

The Manual for IDEOLOG.PAS
A Program for Filtering Econometric Data

by
D.S.G. Pollock

London, April 2018

Contents

Part1: The User's Manual

1. Introduction
2. The Menus of the Program
3. Data Handling
 - 3.1 Data Entry
 - 3.2 Data Plotting
 - 3.3 Data Segmentation
 - 3.4 Data Transformations
 - 3.5 Saving the Data
5. Graph Plotting
 - 5.1 Page Parameters
 - 5.2 PostScript Files
6. Data Characteristics
 - 6.1 The Autocorrelation Function
 - 6.2 The Periodogram
7. Trends and Baselines
 - 7.1 Polynomial Regression
 - 7.2 Weighted Polynomial Regression
 - 7.3 Flexible H-P filtering
4. Data Extrapolation and Tapering
8. FIR Filters
 - 8.1 The The Binomial Lowpass Filter
 - 8.2 The Henderson Filters
 - 8.3 The Baxter-King Filter
 - 8.e The Proposal of Christiano and Fitzgerald
9. Wiener-Kolmogorov Filters
 - 9.1 The Leser or Hodrick-Prescott Filter
 - 9.2 The Butterworth Filter
 - 9.3 A Seasonal-Adjustment Filter
10. Frequency-Domain Filters
 - 10.1 The Ideal Lowpass and Bandpass Filters
 - 10.2 Seasonal Adjustment with the Bandstop Filter
 - 10.3 The Butterworth Frequency-Domain Filter
 - 10.4 The Hordick-Prescott Frequency-Domain Filter
11. Case Studies
 - 12.1 The Consumption Cycle and its Turning Points
 - 12.2 The Long-run Trend of U.K. GDP and its Structural Breaks
 - 12.3 The Trend-cycle Component of Monthly Data on U.S. Retail Sales

Part 2: IDEOLOG, A Program for Filtering Econometric Data A Synopsis of Alternative Method

1. Introduction
 2. The Effects of a Linear Filter
 3. The FIR filters
 - 3.1 Approximation to the Ideal Filter
 4. The Wiener–Kolmogorov Filters
 - 4.1 Filters for Trended Data
 - 4.2 A Filter for Seasonal Adjustment
 5. The Finite-Sample Realisations of the W–K Filters
 - 5.1 Filters for Short Trended Sequences
 - 5.2 The Leser (H–P) Filter and the Butterworth Filter
 - 5.3 A Flexible Smoothing Filter
 - 5.4 A Seasonal-Adjustment Filter
 6. The Frequency-Domain Filters
 - 6.1 Complex Exponentials and the Fourier Transform
 - 6.2 Extrapolations and Detrending
 - 6.3 Anti-Differencing
 - 6.4 Seasonal Adjustment in the Frequency Domain
 7. The Program and its Code
- References

1. INTRODUCTION

The IDEOLOG program originated in the desire to compare some new methods of filtering with existing procedures that are common in econometric analysis. The outcome has been a comprehensive facility that will enable a detailed analysis of univariate econometric time series, as well as time series originating from many other investigations. The program will serve to reveal the extent to which the results of an economic analysis might be the consequence of the choice of a particular filter.

The new filtering procedures are based on the Fourier analysis of the data, and they perform their essential operations in the frequency domain as opposed to the time domain. They depend upon a Fourier transform for carrying the data into the frequency domain and upon an inverse transform for carrying the filtered elements back to the time domain. Filtering procedures usually operate exclusively in the time domain. This is notwithstanding fact that, for a proper understanding of the effects of a filter, one must know its frequency-response function.

The program is operated by typing the numbers and letters that are written beside the menu options, by answering *Yes* or *No* (*Y/N*) to various questions and by entering the required parameters and numerical values. The procedures of the program are either self-explanatory or else they are explained in the appropriate places by short texts displayed on the computer screen.

Many of the procedures require pre-requisite information to be supplied or prior choices to be made. If any such items are missing, then the program will ask for them. Thus, for example, if the procedure to plot a graph is activated before any data has been provided and before the parameters of the graph have been declared, then the program will ask the user to provide these items.

The program is also protected against inappropriate inputs. Thus, if a non-numeric symbol is typed when a bounded integer has been called for, then the program will persist in asking for the integer until one that satisfies the bounds has been provided.

The user of the program should be able to operate it without the aid of this manual, provided that they are familiar with the concepts that are explained in the accompanying document titled *IDEOLOG: A Program for Filtering Econometric Data*, which is in Part 2 of this manual.

The first part of the manual lists the menus of the program, and it gives a thematic account of the facilities of the program. It should serve as a ready reference. The various procedures that are described in this account are accessed via the sub menus of the program. The pathways to these procedures are listed under their headlines in the following manner:

(IDEOLOG.PAS → *Primary Menu* → *Secondary Menu* → *Procedure*)

The program resides in a zip file at the following address:

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

The code of the program, which is in Pascal, can be found within the zip file, which also contains an additional copy of this manual and of the accompanying text, which explains some of the essential concepts of linear filtering.

2. THE MENUS OF THE PROGRAM

The following is a summary of the menus and sub-menus of the program. Most of the items of the main menu of *IDEOLOG.PAS* subsume a sub-menu. Secondary sub-menus are not listed:

IDEOLOG.PAS: Ideal Filters and their Approximations

1. *Page Parameters*
2. *Get the Data, Transform the Data*
3. *Plot the Series*
4. *Periodogram Analysis*
5. *Polynomial Regression*
6. *Variable Smoothing*
7. *FIR Time-Domain Filters*
8. *Wiener-Kolmogorov IIR Filters*
9. *Frequency-Domain Filters*
10. *Save and Exit*

The sub menus associated with the primary menus are as follows:

2. *GET THE DATA, TRANSFORM THE DATA*
 1. *Read the Data*
 2. *Plot the Data*
 3. *Transform the Data*
 4. *Subsample the Data*
 5. *Select or Discard Some Data*
 6. *Segment and Desegment the Data*
 7. *Data Extrapolation and Tapering*
 8. *Plot the Periodogram*
 9. *Return to the Main Menu*
3. *PLOT THE SERIES*
 - d. *Plot the Data*
 - t. *Plot the Trend*
 - b. *Plot the Business Cycle*
 - r. *Plot the Residuals*
 - e. *Return to the Main Menu*
4. *PERIODOGRAM ANALYSIS*
 - d. *Plot the Periodogram of the Data*
 - t. *Plot the the Periodogram of the Trend Component*
 - b. *Plot the Periodogram of the Business Cycle*
 - r. *Plot the the Periodogram of the Trend Residuals*
 - e. *Return to the Main Menu*
5. *POLYNOMIAL REGRESSION*
 1. *Fit by Ordinary Least-Squares Regression*
 2. *Fit by Weighted Least-Squares Regression*
6. *VARIABLE SMOOTHING*
 1. *Specify the Regular Smoothing Parameter*

2. *Specify the Regions of Smoothing Variation*
 3. *Specify the Local Smoothing Parameters*
 4. *Specify the Profile of the Smoothing Variation*
 4. *Smooth the Data*
 5. *Return to the Main Menu*
7. *FIR TIME-DOMAIN FILTERS*
1. *Binomial Lowpass Filter*
 2. *Henderson Lowpass Filter*
 3. *Baxter-King Bandpass Filter*
 4. *Christiano-Fitzgerald Bandpass Filter*
8. *WIENER-KOLMOGOROV IIR FILTERS*
1. *Hodrick-Prescott Filter*
 2. *Butterworth Filter*
 3. *Seasonal Adjustment Filter*
9. *FREQUENCY-DOMAIN FILTERS*
1. *Ideal Frequency Filters*
 2. *Butterworth Frequency Filter*
 3. *Hodrick-Prescott Frequency Filter*
 4. *Return to the Main Menu*
10. *SAVE AND EXIT*
1. *Save the Data*
 2. *Save the Trend Component*
 3. *Save the Business-Cycle Component*
 4. *Save the Trend Residuals*
 5. *Save the Seasonal Component*
 6. *RETURN to the Main Menu*
 7. *EXIT the Program*

3. DATA HANDLING

3.1 Data Entry

(IDEOLOG.PAS→Get the Data, Transform the Data→ Read the Data)

The program relies on data files in the .txt format in which the elements are written on successive lines.

There is an allowance for an accompanying integer index on the same line, which must precede the data value. The index must have a constant increment in passing from one line to the next. The program will check whether or not this is the case and, if it is not, then it will cease to read the data.

Seasonal and Quarterly Data

A further allowance is made for monthly or quarterly data with four or twelve elements per line, respectively, separated by spaces. In such cases, the program will seek conformation that the intention is to read across successive the rows of the data file and to join the elements in a single sequence.

The Headline of the Data File

The first two lines of the data file are permitted to be a description of the data in numerical and non-numerical text. If non-numerical text is found in subsequent lines, then the program will alert the user; and it will cease to read the data.

The Data Frequency

The user is asked to identify the frequency of the data by naming the interval between successive observations, which may be

- (a) *Annual,*
- (q) *Quarterly,*
- (m) *Monthly,*
- (n) *Other.*

The Name of the Data

The user may wish to identify the data by name. Otherwise, the data will be described as an unidentified data series. The naming of the data affects the legends that are displayed below the graphs that may be plotted on the screen. It also affects the tags that are applied to the PostScript code that can be generated while plotting graphs and diagrams on the screen.

3.2 Data Plotting

(IDEOLOG.PAS→ Get the Data, Transform the Data→ Plot the Data)

Before plotting the data, the user is asked to declare whether or not they wish to produce an accompanying PostScript file. If the answer is in the affirmative, then the choice can be made of producing either an *Encapsulated PostScript* code or a code in the *Textures format*.

Thereafter, it will be necessary to specify the page parameters of the graph, which is a matter of choosing the dimensions of the frame that surrounds the graph. See *Page Parameters* in section 5.

3.3 Discard Some Data:

(IDEOLOG.PAS→ Get the Data, Transform the Data→ Discard Some Data)

The program allows the user to discard some of the data that has been read. First, the data will be plotted to allow its inspection. Then, the following options are offered:

- (a) *Discard Data from the Beginning,*
- (b) *Discard Data from the End.*

The number of points to be discarded must be specified; whereafter the shortened data sequence will be plotted.

3.4 Data Transformations

(IDEOLOG.PAS→ Transform the Data→ Transform the Data)

Subtract the Mean from the Data

The program automatically determines whether the data have a significant trend and whether their mean value is significantly different from zero.

If there is a trend in the data, then it may be appropriate, at a later stage, to interpolate a polynomial trend function by a least-squares regression and to replace the data by the residual deviations, which will have a mean of zero. If there is no trend in the data, then it may be appropriate to subtract the mean.

A sequence with a zero mean will be more amenable to tapering the ends to zero than one with a nonzero mean. However, the program also offers a facility for tapering a data sequence to a nonzero mean asymptote, provided that it is free of any significant trend.

Take Differences

The program provides a facility for taking the first and second differences of the data. The differencing operator can be applied a second time to a data sequence to provide its second differences. In that case, the user will be alerted that the first differences have been taken already. Differences of a higher order are prevented.

Take Natural Logarithms

The program enables the user to take the natural logarithms (logarithms to the base e) of the data.

4. DATA EXTRAPOLATION AND TAPERING

4.1 Data Extrapolation and Tapering

(*IDEOLOG.PAS* → *Get the Data, Transform the Data* → *Extrapolate, Taper the Data*)

Some of the filtering procedures require sets of start-up values. To enable a time invariant Finite Impulse Response (FIR) filter of $2m + 1$ coefficients to produce filtered values for the entire data the sequence, m data points must be added at both ends.

In the case of a frequency-domain filtering, the assumption is made that the data sequence represents one cycle of a periodic or circular function. In order to avoid discontinuities in the periodic function, both ends of the data sequence must attain the same level. This can be achieved by tapering the data or by adding appropriate extrapolations to both ends.

Various ways of deriving the additional data points are available. The majority of the methods make the assumption that there is no trend in the data sequence and that its mean value is zero, or close to zero. Before any such method is applied, the program determines whether this is a valid assumption or not.

Reflecting the Data around the Endpoints

The data may be extended by adding the data sequence in reverse to the ends of the data. The endpoint of the original data is not repeated. Nor is the starting point of the data included in the reversed data sequence. The extended data form a circular sequence, in which the same values are to be found at opposite ends of any diameter of the circle.

Tapering the Data

If one is prepared to sacrifice the initial and the final handfuls of data points, then it may be appropriate to taper the ends of the data so that they reach the zero asymptote. The procedure subjects the sample to a tapering operation based on a split cosine bell. The latter is created by separating the two halves of a cosine

bell and by inserting a sequence of units. The operation leaves the bulk of the data unmodified and tapers only the ends. The user is asked to specify the fraction of the sample that is to be tapered.

Extrapolating and Tapering the Data

The data can be extended at both ends by a certain proportion of the original length, before tapering. The extrapolation is achieved by reflecting the sample data around the endpoints. A cosine decrement is applied to the extrapolated portions to reduce their extremities to zero.

Polynomial Extrapolation and Tapering

A polynomial is fitted to the data. It is extrapolated at both ends to a distance that is chosen as a proportion of the length of the original data. The extrapolated segments are modified by the addition of values obtained by reflecting the polynomial residuals around the end points of the sample and by tapering them to zero. The residual deviations from the polynomial trend are plotted, and the user is asked whether they wish them to replace the trended data sequence.

Extrapolation by Circular Morphing

The process of circular morphing works only with detrended data. An extension is devised that avoids any disjunctions when the data is wrapped around the circle. The extension, which is interpolated between the end of the final year of the data and the beginning of the first year, makes a gradual transition between the seasonal patterns of these two years.

5. GRAPH PLOTTING

5.1 Page Parameters

(IDEOLOG.PAS → *Page Parameters)*

The user must first declare whether or not they require a *PostScript* code to be generated to accompany the screen graphics. If the answer is in the affirmative, then the choice can be made of producing an *Encapsulated Postscript* (.eps) code or a code in the *Textures format*. (The latter relates to an implementation of the \TeX typesetting program that is no longer commercially available. It has been used in typesetting this manual.)

