econometrics

MDPI

Article

The Manual for SEASCAPE A Program for Seasonal Adjustment

D. Stephen G. Pollock

Department of Economics, University of Leciceter, Leciceter LE1 7RH, U.K.; stephen_pollock@sigmapi.u-net.com

Received: 29 September 2020; Accepted: 24 December 2020; Published: 5 January 2021

Abstract: This manual should serve as a ready reference to assist in operating the *SEASCAPE* program, which provides some enhanced methods for the seasonal adjustment of economic data. The manual lists the menus of the program, and it gives a thematic account of the facilities of the program. Embedded in the program are brief descriptions of its facilities and its functions. These should also provide guidance for operating the program. The table of contents of this manual contains hypertext links to the sections and subsections.

Contents

1. INTRODUCTION

The *SEASCAPE* program implements two alternative sets of methods for the seasonal adjustment of economic data. The methods operate either in the time domain or in the frequency domain. (See Figure 1.) Some enhanced methods are offered that are intended for the seasonal adjustment of economic data that contain irregular seasonal fluctuations that evade the conventional methods.

1.1. The Seasonal Frequencies

The seasonal fluctuations that affect the data can be construed as a combination of trigonometric functions of varying amplitudes and with frequencies that are in the vicinity of the annual frequency and it harmonic frequencies, which are integer multiples of the annual frequency. Taken together, the annual season frequency and its harmonics may be described as the seasonal frequencies.

The seasonal frequencies are expressed as numbers of degrees or radians traversed by the functions in the period between two successive observations of the data. If the observations are monthly, then the fundamental seasonal frequency has a value of $360^\circ/12 = 30^\circ$ degrees and the harmonic frequencies are its integer multiples, which are $60^\circ, 90^\circ, \ldots, 180^\circ$. In terms of radians, the fundamental monthly frequency is $\pi/6$ and its harmonics are at $\pi/3, \pi/2, \ldots, \pi$. If the observations are quarterly, then the fundamental frequency is at 90° ($\pi/2$ radians) and its harmonic is at 180° (π radians).

The limiting frequency of 180° degrees, or π radians, is described as the Nyquist frequency; and it corresponds to one cycle in every two observations. The interval from zero to the Nyquist frequency is described as the Nyquist range. The Nyquist frequency is the highest frequency that can be detected in sampled data. In a process that is described as aliasing, any elements of the observed signal with frequencies that are in excess of the Nyquist value will by confounded with elements of frequencies that lie within the Nyquist range.

1.2. The Time Domain and the Frequency Domain

The first set of the methods of seasonal adjustment that are provided by the program operate in the time domain, in which the amplitudes of the observations are considered to be functions of time. For sampled data, time is represented by a sequence of integer values. Time-domain methods employ filters that are either ordinary finite moving averages of the data, described as Finite Impulse Response (FIR) filters, or moving averages supplemented by feedback terms, which are described as Infinite Impulse Response (IIR) Filters.

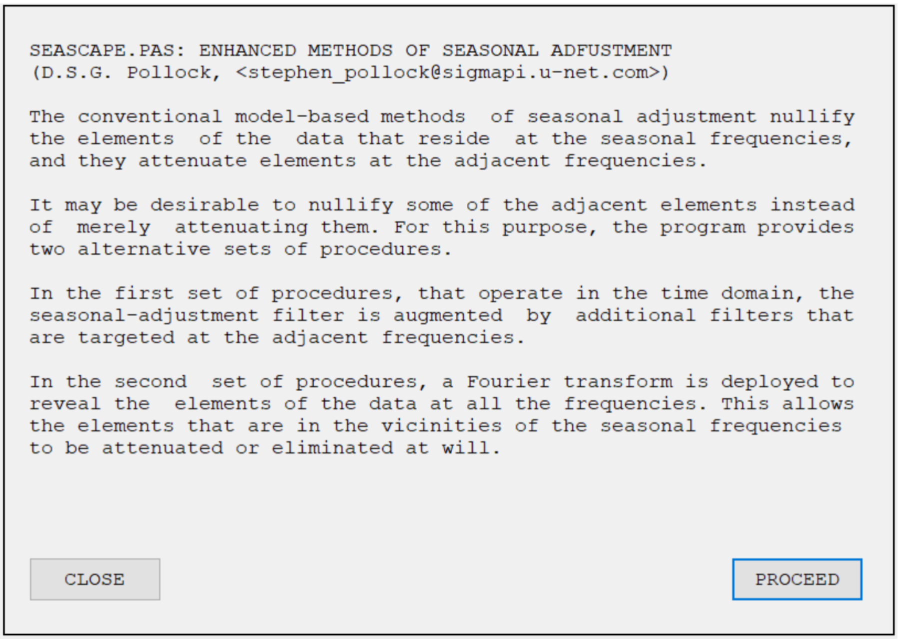

Figure 1. The title page of the *SEASCAPE* program.

The second set of methods operate in the frequency domain, in which attention is focussed on the amplitudes of the trigonometric functions of which the data are composed. In a Fourier analysis of a finite data sequence, there are as many trigonometric functions as there are data points. The observed data can be synthesised from these functions by combining them and by sampling the combination at the integer points in time. The frequency-domain methods of seasonal adjustment operate by modifying or by nullifying the amplitudes of the constituent trigonometric functions, or of the Fourier elements, as they are liable to be described.

A Fourier analysis, which resolves the data into its trigonometric elements, is a natural accompaniment of the methods of seasonal adjustment, whether they operate in the time domain or in the frequency domain; and it will enable the user to judge whether the methods are well targeted at the seasonal fluctuations. To this end, the program displays the frequency response functions of the time-domain and the frequency-domain filters, which show the extent to which they alter the amplitudes of the Fourier elements.

1.3. Enhanced Methods of Seasonal Adjustment

The effect of the conventional methods of seasonal adjustment, which operate in the time domain, is to nullify the amplitudes of the elements at the seasonal frequencies and to attenuate the elements at the adjacent frequencies to an extent that diminishes with the distance from the seasonal frequencies.

If the seasonal fluctuations are of an irregular nature, then they are liable to comprise Fourier elements at the adjacent frequencies that may need to be eliminated in part or in in full from the seasonally adjusted data. This is the purpose of the enhanced methods of seasonal adjustment that are provided by the program.

The enhanced time-domain methods of the program employ multiple filters that are targeted at the seasonal frequencies and at adjacent frequencies. The double filter eliminates elements that are offset, to a fixed number of degrees, at either side of the seasonal frequencies.

The fact that the seasonal frequencies are attenuated rather than totally eliminated may be a disadvantage. Therefore, a triple filter is provided that is targeted at the seasonal frequencies as well as frequencies that are offset on either side. This creates a method that is both more complicated and more flexible than that of the double filter.

The enhanced frequency-domain methods of the program exceed the flexibility of the time-domain methods. They allow the widths of the stop bands on either side of the seasonal frequencies, in which the elements are completely eliminated, to be specified at will. They also allow transition bands with various profiles to be specified that are located on either side of the combined stop band.

1.4. Trended Data

Whereas the theory of filtering has been developed mainly in the context of stationary stochastic processes that have a constant mean value and a constant variance, the typical econometric data sequence will manifest a strong trend and a variance that is systematically related to the level of the data. Both of these effects create difficulties that must be overcome before the filters can be applied successfully.

The inflation of the variance, which is commonly described as heteroskedasticity, may be overcome by taking the logarithms of the data. The effect is that a multiplicative combination of the components of the data becomes a linear combination of their logarithms.

The program removes the trend from the data by interpolating a polynomial trend function. The residuals of the polynomial regression will constitute a sequence of zero mean. By applying a lowpass filter to the residuals, a smooth quasi-cyclical function is extracted and added back to the polynomial to create a trend-cycle function.

