \\ &= \alpha_0 + \sum_{j=1}^{d}\left\{\alpha_j \cos\left(\frac{2\pi j\tau}{N}\right)\right\} + \sum_{j=1}^{d-1}\left\{\beta_j \sin\left(\frac{2\pi j\tau}{N}\right)\right\}. \end{aligned} \tag{32}$$

Here, $\tau = tN/T$ varies continuously in $[0, N)$, whereas t varies continuously in $[0, T)$. Thus, on the RHS, there is a new set of Fourier frequencies $\{\varpi_j = 2\pi j/N; j = 0, 1, \ldots, d\}$, relative to a lower rate of sampling, of which the maximum value is now the limiting Nyquist frequency of $2\pi d/N = \pi$ radians per sample interval.

The N coefficients $\{\alpha_0, \alpha_1, \beta_1, \ldots, \alpha_{d-1}, \beta_{d-1}, \alpha_d\}$ bear a one-to-one correspondence with the set of N ordinates $\{h_\tau = h(\tau T/N); \tau = 0, \ldots, N-1\}$ sampled at intervals of $\pi/\omega_d = T/N$ from $h(t)$. The consequence is that $h(t)$ is fully represented by the resampled data $h_\tau; \tau = 0, \ldots, N-1$, from which it may be derived by Fourier interpolation. Since the spectral support of the resampled data is the full Nyquist interval, the data should now be amenable to ARMA modelling, on the condition that the underlying processes is statistically stationary.

Non-stationary integrated processes of various orders can be derived from a stationary continuous processes. The generic integrals are

$$\int_0^t \cos(\omega_j\tau)d\tau = \left[\frac{\sin(\omega_j\tau)}{\omega_j}\right]_0^t = \frac{\sin(\omega_j\tau)}{\omega_j} \tag{33}$$

and

$$\int_0^t \sin(\omega_j\tau)d\tau = \left[\frac{-\cos(\omega_j\tau)}{\omega_j}\right]_0^t = \frac{1}{\omega_j} - \frac{\cos(\omega_j\tau)}{\omega_j}. \tag{34}$$

Let $g(t)$ of (32) stand for any continuous function derived by the Fourier interpolation of T data points. Then, it follows that the first integral of $g(t)$ is

$$g^{(1)}(t) = \sum_{j=1}^{n-1}\frac{\beta_j}{\omega_j} + \alpha_0 t + \sum_{j=1}^{n}\frac{\alpha_j}{\omega_j}\sin(\omega t) - \sum_{j=1}^{n-1}\frac{\beta_j}{\omega_j}\cos(\omega t). \tag{35}$$

This is a combination of a linear function and a sum of trigonometrical functions. A further integration will produce a combination of a quadratic function and a trigonometrical sum.

In general, an integrated frequency-limited process of any order of integration can be reconstituted for its sampled ordinates by subtracting the appropriate discrete polynomial sequence, by applying the process of Fourier interpolation to the residuals and then by adding the resulting function back to the continuous version of the polynomial function.

For an example of the techniques of this section, we may revert to the logarithmic data on U.K. consumption that are represented in Figure 4. A seasonally-adjusted version of the data sequence is represented in Figure 7 and the periodogram of the residual sequence from a linear detrending of these seasonally-adjusted data is shown in Figure 9.

Figure 9 shows that the essential component of the seasonally-adjusted data resides in the frequency band $[0, \pi/8]$, which is highlighted in the diagram. Elsewhere in the frequency range, there are spectral traces of minor elements of noise. These have the effect of roughening the profile of the seasonally adjusted data of Figure 7. It seems reasonable to cleanse the data of this noise and to see the result as representative both of the non-seasonal component and of the trend-cycle component.

Figure 15 shows the effects of synthesising a continuous function from the Fourier ordinates of the residual sequence that lie in the interval $[0, \pi/8]$. The function is superimposed upon the deviations of the original data from the interpolated linear trend that is depicted in Figure 4. In our judgement, this continuous function is a fair representation of the business cycle, since the linear trend provides a good benchmark against which to measure its fluctuations. The sum of the continuous function of Figure 15 and the trend of Figure 4 is the continuous trend-cycle component that is displayed in Figure 16.

The trend-cycle component will not be affected by varying the way in which its cyclical elements are shared between the trend and the cycle. Therefore, for the purpose of estimating the component, one need not insist on a particular definition of the trend. Nevertheless, there may be firm opinions on how the component should be divided into its two parts.

