

**ONS Methodology and Statistical
Development**

**Guide to Seasonal
Adjustment with
X-12-ARIMA
** DRAFT ****

**TSAB
March 2007**

TABLE OF CONTENTS

1. Introduction	1.1
2. Introduction to seasonal adjustment	2.1
3. Overview of the X12ARIMA Method	3.1
4. How to run X12ARIMA	4.1
7. Length of the series	7.1
8. Consistency across time	8.1
9. Aggregate Series	9.1
10 Revisions and updates	10.1
11. The regARIMA model	11.1
12. Trading day	12.1
13. Easter	13.1
14. Level Shift and Additive Outliers	14.1
15. Which Seasonal Decomposition Model to Use	15.1
16. Moving Averages	16.1
17. Seasonal Breaks	17.1
18. Existence of Seasonality	18.1
19. X12ARIMA Standard Output	19.1
20. Graphs and X-12-graph	20.1
21. Sliding Spans	21.1
22. History Diagnostics.....	22.1
23. Composite spec	23.1

25. Forecasting	25.1
26. Trend Estimation	26.1
27. Quality Measures	27.1

1 INTRODUCTION

1.1 Introduction

Seasonal adjustment is widely used in official statistics as a technique for enabling timely interpretation of time series data.

The X-12-ARIMA seasonal adjustment package has been chosen from the many available seasonal adjustment methods as the standard one for use in official statistics in the United Kingdom (UK). This was agreed by the National Statistics Quality and Methods Programme Board in 2001, and is in line with European best practice and consistent with the Bank of England. The method is used by most of the leading national statistical institutes across the world. X-12-ARIMA is developed by the United States (US) Bureau of the Census. The software is comprehensive, with many options available for tailoring seasonal adjustment to each individual series. It therefore requires many choices to be made by its users.

The main purpose of this guide is to provide practical guidance on seasonal adjustment using X-12-ARIMA. The guide explains what seasonal adjustment means, and addresses many of the statistical, organisational and presentational issues and problems associated with seasonal adjustment.

This guide complements the Office for National Statistics (ONS) training courses in seasonal adjustment. A detailed explanation of the X-12-ARIMA method can be found at www.census.gov/srd/www/x12a or in Ladiray and Quenneville (2001).

1.2 What is in this guide

The guide starts with a brief introduction to seasonal adjustment and to some of the main associated issues (Chapter 2). It then provides an overview of the X-12-ARIMA method (Chapter 3), and a description of how to run the program (Chapter 4). Chapter 5 describes a procedure that should be used when seasonally adjusting a series using X-12-ARIMA, while the prior considerations that users should address when seasonally adjusting a series are discussed in Chapters 6 to 10.

The second part of the guide examines seasonal adjustment in more detail. Here the user will find more detailed descriptions of the regression-ARIMA model, of potential effects such as trading days, and of how X-12-ARIMA deals with them (Chapters 11 to 14). Guidance is also provided on selecting a decomposition model, moving averages, prior adjustments, and on problems such as seasonal breaks (Chapters 15 to 19).

The third part of this guide deals with X-12-ARIMA output and with diagnostics or tools that may be useful in seasonal adjustment (Chapters 20 to 25).

Chapters 26 to 28 look at alternative uses, other than seasonal adjustment, of the X-12-ARIMA software.

1.3 How to use this guide

This guide is not intended to be read from cover to cover. Users are advised to read Part 1 before attempting any seasonal adjustment, and to refer to other parts of the guide as appropriate.

Anyone who is responsible for setting up or reviewing seasonal adjustment should attend a training course aimed at seasonal adjustment practitioners. This guide complements, but is not a substitute for, such a training course. ONS provides short training courses for users of seasonal adjustment, and for

persons responsible for presenting data in seasonally adjusted form. More information on training courses can be obtained from Time Series Analysis Branch (TSAB).

1.4 What this guide is not

This guide does not describe the detailed working of X-12-ARIMA, nor the underlying mathematics. For those who are interested in these matters, the X-12-ARIMA user manual is a good starting point. Additional technical references and papers are listed in the References below.

1.5 What does TSAB do?

The Methodology Directorate (MD) provides leadership, and defines best practice for many methodological aspects of ONS work. MD also provides methodological advice to other parts of ONS and, where resources permit, to the rest of the Government Statistical Service (GSS), on methodological issues. TSAB provides expertise in time series analysis issues, and in particular seasonal adjustment. TSAB has produced this guide, and runs regular courses in seasonal adjustment. TSAB is responsible for the quality of all ONS seasonal adjustment, and manages a rolling programme of annual seasonal adjustment reviews of statistical outputs across ONS. Furthermore, TSAB provides support for practitioners and users of seasonal adjustment, and should be the main point of contact for any questions or queries regarding seasonal adjustment.

1.6 Contact details

If you have a query related to time series analysis, if you require information on seasonal adjustment training, or if you have any comments on this guidance, we can be contacted at:

Office for National Statistics, Time Series Analysis Branch, 1 Drummond Gate, London SW1V 2QQ

or by email at: tsab@ons.gov.uk.

1.7 Acknowledgements

This guide has been developed with input and assistance from the following people:

Claudia Annoni, Simon Compton, Duncan Elliott, Lee Howells, Fida Hussain, Peter Kenny, Jim Macey, Craig McLaren, Ross Meader, Nigel Stuttard, Anthony Szary, and Anita Visavadia.


Please bring any errors or omissions to our attention.

2 INTRODUCTION TO SEASONAL ADJUSTMENT

2.1 What is seasonal adjustment?

Data that are collected over time form a time series. Many of the most well known statistics published by the Office for National Statistics are regular time series, including: the claimant count, the Retail Prices Index (RPI), Balance of Payments, and Gross Domestic Product (GDP). Those analysing time series typically seek to establish the general pattern of the data, the long term movements, and whether any unusual occurrences have had major effects on the series. This type of analysis is not straightforward when one is reliant on raw time series data, because there will normally be short-term effects, associated with the time of the year, which obscure or confound other movements. For example, retail sales rise each December due to Christmas. The purpose of seasonal adjustment is to remove systematic calendar related variation associated with the time of the year, i.e. seasonal effects. This facilitates comparisons between consecutive time periods.

Figure 2-1 - Non-seasonally adjusted (NSA) and seasonally adjusted (SA) data for United Kingdom visits abroad
UK visits abroad



2.2 Components of a time series

Time series can be thought of as combinations of three broad and distinctly different types of behaviour, each representing the impact of certain types of real world events on the data. These three components are: systematic calendar-related effects, irregular fluctuations, and trend behaviour.

Systematic calendar related effects comprise seasonal effects and calendar effects. Seasonal effects are cyclical patterns that may evolve as the result of changes associated with the seasons. They may be caused by various factors, such as:

- Weather patterns: for example, the increase in energy consumption with the onset of winter;
- Administrative measures: for example, the start and end dates of the school year;

- Social / cultural / religious events: for example, retail sales increasing in the run up to Christmas;
- Variation in the length of months and quarters due to the nature of the calendar.

Other calendar effects relate to factors which do not necessarily occur in the same month (or quarter) each year. They include:

- Trading day effects which are caused by months having differing numbers of each day of the week from year to year: for example, spending in hardware stores is likely to be higher in a month with five, rather than four, weekends;
- Moving holidays, which may fall in different months from year to year: for example Easter, which can occur in either March or April.

Taken together these effects make up the **seasonal component**.

Irregular fluctuations may occur due to a combination of unpredictable or unexpected factors, such as: sampling error, non-sampling error, unseasonable weather, natural disasters, or strikes. While every member of the population is affected by general economic or social conditions, each is affected somewhat differently, so there will always be some degree of random variation in a time series. The contribution of the irregular fluctuations will generally change in direction and/or magnitude from period to period. This is in marked contrast with the regular behaviour of the seasonal effects.

The **trend (or trend cycle)** represents the underlying behaviour and direction of the series. It captures the long-term behaviour of the series as well as the various medium-term business cycles.

2.3 The seasonal adjustment process

Although there are many ways in which these components could fit together in a time series, we select one of two models:

- Additive model : $Y = C + S + I$
- Multiplicative model: $Y = C \times S \times I$

where Y is the original series, C is the trend-cycle, S is the seasonal component, and I is the irregular component.

The **seasonally adjusted series** is formed by estimating and removing the seasonal component.

- For the additive model : Seasonally adjusted series = $Y - S = C + I$
- For the multiplicative model: Seasonally adjusted series = $Y / S = C \times I$.

In a multiplicative decomposition, the seasonal effects change proportionately with the trend. If the trend rises, so do the seasonal effects, while if the trend moves downward the seasonal effects diminish too. In an additive decomposition the seasonal effects remain broadly constant, no matter which direction the trend is moving in.

In practice, most economic time series exhibit a multiplicative relationship and hence the multiplicative decomposition usually provides the best fit. However, a multiplicative decomposition cannot be implemented if any zero or negative values appear in the time series.

2.4 Other factors that affect seasonal adjustment

There are a variety of issues that can impact on the quality of the seasonal adjustment. These include:

- Outliers, which are extreme values. These usually have identifiable causes, such as strikes, war, or extreme weather conditions, which can distort the seasonal adjustment. They are normally considered to be part of the irregular component;
- Trend breaks (also known as level shifts) where the trend component suddenly increases or decreases sharply. Possible causes include changes in definitions relating to the series that is being measured, to take account of say a reclassification of products or a change in the rate of taxation;
- Seasonal breaks, where there are abrupt changes in the seasonal pattern.

These issues need to be addressed before the seasonal adjustment process begins, in order to obtain the most reliable estimate of the seasonal component.

2.5 Interpreting time series outputs

The original, seasonally adjusted, and trend estimates are available for users to assess.

In ONS publications the focus is usually on the level of the seasonally adjusted series, the period-to-period change in the level of the seasonally adjusted series, and the period-to-period growth rate for the seasonally adjusted series. A comparison with the same period in the previous year may also be published, but this gives a historical picture of the growth rate. On average, this measure will result in a lag of six months in the identification of turning points in the series. It may be thought that this year-on-year change would be the same for the seasonally adjusted and the non-seasonally-adjusted series, but this will only be the case if the seasonal component is stable. In practice seasonality often evolves over time, and the seasonal factors should reflect this. In general, seasonally adjusted estimates will not sum to the annual estimates. This is due to a variety of factors including moving seasonality, incomplete cycles and outlier treatment.

The ONS currently recommended measure of year-on-year change is based on the seasonally adjusted series, as this takes into account the impact of a change in seasonal patterns over time.

People are often interested in removing the impact of the irregular component from the seasonally adjusted estimates in order to produce a trend (also known as a trend-cycle, or a short-term trend).

Seasonally adjusted and trend estimates are subject to revision as additional raw data become available. As the seasonal component cannot be directly observed, it is estimated using moving averages and these can change as new data points are added to the series.

2.6 X-12-ARIMA

X-12-ARIMA was developed by United States Bureau of the Census (1998) as an extended and improved version of the Statistics Canada X-11-ARIMA method. The program broadly runs through the following steps:

- The series is modified by any user defined **prior adjustments**;
- The program fits a **regression-ARIMA model** to the series in order to: detect and adjust for outliers and other distorting effects, to improve forecasts and seasonal adjustment; detect and estimates additional components such as calendar effects; and extrapolate forwards (forecast) and backwards (backcast) an extra one to three years of data;
- The program then uses a series of **moving averages** to decompose a time series into the three components. It does this in three iterations, getting successively better estimates of the three components. During these iterations **extreme values** are identified and replaced;

- A wide range of **diagnostic statistics** are produced, describing the final seasonal adjustment, and giving pointers to possible improvements which could be made.

2.7 Know your series

In addition to being able to run a seasonal adjustment program, such as X-12-ARIMA, and to understand the outputs, a proficient user should have an appreciation of factors that are likely to affect the series being seasonally adjusted. Knowing the background to a series will give clues as to where to look for likely problems. For example:

- Is the way in which the data are collected likely to lead to any unusual effects? Are they collected on a non-calendar basis, or is there a lag between the activity being measured and when this is recorded?
- Has there been any change to the method or timing of data collection? This may cause trend or seasonal breaks.
- Is the series likely to be affected by trading days or Easter effects?
- Have there been any events which are likely to cause breaks in the series or large outliers? These could include: the Gulf war, the budget moving from March to November, Britain dropping out of the ERM, extreme weather, industrial disputes, or other events which may only affect individual series.

3 OVERVIEW OF THE X-12-ARIMA METHOD

3.1 Introduction

Many users regard the X-12-ARIMA software package as a black box method that can take a time series and automatically derive seasonally adjusted estimates. This chapter outlines the concepts used within X-12-ARIMA. Additional details on the use of X-12-ARIMA can be obtained from: www.census.gov/srd/www/x12a/.

3.2 The X-12-ARIMA method

The X-12-ARIMA method can be described by the flowchart in Figure 3-1.

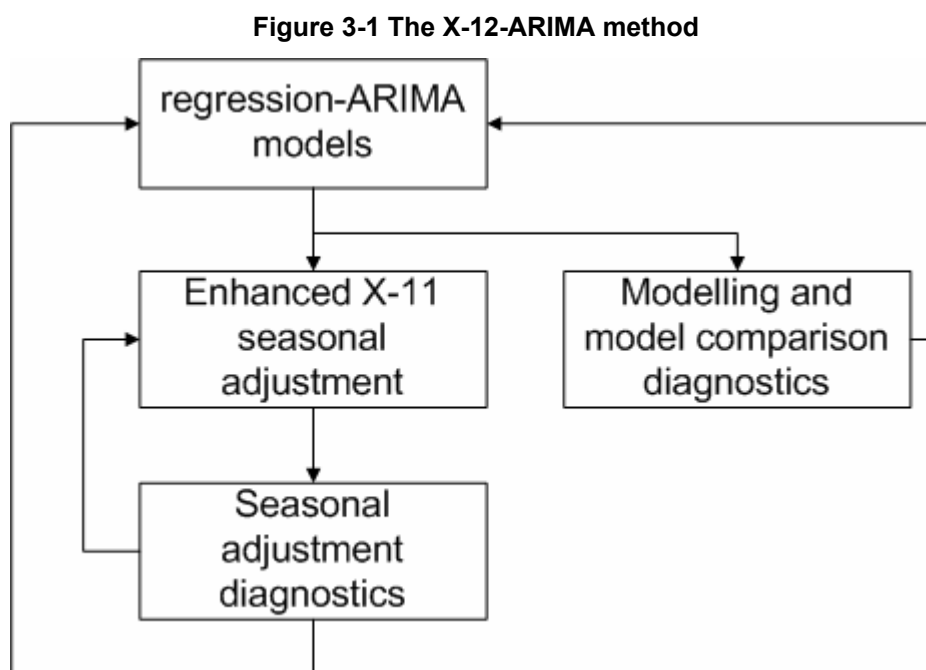


Figure 3-1 shows the chronological sequence in an analysis, but it may be helpful to think in terms of a historical sequence. The X-11 (Shiskin, et. al, 1967) seasonal adjustment method was developed by the United States Bureau of the Census in the 1960s, and has formed the basis for most official government seasonal adjustment since then. The X-11-ARIMA (Dagum, 1980) program is an improved version of the X-11 program and was developed by Statistics Canada in the 1980s. In recent years the United States Bureau of the Census has developed X-12-ARIMA (Findley, et. al, 1998) which uses a modelling strategy called regression-ARIMA, which can be used to identify and quantify factors which will impact on the seasonal factor estimation. If the output of the regression-ARIMA is used as a prior adjustment for the raw series, the result is a series which X-11 can handle effectively. The two parts are described below.

3.3 regression-ARIMA models

This approach can model the time series and estimate any prior adjustments before seasonal adjustment. For example, regression-ARIMA models are able to detect and adjust for outliers and other distorting effects to improve the forecasts and seasonal adjustment. A regression-ARIMA model takes the form:

$$\log\left(\frac{Y_t}{D_t}\right) = \beta' X_t + Z_t$$

where $\log\left(\frac{Y_t}{D_t}\right)$ is an (optional) transformation of the original data Y_t , Z_t = an ARIMA process, X_t = regressors for trading day and holiday or calendar effects, additive temporary changes, level shifts, ramps, user-defined and other effects, and D_t = leap-year adjustment or other modification for known external effects (optional).

There are a range of built-in types of regression variables available in X-12-ARIMA. These include:

- Outlier and trend change effects: Additive (or point) outliers, Temporary change outliers, Level shifts, Ramps
- Seasonal effects: Calendar month indicators, Trigonometric seasonal (sines and cosines)
- Calendar effects: Trading day (flows or stocks), Leap-year February, length of month, Shifting holidays (e.g. Easter)
- User-defined effects

Diagnostics to evaluate the fit of the model or the significance of the regression variables are available in the X-12-ARIMA output. More information on the regression-ARIMA model can be found in chapter 11.

3.4 Description of the X-12-ARIMA method

X-12-ARIMA performs seasonal adjustment using the X-11 method. The following is a brief outline of the X-11 method.

The original series, Y_t , can be decomposed into a trend cycle, C_t ; a seasonal component, S_t and an irregular component, I_t .

Trend-cycle C_t : is defined as the underlying level of the series. It is a reflection of the medium-long term movement in a series and it is typically due to influences such as population growth, general economic development and business cycles. It refers to the generally smooth deterministic movement in a time series.

Seasonal component S_t : includes variations which repeat approximately periodically with a period of one year and which evolve more or less smoothly from year to year. Calendar variations, such as Easter effects, are also included with the seasonal component.

Irregular component I_t : contains those parts of the time series that cannot be predicted and are effectively the residual component after the identification of the trend and seasonal components. It may include sampling errors and unpredictable events like strikes and floods.

The X-11 process can be described by the following sequence of steps. For a full description see Ladiray and Quenneville (2001):

1. Assume that there is a multiplicative relationship between the components. For example: $Y_t = C_t \times S_t \times I_t$. This is the default decomposition model within X-12-ARIMA. An alternative decomposition is the additive decomposition where all the components are related additively. In the multiplicative model S_t and I_t are scale-free numbers varying about a level of 1 (or 100%);

in the additive model S_t and I_t are in the same units as the original series and vary about a level of zero. For more information on the decomposition methods look at chapter 15

2. A preliminary trend cycle (C'_t) is obtained by applying a trend moving average to the original series (Y_t).
3. This initial estimate of the trend is removed from the original estimate to give a “detrended” time series denoted by: $SI_t = S_t \times I_t = Y_t / C'_t$.
4. Outliers are then identified by an automatic process and replaced in the SI_t time series
5. A seasonal moving average is then applied to the modified SI_t time series for each month (or quarter) separately to give a preliminary estimate of the seasonal component, \hat{S}_t .
6. Dividing Y_t by \hat{S}_t gives a preliminary seasonally adjusted series, denoted by \hat{SA}_t^1 .
7. This process is then repeated, using a Henderson moving average (Henderson, 1916) to estimate the trend cycle in step 1.

The following paragraphs describe each step of the X-11 seasonal adjustment process in more detail.

3.4.1 Step 2: Preliminary estimate of the trend

Some moving averages have a smoothing effect and others can estimate and remove seasonality in data. These are applied to the original series to have a preliminary estimate of the trend C; i.e. for a monthly series a 2 x 12 moving average is applied whilst for a quarterly series a 2 x 4 moving average is used.


Moving averages are described in more detail in Chapter 16. A 2x12 moving average is centred at the current month (t, say) and it uses the six months before, the current month and the six months after. In addition, weights are symmetrical i.e. the same weight is given to each month except the first and last. It can be shown that the effect of a 2x12 average on a monthly series is to remove any stable annual variation and to reduce the variance of any purely random component by a factor of 12.5, while leaving any linear trend unchanged.

3.4.2 Step 3: Knowing the trend an estimate of seasonal and irregular can be calculated

Once the trend estimate is known, it is possible to identify an estimate of the remaining SI component. This is done by dividing the original data by the first estimate of C (if a multiplicative model is used, $Y=C \times S \times I$) or by subtracting C from the original data (if an additive model is used, $Y=C+S+I$). Therefore a preliminary estimate of 'SI' is given by Y/C for the multiplicative scenario or $Y-C$ for an additive decomposition. Termed the SI ratios or differences (although they are technically percentages not ratios), they may be plotted for each quarter/month separately.

The following diagram shows an example of SI ratios for January.

Figure 3-2 January SI ratios for Retail Sales in non-specialised stores (predominantly food)

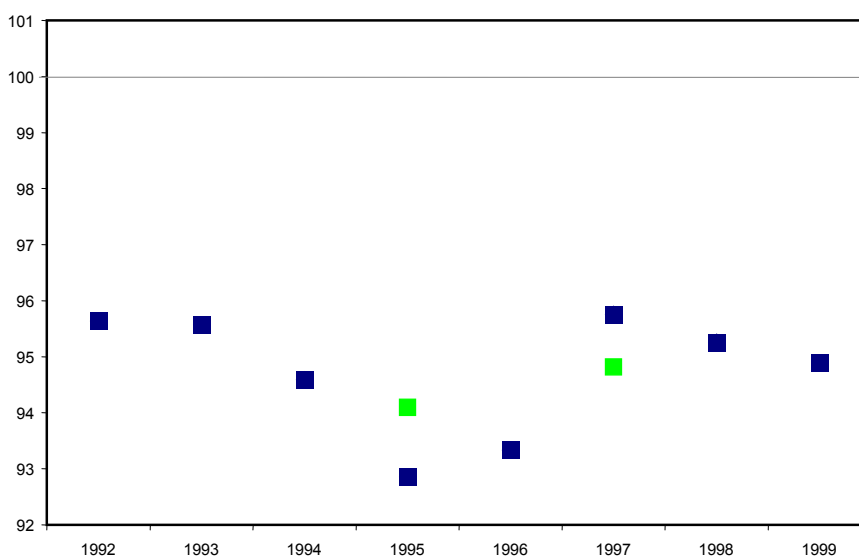


The 100% line at the top of the graph represents the new trend of the series once C has been removed from Y. This diagram clearly shows that the January SI values are all below trend, typical example of a seasonal effect. The graph of the SI ratios is a useful tool that should be used in analysing the seasonal behaviour, especially to identify breaks in the series.

3.4.3 Step 4: Outliers are then identified and replaced in the ‘SI’ series

X-12-ARIMA identifies and temporarily replaces outliers in the SI graphs for each month. This reduces any distortion of the seasonal adjustment. In the following diagram the outliers are temporarily replaced by the points shown in green.

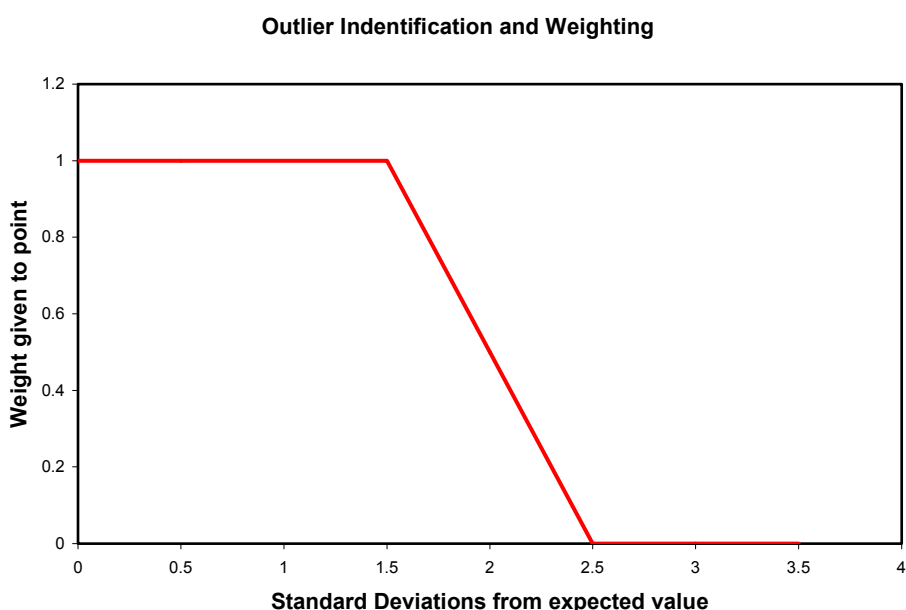
Figure 3-3 January SI ratios with replacement values



The routine to identify and replace outliers is used several times to get progressively better estimates of the size of the outliers. The following steps are used:

- A rough estimate of the irregular is made, from the combined seasonal-irregular component.
- A standard deviation is calculated for each 5 year moving span of the irregular component and used as the standard deviation for the middle year of the five. (For the first two and last two years the standard deviation of the nearest available five year span is used.)
- Where any irregular is more than 1.5 standard deviations from zero (or from 100% for a multiplicative model), the SI value for that point is considered an outlier; it is replaced with an average of that value and the two nearest preceding and following full weight values for corresponding months/quarters. The outlier is given a weight which depends upon how extreme the irregular is, as shown in the following diagram:

Figure 3-4 Outlier Identification and weighting




The modified SI is used for estimating the seasonal component so that outliers do not distort the estimate. However, outliers remain in the final adjustment.

3.4.3.1 A seasonal moving average is applied to the modified ‘SI’ series for each month/quarter separately to give a preliminary estimate of ‘S’ and hence ‘I’

A seasonal moving average (a 3x3 in the first part of iteration B and a 3x5 in the second part of iteration B) is then applied to the SI series to estimate the seasonal factors. If the derived moving averages represent the seasonal component, the irregular component is defined by the deviation of each point from the moving averages. This process is repeated for each month/quarter. An example of the January SI ratios and moving average values is given below.

January SI ratios with moving average values

Figure 3-5 January SI ratios with moving average values

The above diagram illustrates the results of the application of a 3x3 moving average (pink line) to the modified January SI ratios. It shows the graphical representation of the separation of the seasonal and irregular components.

3.4.4 Step 6: Dividing 'Y' by 'S' gives a preliminary seasonally adjusted series, denoted SA1

After the seasonal factors have been estimated, the seasonally adjusted series can be derived. The original series Y is divided by S to give a preliminary seasonally adjusted series $C \times I$ (if a multiplicative model $Y = C \times S \times I$ is used). In case of an additive model $Y = C + S + I$, the original series Y minus the seasonal component S gives an estimate of the seasonally adjusted series $C + I$. This seasonally adjusted series is referred to as SA1 since it is just the result of the first iteration.

3.4.5 Step 7: Repeat seasonal estimation with improved trend estimates

This process from 3.4.1.1 to 3.4.1.5 is then repeated, but this time the trend is estimated from the SA1 series by applying a Henderson moving average (such moving averages are very good at estimating trends, but can only be used on series which do not exhibit seasonality).

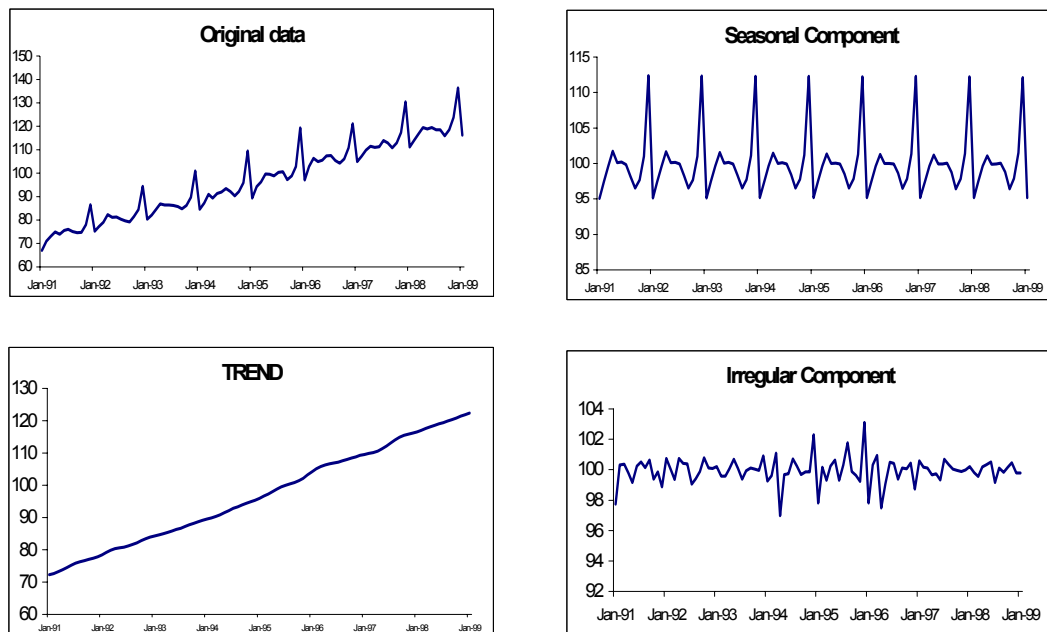
3.4.6 Two more iterations – X-11 Parts C and D

The output of the steps described above is not regarded as the final set of estimates. X-12-ARIMA carries out this entire process (i.e. 3.4.1.1 to 3.4.1.6, including two estimations of the seasonal component) twice more. For these estimations, however, the starting point is not the raw series Y but a modified Y with extremes removed or reduced. The final output of the third iteration (Part D) is the final set of time series estimates.

3.5 Example: Decomposing the original time series

Figure 3-3 shows the original series and the trend-cycle, seasonal and irregular components. The purpose of seasonal adjustment is to estimate and remove the seasonal component from the original data to give a seasonally adjusted estimate.

Figure 3-6 - Graphical representation of time series components



3.6 Diagnostic checking

As the flowchart shows, both the regression-ARIMA modelling and the X-11 decomposition should be checked by the user for accuracy and improved where necessary. This process is carried out by reviewing the outputs produced by X-12-ARIMA. More details can be found in the chapters dealing with those topics.

4 HOW TO RUN X-12-ARIMA

4.1 Introduction

Detailed instructions on how to use X-12-ARIMA are available from <http://www.census.gov/srd/www/x12a/> and in the manuals supplied with X-12-ARIMA. The following is a condensed guide.

To run X-12-ARIMA, the following two files must be located in the same directory of your computer; an executable file, **x12a.exe**, and a model list, **x12a.mdl**.

Two additional files must also be created:

1. One file containing the data that is to be seasonally adjusted using X-12-ARIMA, named, for example, **myfile.dat**.
2. The second file contains the instructions to seasonally adjust the data series, and is known as the specification file. It is given a name such as **myfile.spc**.

Both files must be in text format. They can easily be created using either the Notepad text editor, or Oxedit (an editor that is user friendly). Once these files have been created, X-12-ARIMA can be run to perform the seasonal adjustment.

4.1.1 The data file

The data file is simply a list of the raw data that have to be seasonally adjusted. The data can be either in free format, in date-value format, or in X12-save format.

- **Free format.** This is the more commonly used format. It comprises a single column of data saved with the ".dat" extension. An example of free format data is given below:

```
89814.18
85582.54
85059.04
85880.88
97711.15
78266.99
84479.21
79941.65
```

If the data are in free format, then a "start" statement with the start date of the series is required in the "series" specification.

- **Date-value.** A file in date-value format includes dates. The delimiter that separates each field may be a tab or a space. An example of date-value data is provided below. In this example the delimiter between the year and the month, and between the month and the value, is a tab.

```
1991 1 89814.18
1991 2 85582.54
1991 3 85059.04
1991 4 85880.88
1991 5 97711.15
1991 6 78266.99
1991 7 84479.21
1991 8 79941.65
```

In this case the "start" argument in the "series" specification is not required. Labels are not required within the data file, but they are permitted.

- **X12-save.** A file in X-12-save format includes dates. A file may only be produced in X-12-save format by X-12-ARIMA. An example of X-12-save data is given below:

```
date      mydata.d11
-----
199101    8.98E+04
199102    8.56E+04
199103    8.51E+04
199104    8.59E+04
199105    9.77E+04
199106    7.83E+04
199107    8.45E+04
199108    7.99E+04
```

In this case the "start" argument in the "series" specification is not required. Labels are not required within the data file, but are allowed.

4.1.2 The specification file

To run X-12-ARIMA for any series, a command file, also called the specification file or input file, must be created. A specification file is a text file used to specify program options. The name of the specification file must end with a ".spc" extension.

A specification file contains a set of functional units called specifications that X-12-ARIMA reads to obtain required information for the time series, such as: the decomposition model to be used, the analysis to be performed, and the desired output. Each specification controls options for a specific function. For example, the "series" specification specifies the original data to be input and the span to be used in the analysis, whereas the "transform" specification specifies any transformation and any prior adjustments to be applied. X-12-ARIMA includes four types of specifications:

- Data input and transformation;
- regression-ARIMA modelling;
- Seasonal adjustment;
- Diagnostics.

An X-12-ARIMA specification file must begin with a "series" or a "composite" specification, whilst the other specifications can be entered in any order.

A general input syntax of a spec file is presented below:

```
specificationname{
    argument1 = value
    argument2 = (value1 value2 value3)
    argument3 = "A string value"
    argument4 = 2003.9    # The dates 2003.9 and 2003.SEP are equivalent
}
# This symbol is followed by a comment that is not meant to be executed.
```

When writing a specification file, one should note that:

- Dates are in the format yyyy.period (e.g. 2003.9 or 2003.SEP);
- If an argument has 2 or more values, these must be enclosed in parentheses;
- Character values, such as titles and file names, should be enclosed in quotes (" ");
- Everything on the line following the "#" symbol is treated by X-12-ARIMA as a comment.

Arguments can be set within each specification. Arguments define the function that X-12-ARIMA has to run, and if an argument is not specified, a default value is usually assigned. Arguments can be written in any order within the specification, using either upper, lower, or mixed case (ie. X-12-ARIMA is not case sensitive).

The specification file also includes the name and the path of any optional files containing data for the time series being modelled, data for user-defined and predefined regressors, values for any user-defined prior adjustments, and model types to try with the automatic model procedure. These names and paths are listed in appropriate specifications. An example of a simple specification file is shown below:

```
series{
    title="Example of specification for default seasonal adjustment"
    start=1994.1
    period=4
    name="Default"
    file="c:\research\test.dat"
}

transform{
    function=log
    file="c:\research\testtp.tpp"
    type=temporary
}

arima{
    model=(0,1,1)(0,1,1)
}

regression{
    aictest=(EASTER)
}

x11{}
```

The above specification file includes the **series**, **transform**, **arima**, **regression** and **x11** specifications. The **series** specification provides X-12-ARIMA with a file name and path, indicating where the original data are stored. The **transform** specification takes the logarithms of the series, and provides the name of a file containing the temporary priors to prior adjust the series. The file specifies an ARIMA (0,1,1)(0,1,1) model, using the **arima** specification, and tests the significance of regression variables, known or suspected to affect the series, as indicated in the regression specification. Here aictest=(EASTER) tests whether an Easter variable should be included within the regression-ARIMA model specified. Finally, the specification generates seasonal adjustments from the default selection of seasonal and trend filters.

To set up a more complicated specification file, please refer to the "X-12-ARIMA Reference Manual" or contact TSAB.

4.2 Running X-12-ARIMA

X-12-ARIMA is a DOS program that can be run in one of the following environments:

1. From the DOS prompt;
2. Using an interactive file editor, such as OxEdit.

4.2.1 In DOS

The following commands, detail how to run the X-12-ARIMA program in a DOS environment. The symbol ¶ indicates that the 'return' or 'enter' key should be pressed to run the program.

```
c:\X12a> x12a path\myfile
```

In this statement *path\myfile.spc* is the main input (specification) file used by X-12-ARIMA. The path to the input file is required if the specification file is not in the same directory as the X-12-ARIMA program.

The standard output of the run is written to the file *path\myfile.out*. The "print" optional argument in individual specifications controls what is to be added to (or removed from) this output.

With the standard output, any input errors are printed in appropriate error messages which are stored in a file named *path\myfile.err*. When the program can localise the error, the line in the spec file containing the error will be printed with a caret (^) positioned under the error. If the program cannot localise the error, only the error message will be printed. If the error is fatal, then "ERROR:" will be displayed before the error message, sometimes accompanied by suggestions of what to change. For non-fatal errors, "WARNING:" will be printed before the message. WARNING messages are also sometimes used to call attention to a situation in which no error has been committed, but some caution is appropriate. X-12-ARIMA first reads the whole specification file, reporting all input errors it finds. This enables the user to correct more than one input error per run. The program will stop if any fatal errors are detected. Warnings will not stop the program, but should alert users to check both the input and output carefully, to verify that the desired results are being produced.

For the purpose of examining the effects of different adjustment and modelling options on a given series, it is sometimes desirable to use a different name for the output than was used for the input. The general form for specifying an alternative file name for the output files is:

```
c:\X12a>x12a path\myfile path\newname¶
```

In this case, the X-12-ARIMA run uses the spec file "myfile.spc", but the output will be stored in the file "newname.out" and any error messages will be stored in the file "newname.err".

In a production situation, it is often necessary to run more than one series in a single run. X-12-ARIMA allows for running multiple series in either of two modes:

- (a) multi-spec mode, where there are separate input specification files for each series specified;
- (b) single spec mode, where every series will be run with the options detailed in a single input specification file.

4.2.1.1 Running X-12-ARIMA in multi-specification mode

Before X-12-ARIMA can be run in multi-specification mode, an input metafile must be created. This is a text file, which contains the names of the files to be used by X-12-ARIMA. An input metafile can have up to two entries per line: the file name (and path information, if necessary) of the input specification file for a given series, and an optional output file name for the output of that series. If an output file name is not provided by the user, the path and file name of the input specification file will be used to generate the output files. The input specification files are processed in the order in which they appear in the input metafile. For example, to run the spec files **series1.spc**, **series2.spc** and **series3.spc**, the input metafile should contain the following:

```
series1  
series2  
series3
```

This assumes that all these spec files are stored in the x12a directory. If they were stored in the "c:\research\specs" DOS directory, the metafile ought to read:

```
c:\ research\specs\series1
```

```
c:\research\specs\series2
c:\research\specs\series3
```

To run X-12-ARIMA with an input metafile, use the following command line:

```
x12a -m mygroup¶
```

where **mygroup** is the name of the metafile and **-m** is a flag, which informs X-12-ARIMA of the presence of a metafile. Note that the extension, **.mta**, identifies a file as a metafile, but the extension for the input metafile is excluded from the command line. Thus the input metafile is identified as **mygroup** and not **mygroup.mta** in the command line above. Path information should be included with the input metafile name, if necessary. The names used by X-12-ARIMA to generate output files are taken from the specification files listed in the metafile, not from the name of the metafile. The example shown above would generate output files named **series1.out**, **series2.out** and **series3.out**, corresponding to the individual specification files given in the metafile **mygroup.mta**, rather than a comprehensive output file named **mygroup.out**. To specify alternative output file names, one would simply add the desired output file names to each line of the input metafile, as illustrated below:

```
c:\research\specs\series1 c:\research\alternative\series1
c:\research\specs\series2 c:\research\alternative\series2
c:\research\specs\series3 c:\research\alternative\series3
```

4.2.1.2 Running X-12-ARIMA in single specification mode for multiple time series

To run X-12-ARIMA on many series using the same specification commands for each series, one needs to create a data metafile. A data metafile can have up to two entries per line: the complete file name (and path information, if required) of the data file for a given series, and an optional output file name for the output of that series. If an output file name is not given by the user, the path and file name of the data file will be used to generate the output files. Note that in a data metafile the extensions must be specified, along with the file name, and path (if the data files are not located in the "x12a" directory). The data files are processed in the order in which they appear in the data metafile. The options used to process each data file are provided by a single input specification file identified at run-time. This means that all the data files specified in the data metafile must be in the same format. For example, to process the data files **series1.dat**, **series2.dat** and **series3.dat**, the data metafile should contain the following:

```
series1.dat
series2.dat
series3.dat
```

This assumes that all these data files are in the "x12a" directory. If they are stored in the "c:\research\data" DOS directory, the metafile should comprise:

```
c:\research\data\series1.dat
c:\research\data\series2.dat
c:\research\data\series3.dat
```

To run X-12-ARIMA with a data metafile, use the following command line:

```
x12a myspec -d mygroup¶
```

where **mygroup** is the data metafile, **-d** is a flag which informs X-12-ARIMA of the presence of a data metafile, and **myspec** is the single input specification file used for each of the series listed in the data metafile. Note that input metafiles must be saved using the **.dta** extension (e.g. **mygroup.dta**). Nevertheless, when the name of the data metafile is provided in a command line, as illustrated above, only the file name should be included, without the extension. Path information should be provided with the data metafile name, if appropriate.

The file names used by X-12-ARIMA to generate output files are derived from the data files listed in the metafile, not from the name of the metafile. The example given above would generate output files named

series1.out, **series2.out** and **series3.out**, corresponding to the individual data files given in the metafile **mygroup.dta**, not a comprehensive output file named **mygroup.out**. To specify alternative output file names for this example, one should simply add the desired output file names to each line of the data metafile, as illustrated below:

```
c:\research\data\series1.dat c:\research\alternative\series1
c:\research\data\series2.dat c:\research\alternative\series2
c:\research\data\series3.dat c:\research\alternative\series3
```

Each time X-12-ARIMA is run, a log file is produced. This includes a summary of modelling and seasonal adjustment diagnostics for every series or specification file processed. When X-12-ARIMA is run in multi-specification mode, the log file is stored with the same name and directory as the metafile. Conversely, when X-12-ARIMA is run in single specification mode, the log file is stored with the same name and directory as the data metafile. In both scenarios the log file is stored with the extension, ".log". For example, the command:

```
x12a -m mygroup
```

runs each of the specification files stored in **mygroup.mta** and stores user-selected diagnostics into the log file, **mygroup.log**. If a single series is processed, the output directory and file name is used along with the ".log" file extension to form the name of the log file (e.g. **myseries.log**). Users can specify which diagnostics are stored in the log file using the "savelog" argument found in the series, composite, transform, x11, x11regression, regression, pickmdl, estimate, check, slidingspans, and history specifications. If an error is identified in one of the specification files in a metafile run, a listing of all the input files with errors is provided in the log file.

More details on which diagnostics can be stored in the log file can be found in the "X-12-ARIMA Reference Manual".

4.2.2 In OxEdit

Creation of specification files, running of X-12-ARIMA, and analysis of the analytic output can all be performed using OxEdit

OxEdit is a reliable text editor capable of handling large files. It can run macros keyed to file extensions and it is freeware. OxEdit can be downloaded from the following website:

<http://www.oxedit.com>

Macros have been created for Oxedit to call X-12-ARIMA directly from the editor. Please contact TSAB to obtain the necessary macros.