The program extracts the seasonal component of the data equally from the deviations of the data from the polynomial trend or from their deviations from an estimated trend-cycle function. The choice of which deviations to employ has no bearing on the extraction of the seasonal component, since the low-frequency elements of the polynomial trend or of the trend-cycle function are separate from the seasonal elements.

When the time-domain filters are applied to the polynomial residuals, the program derives a twofold decomposition of the data comprising a seasonal component and a seasonally adjusted data sequence formed by subtracting the seasonal component from the original data.

When the frequency-domain filters are applied to the polynomial residuals, the program will create a similar twofold decomposition. When they are applied to the residual deviations from the trend-cycle trajectory, the frequency-domain filters will give rise to a threefold decomposition of the data, comprising the seasonal component, the trend-cycle function and an irregular noise component.

1.5. The Realisation of the Program

The program has been created using the *Delphi* and the *Lazarus/Free Pascal* compilers, which adopt an object-orientated model. Its functions and procedures are accessed via a set of linked pages or forms, the first of which is displayed as Figure 1. At the top of each form is a brief text, which describes the procedures to which it gives access, via list boxes, edit boxes and radio buttons. There are no pull-down menus, which would be liable to provide the user with a set of unexplained options. Instead, the intention is to guide the user down well-defined paths to the end products, which are liable to be graphical representations of the components extracted from the data. The graphs that are rendered on the computer screen can also be saved as *PostScript* code.

For many of the procedures, pre-requisite information must be supplied or prior choices must 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 have 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 again 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.

It is expected that a user who is sufficiently apprised of the business of seasonal adjustment should be able to use the program without recourse to this manual, which gives a thematic account of the facilities of the program. It should serve as a ready reference.

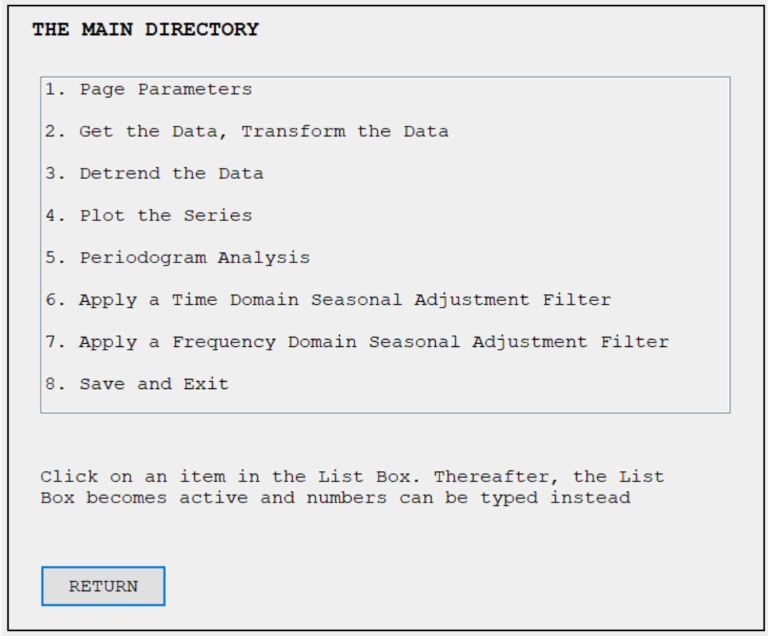

Figure 2. The main directory of the program.

The program and a copy of this manual reside is a zip file at the following address:
//http://www.sigmapi.org.uk/seascape.zip/

The code of the program, which is in Pascal, can be found within the zip file, which also contains a collection of data.

2. THE MENUS OF THE PROGRAM

The following is a summary of the menus and sub-menus of the program. (See Figure 2 for the main directory of the program.) Most of the items of the main menu of *SEASCAPE.PAS* subsume a sub-menu. The capitalised headings are the titles of the program's forms. The secondary sub-menus are not listed:

THE MAIN DIRECTORY
- *1. Page Parameters*
- *2. Get the Data, Transform the Data*
- *3. Detrend the Data*
- *4. Plot the Series*
- *5. Periodogram Analysis*
- *6. Apply a Time-Domain Seasonal-Adjustment Filter*
- *7. Apply a Frequency-Domain Seasonal-Adjustment Filter*
- *8. Save and Exit*

PAGE PARAMETERS
- *Specify the Width*
- *Specify the Height*
- *Generate the PostScript Code*
  - *Encapsulated PostScript*
  - *Textures Format*

GET THE DATA
- *1. Read the Data*

2. Plot the Data
3. Transform the Data
4. Discard Some Data
5. Polynomial Regression
6. Plot The Periodogram

DATA PARAMETERS
Data Frequency
a. Annual
q. Quarterly
m. Monthly
n. Other
Name the Data
Specify the Start Date
Beginning Year
Beginning Quarter
Beginning Month

TRANSFORM THE DATA
1. Subtract the Mean from the Data
2. Take Differences
3. Take Natural Logarithms
4. Take Anti-logarithms

POLYNOMIAL REGRESSION
Fit by Ordinary Least-Squares Regression
Fit by Weighted Least-Squares Regression

WEIGHED LEAST-SQUARES POLYNOMIAL REGRESSION
Weighting Functions
Uniform Function
Cosine Bell Function
Quadratic Function
Crenellated Function

DETREND THE DATA
1. Polynomial Regression
2. Time-Domain Filters
3. Variable Smoothing
4. Frequency-Domain Filters
5. Extrapolate and/or Taper the Data

PLOT THE DATA
d. Plot the Polynomial Residuals
r. Plot the Trend-cycle Residuals
s. Plot the Seasonal Component
a. Plot the Seasonally Adjusted Data
e. Plot the Adjusted Trend Residuals

PERIODOGRAM ANALYSIS
d. Plot the Periodogram of the Data
t. Plot the Periodogram of the Trend-Cycle
c. Plot the Periodogram of the Cycle
s. Plot the Peridogram of the Seasonal Component
r. Plot the Periodogram of the Polynomial Residuals
r. Plot the Periodogram of the Trend-Cycle Residuals
a. Plot the Peridogram of the Seasonally Adjusted Data
e. Plot the Peridogram of the Seasonally Adjusted Trend Residuals

TIME-DOMAIN SEASONAL-ADJUSTMENT FILTERS
The Ordinary Filter
The Filter with Smoothing
The Twopass Offset Filter
The Triple Offset Filter

FREQUENCY-DOMAIN SEASONAL-ADJUSTMENT FILTERS
Transition Profiles
No Transitions
Composite Sigmoid
Upper-half Cosine
Lower-half Cosine
Profile Parameters
Truncation Parameter
Power Parameter

SAVE AND EXIT
d. Save the Data
t. Save the Trend-Cycle Component
s. Save the Seasonal Component
r. Save the Polynomial Residuals
r. Save the Trend-Cycle Residuals
a. Save the Seasonally Adjusted Data
e. Save the Seasonally Adjusted Trend Residuals

3. DATA HANDLING

3.1. Data Entry

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.

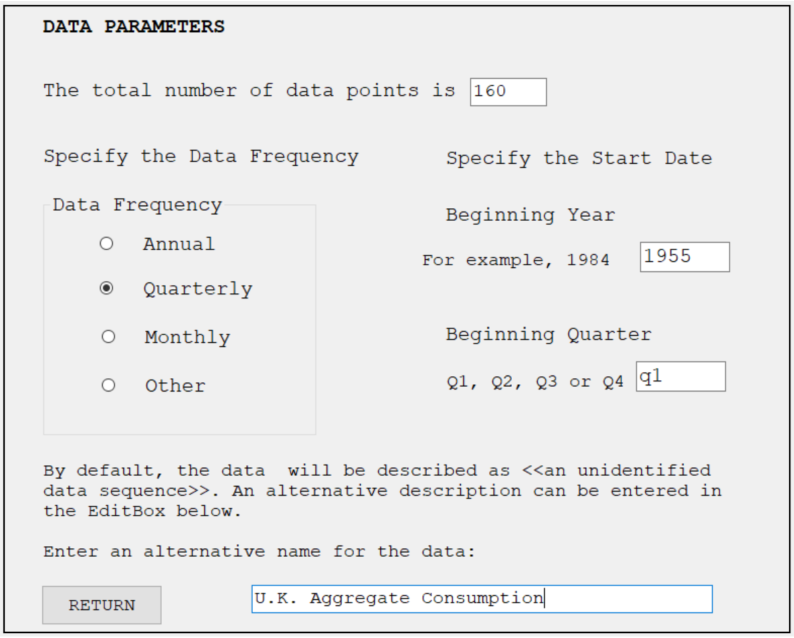

Figure 3. The data parameters for quarterly data.