The graph is plotted in a frame of which the user must supply the dimensions. The available width of the frame is from a minimum of 5cm to a maximum of 13.5 cm and the available height is from 3 cm to 9 cm.

It is assumed that the graph will be placed on a printed page of A4 dimensions. In that case, the advice is given that, to print at most two diagrams per page, the maximum dimensions should be 9 cm width times 6 cm height. For three diagrams per page, the maximum dimensions should be 9 cm times 3.75 cm.

If a graph is to be plotted only on the screen, then the maximum available dimensions may be chosen. The maximum dimensions will be obtained by entering 99, or any other excessive number, in response to the request for the width and the height.

In the *Windows* version of the program, a plotted graph can be dismissed usually by typing a *RETURN*. However, if the graphics screen has been activated

by mouse click, then it will be necessary to dismiss the graph via the check box in the upper right corner of the screen.

5.2 Plot The Series

(IDEOLOG.PAS→ Plot the Series)

The program will plot the primary data sequence that is held in memory or of any other sequence that has been derived in the process of trend estimation or filtering. If the data sequence has been subject to a detrending operation and has not been replaced by the residual deviations from the trend, then the following operations will be available:

- (d) *Plot the Data*
- (t) *Plot the Trend*
- (b) *Plot the Business Cycle*
- (r) *Plot the Trend Residuals*

6. DATA CHARACTERISTICS

6.1 The Periodogram

(IDEOLOG.PAS→ Periodogram Analysis)

The program will plot the periodogram of the primary data sequence that is held in memory and of any other sequence that has been derived in the process of trend estimation or filtering. If the data have been subject to a detrending operation and have not been replaced by the residuals deviations from the trend, then the following operations will be available

- (d) *Plot the Periodogram of the Data*
- (t) *Plot the Periodogram of the Trend Component*
- (b) *Plot the Periodogram of the Business Cycle*
- (r) *Plot the Periodogram of the Trend Residuals*

Once the plotted graph has been dismissed, the user will be asked whether or not they wish to save the sequence of periodogram ordinates to a text file, which they will be asked to name.

In the case of the periodogram of the trend residuals, an option is available for imposing a highlighted band upon the graph, which will serve to emphasise a spectral structure of interest. The lower and the upper frequency limits of the band must be specified in degrees.

6.2 Data Segmentation

(IDEOLOG.PAS→ Get the Data, Transform the Data→ Data Segmentation)

A lengthy data sequence may have evolving characteristics that differ in different periods or epochs. The variability of these characteristics may be assessed by plotting segments of the data and by comparing their periodograms. Also, if the data are too closely clustered within a graph of the entire sequence, then it may be appropriate to expand the sampling intervals by plotting the graphs of a succession of data segments.

The program allows the temporary segmentation of the data to be achieved in either of two ways, by the choice of one of the following options:

- (a) *Specify the length of each segment*
- (b) *Specify the number of segments*

In either case, the segment must have at least 50 data points. After the segments of the desired number or of the desired length have been determined, the remainder of the data will be contained within the final segment.

The segmentation, which affects only the plotting of the data and of the associated periodograms, can be removed by segmenting the data anew and by specifying either that the length of the segment should be equal to the length of the data sequence or that the number of the segments should equal one.

7. TRENDS AND BASELINES

7.1 Polynomial Regression

(*IDEOLOG.PAS* → *Polynomial Regression*)

Ordinary Polynomial Regression

A numerically stable procedure is available for fitting a polynomial function to the data. The procedure avoids using the powers of the polynomial argument, which can give rise to numerical instability. Instead, it employs a sequence of orthogonal polynomials. Once the coefficients associated with the orthogonal polynomials have been determined, they are transformed into the power series coefficients.

Weighted Polynomial Regression

The objective of a weighted regression is to ensure that the fitted polynomial will pass through the midst of the scatters of points at either end of the data sequence. This is to ensure that there are no radical disjunctions in the periodic extension of the residual sequence where the end of one replication of the sequence joins the beginning of the next replication. To achieve this, the polynomial is fitted by a weighted least-squares regression, in which the maximum weight is given to the ends of the data sequence and the minimum weight is given to its mid point.

A Variety of Weighting Functions

The user is first asked to specify the maximum and the minimum weights—the latter being unity by default. They are also asked to specify a U-shaped profile according to which the weights make their transition from the minimum to the maximum values.

Two such choices are a quadratic function and a crenellated function that jumps abruptly from the minimum value, which is maintained over an interval centred on the mid point, to the maximum value, which, by default, is the value in the first and the last quarters of the sample.

There is also an option for splitting the quadratic by interpolating a constant segment into an interval centred on the mid point. The equivalent option in respect of the crenellated function is to allow variations in the length of the segment over which the weighting function assumes its minimum value.

7.2 Interrupted Trends

(*IDEOLOG.PAS* → *Variable Smoothing*)

The program implements a version of Hodrick–Prescott filter in which the value of the smoothing parameter can be varied in as many as three sub intervals of the period spanned by the data. The purpose of reducing the value of the smoothing parameter within a restricted interval is to allow the trend to undergo a change in slope or a rapid shift in level of a sort that corresponds to a structural break in the processes generating the data.

8. FIR FILTERS

8.1 The Binomial Lowpass Filter

(*IDEOLOG.PAS* → *FIR FILTERS* → *Binomial Lowpass Filter*)

The ordinates of the binomial distribution $b(n, p = q = 0.5)$ constitute a lowpass FIR filter. The profile of the binomial distribution approximates that of a Gaussian normal distribution. The frequency response function of the filter also has approximately a normal profile.

The filter has $2m + 1$ coefficients disposed symmetrically around the central element. The greater is the number of binomial coefficients, the less is the dispersion of the frequency response function and lower is the frequency value of the mid point of the transition from passband to stopband.

The product of the dispersion of the normal distribution with that of its Fourier transform (i.e. its characteristic function) attains the Heisenberg minimum limit. Therefore, the binomial filter provides a near-ideal trade-off between the resolutions in the time and frequency domains.

The Frequency Response of the Filter

The frequency response function of the filter shows the positive factors by which the Fourier coefficients of the data are multiplied by the lowpass filter.

The Impulse Response of the Filter

The impulse response of the filter is the sequence of the filter coefficients. A symmetric sequence of filter coefficients is plotted, which would be appropriate to an infinite data sequence, unaffected by any adaptation to cope with the circumstances when filter reaches the ends of the data sequence. The user is asked to specify a data span, within which the coefficients are plotted in a central location. The data span cannot be less than the span of the filter.

The Application of the Filter

When the filter is applied to the data, graphs are generated that show the trend estimate interpolated through the data and the deviations of the data from the trend. The data, the trend component and the deviations from the trend can be saved in separate text files. (See Save and Exit)

8.2 The Henderson Filters

(*IDEOLOG.PAS* → *FIR FILTERS* → *Henderson Lowpass Filter*)

The Henderson moving-average filter is employed in the Census Bureau X-11 program to estimate trends. It can be applied directly to the data to estimate the trend component. A set of supplementary filter weights, which are the so-called

Musgrave weights, are used at the ends of the sample, where there are insufficient data points to apply the normal symmetric filter.

The user is asked to specify the order of the Henderson filter, which is one of the following numbers: 5, 7, 13 or 23.

The Frequency Response of the Filter

The frequency response function of the filter shows the positive or negative multiplicative factors by which the filter alters the cyclical components of a stationary process. The amplitude response, which is the modulus of the frequency response, shows the extent to which the cyclical components are amplified or attenuated, which is described as the gain of the filter. Both the frequency response and the amplitude response can be plotted.

The amplitude response is unity at the zero frequency and it declines with increasing frequency. This corresponds to the fact that the filter will transmit the trend component of the data. By increasing the order of the filter the point of zero gain is brought closer to the zero frequency.

The Impulse Response of the Filter

The impulse response of the filter is the sequence of the filter coefficients. A symmetric sequence of filter coefficients is plotted, which would be appropriate to an infinite data sequence, unaffected by any end effects.

The Application of the Filter

When the filter is applied to the data, graphs are generated that show the trend estimate interpolated through the data and the deviations of the data from the trend. The original data, the trend component and the deviations from the trend can be saved in separate text files. (See Save and Exit)

8.3 The Baxter–King Bandpass Filter

(IDEOLOG.PAS → FIR FILTERS → The Baxter-King Filter)

The bandpass filter of Baxter and King (1999) is obtained by truncating the sequence of the coefficients of an ideal bandpass filter. The user must specify, in degrees, a lower cut-off point and an upper cut-off point. The frequency range of the data, which is from zero to π radians per sampling interval, corresponds to a range from 0 to 180 degrees.

The filter was designed with a view to implementing a definition of the business cycle that is due to Burns and Mitchell (1946). They proposed that it should comprise cyclical components of a duration not exceeding 8 years and of no less than one-and-a-half years. For quarterly data, the period of eight years corresponds to a lower cut off point of 11.5 degrees and the period of one-and-a-half years corresponds to an upper cut-off point of 60 degrees.

8.4 The Christiano-Fitzgerald Filter

(IDEOLOG.PAS → FIR FILTERS → The Christiano-Fitzgerald Filter)

Christiano and Fitzgerald (2003) have proposed to address the end-effects problem affecting the bandpass filter of Baxter and King by a method that amounts to extending the data sequence at the levels that have been reached at either end.

This procedure would be appropriate to data that have been generated by a non stationary process incorporating a single unit root in an autoregressive operator, which gives rise to a first-order random walk. In the case of a mean-zero stationary process, it would be appropriate to extend the data by zero-valued ordinates. In the case where the data embody a drifting random walk, it is appropriate to replace the data with its deviations from a linear trend.

The user is asked to declare which of the following circumstances apply:

- (a) *The sample is from a stationary stochastic process,*
- (b) *The sample is from a process with a single unit root,*
- (c) *The sample is from a process with a drifting random walk.*

9. WIENER–KOLMOGOROV IIR FILTERS

The basic Wiener–Kolmogorov filters, which are the Butterworth filters and the Hodrick–Prescott filters, come in complementary pairs, which comprise a lowpass filter and a highpass filter. Together, the coefficients of the lowpass filter and the highpass filter sum to unity.

That is to say, that the residue from applying the lowpass filter is the sequence of values that would be generated by the highpass filter. Only the lowpass filters are implemented in the program, which also generates the residual sequences.

A bandpass filtering with an upper and a lower cut-off point can be implemented via a succession of lowpass filterings. First, a lowpass filter with the upper cut-off is applied to the data. Then, the resulting sequence is subjected to a second lowpass filter with the lower cutoff. The residue of the second operation corresponds to the result from applying a bandpass filter.

The greater flexibility of the Butterworth filter regarding the choice of a cut-off frequency and the availability of more rapid transitions between the stop band and the pass band makes it more suitable to bandpass filtering than is the Hodrick–Prescott filter.

9.1 The Hodrick–Prescott Filter

(IDEOLOG.PAS → IIR FILTERS → Hodrick–Prescott Filter)

The lowpass filter of Hodrick and Prescott (1980, 1997), which is also attributable to Leser (1961), has a single adjustable parameter λ (lambda), which governs the trade-off between the smoothness of the filtered sequence and the closeness of its fit to the data sequence. The following are the conventional values of the smoothing parameter:

$$\lambda = \begin{cases} 100 & \text{for annual data,} \\ 1,600 & \text{for quarterly data,} \\ 14,400 & \text{for monthly data.} \end{cases}$$

The lowpass Hodrick–Prescott filter is commonly used for extracting the trends from econometric data sequences. It is appropriate to cases where a two-fold differencing of the data would be effective in eliminating the trend. It will transmit a linear function without alteration and it is appropriate to the filtering of a process consisting of an integrated random walk and an added white-noise disturbance.

The Frequency Response of the Filter

The frequency response of the filter shows a gradual transition from unity at the zero frequency to zero at the highest frequency. The greater is the value of λ , the more rapid is the transition and the smoother is the filtered sequence.

The Pole–Zero Diagram

The lowpass Hodrick–Prescott filter has a pair of zeros at the Nyquist frequency of π and a pair of poles (or autoregressive roots) of which the locations are determined by the value of the smoothing parameter

The Impulse Response of the Filter

The impulse response of the filter is the sequence of filter coefficients. In the case of a finite-sample Wiener–Kolmogorov filter such as the Hodrick–Prescott filter, the coefficients vary according to the length of the data sequence and according to the position of the filtered elements within the sequence. Before plotting the impulse response, the user is asked to declare the length data sequence, i.e. the span of the impulse response, and the location of the impulse within this span,

The Application of the Filter

When the filter is applied to a trended data sequence, the program will display a graph of the data with an interpolated smooth trajectory, described as the trend-cycle component. When this graph has been dismissed, the program will display a graph of the residual deviations of the data from the trend-cycle component.

If the original data have been replaced by their deviations from an interpolated polynomial trend, then the program will begin by showing these deviations together with an interpolated smooth low-frequency component, which is described as the cyclical component. Next, it will show the deviations of the polynomial residuals from the cyclical component. Thereafter, if the user wishes, it will add the cyclical component to the polynomial trend to create the trend-cycle component.

9.2 The Butterworth Filter

(*IDEOLOG.PAS* → *IIR FILTERS* → *Butterworth Filter*)

The lowpass Butterworth filter is capable of a rapid transition from a low-frequency pass band to a high frequency stop band. It is specified by declaring the cut-off frequency, which marks the midpoint of the transition and the filter, and the filter order, which governs the rate of transition. The filter is capable of isolating a low-frequency spectral structure with a well-defined upper limit in frequency.

The Frequency Response

The frequency response of the filter is unity at zero frequency and zero at the Nyquist frequency of π radians per sample interval. It displays maximum flatness at these points. The rate of transition in the vicinity of the nominal cut-off increases with the order of the filter.

The Pole-Zero Diagram

The pole-zero diagram shows a multiple zero at the Nyquist frequency of π radians per sampling interval and a cluster of poles on an arc that migrates towards the vertical axis as the cut-off frequency tends to $\pi/2$ radians or 90 degrees.