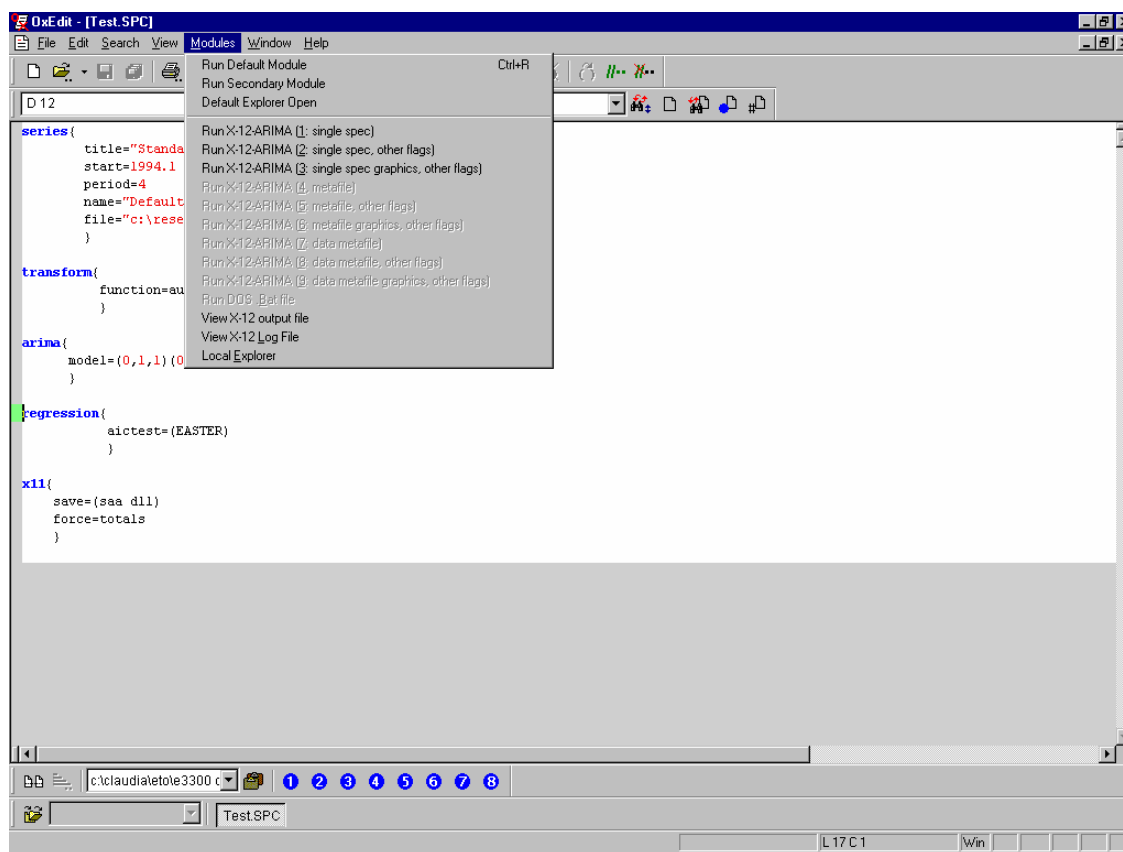
Instructions on how to install OxEdit and the X-12-ARIMA macros are provided below.

NOTE: The installation of programs to ONS computers is governed by ONS policy which must be adhered to.

1. The following instructions are based on the assumption that X-12-ARIMA has already been installed to the directory 'c:\x12a'. If this is not the case, you will need to amend the files below to reflect the path you have X-12-ARIMA installed to.
2. OxEdit has been set-up to run with X-12-ARIMA version 0.3. ONS currently uses version 0.2.8 and this guidance relates to version 0.2.8. If you have version 0.3 the installation is simpler, and you will only need to follow these instructions from step 5.
3. Once OxEdit has been installed, copy 'x12a.tool' and 'x12.def' to OxEdit's installation directory.

4. Copy the remaining files (runx12.bat, runx12v.bat, runx12vb.bat, runx12vl.bat, runxall.bat, viewlog.bat and viewout.bat) to X-12-ARIMA's installation directory. This is typically 'c:\x12a'.
5. You now need to activate the X-12-ARIMA macros. The macros assume you have X-12-ARIMA installed in 'C:\x12a'. If this is not the case, you must edit the files below to reflect this. To do so select 'view->preferences->Add/remove modules'.
6. Click 'load from' and navigate to the ox-edit installation directory (c:\program files\oxedit) and double click 'x12a.tool'.
7. Click 'ok', and close.
8. If the installation has been successful, when you open up a .spc file it should be colour coded.
9. Clicking 'modules->run X-12-ARIMA(single spc)' will bring up options to run X-12-ARIMA.

Figure 4-1 Modules available to run a single spec file



10. Other modules can run additional X-12-ARIMA functions, for which specific knowledge may be required.

4.2.3 In the Office for National Statistics mainframe

An interface will be written on the Central ONS Repository for Data (CORD) for the X-12-ARIMA program. This version will take data stored on the central database and will enable users to apply time series analysis methods using X-12-ARIMA.

7 LENGTH OF THE SERIES

7.1 Introduction

Any series of more than three years in length can be seasonally adjusted by X12ARIMA. However, if the series is very short, the program does not have a large amount of data to work with; thus it may have problems finding significant evidence of, for example, Easter effects or trading days. Very short series are likely to have large revisions as new data has the capacity to greatly change the estimates of the seasonal factors. On the other hand a very long series will have data which will not relate to the pattern of the current series, for example how much would 1970s data reflect the current pattern? What new information would it give which would still be accurate?

7.2 How much data does X12ARIMA need?

X12ARIMA has a few absolute minimum for certain functions to work:

- 3 years of data is the minimum for X12ARIMA to model or do any sort of seasonal adjustment.
- 5 years and 3 months (5 years and 1 quarter) of data is needed for X12ARIMA to automatically fit an ARIMA model and to calculate correctly all the criteria to test the models (especially the average forecast error for the last-3 year). If there is less data then the user can impose their own model.
- There must be at least 5 years of data including the forecast for the automatic selection of the seasonal moving average to work, and to allow for evolving seasonality. Constant seasonality is used if there is less than 5 years data.
- X12ARIMA uses all the data available when fitting an ARIMA model.

7.3 Short series

Short series are considered to be those where less than 5 complete calendar years of data are available. X12ARIMA will seasonally adjust the series assuming strict stable seasonality, i.e. the seasonal factors do not change over the span of the series.

If a short series needs to be seasonally adjusted, please refer to the following guidelines:

- If less than 3 years of data are available, the seasonally adjusted version of the series should not be published until more observations are available. At least 3 years of data are required for X12ARIMA to operate. If it is important to monitor the behaviour of the series, then seasonal adjustment will probably need to be judgmental. Use external information or patterns in related series if possible to seasonally adjust the series.
- If 3-5 years of data are available, consider not publishing the seasonally adjusted version of the series until more observations are available. X12ARIMA can use ARIMA modelling or provide trading day and Easter adjustments, but they are generally of very poor quality and subject to large revisions as future observations become known. The option of constraining of seasonally adjusted series to the raw annual totals is not available

7.3.1 *Methods available to seasonally adjust short series*

Three options are available to seasonally adjust short series with 3-5 years of data. These are:

1. To use the ARIMA option in X12ARIMA to forecast and backcast the series in order to have enough data to run the seasonal adjustment.

2. To seasonally adjust the actual data with no ARIMA modelling (using x11 spec only)
3. To link the series with another that covers a longer time period and then seasonally adjust the resulting linked series.

7.3.1.1 Option 1: the ARIMA extrapolation method

The first method is to use the ARIMA option to forecast and backcast the data. Because the automatic model selection only operates with more than 5 years of data, a simple form of model, such as a $(0,1,1)(0,1,1)$ with a log transformation if appropriate, should be fixed in the "arima" spec. The forecast spec should have 1-2 years of "maxlead" in order to have a series of 5 to 6 years of data becoming available. The "x11" spec can be run with the automatic selection of the seasonal moving average to make sure that the selected moving average is the most appropriate for the behaviour of the seasonal component. Trading day and Easter regressors can be used to estimate those effects, but it should be kept in mind that the adjustments could be of very poor quality and subject to large revisions as future observations become known.

7.3.1.2 Option 2: seasonal adjustment without ARIMA modelling

The ARIMA extrapolation method has some limitations that need to be taken into account. The use of the $(0\ 1\ 1)(0\ 1\ 1)$ model means that the forecasts will have a stable seasonal pattern, which will necessarily be the average seasonal pattern of the actual data. Therefore the effect of adjusting the extrapolated data will be much the same as adjusting the actual data without extrapolation. In this situation, it makes sense to simply use the x11 spec without any modelling. The only drawback to this method is that it is not possible to use ARIMA modelling to identify outliers or calendar effects; however, such identifications are unreliable in short series, as mentioned above.

7.3.1.3 Option 3: Linking

The third option available is to link two series together to provide a longer series of 5 years in length. Linking involves using a related series (the indicator series), which has the same general behaviour as the short series, to extrapolate the series back. This option is only appropriate where an indicator series is available for months/quarters prior to short series becoming available. Where an overlap is available between the two series, they can be linked to provide a longer series that can be adjusted using X12ARIMA. Advice on the linking of series is available from MQ division.

7.3.1.4 Choice of methods

The choice of which approach to follow depends upon a number of factors.

- Options 1 and 2 may be easier to implement; thus if time is limited one of these options may be more appropriate.
- Option 3 can only be used if there is an indicator series which can be used to extend the short series back. The effectiveness of this method will depend upon how similar the behaviour of the indicator series is to the short series being adjusted.

More advice is available from MQ division.

Seasonal adjustments of short series are likely to be subject to large revisions as new data becomes available. Outliers and breaks will have a particularly large effect on short series. So it is especially important that, if they exist, these effects are identified (usually through knowledge of the series) and removed.

7.4 How much data is recommended

Although anything under 5 years is considered short, it is desirable to have more than 5 years to ensure acceptable stability on update. At least 10 years is necessary to ensure that the adjustment of the first year is unlikely to be revised extensively. If the series has between 5 and 10 years of data, it may be modelled and adjusted in the same way as for a longer series, but it must be recognised that there may well be revisions (including perhaps changes of model and recognition of previously undetected breaks) when more data are added.

At the other extreme, a long series may have discontinuities or large changes in seasonal pattern. However, if it is necessary to have a seasonally adjusted series for the whole of a long span, it is generally better to adjust all the data than to divide it into separate sections. The only exception to this is if there is such a radical change of definition or coverage that it is not sensible to regard the series as a single whole.

Regressors should be included in the regARIMA model if there is a trend or seasonal break in the series history. One criterion for using these regressors is whether there is enough data before and after the break to enable reliable identification. Among other things, this depends on how irregular the series is; the greater the volatility, the greater the necessary length of data.

Because there are so many variable factors, it is not possible to lay down any hard and fast rules as to how many years of data are necessary to obtain a publishable adjustment. Anything under 5 years is very risky, and should be published with strong warnings and only if the series is of such high importance that some guidance on its evolution is essential. Between 5 and 10 years the series becomes more reliable, but there are still risks. It is best to base the decision to publish or not on some objective criteria, such as the M and Q quality statistics, with additional warnings if the series is published and there are large revisions on update. With 10 or more years the presumption is that adjustments should be published; this will be overridden only if there is some serious problem indicated by the diagnostics, for instance an indication of serious instability from the sliding spans analysis.

7.5 Cross-references

- Trend and seasonal breaks.

- Additive Outliers.

8 CONSISTENCY ACROSS TIME

8.1 Introduction

When seasonally adjusting a time series it is necessary to decide whether the annual totals of the seasonal adjustment should be constrained to equal the annual totals of the original data. This may depend on the time series, and how the data will be used. Typically this consideration is important for National Accounts data which has specific requirements for accounting purposes.

If the annual totals are to be constrained then the “force” option should be included within the specification file. This adjusts the seasonally adjusted estimates to ensure that yearly totals of the seasonally adjusted estimates equal the yearly totals of the original estimates. A separate problem can sometimes arise when monthly and quarterly versions of the same series are seasonally adjusted independently, resulting in inconsistent seasonally adjusted estimates. There are no obvious solutions to this particular problem, which requires careful consideration on a case-by-case basis.

8.2 Annual Constraining

8.2.1 *What is Annual Constraining?*

In general, the yearly total of seasonally adjusted estimates is not exactly equal to the yearly total of the original estimates. There are procedures in the X-12ARIMA program which ensure that the seasonally adjusted estimates and the original estimates have a similar average level over any 12 month (or four quarter) period, but the agreement is rarely exact. If exact agreement is required, X-12-ARIMA includes a routine which constrains the seasonally adjusted estimates so that the yearly total equals that of the original estimates. It does this in a way that closely preserves the period-to-period changes observed in the unconstrained seasonally adjusted estimates.

There is no mathematical reason why the annual totals of the original estimates and seasonal adjusted estimates should agree. Indeed, if trading day adjustments are made, there are good reasons why they should not (because the seasonally adjusted estimates have been scaled to an “average” pattern of trading days, while the original estimates correspond to the actual pattern of days in the year). However, these constraints are often applied, usually to make series more consistent for users.

Occasionally it may be desirable to constrain the seasonal adjustment so that financial years (or some other yearly totals) of the seasonally adjusted estimates and the original estimates are equal.

8.2.2 *When should Annual Constraining be Applied?*

The main reason for applying annual constraining is so that users are presented with a consistent set of annual totals for a time series, regardless of whether they are looking at the original or seasonally adjusted estimates. The trade-off is that there is a loss of optimality in the seasonal adjustment, particularly when seasonality is evolving quickly, or when trading day effects are large. The following four criteria should be used when deciding whether or not to apply annual constraining:

- 1) **Use:** consider how the data are used and any stated preferences of users. This criterion will sometimes override the other criteria. For example, if the series are part of the National Accounts dataset, they will probably need to be constrained.
- 2) **Concept:** decide whether annual constraining makes conceptual sense for your time series. Consider, for example, how the timing of the observations in your time series relate to calendar years

(thinking carefully about the way the data are collected and compiled) and what they represent. It may be apparent from this that annual constraining is inappropriate for your series. Consider the following examples: the observations in the Labour Force Survey are a sequence of overlapping 3 month averages; Retail Sales data are collected on a 4, 4, 5 week pattern rather than a calendar month basis; prices, earnings, claimant count and many other series are point in time series, for which averaging to the same annual averages might have little presentational importance.

- 3) **Continuity:** if a time series is already being seasonally adjusted, then strong reasons are needed for changing the approach to constraining. Avoid switching between applying and not applying annual constraining every few years.
- 4) **Effect:** look at the E4 table of the X-12-ARIMA output to see if the differences (or ratios if multiplicative) between the constrained and unconstrained data are large. If they are negligible, and the other criteria are satisfied, then annual constraining may be preferable.

8.2.3 Annual Constraining and Revisions

If annual constraining is applied, then it is important to ensure that revision policy take this into account. For example, a revision policy that revises the last month every time a new monthly observation is added to a series would change the additivity of the previous year when January data is published and December is revised. It is primarily for this reason that the revisions policy for the National Accounts datasets only ever involves revising observations in the current year, or, in certain major releases such as the Blue Book, whole years before that.

8.2.4 Annual Constraining in X-12-ARIMA

“Force” is the argument in the X11 specification of X-12-ARIMA that constrains the yearly totals of the seasonally adjusted series and the original series to be equal. By default X-12-ARIMA will not adjust the seasonally adjusted values to force agreement.

If the user wants to force the annual totals, the following specification should be used.

```
series{title="Example of constraining the annual totals"  
      start=1994.1  
      period=4  
      file="mydata.txt"}  
  
transform{function=log}  
  
arima{model=(0,1,1)(0,1,1)}  
  
x11{mode=mult  
  
force=totals}
```

There are two ways of using constraining. If the parameter is specified as “force=totals”, it is guaranteed that the totals will agree, but the totals as printed in the output may not due to rounding effects. If it is required that the totals should agree as rounded and as printed, use “force=both”. (There is another force option, “force=round”, which makes the rounded monthly figures sum to the rounded total, but does not constrain them to the unadjusted total.) Unless it is known in advance how many figures will be required in the final published figures, the recommended option to use is “force=totals”.

At least five years of data are required to use the force argument. Care should be taken when using the force argument in conjunction with the composite specification (see chapter 23).

8.3 Financial Year Constraining

Occasionally it may be desirable to constrain the seasonal adjustment, so that financial year totals (or some other yearly totals) for the seasonally adjusted and the original data are equal. The default operation of the force option uses calendar years, but there is an optional parameter "forcestart" which allows the use of any other 'year'. To use the standard UK financial year for a monthly series, include the specification "forcestart=april" or "forcestart=apr"; for a quarterly series, use "forcestart=q2".

Note that it is never possible to constrain seasonally adjusted estimates to agree with the original estimates over both calendar and financial years. To do so would require that the difference between the seasonally adjusted and original estimates over the first quarter should be exactly the same in every year of the series. Since seasonal factors must be allowed to evolve over time, and in any case may vary with the level of the series, such a restriction is unacceptable. If calendar and financial year constraints are required, two separate constrained time series must be produced. This requires two separate seasonal adjustment runs, with specification files differing only in the value of forcestart. The peculiarity of producing two different constrained series starting from the same unconstrained series should lead to questioning whether such a requirement is sensible.

8.4 Statistics

If the force argument is used, an extra table, D11A, appears in the output. This contains the seasonally adjusted series constrained to annual totals. Compare table A1 (original time series data), with table D11A (seasonally adjusted series with revised yearly totals), and note that the total for each year is the same.

A 1 Time series data

	1st	2nd	3rd	4th	TOTAL
2000	95.54	101.48	115.21	106.12	418.34
2001	113.28	121.72	122.30	114.64	471.95
2002	113.04	131.79	119.39	124.52	488.75

D 11.A Final seasonally adjusted series with forced yearly totals

	1st	2nd	3rd	4th	TOTAL
2000	99.96	99.97	110.77	107.64	418.34
2001	118.59	119.39	117.36	116.61	471.95
2002	118.06	128.97	114.80	126.92	488.75

8.5 Consistency across time

Another way in which non-additivity over time can arise in seasonal adjustment is when monthly and quarterly series are seasonally adjusted independently. Quarterly seasonal adjustment estimates and the sum of the seasonally adjusted estimates of the correspondent months can be very different, giving rise to extreme non-additivity between seasonally adjusted components and totals. This sometimes results in

a phenomenon conceptually similar to the one outlined in the Aggregate Series chapter. The key difference is that rather than a seasonal feature moving between components of a total, it moves between months of a quarter. Hence if all firms brought forward payment of bonuses to February rather than March in a particular year, it would not be recognised as seasonality in the monthly seasonal adjustment, but would in the quarterly seasonality. A similar phenomenon may be observed in energy series, where peaks in energy consumption switch between months according to the weather patterns. The weather also accounts for similar timing differences in harvests, affecting many agricultural statistics.

Under such circumstances there are three methods to seasonally adjust monthly and corresponding quarterly series:

- Indirect seasonal adjustment of the quarterly series by summing the corresponding seasonally adjusted months, which may result in sub-optimal seasonal adjustment of the quarterly series;
- Direct seasonal adjustment of the quarterly and monthly series, which may result in a loss of additivity;
- Direct seasonal adjustment with constraining to ensure additivity, which may lead to a distortion of the monthly series.

A decision about which method to use should be based on careful consideration of the uses of the series and analysis of the technical characteristics of the series involved. It is not possible to constrain monthly to quarterly series within X-12-ARIMA. There is a free-standing program (available from Time Series Analysis Branch) which effectively reproduces the methods of the “force=totals” option in the monthly and quarterly context.

8.6 Cross-references

- Aggregate series
- Revisions and updates
- Composite spec

9 AGGREGATE SERIES

9.1 Introduction

An aggregate series, also known as a composite series, is a series composed of two or more other (component) series. The component series can be combined in a variety of ways to form the aggregate series. An aggregate series itself may also be a component series of another aggregate series, therefore it is possible to have different levels of aggregation.

Many of the time series that are seasonally adjusted by the Office for National Statistics are aggregate series and therefore it is important to understand how to deal with the seasonal adjustment of them. For example:

- the number of visitors to the UK is the sum of the number of visitors from North America, the number from Western Europe and the number from other countries;
- total unemployment is the sum of male and female unemployment and also the sum of unemployment by age groups;
- constant price series in national accounts are equal to the corresponding current price series divided by a suitable deflator.

In practice, as well as in theory, component series can be combined in various ways to form aggregate series.

There are two distinct approaches to the seasonal adjustment of aggregate series:

- **Direct Seasonal Adjustment** - this involves seasonally adjusting the aggregate series without reference to the component series;
- **Indirect Seasonal Adjustment** - this method involves seasonally adjusting the individual component series, and then combining the resulting seasonally adjusted components to obtain the seasonally adjusted aggregate.

The two methods do not usually produce the same results. The direct and indirect methods produce equivalent results only under very restrictive assumptions, i.e. when no calendar or outlier adjustment is made, the decomposition is additive, and no forecasts are used. In practice, such conditions are rarely met and the differences in the series produced under the two approaches can be significant depending on the series concerned. Section 9.2 will discuss this problem of inconsistency between the approaches in more detail. Section 9.3 provides a checklist of the issues that need to be considered when choosing between the options used to perform a seasonal adjustment of an aggregate series. Section 9.4 considers ways of treating the series in order to restore coherence between the components and the aggregate.

9.2 Additivity of components and seasonal adjustment

The issue of consistency between the component parts and the aggregate after seasonal adjustment is usually referred to as one of 'additivity' even though the aggregate may be formed from the components by operations other than addition. It is important to realise that there is an inherent contradiction between the quality of the seasonal adjustment and the consistency across series. The goal of additivity can conflict with the primary purpose of seasonal adjustment, that of helping users to interpret a series' behaviour. For example, it seems reasonable to assume that, for each point in time, seasonally adjusted

male unemployment and seasonally adjusted female unemployment should sum to the seasonally adjusted total unemployment. In practice, using the direct approach to seasonal adjustment – seasonally adjusting the male, female and total series separately – will not normally achieve this additivity. The indirect approach of deriving seasonally adjusted unemployment by adding the seasonally adjusted male and female series will guarantee additivity but will lead to an inferior seasonal adjustment of the total series.

This problem of loss of additivity with the direct seasonal adjustment is particularly striking in situations where a seasonal feature in a time series switches between its components. An example of this in the ONS is the car production series, where seasonal peaks in production can be switched between cars for the home market and cars for export markets according to market conditions and car manufacturers' international production plans. The differences between the direct and indirect seasonal adjustments for these series are large and suggest that this might be happening. A direct seasonal adjustment is therefore used for total car production and no constraining procedures are introduced to reconcile total production with production for the home market and production for export. This lays open the possibility of each of the component series moving in the same direction, but the total series moving in the opposite direction in a particular month – but this is a consequence of the non-additivity and optimising the interpretation of each individual series.

Another situation that gives rise to difficulties is where the total series can be analysed in many ways. This is particularly the case for labour market statistics. The unemployed for example, can be split by sex, age, region, duration of unemployment, ethnic origin, educational attainment, etc. If one were to seasonally adjust series for unemployment by age band and add them together, this would result in a different indirect seasonal adjustment of total unemployment to that derived indirectly from the male and female series and both would be different from that obtained by directly seasonally adjusting the total. Under such circumstances additivity can only be achieved, using standard seasonal adjustment methods, in one of three ways:

- i. indirect seasonal adjustment at the lowest level of multi-dimensional disaggregation, e.g. seasonally adjusting males and females separately in each age band,
- ii. constraining out any non-additivity in the seasonally adjusted series,
- iii. restricting the method of seasonal adjustment so that it is entirely linear and results in a completely additive series.

Each of these is theoretically problematic and in some cases difficult in practice as well. All are potentially distortionary in their effects on the seasonal adjustment of one or more series in the dataset. However, examples of applying each of these exist in official statistics: the first characterises Eurostat's approach to seasonal adjustment of European industrial production, resulting in the seasonal adjustment of thousands of series; there are many examples of the second in the ONS, including extensive constraining of the Labour Force Survey series; and the Bank of England currently uses the third to adjust the monetary aggregates dataset.

The analyst faces a difficult choice. The direct approach is most likely to deliver a seasonally adjusted series that allows users to understand its underlying behaviour. However, for presentational and analytical reasons, users may require seasonally adjusted results that preserve the identities that are present in the non-seasonally adjusted data. In the circumstances the analyst has to choose between three options:

- Indirect seasonal adjustment, which may result in sub-optimal seasonal adjustment of the total series;
- Direct seasonal adjustment, which may result in a loss of additivity;

- Direct seasonal adjustment with constraining to ensure additivity, which may lead to a distortion of the component series.

The next section discusses the circumstances under which the different options should be chosen.

9.3 What type of adjustment should be used.

The decision on the level at which to seasonally adjust an aggregate series needs to be taken on a case by case basis. Listed below are some factors that should be considered when making that decision.

9.3.1 Factors favouring indirect seasonal adjustment

The indirect adjustment may be preferred for one or more of the following reasons:

1. It enables information about series to be used at the level at which it is known, that is to say the lower levels (e.g. the estimation and application of prior adjustments for seasonal breaks for motor car series in the IoS dataset).
2. It enables appropriate filtering of different types of data within a time series dataset (e.g. trade in services where many different data sources, some monthly, some quarterly, some annual, need to be treated differently).
3. It ensures consistency across different datasets (e.g. where a component is used in two different parts of the national accounts).
4. It guarantees additivity between components and totals. Therefore, an indirect adjustment will ensure that the seasonally adjusted components combine to equal the seasonally adjusted aggregate.
5. Disaggregated data often needs to be seasonally adjusted anyway to satisfy user/Eurostat/economic accounts needs.
6. Disaggregated data is sometimes more important to users than an aggregate (e.g. unemployment is more important than the total Labour Force or total working age population)

9.3.2 Factors favouring direct seasonal adjustment

Indirect seasonal adjustment does not necessarily result in a good quality seasonal adjustment at the aggregate level, possible reasons for this might be:

1. That adjustments or other processes occurring between seasonal adjustment and production of the final headline aggregate might re-introduce seasonality.
2. Component time series are not independent of each other and their multi-variate properties generate different seasonal dynamics at an aggregate level (for further explanation see the book fair example - > as this hypothetical situation is an excellent illustration of this), the essential problem is that non-seasonal component series can combine to form a highly seasonal aggregate.
3. Similar to the previous point, is the case where a dataset is all part of the same sample survey. The further the series is disaggregated, the greater the contribution of sampling variability to movements in the series. In this case, seasonality is harder to estimate, with greater potential for revisions. A potential consequence is that seasonal adjustment at a disaggregated level results in much higher I/S and I/C ratios and therefore longer moving averages are used than would be the case at an

aggregate level. The result is that the seasonal adjustment is more sluggish than it should be. The seasonal adjustment is performed to help users interpret short-term movements in the time series ONS presents and indirect adjustment might build in a much slower response to changes than is necessary, resulting in a more volatile seasonally adjusted series and potentially lagging users ability to perceive signals inherent in the data.

4. The more series there are to seasonally adjust, the more time-consuming the monitoring and re-analysing becomes. In extreme cases this becomes completely unmanageable and no attempt can be made to do anything other than run the adjustments on default settings.

9.4 Constraining to preserve additivity

For aggregate series where it is deemed essential to use direct seasonal adjustment and where users demand that additivity be preserved, it is necessary to constrain the seasonally adjusted component series so that they are consistent with the seasonally adjusted total series. One method of ensuring that the difference between the direct and indirect seasonal adjustment is small is to use the same model and consistent prior adjustments when seasonally adjusting the total series and the component parts. The chosen model, the one most appropriate for the total series will not necessarily be the most appropriate model for component series so the seasonal adjustment of the component series may be sub-optimal. Even if this approach is used, total may still not equal the sum of the components. In these circumstances the difference (total – sum of components) needs to be distributed across the component series. There are three ways of doing this:

- All of the difference attributed to the largest series or to the least significant series
- The difference allocated to all series, in proportion to the size of each series
- The difference allocated to all series, in proportion to the size of the irregular components of each series

9.5 Other considerations and related topics

In considering whether or not to undertake a composite analysis of an aggregate series, as well as considering the points discussed in section 3 as they relate to the particular series or data set concerned, a couple of other issues may be important, particularly as each case will be different

Size of the data set and the time available may determine what is feasible in terms of seasonal adjustment.

Does the indirect adjustment have residual seasonality? If so then a simple solution could be direct adjustment? The composite adjustment may also help to identify problems in component series.

Are there any outliers in the aggregate, and/or the component series that have been identified, and replaced that could be adversely effecting the indirect seasonal adjustment of the aggregate series, or are there outliers that have been identified in the direct adjustment that have not been identified in the adjustments of component series?

9.5.1 Using X12-ARIMA to choose the seasonal adjustment

If it is unclear which type of adjustment is needed, and the decision is to be based on the quality of the seasonal adjustment, the X-12-ARIMA program can produce diagnostics to aid in the decision between the direct and indirect methods. In order to compare direct and indirect adjustments in the X-12-ARIMA

program, the **composite** spec can be used to produce diagnostics to evaluate both methods. There is more information about the composite spec in Chapter XX.

Related Topics

- How to run X-12-ARIMA
- Sliding spans
- History diagnostics
- Composite spec

10 REVISIONS AND UPDATES

10.1 Introduction

When a new data point for the original estimates becomes available, more information is available concerning the seasonal pattern and the underlying trend of the time series.

This additional information may lead to revision of the published seasonally adjusted and trend estimates. The first part of this chapter explains how seasonal adjustments can be updated, and the rationale for each method. The second part looks at issues that may influence revisions policies.

10.2 Types of updating

The following methods are widely used for updating the time series outputs from a seasonal adjustment.

Annual updating involves an annual review and assessment of each directly seasonally adjusted time series using the existing prior adjustments. The seasonal adjustment settings and prior corrections are analysed and improved where possible. X-12-ARIMA diagnostics, such as the M and Q statistics, Sliding Spans, and Revision History are used to assess the appropriateness of the seasonal adjustment parameters. Time Series Analysis Branch is responsible for reviewing the seasonal adjustment parameters for all ONS directly seasonally adjusted time series.

Current updating involves running X-12-ARIMA every month or quarter with the latest available time series data to derive the latest seasonally adjusted estimates. When each of these seasonal adjustment runs is performed, the seasonal adjustment options used are not the defaults, but those determined during the most recent annual update. The lengths of the moving averages; prior adjustments for Easter effects and trading days; and the type of ARIMA model used, should all be specified in the specification file. This type of updating is used across the ONS to obtain the latest seasonally adjusted estimates.

Forward Factors involves using forecasted seasonal factors (Table D10a of the X-12-ARIMA output) derived at the time of the annual update. This results in revisions to the seasonally adjusted estimates being applied only once a year, when the new forward factors are estimated. This option is generally adopted by GSS Departments where there are system constraints,

Within ONS annual updates are carried out by Time Series Analysis Branch, while individual branches responsible for the original data carry out current updating.

10.3 Is it better to fix or re-estimate seasonal adjustment parameters?

The following is a general guide to which seasonal adjustment parameters should be fixed and which should be re-estimated each time a current update is run:

- The ARIMA model should be fixed in the X-12-ARIMA “arima” specification;
- The appropriate transformation should be fixed in the “transformation” specification;
- Easter and trading day regressors should be included in the “regression” specification as variables, and not left to the automatic selection using the AIC test.
- The dates for additive outliers and level shifts defined by the automatic outlier detection should be converted into fixed regressors in the regression-ARIMA. In this way the effect of the regressors will be fixed in the model whilst the parameter adjustment is re-estimated every month. An alternative

approach for known impacts is to also fixed the parameter estimate so that it is not re-estimated at each time period;

- User-defined seasonal regressors should be used in the “regression” specification to correct for a seasonal break. These regressors should be converted into permanent priors only if the time series have multiple seasonal breaks;
- The trend and seasonal moving averages should be fixed in the “x11” specification for important time series.

10.4 Revisions

Each new original data point that becomes available will impact on the estimates of the seasonal component and trend component for previous periods. Each seasonal adjustment update will potentially cause revisions along the length of the seasonally adjusted estimates.

Major revisions are typically applied after the addition of new data to the immediately preceding period and to the corresponding period one year prior. For example, when new original data becomes available in December 2006, the seasonally adjusted and trend estimates will be revised from November 2006 backwards in time, with potentially larger revision for November 2006 and December 2005. However, if the annual totals are constrained, revisions cannot be made without revising the whole year, and so the revision would not generally be made.

Often, the nature of the seasonal adjustment will imply that revisions should occur at other time lags. For example, time lags of two and three periods might also be revised, if these show large changes when a new data point is added. In general, the seasonal adjustment must be revised back to any periods where the raw data has to be revised. If revisions to raw data are large then the seasonal adjustment of neighbouring points may also need to be revised. A graph of the adjustment, before and after the revisions, should be checked to see if this is generally the case.

If a problem with seasonal adjustment is found between annual updates, this should be corrected as soon as possible. This will mean that seasonally adjusted and trend estimates may be revised.

If a seasonally adjusted series is particularly smooth, then in the annual update it may be necessary to revise merely the previous two years, or even the previous year only. Conversely, if the adjustment is dominated by the irregular component, it may be necessary to revise up to four years prior.

A revision policy should be determined taking user requirements and revisions of raw data into account. The revision history diagnostic and the revision triangle method are useful indicators for determining which revision policy is best from a seasonal adjustment perspective.

The decision as to when to revise, and how much data to revise, is important. Revisions policy can have a major impact on user confidence and the quality of published data. Therefore, a revisions policy that takes into account all the factors detailed above is an important element for producing high quality seasonally adjusted series. In general, revisions and revisions policy for seasonally adjusted and trend estimates will depend on the nature of the time series and the client area responsible for publishing the data.

11 THE REGARIMA MODEL

11.1 Introduction

The regression-ARIMA part of the X-12-ARIMA program precedes the seasonal adjustment itself. It modifies the time series in such a way that the seasonal adjustment process will produce higher quality estimates. The modification covers three aspects:

- Series extension - the series is extended forwards (adding forecasts) and backwards (adding backcasts). This produces a longer span of data to input in the seasonal adjustment process leading to better quality seasonal adjustment, particularly at the ends of the series. A consequence of this is that revisions, when new observations, become available are likely to be lower.
- Calendar effects – effects associated with the arrangement of the calendar are removed from the series. This improves the estimation of seasonal effects and increases short-term interpretability.
- Outliers, breaks and other changes – the series can be adjusted for unusual and disruptive features such as a sudden and sustained drop in the level of a series (a 'level shift'). Removing such features makes the seasonal adjustment more robust by preventing them from distorting the subsequent estimation of seasonality. However, these features typically represent the real world behaviour of whatever the time series is measuring so the features need to be reinstated into the series after seasonal adjustment is complete.

11.2 What is a regression-ARIMA

Regression-ARIMA is the name of the statistical modelling facility in X-12-ARIMA. It enables two types of models, an ARIMA model and a regression model, to be fitted to a time series.

ARIMA, or Auto-Regressive Integrated Moving Average, is a widely-used family of models for time series, which take account of trend and seasonality in the data. In X-12-ARIMA, the program can be allowed to choose the most appropriate form of ARIMA model for an individual series, using the model fitting criteria built into the program, or the user can specify the form of ARIMA model to be applied.

The regression part of regression-ARIMA refers to the options which enable the basic ARIMA model to be enhanced with additional regression variables. These might take the form of dummy variables set up to model the effect of, for example, a level shift, or might be a direct explanatory variable (e.g. a mean temperature variable).

In technical terms, the regARIMA can be described as a linear regression in which the error terms follow an ARIMA process, rather than a white noise process:

$$y_t = x_t' \beta + z_t$$

where x_t is a vector of regression variables and the error term z_t follows an ARIMA process.

Once the model has been specified, it can be used to produce forecasts and backcasts. The variables of the regression part of the model give estimates of calendar and other effects which we do not want to be present in the data when the series goes through the seasonal adjustment process.

This chapter covers how the regARIMA model should be used set up. The issues are grouped under three main headings:

1. Transformation of the series
2. Specification of the ARIMA part
3. Specification of the regression part

Although it is convenient to separate headings 2 and 3 in this way, they are not separate and sequential processes. If we do not include appropriate regression variables (particularly for effects like level shifts) it may be impossible to produce a satisfactory ARIMA model. In practice there is often an iterative cycle of adding new regression variables and re-specifying the ARIMA model until a satisfactory set of diagnostics is obtained. This process is discussed in more detail in Section 11.6 below.

11.3 Transformation of the series

This is the first step in fitting a regression-ARIMA model to a series. The regression-ARIMA can be fitted either to the original series or to the log transformed original series. Which of these is appropriate depends on whether the series is additive or multiplicative (see chapter 15). When the series is additive, the regression-ARIMA model is fitted directly to the original series. When the series is multiplicative, the model is fitted to the log transformed series. The effect of a log transformation is to change the scale of a series and to turn multiplicative effects into additive ones. This is done purely for the purpose of fitting the model. Any forecasts or prior adjustments estimated from the model are then converted back to the original scale.

Guidance on choosing whether a multiplicative or additive decomposition is most appropriate for a series can be found in chapter 15.

11.4 Specification of the ARIMA part of the model

The purpose of ARIMA modelling is to identify systematic structural features in the history of the series. We assume that these features will continue to be present in the future and can be use to forecast future values. There are many modelling and forecasting techniques; what distinguishes the ARIMA approach is the selection of features it looks for and the systematic method of identifying them. The ARIMA method provides a wide range of possible models, which have been found very effective in modelling typical socio-economic series showing trends, seasonality and business cycle effects.

This section gives only a brief and largely non-technical explanation of ARIMA modelling. TSAB should be consulted if more technical detail is needed.

In the simple case of a non-seasonal series, the form of an ARIMA model is that the current value is expressed as some combination of previous values plus a random quantity called the innovation; innovations are thought of as being independent samples from a normal distribution with zero mean and constant variance. Within this general framework, we can distinguish the three types of effect that the ARIMA model can include:

- Autoregressive: the value depends on some linear combination of previous series values.
- Moving average: the value depends on some linear combination of previous innovations.
- Integrated: the autoregressive and moving average effects apply to differences of the values, rather than the values themselves.

In each of these categories, the number of values involved is referred to as the *order* of the effect.

A few examples may help to illustrate these distinctions. In what follows, the series value at time t is represented by x_t and the innovation by ε_t .

a. A first order autoregressive model:

$$x_t - \phi_1 \cdot x_{t-1} = \varepsilon_t$$

b. A first order integrated moving average model:

$$x_t - x_{t-1} = \varepsilon_t - \theta_1 \cdot \varepsilon_{t-1}$$

c. A second order autoregressive and second order moving average model:

$$x_t - \phi_1 \cdot x_{t-1} - \phi_2 \cdot x_{t-2} = \varepsilon_t - \theta_1 \cdot \varepsilon_{t-1} - \theta_2 \cdot \varepsilon_{t-2}$$

(Note that all terms involving x have been moved to the left of the equals sign and all terms involving ε to the right; this is the conventional arrangement.)

The order of a model is conventionally abbreviated in the form (p, d, q), where p is the order of the autoregressive component, d is the order of differencing and q is the order of moving average. The examples above are represented in this notation as (1, 0, 0), (0, 1, 1) and (2, 0, 2) respectively.

If there is a seasonal component, this means that there is an effect which repeats at annual intervals. This is modelled in the same way as a conventional effect, except that the dependence is on values occurring one, two, etc. years ago rather than one or two periods ago. This means that, for a monthly series, a seasonal autoregressive component will involve a relationship between x_t and $x_{t-12}, x_{t-24}, \dots$, and similarly for seasonal moving averages or differences. A seasonal series will usually have a regular component, defined by (p, d, q) as above, and a seasonal component defined by its own parameters written as (P, D, Q). The combination of these two components is indicated by putting the two brackets together, with a subscript on the seasonal indicating the seasonal period. Thus, the very well-known seasonal model, the “airline model” (so called because Box and Jenkins used it to model monthly numbers of airline passengers), is written (0, 1, 1)(0, 1, 1)₁₂ – though the subscript may be omitted when it is obvious (as is done for example in Section 11.4.2 below).

Because of the effect of sampling variation or irregularity in the observed series, it is often possible to find different ARIMA models that will fit a given time series. One fundamental idea in choosing a model is to take the simplest model which will give a satisfactory fit, where simplest means having the smallest number of parameters. This principle is known as “parsimonious parameterisation”, or parsimony for short. The process of searching for a model often involves adding parameters to avoid specification errors and then seeing which previously chosen parameters may be removed in the interests of parsimony. The principal tool in testing for a satisfactory fit is the autocorrelation of the residuals, which are the estimates of the innovations. The innovations should be independent, so any significant serial correlations in the residuals could be an indication of a deficiency in the model. There are guidelines to indicate how a model should be changed to remove particular patterns of serial correlation, though these will not normally be needed, since X-12-ARIMA provides tools to help automate the process.

There are basically two approaches to automatic model selection:

1. search through a predetermined list of candidate models (usually arranged in order of increasing complexity) to find a satisfactory fit;

2. start with a simple model and allow the program to add terms successively (up to predetermined limits on the number of terms) until the fit meets some criterion of goodness of fit, the terms added at each stage being determined by the deficiencies in the previous fit.

The first approach has been used in all earlier versions of X-12-ARIMA, and is based on work by Statistics Canada. The second approach has been added in the most recent version and is based on the TRAMO method of Gomez and Maravall at the Bank of Spain. The user of the latest X-12-ARIMA has the choice of either method, or may impose a choice of model if the automatic choice is not satisfactory. The first method is provided by the **pickmdl** spec, the second by the **automdl** spec. The two approaches are outlined in the following sections, 11.4.1 and 11.4.2

11.4.1 Automatic model selection with the automdl spec

The **automdl** spec proceeds by successively improving the first simple model. Terms are added or modified when the specification tests show that the current model is inadequate and the modified model is better; terms are removed when they are no longer significant. This continues automatically within user-specified limits on the complexity of permitted models.

The limits on complexity are specified in terms of the maximum values of the parameters p, d, q, P, D and Q. There are default values for these maxima which are built into the program, but these may be overridden by the user if necessary. The maxima for d and D are specified by the **maxdiff** argument, the maxima for p, q, P and Q by the **maxorder** argument. The default values are $d \leq 2$, $D \leq 1$, $(p, q) \leq 2$, $(P, Q) \leq 1$, which may be represented in terms of the arguments as **maxdiff=(2 1)**, **maxorder=(2 1)**. For normal purposes these default values will be adequate; they should not be overridden unless the program fails to produce a satisfactory model using them.

If no satisfactory model is produced with the default values, alternatives may be tried. The only alternative with **maxdiff** is (2 2), which should not be tried until all other possibilities are exhausted. For **maxorder**, the table below shows alternative, higher values which empirical research has shown may sometimes be preferable. Any model produced using these limits should be carefully scrutinised to see that all its diagnostics are satisfactory.

Table 11-1 Alternative values for the maxorder argument

Series length	Monthly series	Quarterly series
<6 years	2,1	2,1
6-10 years	3,1	3,1
10-15 years	4,1	3,1
>15 years	4,2	4,2

Although in most cases the program will select an ARIMA model, sometimes it will not. There are other parameters which may be varied to improve the outcome in difficult cases, but their use requires care and experience. It is recommended that TSAB be consulted if problems arise.

11.4.2 Automatic model selection with the *pickmdl* spec

When this spec is used for automatic model selection, X-12-ARIMA selects a model from a limited list of models, subject to certain model selection tests. The program uses as default a list based on the original research by Statistics Canada, which consists of five models. Three more models have been added to the list by the ONS, so it now includes the following 8 models:

(0,1,1)(0,1,1)
(0,1,2)(0,1,1)
(2,1,0)(0,1,1)
(1,1,1)(0,1,1) *
(0,2,2)(0,1,1)
(1,1,2)(0,1,1) *
(2,1,1)(0,1,1) *
(2,1,2)(0,1,1)

The models that are marked with an asterisk are those added by the ONS. If necessary, the user may add to this list or provide a completely different list; this requires careful testing before an alternative can be considered safe, and it is recommended that TSAB be consulted before undertaking any change.