The Headline of the Data File

The first four 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 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 sequence. (See Figure 3.) 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 attached to the *PostScript* code that can be generated while plotting graphs and diagrams on the screen.

3.2. Data Plotting

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* (.eps) 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

The program allows the user to discard some of the data that have been read. First, the data will be plotted to allow their 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 or their percentage of the total must the specified; whereafter the shortened data sequence can be plotted.

3.4. Data Transformations

The data sequence is held permanently in the memory of the computer. However, several transformations are provided that will have an irreversible effect on the data. To recover its original state, the data sequence must be read afresh.

Subtract the Mean from the Data

The program automatically determines whether the data have a significant trend and whether their mean value differs significantly 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, which will create a sequence of residual deviations of zero mean, which will be held in memory. 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.

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

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 using weighted least-squares regression to remove the trend in the manner described in section 7. The extra weight given to the points at the ends of the sequence will serve to minimise their deviations from the trend. When there are seasonal fluctuations, the data should cover an integral number of years.

The objective of avoiding discontinuities in the circularised data can also be served by extrapolating the data at both ends with segments that converge to zero, or by interpolating a segment between the end of the final year of the data and the beginning of the first year that makes a gradual transition between the seasonal patterns of these two years. (Figure 4 shows the selection of one of these options.)

These facilities are available only in the context of the estimation of a trend-cycle function, which requires the lowpass filtering of a sequence of polynomial residuals to extract a cycle that is added back to the polynomial. They can be by-passed by clicking on the CONTINUE button.

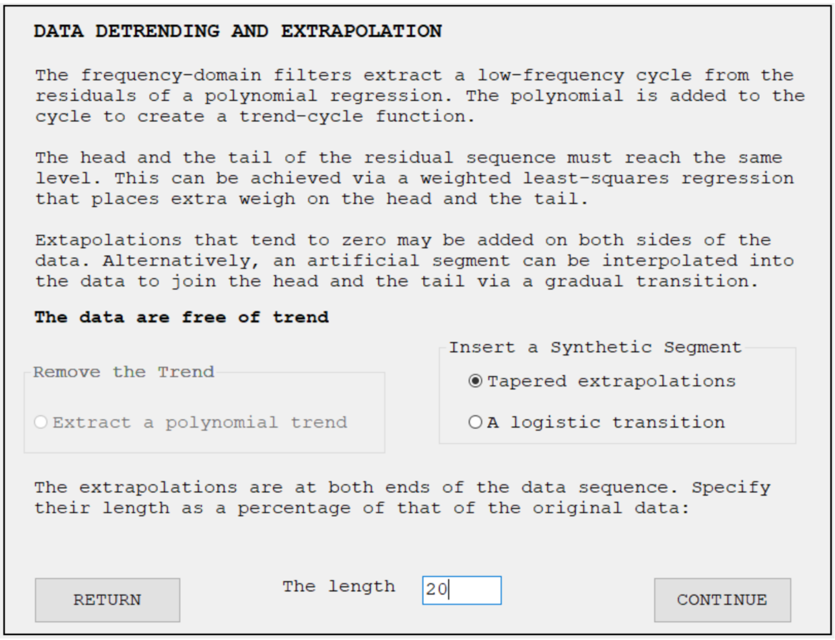

Figure 4. Data de-trending and extrapolation.

Extrapolating and Tapering the Data

The de-trended data can be extended at both ends by a certain proportion of the original length. Then, only the extensions are tapered. 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.

Once the extrapolations have been used in estimating the cyclical component of a trend-cycle function, they will be discarded.

Interpolation of a Logistic Transition

The process of circular morphing works only with de-trended 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.

The transition between the two seasonal patterns is governed by a logistic function, which is formed from the segment of a raised cosine function over the interval that runs from 0° to 180°. the duration of the logistic transition must be specified as an integral number of years.

Once the extension has been used in estimating the cyclical component of a trend-cycle function, it will be discarded.

5. GRAPH PLOTTING

5.1. 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.)

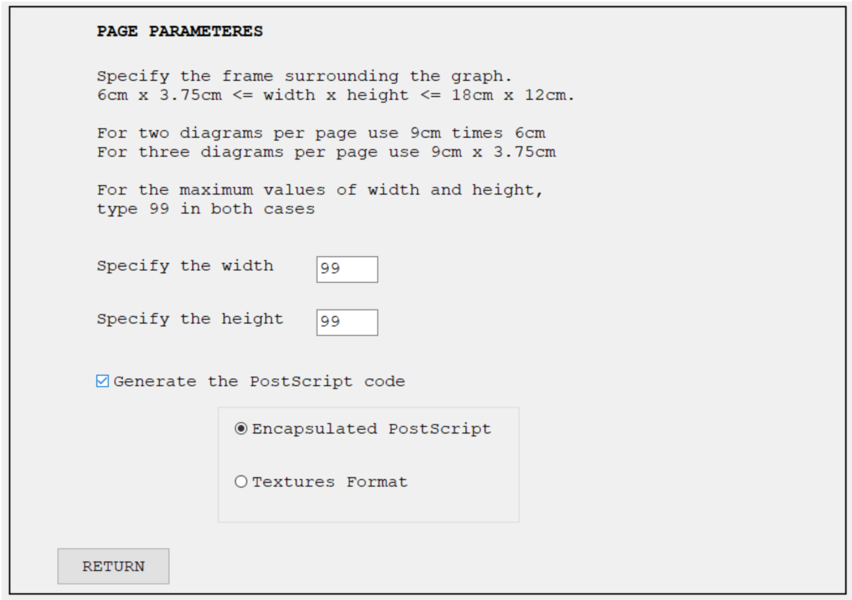

Figure 5. The page parameters form, specifying a graph of maximum size.

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 6 cm to a maximum of 18 cm and the available height is from 3.75 cm to 12 cm. (See Figure 5.)

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. (See Figure 5, where this choice has been made.)

A plotted graph can be dismissed pressing the *RETURN* button.

5.2. Plot The Series

The program will plot the primary data sequence that is held in memory or any other sequence that has been derived in the process of trend estimation or filtering. The following operations will be available, albeit that only one of those listed under *r.* will be available at any one time:

d. Plot the Data
r. Plot the Polynomial Residuals
r. Plot the Trend-Cycle Residuals
c. Plot the Cycle
t. Plot the Trend-Cycle
s. Plot the Seasonal Component
a. Plot the Seasonally Adjusted Data
e. Plot the Adjusted Trend Residuals

6. DATA CHARACTERISTICS

6.1. The Periodogram

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.