The Impulse Response

In the case of the finite-sample Butterworth filter, the response to an impulse at time t differs from the sequence of coefficients involved in generating the output. The former corresponds to the t -th column of the filter matrix that transforms the data vector to the vector of the filtered output, and the latter corresponds to the t -th row. The user is asked to choose whether to plot

- (1) *the impulse response,*
- (2) *the coefficient sequence.*

Before plotting the impulse response or the coefficient sequence, the user is asked to declare the length data sequence, i.e. the span of the impulse response, and the location of the impulse within this span.

The Application of the Filter

When the filter is applied to a trended data sequence, the program will display a graph of the data with an interpolated smooth trajectory, described as the trend-cycle component. When this graph has been dismissed, the program will display a graph of the residual deviations of the data from the trend-cycle component.

If the original data have been replaced by their deviations from an interpolated polynomial trend, then the program will begin by showing these deviations together with an interpolated smooth low-frequency component, which is described as the cyclical component. Next, it will show the deviations of the polynomial residuals from the cyclical component. Thereafter, if the user wishes, it will add the cyclical component to the polynomial trend to create the trend-cycle component.

9.3 A Seasonal-Adjustment Filter

(IDEOLOG.PAS → IIR FILTERS → Seasonal-Adjustment Filter)

A facility is provided for the seasonal adjustment of monthly or quarterly data showing a pattern of seasonal variation. A so-called comb filter is provided, which has a frequency response with notches that extend to zero at the seasonal frequencies. These notches serve to eliminate the elements of the data at the seasonal frequencies and to attenuate the elements at the adjacent frequencies.

The width of the notches is governed by a parameter $\lambda \in (0, 1)$, and they become narrower as $\lambda \rightarrow 0$. A second parameter $\rho \in (0, 1)$ must also be specified, which determines the modulus of the poles of the filter. Values are recommended for both parameters.

In the case of a trended data sequence, the user is constrained to choose one or two differencing operations to be applied to the data prior to filtering. In that case, the filtered sequence will be reinflated by the corresponding number of summation operations.

The Frequency Response

The frequency response of the filter is unity at zero frequency and close to unity at frequencies that lie at the midpoints between adjacent seasonal frequencies. In the case of quarterly data, the seasonal frequencies are at intervals of $\pi/4$ radians or 45 degrees at $\pi/4$, $\pi/2$, $3\pi/4$ and π (45 degrees, 90 degrees, 135 degrees and 180 degrees). In the case of monthly data, the seasonal frequencies are at intervals of $\pi/6$ radians or 30 degrees.

The Pole-Zero Diagram

The pole-zero diagram shows a set of zeros located on the unit circle at angles corresponding to the seasonal frequencies and poles that lie close to the circle at the same angles. The zeros serve to nullify the elements of the data at the seasonal frequencies. The poles ensure that the effects of the zeros are confined to the vicinities of the seasonal frequencies.

The Application of the Filter

The program produces a graph that shows the seasonally adjusted data interpolated through the original data sequence. If the data are deviations from an interpolated polynomial trend, then an option is available for restoring the trend to the data and for plotting a corresponding graph.

10. FREQUENCY-DOMAIN FILTERS

The frequency domain filters operate by multiplying the ordinates of the Fourier transform of the data by weights that are derived from the frequency response of the filter. The weighted ordinates are translated back to the time domain by an inverse Fourier transform to create the filtered sequence.

The data that are subject to a Fourier transform constitute a perpetual periodic or circular sequence. Any trends that are present in the data must first be eliminated. A common way of extracting a trend is to interpolate a trend function and to use the deviations from the trend in place of the original data. If required, the low pass filtered sequence can be added back to the trend function to create a trend-cycle component

Another way to eliminate a trend is to rely on differencing operations. A twofold differencing operator is usually employed, which may be either an ordinary backwards twofold differencing operator or a centralised twofold differencing operator. After filtering, the differenced data will be re-inflated to create a trend-cycle component and a residual component, of which the sum will be equal to the original data sequence.

If the centralised operator has been applied to the data, then the frequency response of the inverse operator will be compounded with the frequency response of the high pass operator. This combined operator is applied to the Fourier ordinates of the data. When the weighted ordinates are translated back to the time domain, the result will be a re-inflated version on the high pass component.

If the backwards difference operator is applied to the data, then they will be subjected to a low pass frequency-domain filter. The result will be carried back to the time domain, where it will be re-inflated by a process of accumulation, which requires the provision of a pair of starting values. These starting values, which impart a linear trend to the accumulated data, are determined so as to minimise the sum of squares of the differences between the points of the resulting trend-cycle component and those of the data sequence.

10.1 The Ideal Lowpass and Bandpass Filters

(IDEOLOG.PAS → Frequency-Domain Filters → Ideal Frequency Filters)

If the data are in the form of deviations from a trend, then an option is offered to add extrapolations to the data based on the reflections of a number of data points from the beginning and the end of the data sequence around their respective

endpoints. These extrapolations will be tapered by a cosine decrement so as to reach the horizontal axis.

The data will be subjected to an ideal lowpass frequency filter, which will preserve the Fourier ordinates associated with frequencies up to and including the chosen cut-off frequency and which will nullify the ordinates that lie beyond it. The complementary highpass component is formed by subtracting the low pass component from the data sequence.

In the case of de-trended data, the option is offered for determining the turning points of the lowpass component, which might represent the business cycle in an index of economic activity. The turning points are the zeros of the first derivative of the low pass trajectory, which is created by a process of Fourier synthesis.

If a trend is detected, then the data will be subject to the centralised twofold differencing operator. Thereafter, the high pass component of the data is isolated by applying a filter that is compounded from the frequency response of the high pass filter and that of the inverse of the difference operator. This will give rise to the re-inflated version of the high pass component

The ideal bandpass filter requires the specification of an upper and a lower cut-off frequency. Both points are specified as a number of degrees in the interval $[0, 90)$.

10.2 Seasonal Adjustment with the Bandstop Filter

(IDEOLOG.PAS → Frequency-Domain Filters → Ideal Frequency Filters)

The de-seasonalising filter is designed to nullify the Fourier ordinates of the seasonal fluctuations within an economic data sequence. If the data are observed at quarterly intervals, then the fundamental seasonal frequency is at 90 degrees and the harmonic frequency is at 180 degrees. If the data are observed at monthly intervals, then the fundamental seasonal frequency is at 30 degrees and the harmonic frequencies, which are integer multiples of the fundamental frequency, are at 60, 90, 120, 150 and 180 degrees.

The seasonal component may comprise frequencies that are adjacent to the fundamental frequency and its harmonics. Their presence will be revealed by the periodogram of the de-trended data, created by extracting a trend function from the original data. (The need to inspect the periodogram virtually compels the user to consider the deviations of the data from an interpolated trend.) The program allows adjacent frequencies to be included in the definition of the seasonal component.

In case the data contain a trend, the following options are provided by the program

- (1) *Obtain the residuals of a polynomial regression*
- (2) *Examine the periodogram of the residuals*
- (3) *Take differences of the data*
- (4) *Continue*

If the data represent the residuals from an interpolated trend function, then only the option under 2. is available

10.4 The Butterworth Filter in the Frequency Domain

(IDEOLOG.PAS → Frequency-Domain Filters → Butterworth Frequency Filter)

The Butterworth filter is realised in the frequency domain by multiplying each of the Fourier ordinates of the data by the value of the relevant Butterworth frequency response function at the corresponding frequency.

If the data are free of trend, then the frequency response is that of a low pass filter specified by its filter order and its cut-off, frequency. Whereas the order of the time-domain filter must be a positive integer, that of the frequency-domain filter can be an arbitrary positive real number.

The option is available for extrapolating the ends of the data by reflecting the initial and the final segments of the data around the end points. A cosine decrement is applied to the extra-sample elements to reduce their values to zeros at the furthest points of the extrapolations.

Observe that the extrapolation may occur either within this routine or outside of it, via the *DATA EXTRAPOLATION AND TAPERING* routine, described in section 4. The absence of a trend in the data is liable to be the result of the prior interpolation of a trend function and the replacement of the data by their deviations from the trend.

If there is a trend in the data, then this may be eliminated by applying the centralised twofold difference operator. Then, a filter is applied to the Fourier ordinates of the differenced data that is formed by dividing the frequency response of the high pass filter by that of the frequency response of the difference operator, which reverses the effect of the time-domain difference operator. The resulting high pass component is subtracted from the data sequence to create the low pass component.

The option is available for extrapolating the ends of the data. This involves fitting a polynomial function to the data of which the extrapolations provide a baseline to which extrapolated residuals are added, which are obtained by reflecting the initial and the final segments of the residual sequence around the end points. A cosine decrement is applied to the extrapolated residuals so that they reach the baseline at the furthest points of the extrapolations.

After filtering, the extrapolations are removed from the data and from the filtered sequences.

10.5 The Hodrick–Prescott Filter in the Frequency Domain

(IDEOLOG.PAS → Frequency-Domain Filters → Hodrick–Prescott Frequency Filter)

The Hodrick–Prescott filter is realised in the frequency domain by multiplying each of the Fourier ordinates of the data by the value of the Hodrick–Prescott frequency response function at the corresponding frequency. The filter is specified by the value of the smoothing parameter, of which the conventional values have been specified already in connection with the time-domain filter in section 9.1.

If the data are free of trend, then the frequency response is that of a low pass filter. The absence of a trend in the data is liable to be the result of the prior interpolation of a trend function and the replacement of the data by their deviations from the trend.

The option is available for extrapolating the ends of the data by reflecting the initial and the final segments of the data around the end points. A cosine decrement is applied to the extra-sample elements to reduce their values to zeros at the furthest points of the extrapolations. These operation may occur either within the routine

or as a result of prior operations via the via the *DATA EXTRAPOLATION AND TAPERING* routine, described in section 4.

If there is a trend in the data, then this will be eliminated by applying the centralised twofold difference operator. Then, a filter is applied to the Fourier ordinates of the differenced data that is formed by dividing the frequency response of the high pass filter by that of the frequency response of the difference operator, which effectively reverses the effect of the time-domain difference operator. The resulting high pass component is subtracted from the data sequence to create the low pass component.

The option is available for extrapolating the ends of the data. This involves fitting a polynomial function to the data of which the extrapolations provide a baseline to which extrapolated residuals are added, which are obtained by reflecting the initial and the final segments of the residual sequence around the end points. A cosine decrement is applied to the extrapolated residuals so that they reach the baseline at the furthest points of the extrapolations.

After filtering, the extrapolations are removed from the data and from the filtered sequences.

11. CASE STUDIES

The user can proceed through the program by selecting the menu items that are marked by numbers or letters, by responding to requests of the program for information or instructions—also by typing numbers or letters—and by typing a `<RETURN>` at points where the execution of the program has halted. The need to type a `<RETURN>` is usually implicit, albeit that it is occasionally indicated by a message on the screen.

The choice to generate the *PostScript* code of the graphics, which is made via the *Page Parameters* menu, has a global effect. To avoid generating the code in a particular instance, one can type a `<RETURN>` in response to a request for the name of the *PostScript* file.

In the listings that follow, which record the steps to be taken in pursuit of specific outputs, the program's requests for information and instructions are displayed in Roman type. Italic type is reserved for non-executable items, which are the information and advice supplied by the program, some of which is omitted from the listings, the headline titles of the program's menus and the captions of the graphics.

12.1 The Consumption Cycle and its Turning Points

The data for this study are quarterly observations on consumers' expenditure in the U.K. in 1990 prices (not seasonally adjusted) running from 1955Q1 to 1994Q4. The purpose is to extract a trend-cycle component from the data.

There are three alternative pathways to the same end result, which are TRACKS A, B and C. The diversions occur after the data has been read and its logarithms have been taken. If track A is pursued, then a facility will be available for marking the turning points in the cycle of consumption.

It appears that, in the U.K., the consumption cycle precedes the main business cycle, for which it provides some of the impetus.

D.S.G. POLLOCK: IDEOLOG—A FILTERING PROGRAM

IDEOLOG.PAS: Ideal Filters and their Approximations

1. Page Parameters
→ 1

SPECIFY THE PAGE PARAMETERS

Do you want PostScript? Y/N
→ N

Specify the frame surrounding the graph. $5\text{cm} \times 3\text{cm} \leq \text{width} \times \text{height} \leq 13.5\text{cm} \times 9.5\text{cm}$.

For two diagrams per page use $9\text{cm} \times 6\text{cm}$. For three diagrams per page use $9\text{cm} \times 3.75\text{cm}$

Specify the width
→ 99 {for maximum width}

Specify the height
→ 99 {for maximum height}

2. Get the Data and/or Transform the Data
→ 2

GET THE DATA

1. Read the Data
→ 1

READ A DATA FILE

Name the Data File =
→ cons.txt

There are 160 data points. What is the interval between observations?

(q) Quarterly
→ q

NAME THE DATA

By default the data will be described as <<an unidentified data series>>

Do you wish to rename the data? Y/N
→ N

GET THE DATA

3. Transform the Data
→ 3

TRANSFORM THE DATA

3. Take Natural Logarithms
→ 3

10. Return to the Main Menu
→ 10

TRACK A. This method reduces the data to stationarity by taking residual deviations of the data from a polynomial trend interpolated by least-squares regression. The linear trend that is evident in the logarithms of the consumption data, is eliminated by twofold differencing. The residual sequence is subjected to a lowpass filter to extract the cyclical component and a facility will be available for marking the turning points in the cycle. To generate the trend-cycle component, the cyclical component is added to the polynomial trend.

IDEOLOG.PAS: Ideal Filters and their Approximations

5. Polynomial Regression

→ 5

FIT A POLYNOMIAL TREND

1. Fit by Ordinary Least-Squares Regression

→ 1

What is the degree of the polynomial?