The selection tests include:

1. A test that the residuals are not serially correlated.
2. Testing the MA part of the model for evidence of overdifferencing.
3. Within-sample forecast error. For this test to pass the absolute average percentage error of within-sample forecasts in the last three years should be less than 15%.

Criteria 1. and 2. aim at misspecification, which if nothing else would give false forecasts even if the fit of the data is satisfactory. Criterion 3 is a quantitative measure of within-sample forecast accuracy in the last three years of the series.

The models are tested sequentially, in the order shown above, which is clearly that of increasing complexity. If no model passes all tests then the first and simplest one - the (0,1,1)(0,1,1) model - is used, but only for the purpose of estimating the regression effects; no forecasts are generated if no model passes the tests. The default operation of the method is that the first model tested which passes all the tests is accepted. This can be specified explicitly by putting **method=first**. The alternative is **method=best**, which means that all models are tested and the acceptable model with the lowest forecast error is accepted.

If the program also performs automatic identification of outliers (using the **outlier** spec), the default is to identify them simultaneously with the estimation of the first ARIMA model that is tested and take them as given for the estimation of the other models. One can alternatively perform outlier identification for each of the tested models by using the optional argument **identify=all**. This is the recommended option, as it entails simultaneous model and regression estimation, which increases the efficiency of both estimations.

Other optional arguments are available, which are detailed in the full X-12-ARIMA documentation. Normally these should be left at their default values.

11.4.3 Fixing the ARIMA model

X-12-ARIMA also provides the option of imposing directly the lag structure of the ARIMA, and even the values of some or all of its parameters. This can be done with the **arima** spec. The model that is specified in this way can be used in any action performed on the series (i.e., outlier detection, regression, etc.). This is likely to be useful in two circumstances:

1. When the user has specific knowledge about the real world processes underlying the production of the data, and this knowledge will not be captured by the automatic modelling procedure. For example, in the case of monthly financial data, it may be known that there is a quarterly cycle due to the way institutions operate as well as the annual cycle. This can be captured by adding an extra bracket to the ARIMA model with a subscript 3. Similarly, the quarterly cycle in the operation of the Labour Force Survey might impose an extra structure on the LFS unemployment series.
2. When we want to ensure that the model identified by the automatic process is not updated during concurrent adjustment. Generally concurrent adjustment uses all available data to adjust the current value, but experience shows that frequent switching between models due to updating can introduce undesirable instability in the seasonally adjusted output. It is therefore general practice to keep the form of the chosen model fixed between annual re-analyses (though the parameter values are re-estimated each period).

The appropriate form of specification for the standard **arima** spec is:

arima{model=(p, d, q)(P, D, Q)}

where p, d, q are the orders of the regular part of the model and P, D, Q are the orders of the seasonal part. The seasonal part may optionally be followed by the seasonal period (either 4 or 12), and if necessary the model may be extended with other brackets followed by an appropriate period. (In common with other X-12-ARIMA specs, the commas in the inner brackets may be omitted provided the figures are separated by spaces.) The other possible arguments of **arima** are concerned with pre-specifying parameter values, and should not be used. Note that if **arima** is used in a spec file there must not also be an automatic modelling spec (**automdl** or **pickmdl**).

11.4.4 Identifying a model manually

Since in most cases the automatic modelling options will produce satisfactory models (at least for the limited forecasting horizons used in seasonal adjustment), it should never be necessary for users of X-12-ARIMA to undertake manual model identification from scratch. The program can produce the necessary autocorrelation outputs, and there are textbooks which describe the procedure, but it requires considerable experience to produce reliable results. This is definitely not recommended.

The one situation in which manual identification may be justified is when automatic identification has produced a model which, while satisfying the tests, still has some unsatisfactory features. For example, although the combined test on the serial correlations of the residuals may be passed, there may still be some individual significant correlations at fairly low lags. It may then be justifiable to try manual refinement of the automatic model. In the example case, it could be worth while adding an extra coefficient at the appropriate lag to the AR or MA component; if the extra coefficient is significant and the significant serial correlation has been removed, the extra term may be justified.

The process of manual refinement needs a degree of skill and care, and advice should be sought if in any doubt. The program provides a number of diagnostics which are helpful in this process, particularly the model identification statistics (AICC, BIC etc.). Details of the use of these may be found in the full X-12-ARIMA documentation.

11.5 Specify the regression part

In addition to the relations within a time series that are captured by the ARIMA model, the series may also be affected by external deterministic factors, such as outliers, trend or seasonal breaks, calendar effects, etc. In this case the only satisfactory approach is first to subtract these deterministic effects, and then fit an ARIMA model to the "linearised" series. However, the exact size of the deterministic effects is usually unknown and needs to be estimated, which cannot be done until the ARIMA structure is determined. The most effective way of doing this is by use of a regression-ARIMA model, which encompasses both the deterministic effects and the ARIMA structure and estimates them simultaneously. The previous section described how to specify the ARIMA part of the model. This section describes how the regression part should be specified.

Specifying the regression part of the model essentially means deciding which variables to include. The recommended option is that a regression variable is included only if theory or knowledge of the series indicates that a regression effect is possible, and this is proved with statistical tests. What is meant by "theory or knowledge of the series" is whether or not the user expects a particular variable to be important, any prior knowledge for the same or a similar series, etc. This theoretical question is answered in the relevant chapters. (For example, the trading day chapter describes in which cases a trading day effect is possible). This chapter concentrates on the statistical tests that should be used to decide whether a particular effect is indeed significant or not.

The rest of this section is organised as follows: Sub-section 11.5.1 considers the question of types of regression variable. Sub-section 11.5.2 describes the built-in regression variables that are available to users and Sub-section 11.5.3 how to specify other variables which are not available in the program. Sub-section 11.5.4 describes how to include variables optionally, basing the decision on statistical tests of significance in X-12-ARIMA.

11.5.1 Types of regression variable

X-12-ARIMA has a wide range of program-specified variables, which will cover many of the common situations; those most often used are described below. Users can also include variables of their own to capture effects which are not provided by the built-in types. One important difference is that users do not need to provide the actual values of any built-in variables, since the program will automatically generate them from the specification. Values for user-defined variables must be input to the program in the same format as series values; this includes values for any forecast periods as well as those corresponding to actual observations.

All the built-in variables are assigned to a variable *type*, and users must specify an appropriate type for any user-defined variables. The concept of variable type is important in two contexts:

- The program output in some areas gives a significance test for the joint effect of all variables of a given type, rather than each variable separately; this is explained in more detail in the discussion of estimation and inference (Section 11.6).
- The type of a variable determines whether its effect is combined with the seasonal effect in adjusting the series.

The latter point needs further explanation. All the regression variables are removed from the original series in order to estimate the "true" seasonal, meaning the effect which is approximately periodic with period one year, but for many variables this removal is not permanent, since we wish to regard their effects as part of the information about real world effects that the adjusted series conveys. For example, we could identify a trading day effect and a level shift in the same series. Both can distort seasonal patterns, and both are removed before estimating the seasonal. However, the trading day effect is part of the variation which we wish to remove, since we regard it as just a calendar-related distortion of the true

seasonal, so it is not replaced when obtaining the seasonally adjusted series. The level shift, on the other hand, represents some real world effect which we wish to be able to observe (and perhaps try to explain) in the seasonally adjusted series; its estimated effect is therefore replaced in the final adjusted series.

As a general rule, all calendar-related variables (trading day, holiday, length of month, leap year etc.) are removed from the final seasonally adjusted series, while seasonal and outlier variables are not. If user-defined variables are assigned the non-specific type **user**, because they do not fit with any of the built-in types, they are not removed from the final series.

It should be noted here that some effects which might be dealt with through regression variables could in some circumstances be handled instead by prior adjustment. For example, if there is a series break of known size, a multiplicative prior adjustment would remove it. The effect of this on the adjustment process is almost exactly the same as that of a level shift regression variable of the same size. The prior adjustment is removed from the series before seasonal estimation, and is restored in the final adjusted series. The one difference is that the prior adjustment is not removed in calculating the final trend, while the level shift is.

11.5.2 Program-specified regression variables

The program-specified variables include:

- a. A **constant term**. This variable will rarely need to be used, as its presence implies a deterministic trend of order equal to the total number of differences, which will in most cases be difficult to justify. A constant term might be used though if the regular part of the ARIMA model includes a zero difference term. Another case is when the ARIMA model was selected with the **automdl** method, and it was found that a constant is significant. On the other hand, it is not advisable to use a constant together with the **pickmdl** model selection routine.
- b. **Deterministic seasonality** (fixed seasonal effects). This can be in the form of either seasonal constants or of trigonometric regression variables, but one cannot include both. The advantage of trigonometric variables is that their use requires fewer parameters.

Although deterministic or almost deterministic seasonality will be the case for many series, the use of seasonal constants is not particularly recommended. This is because deterministic seasonality is equivalent to a seasonal MA coefficient of 1, in a seasonal (0,1,1) ARIMA model. Although in this case it is more efficient to reduce the order of seasonal differencing and capture the seasonal fluctuations with seasonal constants and a stable seasonal ARMA, for the purpose of seasonal adjustment this is not likely to add much value; in practice a seasonal (0,1,1) ARIMA with a MA coefficient equal to, for example 0.95 will be indistinguishable from deterministic seasonality. In addition, if one has to reduce a seasonal overdifferencing to a stable seasonal ARMA with seasonal means as just described this has to be done manually, as X-12-ARIMA does not do it automatically, as it does for non-seasonal overdifferencing. For these reasons the use of seasonal constants is discouraged. Trigonometric regression variables on the other hand can be used occasionally, for example in short monthly series where the alternatives of using seasonal dummies or seasonal differencing cost 12 observations.

It should be noted that there is a difference between constants and deterministic seasonals and all other regression variables, in that these variables are not adjusted out of the original series before estimating the x11 seasonal; these variables are there simply to make it possible to forecast the series.

- c. **Trading day variables**. There is a wide range of alternative variables, which correspond to the various different ways in which the number of occurrences of the days of the week within the month

can affect economic activity. Details on the trading day effect, the alternative variables, and when it is appropriate to use each of them are given in the relevant chapter. What should be emphasised though is that only one trading day variable can be included in a regression.

- d. **Easter effect.** Again, there is a range of built-in variables that correspond to the different ways in which Easter affects economic activity, detailed in the Easter effect chapter. Other holiday variables, aiming in accounting for the effects on economic activity from other holidays, can be used as well.
- e. If the users know that an important event has happened that is likely to have affected the series they should include a regression variable that takes into account this effect. Depending on the nature of the event and the duration of its effects this variable may be an "additive outlier", "level shift", "temporary change", or "ramp". Automatically detected outliers can also be included in the regression, if so requested by the **outlier** spec. In contrast to the above "intervention variables", these outliers are unknown and can be detected only from the data.
- f. Finally, there may be cases where a variable does not affect a series in the same way throughout the sampling period. This can happen in cases of policy changes or other important events. In such cases **change of regime** variables can be used to account for the change. In particular, change of regime variables can be used for either trading day regressors or deterministic seasonality. The case for deterministic seasonality with change of regime is when a **seasonal break** is suspected, of which more will be said in chapter 17.

11.5.3 User-specified variables

Although the list of variables that are built-in in X-12-ARIMA will in most cases cover the user's needs, there will be some special cases for which the program has not provided. In such cases the users can define and use their own regression variables.

One may for instance wish to use variables other than those specified by the program in situations such as:

- a. **Non-calendar data**, that is, data collected on a different basis from that dictated by the monthly calendar. If for instance monthly flow data are collected on the last Sunday of every month, then one has months with 4 or 5 complete weeks, rather than months with 4 or 5 (e.g.) Wednesdays as handled by the standard trading day variable; here a specially designed variable should be used instead of a trading day one.
- b. If the series is affected by a program-specified variable but with a delay then one has to input the **lagged** program-specified variable as a user-specified one, because X-12-ARIMA does not have an option to do it automatically.
- c. A **change of regime** that cannot be captured by the built-in "change of regime" variables.

To specify user-defined variables, the variables must be named and their names listed in the **user** argument of the **regression** spec. Unless all user variables are to have the same type, their types must be listed in the **usertype** argument in matching order to the names in the **user** argument. The values of the variables for all periods (including forecast ones) must be listed in a table, which may be embedded as **data** in the spec file or (preferably) stored in a text file which is referenced by a **file** argument.

11.5.4 Optional variables and statistical tests

The procedure implied above is that variables may be included speculatively in the model, tested for significance and then removed if not significant. This then requires another run to re-estimate the model

without the variables. There is an alternative to this, which is to specify that certain variables are to be tested to see if their inclusion improves the model fit, included if there is improvement and excluded if not. The only variables which may be treated in this way are certain calendar-related variables (specifically trading day and Easter) and variables of type **user**.

The criterion which is used is the overall model fit statistic AICC. The program will estimate the model with and without the optional variables and select the version which gives the better (i.e. smaller) value of AICC. If several variables are specified for this treatment, they are tested in up to three groups: first any trading day variables, next any Easter variables, finally any user variables. All the variables in a group are tested and included or excluded together. The variables to be tested in this way are specified in the **aictest** argument; they must normally also be listed in either the **variables** or **user** arguments, depending on whether or not they are built-in variables. Exceptionally, if **easter** is included in **aictest** it may be omitted from **variables**; in this case, the program automatically considers different lengths $w = 1, 8, 15$ for the **easter** effect, and compares the best of these with the option of no **easter** effect.

Another type of optional variable controlled by statistical tests is an automatic outlier variable. If the spec file includes an **outlier** spec, the program will search each possible time point to see if an outlier included at that point would be significant. All outliers so identified are added to the model. For more details of types of outlier and options on inclusion see Chapter 14 on Outliers.

11.6 Estimation and inference with regression-ARIMA

Three issues are relevant to the estimation of a regARIMA: the order of differencing, regression effects and the ARMA specification of the stationary series. This order corresponds to their order of importance in a seasonal adjustment. If the series is differenced fewer times than required, the regression effects and their statistical significance may be totally wrong. Misspecifying the ARMA of the stationary series, on the other hand, is not as harmful for the purpose of seasonal adjustment as omitting a regression variable.

Once the model is estimated one needs to check whether it is properly fitted or not. In most cases this can be taken for granted for the ARIMA part of the model, as it has most probably been automatically selected. With regard to the regression part of the model though, one needs to test whether the regression variables are significant or not and whether additional variables should be included. The usual "t-statistic", "chi-square" and "AIC" tests are used.

11.6.1 Automatically run tests

Usually all the tests one needs are automatically produced. In particular, the program gives the t-statistics for all regression coefficients. It also gives chi-squares for groups of variables wherever they are more meaningful as a group (for example, it is more meaningful to test whether the six trading day variables together show a significant effect, rather than testing if, say, the number of Tuesdays in a month has any significant effect).

When a variable is tested individually, it should be considered statistically significant if its t-statistic is beyond the critical value for the significance level required. As a rule of thumb, a t-statistic greater than 2 in absolute value is significant at the 5% level. If, on the other hand, a group of variables such as trading day regressors are jointly tested, then the appropriate statistic is the chi-square. For this statistic the program automatically gives the probability (P-value), and any value less than 5% (0.05) is usually considered significant.

In some cases it is not necessary to carry out an explicit test, since the terms have been included automatically because they are significant. This applies to level shifts and other outlier terms selected by the **outlier** spec. In other cases, the appropriate type of test should be carried out on the individual terms or groups of terms, and the terms retained or deleted as appropriate. For example, it is good practice to

verify that the p-value or the t-value associated with trading day or Easter terms selected by the **aictest** are statistically significant.

11.6.2 Non-automatic tests

Although in most cases the automatic tests will be sufficient, one might often wish to test a hypothesis that cannot be tested automatically. This can often be done by some manipulation or combination of the outputs. There are several such examples:

1. Alternative models may require separate runs of X-12-ARIMA, because it is not possible to specify one as a specialised version of another (these are referred to as non-nested models). In such cases the model selection statistics (AIC, AICC, BIC) may be used. The model with the smallest value of the chosen criterion is preferred. There are limitations on the use of these statistics. In particular, the dependent variable in the regression must be identical in all the cases being compared, meaning that the series span, order of differencing, transformation and the outliers included in the model must be the same.
2. The autoregressive and moving average parameters do not have t-statistics attached. They do have a standard error attached to their values, however, and the t value is the ratio of the two. Normally the automatic model selection will have ensured that only significant terms are retained, but if a model has been manually refined such checks are needed.
3. If we wish to test whether the parameters attached to two variables are equal, we can rearrange the model to provide a direct test. We define a new variable which is the sum of the two variables, and include this in the model in place of one of the two. The coefficient of the remaining single variable now represents the difference in the two parameters, and its significance is a test of the difference.
4. If we wish to test the joint significance of a group of variables, and we are not able to define them as a group by type for which a chi-squared value is produced, we can carry out runs with and without the variables concerned and compare the results. The simplest comparison is to use one of the model selection statistics (AICC or BIC). An alternative, which can give a numerical level of significance rather than a yes-no result, is to compare the log likelihood values (from the same table as the model selection statistics); the difference in log likelihood is distributed as chi-squared with degrees of freedom equal to the number of extra parameters.

These are merely examples of the ways in which hypotheses about the model fitting may be constructed and tested to meet particular circumstances. The principles exemplified here may be adapted to other situations. If difficulties arise which users cannot easily solve, in the first instance TSAB should be consulted.

11.7 Summary of implementation instructions

11.7.1 On first analysis

- Consider whether there are special effects which might suggest that automatic model selection is inappropriate. If not, use an automatic model selection with an appropriate forecast horizon. Backcasts may also be used unless the series is very long.
- If the **pickmdl** method is used, backcasts can be produced by specifying **mode=both**
- Include in the regression spec all the variables that might have an effect on the series. For example, consider whether the process being measured is likely to be affected by calendar effects like Easter

or days of the week; see if there are any known events in the history of the series which might have caused breaks; see if any external variables like weather might affect the series.

- Run X-12-ARIMA to generate model selection and regression statistics.
- If the chosen model selection procedure does not give a satisfactory model, consider varying the selection procedure (see 11.4 above) or manually refining the best model produced by the automatic procedure. As a last resort use the default (0 1 1)(0 1 1) model.
- Check the statistical significance of the regression variables, as explained above. If the tests show that the set of variables that should be included is different from the one that was actually included, amend accordingly and run X-12-ARIMA again, repeating the process of ARIMA model selection.
- The simultaneous estimation of the regression and ARIMA parameters may (rarely) lead to convergence problems. If the maximum number of iterations are run without convergence having been achieved, one should increase this maximum number of iterations and try again. This can be done with the **maxiter** argument in the **estimate** spec. If the estimation has not converged, **it is wrong to use the output of the last iteration as the estimate**. If convergence cannot be achieved, more options are given in the X-12-ARIMA manual (or consult TSAB).

11.7.2 At the time of the annual re-analysis

- Use the automatic model selection procedure that was used in the previous analysis. If a different model is chosen, examine the test statistics to see whether the new model is substantially better than the previous one. If the difference is marginal and the statistics for the previous model are still satisfactory, consider imposing the previous model in the interests of stability.
- Include in the regression spec all the variables that were found significant in the previous run, plus any excluded but found to be on the margin.
- Run X-12-ARIMA to generate model selection and regression statistics.
- Check the significance of the regression variables as before. If any which were previously included are now strongly non-significant (with a t-value lower than 1), they can be dropped. If any regression variables that were previously included have a t-value in between 1 and 2 it is better to keep the variable in the model to avoid unacceptable revisions of previously published seasonally adjusted data. That variable should be tested again at the following annual re-analysis.
- All changes to the model should be considered before introduction to see what effect they have on historical data. Stability may often be more important than a marginal improvement in fit.

11.7.3 For production running

- Using the **arima** spec, impose the ARIMA model that was selected at the time of the last re-analysis.
- Use the **forecast** spec to specify the appropriate forecast and backcast horizons.
- Use the regression variables that were selected at the time of the re-analysis. Include automatic outlier identification only for the span of data that became available after the re-analysis. Any effects selected for inclusion by **aictest** should instead be fixed in the **variables** argument, and **aictest** should be removed.
- Fix neither the ARIMA nor the regression parameters.

- Run X-12-ARIMA.

12 TRADING DAY

12.1 Introduction

Trading day effects are those parts of the movements in a time series that are attributable to the arrangement of days of the week in calendar months. For example, a month containing 5 Saturdays is likely to show a higher level of sales than a month containing 4 Saturdays. As with seasonal effects, it is desirable to estimate and remove trading day effects from time series to help interpretation.

X-12-ARIMA estimates trading day effects by adding regressors to the regARIMA model. Section 2 illustrates the problem of the arrangement of the calendar in more detail. Section 3 describes when to adjust for trading day effects, section 4 provides details of the different types of regressors used by the X-12-ARIMA programme to adjust for trading day effects. Section 5 describes the recommended procedure to adjust for trading day effects and explains how to implement the results in a production environment. Section 6 gives details of related options and topics and section 7 gives guidance on dealing with non-calendar data.

12.2 The Arrangement of the Calendar

Trading day effects arise because the number of occurrences of each day of the week, in a month, differs from year to year. An example of the arrangement of the calendar problem is shown in the table 12-1, where the number of occurrences for June is calculated across three years, 2002-2004.

Table 12-1: Day composition for June 2002, 2003 and 2004

Year	Number of days		
	2002	2003	2004
<i>Monday</i>	4	5	4
<i>Tuesday</i>	4	4	5
<i>Wednesday</i>	4	4	5
<i>Thursday</i>	4	4	4
<i>Friday</i>	4	4	4
<i>Saturday</i>	5	4	4
<i>Sunday</i>	5	5	4

These differences will cause regular effects in some series. For example, a production series where no work takes place over weekends will have one less working day in June 2002 than in June 2003, which will have one less working day than June 2004. Thus, it is likely that the series has a slightly lower value in June 2002 than in June 2003 and June 2004 without reflecting the long-term trend.

Those differences are not genuine movements of the production, but are just due to the fact that the numbers of working days in a factory are 20, 21 and 22 respectively. This regular effect can be identified and removed by the use of regressors in the RegARIMA model.

Trading day effects may also reflect how data are recorded more than when the event actually happens, for example sales for the weekend may be recorded on the following Monday. This will still cause a regular effect that should be removed from the series.

If the same analysis is carried out at a quarterly level over a decade, it is possible to comment that:

- There are between 90 and 92 days per quarter with a minimum of 12 occurrences of a particular day and maximum of 14 in any one quarter.
- Quarter 1 is only affected by the leap year effect;
- Quarter 2 is never affected by trading day variation, since the total number of days is always equal to 13;
- Only quarter 3 and quarter 4 are affected by the trading day variation. This implies that trading day effect on quarterly series is less strong than on monthly series.

For the majority of series trading days effects are insignificant for quarterly series.

Another effect due to the arrangement of the calendar is leap year. A leap year is a year with one extra day inserted into February. The leap year is 366 days with 29 days in February as opposed to the normal 28 days. Also this effect will cause regular variation in some series and therefore needs to be removed to make a proper comparison between Februarys or quarter ones.

X12ARIMA enables effects associated with arrangement of the calendar to be removed from the series using regressors (for trading days) or constant variables (for leap year) in the RegARIMA. This adjustment is done before the seasonal adjustment takes place. The adjusted series is usually referred to “trading day adjusted” series if the series is adjusted for trading days and leap year only or “calendar adjusted” series if the series is adjusted for trading days, leap year and moving holidays (i.e. Easter). Trading day adjusted series and calendar adjusted series are sometimes required by Eurostat and are routinely published in some European countries for their national accounts.

Another type of calendar-related effect is usually grouped with trading day effects, namely the effect of variation in the length of calendar periods (months or quarters). To understand the use of this, consider for example the comparison of industrial output in June and July. These may differ systematically for two reasons: firstly, the average output per day may differ from June to July due to differences in demand or other factors; secondly, July has one more day than June. With standard seasonal adjustment both these effects are regarded as part of the seasonal pattern, and the June and July seasonal factors will reflect both. For some purposes it could make the analysis easier to interpret if we regard the first effect as the “true” seasonal and the second as a calendar effect. If we apply a length of month adjustment (see Section 12.6 for details), the seasonal factors will reflect this true seasonal. The overall seasonal adjustment will be almost unchanged, because length of month and “true” seasonal together come to the standard seasonal. Note that the length of month adjustment treats all days as equal, so it is inconsistent with trading day adjustment. Also this is not “length of working month” but “length of calendar month”, so the example above is only realistic for an industry working seven days a week; in other cases an appropriate length of month variable could be defined as a user variable.

12.3 When to adjust for trading day

The X-12-ARIMA programme provides a number of diagnostics to test a series for the presence of trading day effects. Testing the statistical significance of trading day regressors is discussed in greater detail in section 5, which describes the recommended procedure for testing and adjusting for trading day effects. In general, where trading day effects are found to be statistically significant, the series should be adjusted to remove these effects from the final seasonally adjusted series. However, it is always important to look at the results of the trading day regression and try to relate it to the time series itself. If the results are counterintuitive (e.g. they suggest most car production takes place on Saturdays) then it is worth investigating whether there is anything in the recording of the data which is causing this result. It is better not to implement the trading results if they are very counterintuitive.

Trading day effects should not be estimated for the following types of data:

- Data that are collected on a 4,4,5 week pattern. That is to say, the recording periods in a year consists of a four times repeated pattern of a four week recording period, followed by another four week recording period, followed by a five week recording period. This means that the year is divided differently to the calendar periods described by months. This system of data collection does not, in general, exhibit trading day effects.
- Data that is collected at a point in time, for example a particular day in the month. However, if the collection day can occur on different days of the week, it may be that there is an effect depending on the day concerned; this can be estimated using a stock trading variable (see *tdstock* below).
- Data that are not collected in strict calendar months. There may, in this case, be some sort of trading day effect, however, this effect should not be estimated for using the regressors provided by X12ARIMA. For further information about adjusting such data contact the Time Series Branch.
- Quarterly data. Whilst it is possible to estimate trading day effects for quarterly flows data using X12ARIMA, trading day effects will, in general, cancel out within the quarter.

X-12-ARIMA provides different diagnostics to test for the statistical significance of trading day regressors.

12.4 Options available to adjust for trading day effect

As previously noted the X12ARIMA programme can run a regression to model certain effects that result from the arrangement of the calendar, such as trading day effects. It is possible for the user to define regressors in the **regression** spec. The X12ARIMA programme contains a number of predefined variables, contained in the **variables** argument, to adjust for a variety of calendar effects, including five different regressors that specifically adjust for the effects of trading days; four specifically designed for flow series and one for stock series. Three other options not specifically designed to adjust for trading day effects but related to trading day effects are discussed in section 6; they allow the user to adjust for length of month (*lom* option), length of quarter (*loq* option) and leap year (*lpyear* option) effects.

Using the **variables** argument in the **regression** spec allows you to specify one of the following regressors to adjust for trading day effects,

1. *tdnolpyear* this is used to estimate flow trading day effects
2. *td* this is used to estimate flow trading day effects and is a combination of the *tdnolpyear* variable and the *lpyear* variable
3. *td1nolpyear* this is used to estimate flow trading day effects
4. *td1coef* this is used to estimate flow trading day effects and is a combination of the *td1nolpyear* variable and the *lpyear* variable
5. *tdstock[w]* this estimates a day-of-week effect for stock data or inventories that are reported on the *w*-th day of each month.

Each of these options is discussed in more detail in the following table.

Variable name	X12ARIMA command	Comments
'tdnolpyear'	<pre>regression{ variables=tdnolpyear }</pre>	<p>It includes 6 day-of-week contrast variables and it is used for flow series only. The six constant variables are derived in a way that the level of the adjusted series is not distorted when the prior adjustment is estimated. This constraint makes the sum of the coefficients for each day of the week (known as trading day weights) equal to seven. Using the 6 parameters associated with each variable and the constraint it is possible to derive the seventh parameter. tdnolpyear assumes that each day has a different effect.</p> <p>The tdnolpyear regressor cannot be used in conjunction with td, td1coef, td1nolpyear, or tdstock[w] regressors. The following is an example of what to write into the spec file to activate the tdnolpyear option in the regression spec.</p>
'td'	<pre>regression{ variables=td }</pre>	<p>It includes the tdnolpyear regressor as well as estimating the effects of a leap year. The leap year effect is handled either by re-scaling (for transformed series) or by including the lpyear regression variable (for untransformed series).</p> <p>The td regressor cannot be used in conjunction with tdnolpyear, td1coef, td1nolpyear, or tdstock[w] regressors in the regression spec, or the adjust=lpyear, adjust=lom, adjust=loq from the transform spec.</p>
'td1nolpyear'	<pre>regression{ variables=td1nolpyear }</pre>	<p>It is a weekday-weekend contrast variable that can be used for flow series only. This is more parsimonious than the tdnolpyear option, as there is only one variable in the regression. The difference with tdnolpyear is that td1nolpyear assumes one effect for all the weekdays and another for Saturdays and Sundays rather than an effect for each day individually.</p> <p>The td1nolpyear regressor cannot be used in conjunction with td, tdnolpyear, td1coef, or tdstock[w] regressors in the regression spec.</p>
'td1coef'	<pre>Regression{ variables=td1coef }</pre>	<p>It is similar to the td regressor in the same way that td1nolpyear is similar to tdnolpyear. This means that td1coef includes the td1nolpyear regressor as well as estimating the effects of a leap year. The leap year effect is handled either by re-scaling (for transformed series) or by including the lpyear regression variable (for untransformed series).</p> <p>If the td1coef regressor is used neither td, tdnolpyear, td1nolpyear, nor tdstock[w] regressors can be used in the regression spec.</p>

Variable name	X12ARIMA command	Comments
'tdstock[w]'	<pre> Regression{ variables=tdstock[31] } </pre>	<p>It estimates day-of-week effects for inventories or other stocks that are recorded on the <i>w</i>-th day of the month. This allows the user to specify a value for <i>w</i> (from 1 to 31), where specifying 31 will mean that it is an end of month variable, as it will take this to be the last day of the month for those months with less than 31 days. Research suggests that this variable is rarely significant.</p> <p>If the tdstock[w] regressor is used neither td, tdnolpyear, td1coef, td1nolpyear, lom, nor loq regressors can be used in the regression spec. Furthermore the tdstock[w] variable cannot be used with quarterly data.</p>

For further information and description of handling trading day adjustment with regression model used in X-12-ARIMA see Findley, D.F. et al (1998) "New Capabilities and Methods of the X-12-ARIMA Seasonal Adjustment Program" *Journal of Business and Economic Statistics*, Vol.16 No.2.

12.5 How to adjust for trading day effects

Section 2 and 3 described when to adjust for trading day effects, whilst section 4 introduced the options available in X-12-ARIMA to adjust for different trading day effects. This section describes, firstly, how to identify the presence of trading day effects in a series, and secondly, the generally recommended process to adjust for trading day effects.

The general order of testing the significance of regressors is described in chapter 11, which discusses the regARIMA model. Chapter 11 explained how the Chi-squared test should be used to test for the significance of trading day effects. However, other methods also exist for detecting the presence of trading day effects such as the AIC test.

This section will describe three ways of identifying whether or not trading day effects are present in a series, and then the procedure for adjusting for trading day effects in a production run. Two different scenarios will be outlined for setting up the seasonal adjustment for the production run:

- Firstly producing prior adjustments that can be fixed for a year's production run, in both X12ARIMA and X11ARIMA based programmes;
- Secondly the recommended procedure of setting up the spec file for a production run using a regression variable, rather than permanent priors, which is an option that can only be used in X-12-ARIMA based programmes.

12.5.1 Testing for trading day effects with X12ARIMA

12.5.1.1 Spectral plot

The spectral analysis reveals whether or not significant trading day peaks are found in a seasonal adjustment. Two spectral plots are produced, one from the first differences of the adjusted series, adjusted for extreme values from table E2 of the output and a second of the final irregular component, adjusted for extreme values from table E3. From these plots, the X-12-ARIMA programme will estimate

whether any of the peaks at predetermined frequencies (the frequencies are determined by the cyclical nature of the trading day pattern) are significantly different to that of the neighbouring peaks. If the programme finds that peaks exist at the trading day cyclical frequency (see Findley et al (1998) for further information) it will return a warning message in the command prompt such as,

"WARNING: At least one visually significant trading day peak has been found in one or more of the estimated spectra."

This test is very sensitive and has a tendency to show trading day effects when other diagnostics don't (and, no doubt, sometimes when they don't actually exist). Nevertheless if this warning message is returned the series should be tested to estimate the significance of trading day effects. If this warning message is returned when a particular trading day regressor has been used, it may be necessary to test a different trading day regressor to see if that performs better. For example if the **td** regressor has been used, it is possible that using the **td1coef** regressor will perform better and could remove the trading day peaks. The performance of different trading day regressors can be assessed with the two tests described below but can also be assessed by their impact on the overall performance of the seasonal adjustment. If a warning message is still returned after a significant **td** variable has been included and if resources are limited, it may be better not to examine the series in more detail but just to retain the **td** variable.

12.5.1.2 The AIC test

The AIC test can be activated in the **regression** spec to evaluate whether or not a particular regressor is preferred, compared to not having that regressor in the model. For example the following may be specified,

```
regression{aictest=(td1coef)}
```

This will generate the following likelihood statistics in the output,

Likelihood statistics for model without td1coef

Likelihood Statistics

```
-----
Effective number of observations (nefobs)          138
Number of parameters estimated (np)                3
Log likelihood                                    130. 8304
Transformation Adjustment                         -980. 3013
Adjusted Log likelihood (L)                       -849. 4709
AIC                                                1704. 9417
AICC (F-corrected-AIC)                            1705. 1209
Hannan Quinn                                     1708. 5104
BIC                                               1713. 7235
-----
```

Likelihood statistics for model with td1coef

Likelihood Statistics

```
-----
Effective number of observations (nefobs)          138
Number of parameters estimated (np)                4
Log likelihood                                    130. 6730
Transformation Adjustment                         -980. 3013
Adjusted Log likelihood (L)                       -849. 6283
AIC                                                1707. 2565
AICC (F-corrected-AIC)                            1707. 5573
-----
```


Hannan Quinn	1712. 0148
BIC	1718. 9656

***** AICC (with aicdiff= 0.00) prefers model without td1coef *****

In the above example trading day effects do not appear to be present in this particular series, and so the **td1coef** would not be used to adjust this series for trading day effects. The **aicctest** argument compares the AICC statistics (these are in bold in the above example, only to highlight which statistics are compared) and depending on which model has the lower AICC statistic, will return the a line that states which model it prefers. The above example gives an AICC statistic of 1707.5573 for the model with a **td1coef** regressor and an AICC statistic of 1705.1209 for the model without the **td1coef** regressor. Therefore the model without **td1coef** is preferred.

Note that there is an additional option with **aicctest**, namely the argument **aicdiff** mentioned in the tables above. The purpose of this is to avoid extreme sensitivity to minor changes in the AICC criterion. If **aicdiff** is different from the default value of zero, there is a bias in favour of the simpler model (in this case excluding trading day variables); the tested variable is included only if it improves the AICC criterion by at least the amount **aicdiff**. This could be used in a spec which is routinely re-run in the annual re-analysis, to reduce the risk of instability on update.

Only one trading day regressor can be tested at a time with this option, and therefore to compare the performance of different trading day regressors the AICC statistic of each model would have to be saved and the programme re-run using the different regressor. The results of the **aicctest** can be saved in the log file using the **savelog** argument. For example the following **regression** spec would test the **td** variable against no **td** variable and save the results of this test in the log file.

```
regression{ aicctest=(td)
           savelog=aicctest}
```

The AICC statistics and other specified diagnostics that could help evaluate the performance of the seasonal adjustment, with (or without) particular trading day regressors could be compared by saving the log file of each run into another file, for example an excel spreadsheet.

The **aicctest** can be used to test the following regressors, **td**, **tdnolpyear**, **td1coef**, **td1nolpyear**, **tdstock**, **easter**, and **user**. NB the **aicctest** cannot test a specific **tdstock[w]** variable, only **tdstock**, which defaults to an end of month variable. If more than one type of variable is tested, the order in which these tests are carried out is first trading day regressors, second easter regressors, finally user defined regressors.

When resources are sufficient it is recommended to undertake a detailed analysis of the series taking into consideration the nature of the series and whether or not a particular regressor would seem appropriate. For example, in the case of a flow series, the four regressors, **td**, **tdnolpyear**, **td1coef**, and **td1nolpyear** should be tested.

If resources are slightly more limited the **td** regressor should be tested, and where the **aicctest** prefers the model with the **td** variable to use the **td** variable. If the spectral analysis returns a warning that trading day peaks are present when the **td** variable is specified, the user should, if resources permit, test other trading day regressors in particular the **td1coef** regressor.

The AICC statistic is one of the diagnostics that should be considered in choosing between regressors, where they are found to be significant.

12.5.1.3 The Chi-Square test and t-value

The Chi-squared test is used to test the joint significance of a group of trading day regressors. Therefore the Chi-square statistic will only be produced in the cases where the **td**, **tdnolpyear** and the **tdstock[w]** variables are used, as these variables include six day-of-week contrast variables. The **td1coef** and **td1nolpyear** variables are, for the purposes of trading day effects, only using one variable, the weekday-weekend contrast variable and hence the *t*-value will provide an indication of the significance of these variables. When a variable has been specified in the **variable** argument of the **regression** spec then *t*-values will be estimated for the individual regressors and the Chi-squared test will test for the joint significance of those variables that have six day-of-week contrast variables. In the following example the variables included in the **variable** spec are **td** and **easter[1]**. The Chi-squared test is testing the joint significance of the trading day variables only.

Regression Model

Variable	Parameter Estimate	Standard Error	t-value
Trading Day			
Mon	0.0089	0.00581	1.54
Tue	0.0013	0.00591	0.22
Wed	0.0256	0.00601	4.26
Thu	0.0109	0.00590	1.84
Fri	-0.0012	0.00609	-0.19
Sat	-0.0177	0.00605	-2.92
*Sun (derived)	-0.0278	0.00594	-4.68
Easter[1]	0.0281	0.01160	2.42

*For full trading-day and stable seasonal effects, the derived parameter estimate is obtained indirectly as minus the sum of the directly estimated parameters that define the effect.

Chi-squared Tests for Groups of Regressors

Regression Effect	df	Chi-Square	P-Value
Trading Day	6	167.19	0.00

As the above example shows, the *t*-values for some days are not significant, that is the coefficients are not significantly different from zero. However, the Chi-squared test for the trading day effect gives a P-value of 0.00 which is within the recommended 5% limit (i.e. equal or less than 0.05) for accepting significant trading day effects.

The *t*-values should be used as a test of significance for the **td1coef** and **td1nolpyear** variables, whereas the Chi-squared test should be used to check the overall significance of either the **td**, **tdnolpyear**, or the **tdstock[w]** regressors. If a variable is found to be significant it should, in general, be included in the model. If all the trading day regressors are found to be significant, the AICC test and diagnostics for assessing the quality of the seasonal adjustment should be used to determine which of them should be selected.

12.5.2 Adjust for trading day effects

Use one of the following three methods to adjust for trading day effects.

12.5.2.1 Regression variable with test of significance

Once it has been decided which trading day regressor should be used to adjust for the trading day effect, the **aictest** described in section 12.5.1.2 can be used to adjust for the effect. This means that every time X12ARIMA is run to estimate the seasonally adjusted series, AIC values are derived for models with and without the specified trading day variable and the optimal model will be used for forecasting. This option will provide users with the optimal seasonally adjusted series, but will generate higher revisions than the other two methods described in the following sections.

Problems:

- This option should not be used for production run because the trading day regressor variable is not fixed and frequent switching between models due to updating can introduce undesirable instability in the seasonally adjusted output. If there is some special reason why this option must be used, a non-zero value of the **aicdiff** argument should be used to reduce instability.

12.5.2.2 Regression variables fixed in the model

If a particular regression variable has been identified as significant and the best one to use for the seasonal adjustment, then the **regression** spec should contain in the **variable** argument all the regressors that should be included. For example, if, as in the previous example, **easter[1]** and **td** are found to be significant, the spec file, along with the other parameters that have been fixed should include the following arguments in the **regression** spec.

```
regressi on {variables=(easter[1] td)}
```

This is the recommended option for the production run because the form of the chosen regressor is fixed though the parameter values, the trading day factors, are re-estimated each period. When a particular trading day regressor is fixed in the model, this should be included even if it appears to no longer be significant. The decision of whether to remove the variable or not should be taken at the point of a seasonal adjustment review, so as to reduce the number of revisions.

Problems:

- If a strict revision policy is in place it is suggested not to use this option since the estimated trading day factors will change when new data becomes available.

12.5.2.3 Permanent prior adjustments

Once a particular regressor has been chosen (see above and Chapter 11, on regARIMA and Chapter 5 on Procedures for analysing series with X12ARIMA) then the trading day factors can be saved in order to obtain permanent prior adjustments. The following steps should be taken to obtain permanent prior adjustments.

1. Run the spec file with your chosen settings, e.g. de-composition, ARIMA model and so forth. Include in the **regression** spec with the **variable** and **save** arguments activated (NB, in order to save the trading day factors obtained from any trading day variable the appropriate argument is **save = td** for all variables). For example, if the **td** regressor was found to be the most appropriate the **regression** spec would look like this,

```
regression{variables=(td)
           save=(td)}
```

2. The save argument will save a text file with the trading day factors in the same directory that the log and output files are saved. If the name of the spec file was "filename.spc" this would mean that the trading day factors would be saved in a file called "filename.td".
3. If permanent priors are required for more than one year, then the **forecast** spec should be used to set the number of forecast periods and **appendfcst=yes** option should be included in the "x11" spec. For example, the following, assuming monthly data, would provide trading day factors for three years into the future,

```
regression{variables=(td)
           save=(td)}
forecast{maxlead=36}
x11{appendfcst=yes}
```

4. It is essential to save the .td file with a different name (e.g. 'filenamepp.td') otherwise it will be overwritten the next time the spec file would be run with the **variables=(td)** option.
5. In order to set up the spec file for a production run, remove the **regression** spec and use the prior adjustments that have been saved in step 4 by using the **transform** spec. Therefore, in the final spec file, used for production runs, the **regression** spec is not used and the **transform** spec is activated, as in the following,

```
transform {file=("filenamepp.td")
          format="x12save"}
          type=(permanent)}
```

It should be noted that the above example has given a certain method of saving and using the permanent priors in a particular format and so on. There are various ways in which permanent priors can be saved and used to transform the original series. For further information on the different options available see X-12-ARIMA Reference Manual.

If holiday (e.g. Easter) or user defined regressor variables are used to estimate other effects that are required to be used as permanent priors, then the factors must be multiplied together so that they are all in one file to be used in the **transform** spec. For example, if **easter[1]** (see chapter 13) and **td** variables have been estimated and the holiday and trading day factors have been saved (**save=(hol td)** is the appropriate argument in the **regression** spec, which saves the respective factors in separate files with the name of the spec file and the extension hol and .td respectively) then these factors should be multiplied by one another and following steps 4 and 5 will mean that the original series is adjusted by the prior adjustment file so that the seasonal adjustment is then performed on a series with these calendar effects removed, which should improve the quality of the seasonal adjustment.