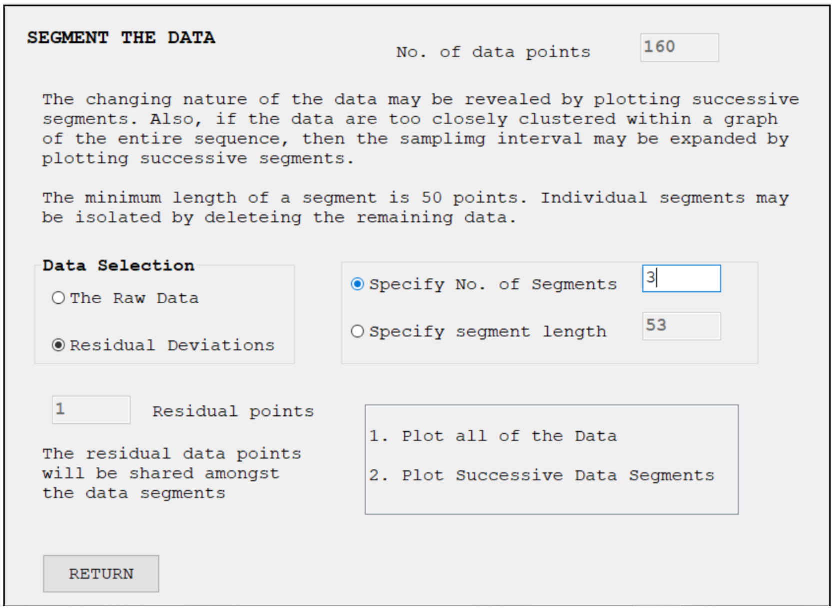

Figure 6. The form for segmenting the data.

An analysis of the periodogram of the de-trended data is the principal means of assessing the requirements for an effective seasonal adjustment of the data. The periodogram will reveal the extent of the elements that are contributing to the seasonal fluctuations and which require to be attenuated or removed from the data.

The prior removal of the trend is necessary, since the periodogram of a trended sequence will be dominated by the elements in the vicinity of the zero frequency, to the extent that the elements at other frequencies may be barely visible. Nevertheless, the program allows the periodogram of the original data, of its seasonally adjusted version and of the estimated trend-cycle function to be displayed. The following commands comprise all of the periodograms that can be plotted, at various stages, by the program, albeit that only one of those listed under *r.* will be available at a time:

d. Plot Periodogram of the Data
t. Plot the Periodogram of the Trend-Cycle
c. Plot the Periodogram of the Cycle
s. Plot the Peridogram of the Seasonal Component
r. Plot the Periodogram of the Polynomial Residuals
r. Plot the Periodogram of the Trend-Cycle Residuals
a. Plot the Peridogram of the Seasonally Adjusted Data
e. Plot the Peridogram of the Seasonally Adjusted Trend Residuals

6.2. Data Segmentation

The evolving nature of the data may be revealed by plotting successive segments. (See Figure 6.) Also, if the data are numerous and too closely clustered within a graph of the entire sequence to be individually identifiable, then the sampling interval may be expanded by plotting successive segments. The minimum length of a segment is 50 points.

Either the raw data or the residual deviations from a polynomial trend may be plotted. Either the number of the segments or their length, subject to the restriction on the minimum length, may be specified. Left-over points will be shared amongst the data segments, if their number has been specified, or else they will be added to the final segment, if the length of the segments has been specified.

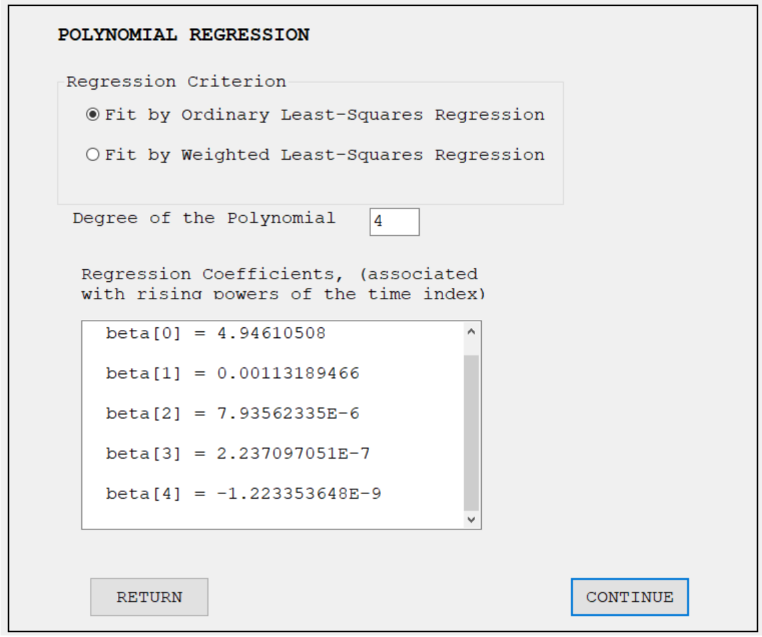

Figure 7. The form displaying the results of a least-squares polynomial regression.

7. TRENDS AND BASELINES

A polynomial regression is the fundamental means by which the program removes a trend from the data. Unless the program has determined that the data are already free of trend, then a polynomial regression is a usual precondition for applying a frequency-domain filter or for pursuing a seasonal adjustment of the data.

The exceptions are the time-domain Wiener–Kolmogorov filters, which can be applied directly to trended data to estimate trend-cycle functions. They deliver residual sequences of zero mean that are free of trends. A polynomial regression will give rise to a periodogram of the residual sequence, which provides a means of determining the frequency range of the elements that contribute to the trend-cycle trajectory.

7.1. 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 a series coefficients associated with powers of the temporal index. (See Figure 7, which shows such coefficients for a polynomial of degree 4.)

Weighted Polynomial Regression

The purpose of a weighted least-squares 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.

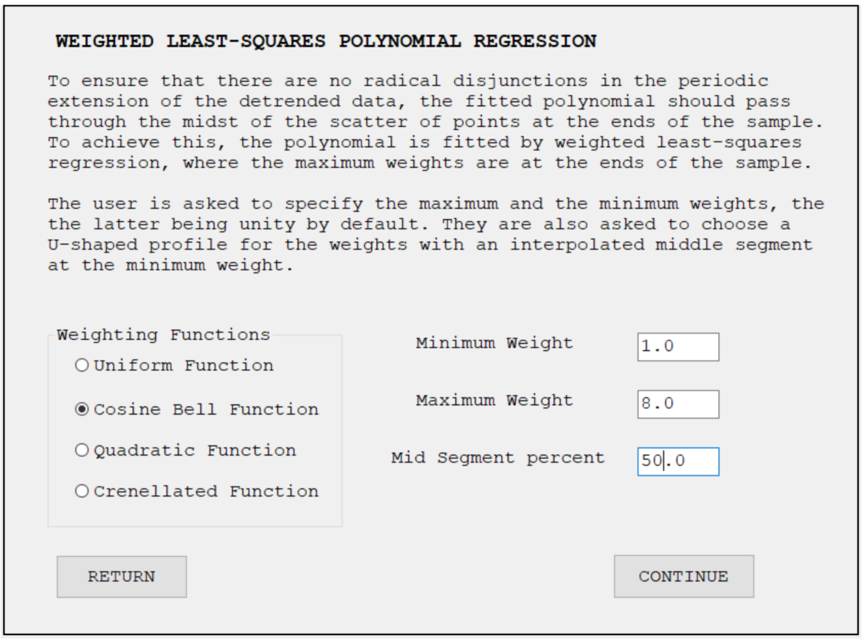

Figure 8. The data parameters for quarterly data.

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. Four choice are offered—see Figure 8.

Two such choices are a quadratic function and a crenellated function that jumps abruptly from the minimum value, which is maintained over an interval, described as the mid segment, which is 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 function by interpolating a constant segment into an interval centred on the mid point. The equivalent option in respect of the crenelated function is to allow variations in the length of the segment over which the weighting function assumes its minimum value.

7.2. Interrupted Trends

It is possible to create a trend function that absorbs prominent breaks in the data, while maintaining a smooth profile elsewhere.

To this end, 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. (See Figure 9, where two such sub intervals have been specified.) 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.

A restricted sub-interval is specified by declaring the index of a *Central point* and by specifying a number of data points above and below that point. These parameters can be determined by viewing a plot of the data and by specifying the points via successive adjustments until an appropriate interval has been located. The interval will be marked by a hatched bar superimposed on the plot of the data.