In the opinion of many economists it is appropriate to represent the trend by a function that is more flexible than a polynomial of low degree. The Leser or Hodrick–Prescott filter is often chosen for the purpose. Figure 17 shows a residual trajectory that has been obtained by applying the filter to data sampled at quarterly intervals from the continuous function of Figure 16. The trajectory of Figure 17 has been created by using a cubic spline to interpolate a continuous function through the points of the resulting residual sequence.

A smoothing parameter with the conventional value of 1600 has been used within the Leser filter. This generates a trend curve that is sufficiently flexible to absorb the greater part of the departures of the trend-cycle trajectory from the linear trend. As a result, the residual trajectory of Figure 17, which purports to represent the business cycles, show a remarkable regularity.

The purpose of this example is to show how easy it is to obtain spurious results in analysing business cycles. The spurious nature of the results would be less evident, and the results would seem more plausible, if we were to allow the cycles of Figure 17 to be obscured by the minor elements of noise that we have chosen to remove from the data at the outset.

There are liable to be problems of spurious regularity in the residue whenever a flexible trend is fitted to the data. The problems are not peculiar to the Hodrick–Prescott or Leser filter, as some authors have suggested. (See

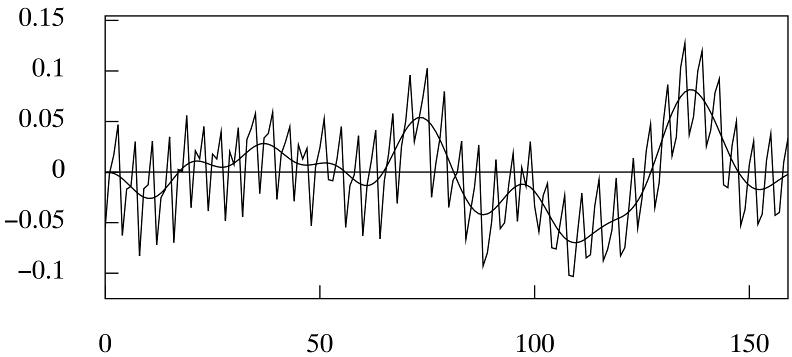

Figure 15: The residual sequence from fitting a linear trend to the logarithmic consumption data with an interpolated function representing the business cycle.

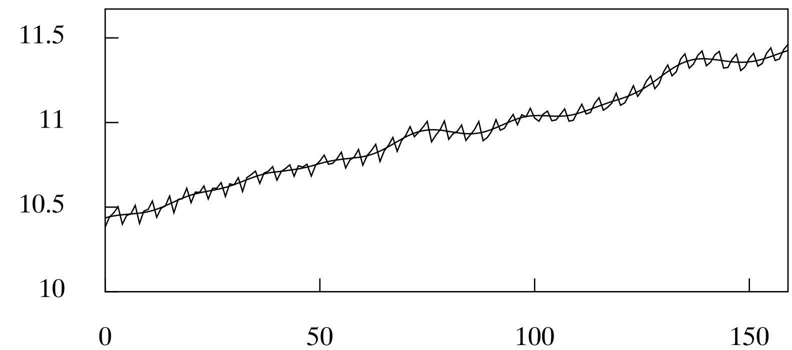

Figure 16: The trend-cycle component of U.K. consumption determined by the frequency-domain method, superimposed on the logarithmic data.

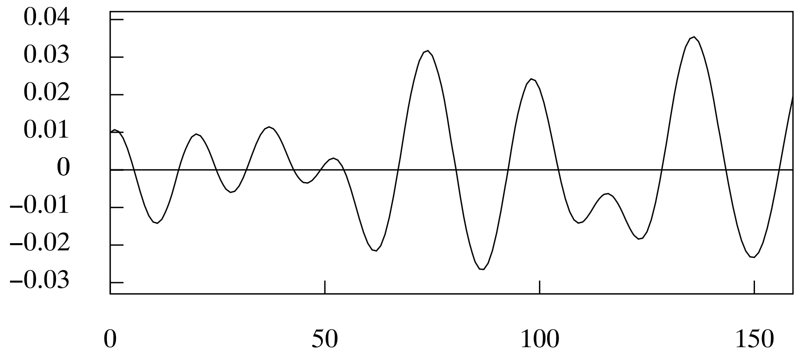

Figure 17: The business cycle determined by the Hodrick–Prescott filter.

Cogley and Nason 1995 and Harvey and Jaeger 1993)

There are undoubtedly instances where the smooth trajectory of a polynomial of low degree is inappropriate to the underlying dynamics of the economy. Examples are provided by the disruptions at the ends of the two world wars, the circumstances at the beginning of the great interwar depression and, perhaps, the effects of the present financial crises.