Problems:

- With the permanent priors option, the adjustment for trading day factors is fixed as they are only based on the data that is available at the point in time when the trading day effects are being estimated (during the annual review). If the trading day pattern changes between annual reviews the

permanent priors will not capture this modification therefore the seasonal adjusted series will not be optimal.

- It is very complicated to set up and keep up-to-date.

Criteria for deciding which of the three methods should be used are as follows:

1. **Revision policy:** How strict the revision policy is determines which approach to use in the production runs. The use of permanent priors is the method that gives the minimum revisions, but also gives the less optimal seasonally adjusted series and it is very complicated to set up and keep up-to-date. The second option in terms of revision size is fixing the regression variable in the model. This is the recommended method since it balances a fairly good quality of seasonal adjustment with revision and practicality. Contrary, the use of a regression variable with a test for significance will provide with the optimal seasonal adjustment but with the biggest revisions of the three methods.
2. **Seasonal Adjustment Review:** for reviewing the parameters the regression variable with a test for significance should be used. It provides users with the optimal model to run the seasonal adjustment.

12.6 Related topics

Section 3 gave specific details on the options available in the X12ARIMA programme, which allow the user to adjust for trading day effects. Three related options that in effect also make adjustments related to average daily effects in flow series are: the length of month (*lom*), length of quarter (*loq*) and leap year (*lpyear*) options. Each of these options is discussed in the following table. These options are very rarely specified in practice.

Variable name	X12ARIMA command	Comments
'lom'	<pre>regression{ variables=lom }</pre>	<p>It can be used to adjust a series for effects resulting from the differing of the length of a month. The regression includes a length-of-month regression variable.</p> <p>It cannot be used with the td, td1coef, or tdstock[w] variables.</p>
'loq'	<pre>regression{ variables=l oq }</pre>	<p>It is the same as the lom option, but is used in the case of quarterly data.</p> <p>It cannot be used with the td, td1coef, or tdstock[w] variables.</p>
'lpyear'	<pre>regression{ variables=l pyear } OR Transform{ adjust=l pyear }</pre>	<p>It is a contrast variable for leap year effects. This variable can only be used with flow series and cannot be used in conjunction with the td, or td1coef variables. Adjustments for leap year effects can be made in the regression spec or the transform spec, but not both.</p> <p>As with the regression spec leap year adjustments cannot be made in the transform spec when the td, or td1coef variables are specified in the variables argument of the regression spec. Furthermore, to use the leap year adjustment in the adjust argument of the transform spec a log transformation must be specified in the power or function arguments of the transform spec.</p>

12.7 Non-Calendar Data

Some data do not align strictly to calendar months and some common problems that can arise under these circumstances are:

- 1) Bank holidays can switch between months. For example, in a 4-4-5 week pattern of collecting data, the late May Bank Holiday might fall in the 'May' collection period in some years and 'June' in others. These can sometimes be modelled in X12ARIMA. TSAB has more details.
- 2) Conventional trading day adjustments are usually inappropriate, but, particularly if the recording periods are not the same length between years, they might have a large impact on the series. Sometimes survey questionnaires ask for respondents to record the exact period their response refers to, and crude adjustments are made to calendarise the data at respondent level in the ONS. This is a difficult area; TSAB has some experience of dealing with this problem and can be contacted for advice.
- 3) 'Phase shift' effects might occur, particularly within a 4-4-5 week data collection pattern. This is because a pattern of 4 weeks, 4 weeks, 5 weeks repeated throughout the year adds up to only 52 weeks, or 364 days, but there are 365 days in a normal calendar year and 366 in a leap year. This means that the collection periods will be earlier and earlier in each successive year. Indeed, in surveys which operate under this regime, such as the Retail Sales Inquiry or the Labour Force Survey, either a 'survey holiday week' is taken or an extra week's data collection is added to a 4 week collection period every 4 or 5 years in order to roughly realign the collection periods with calendar months. However, these moving reference periods can have an impact on the time series, particularly in the presence of strong seasonal patterns. These can sometimes be modelled in X12ARIMA and removed from the data. Ask TSAB for more details, if you encounter this problem.

12.8 Cross-references

Other topics related to trading days include,

- Easter
- RegARIMA
- Prior adjustments

13 EASTER

13.1 Introduction

Easter is a moving holiday that can occur in either March or April, for monthly data, or in the first or second quarter for quarterly data. The effect on series caused by movements in the date of Easter needs to be removed from the seasonally adjusted series that we produce.

X12ARIMA estimates calendar effects, such as Easter effects, by adding regressors to the regARIMA model. This chapter explains when and how to adjust for the effects of Easter. Section 2 illustrates the Easter effects, section 3 describes when to adjust for Easter effects, section 4 provides details on the options available in the X12ARIMA programme to adjust for Easter effects, section 5 explains the recommended procedure to adjust for Easter effects and section 6 lists some related topics.

13.2 The Easter effects

The date of Easter Sunday is the first Sunday after the first full moon of the spring equinox, and based on the Georgian calendar can be anywhere between the 22 March and 25 April. As with seasonal effects, it is desirable to estimate and remove Easter effects from time series to help interpretation. The effect of Easter can be best understood by considering an example. We would expect sales of chocolate to be higher in the days and weeks before Easter. If Easter occurs in March all of the additional expenditure will occur in March. If however, Easter falls on the 25 April, the sales in March will be much lower than the previous case, and much higher in April. These effects need to be removed from our seasonally adjusted series. A similar but opposite effect may be present in series of industrial production, where there may be lower production in months in which Easter falls due to fewer days being worked. This effect can be removed by the X12ARIMA program.

If no Easter correction is made by X12ARIMA, the additional sales will be initially put in the SI ratios. The monthly SIs are then smoothed, leaving the Easter effect in the irregular component. The final seasonal adjusted series will thus show peaks and troughs due to the effects of Easters.

In this respect, Easter effects can be considered similar to trading day effects, and removing the Easter effect from the seasonally adjusted series improves the quality of the seasonal adjustment. As with trading day effects, Easter effects can be estimated in the X12ARIMA programme using the **regression** spec. There are two options within this spec that provide pre-defined regressors that will adjust for the effects of Easter. The **regression** spec also allows one further option, a user-defined regressor, which allows user to define their own Easter holiday regressor.

13.3 When to adjust for Easter effect

This section examines some methods that can be used to identify the presence of Easter effects. An advantage of X12ARIMA over X11ARIMA is the use of the AIC test in the **regression** spec that provides diagnostics for assessing whether an Easter regressor should be included or not. However, there are other indicators that can show the presence of an Easter effect, such as a graphical analysis of the non-seasonally adjusted and seasonally adjusted (without an Easter adjustment) series, the SI ratios plot, and the E5 and E6 tables given in the output. Each of these approaches is discussed in turn. The AIC test for Easter regressors is discussed in more detail in section 5, where a brief introduction to the test is provided.

The X12ARIMA programme provides a number of diagnostics to test a series for the presence of Easter effects. Testing the statistical significance of Easter regressors is discussed in greater detail in section 5, which describes the recommended procedure for testing and adjusting for Easter effects. In general,

where Easter effects are found to be statistically significant, the series should be adjusted to remove these effects from the final seasonally adjusted series. However, it is always important to look at the results of the build up period (w) and try to relate it to the time series itself. If the results are counterintuitive (e.g. they suggest most passengers travel the day before Easter, i.e. $w=1$) then it is worth investigating whether there is anything in the recording of the data which is causing this result. It is better not to implement the Easter results if they are very counterintuitive.

Easter effects should not be estimated with the default X12ARIMA regressors for the following types of data:

- Data that are collected on a 4,4,5 week pattern. That is to say, the recording periods in a year consists of a four times repeated pattern of a four week recording period, followed by another four week recording period, followed by a five week recording period. This means that the year is divided differently to the calendar periods described by months. There may, in this case, be some sort of Easter effect, however, this effect should not be estimated for using the regressors provided by X12ARIMA.
- Data that is collected at a point in time, for example a particular day in the month. There may, in this case, be some sort of Easter effect, however, this effect should not be estimated for using the regressors provided by X12ARIMA.
- Data that are not collected in strict calendar months. Also in this case there may be some sort of Easter effect, however, this effect should not be estimated for using the regressors provided by X12ARIMA. For further information about adjusting such data contact the Time Series Branch.

X12ARIMA provides different diagnostics to test for the statistical significance of Easter regressors.

13.4 Options available to adjust for Easter effects

The X12ARIMA programme runs a regression to model certain effects that result from the arrangement of the calendar, such as Easter effects and trading day effects (see Chapter 11). It is possible for the user to define regressors in the **regression** spec, using the **user** and **usertype** arguments. However, the X12ARIMA programme also contains pre-defined variables in the **variables** argument, to adjust for a variety of calendar effects, including forty nine different regressors that specifically adjust for the effects of Easter. The forty nine regressors can be described in two groups; the **easter[w]** option or **sceaster[w]** option in the **variables** argument. The following sub-sections describe the **easter[w]** option, **sceaster[w]** option and **user** option methods for adjusting for the effects of Easter in monthly or quarterly flow data only.

Variable name	X12ARIMA command	Comments
'easter[w]'	<pre> regression{ variables=easter[w] } </pre>	<p>The value of w must be supplied and states the number of days before Easter for which the level of daily activity changes, due to the effects of Easter. The new level of activity remains the same from the w-th day to the day before Easter. w can take any value from 1 to 25. Hence, easter[1] would mean that the holiday effect is estimating the change in the level of daily activity for the Saturday before Easter, whereas easter[8] assumes the change in the level of activity occurs for the whole week prior to Easter Sunday i.e. all 8 days before.</p>

Variable name	X12ARIMA command	Comments
		It is possible to specify more than one of these variables, with different values for w in order to estimate complex effects.
'sceaster[w]	<pre>regression{ variables=sceaster[w] }</pre>	<p>The regression variable in this case is a Statistics Canada holiday regression variable. This regression assumes that the level of daily activity changes on the (w-1)-th day and remains at the new level through Easter day. w must be supplied and can take any value from 1 to 24.</p> <p>It is possible to specify more than one of these variables, with different values for w in order to estimate complex effects.</p>
'user'	<pre>regression{ user=easter file="easter.txt" usertype=holiday }</pre>	<p>The user argument must be used in conjunction with the usertype and data or file arguments in the regression spec. This allows a user-defined regression variable to be used, i.e. the UK equivalents to the pre-defined regressors in X12ARIMA. If the file is held in a different directory then the path should also be specified, The values in this file should cover the time frame of data including forecasts and backcasts specified in the forecast spec.</p> <p>The file argument cannot be used if the data argument is specified and vice versa.</p>

For further information and description of handling Easter adjustment with regression model used in X-12-ARIMA see Findley, D.F. et al (1998) "New Capabilities and Methods of the X-12-ARIMA Seasonal Adjustment Program" *Journal of Business and Economic Statistics*, Vol.16 No.2.

13.5 How to adjust for Easter effect

Section 2 and 3 described when to adjust for Easter effects, whilst Section 4 introduced the different options available in X12ARIMA to adjust for Easter effects. This section describes in more detail the process of using the **aictest** argument in the **regression** spec to test the **easter[w]** and **user** options and the generally recommended procedure to adjust for Easter effects.

The general order of testing the significance of regressors is described in chapter 11, which discusses the regARIMA model. Chapter 11 explained how the *t*-value should be used to test for the significance of the chosen Easter regressor. However, other methods, such as the AIC test, also exist for detecting the presence of Easter effects and determining the build up period (i.e. the value of **w**).

This section will describe three ways of identifying whether or not Easter effects are present in a series, and then the procedure for adjusting for Easter effects in a production run. Two different scenarios will be outlined for setting up the seasonal adjustment for the production run:

- Firstly the recommended procedure of setting up the spec file for a production run using a regression variable which is an option that can only be used in X12ARIMA based programmes;

- Secondly producing prior adjustments that can be fixed for a year's production run in both X12ARIMA and X11ARIMA based programmes.

13.5.1 Testing for Easter effects with X12ARIMA

13.5.1.1 The AIC test

The AIC test can be activated in the **regression** spec to evaluate whether or not a particular regressor is preferred, compared to not having that regressor in the model. For example the following may be specified,

```
regression{aictest=(easter)}
```

This will test four models, respectively, no Easter regressor, **easter[1]**, **easter[8]**, and **easter[15]**. The model with the lowest AICC value will be the one that is preferred.

This generates the following tables in the output

```
MODEL ESTIMATION/EVALUATION
Exact ARMA likelihood estimation
Max total ARMA iterations          200
Convergence tolerance              1.00E-05
```

Likelihood statistics for model without Easter

Likelihood Statistics

```
-----
Effective number of observations (nefobs)          119
Number of parameters estimated (np)                3
Log likelihood                                    127.9166
Transformation Adjustment                         -391.0636
Adjusted Log likelihood (L)                       -263.1470
AIC                                                532.2940
AICC (F-corrected-AIC)                          532.5027
Hannan Quinn                                     535.6796
BIC                                               540.6314
-----
```

Likelihood statistics for model with Easter[1]

Likelihood Statistics

```
-----
Effective number of observations (nefobs)          119
Number of parameters estimated (np)                4
Log likelihood                                    173.5782
Transformation Adjustment                         -391.0636
Adjusted Log likelihood (L)                       -217.4854
AIC                                                442.9707
AICC (F-corrected-AIC)                          443.3216
Hannan Quinn                                     447.4848
BIC                                               454.0872
-----
```

Likelihood statistics for model with Easter[8]

Likelihood Statistics

```
-----
Effective number of observations (nefobs)          119
Number of parameters estimated (np)                4
```

```

Log likelihood                147.1087
Transformation Adjustment     -391.0636
Adjusted Log likelihood (L)   -243.9549
AIC                           495.9098
AICC (F-corrected-AIC)      496.2607
Hannan Quinn                  500.4238
BIC                           507.0263
-----

```

Likelihood statistics for model with Easter[15]

Likelihood Statistics

```

-----
Effective number of observations (nefobs)      119
Number of parameters estimated (np)           4
Log likelihood                                140.6211
Transformation Adjustment                     -391.0636
Adjusted Log likelihood (L)                   -250.4425
AIC                                             508.8851
AICC (F-corrected-AIC)                    509.2359
Hannan Quinn                                 513.3991
BIC                                             520.0015
-----

```

***** AICC (with aicdiff= 0.00) prefers model with Easter[1] *****

The models are compared on the basis of the AICC values, and the one with the lowest chosen - in the above example **easter[1]** returns the lowest AICC value, at 443. In the above example "aicdiff= 0.00" which means that the Easter regressor used has the lowest AICC value. By specifying **aicdiff = n** where **n** can take any value it is possible, for example if **n = 100**, to make the programme accept using the regressor only if the AICC value associated with the specified regressor is less than the AICC value for the model without that regressor by at least the value **n**. Therefore, in the above example if **aicdiff = 100** then none of these Easter regressors would have been chosen.

This test only checks the three Easter regressors described above. If other regressors are also specified in the **aicctest** argument, then the tests are performed sequentially; trading day regressors, then easter regressors and then user-defined regressors. In order to test a specific **easter[w]** or **sceaster[w]** regressor, X12ARIMA should be run once with and once without the regressor, each time saving the AICC values using the **savelog** argument in the **regression** spec as shown below

```

regression{ variable=easter[w]
             savelog=aicctest}

```

The AICC values can then be compared across different regressors or no regressor. This option is significantly more time consuming and would only be necessary if further analysis was deemed appropriate, if for example, the regressor automatically chosen appeared to perform poorly.

The recommended procedure is to use the **aicctest=easter** argument and use the chosen regressor either to estimate permanent priors or to set regression variables for a production run (see below).

If the preferred model is any of **easter(1)**, **easter(8)**, or **easter(15)** then, in general the chosen regressor should be used to adjust the series for Easter effects. A final check should be done to verify if the t-value of the selected regressor is statistically significant.

The AIC test for a user-defined regressor variable is similar to that described above, but the AICC values compared are for a model with and for a model without the specified regressor. The AIC test is activated as shown below:


```
regression{aictest=(user)}
```

This should be used if a user-defined regressor has been constructed to adjust for Easter effects. Again the model with the lowest AICC value is chosen, and again it is possible to use the **aicdiff** argument. If the user-defined Easter regressor is preferred, then it should be used to adjust for the effects of Easter.

13.5.1.2 Graphical analysis of the non-seasonally and seasonally adjusted data


If a series has been seasonally adjusted without any adjustment for the effects of Easter a plot of the original, non-seasonally adjusted (NSA) and seasonally adjusted (SA) series may reveal the presence of an Easter effect. The graph below shows residual seasonality, as this series has an Easter effect that has not been adjusted for. The effect is particularly noticeable in 1997, where there is an early Easter, as Easter Sunday was on the 30th of March. But as the graph shows this effect is compensated for in the March of other years, which results in residual seasonality in the adjusted (SA) series. If an adjustment were made to account for the Easter effects this would significantly improve the seasonal adjustment.

Figure 13.1 – Easter effects



13.5.1.3 Graphical analysis of the unmodified SI ratios

If a series has been seasonally adjusted without any adjustment for the effects of Easter a plot of the unmodified SI ratios for March and April may reveal the presence of an Easter effect. The two graphs below show that the SI ratios for March and April in both 1991 and 1997 are significantly lower than those for other years. Therefore, when the seasonal factors are estimated with a moving average, this will distort them, as a prior adjustment for the effect of Easter would mean the SI ratios for March and April 1991 and 1997 would be similar to those in the other years.



13.5.1.4 E5 and E6 Output Tables

The E5 table produced in the output gives the month-to-month or quarter-to-quarter percentage change in the original series. Therefore, if an Easter effect is present, it may be possible to see a different changes in March and April (or quarter 1 and quarter 2) of those years where there is an early Easter. For example, the table below shows the figures from an E5 table. There is a noticeable change in the month-to-month changes for March in 1991 and 1997 (early Easters), i.e. over 400 whereas other years the March value tends to be just over 200.

E5 Month-to-month percent change in the original series
From 1991.Feb to 2001.Dec
Observations 131

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	AVGE
1991		-37.09	-64.81	447.67	15.58	1.06	-27.46	-10.31	-19.22	-11.14	51.73	35.92	34.72
1992	23.64	-45.27	-42.53	231.79	17.24	5.12	-30.53	-10.70	-14.30	-25.99	71.59	35.19	17.94
1993	16.33	-45.39	-41.69	230.42	12.55	3.79	-27.78	-16.67	-16.79	-19.63	62.31	38.61	16.34
1994	11.00	-44.31	-36.94	208.54	12.73	-0.04	-28.12	-14.66	-13.69	-20.75	54.83	34.55	13.59
1995	16.19	-45.45	-41.66	227.22	12.25	1.33	-29.19	-14.80	-12.55	-22.50	54.29	35.31	15.04
1996	16.39	-45.98	-40.50	220.36	12.14	1.77	-29.70	-12.88	-17.52	-18.28	50.50	35.62	14.33
1997	16.33	-43.70	-65.14	427.29	12.31	-0.31	-28.70	-15.39	-14.81	-20.64	53.54	34.36	29.59
1998	13.95	-44.52	-40.43	211.56	12.70	1.57	-29.13	-15.59	-16.50	-16.83	48.82	35.76	13.45
1999	14.16	-44.15	-41.11	208.59	13.20	0.03	-28.69	-15.48	-14.72	-20.77	54.37	34.89	13.36
2000	13.89	-45.28	-39.66	205.33	12.69	0.64	-29.89	-14.10	-15.72	-19.20	52.52	32.60	12.82
2001	14.03	-44.20	-38.97	200.60	11.44	0.65	-29.78	-13.21	-15.83	-20.25	52.00	34.00	12.54
AVGE	15.59	-44.12	-44.86	256.31	13.17	1.42	-29.00	-13.98	-15.60	-19.63	55.14	35.17	
Table Total-	2289.87		Mean-	17.48		Std. Deviation-	83.77	Min -	-65.14		Max -	447.67	

Table E6 shows the month-to-month or quarter-to-quarter percentage change in the seasonally adjusted series. In the same way as above for the E5 table, it may be possible to see a change in the values for March or April in years where there is an early Easter relative to years without an early Easter.

13.5.2 Adjust for Easter effects

Use one of the following three methods to adjust for Easter effects.

13.5.2.1 Regression variable with test of significance

Once it has been decided which Easter regressor should be used to adjust for the Easter effect, the **aictest** described in section 13.5.1.1 can be used to adjust for the effect. This means that every time X12ARIMA is run to estimate the seasonally adjusted series, AIC values are derived for models with and without the specified Easter variable and the optimal model will be used for forecasting. This option will provide users with the optimal seasonally adjusted series, but will generate higher revisions than the other two methods described in the following sections.

Problems:

- This option should not be used for production run because the Easter regressor variable is not fixed and frequent switching between models due to updating can introduce undesirable instability in the seasonally adjusted output. If there is some special reason why this option must be used, a non-zero value of the **aicdiff** argument should be used to reduce instability.

13.5.2.2 Regression variables fixed in the model

If a particular regression variable has been identified as significant and the best one to use for the seasonal adjustment, then the **regression** spec should contain in the **variable** argument all the

regressors that should be included. For example, if, as in the previous example, **easter[1]** and **td** are found to be significant, the spec file, along with the other parameters that have been fixed should include the following arguments in the **regression** spec.

```
regression {variables=(easter[1] td)}
```

As with the trading day regressors, this is the recommended option for the production run because the form of the chosen regressor is fixed though the parameter values, the Easter factors, are re-estimated each period. When a particular Easter regressor is fixed in the model, this should be included even if it appears to no longer be significant. The decision of whether to remove the variable or not should be taken at the point of a seasonal adjustment review, so as to reduce the number of revisions.

Problems:

- If a strict revision policy is in place it is suggested not to use this option since the estimated Easter factors will change when new data becomes available.

13.5.2.3 *Permanent prior adjustments*

Once a particular regressor has been chosen (see above and Chapter 11, on regARIMA and Chapter 5 on Procedures for analysing series with X12ARIMA) then the Easter factors can be saved in order to obtain permanent prior adjustments. The following steps should be taken to obtain permanent prior adjustments.

1. Run the spec file with your chosen settings, e.g. de-composition, ARIMA model and so forth. Include in the **regression** spec with the **variable** and **save** arguments activated (NB, in order to save the Easter factors obtained from any Easter variable the appropriate argument is **save = hol** for all variables). For example, if the **Easter[1]** regressor was found to be the most appropriate the **regression** spec would look like this,

```
regression{variables=(Easter[1])
           save=(hol)}
```

2. The save argument will save a text file with the Easter factors in the same directory that the log and output files are saved. If the name of the spec file was "*filename.spc*" this would mean that the Easter factors would be saved in a file called "*filename.hol*".
3. If permanent priors are required for more than one year, then the **forecast** spec should be used to set the number of forecast periods and **appendfcst=yes** option should be included in the "x11" spec. For example, the following, assuming monthly data, would provide Easter factors for two years into the future,

```
regression{variables=(Easter[1])
           save=(hol)}
forecast{maxlead=24}
x11{appendfcst=yes}
```

4. It is essential to save the hol file with a different name (e.g. "*filenamepp.hol*") otherwise it will be overwritten the next time the spec file would be run with the **variables=(Easter[1])** option.

5. In order to set up the spec file for a production run, remove the **regression** spec and use the prior adjustments that have been saved in step 4 by using the **transform** spec. Therefore, in the final spec file, used for production runs, the **regression** spec is not used and the **transform** spec is activated, as in the following,

```
transform {file=("filenamepp.hol")
          format="x12save"}

          type=(permanent)}
```

It should be noted that the above example has given a certain method of saving and using the permanent priors in a particular format and so on. There are various ways in which permanent priors can be saved and used to transform the original series. For further information on the different options available see X-12-ARIMA Reference Manual.

If trading day (e.g. td) or user defined regressor variables are used to estimate other effects that are required to be used as permanent priors, then the factors must be multiplied together so that they are all in one file to be used in the **transform** spec. For example, if **easter[1]** and **td** variables have been estimated and the holiday and trading day factors have been saved (**save=(hol td)** is the appropriate argument in the **regression** spec, which saves the respective factors in separate files with the name of the spec file and the extension .hol and .td respectively) then these factors should be multiplied by one another and following steps 4 and 5 will mean that the original series is adjusted by the prior adjustment file so that the seasonal adjustment is then performed on a series with these calendar effects removed, which should improve the quality of the seasonal adjustment.

Problems:

- With the permanent priors option, the adjustment for Easter factors is fixed as they are only based on the data that is available at the point in time when the Easter effects are being estimated (during the annual review). If the Easter pattern changes between annual reviews the permanent priors will not capture this modification therefore the seasonal adjusted series will not be optimal.
- It is very complicated to set up and keep up-to-date.

Criteria for deciding which of the three methods should be used are as follows:

1. **Revision policy:** How strict the revision policy is determines which approach to use in the production runs. The use of permanent priors is the method that gives the minimum revisions, but also gives the less optimal seasonally adjusted series and it is very complicated to set up and keep up-to-date. The second option in terms of revision size is fixing the regression variable in the model. This is the recommended method since it balances a fairly good quality of seasonal adjustment with revision and practicality. Contrary, the use of a regression variable with a test for significance will provide with the optimal seasonal adjustment but with the biggest revisions of the three methods.
2. **Seasonal Adjustment Review:** for reviewing the parameters the regression variable with a test for significance should be used. It provides users with the optimal model to run the seasonal adjustment.

13.5.3 The US/UK problem

As the US Census Bureau developed the X12ARIMA programme, the regressors that are defined in the programme have been constructed to account for the holiday as it occurs in the United States. When these regressors are used on UK data, they are theoretically wrong. For example, the **Easter[1]** regressor

in X12ARIMA accounts only for Easter Saturday, while the UK one should account for three extra public holidays, Easter Friday, Easter and Easter Monday.

It is possible to construct a regressor that follows the public holidays associated with Easter in the UK. The UK regressor can be computed so as to be a UK equivalent to the **easter[w]** regressors discussed above. If a UK regressor is deemed necessary, then it must be specified using the user-defined regressor, as shown above, and can be compared using the **aictest=user** argument in the regression spec.

As empirical testing suggests that there is little difference between UK and US regressors, it is sufficient to use the pre-defined regressors. However, if a detailed analysis is required, or there is some question over the performance of the pre-defined regressors, then using a UK defined Easter regressor may be appropriate. For further information contact Time Series Analysis Branch.

13.6 Related topics

Other topics related to Easter include,

- Trading days
- RegARIMA
- Prior adjustments

14 LEVEL SHIFT AND ADDITIVE OUTLIERS

14.1 Introduction

When analysing time series, analysts may need to identify possible inconsistencies or effects in the data. In some cases, these problems will have an identifiable and real world cause. For example, promotion activities to achieve a sales goal; change in business rules, regulations and policies; weather changes; natural disasters; international conflicts and wars; financial market crashes. The presence of such effects can affect the qualities of the seasonal adjustment. In these situations, the analyst will normally try to remove the inconsistency before the seasonal adjustment of the series.

This chapter will deal with two types of effects: level shifts and additive outliers.


14.2 Level Shifts

14.2.1 What is a level shift

A level shift (or trend break) is defined as an abrupt but sustained change in the underlying level of the time series. The annual seasonal pattern is not changed. There are many potential causes of level shifts in time series, including change in concepts and definitions of the survey population, collection method change, change economic behaviour, in legislation or in social traditions.

The following graph provides an example of level shift in the Export of Meat.

Figure 14-1 Example of level shift: Export of Meat



In this case the level shift is obvious at April 1996 and was due to the ban in export of beef.

14.2.2 Why adjust for a level shift

Level shifts are a problem for seasonal adjustment because if these effects are not appropriately corrected for they will distort the estimation of the seasonal factors.


Within X-12-ARIMA, the initial trend-cycle is calculated by applying a centred 12-term moving average to the series after trading day effects, Easter, and other prior adjustments have been applied. The trend-cycle is removed from the original estimates to give an estimate of the seasonal times irregular component (SI Ratios) and then the seasonal factors are estimated after the replacement of extreme values. The Henderson filters are then applied to the seasonally adjusted estimates to produce the final trend-cycle estimate.

If there is an abrupt change in the level of the series, when the moving averages are applied to the series to calculate the preliminary trend-cycle, the estimates will be distorted. The estimates of the trend-cycle prior to the abrupt change will be underestimated, and those after the abrupt change will be overestimated. As the calculation of the irregular and seasonal components follow on from this initial trend-cycle estimation, they may be distorted. The resulting seasonal adjusted estimates will be more volatile.

Figure 14-2 shows an example where a level shift has not been corrected.

Figure 14-2 Example of level shift

**Example of Level Shift - Export of Meat Series
No prior adjustment for the Level Shift**



The trend-cycle (dashed line) shows a rapid decrease from mid 1995 to mid 1996. The trend estimates prior to the level shift are lower than expected, and those following the shift are higher than expected. The result for the seasonally adjusted series is an increased level of irregularity around the level shift, increasing the volatility of the seasonally adjusted series. The final seasonally adjusted and trend estimates will be misleading.


14.3 Additive Outliers

14.3.1 What is an additive outlier

An extreme value or additive outlier is a data point which falls out of the general pattern of the trend and seasonal component. Outliers may either be caused by random effects, ie an extreme irregular point, or due to an identifiable cause such as a strike or bad weather. This chapter will deal only with the second group of outlier, where there is an underlying economic reason that explains the unusual behaviour of the data point. For more information about how X-12-ARIMA deals with extreme values see Chapter 8.

Figure 14-3 provides an example of an additive outlier.

Figure 14-3 Example of additive outlier



In this case the additive outlier is in Q2 2002 when expenditure was very high because the World Cup was in Japan.

14.3.2 Why adjust for an additive outlier

Additive outliers are a problem for seasonal adjustment, because the method of seasonal adjustment is based on moving averages. The problem with averages is that they are affected by the presence of extreme values or outliers, for example, the presence of unusual data points can make the average unrepresentative of the pattern of the series. If some adjustment or allowance is not made for outliers then these will cause distortion in the estimates of all the components in a time series.

Seasonally adjusting the data above, without allowing for the outliers, gives the following results:

Figure 14-4 example of outlier




Figure 14-4 shows that the time series is very erratic by nature, but the presence of the outlier makes the seasonally adjusted series more irregular, especially in the span of data around the outlier. In addition, the fact that additive outlier is at the end of the time series means that the time series will be prone to large revisions when new data points become available.

14.4 How to identify and adjust for level shifts and additive outliers

The following paragraph describes how to identify and adjust for both level shifts and additive outliers, since the procedure to adjust for these two discontinuities is the same. In practice, graphing the series, and the results of the regression-ARIMA test for outliers and level shifts will provide the best ways of identifying these features.

14.4.1 Run the series in X-12-ARIMA and look at the output diagnostics

- Run the series using the standard spec file in X-12-ARIMA without checking for level shifts and additive outliers.
- Look at the graph of the NSA and SA series and the X-12-ARIMA output.

The level shifts and additive outlier identification is not always as obvious as the examples shown in this chapter. In cases where the level shifts and additive outlier can not be spotted from the graphical representation of the series, an analysis of the X-12-ARIMA output can help in the detection of level shifts and additive outliers.

Level shifts can be detected by analysing Table E5 and E6 in the output. These tables show the month-on-month (or quarter-on-quarter) changes in the original and seasonally adjusted estimates. In fact, a level shift will appear as a sudden large increase which is not followed by a corresponding decrease (or

vice versa) in Table E5 and E6. The adjustment for the level shift effectively attempts to remove this sudden change in the level.

Additive outliers can be detected by analysing Table C17 and D13 in addition to Table E5 and E6 in the output. In fact, an additive outlier will appear as one or more zeros in the period of the outlier in Table C17, which gives the final weights for the irregular component, and as a residual pattern of the irregular component presented in Table D13. In addition to inspecting these four output tables, M1, M2 and M3 statistics should be checked. M1, M2 and M3 each measure the level of the irregular component in the series compared to the trend and the seasonal components. If these fail it may indicate that outliers need to be replaced as temporary prior adjustments.

14.4.2 Length of the series before and after the level shift or outlier

1. The entire length of the series in total needs to be at least 5 years to use regression-ARIMA.
2. The outlier correction analysis does not have restrictions in terms of data available before and after the outlier.
3. The level shift cannot occur on the start date of the series since the level of the series prior to the given data is unknown.
4. A level shift at the second data point cannot be distinguished from an outlier at the start date of the series, whilst a level shift at the last date of the series cannot be distinguished from an additive outlier since the level of the series after the discontinuity is unknown. In these situations the knowledge of the series or the use of external information are of primary importance to prior adjust the original series.

14.4.3 Testing for level shifts and additive outliers in X-12-ARIMA

If the user wants to search simultaneously for level shifts and additive outliers within a span of data, the following spec should be used.

```
series{title="Example of level shift and outlier search"  
  start=1994.1  
  period=4  
  file="mydata.txt"}  
  
arima{model=(0,1,1)(0,1,1)}  
  
outlier{types=(ao ls)  
  span=(2002.2, )  
  
x11{mode=mult}}
```

The **outlier** specification performs automatic detection of additive outliers and level shifts or any combination of the two using the model specified in the arima specification. After outliers (referring to any of the outlier types mentioned above) have been identified, the appropriate regression variables are incorporated into the model as "Automatically Identified Outliers".

It is important to notice that this specification will detect only additive outliers and level shifts with a particularly high critical value. In fact, the value to which the absolute values of the outlier t-statistics are compared depends on the number of observations included in the interval searched for outliers. For example, if the outlier search is run for the last year of data, the critical value will be 3.16 (compared with

the standard critical value of 1.96). This means that all the outliers with t-statistic values between 1.96 and 3.16 will not be detected using this option and that the visual check of the NSA and SA series remain an important diagnostic in the identification of outliers. For more information on the default critical values for outlier identification, see the X-12-ARIMA Reference Manual on page 107/108.

14.4.4 Confirming the reason for the level shifts or the additive outliers

If the X-12-ARIMA test and the graphs of the NSA and SA series lead you to suspect a level shift or an outlier, then check which months/quarters it appears in. Question if there is any evidence for suspecting a level shift or outlier at this time point.

If there is not a methodological, social or economic reason to explain the presence of the suspected level shift or outlier then **do not adjust** for it.

14.4.5 Adjust for the level shifts or the additive outliers

14.4.5.1 Use the regression-ARIMA

If you suspect that a level shift or an outlier is present at a particular time point (either before or after the graphical analysis) then it should be specifically tested for using X-12-ARIMA by including the following command line in the X-12-ARIMA regression specification.

```
regression{  
    variables=(ao2002.2 ls1996.4)  
}
```

The **variables=(ao2002.2 ls1996.4)** command instructs X-12-ARIMA to analyse the level shift or outlier and to create temporary priors to adjust for the discontinuity.

This command allows the user to test and adjust for a level shift of an additive outlier. More than one **ao** and/or **ls** may be specified in the model. When the command is included in the regression spec, the output includes t-tests for each regressor (**aos** and **lss**) included in the model. The definition of this test is in the regression-ARIMA chapter.

The use of regression-ARIMA **ao** and **ls** regressors is only recommended if the presence of an outlier or level shift has been previously confirmed and the t-test absolute value is greater than 2.

The temporary prior adjustments derived by X-12-ARIMA are shown in Table A8. These prior adjustments are the result of the multiplication (or addition) of all the temporary priors derived by the **aos** and **lss** included in the model. The temporary priors adjust the level of the series before the point at which the level shift occurs without altering the annual seasonal pattern.

The use of temporary priors derived by the regression-ARIMA approach avoid any distortion in the estimation of the components of the series, and lead to a seasonally adjusted series that shows the discontinuity.

Figure 14-5 Example of level shift

**Example of Level Shift - Export of Meat Series
SA with and without adjustment for the Level Shift**




Figure 14-6 Example of additive outlier

**Example of Additive Outlier - UK expenditure in Japan
SA with and without adjustment for the outlier**



The two graphs above show that the use of regressors to correct the discrepancies in the series improve the quality of the seasonal adjustment.

Problems:

- High degree of irregularity in the series may hide the presence of level shift.
- High degree of irregularity can make the estimation of the magnitude and timing of the level shift less accurate.
- Outliers regressors correct for discontinuities even if they are not significant. The t-value of each regressor (or the Chi-squared of a group of regressors) should be checked to be sure of its (their) significance.

14.4.5.2 Use of automatic outlier detection

If the user wants to search simultaneously for level shifts and additive outliers within a span of data, the following spec should be used.

```
series{title="Example of level shift and outlier search"  
  start=1994.1  
  period=4  
  file="mydata.txt"}
```

```
arma{model=(0,1,1)(0,1,1)}
```

```
outlier{types=(ao ls)  
  span=(2002.2, )  
  lsrn=n}
```

```
x11{mode=mult}
```

The **outlier** specification performs automatic detection of additive outliers and level shifts or any combination of the two using the model specified in the arima spec. After outliers (referring to either of the outlier types mentioned above) have been identified, the appropriate regression variables are incorporated into the model as "Automatically Identified Outliers" and the model is re-estimated.

The **lsrun** command computes t-statistics to test null hypotheses that each run of 2, ..., n, ..., 5 successive level shifts cancels to form a temporary level shift. If one or more level shift t-tests indicate that a run of 2 or more successive level shifts cancels, a user-defined regressor can be used to capture the temporary level shift effect. In this way two or more level shifts can be replaced by one user-defined regressors. The "usertype" argument should be set to ls for this regressor, so the user defined regressor is treated as level shift. An example of this is given below:

```
regression{user=(myls)  
  usertype=ls  
  file=myls.rmx  
  format="x12save"}
```

where **myls** is the user-defined regressor created to take into account of the temporary level shift. For example if there are two level shifts that cancel out, one in January 2002 and the other in March 2003, the regressor has the following form:

$$\begin{cases} 0 & t < \text{January 2002} \\ -1 & \text{January 2002} \leq t \leq \text{March 2003} \\ 0 & t > \text{March 2003} \end{cases}$$

where t define the date of the observation.

The **span** command specifies starts and end dates of a span of the time series to be checked for outliers. The two dates must both lie within the series and within the "modelspace" if one is specified (see chapter 10 for more information about "modelspace"). If one of the two dates is missing, e.g. **span=(2002.2,)**, X-12-ARIMA sets the missing date on the start date or end date of the series. For example, using the outlier specification defined above, X-12-ARIMA searches for outliers between February 2002 and the last data point of the series.

Problems:

- The automatic outlier detection is not always stable. Sometimes outliers switch from being significant one month to being not significant the following month and going back to be significant the next month. This instability has an effect on the revision pattern. In fact, the revision history of series with automatic outlier detection in place is more erratic than the revision history without automatic outlier detection or with outlier regression variables in the reg-ARIMA model. This problem do not apply to series with stable seasonality. This stability problem can be reduced if a shorter span is considered in the search of outliers, e.g. if the outlier specification is used in the last year of data.
- It only detects outliers with an high critical value. The automatic outlier detection do not consider outliers with t -value between 1.96 and the critical value set by X-12-ARIMA.
- The test for automatic outlier detection may have less power in the last data point. A level shift at the last date of the series cannot be distinguished from an additive outlier since the level of the series after the discontinuity is unknown. In these situations the knowledge of the series or the use of external information are of primary importance to prior adjust the original series.

Criteria for deciding which of the two methods should be used are as follows:

- **Size of the dataset.** The size of the dataset being analysed influences the decision about which of the two methods to use. If a large dataset is analysed, method 4.5b is preferable as it is less time consuming than method 4.5a, easier to update, but still accurate. The use of regressors in the model (method 4.5a) is preferable for small datasets.
- **Importance of the series.** For important series, method 4.5a is more accurate and produces more stable results. In this situation, though, method 4.5b can be used to monitor the series between seasonal adjustment reviews. This will keep users informed of possible problems that have an effect on the seasonal adjustment of the latest data points.
- **Size of revisions.** If the size of revisions is important, then method 4.5a should be used. Although again in this situation, method 4.5b could be used to monitor the series between seasonal adjustment reviews.

15 WHICH SEASONAL DECOMPOSITION MODEL TO USE

15.1 Definition

X-12-ARIMA seasonally adjusts a time series by modelling it as at least three unobserved components. The process of breaking down a series into these components is known as decomposition.

There are two basic ways in which X-12-ARIMA can model seasonality in order to identify and remove it. The first, and most common, is the **multiplicative** model, which is of the form:

$$Y = C \times S \times I$$

where Y is the original series, C is the trend-cycle, which includes the medium and long term movements in the series, S is the seasonal component and I is the irregular component .

This decomposition model is used for most economic time series. The decomposition model cannot be used when there are negative numbers or zeros in the series.

For other series where the seasonal changes are independent of the level of the series, the **additive** model will be more appropriate. It decomposes the series as follows:

$$Y = C + S + I$$

Another decomposition model is the pseudo-additive model. It is available within X-12-ARIMA although in practice it is only used for time series where there are non-negative values with regular zeros. For example, this could be applied to agricultural time series. ..

All prior adjustments will be of the same type as the decomposition model, i.e. for an additive model all prior adjustments are additive and for a multiplicative model all priors are multiplicative. The above is a slight simplification: trading day variation and Easter effects can both be seen as additive components entering into the decomposition model.

15.2 The options

Two specs are involved in the selection of the model to be used in the seasonal adjustment decomposition, "transform" and "x11". Those two specs **may** be written as:

```
transform{function=log}  
x11{mode=mult}
```

for a multiplicative decomposition or

```
transform{function=none}  
x11{mode=add}
```

for an additive decomposition. Note that in the first case **mode=mult** may be omitted, because it is the default, and in the second the **transform{function=none}** may be omitted, again because it is the default. However, the explicit statement of a default value does no harm and may make the process clearer.

15.3 How to decide which seasonal decomposition to use

There are two ways of identifying which model is the most appropriate: by inspecting the graph or by analytical means.

15.3.1 Graphical inspection

In many cases, inspection of the graph of the time series and knowledge of the data will make it clear which decomposition model to choose. Under the multiplicative model, the seasonality of the series is affected by the level of the series. So when the graph of the series shows that the size of the seasonal peaks and troughs increase (or decrease) as the trend rises (or falls), a multiplicative decomposition model is appropriate. Alternatively, if the size of the seasonal peaks and troughs are independent of the level of the trend, then an additive decomposition model is more appropriate. For example:

Table 15-1 Example of multiplicative decomposition

Example of Multiplicative Model - Airline Passengers




Table 15-2 Example of additive decomposition


It may not be clear from a plot of the series which of the two models is the most appropriate; for example, when the trend is fairly flat, both models may produce satisfactory results. A way of identifying the most appropriate model is to ask X-12-ARIMA to run an AIC test on both models and automatically select the one with the lowest statistic.

If a series or any prior adjustments contain negative numbers or zeros then the additive model must be used. In this case if a multiplicative model is specified the program will generate an error and will interrupt the seasonal adjustment run.

The user may have some knowledge about the series that may help to identify which model is most likely to be appropriate. In general all series in subsystem of accounts will normally be adjusted using the same decomposition model, unless a series clearly follows the other model.