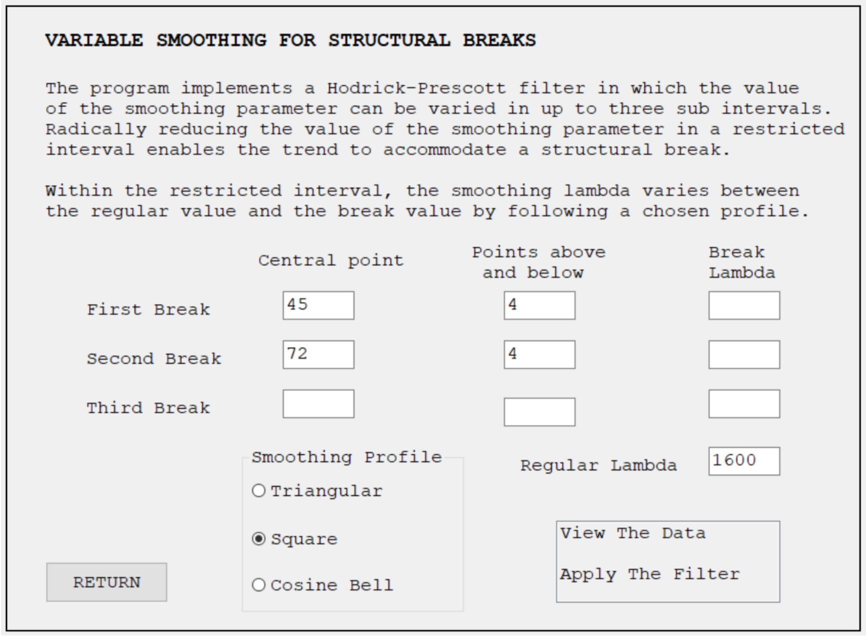

Figure 9. The form for specifying the intervals for adapted smoothing.

The Smoothing Parameters

Within a restricted sub-interval, the smoothing parameter *lambda* varies between the value of the *Regular Lambda,* which would prevail in the absence of the breaks, and the value of the *Break Lambda.* A profile must be chosen to govern this variation, which may be a Triangular Profile, a Square Profile or a Cosine Bell Profile. The Regular Lambda can be determined in view of the conventional recommendations for a Hodrick–Prescott filter, which are listed in section 8.1. The Break Lambda may be given a value in single digits, including zero.

The Application of the Filter

The filter may be applied without specifying restricted sub-intervals or break parameters. Only the regular smoothing lambda is required in order to apply the filter. In that case, the Hodrick–Prescott filter will determine a smooth trajectory that will be displayed as a red line superimposed of the data. On closing the graph, a plot of the residual deviations the data from this trajectory will be displayed. This plot will serve to highlight the breaks that may be absorbed, thereafter, in a modified trajectory, which will also be marked by a red line.

8. TREND-CYCLE ESTIMATION

The program provides lowpass versions of the Hodrick–Prescott and the Butterworth filters, which may be used in estimating a trend-cycle function. (Figure 10 shows the selection of the Butterworth filter.) These Wiener–Kolmogorov filters, which operate in the time domain, may be applied directly to the original data sequence or to its logarithms.

If a polynomial function has been interpolated into the trended data, then the periodogram of the residual deviations will be plotted in the background of the graphs of the frequency response functions of the filters. This should assist in determining appropriate parameters for the filters.

The program also provides lowpass filters for estimating a trend-cycle function that operate in the frequency-domain. These are to be applied to the residual deviations from a polynomial trend function to create estimates of the cyclical component. The estimates are added back to the polynomial to create a trend-cycle trajectory.

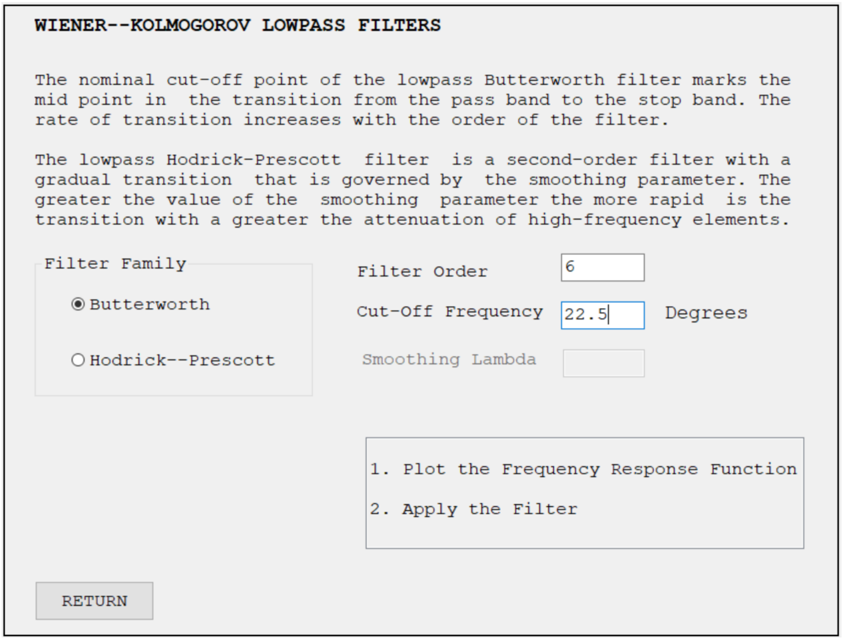

Figure 10. The Wiener–Kolmogorov filters and their parameters.

8.1. *The Leser or Hodrick–Prescott Filter*

The lowpass filter of Hodrick and Prescott (1997), which is also attributable to Leser (1961), has a single adjustable smoothing 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:

100 *for annual data*
$1,600$ *for quarterly data*
$14,400$ *for monthly data*

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 with 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 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, superimposed and marked in red. When this graph has been dismissed, the program will display a graph of the residual deviations of the data from the trend-cycle component.

8.2. *The 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 of the filter, and the filter order, which governs the rate of transition.

An appropriate cut-off point may be determined by inspecting the periodogram of the residual deviations from an interpolated polynomial. 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 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.

8.3. The Ideal Frequency-Domain Filter

The ideal lowpass filter makes an instantaneous transition from the pass band to the stop band. The filter preserves the Fourier elements that fall within the stop band, without altering their amplitudes, and it nullifies the elements that fall in the stop band. The only parameter that requires to be specified is the cut-off point, which is denominated in degrees, and which marks the boundary between the bands.

The Frequency Response of the Filter

The frequency response is displayed superimposed on the periodogram of the de-trended data, which are the residual deviations from an interpolated polynomial. This should assist in determining an appropriate cut-off frequency.

The Application of the Filter

When the filter is applied to the sequence of polynomial residuals, a graph is displayed of the filtered sequence, which contains low-frequency cyclical fluctuations, superimposed on the residuals as a red line. Thereafter, on closing the graph, a trend-cycle function is displayed, which is obtained by adding the low-frequency cycle to the trajectory of the de-trending polynomial. This is superimposed on the original data, as a red line.

8.4. Sigmoid Frequency-Domain Filter

The sigmoid filter makes the transition from pass band to the stop band within a transition band of which the bounds are specified by the user. The trajectory of the transition is determined by a logistic function which is anti-symmetric about its midpoint (i.e. invariant with respect to a rotation of 180°) and which resembles the segment from 0° to 180° of the cosine function.

Having determined the upper and lower bounds of the transition band, the user can confirm the value of its mid point by clicking on the corresponding label. (See Figure 11.)

The slope of the transition function is governed by the power parameter, which takes a real value in the interval [1, 6] and which has been given a default value of 3. Increasing the value of the parameter delays the onset of the transition and it steepens its decent.

The Frequency Response of the Filter

The frequency response is displayed superimposed on the periodogram of the de-trended data, which are the residual deviations from a interpolated polynomial. The bounds of the transition band are plotted on the graph.