Such disruptions can be reflected in the trend function by making it sufficiently flexible in the appropriate localities. A device for this purpose, which is available in th IDEOLOG computer program, takes the form of a Leser filter with a variable smoothing parameter.

5 Detecting Turning Points

The choice of a method for detrending a macroeconomic data sequence greatly affects the determination of the turning points of the business cycle. So crucial is this choice that some economists have advised that it is best to avoid detrending altogether and to concentrate on the task of finding the turning points in the original data. (See, for example, Harding and Pagan 2002.)

The locations of the turning points in a trending sequence are affected by its rate of growth. In an economy that is growing more rapidly than another that is subject to the same cyclical influences, the peaks of the business cycle will be delayed and the troughs will be advanced, thereby shortening the periods of recession. Therefore, if the timing of the cycles of economies that are growing at different rates are to be compared, then it will be necessary to remove the trends from their data.

If detrending the data is unavoidable in a comparative analysis, then it is important to have a clear idea of the effects of the different methods of detrending. One advantage of representing the business cycle by the band-limited function of equation (31) lies in the fact that the function is differentiable. Therefore, the turning points of the business cycle are identified by the points where the derivative function cuts the horizontal axis.

Figure 18 shows the function that is obtained by differentiating the business cycle function of Figure 15. 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 at a distance that corresponds to the slope of the log-linear trend line of Figure 6, which represents the underlying rate of growth of U.K. consumption. The intersection of the derivative function with this line indicates the turning points in the seasonally-adjusted data of Figure 16, which represent the trend-cycle component.

The derivative of the Leser (Hodrick–Prescott) trend function that underlies Figure 17 is available from the derivatives of the cubic spline. Its intersections with the derivative business-cycle function would serve to identify the turning points of Figure 18.

Figure 19 shows a sequence that has been derived from the detrended

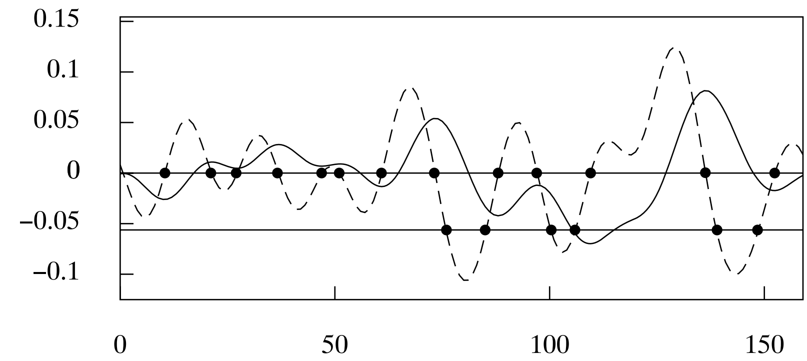

Figure 18: The turning points of the business cycle marked on the horizontal axis by black dots. The solid line is the business cycle of Figure 15. The broken line is the derivative function.

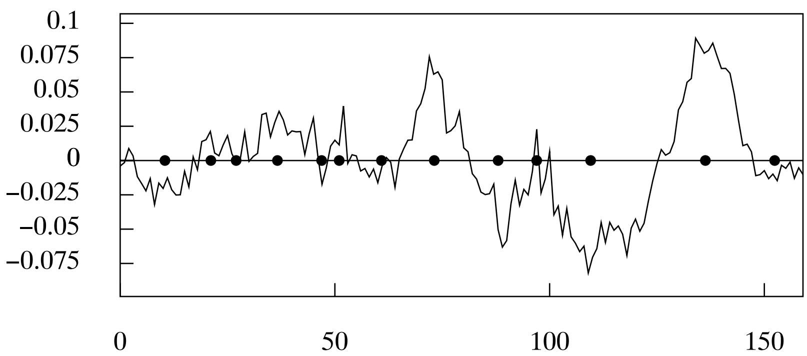

Figure 19: A sequence derived from the detrended data of Figure 15 via the time-domain method of seasonal adjustment.

data of Figure 15 via the time-domain method of seasonal adjustment. This is, in effect, a linearly detrended version of the sequence of Figure 7. On the horizontal axis are the turning points from Figure 18. Although a careful analysis of the data of Figure 19 would probably identify similar turning points, the potential for being misled by the such data, which is full of pitfalls and spikes, is manifest.