The model used also affects how the seasonal and irregular components are presented. Under a multiplicative model both these components are given as ratios (expressed in percentage terms) which vary around 100. With an additive model the components are differences which vary around zero.

Example of additive model - Airport Output



15.3.2 Analytical approach

It was said earlier that initially a graph should be examined to identify which model should be used, but in the case that knowledge of the series and the graph do not help in the decision, an automatic procedure can be used. This is done by invoking the automatic transformation in the **transform** spec, i.e.

transform{function=auto}

Provided the series being processed has all positive values, X-12-ARIMA performs an AIC-based selection to decide between a log transformation and no transformation. In this case the **mode** in the **x11** spec is not necessary, since it is automatically selected by X-12-ARIMA in order to match the transformation selected for the series (**mult** for the log transformation and **add** for no transformation).

To decide which transformation to use, and therefore which decomposition model to use, X-12-ARIMA fits a regARIMA model to the untransformed and transformed series and chooses the log transformation except when

$$AICC_{no\log} - AICC_{log} < \Delta_{AICC}$$

where $AICC_{no\log}$ is the value of AICC from fitting the regARIMA model to the untransformed series, $AICC_{log}$ is the value of AICC from fitting the regARIMA model to the transform series, and Δ_{AICC} is the value entered for the **aicdiff** argument, with a default of -2. For more information on **aicdiff**, see the AIC, AICC and BIC for Modelling Comparison chapter. Negative values of Δ_{AICC} bias the selection in favour of

the log transformation. The default -2 is used not for statistical reasons but for convenience. Multiplicative adjustment is appropriate for the great majority of ONS series and the use of additive decomposition is suggested only when statistical support for additive adjustment is rather strong.

If a regARIMA model has been specified in the **regression** and/or **arima** specs, then the procedure uses this model to generate the AICC statistics needed for the test. If no model is specified, the program uses the so-called airline model ((011)(011)) to generate the AICC statistics.

15.4 Updating

The same decomposition model should be used throughout the year. It should not need to be changed at the annual update except in very exceptional circumstances. If the model does look as if it should be changed then TSAB should be consulted.

16 MOVING AVERAGES

16.1 Introduction

The X12ARIMA program uses moving averages throughout its iterations in order to decompose the original time series into trend-cycle, seasonal and irregular components, as described in chapter 8. The program uses moving averages for two different purposes: to estimate the trend component and to estimate the seasonal component. It uses two different types of moving averages to estimate the trend component – one before the seasonal component has been removed and one after. This chapter will describe what moving averages are, the different types of moving averages that are used, the options available in the program to select specific types of moving averages and why it is that these may be of use.

16.2 What are Moving Averages

A moving average is a weighted average of a moving span of fixed length of a time series. They can be used to modify a time series to produce a smoother version. In the simplest cases the objective is to remove or filter out as much as possible of the irregular component of the series, leaving the smooth component or trend. Besides removing or reducing irregularity, a moving average may also be chosen because of its ability to remove a fixed seasonal pattern. There are many forms of moving average, which differ in their ability to remove irregularity while preserving the trend as accurately as possible. X12ARIMA uses several of these forms, according to the needs of the different stages of the program.

A simplest moving average is a symmetric moving average (i.e. an equal number of data points either side of the target data point are included in the average) with equal weights applied to each data point.

Consider the following time series:

$$a_1, a_2, a_3, a_4, \dots, a_{t-1}, a_t, a_{t+1}, \dots$$

where a is the value of the series and the subscript denotes the time period.

If the span of data for the symmetric moving average is equal to 3 and equal weight is given to each point then the moving average at time period 2 is equal to:

$$u_2 = \frac{a_1 + a_2 + a_3}{3} = \left(\frac{1}{3} \times a_1\right) + \left(\frac{1}{3} \times a_2\right) + \left(\frac{1}{3} \times a_3\right)$$

Where u_2 is the value of the moving average at time period 2. The weights applied to each data point are equal to $1/3$. As this is a moving average the average is calculated for the next time period. Hence the moving average at time period 3 is:

$$u_3 = \frac{a_2 + a_3 + a_4}{3} = \left(\frac{1}{3} \times a_2\right) + \left(\frac{1}{3} \times a_3\right) + \left(\frac{1}{3} \times a_4\right)$$

More generally, the moving average at time period t is:

$$u_t = \frac{a_{t-1} + a_t + a_{t+1}}{3} = \left(\frac{1}{3} \times a_{t-1}\right) + \left(\frac{1}{3} \times a_t\right) + \left(\frac{1}{3} \times a_{t+1}\right)$$

This is called a 3x1 moving average. It cannot be calculated for the first and last points of the series, as we do not have equal number of data points either side of the target data points. The 3x1 moving average can be generalised to take any number of data points, each with an equal weight.

In a 5x1 moving average each data point will have a weight of 1/5 and in a 9x1 each will have a weight of 1/9. The longer the moving average the smoother the resultant series, as it takes information from a greater number of data points. However, with a longer moving average, the value can be estimated for fewer points. In the case of a 9x1 moving average, the average cannot be calculated for the first four and the last four terms of the series.

A symmetric moving average with equal weights where the span has an even number of points would estimate an average at the mid-point between two time periods. By not applying the same weight to each point, it is possible to produce a symmetric moving average with an even length of span. Using the notation introduced earlier, a 4x1 moving average first four points of a series

$$\frac{a_1 + a_2 + a_3 + a_4}{4} = \frac{1}{4}a_1 + \frac{1}{4}a_2 + \frac{1}{4}a_3 + \frac{1}{4}a_4$$

would estimate an average between the second and third time periods and a moving average of the next four points

$$\frac{a_2 + a_3 + a_4 + a_5}{4} = \frac{1}{4}a_2 + \frac{1}{4}a_3 + \frac{1}{4}a_4 + \frac{1}{4}a_5$$

would estimate an average between the third and fourth time periods.

An average of these two averages would then be centred at the third time period. This is called a 2x4 moving average and takes the general form

$$u_t = \frac{1}{2} \left(\left(\frac{a_{t-2} + a_{t-1} + a_t + a_{t+1}}{4} \right) + \left(\frac{a_{t-1} + a_t + a_{t+1} + a_{t+2}}{4} \right) \right)$$

$$= \frac{1}{8}a_{t-2} + \frac{2}{8}a_{t-1} + \frac{2}{8}a_t + \frac{2}{8}a_{t+1} + \frac{1}{8}a_{t+2}$$

This sort of average is used to remove the seasonally effect from a quarterly series. Although it does not give equal weights to all the data points, it is symmetric and it gives equal weights to each quarter. For example, the 2x4 moving average for quarter 2 will give a weight of 1/4 to quarters 1,2 and 3 of the current year, 1/8 to quarter 1 of the current year and the previous year. In the same way, a 2x12 moving average can be used to remove the seasonal effect and, hence, give an estimate of the trend for a monthly series.

The X12ARIMA program also uses combinations of simple (equally weighted) moving averages to estimate the seasonal component. An example of this is the 3x3 moving average, that an average of three consecutive 3x1 moving averages (i.e. the 3x1 moving averages for a_{t-1} , a_t and a_{t+1}). The structure of weights for this is:

$$u_t = \frac{1}{3} \left(\left(\frac{a_{t-2} + a_{t-1} + a_t}{3} \right) + \left(\frac{a_{t-1} + a_t + a_{t+1}}{3} \right) + \left(\frac{a_t + a_{t+1} + a_{t+2}}{3} \right) \right)$$

$$= \frac{1}{9}a_{t-2} + \frac{2}{9}a_{t-1} + \frac{3}{9}a_t + \frac{2}{9}a_{t+1} + \frac{1}{9}a_{t+2}$$

All these examples have used simple (equally weighted) moving averages or combinations of them. There are other moving averages, which use more complex patterns of unequal weights. All these forms have been designed with particular objectives in view, which can be described in terms of how they affect particular simple models of series. For example, if we suppose the series to be a linear trend with an irregular added, the simple moving averages will give a value, which is the linear trend, plus an irregular which is smaller because of the averaging. The length of average chosen will depend on how much smoothing we want to produce, which will in turn depend on how irregular the original series is.

There is a penalty in making the average longer, however, because the linear trend will seldom continue for any long time, although it may be an adequate approximation over a short period. If the trend has some curvature, using too long an average will distort it. Thus the choice of a length of average will be a matter of balancing the objectives of smoothing and following the trend.

Another objective may be exemplified by the 2x4 average. If a quarterly series has a stable seasonal component, it is obvious that averaging any four successive terms will produce a series in which the seasonal has been cancelled out. As shown above, the 2x4 average is equivalent to taking a 4-term average and then a 2-term average, the first step will remove the seasonal. Thus, this average applied to a seasonal quarterly series will reduce the irregular, remove the seasonal and still preserve the trend (provided it does not curve too much). Similarly, a 2x12 moving average will give an estimate of the trend of a seasonal monthly series.

In many cases it is found that the trend of a series has too much curvature to be adequately represented by these simple averages. To avoid this problem, X12ARIMA uses a family of moving averages called Henderson averages. These have the property that they will reproduce exactly a trend which can be represented as a cubic polynomial, while producing an output which has maximum smoothness for their given length. However, it should be noted that they cannot remove a seasonal effect. The formulae for the weights of the Henderson averages are rather complex, and so they are not reproduced here. For more information on their design and performance, please consult Time Series Analysis Branch.

All the examples quoted assume that the moving average is calculated at a point at which enough time series values are available before and after the point to apply the formula. Obviously, if we are calculating an average at or near the end of the series this will not be possible. In such cases the x11 part of the program provides asymmetrical approximations to the symmetric weights, which are generally constructed by assuming that the series may be forecast in some simple way, such as fitting a straight line by regression to the last few points and extrapolating it.

One of the improvements in X12ARIMA is to provide a better way of dealing with this so-called "end weight problem". It is worth noting that the ARIMA part of the program, which allows forecasting and backcasting of the series, in turn allows centred moving averages to be calculated for the first and final data points. The moving average for these points would otherwise use asymmetric weights, which gives a greater weight to the point itself and would therefore not generate such stable estimates of the separate components (trend-cycle, seasonal and irregular) at those points.

16.3 Trend Moving Averages

The trend moving averages are weighted arithmetic averages of data along consecutive points, as described in the previous section. There are two types of trend moving averages used by X12ARIMA:

- Centred simple moving averages (e.g. 2x12 or 2x4)

- Henderson moving averages

The simple moving average is applied at the first stage of the seasonal adjustment process, before the seasonal component has been removed, and will give a preliminary estimate of the trend. Henderson moving averages are used later in the process, after an estimate of the seasonal component has been removed, and are applied to give refined estimates of the trend and hence also the seasonally adjusted series.

16.3.1 Options for Trend Moving Averages

1. *Centred simple moving averages*: There are two types of centred simple moving averages available that can be manually chosen using the **x11** spec, a centred one year moving average and a two year moving average with a modified Leser filter. The program will default to a centred one year moving average, which is generally recommended. If, however the series that is being adjusted has short cyclical fluctuations or sudden changes in the trend level the two-year option may be preferred. The arguments in the **x11** spec for these one or two year options are respectively
 - `itrendma=centered1yr` (i.e. a centred 12-term moving average for monthly series or a centred 4-term for quarterly series)
 - `itrendma=cholette2yr` (i.e. a centred 24-term moving average for monthly series or a centred 8-term for quarterly series with a modified Leser filter)

2. *Henderson moving averages*: There are 50 types of Henderson moving averages that can be manually chosen using the **x11** spec. Any odd number greater than one and less than or equal to 101 can be specified. In general the ones that will be used are either a 9-, 13- or 23-term for monthly data or a 5- or 7-term for quarterly data, as by default the program will choose one these, if the **trendma** argument is not used in the **x11** spec. That is to say the default option is that the program will choose one of these term Henderson moving averages. As mentioned above, the choice of a trend filter is a matter of balancing smoothing power and flexibility of the trend. Here the choice is based upon the I/C ratio i.e. the size of the irregular variations relative to those of the trend (more detail is given below). The default choice should be accepted in all normal cases; the only common situation in which it may be overridden is if it is necessary to ensure that a group of series all use the same trend filter. To choose a specific length of Henderson trend moving average using the **x11** spec the following argument should be used
 - `trendma=n` (where *n* is any odd number from 3 to 101)

16.3.2 Trend Moving Average selection by the X12ARIMA program

If neither moving average is specified by the user the default is that the first estimate of the trend will use a centred one year moving average and the Henderson trend moving average is selected by the "expert" algorithm using the following criterion.

<u>I/C</u>	<u>Quarterly Data</u>	<u>Monthly Data</u>
0 to 0.99	5-term	9-term
1 to 3.49	5-term	13-term
3.5 and over	7-term	23-term

The I/C ratio is shown in table D12 of the output (the trend estimate) and in table F2H. The program makes these selections because if there is a large irregular component relative to the trend, a longer span of data is required to estimate the trend. If, for example, a 23-term moving average is applied this uses 23 consecutive data points, the 11 months prior to the data point, the point itself and the 11 months after the data point. As discussed above, where an ARIMA model is applied a symmetric Henderson moving average can be used, as the program can use the forecast data from the model. If, however, no ARIMA model is applied, then asymmetric moving averages will be used, the weights of which are available from the Time Series Analysis Branch.

16.3.3 When to change the Trend Moving Average

In general it is recommended that the trend moving average is selected automatically during a seasonal adjustment review via the default option and then fixed for production run purposes. However, in some situations it may be necessary to change the moving averages manually. If for example, the M4 summary statistic, given in table F3 (see chapter 19 for further details), fails it may be necessary to use a shorter moving average, as a longer moving average may be removing some pattern in the irregular. For more information on manually changing the trend moving averages contact the Time Series Analysis Branch.

16.4 Seasonal Moving Averages

Seasonal moving averages are weighted arithmetic averages applied to a each month (quarter) over all the years in the series i.e. a particular seasonal moving average is applied to each column of data, as it is presented in the output. They are used by the X12ARIMA program to estimate the seasonal component of the series. The moving averages are applied to the SI series, that is the series with the trend component removed; thus the seasonal factors for January, for example, are obtained by smoothing the SI values for January in successive years. This procedure is a way of estimating both the seasonal and irregular components. As with trend estimation, the choice of moving average is a matter of balancing smoothing power against flexibility. The moving averages that are applied are a combination of simple moving averages, for example a 3x3 moving average as shown in section two or a 3x5 moving average (three applications of a 5-term moving average).

16.4.1 Options for Seasonal Moving Averages

As with the trend moving averages, the program has a default rule to select the length of average automatically based on the irregularity of the SI series, while it is also possible to manually choose the seasonal moving averages applied to each month. The seasonal filters (seasonal moving averages) that are available include a 3x1, 3x3, 3x5, 3x9 and a 3x15.

If no seasonal filter is specified in the **x11** spec then the program will choose the filter automatically, as described in the following section. The filter chosen by the program is then applied to every month. The options available using the **seasonalma** argument in the **x11** spec are shown below.

x11{seasonalma=name} #(options are shown below, with their respective descriptions)

<u>Name</u>	<u>Description of the option</u>
1. msr	Applies default option
2. s3x1	Applies a 3x1 moving average to all months/quarters.

3. s3x3 Applies a 3x3 moving average to all months/quarters.
4. s3x5 Applies a 3x5 moving average to all months/quarters.
5. s3x9 Applies a 3x9 moving average to all months/quarters.
6. s3x15 Applies a 3x15 moving average to all months/quarters.
7. stable Applies a stable seasonal filter to all months/quarter.
8. x11default Applies a 3x3 moving average to all months to calculate the initial seasonal factors in each iteration, and a 3x5 moving average to calculate the final seasonal factors.
9. (snxm snxm
snxm snxm) Allows the application of a specific *nxm* moving average for quarter 1, 2, 3, and 4
10. (snxm snxm
snxm snxm
snxm snxm
snxm snxm
snxm snxm) Allows the application of a specific *nxm* moving average for month 1, 2, 3, ..., 11, and 12

The criterion for selection of the seasonal moving average when the default is activated (either by not using the **seasonalma** argument or by specifying **x11{seasonalma=msr}**), is based on the global I/S ratio, which is shown in table D10 and F2H of the output. The global I/S ratio measures the relative size of irregular movements and seasonal movements averaged over all months or quarters. It is used to determine what seasonal moving average is applied using the following criteria:

<u>I/S</u>	<u>Seasonal moving average applied</u>
0 to 2.5	3x3
3.5 to 5.5	3x5
6.5 and over	3x9

The global I/S ratio is calculated using data that ends in the last full calendar year available. If the global I/S ratio is greater than 2.5 but less than 3.5 or greater than 5.5 but less than 6.5 (i.e. it does not fall in the bands specified above) then the I/S ratio will be calculated using one year less of data to see if the I/S ratio then falls into one of the ranges given above. If it still does not fall into one of these ranges, the I/S ratio is calculated with another year removed. This is repeated either until the I/S ratio falls into one of the ranges or after five years of data have been removed in the calculation of the I/S ratio. If after five years have been removed and the resulting I/S ratio still does not fall into the above range then a 3x5 moving average will be used. The purpose of this procedure is to avoid instability on update; for example, if a 3x5 average has been used but the updated I/S ratio falls below the lower limit for 3x5, the choice will be changed to 3x3 only if the ratio falls clearly within the band for 3x3.

As can be seen in the selection criteria, the larger the I/S ratio - indicating that the irregular component is large relative to the seasonal component - the longer the moving average applied.

16.4.2 Manual selection of Seasonal Moving Averages

Section 4.1 showed two options for using different moving averages (options 9 and 10), depending on whether the series is quarterly or monthly. If for instance we know that there is a systematically low level of activity in a particular month (or quarter) then a different seasonal moving average may be applied for that month (or quarter). A good example of this might be an average wage series, which includes bonus payments every April. If the default returned a 3x3 moving average, based on the global I/S ratio, but from table D9A in the output the I/S ratio for April was greater than 6.5 (indicating that this month is volatile, i.e. the irregular component is large relative to the seasonal component) then as there is an identifiable reason for this volatility (bonus payments tend to be very volatile relative to the average wage), it is valid to specify the use of a 3x9 average for April only by including the following line in the **x11** spec.

```
x11{seasonal ma=(s3x3 s3x3 s3x3 s3x9 s3x3 s3x3 s3x3 s3x3 s3x3 s3x3 s3x3 s3x3)}
```

Another option described in section 4.1 was selection of a 3 term moving average (3x1). This option uses only three years of data, which allows seasonality to change very rapidly over time. This option should not normally be used, as it will usually lead to large revisions as new data becomes available. If this is to be used it will usually be in one month, and because there is a known reason for wanting to track fast changing seasonality. For example, if a new data collection method has been introduced but a seasonal break has not been identified, it may be beneficial for the first few years to track this changing seasonality closely with a shorter moving average.

The option **stable** means that all of the values from the month or quarter are used to calculate the average. This average would only be suitable where the I/S ratio is very high, and in general should not be used. If there are less than five full years of data in the series, then the program will by default use the stable moving average option. However, it is possible using the **x11** spec to specify that a seasonal moving average, specified in the **seasonalma** argument by using the **sfshort** argument with less than five years of data. There are two options with this argument either **sfshort=no** (default), which means that a stable seasonal filter will be used, or **sfshort=yes**, which will then use the seasonal moving average specified in the **seasonalma** argument.

A final option that is perhaps not so self-explanatory is **x11default**. This specifies an option for the program to use moving averages as they were used in previous version of X-11 and X-11-ARIMA. If no seasonal moving average is specified then a 3x3 moving average is used to calculate the initial seasonal factors in the first two iterations, and then a 3x5 moving average is used to calculate the final seasonal factors. This option is not recommended.

16.4.3 When to change the Seasonal Moving Averages

The default options will, in general, select the most appropriate moving average. However there will be occasions when the user will need to specify a different seasonal moving average to that identified by the program. In order to identify those occasions when user intervention is necessary the following guidelines provide some useful indicators.

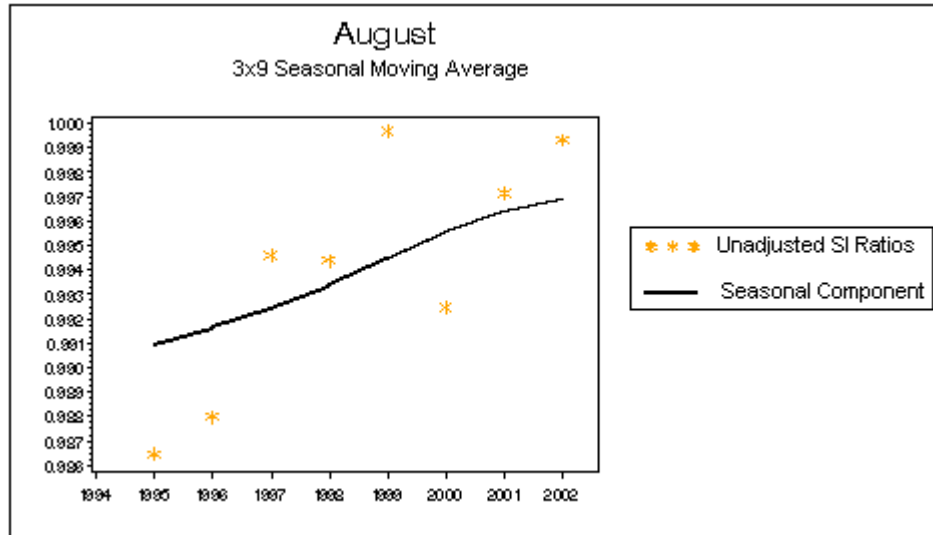
Some of the quality statistics given in table F3 of the output, particularly M4, M6, M8 and M9, can indicate that changing the seasonal moving average may be beneficial (see chapter 19 for more details).

SI ratios given in table D8 and/or these ratios plotted against the seasonal components given in table D10 (see below for more detail and Chapter 20).

If SI ratios are plotted with the seasonal component of the series for each month or quarter, then if the SI values do not closely follow the seasonal component, it may be appropriate to use a shorter moving average, i.e. a more responsive moving average. For example, the graph below plots the SI ratios and


the seasonal component for the month of August. In this case a 3x9 moving average has been selected by the global I/S ratio, but this does not track the seasonality so well.

Figure 16-1 3x9 seasonal moving average



If, however a 3x3 moving average is used this tracks the SI ratios more closely and will improve the seasonal adjustment

Figure 16-2 3x3 seasonal moving average



An alternative to looking at these graphs is to look at the I/S ratios for each month; these are shown in table D9A of the output. Caution should be exercised when selecting a different seasonal moving average for a particular month, as this may cause unwanted revisions. In the majority of cases the automatic selection procedure of the X12ARIMA program - based on the global I/S ratio, estimated from the entire series - is sufficient.

If the size of the irregular component has changed through the series, for example, if the sample size of the survey has changed, then different moving averages may be more appropriate for different sections of

the series (as highlighted in the previous section). In this case it will usually be best to use the length of moving average most appropriate for the recent section of the series. Seasonal breaks can also distort the automatic selection of the most appropriate seasonal moving average (see chapter 17).

16.5 Updating

When running X12ARIMA for an annual update in the seasonal adjustment review, the following steps are recommended.

1. Run the seasonal adjustment on the default trend moving average options (i.e. the program selects the trend and seasonal moving average);
2. From the output find the Henderson trend moving average and the seasonal moving average that are selected;
3. Plot the SI ratios and seasonal components, check the I/S ratios in table D9A and ask the producer of the series about any patterns that could be adjusted for;
4. In the **x11** spec specify the argument {trendma=*n* and seasonalma=*name*} (where *n* is the Henderson trend identified in step 2 and *name* is the seasonal moving average option chosen from Steps 2 and 3);
5. Use this spec file for seasonal adjustment of the series for the following year.

16.6 Cross-references

- Overview of the X12ARIMA method
- Procedures for analysing series with X12ARIMA
- Output and diagnostics
- Graphs and SAS graphics
- Implementation and monitoring

17 SEASONAL BREAKS


17.1 What is a seasonal break?

A seasonal break is defined as a sudden and sustained change in the seasonal pattern of a series. There are many potential causes of seasonal breaks in series, including changes in the data source, methodological changes, or administrative changes. Seasonal breaks will often be accompanied by a level shift.

The following graph provides an example of a seasonal break:

Figure 17-1 Example of seasonal break

Example of Seasonal Break - Car Registration Series



There is a stable pattern of one major peak and one major trough per year prior to January 1999. After January 1999, there is a sudden and permanent change in the seasonal pattern to two peaks and two troughs per year. This particular change resulted from an alteration in the car registration system.


17.2 Why adjust for a seasonal break?

Seasonal breaks are a problem for seasonal adjustment because the methodology is based on moving averages. A moving average is applied to the seasonal irregular component (SI ratios, which can be found in the D8 table within X-12-ARIMA output) of the series to obtain the estimate of the seasonal component (that will be removed by seasonal adjustment). But the moving average used for this purpose is designed to deal with series which have a smoothly evolving 'deterministic' seasonal component plus an irregular component with stable variance. If there is a seasonal break in the series it will be reflected in the SI ratios. When the moving averages are applied to the SI ratios to estimate the seasonal component, the estimate of the seasonal component will be distorted. The result is "leakage".

Leakage occurs when part of the variation of one component has been incorporated into the variation of another component. The result could be that either some seasonal variation is left in the irregular component (and so not all the seasonal variation is removed from the seasonally adjusted series); or some variation that is not caused by seasonality is removed from the series. The result in both cases is increased volatility in the seasonally adjusted series.

In the case of seasonal breaks, the leakage is of both types. The break in the seasonality causes distortion in the estimation of the seasonal factors (D8 table). Specifically, part of the step change at the break will leak into the irregular component. This has the effect of making the irregular in the vicinity of the break look larger, while turning the step change in the seasonal factor into a smooth transition. Besides the distorted seasonal factors, this has an important effect on the main diagnostic for stable seasonality (the F-test in Table D8A); the artificially large irregulars will inflate the residual mean square, and so reduce the significance of the F value.

Figure 17-2 SI ratios



In the example above, there has been a sudden drop in the level of the SI ratios for August between 1998 and 1999. Moving averages are applied to these ratios in order to estimate the seasonal component of the series for August. In this example, for the years prior to 1999, the estimates for the seasonal factors will be lower than they should be. When these seasonal factors are applied to the original data to produce the seasonally adjusted series, some of the seasonal variation will remain in the irregular component (it has leaked into the irregular series), resulting in residual seasonality in the seasonally adjusted series. Conversely, after 1999 the seasonal factor estimates will be higher than they should be. The result of this is that variation is removed from the seasonally adjusted series that is not seasonal. This means that the seasonally adjusted series will not be reflecting the economic behaviour of the series. The result in both cases is a higher level of volatility in the seasonally adjusted series (as shown in the graph below), and a greater likelihood of revisions.


Seasonal breaks cause distortion in the estimation of seasonal factors before and after the break. This distortion is often quantitatively significant. It can influence several years of data, depending on the length of the moving average, and, if not corrected for, can generate a seasonal pattern in the seasonally

adjusted series. In the example above, where a 3x3 moving average was applied, the estimates of the seasonal factors will be distorted for 5 years. If the seasonal moving average applied to the majority of series were a 3x5 moving average, the distortions in the seasonal factors would occur for seven years. As a result, the seasonally adjusted series would appear to be more volatile near the break and would not reflect the underlying behaviour of the series.

Furthermore, as with trend breaks, seasonal breaks create problems in the identification of trading day and Easter effects and in fitting an ARIMA model. Hence, adjusting for seasonal breaks improves other parts of the seasonal adjustment process.

The graph below shows how the seasonal adjustment for the Car Registration series is very volatile and does not reflect the underlying behaviour of the series when no adjustment is made for the seasonal break.

Figure 17-3 Example of seasonal break



17.3 How to identify and adjust for a seasonal break?

The following paragraph describes how to identify and adjust for a seasonal break. The paragraph has been structured following a logical order that should be used in the process of analysing the seasonal break.

17.3.1 Run the series in X-12-ARIMA and look at the output diagnostics

In some cases, there will be prior knowledge that a seasonal break is expected. In the case of the car registration series, for example, analysts knew of the administrative change in registration procedures in advance and that its purpose was to change the seasonal pattern. In such cases we can go straight to the tests described in 17.3.3 below. In other cases we should carry out a preliminary screening to see if there

is any indication of a seasonal break. The following steps suggest a systematic way of carrying out this screening.

- Run the series using the standard spec file in X-12-ARIMA without testing for a seasonal break at any particular point.
- As a first test, look at the diagnostics in Table D8A (F-tests for seasonality). If the stable seasonality test is significant but the moving seasonality test is non-significant, it is unlikely that there is a seasonal break. If the moving seasonality F-value is large, the seasonality is changing considerably, and it is worth checking to see if some of this change occurs in steps – proceed to the next check.
- Graph the series and look at the graphs of SI ratios against the seasonal factors (the output section of this guide gives instructions on how to do this). A seasonal break will appear as a sudden change in the level of the SI ratios for a particular month/quarter. The adjustment for the seasonal break effectively attempts to remove this sudden change in the level of the SI ratios.
- The seasonal break identification is not always as obvious as the example reported in this chapter. In those cases where the seasonal break cannot be spotted from the graphical representation of the series, an analysis of the X-12-ARIMA output can help in the breaks detection. In addition to the inspection of the SI ratios, seasonal breaks can be detected by analysing tables C17 and E5 in the output: when a seasonal break occurs, C17 may show a concentration of outliers within a particular month/quarter and E5 may show a sudden change in the size of the month-to-month/quarter-to-quarter changes for one or more months/quarters.

17.3.2 Testing for seasonal break using X-12-ARIMA

The regARIMA modelling stage of the program provides a function called a seasonal change of regime variable, which is specifically designed to model a seasonal break. If there is reason to suspect a seasonal break is present at a particular time point (either before or after the SAS graphics analysis) then it can be specifically tested for using X-12-ARIMA by including a command line of the following form in the X-12-ARIMA regression spec.

```
regression{
    variables=(seasonal/1999.jan//)
    save=(rmx)
}
x11{mode=mult}
```

This command is appropriate to the car registration series above, where the presumed date of the break is Jan 1999. The **variables=(seasonal/1999.jan//)** command instructs X-12-ARIMA to analyse for a change in the seasonal component; the regressors on and after the change date are equal to zero. When it is included in the regression spec, the output includes t-tests for each month/quarter regressor and a Chi-Square test to verify the significance of the regressors as a group. The definition of these tests is in the regARIMA chapter.

This command allows the user to test, but not adjust, for a change in the pattern of the seasonal component. This means that the data passed to the x11 spec will still be affected by the break, and so the calculated seasonals and the D8A F-test will be distorted in the way described. Hence it is unwise to place any reliance in the output of the x11 spec in the sequence above. Indeed, it may be sensible to run just the regARIMA part at this stage, without any x11, until decisions have been made on whether a seasonal break is present.

The date specified for the change of the seasonal component (e.g. Jan 1999 in the example above) divides the series into two spans. The first span contains the data for periods prior to this date and the second span contains data for periods on and after this date. Including the command **variables=(seasonal/1999.jan//)** in the regression spec means that X-12-ARIMA has been asked to estimate partial change of regime variables for the early span, where "partial" means that the change of the seasonal component is restricted to the early span.

This type of partial change of regime can only be used in conjunction with another component which models the seasonality over the whole series. This could be either a fixed seasonal or a seasonal difference term in the ARIMA model. With this combination, the whole series model deals with seasonality on both sides of the break, while the partial change of regime estimates how much must be added to the whole series model to represent seasonality before the break. The result of this analysis can be used to estimate permanent prior adjustments for the data points of the first period.

The effect of the change of regime variables is to include in the regARIMA model a set of regressor variables (11 in the case of a monthly series, 3 for quarterly) which show the contrast between one month and another. If these regressor variables can be saved or generated by some other means, they can be included as user variables with exactly the same effect as the change of regime. In the example above, the **save=(rmx)** command has been included in the regression spec to save the regressor variables with the associated dates so that they can be used later for the seasonal adjustment.

The use of regARIMA regressors is only recommended if the presence of a seasonal break has been previously confirmed and the Chi-square p-value is less than 0.05. This is due to the chi-square test being biased in defining a seasonal break as significant (i.e. it has a high type I error). In addition, it is recommended that the seasonal break analysis is not used if the series is judged to be not-seasonal.

One possible drawback of the **variables=(seasonal/1999.jan//)** command is that the estimated effects of the permanent priors do not always balance out, leading to the problem that the level of the series does not remain constant. This means that the permanent priors should be checked to see if they compensate (that is any increase must be matched by an equal decrease), especially in current price series. [The change of seasonal will be self-compensating automatically for an additive model, because only 11 factors are estimated, the twelfth being calculated to make them sum to zero. For multiplicative this applies to the logs of the seasonals, which may not cancel out exactly when transformed back to the original scale, but unless the seasonal changes are huge this should not be a major effect. Do we have evidence that this effect is big enough to be visible against the background of the irregularity?]

17.3.3 Confirming a reason for a seasonal break with the host branch

If the X-12-ARIMA test and the graphs from SAS graphics lead you to suspect a seasonal break, then check which months/quarters it appears in. Ask the host branch if there is any reason for suspecting a seasonal break at this time point.

If there is not an economic reason to explain for the presence of the suspected seasonal break then **do not adjust** for it.

17.3.4 Length of the series before and after the seasonal break

1. The entire length of the series in total needs to be at least 5 years to use regARIMA.
2. At least 1 year of data either side of the break is needed to be able to use regARIMA as a tool for analysing the break.

3. If less than 1 year is available after the break, consider not publishing the seasonally adjusted version of the series until more observations are available. This might not be possible if the series is a component of an aggregate, in which case adjustments will probably need to be judgmental. In that case use external information (e.g. forecasts of sales patterns produced by the car industry ahead of registration change) or patterns in related series if possible.
4. If 1-2 years (12-24 months or 4-8 quarters) of data are available after the break, consider not publishing the seasonally adjusted version of the series until more observations are available. X-12-ARIMA can provide adjustments, but they are generally of very poor quality and subject to large revisions as future observations become known. X-12-ARIMA estimates can therefore either replace, be used in combination with, or validate ad hoc attempts to adjust for the break.
5. If 2-3 years (24-36 months or 8-12 quarters) of data are available after the break, consider forecasting 1 year of data beyond the end of the series using regARIMA. If the seasonal pattern after the break looks very regular (both the model and the SI ratios are stable), then publication of the seasonal adjustment can be reinstated at this point.
6. If more than 3 years of data are available after the break, publication of the seasonal adjustment can generally be reinstated. But if the model is exceptionally poor, the series is exceptionally erratic or SI ratios are exceptionally inconsistent between years, do not reinstate publication at this point.

17.3.5 Adjust for the seasonal break

Use one of the three following methods to adjust for the seasonal break.

17.3.5.1 Use of regARIMA to derive permanent priors

It is possible to use the regARIMA modelling capabilities of X-12-ARIMA to calculate the permanent prior adjustments. As discussed above, the effect of a change of regime can be reproduced by including appropriate user variables. These variables can be generated by formula or by saving the regression matrix in a run with a change of regime: the latter is probably the easier approach.

This means that it is necessary to re-run the spec file removing the **variables=(seasonal/1999.jan//)** and **save=(rmx)** arguments from the regression spec and including the variables previously saved in the .rmx file as user-defined regression variables. For a monthly series there should be 11 variables, so the spec could be of the following form:

```
regression{user=(M1, M2, M3, M4, M5, M6, M7, M8, M9, M10, M11)
           file="name.rmx"
           format="x12save"}

x11{mode=mult
    final=user}
```

where name.spc is the spec file used in generating the saved variables. One important point is to ensure that the change of regime variables are the first 11 variables in the regression matrix, which will contain all the variables used in the ARIMA model, including Easter, trading day, level shift etc. The order of variables will normally be the order in which they were mentioned in the spec which generated them. This can be checked by reading the saved rmx file into Excel and examining the column headings; the change of regime variables will have headings reading "Jan I", "Feb I"... "Nov I". If any columns other than the date precede these, they should be deleted and the file saved (as a text file).


There are two ways of ensuring that the permanent priors are reflected in the final seasonal adjustment. One is shown above, the inclusion of **final=user** in the **x11** spec. The alternative is to include the argument **usertype=seasonal** in the **user** argument of the **regression** spec. These should **not** be used together. Thus if **usertype=seasonal** is invoked, the command **final=user** should be omitted from the X11 spec. The **usertype** method includes the prior adjustment as part of the seasonal, while the **final** method treats it as an ordinary prior. As a result, although the two methods give the same final seasonally adjusted series they give somewhat different quality diagnostics (the M and Q statistics) because of the different breakdown of the original series.

The permanent priors derived by X-12-ARIMA are shown in table A9 (or A10 if the **usertype** argument is used in the regression spec). These prior adjustments can also be derived from the parameter estimates of the regression model used to parameterise the seasonal component. In other words the permanent priors are equal to the parameter estimates themselves if the series is additive and a seasonal variable has been included in the regression spec in conjunction with a change in regime option. The permanent priors, over any single year, should average out to approximately 100, for multiplicative cases.

The permanent priors derived by the regARIMA regression accurately adjust for the leakage of the seasonal variation into the irregular component, and the estimated seasonally adjusted series and trend component are more robust than the ones estimated using X11ARIMA. They also cope well in case of fast moving seasonality.

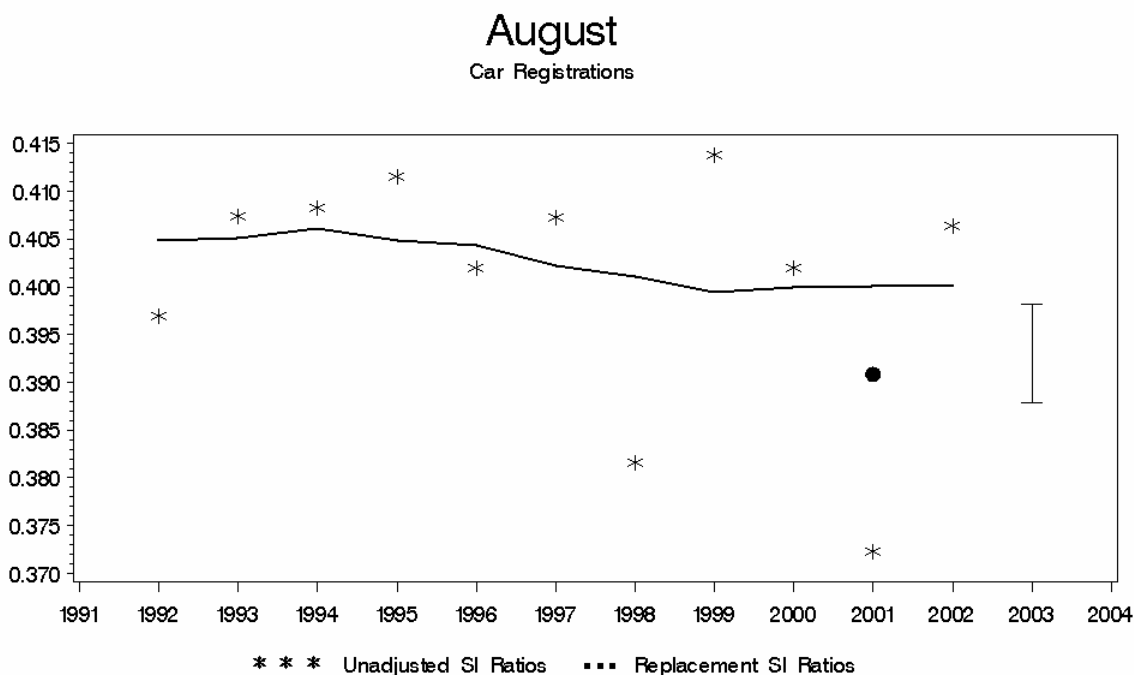
The graph below shows the improvements in the result of seasonal adjustment. In fact, the SA series does not show any residual seasonality either before and after the break.

Figure 17-4 Example of seasonal break



The improvement in the seasonal adjustment can also be seen in the SI ratios. The SI ratios for August after the seasonal break adjustments derived by regARIMA are much flatter, so an improved estimation of the seasonal factors can be calculated.

Figure 17-5 SI Ratios



In addition, if a stability analysis is conducted, the sliding spans statistics give good results. This indicates that even in case of short series, the use of seasonal regressors provides stable and possibly good quality seasonal adjustments. For more information on sliding spans statistics, see chapter 21.

Problems:

- According to the X-12-ARIMA documentation, the **usertype** arguments that can be used are **seasonal** and **userseasonal**. The **seasonal** option should automatically consider the A10 table as permanent priors, but only after the calculation of the x11 moving averages, while **userseasonal** should be used in order to have the original series prior adjusted before the seasonal adjustment specified in the x11 phase. However, contrary to the documentation, at present the program treats **userseasonal** as an invalid argument, while **seasonal** operates in the way that **userseasonal** is supposed to. The procedures described above have been tested and found to work on the current version of the program. Presumably the program and documentation will be brought into line; depending on how this is done, it may at some time be necessary to replace **seasonal** by **userseasonal** in the specifications above.
- The interpretation of the seasonal diagnostics needs to be done with care if this approach is used. The usual Table D8a test for seasonality is based on the series input to the x11 phase, which will have been prior adjusted for the seasonal break. Therefore this test may not give a true picture of seasonality over the whole series; it will be dominated by the seasonal behaviour after the break. Also, the break adjustment applies a fixed seasonal effect to the series before the break; if the seasonal pattern was evolving before the break, the effect of the break adjustment may be to magnify the relative effect of the movement. In some circumstances, these two effects may move the diagnostic for identifiable seasonality towards “not present”.
- If the break is non-compensating, the estimated permanent priors have to be modified so that any adjustment in one month can be balanced by a compensating adjustment in another month (or spread across several months). In this case the seasonal adjustment of the series should be run using the modified adjustments as permanent priors in the transform spec.

- The analysis of seasonal breaks can be difficult in series with $MCD = 1$ or 2 . In fact, in those series the influence of the irregular component is very small and the trend and seasonal components mainly drive the behaviour of the raw data. In this situation any small change in the seasonal factors is considered as a significant change in the seasonal pattern and produces a chi-square less than 0.05 . Therefore for these series the chi-square test is less useful, meaning that a seasonal break analysis should rely more on the SI graphs analysis, with a seasonal break only being adjusted for if an economic reason has been confirmed.
- The presence of other discontinuity factors (i.e. additive outliers or level shifts) not already adjusted for using regressors in the regARIMA may affect the performance of the chi-square in the analysis of seasonal breaks. For this reason, it is recommended to check the significance of the seasonal break and to use the estimated permanent priors calculated after the adjustment for outliers or level shifts has taken place.
- Depending on the nature of the seasonal pattern and how it changes, it may be difficult to say from the data exactly when the break occurs. The chi-square can be significant (less than 0.05) in a range between 7 months before and 9 months after the actual date of the seasonal break. Therefore it is useful to check on the timing of the seasonal break in the C17 output table. In fact, if C17 presents, within this range, one or more zeros before or after the date indicated in the regression spec (i.e. variables=(seasonal/1999.jan/)), it is necessary to change the date defined in the regression spec, anticipating or postponing a number of months/quarters equal to the number of zeros. Of course, all this should be considered in conjunction with the views of the host branch on when the events causing the break may have occurred.