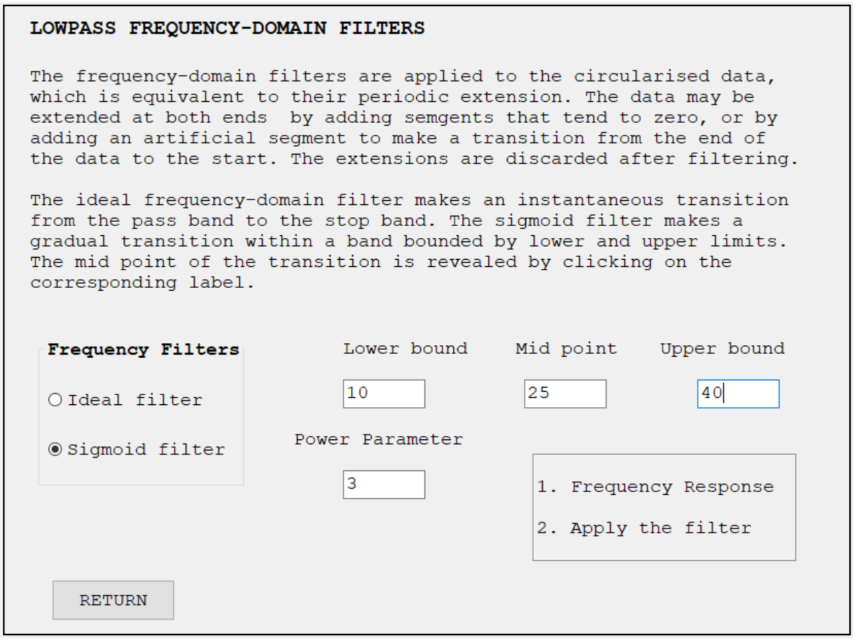

Figure 11. The lowpass frequency-domain filters and their parameters.

The Application of the Filter

The filter is applied to the residuals from a polynomial regression. The filtered sequence is displayed in a graph as a red line superimposed on the residuals. On closing the graph, a trend-cycle function is created by adding the filtered residuals to the polynomial function. It is displayed as a red line superimposed on the original data sequence.

9. SEASONAL ADJUSTMENT IN THE TIME DOMAIN

The program provides a variety of filters for the seasonal adjustment of the data that operate in the time domain, including a basic filter that is designed to imitate the affects of the conventional model-based filters that are exemplified, for example, by the *SEATS–TRAMO* program—see Gómez and Maravall (1997, 2001) and also Grudkowska (2017), who describes a re-engineered version of the program.

The filters operate on the residuals from a polynomial de-trending. They nullify the seasonal component within the residual sequence, leaving a vector of seasonally adjusted residuals. The sequence of seasonal fluctuations is the complement, within the sequence of polynomial residuals, of the sequence of seasonally adjusted residuals. Subtracting it from the data sequence creates the seasonally adjusted version.

In effect, time-domain filters effect a twofold decomposition of the data to create a sequence of seasonal fluctuations and a sequence of seasonally adjusted data.

The basic filter eliminates the elements at the seasonal frequencies and it attenuates the adjacent elements to an extent that decreases with the distance from the seasonal frequencies.

The enhanced time-domain filters are intended to suppress the adjacent elements more firmly. In the case of the two-pass filter, this is achieved by eliminating elements at the adjacent frequencies on either side of the seasonal frequencies. In the case of the triple filter, the user targets the seasonal frequencies together with adjacent frequencies on either side that are at distances that can be determined individually in every instance.

A trend-cycle function is created by applying a smoothing filter to the seasonally adjusted residuals created by the basic filter. Adding the resulting sequence back to the polynomial function creates a trend-cycle function. This represents an attempt to mimic one of the products of a conventional method of seasonal adjustment.

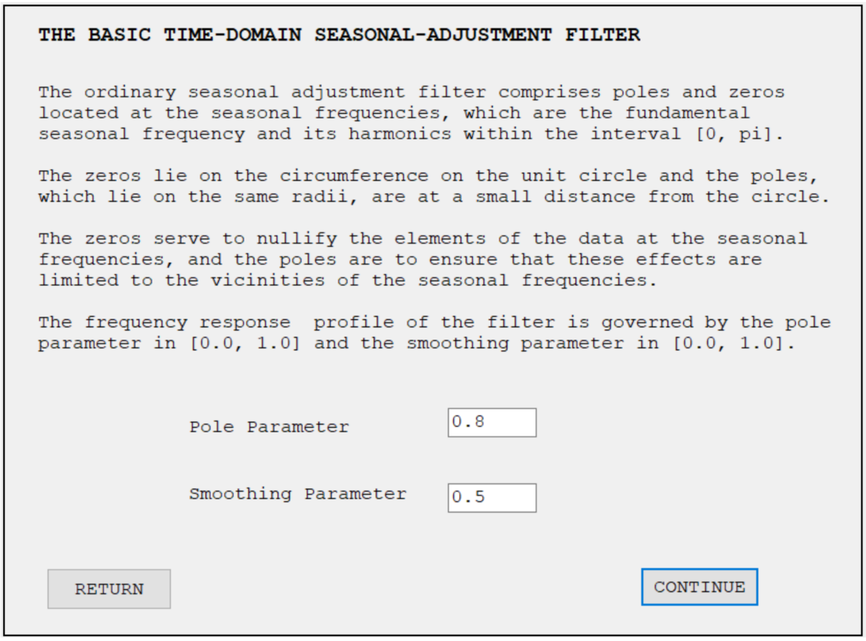

Figure 12. The parameters of basic time-domain seasonal-adjustment filter.

9.1. The Basic 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 pole parameter $\rho \in (0, 1)$, which is the modulus of the poles of the filter, is the primary means of adjusting the width of the notches. The notches become narrower as $\rho \to 1$. A default value of 0.8 has been has been given to this parameter. The poles serve to counteract the effects of the zeros at frequencies that lie between the seasonal frequencies.

The width of the notches is also affected by a smoothing parameter $\lambda \in (0, 1)$, and they become narrower as $\lambda \to 0$. A default value of 0.5 has been given to this parameter. (See Figure 12.)

The user is constrained to remove any trend in the data by a polynomial regression.

The Frequency Response of the Filter

The program displays the frequency response of the filter superimposed on the periodogram of the residual sequence.

The Application of the Filter

On closing the graph that displays the frequency response of the filter, a plot of the seasonal component that has been extracted from the polynomial residuals is displayed. This is followed by a graph that shows the seasonally adjusted data, represented by a red line, superimposed on original data sequence.

9.2. The Basic Seasonal-Adjustment Filter with Smoothing

A trend-cycle function, resembling a function generated by the conventional seasonal-adjustment procedures, is created by applying a lowpass smoothing filter to the seasonally adjusted residuals, generated by the basic filter, before they are added back to the polynomial function that has served to de-trend the data. A binomial filter is employed in smoothing the seasonally adjusted residuals.

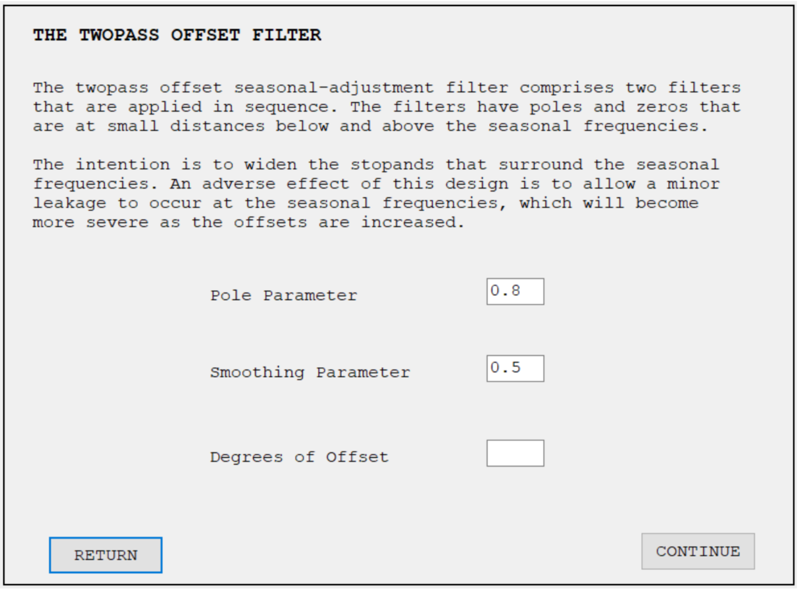

Figure 13. The parameters of the two-pass seasonal-adjustment filter.