→ 1

<<Figure. A polynomial of degree 1 interpolated through 160 points of an unidentified data series>>

→ <RETURN>

<<Figure. The residual sequence from fitting a polynomial of degree 1 to 160 points of an unidentified data series>>

→ <RETURN>

FIT A POLYNOMIAL TREND

Do you wish to replace the data series by the residual series Y/N

→ Y

IDEOLOG.PAS: Ideal Filters and their Approximations

9. Frequency-Domain Filters

→ 9

SPECIFY AND APPLY A FREQUENCY-DOMAIN FILTER

1. Ideal Frequency Filters

SPECIFY AND APPLY AN IDEAL FREQUENCY-DOMAIN FILTER

1. Lowpass Filter

→ 1

SPECIFY THE IDEAL FREQUENCY-DOMAIN LOWPASS FILTER

The data consist of the residual deviations from an interpolated trend. After the low-frequency cycle has been extracted from the data, the option provided for identifying its turning points by locating the points where the derivative function cuts the horizontal axis

Do you wish to investigate the turning points? Y/N

→ Y

Specify the cut-off point in degrees

→ 22.5

ENSURE INTEGRAL CYCLES

Do you wish ensure that the data comprise an integral number of seasonal cycles? Y/N

→ N

⟨⟨Figure. A low-frequency trend/cycle component extracted from 160 points of an unidentified data series by a Fourier method⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. The residual sequence from detrending 160 points of an unidentified data sequence via the Fourier Method⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. The turning points of the low-frequency cycle marked by the vertical lines at the points where the derivative function cuts the horizontal axis⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. The residual sequence from detrending 160 points of an unidentified data sequence via the Fourier Method⟩⟩

→ ⟨RETURN⟩

THE TREND CYCLE COMPONENT

The trend-cycle component can be derived by adding the filtered version of the residual sequence, which has been subjected to a lowpass filter, to the polynomial function interpolated through the data.

Do you wish to plot the Trend/Cycle component? Y/N

→ Y

⟨⟨Figure. The trend-cycle component derived by adding to the interpolated polynomial the low-frequency components of the residual sequence⟩⟩

→ ⟨RETURN⟩

TRACK B. This method reduces the data to stationarity via the difference operator. Then, the differenced data are translated to the frequency domain, where the low-frequency ordinates are selected and the high-frequency ordinates are rejected. Within the frequency domain, the selected ordinates are re-scaled using the values of the ordinates of the frequency response function of the summation operator. Then, the re-scaled ordinates are translated to the time domain and are used for synthesising the trend/cycle sequence.

IDEOLOG.PAS: Ideal Filters and their Approximations

9. Frequency-Domain Filters

→ 9

SPECIFY AND APPLY A FREQUENCY-DOMAIN FILTER

1. Ideal Frequency Filters
→ 1

SPECIFY AND APPLY AN IDEAL FREQUENCY-DOMAIN FILTER

1. Lowpass Filter
→ 1

SPECIFY THE FREQUENCY-DOMAIN LOWPASS FILTER

- B. *A lowpass filter can be realised by subtracting a highpass component from the data. The differencing factor that is contained within the highpass filter can be cancelled with the summation operator to avoid the need for reinflating the filtered sequence.*

→ B

Specify the cut-off point in degrees

→ 22.5

The cut-off demarcates cycles of duration larger than 16 quarters

EXTRACT LOWPASS TREND COMPONENT

The trend comprises a set of sinusoidal components of frequencies that range from zero to the Fourier frequency nearest to a specified cut-off point

The cut-off demarcates cycles of duration larger than 16 quarters

TREAT THE END-OF-SAMPLE PROBLEM

Do you wish to extrapolate the data? Y/N

→ N

EXTRACT LOWPASS TREND COMPONENT

Before estimating the trend, you may require to reduce the data to stationarity a differencing operation

Do you wish to take centralised second differences? Y/N

→ Y

⟨⟨Figure. A trend comprising Fourier components of frequencies less than 22.5 degrees per period interpolated through 60 points of an unidentified data series⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. The residual sequence from detrending 160 points of an unidentified data sequence via the Fourier Method⟩⟩

→ ⟨RETURN⟩

TRACK C. In this method, the data are reduced to stationarity by the difference operator. Then, the differenced data are subjected to a lowpass frequency-domain filter. Thereafter, the filtered sequence is re-inflated via the summation operator, given the appropriate initial conditions.

IDEOLOG.PAS: Ideal Filters and their Approximations

9. Frequency-Domain Filters
→ 9

SPECIFY AND APPLY A FREQUENCY-DOMAIN FILTER

1. Ideal Frequency Filters

SPECIFY AND APPLY AN IDEAL FREQUENCY-DOMAIN FILTER

1. Lowpass Filter
→ 1

SPECIFY THE FREQUENCY-DOMAIN LOWPASS FILTER

- A. *The frequency filter may be applied to the differenced data and, thereafter, the filtered sequence can be reinflated via a summation operation.*
→ A

Specify the cut-off point in degrees
→ 22.5

TREAT THE END-OF-SAMPLE PROBLEM

Do you wish to extrapolate the data? Y/N
→ N

EXTRACT LOWPASS TREND COMPONENT

Before estimating the trend, you may require to reduce the data to stationarity by at most two differencing operations.

Specify the number of differencing operations
→ 2

ENSURE INTEGRAL CYCLES

Do you wish ensure that the differenced data comprise an integral number of seasonal cycles? Y/N
→ Y

<<Figure. A low-frequency trend/cycle component extracted from 158 points of an unnamed data sequence by a Fourier Method>>
→ <RETURN>

<<Figure. The residual sequence from detrending 158 points of an unnamed data sequence via the Fourier Method>>
→ <RETURN>

12.2 The Long-run Trend of U.K. GDP and its Structural Breaks

The data are the logarithms of 129 annual observations on U.K. GDP at constant prices, from 1873 to 2001. The object is to fit a flexible trend that will accommodate the structural breaks that followed two world wars. It should also be possible to accommodate the structural break occasioned by the financial crisis of 1929.

D.S.G. POLLOCK: IDEOLOG—A FILTERING PROGRAM

IDEOLOG.PAS: Ideal Filters and their Approximations

1. Page Parameters
→ 1

SPECIFY THE PAGE PARAMETERS

Do you want PostScript? Y/N
→ N

Specify the frame surrounding the graph. $5\text{cm} \times 3\text{cm} \leq \text{width} \times \text{height} \leq 13.5\text{cm} \times 9.5\text{cm}$.

For two diagrams per page use $9\text{cm} \times 6\text{cm}$. For three diagrams per page use $9\text{cm} \times 3.75\text{cm}$

Specify the width
→ 99 {for maximum width}

Specify the height
→ 99 {for maximum height}

2. Get the Data and/or Transform the Data
→ 2

GET THE DATA

1. Read the Data
→ 1

READ A DATA FILE

Name the Data File =
→ hendata.txt

Do you wish to construct a univariate series by joining the lines? Y/N
→ Y

READ A DATA FILE

There are 129 data points What is the interval between observations?

- (a) Annual
→ a

NAME THE DATA

By default the data will be described as <<an unidentified data series>>

Do you wish to rename the data Y/N?
→ N

10. Return to the Main Menu
→ 10

IDEOLOG.PAS: Ideal Filters and their Approximations

6. Variable Smoothing
→ 6

SMOOTHING WITH LOCAL VARIATIONS

1. Specify the Regular Smoothing Parameter
→ 1

SPECIFY THE FREQUENCY REGULAR SMOOTHING PARAMETER

What is the value of the smoothing parameter?
→ 16000

SMOOTHING WITH LOCAL VARIATIONS

2. Specify the Regions of Smoothing Variation
→ 2

The user will be asked to specify no more than three sub intervals within the period spanned by the data, wherein the values of the smoothing parameter can be reduced in order to enhance the flexibility of the trend so as to allow it to accommodate structural breaks.

SPECIFY THE FIRST REGION OF ADAPTED SMOOTHING

Specify the data central point
→ 46

Specify the number of points above and below the central point
→ 5

Do you accept this region Y/N?
→ Y

SPECIFY THE SECOND REGION OF ADAPTED SMOOTHING

Specify the data central point
→ 72

Specify the number of points above and below the central point
→ 5

Do you accept this region Y/N?
→ Y

Do you wish to specify additional regions Y/N?
→ N

SMOOTHING WITH LOCAL VARIATIONS

3. Specify the Local Smoothing Parameters
→ 3

2. Specify the Smoothing Parameters
→ 2

Specify the value of the smoothing parameters in the FIRST subinterval
→ 5

Specify the value of the smoothing parameters in the SECOND subinterval
→ 5

SMOOTHING WITH LOCAL VARIATIONS

4. Specify the Profile of the Smoothing Variation
→ 4

SPECIFY THE PROFILE OF THE SMOOTHING PARAMETER

2. Square Trough
→ 2

SMOOTHING WITH LOCAL VARIATIONS

5. Smooth the Data
→ 5

⟨⟨Figure. A trend determined by a filter with a variable smoothing parameter applied to 129 observations of an unidentified data series ⟩⟩
→ ⟨RETURN⟩

⟨⟨Figure. The residual sequence from fitting a trend via a filter with a variable smoothing parameter to 129 observations of an unidentified data series⟩⟩
→ ⟨RETURN⟩

12.3 The Trend-cycle Component of Monthly Data on U.S. Retail Sales

The data for this study are 144 monthly observations on the total retail sales in the US in millions of dollars, running from January 1953 to December 1964.

The object is to fit a trend to the data and to extract a seasonal component from the data, which will characterise the annual variation in the sales.

IDEOLOG.PAS: Ideal Filters and their Approximations

1. Page Parameters
→ 1

SPECIFY THE PAGE PARAMETERS

Do you want PostScript? Y/N
→ N

Specify the frame surrounding the graph. $5\text{cm} \times 3\text{cm} \leq \text{width} \times \text{height} \leq 13.5\text{cm} \times 9.5\text{cm}$.

For two diagrams per page use $9\text{cm} \times 6\text{cm}$. For three diagrams per page use $9\text{cm} \times 3.75\text{cm}$

Specify the width
→ 99 {for maximum width}

Specify the height
→ 99 {for maximum height}

2. Get the Data and/or Transform the Data
→ 2

GET THE DATA

2. Get the Data and/or Transform the Data

→ 2

1. Read the Data

→ 1

Name the new data file =

→ salesus.txt

Do you wish to construct a univariate series By joining the lines? Y/N

→ Y

READ A DATA FILE

There are 144 data points What is the interval between observations?

(m) Monthly

→ m

NAME THE DATA

By default the data will be described as <<an unidentified data series>>

Do you wish to rename the data Y/N?

→ N

GET THE DATA

3. Transform the Data

→ 3

TRANSFORM THE DATA

3. Take Natural Logarithms

→ 3

10. Return to the Main Menu

→ 10

IDEOLOG.PAS: Ideal Filters and their Approximations

9. Frequency-Domain Filters

→ 9

SPECIFY AND APPLY A FREQUENCY-DOMAIN FILTER

1. Ideal Frequency Filters

SPECIFY AND APPLY AN IDEAL FREQUENCY-DOMAIN FILTER

1. Lowpass Filter

→ 1

SPECIFY THE FREQUENCY DOMAIN LOWPASS FILTER

A. *The frequency filter may be applied to the differenced data and, thereafter, the filtered sequence can be re-inflated via a summation operation.*

D.S.G. POLLOCK: IDEOLOG—A FILTERING PROGRAM

→ A

Specify the cut-off point in degrees

→ 25

TREAT THE END-OF-SAMPLE PROBLEM

Do you wish to extrapolate the data? Y/N

→ Y

POLYNOMIAL EXTRAPOLATION AND TAPERING

What proportion of the original data in each extensions? Specify the proportion as a percentage

→ 25

FIT A POLYNOMIAL TREND

2. Fit by Weighted Least-Squares Regression

→ 2

SPECIFY THE WEIGHTING FUNCTION

Do you accept that the minimum weight should be unity Y/N?

→ Y

Specify the maximum value for the weights

→ 8

SPECIFY THE WEIGHTING FUNCTION

3. Cosine-Bell Weighting Function

→ 3

COSINE-BELL WEIGHTING FUNCTION

Do you wish to interpolate a middle segment Y/N?

→ Y

Specify the relative length of the middle range as a percentage

→ 75

⟨⟨Figure. The profile of the weights applied to the data points in the polynomial regression⟩⟩

→ ⟨RETURN⟩

What is the degree of the polynomial?

→ 1

⟨⟨Figure. The plot of 216 points of an unidentified data series⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. A polynomial of degree 1 interpolated by weighted least squares⟩⟩

→ ⟨RETURN⟩

EXTRACT LOWPASS TREND COMPONENT

Before estimating the trend, you may require to reduce the data to stationarity by at most two differencing operations.

D.S.G. POLLOCK: IDEOLOG—A FILTERING PROGRAM

Specify the number of differencing operations

→ 2

DISCARD THE EXTRAPOLATIONS

The data have been supplemented by 36 points at each end Do you wish to discard the extrapolated data points? Y/N

→ Y

⟨⟨Figure. A low-frequency trend/cycle component extracted from 144 points of an unidentified data sequence by a Fourier Method⟩⟩

→ ⟨RETURN⟩

⟨⟨Figure. The residual sequence from detrending 144 points of an unidentified data series via the Fourier Method⟩⟩

→ ⟨RETURN⟩

Now we may proceed to extract the seasonal component from the logarithmic data sequence

SPECIFY AND APPLY THE FREQUENCY DOMAIN FILTER

3. De-seasonalising BandStop Filter

→ 3

REMOVE SEASONAL FLUCTUATIONS FROM THE DATA

3. Take Differences of the Data

→ 3

Specify the order of the differencing operator

→ 2

REMOVE SEASONAL FLUCTUATIONS FROM THE DATA

4. Continue

→ 4

ESTIMATE THE SEASONAL COMPONENT

We include, within the seasonal component, the elements at the seasonal frequency of 30 degrees and at the harmonic frequencies of 60 degrees, 90 degrees, 120 degrees, 150 degrees and 180 degrees. In addition, we include two elements at the points immediately above 120 degrees.

Do you wish to add some supplementary non-harmonic frequencies? Y/N

→ Y

How many additional Fourier points below 30 degrees

→ 0

How many additional Fourier points above 30 degrees

→ 0

ETC, ETC.

How many additional Fourier points below 120 degrees

D.S.G. POLLOCK: IDEOLOG—A FILTERING PROGRAM

→ 0

How many additional Fourier points above 120 degrees

→ 2

ETC, ETC.