17.3.5.2 Adjust the two parts of the series separately

Although it is often easier to use the regressors, if there is more than five years of data before and after the break the two parts of the series can be seasonally adjusted separately.

This is probably most appropriate for "historic" seasonal breaks where the period of the break is not going to be revised in published data. If the path through the period of the break is important, treating it as one adjustment with a break is probably preferable. Furthermore, if it is a more recent break, adjusting the two parts of the series separately is not an option.

Problems:

- Adjusting the two parts of the series separately can create discontinuities in the seasonally adjusted series.
- It is possible to use this method only with long series that contain more than five years of data before and after the break.

17.3.5.3 Use of output tables to manually calculate the permanent priors

It is possible to calculate manually permanent priors for seasonal break from table D8 of the output using the following procedure:

- Run the monthly/quarterly series in X-12-ARIMA using a default spec;
- Use permanent, temporary priors and Easter and trading day adjustment if they have already been defined, otherwise, use default settings for all options.

- This default run should include all the data. From Table D8 take the average of the values within each month/quarter before the break, and similarly the average of the values after the break.
- For each month/quarter divide the average before the break by the average after the break to give the permanent prior for that month/quarter
- Apply the permanent priors for the appropriate month/quarter to all data points before the break

Permanent priors may already be pre-defined for Easter or other reasons such as errors in the data. To incorporate these with the permanent priors for the seasonal break multiply them together and divide by 100 for multiplicative models or add them together for additive models.

Problems:

- This method can be used to validate the permanent priors derived by the regARIMA model, but it falls down under fast moving seasonality.
- It may create additivity problems.

Criteria for deciding which of the three methods should be used are as follows:

1. **Length of the series before and after the break.** The length of the series before and after the break influences the decision of which of the three methods should be used. In fact, if 1-3 years of data are available either side of the break, method 3.5.1 should be used in conjunction with method 3.5.3 to validate the quality of the derived permanent priors. If more than 3 years of data are available after the break, method 3.5.1 should be used. If more than 5 years of data are available after the break, method 3.5.1 and 3.5.2 can be used, although method 3.5.1 is preferable since it make the process of updating the prior adjustments easier.
2. **Ease of use/updating.** The use of regressors in the model (method 3.5.1) is easy to do and facilitates the update of the permanent priors. This is particularly useful when less than three years of data are available after the break and the parameters need to be re-estimated frequently until they become stable. Also the method for adjusting the two parts of the series separately (3.5.2) is easy to use, but requires at least 5 years of data either side of the break. On the other hand the manual calculation of the permanent priors (method 3.5.3) is more elaborate and makes the update more difficult.
3. **Multiple seasonal breaks.** A particular strategy needs to be adopted in this case, since it is not possible to define two or more sets of change of regime variables in a regARIMA model to correct for seasonal discontinuities. In case of multiple seasonal breaks, it is necessary to analyse the series in stages. Starting from the first seasonal break, test for significance using a change of regime specification, and if significant include user variables from the saved .rmx file in all later stages. Repeat the process for each of the subsequent breaks, including in the regARIMA model change of regime regressors only for the latest break. Finally all identified breaks will have user defined variables, which will be used for production running.

Related Topics

- Length of the series
- The regARIMA model

- Moving Averages

18 EXISTENCE OF SEASONALITY

18.1 Introduction

The purpose of seasonal adjustment is to remove seasonality from an observed time series. For this, it is first necessary to decide whether the series is seasonal (i.e. whether there is a seasonal component to remove or not) and how accurately the seasonal component is estimated. In fact, X12ARIMA will always estimate a seasonal component and remove it from the series even if no seasonality is present, but not all the estimates of the seasonally adjusted series will be good.

18.2 A general criterion for existence of seasonality

Empirical research showed that the most appropriate test for seasonality is the "Combined test for the presence of identifiable seasonality", given after table D8 of the output. In particular, one of the following statements will always appear:

1. IDENTIFIABLE SEASONALITY PRESENT
2. IDENTIFIABLE SEASONALITY PROBABLY NOT PRESENT
3. IDENTIFIABLE SEASONALITY NOT PRESENT

It is recommended that a series is adjusted in the first two cases and not adjusted in the last one. However there are two cases where one might need to deviate from this practice:

18.3 Marginal seasonality

When a series is only marginally seasonal, the criterion might swing from "non seasonal" to "seasonal" as new data become available. It might however not be desirable to change the policy of seasonally adjusting or not all the time, as this causes revisions to the published seasonally adjusted series and revisions back (i.e., when a series is deemed non-seasonal its seasonally adjusted version is same as the unadjusted, but when one actually adjusts the series the adjusted version is a different series). This problem is approached by introducing a "same as last time" area between the "seasonal" and "non seasonal". The alternative diagnostic that is used for this purpose is the M7 statistic, which contrary to the "identifiable seasonality test" (IDS) is a continuous variable. Furthermore, it is very highly related to the test: Although there is no clear-cut value of M7 for which the IDS test passes or fails, the area of M7 for which the IDS can go either way is rather narrow. In other words, the M7 diagnostic is approximately a continuous variable for the IDS test, and for this reason it is very useful in defining the "same as last time" area for series that are only marginally seasonal.

18.4 Strategy for deciding whether to seasonally adjust

The following guidelines should be used:

1. If a series is being looked at for a first time then adjust it if the seasonality test after table D8 gives either "identifiable seasonality present" or "identifiable seasonality probably not present". Do not adjust if the test gives "identifiable seasonality not present".
2. If the series was adjusted in last year's re-analysis then continue adjusting it, unless the seasonality test shows absence of seasonality **and** the M7 is higher than 1.250 (monthly series) or 1.050 (quarterly series).

3. If the series was not adjusted in last year's re-analysis then continue not adjusting it, unless the seasonality test shows that there is seasonality present **and** the M7 is lower than 1.150 (monthly series) or 0.900 (quarterly series).

It must however be emphasised that these margins may not be appropriate for all series. It is therefore recommended that one also uses judgement, especially when the decision of whether to adjust a series or not has already changed at least once in the past. Finally, it should be mentioned that both the IDS test and the M7 are saved in the "log file", if this is explicitly specified in the command file. This makes it easier to decide for many series at the same time, i.e., by using a data metafile.

18.5 Indirect Seasonal Adjustment

For many data sets it is quite common to do the seasonal adjustment at a low level of aggregation and add the seasonally adjusted series together to get the seasonally adjusted version of an aggregate series. The problem with this approach is that there is often residual seasonality in the aggregate series. If one attempts to adjust the (indirectly) seasonally adjusted aggregate series, the IDS test often shows that seasonality is present.

The cause of this problem is that the more volatile the irregular component is, the more difficult for any seasonality in the series to be detected. As the irregular component is more volatile for the lower level series, some of them are deemed non-seasonal. Although this is acceptable for the individual series, when added together their suppressed seasonal components become significant.

X12ARIMA provides assistance in investigating direct versus indirect adjustment in the form of the **composite** spec. In this case, a conventional spec file is provided for each component series, while for the aggregate the **series** spec is replaced by a **composite** spec. The link between the components and the aggregate is provided by an input metafile, with filename extension .mta. The program will run a conventional adjustment for each component and for direct adjustment of the aggregate, while also calculating the indirectly adjusted aggregate. The full range of diagnostics is provided for each case, including the indirect aggregate; thus it is easy to compare the quality of the two aggregate adjustments, and in particular to see whether there is residual seasonality in the indirect adjustment. The case in which some of the components are not adjusted is catered for by setting **type=summary** in the **x11** component of their spec files.

Ideally one would prefer to adjust each series individually, but it is often the case that certain additivity constraints need to be respected. One option is to use a direct adjustment and apportion the differences between the direct and indirect adjustments to the component series. This might not be easy though, on the one hand because of the amount of calculations required and on the other because it is very possible that the constraining process will generate residual seasonality in some of the component series. An alternative approach is therefore recommended, which requires the seasonal adjustment of more of the component series in order to exactly capture the part of the seasonality that is negligible for an individual component series but significant in the aggregate. The decision as to which component series should be seasonally adjusted can be based on IDS test. Whereas the general IDS test is based on three different tests (which are listed just after table D8):

- **Test for seasonality assuming stability;**
- **Nonparametric test for the presence of seasonality assuming stability;**
- **Moving seasonality test;**

the alternative test consists of seasonally adjusting a series if just one of the first two tests passes.

As more of the component series are now adjusted the (indirectly adjusted) aggregate series is unlikely to have any residual seasonality. If however it does, then the best option is to adjust all component series in the data set regardless of seasonality tests.

Of course when this alternative test is used, the guidelines of the previous section are not applicable anymore. An alternative approach should therefore be used for series that, by the IDS criterion, swing between "pass" and "fail". The answer to this problem is that it is not very important whether to adjust or not, as these series would not be adjusted under normal circumstances anyway. Judgement should therefore be used, aimed on the one hand at eliminating any residual seasonality in the aggregate series and on the other hand at avoiding unnecessary revisions caused by frequent changes of the decision to adjust or not.

18.6 Testing a large dataset for seasonality

The following steps should be followed in order to test a large data set for seasonality:

1. Specify the X12ARIMA command file as follows:

```
x11{
  save log=(i ds, M7)
}
```

2. Run seasonal adjustment of all series in the data set, using a data metafile.
3. Check the log file, to see for which series the IDS test is "yes" and for which it is "no".
4. Use the strategy in section 3 to decide which of the series should be adjusted. The M7 diagnostic is given in the log file.
5. Keep the seasonal adjustment of those series that in the previous stage it was decided to adjust. For the other series, the seasonal adjustment is same as the original series, minus any possible calendar effects.
6. Derive the indirect seasonal adjustment for all aggregate series. Run it through X12ARIMA to check if there is any residual seasonality (i.e., if the IDS is "yes"). If there is no residual seasonality in any of the indirectly adjusted series then the seasonal adjustment is completed.
7. If there is residual seasonality in any of the indirectly adjusted series, then study those component series that were not adjusted in the first place (or were adjusted only for calendar effects). Adjust those that pass at least one of the two "stable seasonality tests", as described in section 3.
8. Derive the indirect seasonal adjustment of the aggregate series again, also using the adjusted versions of the series that were adjusted in stage 7. Run the indirect seasonal adjustment through X12ARIMA again to check for residual seasonality.
9. If the IDS is "no" then the seasonal adjustment is complete. If it is not, then adjust all component series, regardless of the IDS. The new indirect seasonal adjustment will not have any residual seasonality.

18.7 Limitations

The above results rely on the assumption that all processing that is required before testing for seasonality was conducted. In particular, several things can obscure seasonal effects. The non-exclusive list includes

- outliers
- trend or seasonal breaks
- large calendar (Easter or trading day) effects

All such effects should be investigated individually and any resultant prior adjustments should be made before testing for seasonality or running seasonal adjustment.

A further limitation is that a series may have identifiable seasonality, according to the criteria in this section, but nevertheless the adjustments may be too unstable for publication. The usual test for instability is the sliding spans analysis (see the relevant chapter for details). If a series with identifiable seasonality has unstable adjustments there is a dilemma; either publish unadjusted, with possible confusion for users, or publish the adjustments and recognise that they will be subject to large revisions on update. There is no hard and fast rule for these cases; the size of the seasonality and the importance of the series to users will have to be taken into account in forming a judgment.

19 X12ARIMA STANDARD OUTPUT

19.1 Introduction

An X12 run generates the following standard files:

An output (*.out) file for each series that is processed.

An error (*.err) file for each series that is processed.

One log (*.log) file for all series processed, where certain diagnostics requested by the user are stored.

Of them, most important is the output file, which gives the results of the seasonal adjustment as well as useful quality diagnostics. The output file consists of a fairly long list of Tables, which are identified by a capital letter followed by a number, and they are organised as follows:

"Tables A" (i.e., preceded by the letter "A"): These tables show the prior adjustment of the series. These include any prior adjustments that are specified by the user in the "transform" spec, but also the effects of the regression part of the regARIMA model and any automatically identified outliers or breaks.

Tables "B", "C" and "D": A seasonal adjustment run consists of 3 iterations of the "X11 method". The output of the first iteration is saved on the B-tables, the output of the second is saved on the C-tables, and the output of the final iteration is saved on the D-tables. Of course only the D-tables are final.

Tables "E" and "F": These tables provide diagnostics of the seasonal adjustment.

Tables "G": These are graphics (more in chapter 20).

Tables "R": These tables show the **revision histories** (more in chapter 22).

Tables "S": These tables show the **sliding spans** diagnostics (more in chapter 21).

Some of these tables are more important than the others and in fact the majority of the tables are not even printed in the output unless specifically requested by the user. Therefore it is deemed more appropriate to organise the presentation of the output tables not by their numbers but by the purpose they serve. This is done in sections 2 and 3, while section 4 describes the log and error files.

19.2 Output Diagnostics

Before using the results of the seasonal adjustment the users need to satisfy themselves that the adjustment is of a good quality, or in other words to check the **diagnostics**. An X12 output starts by repeating the specification file that was used to generate it. It is not a bad idea to check that the specifications are those intended by the user, especially if alternative specifications are tried.

Table A1 follows immediately after and shows the original series. Again, it is not a bad idea to have a quick look to make sure that it is the correct data, especially when it is imported from another file. Usually having the first and last data points correct ensures that the intermediate points have been correctly imported as well.

After these basic checks the output diagnostics are looked at. These can be classified as diagnostics of the prior adjustment and diagnostics of the seasonal adjustment. Since good prior adjustment is a

prerequisite for good seasonal adjustment the users should first satisfy themselves with the quality of the prior adjustment.

19.2.1 Diagnostics of the prior adjustments

Good prior adjustment implies appropriate treatment of any external factors that affect the series. These effects are captured either through the regression part of the regARIMA model, as it was described in chapter 10, or by putting in prior adjustments directly by means of the "transform" spec.

Table A2 shows the total prior adjustments specified by the user. As with Table A1, it's not a bad idea to check that they are indeed those intended by the user. However, if both permanent and temporary priors are used Table A2 will show their combined effect, thus it will be different from each of them taken individually.

For effects captured with the regression model one should look at the regARIMA fit, following the guidance of chapter 10. However, the prior adjustment might be inadequate even if the regARIMA diagnostics are good. This would be the case when significant regression variables are missing. For example, there might be a strong trading day effect that the user did not include or included incorrectly. But the most common cases where a significant variable is omitted are breaks and outliers. These can be identified by means of the following diagnostics:

- Automatic outlier identification: By specifying the "outlier" spec in the spec file, one can automatically identify any unaccounted outliers or level shifts. (Of course only those among them that are "valid" should be included in the regression- see chapter 13). It should also be mentioned that if outliers are automatically detected in the same month or quarter it could be evidence of a seasonal break.
- Table D8 (SI ratios): If for one or more months / quarters there is a sudden increase / decrease this might imply a seasonal break. This is more likely if it occurs for many months / quarters and at the same time.

Example:

D 8 Final unmodified SI ratios

From 1991. 1 to 2002. 4

Observations 48

	1st	2nd	3rd	4th	AVGE
1991	83. 7	100. 6	104. 1	111. 9	100. 1
1992	83. 2	100. 8	104. 5	110. 6	99. 8
1993	84. 7	100. 0	104. 0	112. 1	100. 2
1994	83. 1	101. 6	103. 9	110. 7	99. 8
1995	80. 9	101. 1	104. 9	112. 9	99. 9
1996	81. 8	100. 3	104. 7	107. 0	98. 5
1997	89. 1	100. 1	104. 9	103. 2	99. 3
1998	92. 4	101. 6	104. 2	101. 1	99. 8
1999	94. 1	102. 8	103. 6	100. 1	100. 1
2000	93. 8	102. 3	104. 8	98. 2	99. 8
2001	94. 1	104. 5	102. 4	98. 6	99. 9
2002	94. 2	104. 5	103. 4	97. 5	99. 9

The fourth quarter is persistently higher than 110 up to 1995, falling to 107 in 1996, to 103 in 1997, and finishing below 100. This could be due either to evolving seasonality or to a seasonal break. However, the speed of this evolution (completed in just 3 periods with stability before and after), implies that it is probably a seasonal break. Inspection of the first quarter supports this suspicion; whilst for the period up to 1996 the SI ratio was between 80 and 85, in 1997 it jumps to 89 and there after it is persistently above 90. Noteworthy also is that this rapid evolution of the 1st and 4th quarters takes place at the same time, which again supports the suspicion of a seasonal break. Note finally that one does not need all quarters to be affected to argue that there is a seasonal break; indeed, in the example above the 2nd and 3rd quarters are fairly stable- the break basically consists of shifting activity between the 4th and the 1st quarter.

- Table E5 (month-to-month or quarter-to-quarter changes in the unadjusted series). These are expected to be dominated by the seasonal component, thus it is expected that at least some months / quarters have the same sign all the time. If this pattern changes this implies a seasonal break.

Example:

E 5 Quarter-to-quarter percent change in the original series
 From 1991. 2 to 2002. 4
 Observations 47

	1st	2nd	3rd	4th	AVGE
1991		21. 3	4. 6	8. 3	11. 4
1992	-24. 3	24. 7	6. 1	8. 1	3. 7
1993	-22. 3	18. 6	5. 6	10. 1	3. 0
1994	-24. 7	24. 1	4. 3	8. 3	3. 0
1995	-24. 8	28. 9	6. 0	8. 2	4. 6
1996	-26. 6	26. 9	9. 7	5. 3	3. 8
1997	-18. 2	7. 9	2. 7	-0. 5	-2. 0
1998	-8. 7	10. 5	1. 7	-1. 9	0. 4
1999	-4. 4	11. 7	1. 7	-2. 6	1. 6
2000	-6. 5	9. 5	3. 9	-4. 6	0. 6
2001	-2. 6	11. 9	-1. 6	-1. 9	1. 5
2002	-2. 1	12. 5	-0. 6	-5. 7	1. 0

The above Table is from the same output as the previous example. One can immediately see that the sign for the 4th quarter is positive till 1996 and negative afterwards. Further, although the sign in the other quarters does not change there are some significant changes in the size of the figures: The 1st quarter rises from below -20 to single digits, the 2nd quarter falls from 20-30 to around 10, while the 3rd quarter is affected too. This is good evidence of a seasonal break, which probably occurred some time between 1996.3 and 1997.1.

Further, Table E5 is also useful in identifying outliers or level shifts. Outliers manifest themselves with a large value which is followed by a large change of the opposite sign in the following month; level shifts manifest themselves with large changes which are not balanced with subsequent opposite changes.

- Table E6 (month-to-month or quarter-to-quarter changes in the adjusted series). As with Table E5, outliers show up as large numbers which are followed by a change back, while level shifts don't change back.
- Table C17 (final weights of the irregular component): The weight of each point in table C17 is 100, unless this point has been picked up as an outlier in the irregular component, during the X11

iterations. Thus Table C17 is very useful in identifying problems such as breaks or outliers. For example, if there is a concentration of outliers (values lower than 100) within a particular year or month this could be indicative of a seasonal break; if outliers dominate March and April this may indicate a need for Easter adjustments.

To sum up this sub-section, good prior adjustment consists of the following:

- Appropriate fit of the regression model
- No significant variable is omitted
- All outliers, level shifts and seasonal breaks have been adjusted for.

Finally, it must be emphasised that the diagnostics above listed should be used to help identify a *potential* problem, but should not be entirely relied upon to decide on whether the problem is significant or not. Instead, once a problem is identified it should be modelled in the regARIMA model, and it is through the regression diagnostics that it will be eventually decided whether it is insignificant or it is significant and needs to be adjusted for.

Once the user is satisfied with the quality of the prior adjustment the next step is to check the quality of the seasonal adjustment itself:

19.2.2 Diagnostics of the seasonal adjustment

The first thing to check is the seasonality tests in table D8a and the M7 diagnostic. This will determine whether the series is seasonal or not, in the way that is described in chapter 18. If the conclusion is that the series is not seasonal then one should not adjust the series and needs to look at nothing else.

Assuming that the series is indeed seasonal, one should proceed to check other diagnostics of the seasonal adjustment. These are listed below:

- The M-diagnostics in [Table F3](#). These take values from 0 to 3, and a value higher than 1 indicates a source for potential problems for the seasonal adjustment. In particular:
- M7 is the most important among the M-diagnostics, showing the amount of moving comparing to stable seasonality, or in other words how regular the seasonal pattern is. Although M7 is also used as a test for existence of seasonality, it is important to remember that it is not a binary (existent / non existent) test but it takes continuous values.
- Next most important is M1, which shows how large the irregular component is compared to the seasonal. Failure (i.e., a value higher than 1) of M1 implies that the irregular component is large and therefore it might be difficult to estimate the seasonal component accurately.
- M6 measures the irregular too, but is valid only when a 3x5 seasonal filter is used. Failure of M6 means that a shorter than 3x5 filter should be used.
- Next most important are the M8 through M11 diagnostics, which show the fluctuations in the seasonal component. These diagnostics are useful in showing potential problems in the seasonal component, such as seasonal breaks. M10 and M11 are the same as M8 and M9, but only for the end of the series. Thus comparison of M10 and M11 with M8 and M9 can also help identify problems at the end of the series. It should be also noted that M10 and M11 might fail even if there are no problems with the seasonal adjustment, for example if the fit of the ARIMA model that was used to generate

forecasts is poor. Finally, it must be mentioned that if M7 is high then M8-M11 are likely to be high as well.

- M4 is a measure of autocorrelation in the irregular component. It is a less important diagnostic, as good quality of seasonal adjustment does not require an uncorrelated irregular.
- M2 measures the amount of the irregular compared to a straight-line trend. As a consequence, M2 is misleading if the series has a trend that is not well-approximated by a straight line. (In the US Bureau of Census they disregard M2 completely).
- M3 and M5 measure the irregular compared to the trend. They are not important diagnostics.

Other diagnostics include:

- Tests for residual seasonality, shown after table D11. If the tests show that "residual seasonality is present", it can be eliminated by one of the following means:
 1. If the series is very long one might consider cutting off the first few years, as the cause of the problem might be that the seasonal pattern has changed with time.
 2. Alternatively one might wish to change the lengths of the seasonal filters. This also includes the option of using different filters for different months / quarters, if appropriate.
- Table E6: This Table, previously mentioned as useful for checking for outliers and breaks, can be also used to check for residual seasonality. In particular, the seasonally adjusted series is not supposed to have any seasonal elements. Consequently, if the month-to-month change always has the same sign for some months this might imply residual seasonality. (However if the same sign is present for most of the values of the table this is probably due to steady trend, and not to residual seasonality).
- Table D9a (moving seasonality ratio) gives the annual change of the (preliminary) seasonal and irregular component for each month / quarter. This is used by the programme to automatically select the appropriate seasonal moving average (unless a particular filter was specified by the user). However, sometimes one might need to use different seasonal filters for some months or quarters. This is the case when in a particular month / quarter the seasonal component evolves too fast, or the irregular component fluctuates too much. In such cases one might wish to use a shorter filter for the month / quarter in question. Table D9a can be used to identify a problem of this kind. Of interest is the second line of the Table, which gives the fluctuation of the seasonal component. If this is much higher for one month than the others it probably implies the need for a shorter filter.
- Another criterion for **heteroskedasticity** is the SI ratios (from Table D8 or from the corresponding graph). If they fluctuate too much for one month (yet not in the way that implies a seasonal break), this is also evidence that a shorter filter might be needed.
- Graphical output of the spectrum of some components is produced immediately after the F tables. These graphs (labelled G.0 to G.2) are the only parts of the G outputs which are produced by default. The effect to be looked for is a seasonal or trading day peak in either the seasonally adjusted series or the irregular. A peak is defined as a value which exceeds the adjacent values by at least six "stars" on the plot. Any such peaks are mentioned in a brief note just above the G.0 plot. (If regARIMA modelling has been carried out, a similar plot of the spectrum of the regARIMA residuals appears after the modelling output, with a similar note if any peaks are found.) Any peaks found are also mentioned in the console output and the error file.

- Finally, useful diagnostics are **History analysis**, **Sliding spans**, and **X12 graphs**, which are described in separate chapters.

19.3 The results of the seasonal adjustment

Once the users are satisfied that the seasonal adjustment is of good quality (or at least as good as possible) then they can proceed and use the results of the seasonal adjustment. Before describing the relevant output tables it is useful to mention that there are different lengths of output available depending on how much detail is required. The user can specify the desired output length with the "print" argument in the "x11" spec. The options, from the shortest to the longest output, include:

- print=none. With this option X12ARIMA gives only the results from the estimation of the regARIMA model. In particular, it gives the A-tables, regARIMA coefficients, residual checking (if requested) and forecasts. Nothing related to the seasonal adjustment is printed.
- print=brief. With this option the X12ARIMA output includes everything that is printed with the "none" option, as well as basic seasonal adjustment tables and diagnostics.
- print=default gives the same output as above, only with more seasonal adjustment tables and diagnostics.
- With print=alltables all output tables are printed. Additionally to the previous option, these include the "intermediate" tables from the first and second iteration of the X11 method.
- Print=all gives additionally all graphs that are produced. It is exactly for this reason that no-one will ever need this option, as the quality of the graphs is very bad and there are other graphical tools available.

In fact for most seasonal adjustment reviews the default option is more than adequate. Next the output tables are described, with more emphasis given to the most important of them:

19.3.1 *Most important seasonal adjustment tables*

Tables A6, A7, A8 and A9 show the trading day, holiday, outlier and user-defined regression components respectively, as they are estimated from the regARIMA model. These components include the effects of both programme and user-specified variables, as long as the latter have been assigned the appropriate "usertype" (see chapter 10). The effects given in tables A6-A9 can be used not only for analysis, but also as prior adjustments for production running- as an alternative to re-estimating these effects every time a new data point becomes available.

Table B1 shows the original series, after all prior adjustments- including the regression model and automatically detected outliers. It is effectively the series in B1 that goes through the "X11 box". Table B1a shows the forecasts of the prior-adjusted original series.

Table D10 shows the final seasonal component. On the top of the table the seasonal moving average filter that was used is shown, which should be fixed and used till the next re-analysis.

Table D11 shows the seasonally adjusted series. However, if one wants to constrain the annual totals then Table D11A is the appropriate table.

Table D12 shows the final trend component. On the top of the table the Henderson moving average filter that was used is shown. This should be fixed and used for production running till the next re-analysis.

Table D18 shows the combined trading day and holiday factors that are used in the seasonal adjustment and it is equal to the sum (or product for multiplicative series) of tables A6 and A7. As with those tables, it includes the effects of both programme and user-defined calendar effect variables, and it can be used for analysis or as prior adjustment for production running- if the calendar effects are kept constant between seasonal adjustment reviews.

Also important are Tables C17, D8, D8a, D9, D9a, E5, E6, and F3 which were mentioned in the previous section as useful diagnostics.

19.3.2 Less important seasonal adjustment tables

Few are the times where the user would need any of the tables of this section. Nevertheless, it was deemed useful to briefly present them here:

Table D13 shows the final irregular component.

Table D16 is the sum (or product) of tables D10 and D18 and shows the total calendar and seasonal adjustment to the series.

Table E4 shows the annual totals of the original series divided by the annual totals of the seasonally adjusted series. (In the case of additive seasonal adjustment it is subtraction rather than division). The second column of Table E4 shows the same ratio (or difference) but for the extreme-values-adjusted original and seasonally adjusted series. This table can be used as a quality diagnostic, especially when the annual totals are constrained; in that case, the more the ratio is away from 100 (or the more the difference is away from 0 in case of an additive model), the higher the compromise to the seasonal adjustment that is caused by the annual constraining.

Table E7 shows the month-to-month change in the trend-cycle component (Table D12).

Tables F2A-F2I show certain diagnostics in more detail. These include changes, duration of run, or analysis of variance for certain components of the series. Most of these diagnostics are used to derive the single-value M-diagnostics of table F3.

Of the F2 tables it is deemed useful to mention Table F2E which gives the months for cyclical dominance (MCD), that is, the number of months it takes for the variation of the trend-cycle to become larger than the variation of the irregular component. This is very useful for presentation purposes, because if changes during spans shorter than the MCD are presented they will be dominated by the irregular, thus they will be uninformative and perhaps misleading. Further, if the MCD is greater than 6 this indicates that the series is very volatile and the quality of the seasonal adjustment is not likely to be good.

19.3.3 Tables not printed by default

The user will only need tables in this subsection where an exceptionally deep analysis is needed. Nevertheless, it was deemed useful to briefly present them here. For a detailed presentation one should look at the X12ARIMA or X11ARIMA literature.

Tables C1 and D1 show the series that go through the "X11 box" in the second and third X11 iteration. These series are the original series adjusted for prior adjustments, regression effects, and extreme values identified in the previous X11 iteration.

Tables D2, D4, D5, D6 and D7 show the preliminary estimation of the components of the series, in the final X11 iteration.

Tables B2-B13 and C2-C13 are basically the same as the corresponding D-Tables, but for the first and second X11 iteration.

Table B17 is same as C17, but for the first X11 iteration.

Tables B20 and C20 show the factors by which the extreme values identified in the first and second X11 iteration are adjusted before the following (second or third) iteration.

Tables C15 & C16 are generated only when regression on the irregular is run (X11regression). C15 gives the regression output while C16 gives the resulting trading day, holiday, etc., components.

Table D8B is same as D8, only it marks any extreme values according to whether they were identified during the regression / automatic outlier procedure, or during the X11 iterations. This can be useful information for analysis purposes.

Tables D12B and D13B show the trend and irregular components net of extreme values identified in the regression part of the model (the users are reminded that additive outliers and temporary changes are assigned to the irregular component, while level shifts and ramps are assigned to the trend-cycle).

Tables E1, E2 and E3 show the original series, the seasonally adjusted series, and the irregular component, corrected for extreme values.

Table E11 shows a robust estimate of the final seasonally adjusted series. It is equivalent to table E2, except for those points considered extreme, i.e. those which have been assigned zero weight in table C17.

Table E16 shows the final adjustment ratios.

Table F1 is a "smooth" seasonal adjustment, derived by smoothing the original seasonal adjustment by means of a simple moving average, the length of which depends on the "months for cyclical dominance" diagnostic.

Table F4 is produced only if a trading day component has been included in the model of a monthly series. It shows the effect of the trading day component on the monthly adjusted figures, as a function of the length of month and the day of the week on which the month starts. It may be used as a second check on the plausibility of the trading day effects; the user should ask whether there is some known reason in the pattern of weekly activity which would explain why some figures are high and some low.

19.4 Error and Log files

One **error file** is generated for each series that is processed through X12ARIMA. This file stores the following information:

- Errors in the input file. Examples include syntax errors, internal inconsistencies of the input file, or problems with reading data files.
- Problems encountered during processing, which halted the procedure. Singularity of the regression matrix and non-convergence are such examples.
- Problems encountered during the processing, for which an automatic fix was put in place, of which the user must be warned. Examples include changes to certain procedures due to insufficient data to run them with the options specified by the user (i.e., shorter spans for history analysis).

- Properties of the series, identified during the processing and for which no specific action has been taken. For example, trading day or seasonal peaks identified in the spectrum.

When running a seasonal adjustment of many series with a data or an input metafile one might wish to save certain information for all series in one file, instead of checking a large number of output files. To achieve this, one has to specify the requested information with the **savelog** argument of the appropriate spec. The following is an example of a spec file and the log file that was generated when it was run for a data metafile:

```
series{
    start=1991. 1
    period=12
}
transform{
    function=AUTO
}
automdl{
    MAXORDER=(4 1)
    MAXDIFF=(2 1)
    savelog=(AMD MU)
}
REGRESSION{
    AICTEST=(EASTER TD)
    SAVELOG=AICTEST
}
X11{
    SAVELOG=(IDS M7)
}
```

This spec file saves in the log file the automatic ARIMA model selected, the results of the AIC tests for Easter and trading day, and two diagnostics of the seasonal adjustment (identifiable seasonality and M7), as specified by the **savelog** argument in the appropriate specs. This information is summarised for all series in the log file, an extract of which is next presented:

M-ADD NI ----- X-12-ARIMA run of NI

automean: not significant.

Automatic model chosen : (0 0 0)(0 1 1)

AICtd : rejected

AICeaster : rejected

Identifiable seasonality : no

M07 : 1.907

M-ADD SCOTLA ----- X-12-ARIMA run of SCOTLAND

automean: not significant.

Automatic model chosen : (0 0 0)(0 1 1)

AICtd : rejected

AI Ceaster : rejected
Identifiable seasonality : yes
MD7 : 0.676

M-ADD WALES ----- X-12-ARIMA run of WALES

automean: not significant.

Automatic model chosen : (0 0 0)(0 1 1)

AICtd : rejected
AI Ceaster : rejected
Identifiable seasonality : yes
MD7 : 0.787

M-AUTO UK ----- X-12-ARIMA run of UK

automean: not significant.

Automatic model chosen : (0 1 1)(0 1 1)

AICtd : rejected
AI Ceaster : accepted
Identifiable seasonality : yes
MD7 : 0.296

20 GRAPHS AND X-12-GRAPH

20.1 Introduction

When analysing data one of the most basic but powerful tools is to graph the time series. The X-12-ARIMA program produces several graphs at the end of the printout (Table G) however these charts are poor quality character graphs. If possible, it is preferable to import the output into an external package for high resolution graphics or use the available X-12-Graph interface for SAS.

This chapter is a guide to the graphs to use in a seasonal adjustment analysis and to the functions within X-12-Graph, which are a useful aid in the analysis of a time-series. X-12-Graph can be used in conjunction with the MS-DOS version of X-12-ARIMA. Full details on how to install and use X-12-Graph can be found at www.census.gov/srd/www/x12a

20.2 SAS Graphics

To run the program in graphics mode, users need to use the "-g" flag during the run of X-12-ARIMA and supply the name of an existing directory where X-12-ARIMA will store the graphics files. The full path of the directory needs to be used, remembering that it must be different from the directory where the output (.out) is stored. An example of the command line to run X-12-ARIMA in graphics mode is given below:

```
x12a myspec -g c:\x12a\graphics
```

This command will create graphics output files for the specification file "myspec" in the graphics subdirectory. The seasonal adjustment diagnostics file and the model diagnostics file produced using the "-g" flag store only essential information about the seasonal adjustment and model run needed for the SAS/Graph external graphics procedure.

A description of how to start X-12-Graph in SAS and of the capability of X-12-Graph can be found at www.census.gov/srd/www/x12a

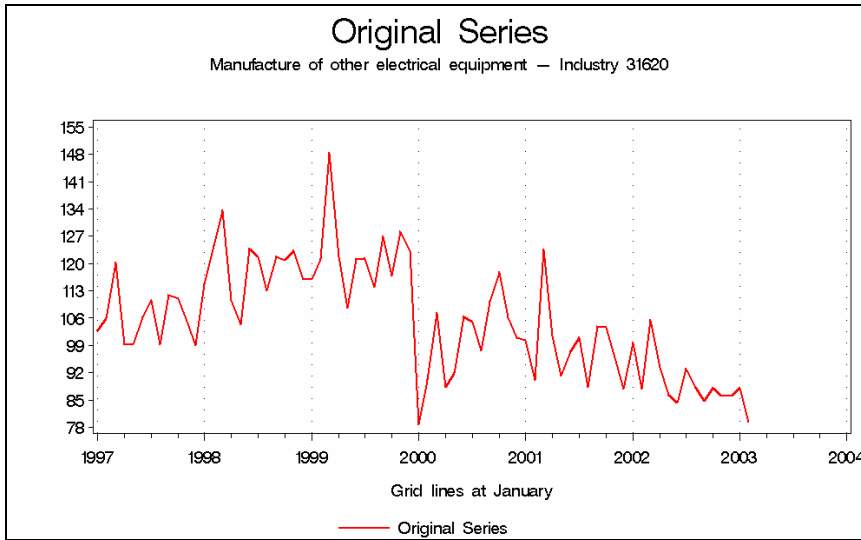
20.3 The raw data

Before seasonally adjusting a series, looking at a graph of the original estimates will enable the user to identify possible problems. For example, look for:

- sudden changes in the general level of the series (trend breaks);
- sudden changes in the months in which peaks and troughs occur (seasonal breaks);
- large extreme values - outliers.



Figure 20-1 Graph of the original series



For example in the above series it can be seen that the series is seasonal and that it has one seasonal peak in March. There is a change in the level of the series between 1999 and 2000, indicating the possibility of a trend break.

20.4 The seasonally adjusted estimates

A graph of the original and seasonally adjusted estimates can show where breaks or outliers may have affected the series.

Figure 20-2 Graph of the seasonally adjusted series




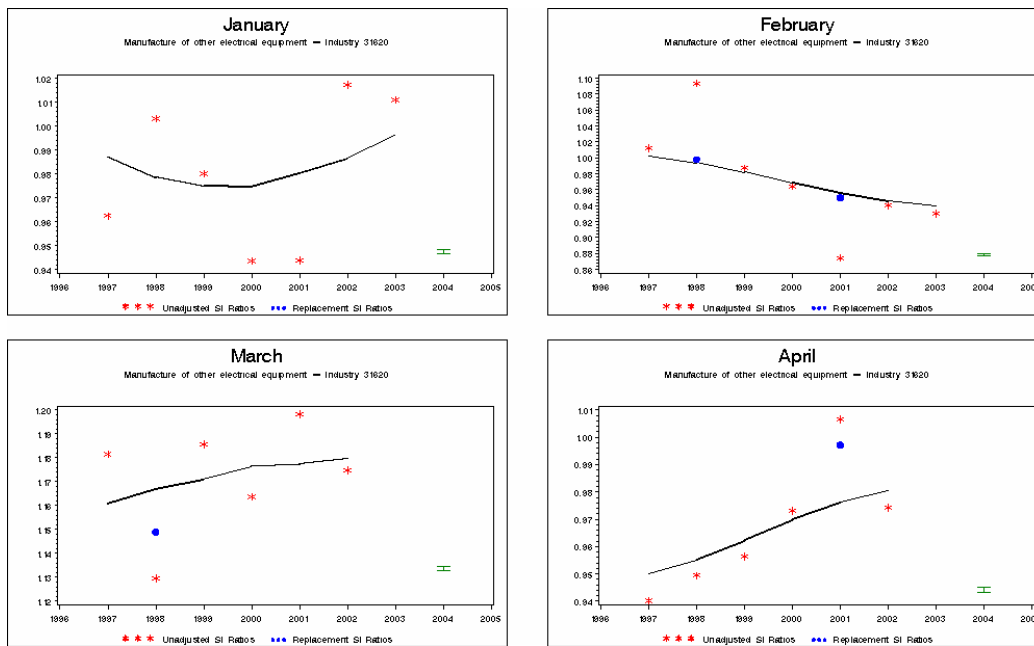
Figure 20-2 shows a steep fall between 1990 and 1991 (a possible trend break), and a possible outlier in February 1998. There is an option within X-12-Graph to highlight any span of data.

20.5 The seasonal times irregular ratios against the seasonal factors

This type of chart plots the Seasonal x Irregular (SI) ratios (de-trended series found in the D8 table of the output) and the seasonal factors (D10) over years where the information has been grouped by different periods. See Figure 20-3 for an example.

A moving average is applied to the SI ratios to obtain the estimate of the seasonal component (which will be removed later by seasonal adjustment). Both of these series are useful for helping decide whether there are any seasonal breaks within the time series or whether there is any need to change the seasonal moving averages for any particular month or quarter. They can also indicate, if there is any one month(quarter) with more statistical variability than the other months(quarter). If the SI ratios are in an approximate straight line, then the seasonal component should follow this e.g. a short moving average is appropriate. If the SI ratios appear to be very erratic, the seasonal factors will try to follow too closely to the SI ratios, producing an erratic seasonal factor line. In this case a long moving average is appropriate. This will remove any unwanted variation in the seasonal without distorting important patterns in the SI ratios.

Figure 20-3 Graph of the SI Ratios and seasonal factors for January to April




X-12-ARIMA replaces some outliers with an automatic process. Where there are lots of replaced SI values in a particular month (denoted by blue points) this could be because the month is particularly volatile. This may indicate evidence of heteroskedasticity or non-constant variance.

SI ratios are a useful guide to the presence of seasonal breaks. These would show up as sudden changes to the level of the SI ratios. A seasonal break in the time series will distort the estimation of the seasonal component. Seasonal breaks may result in 'leakage' in the variation of one component into the variation of another. Where the seasonal factor (D10) is much lower than the SI ratio, some of the seasonality could have been included as part of the irregular, this may result in residual seasonality in the seasonally adjusted series. Where the seasonal factor is much higher than the SI ratio, the seasonal factor could have included too much irregular variation. In this case, some variation may have been removed from the seasonally adjusted series that is not seasonal.

Another useful graph that can be produced in X-12-Graph is the seasonal factors (table D10 in the output) plotted against the mean seasonal factor for each month.

Seasonality is the regular intra-year movement of the series. It is relatively stable and repeats itself every year. Change in seasonality over time can be shown as a gradual movement in the seasonal factors. This is usually not a problem unless there is a pronounced change from below to above the mean, and this is concentrated in one month or quarter. Seasonal moving averages are selected on the basis of the SI ratios, using information from the whole series. Where one month or quarter is not being tracked well by the seasonal factors (it may be more volatile than other months or quarters) the moving average can be tailored to that month or quarter. E.g. in Figure 20-4, a 3x3 moving average can be selected for February when the program originally selected a 3x5 moving average to be applied to all months.

Figure 20-4 Seasonal factors



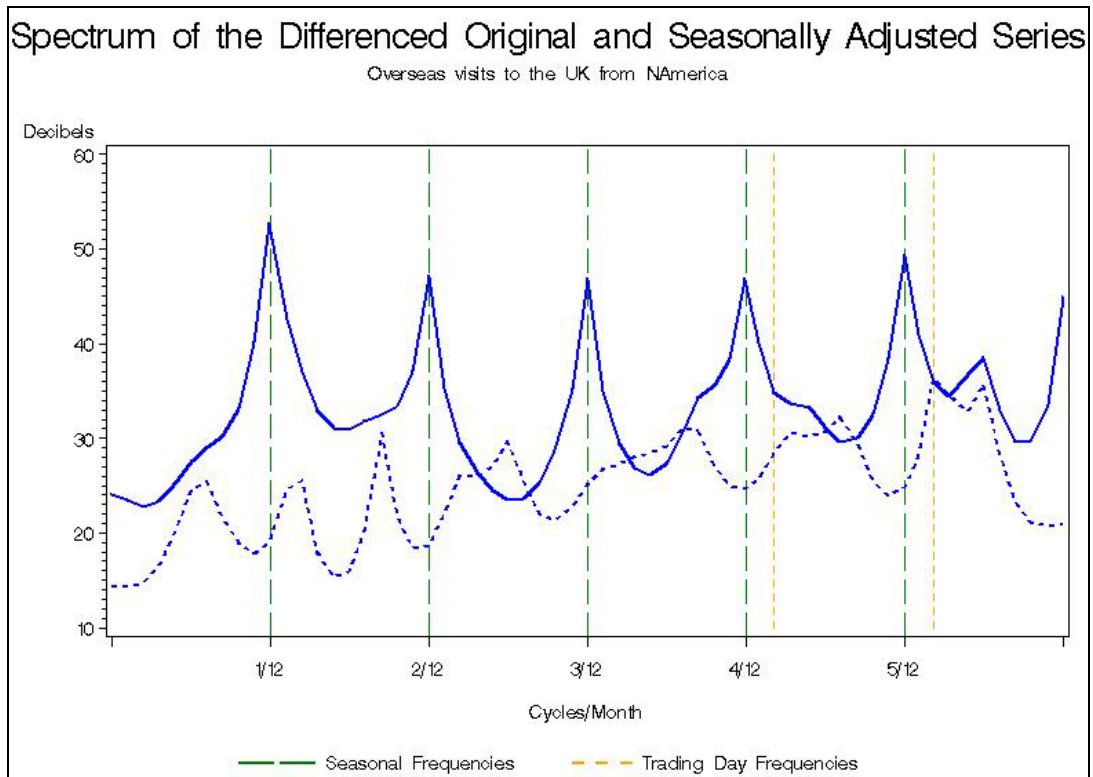
Since all the values are on the same scale it can be seen whether any change in any month is large or small compared with the overall pattern of the series. The overall pattern of the series shown above is one in which the March is the highest month. The most unstable month is February, where the values change substantially through time. The I/S ratios (in the D9A table) can be used in deciding whether distinct moving averages need to be applied to any one month or quarter.