The combination, in series, of the basic filter and the binomial smoothing filter may be described as a tapered filter

Three parameters require to be specified. In addition to the smoothing parameter and the pole parameter, which are given default values of 0.5 and 0.8 respectively, the user must specify the span of the binomial filter defined as the total number of coefficients, which must be an odd number. The greater the span, the greater is the smoothing effect.

A trend-cycle function, comprising elements that fall short of the fundamental seasonal frequency, can be created by the frequency-domain methods of section 8. Such methods are liable to be preferred.

The Frequency Response of the Filter

The frequency response of the smoothing binomial filer is displayed. This is compounded with the frequency response of the ordinary seasonal-adjustment procedure to produce the frequency response of a tapered filter, to be used in estimating the trend-cycle function. This response is displayed superimposed on the periodogram of the residuals of the polynomial regression, which will have been used beforehand to remove the trend from the data.

The Application of the Filter

The program will display, in succession, the frequency response of the binomial filter and that of the tapered filter. The latter will be superimposed on the periodogram of the residual sequence from a polynomial de-trending of the data.

Thereafter, the program will apply the tapered filter to the residual sequence to create a cyclical component, which is displayed in red superimposed on the residual sequence. Subtracting this from the residual sequence will create a seasonal component, which is also displayed. Adding the cyclical component to the polynomial creates the trend-cycle function of the final display. This may be regarded as a substitute for the seasonally adjusted data.

9.3. The Two-pass Offset Seasonal-Adjustment Filter

The two-pass offset seasonal-adjustment filter comprises two filters that are applied in sequence. The filter has poles and zeros that are at small distances below and above the seasonal frequencies.

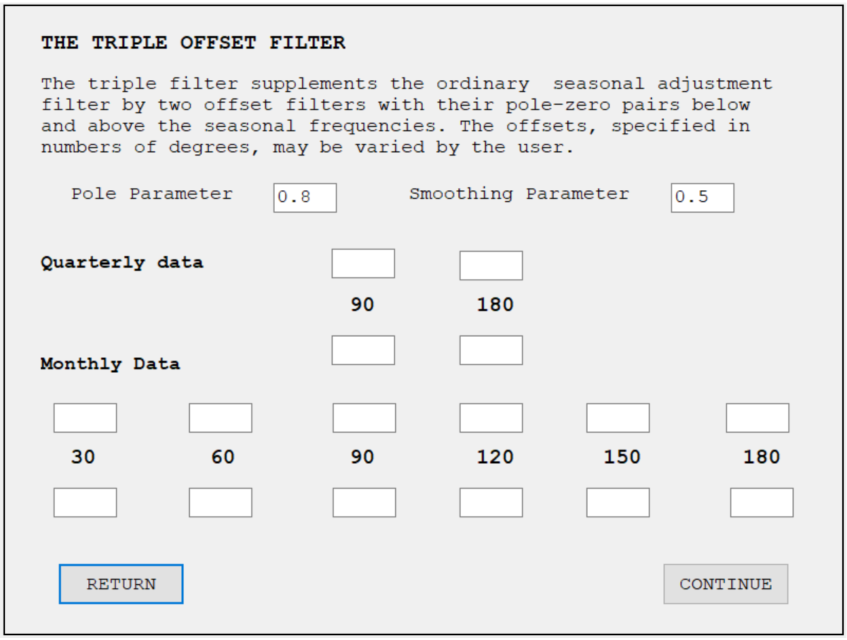

Figure 14. The form for locating the zeros of the three-pass seasonal-adjustment filter. The form serves both for monthly and for quarterly data.

The intention is to widen the stop bands that surround the seasonal frequencies. An adverse effect of this design is to allow a minor leakage to occur at the seasonal frequencies, which will become more severe as the offsets are increase.

Three parameters require to be specified. In addition to the smoothing parameter and the pole parameter, which are given default values of 0.5 and 0.8 respectively, the user must specify the offset of the filters as a number of degrees. (See Figure 13.)

The same offset is applied to both filters in respect of all of the seasonal frequencies. The first filter is offset below the seasonal frequencies and the second filter is offset above the seasonal frequencies. The maximum allowable offset is 4° degrees.

The Frequency Response of the Filter

The program displays the frequency response of the filter superimposed on the periodogram of the residual sequence. The leakage of the filter at the seasonal frequencies may be noted.

The Application of the Filter

After displaying, the frequency response of the twofold filter, superimposed on the periodogram of the residual sequence from a polynomial de-trending of the data, the program will display the seasonal component that has been extracted from the residuals. Subtracting the seasonal component from the original data sequence will create a seasonally adjusted sequence, which is plotted in red superimposed on the data.

9.4. The Triple Offset Seasonal-Adjustment Filter

The triple filter supplements the ordinary seasonal adjustment filter by two offset filters with their pole–zero pairs below and above the seasonal frequencies. The offsets above and below each of the seasonal frequencies may be specified individually in numbers of degrees. (Figure 14 shows the form displaying the cells for entering these values for both monthly and quarterly data. In practice, the pre-existing choice of the data frequency will deactivate one set of cells. The remaining cells will be filled by default values.)

The default values for the offsets, which are provided in every case, are liable to be changed by the user in view of the periodogram of the de-trended data, which are the residuals from a polynomial regression. The objective is to create stop bands and clefts surrounding the seasonal frequencies that are wide enough to encompass all of the elements that contribute to the seasonal fluctuations.

In addition to the values of the offsets, the user must specify the smoothing parameter and the pole parameter, which are given default values of 0.5 and 0.8.

The Frequency Response of the Filter

The program displays the frequency response of the filter superimposed on the periodogram of the residual deviations from the interpolated polynomial. The graph should indicate whether the clefts in the frequency response will be sufficient to eliminate the seasonal fluctuations from the data.

The Application of the Filter

After displaying the frequency response of the filter, the program proceeds to display the estimated seasonal component that has been extracted from the residuals. The seasonally adjusted data are obtained by subtracting the elements of the seasonal component from the original data. It is displayed in a graph as a red line superimposed on the data.

10. SEASONAL ADJUSTMENT IN THE FREQUENCY DOMAIN

The methods of seasonal adjustment that operate in the frequency domain effect either a twofold or a threefold decomposition of the data.

The twofold decomposition gives rise to a sequence of seasonal fluctuations and a sequence of seasonally adjusted data. It arises when a polynomial regression is used in extracting a trend from the data, prior to the extraction of the seasonal component.

The threefold decomposition gives rise to a sequence of seasonal fluctuations, a trend-cycle trajectory and a sequence of residual disturbances. It entails a prior estimation of the trend-cycle function.

There is no spectral overlap between the trend component or the trend-cycle component and the seasonal component. Therefore, identical seasonal components are extracted from the residual deviations from a polynomial regression and from the residual deviations from a trend-cycle function.

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 polynomial trend function is interpolated and the deviations from the trend are used in place of the original data. To create a trend-cycle function, a lowpass filter is applied to the residuals of the polynomial regression and the resulting sequence, designated the cyclical component, is added back to the polynomial.