How many additional Fourier points below 180 degrees

→ 0

⟨⟨*Figure. The seasonal component extracted from 144 points of an unidentified data series*⟩⟩

→ ⟨RETURN⟩

⟨⟨*Figure. A seasonally adjusted version of a sequence of 144 points of an unidentified data series*⟩⟩

→ ⟨RETURN⟩

IDEOLOG: A PROGRAM FOR FILTERING ECONOMETRIC DATA—A SYNOPSIS OF ALTERNATIVE METHODS

An account is given of various filtering procedures that have been implemented in a computer program, which can be used in analysing econometric time series. The program provides some new filtering procedures that operate primarily in the frequency domain. Their advantage is that they are able to achieve clear separations of components of the data that reside in adjacent frequency bands in a way that the conventional time-domain methods cannot.

Several procedures that operate exclusively within the time domain have also been implemented in the program. Amongst these are the bandpass filters of Baxter and King and of Christiano and Fitzgerald, which have been used in estimating business cycles. The Butterworth filter and the Leser or Hodrick–Prescott are described and it is shown that they can be implemented both in the time domain and in the frequency domain. The Henderson filter, which has been used in the X-11 program for seasonal adjustment, has also been implemented in the IDEOLOG program

Econometric filtering procedures must be able to cope with the trends that are typical of economic time series. If a trended data sequence has been reduced to stationarity by differencing prior to its filtering, then the filtered sequence will need to be re-inflated. This can be achieved within the time domain via the summation operator, which is the inverse of the difference operator. The effects of the differencing can also be reversed within the frequency domain by recourse to the frequency-response function of the summation operator.

1. Introduction

This paper gives an account of some of the facilities that are available in a new computer program, which implements various filters that can be used for extracting the components of an economic data sequence and for producing smoothed and seasonally-adjusted data from monthly and quarterly sequences.

The program can be downloaded from the following web address:

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

It is accompanied by a collection of data and by three log files, which record steps that can be taken in processing some typical economic data. Here, we give an account of the theory that lies behind some of the procedures of the program.

The program originated in a desire to compare some new methods with existing procedures that are common in econometric analyses. The outcome has been a comprehensive facility, which will enable a detailed investigation of univariate econometric time series. The program will also serve to reveal the extent to which the results of an economic analysis might be the consequence of the choice of a particular filtering procedure.

The new procedures are based on the Fourier analysis of the data, and they perform their essential operations in the frequency domain as opposed to the time domain. They depend upon a Fourier transform for carrying the data into the frequency domain and upon an inverse transform for carrying the filtered elements back to the time domain. Filtering procedures usually operate exclusively in the

time domain. This is notwithstanding fact that, for a proper understanding of the effects of a filter, one must know its frequency-response function.

The sections of this paper give accounts of the various classes of filters that have been implemented in the program. In the first category, to which section 3 is devoted, are the simple finite impulse response (FIR) or linear moving-average filters that endeavour to provide approximations to the so-called ideal frequency-selective filters. Also in this category of FIR filters is the time-honoured filter of Henderson (1916), which is part of a seasonal-adjustment program that is widely used in central statistical agencies.

The second category concerns filters of the infinite impulse response (IIR) variety, which involve an element of feedback. The filters of this category that are implemented in the program are all derived according to the Wiener–Kolmogorov principle. The principle has been enunciated in connection with the filtering of stationary and doubly-infinite data sequences—see Whittle (1983), for example. However, the purpose of the program is to apply these filters to short non stationary sequences. In section 4, the problem of non stationarity is broached, whereas, in section 5, the adaptations that are appropriate to short sequences are explained.

Section 6 deals with the new frequency-domain filtering procedures. The details of their implementation are described and some of their uses are highlighted. In particular, it is shown how these filters can achieve an ideal frequency selection, whereby all of the elements of the data that fall below a given cut-off frequency are preserved and all those that fall above it are eliminated.

2. The Effects of a Linear Filter

A stationary data sequence can be resolved into a sum of sinusoidal elements with frequencies that range from zero up to the Nyquist frequency of π radians per sample interval, which represents the highest frequency that is observable in sampled data.

A data sequence $y(t) = \{y_t, t = 0, 1, \dots, T-1\}$ comprising $T = 2n$ observations has the following Fourier decomposition:

$$y(t) = \sum_{j=0}^n \{\alpha_j \cos(\omega_j t) + \beta_j \sin(\omega_j t)\}. \quad (1)$$

Here, $\omega_j = 2\pi j/T; j = 0, \dots, n$, are the Fourier frequencies, which are equally spaced in the interval $[0, \pi]$, whereas α_j, β_j are the associated Fourier coefficients, which indicate the amplitudes of the sinusoidal elements of the data sequence. Since the trigonometric functions constitute perpetual circular sequences, it is appropriate to regard the data also as elements of a circular sequence.

A more convenient expression for the data is obtained by using Euler's equations to express the trigonometric functions in terms of complex exponential functions. Thus

$$\begin{aligned} \cos(\omega_j t) &= \frac{e^{i\omega_j t} + e^{-i\omega_j t}}{2} \quad \text{and} \\ \sin(\omega_j t) &= \frac{-i}{2}(e^{i\omega_j t} - e^{-i\omega_j t}) = \frac{1}{2i}(e^{i\omega_j t} - e^{-i\omega_j t}). \end{aligned} \quad (2)$$

The inverse relationships are

$$e^{i\omega t} = \cos(\omega t) + i \sin(\omega t) \quad \text{and} \quad e^{-i\omega t} = \cos(\omega t) - i \sin(\omega t). \quad (3)$$

Using (2) in equation (1) gives

$$y(t) = \sum_{j=-n}^n \zeta_j e^{-i\omega t}, \quad \text{with} \quad \zeta_j = \frac{\alpha_j - i\beta_j}{2} \quad \text{and} \quad \zeta_{-j} = \frac{\alpha_j + i\beta_j}{2}. \quad (4)$$

In representing the properties of a linear filter, it is common to imagine that it is operating on a doubly-infinite data sequence of a statistically stationary nature. Then, the Fourier decomposition of the data comprises an infinity of exponential elements of negligible amplitudes of which the frequencies form a continuum in the interval $[-\pi, \pi]$.

The effects of a linear filter on the exponential elements are specific to the frequencies. In representing these effects, it is appropriate to consider the case of an elementary complex exponential function $x(t) = \exp(-i\omega t) = \cos(\omega t) - i \sin(\omega t)$. A filter defined by the coefficients $\{\psi_j\}$ forms a weighted combination of the values of the exponential function over an interval of time that is spanned by the coefficients. Its output is given by

$$y(t) = \sum_j \psi_j e^{-i\omega(t-j)} = \left\{ \sum_j \psi_j e^{-i\omega j} \right\} e^{i\omega t}. \quad (5)$$

The effects of the filter are summarised by the complex frequency response function

$$\psi(\omega) = \sum_j \psi_j e^{-i\omega j} = |\psi(\omega)| e^{-i\theta(\omega)}. \quad (6)$$

On the RHS, there is the gain factor $|\psi(\omega)|$, which corresponds to the modulus of the complex function, and the phase angle $\theta(\omega)$, which corresponds to its argument. The gain effect serves to increase or to diminish the amplitudes of the sinusoidal components of the exponential function and the phase effect advances them or delays them in time.

The frequency response function may be separated into its real and imaginary components, whereby

$$\psi(\omega) = \psi^{\text{re}}(\omega) + i\psi^{\text{im}}(\omega) \quad (7)$$

with

$$\psi^{\text{re}}(\omega) = \sum_j \psi_j \cos(\omega j) \quad \text{and} \quad \psi^{\text{im}}(\omega) = - \sum_j \psi_j \sin(\omega j). \quad (8)$$

Then, the squared gain of the filter is given by

$$|\psi(\omega)|^2 = \{\psi^{\text{re}}(\omega)\}^2 + \{\psi^{\text{im}}(\omega)\}^2, \quad (9)$$

while its phase displacement is given by

$$\theta(\omega) = - \tan^{-1} \left\{ \frac{\psi^{\text{im}}(\omega)}{\psi^{\text{re}}(\omega)} \right\}. \quad (10)$$

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

Within the context of a doubly infinite data sequence, all of the filters that are considered in this paper, and which are implemented in the IDEOLOG program, are of the symmetric variety. Therefore, their frequency response functions are devoid of phase effects. However, in order to accommodate finite data sequences, various adaptations are made to enable the filters to reach the ends of the samples. The phase effect of such adaptations are not readily apparent.

In analysing a filter, it is helpful to employ the z -transform of the coefficient sequence, which is

$$\psi(z) = \sum_j \psi_j z^j, \quad (11)$$

where z , which is described as an algebraic indeterminate, may be regarded as an arbitrary complex number. However, on setting $z = \exp\{-i\omega\}$ the z -transform becomes the frequency response function. It is also helpful to form the z -transform of the data sequence, which is $y(z) = \sum_t y_t z^t$. Then, there is

$$\psi(z)y(z) = \left\{ \sum_j \psi_j y_{t-j} \right\} z^t, \quad (12)$$

from which it can be seen that the convolution of the sequences corresponds to a multiplication of their z -transforms.

3. The FIR Filters

One of the purposes in filtering economic data sequences is to obtain a representation of the business cycle that is free from the distractions of seasonal fluctuations and of high-frequency noise. According to Baxter and King (1999), the business cycle should comprise all elements of the data that have cyclical durations of no less than of one and a half years and not exceeding eight years, which was the definition of Burns and Mitchell (1946). For this purpose, they have proposed to use a moving-average bandpass filter to approximate the ideal frequency-selective filter. An alternative approximation, which has the same purpose, has been proposed by Christiano and Fitzgerald (2003). Both of these filters have been implemented in the program.

An ideal filter is one that transmits the elements that fall within a specified frequency band, described as the pass band, and which blocks elements at all other frequencies, which constitute the stop band.

For such a filter, the gain is unity within the pass band and zero within the stop band. Such a response is depicted in Figure 1, where the pass band, which runs from $\pi/16$ to $\pi/3$ radians per sample interval, is intended to transmit the elements of a quarterly econometric data sequence that constitute the business cycle.

To achieve an ideal frequency selection with a linear moving-average filter would require an infinite number of filter coefficients. This is clearly impractical; and so the sequence of coefficients must be truncated, whereafter it may be modified in certain ways to diminish the adverse effects of the truncation. Nevertheless, it




Figure 1. The frequency response of the truncated bandpass filter of 25 coefficients superimposed upon the ideal frequency response. The lower cut-off point is at $\pi/16$ radians (11.25°), corresponding to a period of 32 quarters, and the upper cut-off point is at $\pi/3$ radians (60°), corresponding to a period of the 6 quarters.

is possible to implement an ideal filter in the frequency domain. as is described in section 6.

3.1 Approximation to the Ideal Filter

Figure 1 also shows the frequency response of a filter that has been derived by taking twenty-five of the central coefficients of the ideal filter and adjusting their values by equal amounts so that they sum to zero. This is the filter that has been proposed by Baxter and King (1999) for the purpose of extracting the business cycle from economic data. The filter is affected by a considerable leakage, whereby elements that fall within the stop band are transmitted in part by the filter.

The z -transform of a sequence $\{\psi_j\}$ of filter coefficients is the polynomial $\psi(z) = \sum_j \psi_j z^j$. Constraining the coefficients to sum to zero ensures that the polynomial has a root of unity, which is to say that $\psi(1) = \sum_j \psi_j = 0$. This implies that $\nabla(z) = 1 - z$ is a factor of the polynomial, which indicates that the filter incorporates a differencing operator.

If the filter is symmetric, such that $\psi(z) = \psi_0 + \psi_1(z + z^{-1}) + \dots + \psi_q(z^q + z^{-q})$ and, therefore, $\psi(z) = \psi(z^{-1})$, then $1 - z^{-1}$ is also a factor. Then, $\psi(z)$ has the combined factor $(1 - z)(1 - z^{-1}) = -z\nabla(z)^2$, which indicates that the filter incorporates a twofold differencing operator. Such a filter is effective in reducing a linear trend to zero; and, therefore, it is applicable to the logarithms of an econometric data sequences with an underlying geometric or exponential trend.

The filter of Baxter and King (1999), which fulfils this condition, is appropriate for the purpose of extracting the business cycle from a trended data sequence. Figure 2 shows the logarithms of data of U.K. real domestic consumption for the years 1955–1994 through which a linear trend has been interpolated. Figure 3 shows the results of subjecting these data to the Baxter–King filter. A disadvantage of the filter, which is apparent in Figure 3, is that it is incapable of reaching the ends of the sample. The first q sample values and the last q remain unprocessed.




Figure 2. 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 3. The sequence derived by applying the truncated bandpass filter of 25 coefficients to the quarterly logarithmic data on U.K. Consumption.




Figure 4. The sequence derived by applying the bandpass filter of Christiano and Fitzgerald to the quarterly logarithmic data on U.K. Consumption.

To overcome this difficulty, Christiano and Fitzgerald (2003) have proposed a filter with a variable set of coefficients. To generate the filtered value at time t , they associate the central coefficient ψ_0 with y_t . If y_{t-p} falls within the sample, then they associate it with the coefficient ψ_p . Otherwise, if it falls outside the sample, it is disregarded. Likewise, if y_{t+p} falls within the sample, then it is associated with ψ_p , otherwise it is disregarded. If the data follow a first-order random walk, then the first and the last sample elements y_0 and y_{T-1} receive extra weights A and B , which correspond to the sums of the coefficients discarded from the filter at either end. The resulting filtered value at time t may be denoted by

$$x_t = Ay_0 + \psi_t y_0 + \cdots + \psi_1 y_{t-1} + \psi_0 y_t + \psi_1 y_{t+1} + \cdots + \psi_{T-1-t} y_{T-1} + By_{T-1}. \quad (13)$$

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

For data that appear to have been generated by a first-order random walk with a constant drift, it is appropriate to extract a linear trend before filtering the residual sequence. Figure 4 provides an example of the practice. In fact, this has proved to be the usual practice in most circumstances.

Within the category of FIR filters, the program also implements the time honoured smoothing filter of Henderson (1916), which forms an essential part of the detrending procedure of the X-11 program of the Bureau of the Census. This program provides the method of seasonal adjustment that is used predominantly by central statistical agencies.