20.6 The Spectrum

The spectrum graph is a useful graphical tool. It can aid the user in establishing whether a series is seasonal or not. Figure 20-5 shows the spectrum of the differenced original and seasonally adjusted series. Spectral peaks, occurring at least at one of the seasonal frequencies provide evidence of seasonality. For example, in Figure 20-5 there are spectral peaks at all six seasonal frequencies which indicates strong seasonality (solid line). The seasonally adjusted series are non-seasonal by virtue that there are no spectral peaks evident (dashed line). Note for quarterly series, there will be two spectral frequencies.

Marginally seasonal series can be identified from the graph as they will have spectral peaks at at least one frequency (most probably at 1/12). However, the user should be aware that in some marginally seasonal cases, the peaks may not be so clearly defined as the ones shown in Figure 20-5.


Figure 20-5 Spectrum Graph



20.7 Other useful graphs

Two other graphical options of X-12-Graph should be used whilst checking the quality of the seasonal adjustment. Those options are the "Component Graphs" and the "Comparison Graphs for Two Adjustments or Two Models".

Figure 20-6 Comparison graph



Related Topics

- Existence of Seasonality

21 SLIDING SPANS

21.1 What is Sliding Spans?

When a series is seasonally adjusted, an important property is stability. A series is defined to be stable if removing or adding data points at either end of the series does not overly affect the seasonal adjustment. If it does, any interpretation of the seasonally adjusted series would be unreliable. The **Sliding Spans** diagnostic is a way of deciding if a seasonal adjustment is stable.

Stability and quality are linked but are not the same. Most of the time, the stability of a seasonal adjustment is a good indicator of its quality. However, it will sometimes be the case that the “best” seasonal adjustment is an unreliable one. In this case, the needs of the user will dictate whether stability or quality of seasonal adjustment is more important.

Definitions: A *span* is a range of data between two dates. A *span length* is the number of data points within that specific range. *Sliding Spans* are a series of 2, 3 or 4 spans that overlap. For example, a span of data could be from January 1990 to December 2000. It has a *span length* of 132 months. An example of a span within this last series is March 1991 to November 1999, which has a span length of 105 months.

21.2 How Sliding Spans work

The Sliding Spans diagnostic works by separately seasonally adjusting each of 2, 3 or 4 overlapping spans of an NSA series. Where two or more sliding spans overlap, the diagnostic compares the different seasonal adjustments. The full process is:

1. The program selects a span length, based on the seasonal moving averages used, the length of the series and whether the data is monthly or quarterly. The span lengths that X12 ARIMA defaults to are:
 - 6 years, if a 3 by 1 seasonal moving average is used
 - 7 years, if a 3 by 3 seasonal moving average is used
 - 8 years, if a 3 by 5 seasonal moving average is used
 - 11 years, if a 3 by 9 seasonal moving average is used

If the data series is not a whole number of years long, the span length increases by the part year length. For example, a series running from January 1990 to March 2001 and using a 3 by 3 seasonal filter would have spans (7 years + 3 months) = 87 months long. Note that in this case,

- The first span would run between January 1991 and March 1998.
- The second span would run between January 1992 and March 1999.
- The third span would run between January 1993 and March 2000.
- The final span would run between January 1994 and March 2001 (the most recent data point).

If different months or quarters have different seasonal filters, X12 ARIMA uses the *longest* seasonal filter to set span length.

2. The program sets up a maximum of 4 spans. The spans start in 1 year intervals. So if the first span starts in Q1 1991, the second span would start in Q1 1992. If there isn't enough data to create 4 spans, X12 ARIMA will produce 2 or 3 sliding spans. In any case, all spans start at the beginning of a year (January or Q1) and the last span includes the most recent data point. If there is not enough data for X12 ARIMA to produce at least 2 spans of the preferred length the diagnostic is suppressed.
3. The program seasonally adjusts each span separately.
4. Where the spans overlap, the program compares the seasonal factors and the month on month and year on year percentage changes in each span for each data point. This is shown in figure 1 below.

Figure 21-1 Illustration of 4 sliding spans, length 7 years

	1st span	2nd span	3rd span	4th span
Q1 1991	100			
Q2 1991	106			
Q3 1991	103			
Q4 1991	110			
Q1 1992	105	105		
Q2 1992	110	110		
Q3 1992	108	108		
Q4 1992	115	115		
Q1 1993	110	110	110	
Q2 1993	114	114	114	
Q3 1993	111	111	111	
Q4 1993	118	118	118	
Q1 1994	107	107	107	107
Q2 1994	116	116	116	116
Q3 1994	114	114	114	114
Q4 1994	121	121	121	121
Q1 1995	115	115	115	115
Q2 1995	124	124	124	124
Q3 1995	120	120	120	120
Q4 1995	125	125	125	125
Q1 1996	120	120	120	120
Q2 1996	128	128	128	128
Q3 1996	123	123	123	123
Q4 1996	129	129	129	129
Q1 1997	124	124	124	124
Q2 1997	133	133	133	133
Q3 1997	127	127	127	127
Q4 1997	135	135	135	135
Q1 1998		128	128	128
Q2 1998		139	139	139
Q3 1998		132	132	132
Q4 1998		138	138	138
Q1 1999			131	131
Q2 1999			143	143
Q3 1999			138	138
Q4 1999			144	144
Q1 2000				134
Q2 2000				149
Q3 2000				138
Q4 2000				150

Note that the Q1 1997 value, 124, is the same in all four spans. However, as each span is seasonally adjusted individually, the seasonally adjusted Q1 1997 values will probably be different. This difference is the basis for sliding spans analysis.

5. For each month (or quarter) that is covered by 2 or more spans, the program works out:

- The percentage difference between the largest and smallest seasonal factor:

$$S_t^{\max} = \frac{(\text{Largest SF}) - (\text{Smallest SF})}{(\text{Smallest SF})}$$

- The difference between the greatest and smallest percentage change in the seasonally adjusted series since the previous month (or quarter).

$$MM_t^{\max} = (\text{Largest Month on Month change}) - (\text{Smallest M on M change})$$

$$QQ_t^{\max} = (\text{Largest Quarter on Quarter change}) - (\text{Smallest Q on Q change})$$

- The difference between the greatest and smallest percentage change since the same month in the previous year.

$$YY_t^{\max} = (\text{Largest Year on Year change}) - (\text{Smallest Y on Y change})$$

In all cases, a value of more than **3%** is regarded as unstable.

6. The program then works out what proportion of data points in the series, where two or more spans overlap, qualify as unstable for each statistic given above. These proportions are called S (seasonal factors), MM (month on month change), QQ (quarter on quarter change) and YY (year on year change).

21.3 When to use Sliding Spans

Sliding Spans can be used when there is a need to know how good a seasonal adjustment has been performed. Sliding Spans analysis is particularly interesting if:

- Seasonal breaks, outliers or fast moving seasonality in the series are suspected. Often, the Q statistic gives good results for series with clear seasonal breaks, but the sliding spans statistics will fail. Looking at the graph of original series and seasonally adjusted series will confirm this.
- Two options for a seasonal adjustment have to be compared, and one wishes to know which would produce the most stable seasonal adjustment estimates. For example, comparing direct and indirect seasonal adjustment. When comparing direct and indirect seasonal adjustments, Sliding Spans statistics are produced for both methods. The series with the lower stability statistics is more stable, and therefore likely to be a better adjustment. However, it is important to make sure the lengths of the sliding spans of the component series are the same. See chapter 6 for more information on direct and indirect seasonal adjustment, and the section “Length and number of spans” in this chapter for more details on the length of the series problem.
- When one or two months a year are unstable, for example a sales series where November, December and January are particularly volatile or when there has been a period of instability in the data. One of the output tables (S3) can be useful for tracing particular months that are susceptible to producing unstable seasonal factors. For example, some series (typically in the USA) are very

sensitive to winter temperatures. The series is still seasonal – but getting a stable seasonal adjustment is difficult.

- For many other different comparisons, such as with or without a trading day effect. See “Sliding Spans Diagnostics for Seasonal and Related Adjustments”, by Findley, Monsell, Shulman and Pugh (1990) for more examples.

In all cases, the question to ask is if stability, or the best use of available information and options, is more important to users.

21.4 How to use Sliding Spans

Most of the time, just using the specification `slidingspans{}` is enough to use the diagnostic. Normally, the `slidingspans` specification is included at the bottom of the spec file. A typical example of part of an X12 ARIMA spec file with a sliding spans analysis is:

```
series{title="Example of sliding spans spec"  
      start=1996.1  
      period=4  
      file="mydata.txt"}
```

```
arima{model=(0, 1, 1) (0, 1, 1)}
```

```
x11{mode=mult}
```

```
slidingspans{}
```

21.4.1 The output

The following table presents a short description and some uses for each of the sliding spans output tables.

Table 21-1 The sliding spans output tables

Table	Description	Uses
S0	Summary of options	
S1	Period means of seasonal factors	
S2	Percentage of periods unstable	Main output table of Sliding Spans
S3	Breakdown of unstable periods	Highlights months and years that are particularly unstable
S7	Full Sliding Spans analysis	Gives an idea of the distribution of unstable months. Also contains SA spans for analysis.

Tables S2, S3 and S7 are produced several times, labelled a) to e):

- Table a) represents seasonal factors.
- Table b) covers trading days. Only prints out if **fixmdl** is set to **no** or **clear**.
- Table c) is for the final seasonally adjusted series. This only prints out if one or more regression variables, such as trading days or an Easter effect, are included in the model and are not fixed.
- Table d) is for the month to month or quarter to quarter changes.
- Table e) is for year on year changes.

In table S2 guideline percentages are only produced for seasonal factors, month on month (or quarter on quarter) and year on year changes. Looking at the values of, in particular, S and MM (or QQ if quarterly data) of table S2 gives an idea as to how good the adjustment is. The guidelines given in Findley, Monsell, Shulman and Pugh (1990) for interpreting S, MM (or QQ) and YY are:

- For S, series with stable seasonal adjustments usually have $S < 15\%$. Series with $S > 25\%$ almost never have good seasonal adjustments. Those for which S falls between 15% and 25% should cause concern, but may give acceptable seasonal adjustments.
- Series for which $MM > 40\%$ ($QQ > 40\%$ for quarterly data) should not be seasonally adjusted.
- $YY < 2\%$ is common with good series. While there is no recommended upper limit for YY above which a series should not be seasonally adjusted, $YY=10\%$ is quite high. However, most series with high YY values also have poor results for S or MM, so the YY statistic is of less importance.

These guidelines are summarised under table S2 in the X12 ARIMA output, when Sliding Spans is run.

The output of tables S3 can be particularly useful in tracking down areas of instability. It is not automatically printed with the **print = brief** option, so the default or **print = all** should be used when a single series is being analysed in more detail.

21.4.2 Sliding spans options

It is not necessary to use any of the available options to use Sliding Spans. However, some of the options (such as **length**, **print** and **outlier**) can allow the user to be able to apply the sliding spans diagnostic in the widest range of situations.

21.4.2.1 Length and number of spans

The span length is of prime importance, since it can affect the performance of the sliding spans diagnostics, especially in case that less than 4 spans are created.

The **length** argument allows the user to select the number of data points included in the spans and used to generate output for comparisons. The **length** command is needed:

- To bring all the span lengths into line with the *longest* span used.
- To run an indirect seasonal adjustment if different seasonal moving averages are used for the component series. In this case, the span lengths for each series will be incompatible. For example, a composite series is made up of monthly male and female seasonally adjusted series. The adjustment for men uses a 3 by 5 seasonal moving average, while the adjustment for women uses a 3 by 3

seasonal filter. The two series will use different span lengths, so we need to change the span length. We would use the command **length = 96** in the spec file for adjusting the series for women.

In these situations the guide percentages given by Findley, Monsell, Shulman and Pugh become too high and cannot be reliably used. However, the results can still be used to say something useful about the adjustment.

When the `slidingspans` spec is included in the analysis of a series, users have to take into account that:

- Longer span lengths lead to a *smaller* proportion of unstable months. Therefore, using longer spans will, if anything, bias the stability statistics *downwards*. Therefore, any Sliding Spans statistics which are *higher* than the guideline percentages are still more unstable than is desirable.
- When two or three spans are produced, the program will say that the guideline percentages are now too high and should be lowered. If the values of S and MM (or QQ) are too high when less than four spans are produced, the series is still unstable.

In both these cases, there is a chance that a series may be not be classified as unstable, when in fact it is. This is because the series would have been close to being classified as unstable under the guideline percentages, and either the length or number of spans has been changed. As it is not known how much to change the guideline percentages by, it cannot be said for sure whether the series is stable or unstable. However, this is a small price to pay for gaining the use of Sliding Spans in a large number of cases where otherwise it could not possible.

21.4.2.2 *Other important options in the sliding spans analysis*

There are other two important options that can be useful in the series analysis:

- The argument **print = all** allows to view the full Sliding Spans output, which can be a useful additional source of information when analysing a series in more detail. However, the default output is usually enough.
- If there are a lot of outliers in the data, it is recommended to use the command **outlier = keep** to allow the Sliding Spans analysis to use them. This may improve the stability statistics (see below), and provide further evidence for including them in the adjustment.

21.4.3 *When Sliding Spans will not work*

There are several situations in which Sliding Spans will either not work, or produce output that is not useful in assessing the stability of the data. They are:

1. When an additive seasonal adjustment is selected. Sliding Spans will not produce tables S1 and S2 in this case, and the remaining output should not be used in the same way as for multiplicative adjustments.
2. If the series is too short for X12 ARIMA to construct at least 2 spans, X12 ARIMA will not carry out the Sliding Spans analysis, and display an error message in MS-DOS. However, the rest of the seasonal adjustment will be carried out.
3. If X12 ARIMA produces fewer than four spans, it will recommend that the threshold values of S, MM (or QQ) and YY be lowered. See the section "Length and number of Sliding Spans" for more details.

4. If all the seasonal factors are close to 100. In this case, the threshold value of 3% is too large, and no S2 table will be produced. This is quite common, and unfortunately renders the Sliding Spans diagnostic useless when it happens. Findley & Monsell have carried out a study as to how the threshold values should be lowered, but without success (see P148 of the X12 ARIMA manual).
5. If adjusting a composite series indirectly, and the span lengths of the component series are different, a period in each span will have unrealistic seasonal factors. In this case, using the **length** argument as detailed earlier to even up the span lengths will produce more accurate seasonal factors. However, care must still be taken when interpreting the stability statistics – see “Length and number of Spans” for more details.

21.5 Summary

The Sliding Spans diagnostic is one of the more useful statistics in the X12 ARIMA output to look at in the “first wave”. The Sliding Spans diagnostic measures the *stability* of a seasonal adjustment, which is related to quality but not necessarily the same thing.

A series that fails Sliding Spans is likely to have a seasonal break or a period of instability. Sliding Spans is particularly useful in these situations, as the M and Q statistics can misleadingly suggest the seasonal adjustment is good. Sliding Spans can also be used to compare different methods of seasonal adjustment, and give an indication of which option leads to a more stable adjusted series.

References

“Sliding Spans Diagnostics for Seasonal and Related Adjustments.” Findley, Monsell, Shulman and Pugh (1990): *Journal of the American Statistical Association* **85**, 345–355.

22 HISTORY DIAGNOSTICS

22.1 What is Revision History?

'Revision histories' is one of the stability diagnostics available in X12ARIMA. It considers the revision of continuous seasonal adjustment over a period of years, and is therefore a way of seeing how a time series is affected when new data points are introduced. The basic revision calculated by X12ARIMA, to enable visualisation of these effects, is the 'difference between the earliest adjustment of a month's datum' - obtained when that month is the latest month in the series - and 'a later adjustment based on all future data available at the time of the diagnostic analysis'.

When a new data point becomes available for a series, more is learnt about the behaviour of that series. This is especially true of seasonal series. With that new data point, more is found out about the seasonal pattern, as well as the underlying trend. However, when the series is updated with the extra information, using the extra information creates revisions to the seasonally adjusted series. Generally, users want to minimise the number and size of these revisions. How important this is depends on the user. Some users prefer to have the most up-to-date figures possible, regardless of the size of revisions. Other users prefer limited revisions, to make the data they work with more consistent. The choice of revision policy is also influenced by the nature of the data.

The revision history diagnostic is a way of seeing how a time series is affected when new data points are introduced. Obviously, it is not possible to predict the future. Instead, the diagnostic works by taking a period of existing data, and adding the data points one by one as if they were new observations. This creates revisions in the data, and it is these revisions that users are interested in.

22.2 How the Revision History Diagnostic works

The process by which the diagnostic works is:

1. The start date of the history diagnostic is chosen. This can either be the default X12 ARIMA choice, or specified by the user with the **start** argument.
2. The program seasonally adjusts the series, up to and including the start point of the revision history.
3. X12 ARIMA then seasonally adjusts the whole series again, but up to the observation after the revision history start date. The program repeats this process, including one extra data point each run, until all the data has been added and the entire series has been adjusted.
4. The final seasonal adjustment, which covers all of the data, is the most recent adjustment available, and represents the best estimate of the seasonal factors.
5. The program works out the percentage difference between the first seasonal adjustment of the starting data point (worked out in step 2) and the final adjustment for the same month (as in step 4). The program repeats this for every data point up to, but not including, the final data point in the series. The program also calculates any other histories specified by the user.
6. The program then produces the appropriate tables for the series.

22.3 When to use Revision History

The diagnostic is most often used when there are two competing methods for seasonally adjusting a series, both of which are acceptable in terms of other diagnostics, e.g., the 'Q' statistics.

For example,

- Use of two different ARIMA models;
- Use of different seasonal moving averages;
- Choosing between direct and indirect seasonal adjustment of a composite series.

To compare how prone two different adjustments are to revisions, run the history diagnostic on both series. Then compare the revisions produced over the test period. In general, smaller is better. An adjustment with small average annual revisions will be more reliable, than an adjustment with larger revisions. Unlike some other diagnostics, such as the Q-statistic and sliding spans, there is no absolute measure of what are acceptable revisions.

There are eight different variables that X12 ARIMA can produce a history for, and many of these tables can be used in more than one way. Any full explanation of the diagnostics and options available would be long and unwieldy. Therefore, this paper will concentrate on three functions of the history diagnostic.

1. Using the revision history to compare two competing adjustments
2. Comparing revisions in direct and indirect seasonal adjustments
3. Using AIC histories to choose between two adjustments

Therefore, the revision history diagnostic is one of the later tests to apply to a series.

22.4 How to use Revision History

The history specification can be used with different sets of arguments and the choice of the arguments depends on the scope of the revision analysis. Normally, the **history** spec is included at the bottom of the spec file. An example of part of an X12ARIMA spec file with a revision history analysis is:

```
series{title="Example of sliding spans spec"  
      start=1996.1  
      period=4  
      file="mydata.txt"}
```

```
arima{model=(0, 1, 1) (0, 1, 1)}
```

```
x11{mode=mult}
```

```
history{start=1999.1  
        sadjlags=(1, 2)}
```

22.4.1 Comparing two competing adjustments

This is the simplest use of the revision history diagnostic. It is used to produce just table R1, which gives the revision between the first and final estimates of the seasonally adjusted series. The idea of this diagnostic is to compare the stability of two competing adjustments when new data is introduced.

The **history** specification is typically near the end of a spec file. The bottom of a spec file in which a history analysis is being run might look like this:

```
x11{
  mode=mult
  seasonal ma=s3x3
  trendma=5
}

history{
  start=1999. 1
  sadj lags=(1, 2) }
```

sliding spans{ }

In this example the program is being run on the series NZJW, which is from the Motor Vehicle Production Index data set. In this spec file, two options have been included within the history specification. The **start** argument determines the start date of the series. Users are most interested in how the seasonal adjustment performs with current data, rather than data that is several years old. It is recommended that the **start** argument be used to limit the length of the history diagnostic to the last couple of years. If the **start** argument is not used, X12 ARIMA will select the start date – see the technical summary to see how the program does this.

The other argument used in the **history** specification above is **sadjlags**. This argument allows the user to choose up to 5 lags to be covered in the revisions history. See below for details as to what this means.

Here is table R1 from the output obtained when this spec file was run on the series NZJW.

R 1 Percent revisions of the concurrent seasonal adjustments

From	1999. 1 to 2000. 2		
Observations	6		
1999			
1st	0. 89	- 2. 08	- 2. 27
2nd	- 9. 11	- 2. 88	- 1. 77
3rd	4. 00	3. 13	2. 61
4th	1. 08	- 0. 25	- 0. 52
2000			
1st	- 1. 00	- 1. 03	*****
2nd	- 2. 55	*****	*****

What do the figures in table R1 mean? In this case, the data is quarterly and the final adjustment (with all the data up to Q3 2000 available) is the target. The first column is the revision history of the initial adjustment. It is the percentage difference between the *first* vintage (including all the data up to the point in question in the adjustment) and the *final* vintage (including all the data in the series).

For example, the value 1.08 (shown in bold) has been calculated by working out the percentage difference between the seasonally adjusted value for Q4 1999 given *all* the data, and the value for Q4 1999 given the data *up to Q4 1999*. In other words, it is the percentage difference between the Q3 2000 vintage (the final vintage) and the Q4 1999 vintage (the first vintage) of Q4 1999.

The next column is the revision at lag 1 – this is stated in the output in R0. In this case, the value for Q4 1999 is -0.25. This has been calculated by working out the percentage difference between the seasonally adjusted value for Q4 1999 given *all* the data, and the value for Q4 1999 given the data *up to Q1 2000*.

This is the percentage difference between the *final* vintage and the Q1 2000 vintage (the second vintage) of Q4 1999.

The final column is the revisions at lag 2. This is the percentage difference between the *final* vintage and the *third* vintage. In general, for a column representing the revision at lag n , the figure quoted is the percentage difference between the *final* vintage of the data and the $(1+n)$ vintage.

Why are the revisions at various lags of interest? They give some idea as to the size of revisions at each lag, which can say a lot about the data. For example, the relative sizes of the revisions at each lag give you some information as to how quickly adjusted data converges to its final values. This information can be used to inform decisions on a revision policy for the series. Up to five lags at a time can be specified, using the **sadjlags** argument.

The diagnostic also prints out summary tables, shown below.

R 1.5 Summary statistics : average absolute percent revisions of the seasonal adjustments

Quarters:			
1st	0.95	1.55	2.27
2nd	5.83	2.88	1.77
3rd	4.00	3.13	2.61
4th	1.08	0.25	0.52
Years:			
1999	3.77	2.08	1.79
2000	1.78	1.03	*****
Total:	3.11	1.87	1.79
Hinge Values:			
Min	0.89	0.25	0.52
25%	1.00	1.03	1.14
Med	1.82	2.08	2.02
75%	4.00	2.88	2.44
Max	9.11	3.13	2.61

Of these statistics, the most useful value is the "Total" revision. This is the average absolute percentage revision from the first (or second, third... etc, depending on the **sadjlags** option used) seasonal adjustment, and the adjustment using all the data available. This is a single figure that can be used to compare the quality of competing adjustments.

The other values can be useful too. Data users will be more interested in recent data. Therefore, the size of revisions with recent data may be more important in deciding which adjustment to use than earlier years.

So, how do users actually use the diagnostic? An example would be to check the revisions caused by different ARIMA models of the series NZJW. To do this, run the series repeatedly through X12 ARIMA, once for each ARIMA model you want to compare. No other options should be changed. The results were:

Table 22-1 Revision History

Model	Average absolute revision percentage		
	First estimate	At Lag 1	At Lag 2
<i>LOG (0 1 1)(0 1 1)</i>	3.11	1.87	1.79
<i>LOG (0 1 2)(0 1 1)</i>	3.11	1.92	1.90
<i>LOG (2 1 0)(0 1 1)</i>	3.17	1.91	1.89
<i>LOG (0 2 2)(0 1 1)</i>	3.11	1.86	1.79
<i>LOG (2 1 2)(0 1 1)</i>	3.20	1.83	1.81

The table shows that changing the ARIMA model used does not have a very large effect on the size of revisions. The first and fourth model appear to be best, and in practice a (0 1 1)(0 1 1) model would probably be selected as it uses the fewest parameters.

Other options

Tables R2, R4 and R5 are all similar to table R1, but refer to slightly different target variables. The table below summarises the target variables available, and the commands needed to produce each variable.

Table 22-2 Other options

Number	Command	Details
R1	Sadj	Final seasonally adjusted series (the default)
R2	Sadjchng	Period to period changes in final seasonally adjusted series
R3	Trend	Final Henderson trend component
R4	Trendchng	Period to period changes in Henderson trend component

The variable has to be included in the **estimate** argument. So, if you were interested in looking at lags 1, 2, 3 and 12 in both the level and month to month changes in the final trend component, the history part of the spec file would be:

```
history{
  start=1998. 1
  estimates=(trend trendchng)
  trendlags=(1, 2, 3, 12)
}
```

Full explanations of what the variables are can be found in the technical summary. More complicated growth rates, such as the change of most recent 3 months on previous 3 months (for example, Apr-May-Jun compared with Jan-Feb-Mar) need a different approach. See the section on revision triangles for details on how to construct revision histories manually for this purpose.

22.4.2 Comparing direct and indirect seasonal adjustment

A particular comparison it is wanted to run is between direct and indirect seasonal adjustment.

To get table R3, you need to be carrying out a composite adjustment. This table then gives the revisions obtained by adjusting all the components individually, in other words, indirect seasonal adjustment. In the same output, Table R1 gives the revisions to the direct seasonally adjusted series. Other than setting up the composite adjustment (covered in chapter 6) and including the **revision** specification in the individual component series, no other changes to the spec files are needed.

As an example, here are the summary tables produced when the series GMAC, GMAD and GMAE, overseas visitors to the UK (from different locations), are run both directly and indirectly:

R 1.S Summary statistics : average absolute percent revisions of the seasonal adjustments

Years:

1999	0.73
2000	0.50

Total: 0.65

R 3.S Summary statistics : average absolute percent revisions of the concurrent indirect seasonal adjustments

Years:

1999	1.06
2000	0.67

Total: 0.92

The tables above show that since 1999, the indirect adjustment has produced greater revisions than the direct adjustment. This might not be a problem, and the indirect adjustment may have other virtues, for example, it is more stable under sliding spans. While the direct adjustment is better than the indirect in terms of the size of revisions it suffers, the table shows how big the difference is and allows it to be set against any perceived advantages of indirect adjustment.

22.4.3 Using AIC History to choose between two adjustments

Table R7 is a history of the model's AIC values over the period of the history. Unlike the other tables talked about so far, this is not a revision history. It is a history of the AIC values calculated when fitting the model. **Lower AIC values indicate a better model fit.** The AIC history is a useful and flexible way to choose between two different models for seasonal adjustment.

R 7. Likelihood statistics from estimating regARIMA model over spans with ending dates 1: 1996 to 3: 2000

Span End	Log Likelihood	AICC
-----	-----	----
1999. 1	- 337. 006	676. 100
1999. 2	- 345. 626	693. 338
1999. 3	- 358. 060	718. 205
1999. 4	- 364. 566	731. 214
2000. 1	- 371. 809	745. 700
2000. 2	- 378. 743	759. 565
2000. 3	- 386. 011	774. 101

This is an example AIC history for the example series used above. The method is to produce two AIC histories, for two competing methods of seasonal adjustment. Then, calculate for each point in the history, the difference between the first and second model's AICC value, including any corrections you may need to apply. For example, when choosing between an additive or multiplicative model decomposition, X12 ARIMA penalises the additive model by adding 2 to its AICC value, before comparing it against the multiplicative adjustment's AICC value.

As stated earlier, a small AICC value is more desirable. Therefore, a *negative* difference in AICC values means that the *first* model is better, and a *positive* difference favours the *second* model.

The following example uses a monthly series showing the number of men claiming Job Seekers' Allowance, as of the second Thursday in the month. There is an issue as to whether or not this series is best modelled with an additive or a multiplicative model. Therefore, looking at the AIC histories of both models might be useful. The **history** specification of the spec file needed to run this diagnostic is:

```
history{
  start=1999. 1
  estimates=aic
  fixreg=holiday
  save=lkh}
```

In this spec file, there are three options not previously mentioned. The first is the option **estimates = aic**. This command produces the AIC history. The option **save = lkh** means that the AIC history will be saved as a separate file, *.lkh.

The option **fixreg = holiday** needs a little more explaining. When branches carry out seasonal adjustment during the year, they use fixed prior adjustments derived by TSAB to account for Easter effects, trading day effects and so on. When analysing series in X12 ARIMA, the **regression** specification is often used for these effects.

To be consistent with the way the data is actually treated during the year, use **fixreg = holiday** for Easter, **fixreg = td** for trading days and **fixreg = user** for other variables not assigned a type by the **usertype** argument in the **regression** spec. It is not necessary to use **fixreg** when prior adjustments, rather than the **regression** specification, are used to build in effects like Easter and trading days.

Table 22-3 Table of selected values from R7

Span end date	AICC (mult. Model)	AICC (add. model)	Difference	Model choice
Jan 1999	2174	2166	8	Additive
Apr 1999	2241	2230	11	Additive
Jul 1999	2307	2292	15	Additive
Oct 1999	2372	2354	18	Additive
Jan 2000	2438	2421	17	Additive
Apr 2000	2504	2488	16	Additive
Jul 2000	2569	2549	20	Additive
Oct 2000	2635	2613	22	Additive
Nov 2000	2657	2633	24	Additive
Dec 2000	2678	2654	24	Additive

In each case, the additive model is selected if $AICC(\text{multiplicative}) - AICC(\text{additive}) > 2$.

The table above shows that the choice of an additive model is stable. Throughout the last two years, X12 ARIMA would have chosen an additive model in any month, even though the additive model is penalised in the model selection process.

A further example can be found in section 4.4 of "New Capabilities of X12 ARIMA" by Findley et al., (1998).

22.5 When Revision History will not work

The Revision History diagnostic needs a minimum of 5 years of data to work. This is because it needs 5 years of reference data to make a reasonable series of seasonal adjustments during the revision history.

If there are less than 5 years of data, X12ARIMA will fail to run. If the **start** argument is used to specify a date that is less than 5 years from the beginning of the series, X12 ARIMA will run. However, it will move the start date of the revision history to 5 years after the series start date.

In some cases, the revision history of an additive seasonal adjustment can be difficult to interpret. If an additive series has just positive or just negative values, the revisions in table R1 are calculated in just the same way as for a multiplicative seasonal adjustment. However, if a series has both positive and negative values, then the revision, rather than the percentage revision, is calculated. For example,

R 1 Revisions of the concurrent seasonal adjustments

From 2000. Jan to 2001. May
Observations 17

2001

Jan	- 53. 96
Feb	33. 89
Mar	13. 27
Apr	22. 06
May	25. 18

Most other variables are not affected by this effect. Details of which other variables are affected and how are given in the technical summary.

22.6 Summary

The revision history diagnostic is most useful when comparing two methods for seasonal adjustment that are both usable in terms of M and Q statistics, sliding spans and so on. There is no absolute measure of what is an acceptable level of revisions, so the diagnostic is of limited use on a single series.

There are several histories available, all of which can be useful in observing the performance of the seasonal adjustment. Of these, those in tables R1 and R3 are probably most useful. If month-to-month changes, the trend component, or the month-to-month trend changes are of interest, then tables R2, R4 and R5 respectively will be worth looking at.

Table R7 is quite easy to use to compare two competing models for a series, such as two ARIMA models or two different prior adjustments. Table R8 is for the more advanced user, but can be used to compare almost anything, as good forecasts are important in getting a good initial seasonal adjustment of a new data point.

References

“New Capabilities and Methods of the X12ARIMA Seasonal Adjustment Program.” Findley, D., Monsell, B.C. et al (1998): *Journal of Business and Economic Statistics* **16**, 127–77.

23 COMPOSITE SPEC

23.1 Introduction

The composite spec is used as part of the procedure for obtaining both direct and indirect adjustment of a composite series. This chapter describes the spec files of the composite series and of the components and how to run the seasonal adjustment.

Section 2 will describe how the **composite** spec can be used in the X12ARIMA to compare the performance of direct and indirect seasonal adjustments. Section 3 will discuss the diagnostics that are produced and how these can be used to inform the best method of seasonally adjusting an aggregate series. Section 4 outlines other considerations that should be taken into account when choosing the level of seasonal adjustment and lists some topics related to aggregate series.

23.2 The spec files

In order to run a composite seasonal adjustment (i.e. direct and indirect) you require spec files for the component series (one spec for each series), a spec file for your aggregate series, and a metafile that contains the names of all the component series that form the aggregate series and the name of the aggregate series itself.

The individual spec files must define how they are combined to form the aggregate, using the **comptype** and the **compwt** arguments in the **series** spec. The **comptype** argument specifies whether that particular component is added, subtracted, multiplied or divided, whereas the **compwt** argument specifies the size of the constant used to multiply that particular component before it is combined to form the aggregate series (i.e. the weight of the component series).

A composite adjustment run produces an indirect seasonal adjustment of the composite series as well as a direct seasonal adjustment. The indirect seasonal adjustment is the combination specified by the **comptype** of the components. If one of the component series is not seasonal then specifying the summary measures option by setting **type=summary** in the **x11** spec of that component will include the component in the indirect adjustment without seasonally adjusting that particular series.

A sliding span or revision history analysis (see chapters 21 and 22) of the direct and indirect adjustments can be obtained but the options must be specified in each of the component series as well as the composite series.

When a composite series is adjusted using X-12-ARIMA the outputs are:

- A direct seasonal adjustment of the aggregate, with normal output.
- An indirect adjustment of the aggregate found by combining the seasonal adjustment of the components, with an output that includes the D-tables onwards only.

23.2.1 EXAMPLE

The following example illustrates all the steps of a composite adjustment

The series with file names "enj.q.txt" and "enjl.txt" are two component series that sum to an aggregate series called "trade".

Step 1 Create a spec file for each of the component series, e.g.:

ENJQ. SPC

```
series{title="Food & Beverages Imports CP"
```

```
  file="enj q. txt"
```

```
  name="enj q"
```

```
  comptype=add}
```

```
x11{}
```

A spec file for a component series that is not seasonally adjusted is given below

ENJL. SPC

```
series{title="Crude Oil Imports CP"
```

```
  file="enj l. txt"
```

```
  name="enj l "
```

```
  comptype=add}
```

```
x11{type=summary}
```

This will add the unadjusted series to the composite seasonally adjusted series

Step 2 Create a spec file for the composite series

TRADE. SPC

```
Composite{title="Total Imports CP"
```

```
  name="Trade"}
```

```
x11{}
```

Step 3 Create a **metafile** with details of the component and composite series. The metafile saved as **imports.mta** is shown below:

```
Enjq
```

```
Enjl
```

```
Trade
```

NB the spec file for the composite series must be listed last.

Step 4 To run X12ARIMA in DOS, enter

```
x12a -m imports -w -p
```

and press Enter.

To run X12ARIMA in OxEdit,

1. Have the metafile open and active
2. Select from the Modules menu to run metafile with appropriate module (4: metafile, 5:metafile other flags or 6:metafile graphics other flags)

The resulting output (of the aggregate series) will provide diagnostics from which you can compare the performance of the direct and indirect seasonal adjustment of an aggregate series.

23.3 Comparing adjustments using X12ARIMA diagnostics

If there are no prior reasons for which level of seasonal adjustment should be used then the composite adjustment spec provides a powerful tool for deciding what is the optimal level of seasonal adjustment. The following are some criteria to help choose between the direct and indirect approach and may be put in order of priority as follows:

23.3.1 Residual seasonality in the seasonally adjusted series

The estimated seasonally adjusted series should not have any residual seasonality. X-12-ARIMA provides a set of spectral diagnostics what could be used as a quick check for the existence of significant seasonal peaks in one of the estimated spectra. Note that if a warning message for significant seasonal peaks is printed out while running the composite spec file, then it refers to the direct seasonal adjustment of the composite series. In order to check whether there are significant peaks in the indirect seasonal adjusted series, the spectral plots in the output should be examined. To be visually significant, the spectral amplitude at a seasonal frequency must exceed both of its neighbours by at least "6-stars". If there are any significant peaks in one of the two options then the alternative approach should be chosen.

23.3.2 Revision errors

X-12-ARIMA produces a set of revision history diagnostics (see chapter 22). In general, the preferred alternative is the one that minimises the average percentage of revisions in the seasonally adjusted series.

23.3.3 Stability

X-12-ARIMA produces a set of sliding span diagnostics (see chapter 21) as a way of deciding if a seasonal adjustment is stable. In general, the preferred alternative is the one producing a more stable seasonally adjusted series.

23.3.4 Interpretability of seasonally adjusted series

The M-statistics measure various areas in the quality of seasonal adjustment while the Q-statistic is a weighted average of all M-statistics. If the Q statistic fails then there might be problems interpreting short-term movements in the seasonally adjusted series. In general the chosen approach should minimise the Q-statistic. Note however that the Census Bureau have found evidence that, in an indirect analysis, some of the statistics M8 to M11 are misleadingly high. If it appears that the indirect Q statistic is unduly high mainly because of these statistics, it may be justifiable to give less weight to this fact.

23.3.5 Smoothness of seasonally adjusted series

The choice between direct and indirect methods is based on the comparison between the roughness measures, computed for the two series derived under the two different approaches. In general, the method of adjustment which gives the smoother series should be used. The measure of roughness is given at the end of the output (for the aggregate series), or before the sliding spans and history diagnostics if these have been activated. The following is an example

MEASURES OF ROUGHNESS R1 AND R2 FOR SEASONALLY ADJUSTED SERIES

	DIRECT		INDIRECT		PERCENTAGE CHANGE	
	FULL SERIES	LAST THREE YEARS	FULL SERIES	LAST THREE YEARS	FULL SERIES	LAST THREE YEARS
R1-MEAN SQUARE ERROR	1.453	2.692	1.663	3.366	-14.466%	-25.042%
R1-ROOT MEAN SQUARE ERROR	1.205	1.641	1.290	1.835	-6.989%	-11.822%
R2-MEAN SQUARE ERROR	0.000	0.000	0.000	0.000	-25.768%	-34.109%
R2-ROOT MEAN SQUARE ERROR	0.008	0.011	0.008	0.013	-12.146%	-15.805%

POSITIVE PERCENTAGE CHANGES INDICATE THAT THE INDIRECT SEASONALLY ADJUSTED COMPOSITE IS SMOOTHER THAN THE DIRECT SEASONALLY ADJUSTED COMPOSITE.

The measures of roughness describe the size of the deviations from a smooth trend of the adjusted series (with R1 and R2 using different methods of trend estimation. The test results reported above suggest that the direct seasonal adjustment is smoother than the indirect adjustment. However it should be noted that smoothness is not necessarily the most desirable characteristic of a seasonally adjusted series.

23.4 Problem associated with the composite spec

A partial year at the beginning or end of the series creates problems in the constraining process. When a composite spec is run to seasonally adjust the aggregate of a group of series and the "force=totals" command is used in the x11 spec, X12ARIMA does not calculate the indirect seasonal adjustment for the partial year data. For this reason the measures of roughness will not be calculated

25 FORECASTING

25.1 Introduction

Forecasting techniques are used extensively in National Accounts to extend series beyond the most recent period for which data are available. This document describes the options to use for forecasting in seasonal adjustment and how to use X12ARIMA to forecast series for a purpose different from seasonal adjustment. This chapter is intended to be a reference document to allow users to make sound decisions without needing to understand the forecasting procedure in much depth, although a more technical section, which explains the regARIMA method in more detail, can be found in Chapter 11.

25.2 Forecasting in Seasonal Adjustment

The regARIMA part of the X12ARIMA program enables time series to be modelled to extend the series forwards (adding forecasts) and backwards (adding backcasts), to estimate calendar effects and to enable the series to be adjusted for unusual and disruptive features such as outliers or breaks.

Chapter 11 described the regARIMA method used in X12ARIMA and explained the options available to fit a model before the seasonal adjustment itself is performed. This section will go back to the regARIMA and will give details on the method to use to forecast series for seasonal adjustment purposes. Section 3 will deal with the problem of model selection for forecasting series for a purpose different from seasonal adjustment.

A question one might wish to ask is, under what circumstances should each of the methods described in Chapter 11 be used for extending a series before seasonally adjusting it? In particular, should **automdl** or **pickmdl** be chosen for automatic model identification? A related question is, what forecasting horizon should be chosen for the forecasts and backcasts whose generation is one of the main purposes of modelling? The apparent assumption in X12ARIMA is that **automdl** is the norm (in using **pickmdl** we are reverting to the method of the former X11ARIMA). The default horizon for forecasts, if none is specified, is one year (the default for backcasts is none). Thus the 'standard' approach envisaged by the designers of X12ARIMA is to use the model selected by **automdl** with a year of forecasts.

Extensive empirical analysis of ONS data has been carried out to verify this view. It is found that performances of **pickmdl** and **automdl** are not very dissimilar. If a uniform approach is desired, use **pickmdl** and one year forecast horizon. In fact, using **pickmdl** with "*method=firsl*", with the extended list of model (reported in chapter 11, section 4.2) and a year of forecasts will result in a substantial reduction of revisions compared with no forecasts at all – typically the mean absolute revision is reduced by 10-15%. If the recommended method fails to select a model, the appropriate alternative forecasting method should be adopted, as highlighted in Table 25-1.

Table 25-1 Appropriate modelling method for forecasting

	Recommended method		Alternative method	
Series length	Monthly series	Quarterly series	Monthly series	Quarterly series
<6 years	▪ Pickmdl (first, extended list);	▪ pickmdl (first, extended list);	▪ automdl (maxlag=(2,1), no constant);	▪ automdl (maxlag=(2,1), no constant);

	Recommended method		Alternative method	
Series length	Monthly series	Quarterly series	Monthly series	Quarterly series
6 – 10 years	▪ pickmdl (first, extended list);	▪ Pickmdl (first, extended list);	▪ automdl (maxlag=(3,1), no constant);	▪ automdl (maxlag=(3,1), no constant);
10 - 12 years	▪ pickmdl (first, extended list);	▪ pickmdl (first, extended list);	▪ automdl (maxlag=(4,1), no constant);	▪ automdl (maxlag=(3,1), no constant);
>12 years	▪ pickmdl (first, extended list).	▪ pickmdl (first, extended list).	▪ automdl (maxlag=(4,2), constant).	▪ automdl (maxlag=(4,2), no constant).