The frequency-domain filters impose stopbands surrounding the seasonal frequency and its harmonic frequencies. The stop bands may be bordered by transition bands. The width of the bands is a matter of choice. The profiles of the transition bands can also be chosen from amongst several options, after which their widths can been chosen. (Figure 15 shows the form for specifying the widths of the stop bands and transition bands in the case of monthly data. The cells contain the default values.)

The Twofold Decomposition of the Data

When the filter is applied to the residuals from a polynomial de-trending, the program will effect a twofold decomposition of the data, comprising a seasonal component and a seasonally adjusted data sequence.

The program will display the periodogram of the residuals upon which frequency response function of the filter is superimposed. This will be followed by a graph of the seasonal component

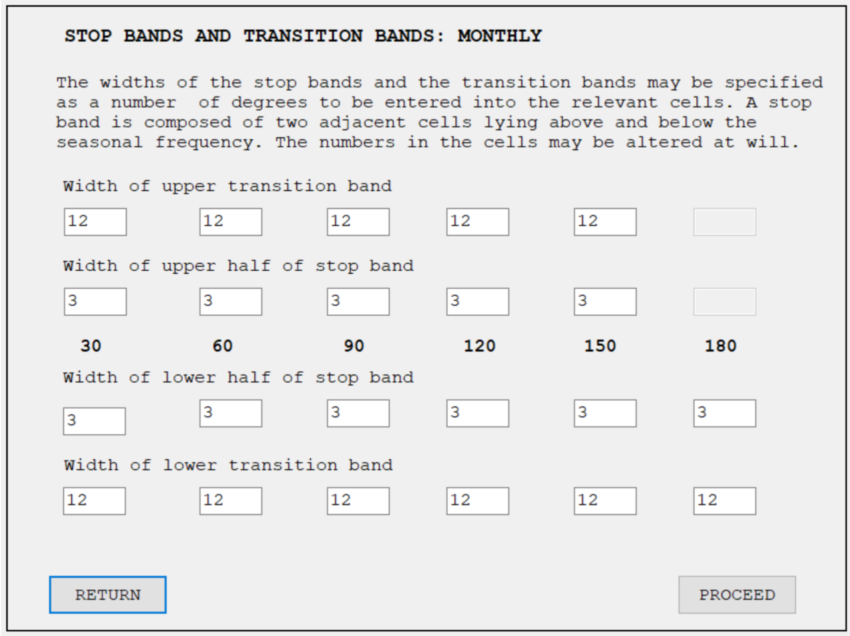

Figure 15. The stop bands and transition bands for a frequency-domain seasonal-adjustment filter for monthly data.

that has been extracted by the filter. Then, the seasonal component will be subtracted from the data to produce a rough seasonally adjusted sequence, which is plotted in red and superimposed on the data.

The Threefold Decomposition of the Data

If a trend-cycle has been extracted already from the data, then program will effect a threefold decomposition of the data, comprising the seasonal component, the trend-cycle function and a residual noise sequence. It will also create a sequence of seasonally adjusted data that is identical to the sequence produced by a twofold decomposition of the data.

The program will display the periodogram of the residual deviations of the data from the trend-cycle function, upon which the frequency response function of the filter will be superimposed. It will be observed that this periodogram lacks the spectral structure of the low-frequency fluctuations.

The graph of the periodogram will be followed by a graph of the seasonal component extracted by the filter. Then, the seasonal component will be subtracted from the data to produce a seasonally adjusted sequence, which is plotted in red and superimposed on the data. In the next display, the trend-cycle function is plotted, superimposed of the data, in place of the seasonally adjusted data sequence of the previous display. Finally, the residual sequence, which is the third component of the threefold decomposition, is plotted.

10.1. A Seasonal-Adjustment Filter without Transition Bands

A seasonal adjustment filter without transition bands cuts vertical shafts in an otherwise constant frequency response function of unit gain. The stop bands are composed of adjacent intervals lying below and above the seasonal frequencies.

The Frequency Response of the Filter

The frequency response function takes the value of unity in the pass band and zero in the stop bands. A default value, specified in degrees, is supplied for the widths of all of the stop bands, which the user may alter in view of the periodogram of the residuals of a polynomial de-trending of the data. The widths of the stop bands are specified in numbers of degrees. The objective is to encompass, within the stop bands, all of the elements that contribute to the seasonal fluctuations of the data.

10.2. A Seasonal-Adjustment Filter with Sigmoid Transitions

The transition from a pass band to a stop band has the same profile as that of the transition of the lowpass Sigmoid Filter employed in estimating a trend-cycle function. The profile of the transition from a stop band to the pass band is liable to be a mirror image of the former profile, albeit that the widths of the two transition bands may differ.

The Frequency Response of the Filter

The profile of the transition from pass band to stop band resembles that of the cosine function over the interval $[0°, 180°]$, which it does exactly when the Power Parameter is unity. The Power Parameter, which is an integer value in the interval $[1, 6]$, is given a default value of 3. The effect of raising its value is to defer the onset of the transition from pass band to stop band and, thereafter, the rate of descent is increased. As the stop band is approached, the effect is to reduce the slope of the curve more rapidly as it tends to zero.

10.3. A Seasonal-Adjustment Filter with Upper-half Cosine Transitions Bands

In the absence of stop bands, the transition bands incorporated in the Upper-half Cosine filter enable it to mimic closely the effects of the ordinary time-domain filter. The power of the filter to suppress elements adjacent to the seasonal frequencies is enhanced by adding stop bands on either side of these frequencies. The extent of each of these stop bands can be specified individually as a number of degrees.

The Frequency Response of the Filter

The frequency response of the filter is based on the trajectory of the cosine function over the interval $[0°, q \times 90°]$, where $q \leq 1$ is described as the Truncation Parameter. This parameter takes a value in the interval $[0.25, 1.0]$ and it is given a default value of 0.5. By taking a power of the argument of the cosine function, the onset of the descent of function from the pass band to the zero of the stop band is delayed whereafter the descent becomes steeper. The Power Parameter, which is an integer value in the interval $[1, 6]$, is given a default value of 3.

10.4. A Seasonal-Adjustment Filter with Lower-half Cosine Transition Bands

The Lower-half Cosine filter has transition bands of which the profiles resemble that of a hockey stick. The descent from the pass band to the stop band begins abruptly and steeply. It flattens as the stop band is approached, with a slope that is tending to zero. The widths of the stop bands on either side of the seasonal frequencies may be specified individually in all cases by the user, albeit that default values are provided.

The Frequency Response of the Filter

The frequency response of the filter in descending from the pass band to the stop band is based on the trajectory of cosine function in the interval $[q \times 90°, 180°]$, where $q \leq 1$ is described as the Truncation Parameter. The initial rate of descent is increased by taking a power of the argument of the cosine function, which also flattens the trajectory as it approaches the stop band. The Power Parameter, which is an integer value in the interval $[1, 6]$, is given a default value of 3.

References

Gómez, Victor, and Augustin Maravall.l. 1997. *TRAMO (Time Series Regression with ARIMA Noise, Missing Observations, and Outliers) and SEATS (Signal Extraction in ARIMA Time Series) Instructions for the User,* Banco de España, Madrid.

Gómez, Victor, and Augustin Maravall. 2001. Seasonal Adjustment and Signal Extraction in Economic Time Series. In Daniel. Peña, George C. Tiao, and Ruey .S. Tsay (eds.), *A Course in Time Series Analysis*, John Wiley and Sons, New York.

Grudkowska, Sylvia. 2017. *JDemetra*+ Reference Manual Version 2.2. Available from URL: https://jdemetradocumentation.github.io/JDemetra-documentation/ (accessed on 1 January 2021.)

Hodrick, Robert J. and Edward C. Prescott. Postwar U.S. Business Cycles: An Empirical. Investigation, *Journal of Money Credit and Banking*, 29, 1–16

Leser, Conrad Emmanuel Victor. 1961. A Simple Method of Trend Construction, *Journal of the Royal Statistical Society*, Series B, 23, 91–107.