Here, the end-of-sample problem is overcome by supplementing the Henderson filter with a set of asymmetric filters that can be applied to the elements of the first and the final segments. These are the Musgrave (1964) filters. (See Quenneville, Ladiray and Lefranc, 2003 for a recent account of these filters.) In the X-11 ARIMA variant, which is used by Statistics Canada, the alternative recourse is adopted of extrapolating the data beyond the ends of the sample so that it can support a time-invariant filter that does run to the ends.

4. The Wiener–Kolmogorov Filters

The program also provides several filters of the feedback variety that are commonly described as infinite-impulse response (IIR) filters. The filters in question are derived according to the finite-sample Wiener–Kolmogorov principle that has been expounded by Pollock (2000, 2007).

The ordinary theory of Wiener–Kolmogorov filtering assumes a doubly-infinite data sequence $y(t) = \xi(t) + \eta(t) = \{y_t; t = 0, \pm 1, \pm 2, \dots\}$ generated by a stationary stochastic process. The process is compounded from a signal process $\xi(t)$ and a noise process $\eta(t)$ that are assumed to be statistically independent and to have zero-valued means. Then, the autocovariance generating function of $y(t)$ is given by

$$\gamma_y(z) = \gamma_\xi(z) + \gamma_\eta(z), \quad (14)$$

which is sum of the autocovariance functions of $\xi(t)$ and $\eta(t)$.

The object is to extract estimates of the signal sequence $\xi(t)$ and noise sequence $\eta(t)$ from the data sequence. The z -transforms of the relevant filters are

$$\psi_\xi(z) = \frac{\gamma_\xi(z)}{\gamma_\xi(z) + \gamma_\eta(z)} = \frac{\theta_\xi(z^{-1})\theta_\xi(z)}{\phi(z^{-1})\phi(z)}, \quad (15)$$

and

$$\psi_\eta(z) = \frac{\gamma_\eta(z)}{\gamma_\xi(z) + \gamma_\eta(z)} = \frac{\theta_\eta(z^{-1})\theta_\eta(z)}{\phi(z^{-1})\phi(z)}. \quad (16)$$

It can be seen that $\psi_\xi(z) + \psi_\eta(z) = 1$, in view of which the filters can be described as complementary.

The factorisations of the filters that are given on the RHS enable them to be applied via a bi-directional feedback process. In the case of the signal extraction filter $\psi_\xi(z)$, the process in question can be represented by the equations

$$\phi(z)q(z) = \theta_\xi(z)y(z) \quad \text{and} \quad \phi(z^{-1})x(z) = \theta_\xi(z^{-1})q(z^{-1}), \quad (17)$$

wherein $q(z)$, $y(z)$ and $x(z)$ stand for the z -transforms of the corresponding sequences $q(t)$, $y(t)$ and $x(t)$.

To elucidate these equations, we may note that, in the first of them, the expression associated with z^t is

$$\sum_{j=0}^m \phi_j q_{t-j} = \sum_{j=0}^n \theta_{\xi,j} y_{t-j}. \quad (18)$$

Given that $\phi_0 = 1$, this serves to determine the value of q_t . Moreover, given that the recursion is assumed to be stable, there need be no restriction on the range of t . The first equation, which runs forward in time, generates an intermediate output $q(t)$. The second equation, which runs backwards in time, generates the final filtered output $x(t)$.

4.1 Filters for Trended Data

The Wiener–Kolmogorov (W–K) methodology can be applied to nonstationary data with minor adaptations. A model of the processes underlying the data can be adopted that has the form of

$$\begin{aligned} y(z) &= \xi(z) + \eta(z) \\ &= \frac{(1+z)^n}{\nabla^d(z)} \zeta(z) + \frac{(1-z)^m}{\nabla^d(z)} \varepsilon(z), \end{aligned} \quad (19)$$

where $\zeta(t)$ and $\varepsilon(t)$ are mutually independent white-noise sequences of independently and identically distributed random variables, and where $\nabla(z) = 1 - z$ is the z -transform of the difference operator. The condition $m \geq d$ is necessary to ensure the stationarity of $\eta(t)$, which is obtained from $\varepsilon(t)$ by differencing $m - d$ times.

It must be conceded that a nonstationary process of infinite duration, such as $\xi(t)$, is a mathematical construct of doubtful reality, since its values will be




Figure 5. The frequency-response function of the Hodrick–Prescott smoothing filter for various values of the smoothing parameter λ .




Figure 6. The frequency-response function of the lowpass Butterworth filters of orders $n = 6$ and $n = 12$ with a nominal cut-off point of $2\pi/3$ radians.

unbounded, almost certainly. Nevertheless, to deal in these terms is to avoid the complexities of the finite-sample approach, which will be the subject of Section 5.

The filters that are applied to $y(t)$ to estimate $\xi(t)$ and $\eta(t)$ takes the form of

$$\psi_{\xi}(z) = \frac{\sigma_{\zeta}^2(1+z^{-1})^n(1+z)^n}{\sigma_{\zeta}^2(1+z^{-1})^n(1+z)^n + \sigma_{\varepsilon}^2(1-z^{-1})^m(1-z)^m}, \quad (20)$$

and

$$\psi_{\eta}(z) = \frac{\sigma_{\varepsilon}^2(1-z^{-1})^m(1-z)^m}{\sigma_{\zeta}^2(1+z^{-1})^n(1+z)^n + \sigma_{\varepsilon}^2(1-z^{-1})^m(1-z)^m}. \quad (21)$$

It will be observed that the operator $\nabla^d(z)$ is absent from the numerator and the denominator on account of its cancellation.

Within the program, two such filters have been implemented. The first is the filter of Leser (1961) and of Hodrick and Prescott (1980, 1997), which is obtained by setting $d = m = 2$ and $n = 0$ in equation (10). This gives rise to a filter that is




Figure 7. The pole–zero diagrams of the lowpass Butterworth filters for $n = 6$ when the cut-off is at $\omega = \pi/2$ (left) and at $\omega = \pi/8$.

appropriate for extracting the non stationary signal or trend component when the data are generated according to the equation

$$y(z) = \frac{\zeta(z)}{\nabla^2(z)} + \varepsilon(z). \quad (22)$$

The z -transforms of the relevant filters become

$$\psi_\xi(z) = \frac{1}{1 + \lambda \nabla^2(z^{-1}) \nabla^2(z)}, \quad (23)$$

and

$$\psi_\eta(z) = \frac{\nabla^2(z^{-1}) \nabla^2(z)}{\lambda^{-1} + \nabla^2(z^{-1}) \nabla^2(z)}, \quad (24)$$

where $\lambda = \sigma_\varepsilon^2 / \sigma_\zeta^2$, which is described as the smoothing parameter.

The frequency-response functions of the filters for various values of λ are shown in Figure 5. These are obtained by setting $z = \exp(-i\omega) = \cos(\omega) - i \sin(\omega)$ in the formula of (10) and by letting ω run from 0 to π . (In the process, the imaginary quantities are cancelled so as to give rise to the real-valued functions that are plotted in the diagram.)

It is notable that the specification of the underlying process $y(t)$, in which both the signal component $\xi(z) = \zeta(z)/\nabla^2(z)$ and the noise component $\eta(z)$ have spectral density functions that extend over the entire frequency range, precludes the clear separation of the components. This is reflected in the fact that, for all but the highest values λ , the filter transmits significant proportions of the elements at all frequencies.

The second Wiener–Kolmogorov filter that is implemented in the program is the Butterworth filter that is obtained by setting $m = n$ in equation (20). This filter is capable of a much firmer discrimination between the signal and noise than is the Leser (1961) filter.

The Butterworth (1930) filter was originally devised as an analogue filter. It can also be rendered in digital form—See Pollock (2000). The filters that are

appropriate for extracting the components $\xi(z)$ and $\eta(z)$ from

$$\begin{aligned} y(z) &= \xi(z) + \eta(z) \\ &= \frac{(1+z)^n}{\nabla^d(z)} \zeta(z) + \frac{(1-z)^n}{\nabla^d(z)} \varepsilon(z), \end{aligned} \quad (25)$$

are

$$\psi_\xi(z) = \frac{(1+z^{-1})^n(1+z)^n}{(1+z^{-1})^n(1-z)^n + \lambda(1-z^{-1})^n(1+z)^n}, \quad (26)$$

and

$$\psi_\eta(z) = \frac{(1-z^{-1})^n(1-z)^n}{(1+z^{-1})^n(1-z)^n + \lambda^{-1}(1-z^{-1})^n(1+z)^n}, \quad (27)$$

where $\lambda = \sigma_\kappa^2/\sigma_\delta^2$, and where it is assumed the $n > d$.

It is straightforward to determine the value of λ that will place the cut-off of the filter at a chosen point $\omega_c \in (0, \pi)$. Consider setting $z = \exp\{-i\omega\}$ in the formula of (15) of the lowpass filter. This gives the following expression for the gain:

$$\begin{aligned} \psi_\xi(e^{-i\omega}) &= \frac{1}{1 + \lambda \left(i \frac{1 - e^{-i\omega}}{1 + e^{-i\omega}} \right)^{2n}} \\ &= \frac{1}{1 + \lambda \{ \tan(\omega/2) \}^{2n}}. \end{aligned} \quad (28)$$

At the cut-off point, the gain must equal $1/2$, whence solving the equation $\beta_\xi(\exp\{-i\omega_c\}) = 1/2$ gives $\lambda = \{1/\tan(\omega_c/2)\}^{2n}$.

Figure 6 shows how the rate of the transition of the Butterworth frequency response between the pass band and the stop band is affected by the order n of the filter. Figure 7 shows the pole-zero diagrams of filters with different cut-off points. As the cut-off frequency is reduced, the transition between the two bands becomes more rapid. Also, some of the poles of the filter move towards the perimeter of the unit circle.

4.2 A Filter for Seasonal Adjustment

The Wiener-Kolmogorov principle is also used in deriving a filter for the seasonal adjustment of monthly and quarterly econometric data. The filter 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 corresponding zeros at the same angles but located inside the circle. The z -transform of the output sequence gives

$$\begin{aligned} y(z) &= \eta(z) + \frac{R(z)}{S(z)} \nu(z) \quad \text{or} \\ S(z)y(z) &= S(z)\eta(z) + R(z)\nu(z), \end{aligned} \quad (29)$$

where

$$R(z) = 1 + \rho z + \rho^2 z^2 + \dots + \rho^{s-1} z^{s-1} \quad (30)$$




Figure 8. The gain of a filter for the seasonal adjustment of monthly data. The defining parameters are $\rho = 0.9$ and $\lambda = \sigma_\eta^2/\sigma_\nu^2 = 0.125$

with $\rho < 1$, and

$$S(z) = 1 + z + z^2 + \cdots + z^{s-1}. \quad (31)$$

The z -transform of the seasonal-adjustment filter is

$$\psi(z) = \frac{\sigma_\eta^2 S(z)S(z^{-1})}{S(z)S(z^{-1})\sigma_\eta^2 + \sigma_\nu^2 R(z)R(z^{-1})}. \quad (32)$$

Setting $z = \exp\{-i\omega\}$ 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 8 for the case where $\rho = 0.9$ and $\lambda = \sigma_\eta^2/\sigma_\nu^2 = 0.125$.

5. The Finite-Sample Realisations of the W–K Filters

To derive the finite-sample version of a Wiener–Kolmogorov filter, we may consider a data vector $y = [y_0, y_1, \dots, y_{t-1}]'$ that has a signal component ξ and a noise component η :

$$y = \xi + \eta. \quad (33)$$

The two components are assumed to be independently normally distributed with zero means and with positive-definite dispersion matrices. Then,

$$\begin{aligned} E(\xi) &= 0, & D(\xi) &= \Omega_\xi, \\ E(\eta) &= 0, & D(\eta) &= \Omega_\eta, \\ & & \text{and } C(\xi, \eta) &= 0. \end{aligned} \quad (34)$$

A consequence of the independence of ξ and η is that $D(y) = \Omega_\xi + \Omega_\eta$.

The estimates of ξ and η , which may be denoted by x and h respectively, are derived according to the following criterion:

$$\text{Minimise } S(\xi, \eta) = \xi' \Omega_\xi^{-1} \xi + \eta' \Omega_\eta^{-1} \eta \quad \text{subject to } \xi + \eta = y. \quad (35)$$

Since $S(\xi, \eta)$ is the exponent of the normal joint density function $N(\xi, \eta)$, the resulting estimates may be described, alternatively, as the minimum chi-square estimates or as the maximum-likelihood estimates.

Substituting for $\eta = y - \xi$ gives the concentrated criterion function $S(\xi) = \xi' \Omega_\xi^{-1} \xi + (y - \xi)' \Omega_\eta^{-1} (y - \xi)$. Differentiating this function in respect of ξ and setting the result to zero gives a condition for a minimum, which specifies the estimate x . This is $\Omega_\eta^{-1} (y - x) = \Omega_\xi^{-1} x$, which, on pre multiplication by Ω_η , can be written as $y = x - \Omega_\eta \Omega_\xi^{-1} x = (\Omega_\xi + \Omega_\eta) \Omega_\xi^{-1} x$. Therefore, the solution for x is

$$x = \Omega_\xi (\Omega_\xi + \Omega_\eta)^{-1} y. \quad (36)$$

Moreover, since the roles of ξ and η are interchangeable in this exercise, and, since $h + x = y$, there are also

$$h = \Omega_\eta (\Omega_\xi + \Omega_\eta)^{-1} y \quad \text{and} \quad x = y - \Omega_\eta (\Omega_\xi + \Omega_\eta)^{-1} y. \quad (37)$$

The filter matrices $B_\xi = \Omega_\xi (\Omega_\xi + \Omega_\eta)^{-1}$ and $B_\eta = \Omega_\eta (\Omega_\xi + \Omega_\eta)^{-1}$ of (36) and (37) are the matrix analogues of the z -transforms displayed in equations (15) and (16).

A simple procedure for calculating the estimates x and h begins by solving the equation

$$(\Omega_\xi + \Omega_\eta) b = y \quad (38)$$

for the value of b . Thereafter, one can generate

$$x = \Omega_\xi b \quad \text{and} \quad h = \Omega_\eta b. \quad (39)$$

If Ω_ξ and Ω_η correspond to the narrow-band dispersion matrices of moving-average processes, then the solution to equation (38) may be found via a 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 procedure has been described by Pollock (2000).