The first two columns show the recommended method for forecasting monthly and quarterly series. If the method recommended fails to select a model, then the alternative method should be used.

For monthly series with more than 12 years of data available the alternative method allows to test the significance of a constant term in the model. Particular attention should be paid when interpreting the results. In fact, a significant constant should be maintained only if the selected model does not present regular differencing, i.e. the model picked is of the type $(p,0,q)(P,D,Q)$. This is because when regular differencing is selected in the model, the constant term corresponds to an exogenous trend, which is not so common in real life series. On the contrary without regular differencing in the model the constant term would refer just to the mean of the series, useful to set the long-term behaviour of the series.

Once a model has been selected, it should be specified in the **arima** spec in the following way:

arima{model=(p, d, q)(P, D, Q)}

where p, d, q are the orders of the regular part of the model and P, D, Q are the orders of the seasonal part. The **arima** specification should be used for the production runs of the forthcoming year. The form of the model should be re-estimated once a year during the seasonal adjustment review process.

Finally, it must be emphasised that if the model used poorly describes the series, then the forecasts and backcasts that are generated are also of poor quality. However, for seasonal adjustment purposes, the gain in using the forecasts is greater than using asymmetric weights in the moving average calculation. Therefore if the **automdl** or the **pickmdl** methods fail to select a model, a simple model (i.e. 0 1 1 0 1 1) should be fixed in the arima spec so that forecasts can be used in the calculation of the moving averages.

25.3 Forecasting for a purpose different from seasonal adjustment

Forecasting techniques are used extensively in National Accounts to extend series beyond the most recent period for which data are available. The regARIMA part of the X12ARIMA program can be also used for this aim to extend the series forwards (adding forecasts) and backwards (adding backcasts). If the forecasting performance of the methods, looked at in section 2, are not different to when forecasting for a purpose other than for seasonal adjustment, other considerations will need to be made on the forecast horizon and on the length of the series.

25.3.1 Model selection to generate forecasts

There are two criteria to select a model with X12ARIMA to generate forecasts:

1. Manual model selection using autocorrelation and partial autocorrelation functions;
2. Automatic model selection.

The model selection criteria depend on the length of the series. In fact:

- If less than 3 years of data are available, other methods of forecasting should be used (e.g. simple extrapolation methods). At least 3 years of data are required for X12ARIMA to operate.
- If 3-5 years of data are available, X12ARIMA can use ARIMA modelling or provide trading day and Easter adjustments, but they are generally of very poor quality and subject to large revisions as future observations become known. At least 5 years of data are required for any automatic model selection to operate.
- If more than 5 years of data are available, X12ARIMA can use automatic model selection and provide good estimates for trading day and Easter adjustments.

25.3.1.1 *Identifying a model manually to generate forecasts*

If the series is less than five years long but has more than 3 years of data available, it is not possible to use the automatic model selection. In this case, users could undertake manual model identification from scratch to identify the model to generate forecasts.

The program can produce the necessary autocorrelation outputs, and there are textbooks that describe the procedure, but it requires considerable experience to produce reliable results. In fact, the process of manual refinement needs a degree of skill and care and advice should be sought if in any doubt. The program provides a number of diagnostics that are helpful in this process particularly the model identification statistics (AICC, BIC etc.). Details of the use of these may be found in the full X12ARIMA documentation.

Once the appropriate model specification has been found, it should be specified in the **arima** spec in the following way:

arima{model=(p, d, q)(P, D, Q)}

where p, d, q are the orders of the regular part of the model and P, D, Q are the orders of the seasonal part. The seasonal part may, optionally, be followed by the seasonal period (either 4 or 12), and if necessary the model may be extended with other brackets followed by an appropriate period. (In common with other X12ARIMA specs, the commas in the inner brackets may be omitted provided the figures are separated by spaces.) The other possible arguments of **arima** are concerned with pre-specifying parameter values and should not be used.

It is important to keep in mind that if 3-5 years of data are available X12ARIMA can use ARIMA modelling or provide trading day and Easter adjustments but they are generally of very poor quality and subject to large revisions as future observations become known.

Another approach to X12ARIMA can be used if only 3-5 years of data are available: the Holt-Winters method. Holt-Winters basic idea involves a procedure known as exponential smoothing. This is a process in which a predicted value is updated each time new information becomes available at the end of a series. It takes its name from the use of moving averages with exponentially declining weights that ensure that the most recent and relevant data points in the series supply the most information to the predictions.

In the Holt-Winters method, predictions of the level, slope, and seasonality of a series are updated using exponential smoothing. These predictions are then combined to give a forecast of the next observation in the series. Holt-Winters has additive and multiplicative forms, which determine whether the seasonal component is added to, or multiplied by, the level and slope components.

For example, if simple forecasts of a large number of series are required and between 3-5 years of data are available then the most sensible approach would be to use "Holt-Winters" across the board. This would provide good quality forecasts of most series without the lengthy manual modelling of the "ARIMA" approach.

25.3.1.2 Using RegARIMA model to generate forecasts

Section 2 described the model selection criteria and the options that should be used to fit a model before the seasonal adjustment itself is performed. This section will go back to those suggestions and will give details on the method and forecast horizon to use to forecast series for a purpose other than seasonal adjustment.

Extensive analysis of the ONS data shows that the appropriate forecasting methods derived in section 2 for seasonal adjustment purposes are still appropriate and that forecast performances for the recommended method are not dissimilar from the performances of the alternative methods. This suggests that if the recommended method fails to select a model, users can equally rely on the alternative approaches in terms of the quality of forecasts. A major difference from the methods described in section 2 is the forecast horizon. In fact, the forecast horizon can now be increased to the maximum reported in the table below if the fit of the selected model is satisfactory.

Table 25-2 Appropriate modelling method and forecast horizon

	Recommended method / horizon		Alternative method / horizon	
Series length	Monthly series	Quarterly series	Monthly series	Quarterly series
<6 years	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 1 year 	<ul style="list-style-type: none"> ▪ pickmdl (first, default list); ▪ 1 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(2,1), no constant), ▪ 1 year 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(2,1), no constant),, ▪ 1 years
6 – 10 years	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 2 years 	<ul style="list-style-type: none"> ▪ Pickmdl (first, extended list); ▪ 2 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(3,1), no constant), ▪ 2 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(3,1), no constant), ▪ 2 years
10 - 12 years	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 3 years 	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 3 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(4,1), no constant), ▪ 3 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(3,1), no constant), ▪ 3 years
>12 years	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 3 years 	<ul style="list-style-type: none"> ▪ pickmdl (first, extended list); ▪ 3 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(4,2), constant), ▪ 3 years 	<ul style="list-style-type: none"> ▪ automdl (maxlag=(4,2), no constant), ▪ 3 years

The first entry in each cell is the recommended method, and the second is the maximum number of forecasts that could be generated. The same attention mentioned in section 2 on the constant term should be paid for series with more than 12 years of data available.

Another major difference between forecasting for seasonal adjustment purposes and forecasting for other purposes is that for the former, the focus is not on the quality of the forecast but on the quality of the moving average whilst for the latter purpose the accuracy of the forecast is very important. For this reason, when forecasting for purposes other than seasonal adjustment, the model should not be fixed and the recommended automatic model selection should be run every time a new forecast is required. If the automatic model selection procedure fails to select a model, whether this automatic procedure is the **automdl** or the **pickmdl** one, no arima forecasts should be generated. In this situation users should not use X12ARIMA for forecasting and another forecasting technique should be used, such as the Holt-Winters one.

It is important to keep in mind that forecasts estimated beyond the forecast horizon reported in table 25-2 have a high forecast error and therefore are not reliable. In addition, if a series contains more forecasts than the one suggested in table 25-2, that series should be considered as unreliable and it should not be used in any system. The reasons as to why that series required a longer forecast horizon should be considered (i.e. unavailability of recent data points) and a different solution should be implemented (i.e. investigate the issue with data compiler and agree, if possible, on an alternative way to get the recent data points).

If automatic identification has selected a model that, while satisfying the tests, still has some unsatisfactory features then the following refinements can be used to improve the quality of the model selected by X12ARIMA:

- If the X12 output states that there is evidence of regular (non-seasonal) overdifferencing, **automdl** approach should be used allowing for a constant term to be selected if significant. However, if the order of the non-seasonal difference is not reduced (e.g. from 1 to 0), it is better to use the initially selected model.
- If X12 output states that there is evidence of seasonal overdifferencing, **automdl** approach should be used with seasonal dummies in the regression part of the model. However, the seasonal dummies should be removed if they are statistically not significant. This could be the case when the series is not seasonal.

The use of these refinements is a matter of judgement and should be used together with other manual refinements when the automatic model selected seems counterintuitive, when the series to be forecasted is important and when experienced resources are available to carry out the analysis.

25.3.2 Model validation

The Box-Jenkins forecasting methodology used by X12ARIMA carries with it certain hazards. In particular, since model selection is an iterative procedure, it is possible to end up with a model which describes the data fairly well which could be due to one of the two following reasons:

1. The model is a good model;
2. So many different models have been tried that users have, quite by chance, found a bad model that just happened to fit their particular data reasonably well.

If the model is a good model, then presumably it will produce good forecasts. But, if the model used describes the series poorly then the forecasts and backcasts that are generated are also poor. Using the

quality measures reported in Chapter 27, for model validation, can reduce the probability of selecting a bad model by chance. This means before using a model for forecasting users should verify the accuracy of the model and manually refine it where necessarily.

For example, although the test on the serial correlations of the residuals may be passed, there may still be some individual significant correlations at fairly low lags. It may then be justifiable to try manual refinement of the automatic model. In this example, it could be worthwhile adding an extra coefficient at the appropriate lag to the AR or MA component; if the extra coefficient is significant and the significant serial correlation has been removed, the extra term might be justified.

25.3.3 Considerations

- How far ahead to forecast? This depends on individual user requirements although it is often found that only one or two future points are estimated. Users should be aware that the further beyond the end of the real data a point being forecast is, the less reliable is the estimate (i.e. the less likely it is to be correct when the real data for that point becomes available). For this reason, it is recommended that forecasts of data points further than the recommended forecast horizon are not used as estimates.
- Is the data non-seasonal? If the time series being forecast is already seasonally adjusted then the alternative **automdl** method should be used to select the model and check that no seasonal coefficients have been included in the model. The same applies for series for which it is known from the seasonal adjustment re-analysis that are not seasonal (check the seasonal adjustment analysis report).
- If the data is seasonal, then is the time series additive or multiplicative? This will involve using "log" or "none" function in the **transform** spec to invoke either the no transformation (additive) model or log transformation (multiplicative) model. This is determined by whether the seasonal component of the series is in the form of a factor or an additive quantity. The decision between additive and multiplicative is similar to that in seasonal adjustment:
 - ✓ Graph the series to check whether it has additive or multiplicative properties. (see "Guide to seasonal adjustment with X11ARIMA" Red Book for details)
 - ✓ Check the seasonal adjustment review report if available. This will state whether the series is seasonally adjusted on an additive or multiplicative basis.
 - ✓ If a series contains any zero or negative values, then the additive method must be used to prevent any division by zero.
 - ✓ If it is still unclear, it is recommended that the multiplicative method be used.
- Is the series long enough? ARIMA models requires a minimum of five years to automatically select a model but a minimum of three years to fix a regARIMA model to carry out a basic forecast. Less than three years of data would not provide enough information about the seasonal pattern of the series. Even with three years though, the forecast would be of poor quality and ideally the series should be as long as possible. If three years of data is not available, it would not be possible to use the "ARIMA" specification and other methods should be used (i.e. Holt-Winters or extrapolation methods).
- Is the series being forecast an interpolated series? Interpolated data should not normally be forecast. Forecasting should be conducted prior to interpolation wherever possible.
- Outliers, Level Shifts and Seasonal Breaks. These are all unusual features which can occur in time series data but which have the potential to undermine or distort the forecasts produced by the regARIMA. This is particularly true where the feature is close to the end of the series. There are two possible ways of approaching this problem:

1. An easy solution is to ignore the series up to the point of the break/outlier and just use the subsequent part of the series for the forecast. The effectiveness of this will obviously depend on the length of series available after the event (see comment above on length of series) and the irregularity over that period.
2. A more reliable way to treat outliers, level shifts or seasonal breaks is to treat them in the same manner as in seasonal adjustment procedures, that is, with regressors in the regARIMA model. This involves setting up a regressor variable as described in Chapter 14 for a level shift and/or an outlier, or in Chapter 17 for a seasonal break.

25.4 Cross-references

- The RegARIMA model.
- Level shift and additive outliers.
- Seasonal Breaks.
- Quality Measures.

26 TREND ESTIMATION

26.1 Introduction

Many of ONS time series outputs are presented in seasonally adjusted form. However, a seasonally adjusted series still contains irregular movements and volatility, which can sometimes obscure the underlying behaviour in the data that many users are primarily interested in. Trend estimation is an approach used to reduce the impact of the irregular movement in the seasonally adjusted estimates so as to highlight the medium- to long-term underlying behaviour.

Following a research project conducted in 1997, the ONS adopted a standard method for estimating trends for monthly time series published in First Releases. This method uses X-11-ARIMA or X-12-ARIMA and is the only method currently recommended for estimating trend-lines (see, Knowles and Kenny, 1997).

26.2 The I/C ratio and the MCD

When estimating trends, there are two related characteristics of the seasonally adjusted series that are important: the noise to trend ratio (commonly referred to as the I/C ratio), and the Months for Cyclical Dominance (MCD). These can both be determined by performing a seasonal adjustment in X11ARIMA or X-12-ARIMA and referring to the "summary measures" option.

The I/C ratio can be found in Table F2H of the analytic output of X11ARIMA or X-12-ARIMA, as in the following example:

```
F 2.H: The final I/C Ratio from Table D12: 0.93
      The final I/S Ratio from Table D10: 5.92
```

The MCD can be found in table F2E of the analytic output, along with the I/C ratio for certain spans, as shown in the example output below.

```
F 2.E: I/C RATIO FOR MONTHS SPAN
1      2      3      4      5      6      7      8      9      10     11     12
0.79   0.39   0.23   0.17   0.14   0.12   0.10   0.09   0.08   0.07   0.06   0.05

MONTHS FOR CYCLICAL DOMINANCE: 1
```

The I/C ratio represents the average percentage change in the irregular (I) component (the noise) in the series in comparison to the average percentage change in the trend (C) (the signal) over a certain span. This gives a measure of the volatility of a series. The output has a number of different I/C ratios, all for different spans of data. The very left hand figure (0.79), is the I/C ratio for a one month span, so is equal to the average percentage change in the irregular from one month to the next divided by the average percentage change in the trend from one month to the next. The figure of 0.79 suggests that the change in the trend dominates the movement of the series over a one month span meaning that the series is not particularly volatile.

The next figure in the line (0.39) is the I/C ratio for a two month span. It is equal to the average percentage change in the irregular from one month to two months later divided by the average percentage change in the trend from one month to two months later. This figure of 0.39 is less than that for a one-month span, as the longer the span over which we look, the more dominant the long-term trend

of the series becomes, and the less influence the irregular has. The figures continue all the way up to a 12-month span.

The MCD represents the span of months it takes for the I/C ratio to fall below 1, that is, for the trend to dominate the irregular. It is useful as it gives an indication of the number of months of the seasonally adjusted time series, on average, over which you will need to look to get an indication of the underlying trend without the irregular volatility of the series getting in the way. In this example, the I/C ratio is less than 1 for a one month span, so it is possible to get a good indication of the underlying trend of the series by looking at the movement from one month to the next. Hence the Months for Cyclical Dominance is 1. The Months for Cyclical Dominance is given as a rounded number to the nearest month. The exact number could be calculated.

In the example below, the situation is slightly more complex.

F 2.E: I/C RATIO FOR		MONTHS SPAN									
1	2	3	4	5	6	7	8	9	10	11	12
2.47	1.28	0.81	0.59	0.44	0.42	0.35	0.31	0.27	0.23	0.21	0.20
MONTHS FOR CYCLICAL DOMINANCE:		3									

In this case the I/C ratio does not drop below 1 until a three-month span is used, hence the MCD is 3. This means that in the seasonally adjusted series, we will need to look across a 3 month span (on average) in order to discern the underlying trend.

It is worth noting that problems arise with series that do not exhibit a long-term upwards trend (i.e. are quite flat). In these instances, the average percentage change in the trend will be very small even over longer spans. As we are dividing by this value to calculate the I/C ratio, I/C could become very large even if there is not much volatility in the series. This would lead to a very high MCD even though the series isn't volatile.

A more formal explanation can be found in Economic Trends (1972). A copy of this can be obtained from Time Series Analysis Branch.

26.3 The standard trend estimation method

Trend estimates are produced as part of the standard seasonal adjustment process. The trend estimates that are generated by X-12-ARIMA can be extracted and used. This is the easiest option to obtain trend estimates.

An alternative approach is to use the derived seasonally adjusted estimates and apply a user defined smoother to the seasonally adjusted estimates. This can give greater flexibility in calculating a trend estimate as the user has more control of the smoothing functions that will be applied.

26.4 Presentation of trends

The following extract is taken from the Economic Trends article "Estimating and Presenting Short-Term Trends" and sets out the standards for the presentation of trends in First Releases.

"Trend estimates should not be quoted as headline figures; they should always be given less emphasis than the seasonally adjusted series. They should be presented in a graphical form on the front page of a First Release and numbers should be made available on request. The graph on the front page should show the last 15 months seasonally adjusted data and trend. The trend should be represented by a solid

line with a dashed end to reflect the relative uncertainty of the trend at the end of the series. The length of the dashed part of the line should be determined by the following:

For a Months for Cyclical Dominance (MCD - the number of months, on average before movements in the trend dominate irregular movements in the series) of:

- 1 - use a dashed line for the most recent month;
- 2 - use a dashed line for the most recent two months;
- 3+ - use a dashed line for the most recent three months;

All commentary should be written in the past tense.

As an optional addition, 'what-if' graphs can be shown in the background notes of the First Release. These give a clearer indication of the degree of uncertainty of trend estimates at the end of the series."

Chart 2.2a shows an example of the standard presentation of the headline seasonally adjusted data with the trend, while Chart 2.2b gives an example of a "what-if" or "Trumpet" graph (note: the trumpet begins 4 points from the end of the series).

Chart 2.2a - Front Page, First Release

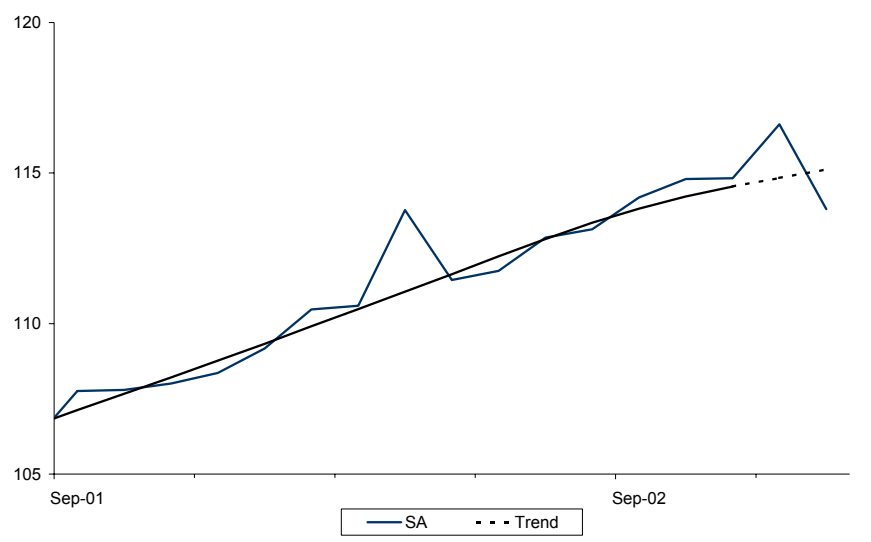



Chart 2.2b - "What-If" Graph, Background Notes

26.5 Considerations

When using X-12-ARIMA to estimate the trend component, the following issues need to be taken into account:

- **What is the form of the input series?** The data that feeds into the trend estimation method should be the published seasonally adjusted series from the first release.
- **Should the additive or multiplicative decomposition be used?** The two approaches to estimate the trend component depend on the decomposition method in the "x11" argument of the specification file. The choice between additive and multiplicative should already have been made when the input series was seasonally adjusted. This information can be found through a check of the X-12-ARIMA parameter file used to seasonally adjust the data or a look at the Annual Seasonal Adjustment Review report. If the input series was not seasonally adjusted directly (i.e. it is an aggregate series constructed by summing lower level seasonally adjusted composite series), then this information will obviously not be available. The choice should then be made using the following criteria:
 - ✓ Does the series have negative values? If so, use "add" in the **mode** option of the **x11** specification. If not, proceed to the next step.
 - ✓ When viewing a graph of the time series, does the irregular volatility increase as the level of the series increases? If yes, use "mult" in the **mode** option of the **x11** specification. If no, use "add". If unsure, proceed to next step.
 - ✓ What is the nature of your data? If the series is an index, use "mult" in the **mode** option of the **x11** specification. If the series is a difference (eg Balance of Payments = Exports - Imports), use "add". If a choice has still not been made, run the function first with "add", then with "mult" and compare the quality of the two resulting trends.

27 QUALITY MEASURES

27.1 Introduction

The Government Statistics Service (GSS) has long recognised the need to provide users with information about the quality of statistics and about the analytical techniques used to derive the figures. This chapter provides information about quality measures and indicators that can be used for considerations when measuring and reporting on the quality of time series outputs.

More information on quality measures can be found in “Guidelines for measuring statistical quality”.

27.2 What is “quality”

Quality can have many different meanings, for example, luxury; good value for money and convenience. Its meaning will often depend upon the context in which it is used. Quality, in terms of statistical outputs, can generally be thought of as a degree to which the data meet user needs, or simply put, the degree to which the data are “fit for purpose”.

Quality has often been associated with accuracy and timeliness. But even if statistics are accurate and timely, they cannot be deemed to be good quality if they are not based on concepts which are meaningful and relevant to the users. In addition, different users will have different needs. Quality measurement and reporting for statistical outputs is therefore concerned with providing the user with sufficient information to judge for themselves whether or not the data are of sufficient quality for their intended use.

The quality of a statistical output should be determined by its performance against a range of attributes that together can be used to assess whether an output meets users’ quality criteria.

The office for National Statistics has adopted the data quality attributes defined for the European Statistical System (ESS), which are shown in the following table:

Definition	Key components
1. RELEVANCE	
The degree to which the statistical product meets user needs for both coverage and content.	Any assessment of relevance needs to consider: <ul style="list-style-type: none"> • who are the users of the statistics? • what are their needs? • how well does the output meet these needs?
2. ACCURACY	
The closeness between an estimated result and the (unknown) true value.	Accuracy can be split into sampling error and non-sampling error, where non-sampling error includes: <ul style="list-style-type: none"> • coverage error; • non-response error; • measurement error; • processing error; and • model assumption error.

Definition	Key components
3. TIMELINESS AND PUNCTUALITY	
Timeliness refers to the lapse of time between publication and the period to which the data refer. Punctuality refers to the time lag between the actual and planned dates of publication.	An assessment of timeliness and punctuality should consider the following: <ul style="list-style-type: none"> • production time; • frequency of release; and • punctuality of release.
4. ACCESSIBILITY AND CLARITY	
Accessibility is the ease with which users are able to access the data, also reflecting the format(s) in which the data are available and the availability of supporting information. Clarity refers to the quality and sufficiency of the metadata, illustrations and accompanying advice.	Specific areas where accessibility and clarity may be addressed include: <ul style="list-style-type: none"> • needs of analysts; • assistance to locate information; • clarity; and • dissemination.
5. COMPARABILITY	
The degree to which data can be compared over time and domain.	Comparability should be addressed in terms of comparability over: <ul style="list-style-type: none"> • time; • spatial domains (sub-national, national, international); and • domain or sub-population (industrial sector, household type).
6. COHERENCE	
The degree to which data that are derived from different sources or methods, but which refer to the same phenomenon are similar.	Coherence should be addressed in terms of coherence between: <ul style="list-style-type: none"> • data produced at different frequencies; • other statistics in the same socio-economic domain; • sources and outputs.

Quality of data can rarely be explicitly “measured”. For example, in the case of accuracy, it is almost impossible to measure non-response bias as the characteristics of those who do not respond can be difficult to ascertain. Instead, certain information can be provided to help indicate quality. Quality indicators usually consist of information that is a by-product of the statistical process. They do not measure quality directly, but can provide enough information to make inferences about the quality. Section 3 includes both quality measures and suitable quality indicators that can either supplement or act as substitute for the desired quality measure.

27.3 Time series quality measures


The following table reports the quality measures specific for time series analysis together with the output for which they are relevant and the European Statistical System they refer to.

N.	Quality Measure	Outputs applicable to	ESS Dimension
1	Original data visual check	All time series	Accuracy

N.	Quality Measure	Outputs applicable to	ESS Dimension
2	Comparison of the original and seasonally adjusted data	All seasonal adjustments	Accuracy
3	Graph of Seasonal-Irregular (SI) ratios	All time series	Comparability
4	Analysis Of Variance (ANOVA)	All seasonal adjustments and trend estimations	Accuracy
5	Months (or Quarters) for Cyclical Dominance	All seasonal adjustments and trend estimations	Comparability
6	The M7 statistic	All seasonal adjustments	Accuracy
7	Contingency Table Q	All seasonal adjustments and trend estimations	Comparability
8	Stability of Trend and Adjusted Series Rating (STAR)	All seasonal adjustments and trend estimations	Accuracy
9	Comparison of annual totals before and after seasonal adjustment	All seasonal adjustments	Accuracy
10	Normality test	All forecasts	Accuracy
11	P-values	All forecasts	Accuracy
12	Percentage standard forecast error	All forecasts	Accuracy
13	Graph of the confidence intervals	All forecasts	Accuracy
14	Percentage difference of unconstrained to constrained values	All constrained series	Accuracy


Each quality measure is presented below with an example or formula describing their use and notes providing more detail on the measure/indicator.

1. **Original data visual check.** The graph below shows the Airline Passengers original series. From the graph it is possible to see that the series has repeated peaks and troughs that occurs at the same time each year. This is a sign of seasonality. It is also possible to see that the trend is affecting the impact of the seasonality. In fact the size of the seasonal peaks and troughs are not independent of the level of the trend suggesting that a multiplicative decomposition model is appropriate for the seasonal adjustment. This series does not show any particular discontinuities (outliers, level shifts or seasonal breaks).




Notes: Graphed data can be used in a visual check for the presence of seasonality, decomposition type model (multiplicative or additive), extreme values, trend breaks and seasonal breaks.

2. **Comparison of the original and seasonally adjusted data.** By comparing the original and the seasonally adjusted series, in the graph below, it can be seen that the seasonal adjustment performs well until January 2002 where the presence of an outlier distorts the pattern of the series.



Notes: By graphically comparing the original and seasonally adjusted series, it can be seen whether the quality of the seasonal adjustment is affected by any extreme values, trend breaks or seasonal breaks and whether there is any residual seasonality in the seasonally adjusted series.

3. **Graph of Seasonal-Irregular (SI) ratios.** In the example below, there has been a sudden drop in the level of the Seasonal-Irregular component (called SI ratios or detrended ratios) for August between 1998 and 1999. This is caused by a seasonal break in the Car Registration series which was due to the change in the car number plate registration legislation. Permanent prior adjustments should be estimated to correct for this break. If no action is taken to correct for this break, some of the seasonal variation will remain in the irregular component resulting in residual seasonality in the seasonally adjusted series. The result would be a higher level of volatility in the seasonally adjusted series and a greater likelihood of revisions.



Notes: It is possible to identify a seasonal break by a visual inspection of the seasonal irregular graph (graph of the SI ratios). Any change in the seasonal pattern indicates the presence of a seasonal break.

4. **Analysis Of Variance (ANOVA).** The formula to calculate the Analysis Of Variance statistic is as follows:

$$ANOVA = \frac{\sum_{t=2}^n (D12_t - D12_{t-1})^2}{\sum_{t=2}^n (D11_t - D11_{t-1})^2}$$

Where:

D12t = data point value for time t in table D12 (final trend cycle) of the analytical output;

D11t = data point value for time t in table D11 (final seasonally adjusted data) of the output;

When calculating the ANOVA the following considerations should be take into account:

- If constraining is used and hence the D11A table is produced D11A point values are used in place of D11 in the above equation.
- If the statistic is used as a quality indicator for trend, table A1 (final seasonally adjusted series from which the final trend estimate is calculated) should be used instead table D11. The D12t values should be taken from the final trend output.

Notes: The ANalysis Of VAriance (ANOVA) compares the variation in the trend component with the variation in the seasonally adjusted series. The variation of the seasonally adjusted series consists of variations of the trend and the irregular components. ANOVA indicates how much of the change of the seasonally adjusted series is attributable to changes in primarily the trend component. The statistic can take values between 0 and 1 and it can be interpreted as a percentage. For example, if ANOVA=0.716, this means that 71.6% of the movement in the seasonally adjusted series can be explained by the movement in the trend component and the remainder is attributable to the irregular component.

This indicator can also be used to measure the quality of the estimated trend.

5. **Months (or Quarters) for Cyclical Dominance.** The months for cyclical dominance (MCD) or quarters for cyclical dominance (QCD) are measures of volatility of a monthly or quarterly series respectively. The formula to derive the statistic is as follows:

$$MCD \text{ (or QCD)} = d \text{ for which } \frac{I_d}{C_d} < 1$$

where I_d is the final irregular component at lag d and C_d is the final trend component at lag d

Notes: This statistic measures the number of periods (months or quarters) that need to be spanned for the average absolute percentage change in the trend component of the series to be greater than the average absolute percentage change in the irregular component. For example, an MCD of 3 implies that the change in the trend component is greater than the change in the irregular component for a span at least 3 months long. The MCD (or QCD) can be used to decide the best measure of short-term change in the seasonally adjusted series - if the MCD =3 the three months on three months growth rate will be a better estimate than the month on month growth rate. The lower the MCD (or QCD), the less volatile the seasonally

adjusted series is and the more appropriate the month on month growth rate is a measure of change. The MCD (or QCD) value is automatically calculated by X12ARIMA and is reported in table F2E of the analytical output.

For monthly data the MCD takes values between 1 and 12, for quarterly data the QCD takes values between 1 and 4.

6. **The M7 statistic.** The formula behind M7 statistic is as follows:

$$M7 = \sqrt{\frac{1}{2} \left(\frac{7}{F_s} + \frac{3F_m}{F_s} \right)}$$

Where:

$$F_M = \frac{S_B^2 / (N-1)}{S_R^2 / (N-1)(k-1)} \quad F_S = \frac{S_A^2 / (k-1)}{S_R^2 / (n-k)}$$

F_M = F-statistic capturing moving seasonality.

S_B^2 is the inter-year sum of squares, S_R^2 is the residual sum of squares, N is the number of observations.

F_S = F-statistic capturing stable seasonality.

S_A^2 is the variance due to the seasonality factor and S_R^2 is the residual variance

M7 compares on the basis of F-test statistics the relative contribution of stable (statistic F_M) and moving (statistic F_S) seasonality.

Notes: This indicates whether the original series has a seasonal pattern or not. Low values indicate clear and stable seasonality has been identified by X12ARIMA. The closer the M7 statistic is to 0, the better the seasonal adjustment. Values around 1 imply that the series is marginally seasonal. An M7 value close to 3 indicates that the series is non-seasonal. The M7 statistic is calculated by X12ARIMA and can be obtained from the F3 table of the analytical output.

7. **Contingency Table Q.** The formula to analyse the Contingency Table Q (also called CTQ in the example below) is as follows:

$$CTQ = \frac{U_{11} + U_{22}}{U_{11} + U_{12} + U_{21} + U_{22}}$$

U_{kl} = value for contingency table cell with row k and column l

where

	Delta C > 0	Delta C <= 0
Delta SA > 0	U11	U12
Delta SA <= 0	U21	U22

$\Delta SA = D11_t - D11_{t-1}$ (Delta SA is the change in the SA data)

$\Delta C = D12_t - D12_{t-1}$ (Delta C is the change in the trend)

Notes: The Contingency Table Q (CTQ) shows how frequently the gradient of the trend and the seasonally adjusted series over a one period span have the same sign. CTQ can take values between 0 and 1. A CTQ value of 1 indicates that historically the trend component has always moved in the same direction as the seasonally adjusted series. A CTQ value of 0.5 suggests that the movement in the seasonally adjusted series is likely to be independent from the movement in the trend component, this can indicate that the series has a flat trend or that the series is very volatile. A CTQ value between zero and 0.5 is unlikely but would indicate that there is a problem with the seasonal adjustment.

If constraining is used point values from table D11A are used in place of D11 in the above equation.

8. **Stability of Trend and Adjusted Series Rating (STAR).** The formula to calculate the Stability of Trend and Adjusted Series Rating (STAR) is as follows:

$$STAR = \frac{1}{N-1} \sum_{t=2}^N \left| \frac{(D13_t - D13_{t-1})}{D13_{t-1}} \right|$$

Where D13_t = data point value for time t in table D13 (final irregular component) of the output, and N = number of observations in table D13.

Notes: This indicates the average absolute percentage change of the irregular component of the series.

The STAR statistic is applicable to multiplicative decompositions only. The expected revision of the most recent estimate when a new data point is added is approximately half the value of the STAR value e.g. a STAR value of 7.8 suggests that the revision is expected to be around 3.9%.

9. **Comparison of annual totals before and after seasonal adjustment.** The formula for multiplicative models is as follows:

$$\frac{1}{n} \sum_{i=1}^n \left| (E4_{total(t)}^{D11} - 100) \right|$$

Where $E4_{total(t)}^{D11}$ = unmodified ratio of annual totals for time t in table E4.

E4 is the output table that calculates the ratios of the annual totals of the original series to the annual totals of the seasonally adjusted series for all the n years in the series.

The formula for additive models is:

$$\frac{1}{n} \sum_{i=1}^n \left| \frac{E4_t^{D11}}{D11_{total(t)}} \right|$$

$E4_{total(t)}^{D11}$ = unmodified difference of annual totals for time t in table E4.

For additive models E4 is the output table that calculates the difference between original annual total and seasonally adjusted annual totals for all the n years in the series.

Notes: This is a measure of the quality of the seasonal adjustment and of the distortion to the seasonally adjusted series brought about by constraining the seasonally adjusted annual totals to the annual totals of the original series. It is particularly useful to judge if it is appropriate for the seasonally adjusted series to be constrained to the annual totals of the original series.

10. **Normality test.** The two statistics below (Geary's and Sample Kurtosis respectively) test the regARIMA model residuals for deviations from normality.

$$\alpha = \frac{\frac{1}{n} \sum_{i=1}^n |X_i - \bar{X}|}{\sqrt{\frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2}} \quad b_2 = \frac{m_4}{m_2^2} = \frac{n \sum_{i=1}^n (X_i - \bar{X})^4}{\left(\sum_{i=1}^n (X_i - \bar{X})^2 \right)^2}$$

Notes: The assumption of normality is used to calculate the confidence intervals from the standard errors. Consequently, if this assumption is rejected the estimated confidence intervals will be distorted even if the standard forecast errors are reliable. A significant value of one of these statistics (Geary's and Kurtosis) indicates that the standardised residuals do not follow a standard normal distribution, hence the reported confidence intervals might not represent the actual ones. (X12ARIMA tests for significance at one percent level)

11. **P-values.** An example of the p-value quality measure is as follows:

DIAGNOSTIC CHECKING												
Sample Autocorrelations of Residuals												
Lag	1	2	3	4	5	6	7	8	9	10	11	12
ACF	-0.06	0.12	-0.1	-0.31	-0.25	-0.13	0.21	-0.09	0.21	-0.05	0.09	-0.15
SE	0.19	0.19	0.2	0.2	0.21	0.22	0.23	0.23	0.24	0.24	0.24	0.24
Q	0.09	0.56	0.86	4.04	6.2	6.8	8.51	8.86	10.8	10.89	11.31	12.44
DF	0	0	1	2	3	4	5	6	7	8	9	10
P	0	0	0.355	0.133	0.102	0.147	0.13	0.181	0.148	0.208	0.255	0.256

Model fitted is: ARIMA (0 1 1)(0 1 1)

(Quarterly data so 12 lags will be extracted)

The model will not be adequate if any of the p-value's of the 12 lags is less than 0.05 and will be adequate if all the p-value > 0.05.

Because we are estimating 2 parameters in this model (one AR and one MA parameter), no p-values will be calculated for the first 2 lags as 2 DF's will be automatically lost hence we will start at lag 3. Here all the p-value's are greater than 0.05 e.g. at lag 12 Q=12.44, DF=10 and p-value = 0.256, therefore this model is adequate for forecasting.

Notes: This is found in the 'Diagnostic checking Sample Autocorrelations of the Residuals' output table. The p-values show how good the fitted model is. They measure the probability of the ACF occurring under the hypothesis that the model has accounted for all serial correlation in the series up to the lag. P-values greater than 0.05, up to lag 12 for quarterly data and lag 24 for monthly data, indicate that there is no significant residual autocorrelation and that, therefore, the model is adequate for forecasting.


12. **Percentage forecast standard error.** The percentage forecast standard error is given by:

$$\frac{\text{forecast standard error}}{\text{forecast}}$$

Notes: The percentage forecast standard error is required for each forecast produced and can be found in the forecast table. There will be one number for each period that has been forecasted.

The percentage forecast standard error is applicable to multiplicative decompositions only.

13. **Graph of the confidence intervals.** An example of the graph of confidence intervals is reported below.



Notes: Time to be plotted along the x-axis, and forecast estimates and confidence intervals along the y-axis.

14. **Percentage difference of unconstrained to constrained values.** The formula to compute the percentage difference of unconstrained to constrained values is as follows:

$$\left(\frac{\text{constrained data point}}{\text{unconstrained data point}} - 1 \right) \times 100$$

Notes: This indicates the percentage difference between the unconstrained and the constrained value. This will need to be calculated for each data point being constrained. The average, minimum and maximum percentage differences (among all data points of all the series) of the unconstrained should be calculated to give an overview of the effect of constraining.

27.4 How to interpret time series quality measures

The following section provide examples to help users in interpreting the quality measures and the quality indicators when running seasonal adjustment, trend estimation, forecasting and constraining.

When using the quality measure to interpret a time series it is important to keep in mind that none of the diagnostics should be used to explain the quality of a single component of a time series, but to give users information on the characteristics of the series. For example the diagnostics should not be used to explain if the trend is a good or a bad estimate, but to relate the behaviour of the trend with that of the seasonal and irregular components.


27.4.1 Quality measures for seasonal adjustment

The quality indicators for seasonal adjustment should be used to give users information on the characteristics of the series (e.g. to relate the behaviour of the trend with that of the seasonal and irregular components). For example, the quality measure can be used to define:

- if the series has identifiable seasonality (by the size of M7),
- how volatile the series is (by the size of STAR and MCD)' and
- how the trends relates with the seasonally adjusted series (by the size of the ANOVA and CTQ).

The visual inspection of the graph of the original series against the seasonally adjusted series can help users to identify residual seasonality in the seasonally adjusted series.

For example, the graph below shows the original and the seasonally adjusted series.




The M7 statistic is high. This suggests that the series is non-seasonal and it should not be seasonally adjusted. However, as can be seen in the graph, the series is not very irregular and most of the movement in the series is caused by the trend. This is confirmed by the high ANOVA and the Contingency Table Q (CTQ) statistics and by the very low MCD and STAR measures. The low Comparison of Annual Totals (CAT) also confirms that the irregular component does not affect much the series.

27.4.2 Quality measures for forecasting

The quality indicators for the forecast should be used to give users information on the reliability of the projections (e.g. to relate the goodness of the fitted model with goodness of the forecasts). For example, the quality measures can be used to define if the model is suitable for forecasting (by the size of the normality test and the p-value) and the dispersion of the forecasts themselves (by the size of the percentage standard forecast error and the wideness of the confidence intervals). The visual inspection of the confidence intervals can help users understanding how reliable the forecasts are (wider the confidence intervals are and less reliable the projections are).

The following graph shows the confidence intervals of the two-year ahead forecast.



The forecasts are shown by the blue line in the middle. The actual values are expected to lie somewhere between the lower (pink) and upper (yellow) lines, for 95% of the time.

An assumption of normality is used to calculate the confidence intervals. As the kurtosis statistic is significant this assumption is rejected and consequently the reported confidence intervals might not represent the actual ones.


The P-values for this series are high therefore the forecasts and estimated percentage forecast errors are reliable.

The Percentage Forecast Errors are a measure of how much the forecast is expected to differ from the actual value. For January 1998, the first forecasted value, we expect that the actual value will be different than the forecast by 25% on average i.e., Actual value = Forecast(25%). For December 1999, the last forecasted value, we expect the actual value will be different than the forecast by 75% on average.

27.4.3 Quality measures for trend estimates

The quality indicators for the estimated trend should be used to give users information on the characteristics of the series (e.g. to relate the behaviour of the trend with that of the irregular component). For example, the quality measure can be used to define how volatile the seasonally adjusted series is (by the size of STAR and MCD) and how the trends relates with the seasonally adjusted series (by the size of the ANOVA and CTQ). The visual inspection of the graph of the seasonally adjusted series against the trend can help users to identify turning points in the behaviour of the series.

The graph below shows a seasonally adjusted series and the trend, which has been estimated for this series.



The ANOVA statistic is low. This suggests, as can be seen in the graph, that most of the movement from one period to the next in the seasonally adjusted series is caused by the irregular component. Approximately one percent of the month on month change in the seasonally adjusted series is explained by movement in the trend. Also the high MCD statistic confirms that the series is volatile. In fact it takes on average 9 months for the trend component to explain more of the movement in the series than the irregular component. The CTQ statistic of 0.6 is a further indication that the series is volatile. The STAR value of 7.8 suggests that the expected revision of the most recent estimate when a new data point is added will be between 3 and 4%.

27.4.4 Quality measures for constraining

The quality indicators for the constrained series should be used to give users information on the effect of constraining. For example, the quality measures can be used to define how far the constrained series is from the original series. The summary statistics (minimum, average and maximum) give users an indication of the distribution of the percentage differences of unconstrained to constrained values.

For example, after seasonal adjustment, LFS series are constrained so that the property of additivity is fulfilled in every dimension i.e. by age group, employment status etc. This constraining distorts the seasonal adjustment so it is acceptable only if the distortion is small. Eight important high-level aggregation series of total numbers of active, inactive, employed and unemployed for both males and females were selected from the LFS dataset. The average difference (among all data points of all series) was 0.11% of unconstrained. The maximum difference (among all data points of all series) was 0.12% of the unconstrained. For these high-level aggregation series a Quality Measure of 0.11% suggests good constraining thus no problems have been detected with the constraining of these series. The acceptable threshold, for the distortion as a result of constraining, needs to be considered for each individual series.

27.5 Cross-references

- X12ARIMA Standard Output.
- Forecasting.
- Trend.