6 Conclusion

The empirical analysis of this paper has been based on a careful inspection of the periodograms of the detrended data. The periodograms have revealed data components that fall within well segregated frequency bands. Experience had shown that such spectral structures are common amongst economic data sequences.

These findings have clear implications for how one should set about the task of extracting the data components. In particular, they suggest that the estimates of the components should be purged of all elements of noise that lie in the spectral dead spaces. Then, the estimated components will be band-limited in frequency. To make the resulting business cycle component amenable to the usual techniques of ARMA modelling, its sequence will need to be subjected to the resampling procedures that are described in section 4 of the paper.

A discrete Fourier analysis, on which a periodogram is based, embodies the theoretical fiction that the data sequence corresponds to a single cycle of a periodic function. Contrary to a naive supposition, this has no implications for what we expect to observe beyond the limits of the sample. However, a Fourier analysis does require a careful detrending of the data. In particular, when the detrended data sequence is disposed around a circle, there should be no disjunction at the point where its head meets its tail. This can be achieved by the methods of extrapolation, tapering and morphing that have been incorporated in the computer program that is an adjunct of this paper.

The presence of any disjunctions in the data, including those that economists describe as structural breaks, will given rise to a slew of spectral power that will tend to obscure the finer details of their spectral structure. The resulting periodogram is liable to have the *one-over-f* profile of a rectangular hyperbola, which descends from a high point adjacent to the zero frequency to a low point at the limiting Nyquist frequency.

Granger (1966) has described such spectra as typical of econometric time series. His characterisation has greatly influenced the perceptions of econometricians. It has suggested to many of them that the spectral analysis of economic data is liable to be difficult, if not fruitless. However, in many cases, such spectra are the products of inadequate detrending. In this paper, we have reaffirmed the practicality and the importance of the spectral analysis of econometric data.

References

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

Box, G.E.P., and G.M. Jenkins, (1976), *Time Series Analysis: Forecasting and Control, Revised Edition,* Holden Day, San Francisco.

Caporello G., and A. Maravall, (2004), *Program TSW, Revised Reference Manual,* Servicio de Estudios, Banco de España.

Cogburn, R., and H.T. Davis, (1974), Periodic Splines and Spectral Estimation, *Annals of Statistics,* 2, 1108–1126.

Cogley, T., and J.M. Nason, (1995), Effects of the Hodrick–Prescott filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research, *Journal of Economic Dynamics and Control,* 19, 253–278.

Dagum, E.B., (1980), The X-11-ARIMA Seasonal Adjustment Method, *Technical Report 12-564E,* Statistics Canada, Ottawa.

Dagum, E.B., (1988), *The X11-ARIMA/88 Seasonal Adjustment Method—Foundations and User's Manual,* Time Series Research and Analysis Division, Statistics Canada Technical Report.

Findley, D.F., B.C. Monsell, W.R. Bell, M.C. Otto and B. Chen, (1998), New Capabilities and Methods of the X12-ARIMA Seasonal Adjustment Program, *Journal of Business and Economic Statistics,* 16, 127–152.

Findley, D.F., (2005), Some Recent Developments and Directions in Seasonal Adjustment, *Journal of Official Statistics,* 21, 343–365.

Findley, D.F., and D.K. Martin (2006), Frequency Domain Analyses of SEATS and X-11/12-ARIMA Seasonal Adjustment Filters for Short and Moderate-Length Time Series, *Journal of Official Statistics,* 22, 1–34.

Granger, C.W.J, (1966), The Typical Spectral Shape of an Economic Variable, *Econometrica,* 34, 150–161.

Harding, D., and A. Pagan, (2002), Dissecting the Cycle, *Journal of Monetary Economics,* 49, 365–381.

Harvey, A.C., and A. Jaeger, (1993), Detrending, Stylised Facts and the Business Cycle, *Journal of Applied Econometrics,* 8, 231–247.

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, 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.

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.

Ladiray, D., and B. Quenneville, (2001), *Seasonal Adjustment with the X-11 Method,* Springer Lecture Notes in Statistics 158, Springer Verlag, Berlin.

Leser, C.E.V., (1961), A Simple Method of Trend Construction, *Journal of the Royal Statistical Society,* Series B, 23, 91–107.

Monsell, B.C., J.A.D. Aston and S.J. Koopman, (2003), *Toward X-13?* Discussion paper of the U.S. Census Bureau,
`http://www.census.gov/ts/papers/jsm2003bcm.pdf`

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

Pollock, D.S.G., (2007), Wiener–Kolmogorov Filtering, Frequency-Selective Filtering and Polynomial Regression, *Econometric Theory,* 23, 71–83.

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