5.1 Filters for Short Trended Sequences

To adapt these estimates to the case of trended data sequences may require the provision of carefully determined initial conditions with which to start the recursive processes. A variety of procedures are available that are similar, if not identical, in their outcomes. The procedures that are followed in the program depend upon reducing the data sequences to stationarity, in one way or another, before subjecting them to the filters. After the data have been filtered, the trend is liable to be restored.

The first method, which is the simplest in concept, requires the trend to be represented by a polynomial function. In some circumstances, when the economy has been experiencing steady growth, the polynomial will serve as a reasonable characterisation of its underlying trajectory. Thus, in the period 1955–1994 a log-linear trend function provides a firm benchmark against which to measure the cyclical fluctuations of the U.K. economy. The residual deviations from this trend may be subjected to a lowpass filter; and the filtered output can be added to the

trend to produce a representation of what is commonly described as the trend-cycle component.

It is desirable that the polynomial trend should interpolate the scatter of points at either end of the data sequence. For this purpose, the program provides a method of weighted least-squares polynomial regression with a wide choice of weighting schemes, which allow extra weight to be placed upon the initial and the final runs of observations.

An alternative way of eliminating the trend is to take differences of the data. Usually, twofold differencing is appropriate. The matrix analogue of the second-order backwards difference operator in the case of $T = 5$ is given by

$$\nabla_5^2 = \begin{bmatrix} Q'_* \\ Q' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ -2 & 1 & 0 & 0 & 0 \\ \hline 1 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & 1 & -2 & 1 \end{bmatrix}. \quad (40)$$

The first two rows, which do not produce true differences, are liable to be discarded. In general, the p -fold differences of a data vector of T elements will be obtained by pre multiplying it by a matrix Q' of order $(T - p) \times T$. Applying Q' to equation (33) gives

$$\begin{aligned} Q'y &= Q'\xi + Q'\eta \\ &= \delta + \kappa = g. \end{aligned} \quad (41)$$

The dispersion matrices of the differenced vectors are

$$D(\delta) = \Omega_\delta = Q'D(\xi)Q \quad \text{and} \quad D(\kappa) = \Omega_\kappa = Q'D(\eta)Q. \quad (42)$$

The estimates d and k of the differenced components are given by

$$d = \Omega_\delta(\Omega_\delta + Q'\Omega_\eta Q)^{-1}Q'y \quad (43)$$

and

$$k = Q'\Omega_\eta Q(\Omega_\delta + Q'\Omega_\eta Q)^{-1}Q'y. \quad (44)$$

To obtain estimates of ξ and η , the estimates of their difference versions must be re-inflated via an anti-differencing or summation operator. We begin by observing that the inverse of ∇_5^2 is a twofold summation operator given by

$$\nabla_5^{-2} = [S_* \quad S] = \left[\begin{array}{cc|ccc} 1 & 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ \hline 3 & 2 & 1 & 0 & 0 \\ 4 & 3 & 2 & 1 & 0 \\ 5 & 4 & 3 & 2 & 1 \end{array} \right]. \quad (45)$$

The first two columns, which constitute the matrix S_* , provide a basis for all linear functions defined on $\{t = 0, 1, \dots, T - 1 = 5\}$. The example can be generalised to the case of a matrix ∇_T^{-p} of order T . However, in the program, the maximum order of differencing is $p = 2$.

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

$$y = S_*g_* + Sg. \quad (46)$$

In effect, the elements of g_* , which may be regarded as polynomial parameters, provide the initial conditions for the process of summation or integration, which we have been describing as a process of re-inflation.

The equations by which the estimates of ξ and η may be recovered from those of δ and κ are analogous to equation (46). They are

$$x = S_*d_* + Sd \quad \text{and} \quad h = S_*k_* + Sk. \quad (47)$$

In this case, the initial conditions d_* and k_* require to be estimated. The appropriate estimates are the values that minimise the function

$$\begin{aligned} (y-x)' \Omega_\eta^{-1} (y-x) &= (y - S_*d_* - Sd)' \Omega_\eta^{-1} (y - S_*d_* - Sd) \\ &= (S_*k_* + Sk)' \Omega_\eta^{-1} (S_*k_* + Sk). \end{aligned} \quad (48)$$

These values are

$$k_* = -(S'_* \Omega_\eta^{-1} S_*)^{-1} S'_* \Omega_\eta^{-1} Sk \quad (49)$$

and

$$d_* = (S'_* \Omega_\eta^{-1} S_*)^{-1} S'_* \Omega_\eta^{-1} (y - Sd). \quad (50)$$

Equations (49) and (50) together with (43) and (44) provide a complete solution to the problem of estimating the components of the data. However, it is possible to eliminate the initial conditions from the system of estimating equations. This can be achieved with the help of the following identity:

$$\begin{aligned} P_* &= S_*(S'_* \Omega_\eta^{-1} S_*)^{-1} S'_* \Omega_\eta^{-1} \\ &= I - \Omega_\eta Q(Q' \Omega_\eta Q)^{-1} Q' = I - P_Q. \end{aligned} \quad (51)$$

In these terms, the equation of (47) for h becomes $h = (I - P_*)Sk = P_Q Sk$. Using the expression for k from (32) together with the identity $Q'S = I_{T-2}$ gives

$$h = \Omega_\eta Q(\Omega_\delta + Q' \Omega_\eta Q)^{-1} Q'y. \quad (52)$$

This can also be obtained from the equation (44) for k by the removal of the leading differencing matrix Q' . It follows immediately that

$$\begin{aligned} x &= y - h \\ &= y - \Omega_\eta Q(\Omega_\delta + Q' \Omega_\eta Q)^{-1} Q'y. \end{aligned} \quad (53)$$

The elimination of the initial conditions is due to the fact that η is a stationary component. Therefore, it requires no initial conditions other than the zeros that are the appropriate estimates of the pre-sample elements. The direct estimate x of ξ does require initial conditions, but, in view of the adding-up conditions of (33),




Figure 9. The periodogram of the first differences of the U.K. logarithmic consumption data.




Figure 10. The periodogram of the residual sequence obtained from the linear detrending of the logarithmic consumption data. A band, with a lower bound of $\pi/16$ radians and an upper bound of $\pi/3$ radians, is masking the periodogram.

x can be obtained more readily by subtracting from y the estimate h of η , in the manner of equation (53).

Observe that, since

$$f = S_*(S'_*S_*)^{-1}S'_*y \quad (54)$$

is an expression for the vector of the ordinates of a polynomial function fitted to the data by an ordinary least-squares regression, the identity of (39) informs us that

$$f = y - Q(Q'Q)^{-1}Q'y \quad (55)$$

is an alternative expression.

The residuals of an OLS polynomial regression of degree p , which are given by $y - f = Q(Q'Q)^{-1}Q'y$, contain same the information as the vector $g = Q'y$ of the p -th 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. Therefore, the low-frequency spectral structures of the data are not

perceptible in the periodogram of the differenced sequence. Figure 9 provides evidence of this.

On the other hand, the periodogram of a trended sequence is liable to be dominated by its low-frequency components, which will mask the other spectral structures. However, the periodogram of the residuals of the polynomial regression can be relied upon to reveal the spectral structures at all frequencies. Moreover, by varying the degree p of the polynomial, one is able to alter the relative emphasis that is given to high-frequency and low-frequency structures. Figure 10 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.

5.2 The Leser (H–P) Filter and the Butterworth Filter

The specific cases that have been considered in the context of the classical form of the Wiener–Kolmogorov filter can now be adapted to the circumstances of short trended sequences. First, there is the Leser or H–P filter. This is derived by setting

$$D(\eta) = \Omega_\eta = \sigma_\eta^2 I, \quad D(\delta) = \Omega_\delta = \sigma_\delta^2 I \quad \text{and} \quad \lambda = \frac{\sigma_\eta^2}{\sigma_\delta^2} \quad (56)$$

within (93) to give

$$x = y - Q(\lambda^{-1}I + Q'Q)^{-1}Q'y \quad (57)$$

Here, λ is the so-called smoothing parameter. It will be observed that, as $\lambda \rightarrow \infty$, the vector x tends to that of a linear function interpolated into the data by least-squares regression, which is represented by equation (55). The matrix expression $\Psi = I - Q(\lambda^{-1}I + Q'Q)^{-1}Q'$ for the filter can be compared to the polynomial expression $\psi_\eta(z) = 1 - \psi_\xi(z)$ of the classical formulation, which entails the z -transform from (24).

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

$$x = y - \lambda \Sigma Q (M + \lambda Q' \Sigma Q)^{-1} Q' y, \quad (58)$$

where $\lambda = \{1/\tan(\omega_c/2)\}^{2n}$ wherein ω_c is the nominal cut-off frequency. Here, the matrices

$$\Sigma = \{2I_T - (L_T + L'_T)\}^{n-2} \quad \text{and} \quad M = \{2I_T + (L_T + L'_T)\}^n \quad (59)$$

are obtained from the RHS of the equations $\{(1-z)(1-z^{-1})\}^{n-2} = \{2 - (z + z^{-1})\}^{n-2}$ and $\{(1+z)(1+z^{-1})\}^n = \{2 + (z + z^{-1})\}^n$, respectively, by replacing z by L_T and z^{-1} by L'_T . Observe that the equalities no longer hold after the replacements. However, it can be verified that

$$Q' \Sigma Q = \{2I_T - (L_T + L'_T)\}^n. \quad (60)$$

5.3 A Flexible Smoothing Filter




Figure 11. The logarithms of annual U.K. real GDP from 1873 to 2001 with an interpolated trend. The trend is estimated via a filter with a variable smoothing parameter.

A derivation of the estimator of ξ is available that completely circumvents the problem of the initial conditions. This can be illustrated with the case of a generalised version of the Leser (1961) filter in which the smoothing parameter is permitted vary over the course of the sample. The values of the smoothing parameter are contained in the diagonal matrix $\Lambda = \text{diag}\{\lambda_0, \lambda_1, \dots, \lambda_{T-1}\}$. Then, the criterion for finding the vector is to minimise

$$L = (y - \xi)'(y - \xi) + \xi'Q\Lambda Q'\xi. \quad (61)$$

The first term in this expression penalises departures of the resulting curve from the data, whereas the second term imposes a penalty for a lack of smoothness in the curve. The second term comprises $d = Q'\xi$, which is the vector of the p -th order differences of ξ . The matrix Λ serves to generalise the overall measure of the curvature of the function that has the elements of ξ as its sampled ordinates, and it serves to regulate the penalty for roughness, which may vary over the sample.

Differentiating L with respect to ξ and setting the result to zero, in accordance with the first-order conditions for a minimum, gives

$$y - x = Q\Lambda Q'x = Q\Lambda d. \quad (62)$$

Multiplying the equation by Q' gives $Q'(y - x) = Q'y - d = Q'Q\Lambda d$, whence $\Lambda d = (\Lambda^{-1} + Q'Q)^{-1}Q'y$. Putting this into the equation $x = y - Q\Lambda d$ gives

$$\begin{aligned} x &= y - Q(\Lambda^{-1} + Q'Q)^{-1}Q'y \\ &= y - Q\Lambda(I + \Lambda Q'Q)^{-1}Q'y. \end{aligned} \quad (63)$$

This filter has been implemented in the program under the guise of a variable smoothing procedure. By giving a high value to the smoothing parameter, a stiff curve can be generated, which approaches a straight line as $\lambda \rightarrow \infty$. On the other hand, structural breaks can be accommodated by greatly reducing the value of the smoothing parameter in their neighbourhood. When $\lambda \rightarrow 0$, the filter tends to transmit the unaltered data values.

Figure 11 shown an example of the use of this filter. There were brief disruptions to the steady upwards progress of GDP in the U.K. after the two world wars. These breaks have been absorbed into the trend by reducing the value of the smoothing parameter in their localities. By contrast, the break that is evident in the data following the year 1929 has not been accommodated in the trend.

5.4 A Seasonal-Adjustment Filter

The need for initial conditions cannot be circumvented in cases where the seasonal adjustment filter is applied to trended sequences. Consider the filter that is applied to the differenced data $g = Q'y$ to produce a seasonally-adjusted sequence q . Then, there is

$$q = Q'_S(Q'_S Q_S + \lambda^{-1} Q'_R Q_R)^{-1} Q'_S g, \quad (64)$$

where Q'_R and Q'_S are the matrix counterparts of the polynomial operators $R(z)$ and $S(z)$ of (30) and (31) respectively. The seasonally adjusted version of the original trended data will be obtained by re-inflating the filtered sequence q via the equation

$$j = S_* q_* + S q, \quad (65)$$

where

$$q_* = (S'_* S_*)^{-1} S'_*(y - S q) \quad (66)$$

is the value that minimises the function

$$(y - j)'(y - j) = (y - S_* q_* + S q)'(S_* q_* + S q). \quad (67)$$

6. The Frequency-Domain Filters

Often, in the analysis economic data, we would profit from the availability of a sharp filter, with a rapid transition between the stop band and the pass band that is capable of separating components of the data that lie in closely adjacent frequency bands.

An example of the need for such a filter is provided by a monthly data sequence with an annual seasonal pattern superimposed on a trend-cycle trajectory. The fundamental seasonal frequency is of $\pi/6$ radians or 30 degrees per month, whereas the highest frequency of the trend-cycle component is liable to exceed $\pi/9$ radians or 20 degrees. This leaves a narrow frequency interval in which a filter that is intended to separate the trend-cycle component from the remaining elements must make the transition from its pass band to its stop band.

To achieve such a sharp transition, a FIR or moving-average filter requires numerous coefficients covering a wide temporal span. Such filters are inappropriate to the short data sequences that are typical of econometric analyses. Rational filters or feedback filters, as we have described them, are capable of somewhat sharper transitions, but they also have their limitations.

When a sharp transition is achieved by virtue of a rational filter with relatively many coefficients, the filter tends to be unstable on account of the proximity of some its poles to the circumference of the unit circle. (See Figure 7 for an example.)

Such filters can be excessively influenced by noise contamination in the data and by the enduring effects of ill-chosen initial conditions.

A more effective way of achieving a sharp cut-off is to conduct the filtering operations in the frequency domain. Reference to equation (1) shows that an ideal filter can be obtained by replacing with zeros the Fourier coefficients that are associated with frequencies that fall within the stop band.

6.1 Complex Exponentials and the Fourier Transform

The Fourier coefficients are determined by regressing the data on the trigonometrical functions of the Fourier frequencies according to the following formulae:

$$\alpha_j = \frac{2}{T} \sum_t y_t \cos \omega_j t, \quad \text{and} \quad \beta_j = \frac{2}{T} \sum_t y_t \sin \omega_j t. \quad (68)$$

Also, there is $\alpha_0 = T^{-1} \sum_t y_t = \bar{y}$, and, in the case where $T = 2n$ is an even number, there is $\alpha_n = T^{-1} \sum_t (-1)^t y_t$.

It is more convenient to work with complex Fourier coefficients and with complex exponential functions in place sines and cosines. Therefore, we define

$$\zeta_j = \frac{\alpha_j - i\beta_j}{2}. \quad (69)$$

Since $\cos(\omega_j t) - i\sin(\omega_j t) = e^{-i\omega_j t}$, it follows that the complex Fourier transform and its inverse are given by

$$\zeta_j = \frac{1}{T} \sum_{t=0}^{T-1} y_t e^{-i\omega_j t} dt \quad \longleftrightarrow \quad y_t = \sum_{j=0}^{T-1} \zeta_j e^{i\omega_j t}, \quad (70)$$

where $\zeta_{T-j} = \zeta_j^* = (\alpha_j + i\beta_j)/T$. For a matrix representation of these transforms, one may define

$$\begin{aligned} U &= T^{-1/2} [\exp\{-i2\pi t j/T\}; t, j = 0, \dots, T-1], \\ \bar{U} &= T^{-1/2} [\exp\{i2\pi t j/T\}; t, j = 0, \dots, T-1], \end{aligned} \quad (71)$$

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

$$\zeta = T^{-1/2} U y \quad \longleftrightarrow \quad y = T^{1/2} \bar{U} \zeta, \quad (72)$$

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

This notation can be used to advantage for representing the process of applying an ideal frequency-selective filter. 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 $J\zeta$. By an application of the inverse Fourier




Figure 12. The residual sequence from fitting a linear trend to the logarithmic consumption data with an interpolated line representing the business cycle.

transform, the selected elements are carried back to the time domain to form the filtered sequence. Thus, there is

$$x = \bar{U}JUy = \Psi y. \quad (73)$$

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

$$\psi_k^{\circ} = \sum_{q=-\infty}^{\infty} \psi_{qT+k}. \quad (74)$$

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 method of frequency-domain filtering can be used to mimic the effects of any linear-time invariant filter, operating 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 (73) by a diagonal matrix containing the ordinates of the desired frequency response, sampled at points corresponding to the Fourier frequencies.

In the case of the Wiener–Kolmogorov filters, defined by equation (36) and (37), one can consider replacing the dispersion matrices Ω_{ξ} and Ω_{η} by their circular counterparts

$$\Omega_{\xi}^{\circ} = \bar{U}\Lambda_{\xi}U \quad \text{and} \quad \Omega_{\eta}^{\circ} = \bar{U}\Lambda_{\eta}U. \quad (75)$$

Here, Λ_{ξ} and Λ_{η} are diagonal matrices containing ordinates sampled from the spectral density functions of the respective processes. The resulting equations for

the filtered sequences are

$$x = \Omega_{\xi}^{\circ}(\Omega_{\xi}^{\circ} + \Omega_{\eta}^{\circ})^{-1}y = \bar{U}\Lambda_{\xi}(\Lambda_{\xi} + \Lambda_{\eta})^{-1}Uy = \bar{U}J_{\xi}Uy \quad (76)$$

and

$$h = \Omega_{\eta}^{\circ}(\Omega_{\xi}^{\circ} + \Omega_{\eta}^{\circ})^{-1}y = \bar{U}\Lambda_{\eta}(\Lambda_{\xi} + \Lambda_{\eta})^{-1}Uy = \bar{U}J_{\eta}Uy. \quad (77)$$

An example of the application of the lowpass frequency-domain filter is provided by Figure 12. Here, a filter with a precise cut-off frequency of $\pi/8$ radians has been applied to the residuals from the linear detrending of the logarithms of the U.K. consumption data.

The appropriate cut-off frequency for this filter has been indicated by the periodogram of Figure 10. The smooth curve that has been interpolated through these residuals has been constituted from the Fourier ordinates in the interval $[0, \pi/8]$.

The same residual sequence has also been subjected to the approximate band-pass filter of Christiano and Fitzgerald (2003) to generate the estimate business cycle of Figure 4. This estimate fails to capture some of the salient low-frequency fluctuations of the data.

The highlighted region Figure 10 also show the extent of the pass band of the bandpass filter; and it appears that the low-frequency structure of the data falls mainly below this band. The fact that, nevertheless, the filter of Christiano and Fitzgerald does reflect a small proportion of the low-frequency fluctuations is due to its substantial leakage over the interval $[0, \pi/16]$, which falls within its nominal stop band.

6.2 Extrapolations and Detrending

To apply the frequency-domain filtering methods, the data must be free of trend. This can be achieved either by differencing the data or by applying the filter to data that are residuals from fitting a polynomial trend. The program has a facility for fitting a polynomial time trend of a degree not exceeding 15. To avoid the problems of collinearity that arise in fitting ordinary polynomials specified in terms of the powers of the temporal index t , a flexible generalised least-squares procedure is provided that depends upon a system of orthogonal polynomials.

In applying the methods, it is also important to ensure that there are no significant disjunctions in the periodic extension of the data at the points where the end of one replication of the sample sequence joins the beginning of the next replication. Equivalently, there must be a smooth transition between the start and finish points when the sequence of T data points is wrapped around a circle of circumference T .

The conventional means of avoiding such disjunctions is to taper the mean-adjusted, detrended data 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, either before or after it has been detrended.




Figure 13. The trend/cycle component of U.K. Consumption determined by the frequency-domain method, superimposed on the logarithmic data.

In the first case, a polynomial is fitted to the data; and tapered versions of the residual sequence that have been reflected around the endpoints of the sample are added to the extrapolated branches of the polynomial. Alternatively, if the data show strong seasonal fluctuations, then a tapered sequence based on successive repetitions of the ultimate seasonal cycle is added to the upper branch, and a similar sequence based on the first cycle is added to the lower branch.

In the second case, where the data have already been detrended, by the subtraction of a polynomial trend or by the application of the differencing operator, the extrapolations will be added to the horizontal axis.

This method of extrapolation will prevent the end of the sample from being joined directly to its beginning. When the data are supplemented by extrapolations, the circularity of the filter will effect only the furthest points the extrapolations, and the extrapolations will usually be discarded after the filtering has taken place. However, in many cases, extrapolations and their associated tapering will prove to be unnecessary. A case in point is provided by the filtering of the residual sequence of the logarithmic consumption data that is illustrated by Figure 12.

6.3 Anti-Differencing

After a differenced data sequence has been filtered, it will be required to reverse the effects of the differencing via a process of re-inflation. The process can be conducted in the time domain in the manner that has been indicated in section 4, where expressions have been derived for the initial conditions that must accompany the summation operations.

However, if the filtered sequence is the product of a highpass filter and if the original data have been subjected to a twofold differencing operation, then an alternative method of re-inflation is available that operates in the frequency domain. This method is used in the program only if the filtering itself has taken place in the frequency domain.

In that case, the reduction to stationarity will be by virtue of a centralised twofold differencing operator of the form

$$(1 - z^{-1})(1 - z) = -z\nabla^2(z) \tag{78}$$

The frequency-response function of the operator, which is obtained by setting $z = \exp\{-i\omega\}$ in this equation, is

$$f(\omega) = 2 - 2 \cos(\omega). \quad (79)$$

The frequency response of the anti-differencing operator is $v(\omega) = 1/f(\omega)$.

The matrix version of the centralised operator can be illustrated by the case where $T = 5$:

$$N_5 = \begin{bmatrix} n'_0 \\ -Q' \\ n'_4 \end{bmatrix} = - \begin{bmatrix} -2 & 1 & 0 & 0 & 0 \\ 1 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 1 & -2 \end{bmatrix}. \quad (80)$$

In applying this operator to the data, the first and the last elements of $N_T y$, which are denoted by $n'_0 y$ and $n'_{T-1} y$, respectively, are not true differences. Therefore, they are discarded to leave $-Q' y = [q_1, \dots, q_{T-2}]'$. To compensate for this loss, appropriate values are attributed to q_0 and q_{T-1} , which are formed from combinations of the adjacent values, to create a vector of order T denoted by $q = [q_0, q_1, \dots, q_{T-2}, q_{T-1}]'$.

The highpass filtering of the data comprises the following steps. First, the vector q is translated to the frequency domain to give $\gamma = Uq$. Then, the frequency-response matrix J_η is applied to the resulting Fourier ordinates. Next, in order to compensate for the effects of differencing, the vector of Fourier ordinates is premultiplied by a diagonal matrix $V = \text{diag}\{v_0, v_1, \dots, v_{T-1}\}$, wherein $v_j = 1/f(\omega_j)$; $j = 0, \dots, T-1$, with $\omega_j = 2\pi j/T$. Finally, the result is translated back to the time domain to create the vector h .

The vector of the complementary component is $x = y - h$. Thus there are

$$h = \bar{U} H_\eta U q \quad \text{and} \quad x = y - \bar{U} H_\eta U q, \quad (81)$$

where $H_\eta = V J_\eta$. It should be noted that the technique of re-inflating the data within the frequency domain cannot be applied in the case of a lowpass component for the reason that $f(0) = 0$ and, therefore, the function $v(\omega) = 1/f(\omega)$ is unbounded at the zero frequency $\omega = 0$. However, as the above equations indicate, this is no impediment to the estimation of the corresponding component x .

An example of the application of these procedures is provided by Figure 13, which concerns the familiar logarithmic consumption data, through which a smooth trend-cycle function has been interpolated. This is indistinguishable from the function that is obtained by adding the smooth business-cycle of Figure 12 to the linear trend that was subtracted from the data in the process of detrending it. The program also allows the trend-cycle function to be constructed in this manner.

6.4 Seasonal Adjustment in the Frequency Domain

The method of frequency-domain filtering is particularly effective in connection with the seasonal adjustment of monthly or quarterly data. It enables one




Figure 14. The plot of a seasonally adjusted version of the consumption data of Figures 2 and 13, obtained via the time domain filter




Figure 15. The seasonal component extracted from the U.K. consumption data by a time-domain filter.




Figure 16. The seasonal component extracted from the U.K. consumption data by a frequency-domain filter.

to remove elements not only at the seasonal frequencies but also at adjacent frequencies by allowing one to define a neighbourhood for each of the stop bands surrounding the fundamental seasonal frequency and its harmonics.

If only the fundamental seasonal element and its harmonics are entailed in its synthesis, then the estimated seasonal component will be invariant from year to year. If elements at the adjacent frequencies are also present in the synthesis, then it will evolve gradually over the length of the sample period.

The effects of the seasonal-adjustment filters of the program are illustrated in Figures 14–16. Figure 14 shows the seasonally adjusted version of the logarithmic consumption data that has been obtained via the Wiener–Kolmogorov filter of section 4. Figure 15 shows the seasonal component that has been extracted in the process.

The regularity of this component is, to some extent, the product of the filter. Figure 16 shows a less regular seasonal component that has been extracted by the frequency-domain filter described in the present section. This component has been synthesised from elements at the Fourier frequencies and from those adjacent to them that have some prominence in the periodogram of Figure 10.

7. The Program and its Code

The code of the program that has been described in this paper is freely available at the web address that has been given. This code is in Pascal. A parallel code in C has been generated with the help of a Pascal-to-C translator, which has been written by the author. The aim has been to make the program platform-independent and to enable parts of it to be realised in other environments.

This objective has dictated some of the features of the user interface of the program, which, in its present form, eschews such devices as pull-down menus and dialogue boxes etc. Subsequent versions of the program will make limited use of such enhancements.

However, the nostrum that a modern computer program should have a mode-less interface will be resisted. Whereas such an interface is necessary for programs such as word processors, where all of the functions should be accessible at all times, it is less appropriate to statistical programs where, in most circumstances, the user will face a restricted set of options. Indeed, the present program is designed to restrict the options, at each stage of the operations, to those that are relevant.

A consequence of this design is that there is no need of a manual of instructions to accompany the program. Instead, the three log files that record the steps taken in filtering some typical data sequences should provide enough help to get the user underway. What is more important is that the user should understand the nature of the statistical procedures that have been implemented; and this has been the purpose of the present paper.

References

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

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

Burns, A.M., and W.C. Mitchell, (1946), *Measuring Business Cycles*, New York, National Bureau of Economic Research. Available in electronic form at the address <http://papers.nber.org/books/burn46-1>

Butterworth, S., (1930). On the Theory of Filter Amplifiers. *The Wireless Engineer* (From 1923 to 1930, the journal was called *Experimental Wireless and the Radio Engineer*), 7, 536–541.

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

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.

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

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

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

Henderson, R., (1924). A New Method of Graduation. *Transactions of the Actuarial Society of America*, 25, 29–40.

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

Musgrave, J. (1964). *A Set of End Weights to End all End Weights*, Working Paper, US Bureau of the Census, Washington

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

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

Pollock, D.S.G., (2008). Investigating Economic Trends and Cycles, in Palgrave Handbook of Econometrics: Vol. 2 Applied Econometrics, T.C. Mills and K. Patterson (editors). Palgrave Macmillan Ltd, Houndmills, Basingstoke.

Quenneville, B., D. Ladiray and B. Lefranc, (2003). A Note on Musgrave Asymmetrical Trend-cycle Filters. *International Journal of Forecasting*, 19, 727–734.

Whittle, P., (1983). *Prediction and Regulation by Linear Least-Square Methods*, Second Edition, Basil Blackwell, Oxford.