Australian Bureau of Statistics

New Issue

Information Paper

A Guide to Interpreting Time Series — Monitoring Trends

2003

Dennis Trewin
Australian Statistician

INQUIRIES

AUSTRALIAN BUREAU OF STATISTICS

EMBARGO: 11.30AM (CANBERRA TIME) MON 4 AUG 2003

ABS Catalogue No. 1349.0
ISBN 0 642 54286 4

Produced by the Australian Bureau of Statistics

INQUIRIES

- For further information about these and related statistics, contact the National Information and Referral Service on 1300 135 070 or Mark Zhang on Canberra 6252 5132.

CONTENTS

LIST OF TABLES AND GRAPHS

page

GRAPHS

page

GRAPHS *continued*

PREFACE

A GUIDE TO INTERPRETING TIME SERIES — MONITORING 'TRENDS'

To assist users of Australian Bureau of Statistics (ABS) time series data to identify and analyse the underlying movement of the series, the ABS publishes trend estimates along with seasonally adjusted and original estimates. Increasingly, the main features section and commentaries concerning time series emphasise the trend series, rather than the seasonally adjusted or original data.

This Information Paper explains the statistical concepts and notions that underpin the statistical procedures employed by the ABS to obtain the trend estimates. It is not intended as a complete technical reference, but it does describe in some detail the smoothing techniques used to produce the trend estimates. Details of the smoothing filters used to produce trend estimates are given, including the preliminary filters used towards the end of the time series. The unavoidable use of these preliminary filters at the end of the series is one of the reasons that the most recent trend estimates are subject to subsequent revision.

I would particularly like to thank Mr. John Zarb, former director of Time Series Analysis Section, ABS, who over many years worked to improve the understanding and use of official time series, particularly the trend estimates, and who drafted the original manuscript. Thanks are also due to several other staff ABS staff members for their valuable contribution to the final manuscript, most noteably to Dr. Katrina Hicks and Dr. Mark Zhang.

We hope the Information Paper will help users appreciate the usefulness of trend estimates, and interpret the published time series effectively.

Dennis Trewin
Australian Statistician

ABBREVIATIONS

'000	thousand
$b	billion (thousand million) dollars
$m	million dollars
ABS	Australian Bureau of Statistics
ARIMA	autoregressive integrated moving average
BLUE	best linear unbiased estimate
CLI	Composite Leading Indicator
CPI	Consumer Price Index
EWP	end weight parameter
GDP	gross domestic product
GOS	gross operating surplus
I	irregular
Inf	infinity
O	original series
RBA	Reserve Bank of Australia
SA	seasonally adjusted series
T	trend series

CHAPTER 1 INTRODUCTION

BACKGROUND

Regular readers of ABS publications will have noticed that the Main Features section and the commentaries concerning time series emphasise the trend series rather than the seasonally adjusted or original series. This Information Paper explains this emphasis, before going on to discuss the derivation and properties of the ABS trend estimates, and some of the practical considerations that must be taken into account when interpreting such estimates. Details of the procedure used to derive the ABS trend estimates have been discussed previously in the *Information Paper: A Guide to Smoothing Time Series — Estimates of 'Trend'* (cat. no. 1316.0), 1987. The companion paper *Information Paper: Time Series Decomposition — An Overview* (cat. no. 1317.0), 1987, also provides some motivation for the derivation of trend estimates.

Time series data enable us to assess how real world forces affect levels of social and economic activities over time. However, looked at on a period to period basis, these series are typically volatile, and to understand the underlying level of socioeconomic activity a smoothing mechanism is required. Analysts frequently use sub-optimal smoothing mechanisms that can distort the signals in the data. The ABS produces a trend series which is designed to reflect the longer term behaviour of the series, as a better alternative to the many univariate smoothing procedures in general use.

TREND ESTIMATES

When interest centres on the fundamental direction of a series, ignoring short-term irregular fluctuations, the ABS trend estimates provide smoothed figures with a number of desirable properties, as will be described in detail in this paper. While these trend estimates are generally reliable, there are nevertheless instances where the usefulness of the trend estimates is reduced, either by the high degree of irregularity contained in the original and seasonally adjusted series, or where the time series characteristics of the original data change abruptly. In such cases care must be exercised when using the trend series to make interpretations about real world events, although this consideration also applies to the use of original and seasonally adjusted series in these situations, even to a larger extent.

The ABS generally recommends that those using its time series for business or policy decisions concerned with non-seasonal and non-irregular matters, should give the greatest emphasis to the relevant trend series, bearing in mind the strengths and limitations that will be discussed in this Information Paper.

CHAPTER 2 THE NATURE OF TIME SERIES

CONTENTS

2.0 OVERVIEW

The time series collected by the ABS are statistical records of real world phenomena, measured at regular intervals of time, over relatively long periods. For example, the Employed Persons figures collected monthly since February 1978 form a time series, as do the monthly Retail Trade figures, and the quarterly Balance of Payment figures. However, many other collections run by the ABS do not give rise to time series. Such examples include the Time Use Survey, which is run irregularly, and the Survey of Aspects of Literacy, which has been run only once. The advantage in considering data that forms a time series is that patterns of behaviour can be revealed over time, enabling current estimates to be placed into a more meaningful historical perspective. This in turn permits analysts to describe the current state of play, relative to earlier outcomes.

Monthly series will be used primarily for illustration throughout this paper, although the concepts, principles and methods discussed, also apply to other time series whose data are collected at regular intervals of time, such as half yearly, quarterly, fortnightly, weekly, etc. In those instances when time series are discussed without reference to a specific collection frequency, the general term 'period' will be used to indicate the regular interval of time between data collections. For monthly series the period between data collections is a month, for quarterly series the period between data collections is a quarter, and so on.

2.1 COMPONENTS OF A TIME SERIES

Any time series can be thought of as a combination of three broad and distinctly different types of behaviour, each representing the impact of certain types of real world events on the information being collected. The three components of behaviour are: the systematic calendar related effects, the irregular fluctuations and the trend behaviour. The contribution of each of these behaviours varies from series to series, and indeed throughout time for a given series, depending on the nature of the interactions between real world events and the data of interest. Fortunately, however, these behaviours display distinctly different characteristics, and this enables them to be identified and quantified for a given series, as long as sufficient data is available.

2.1.1 Systematic calendar related effects

Many series react to regular calendar related events in a consistent and predictable manner. The most commonly observed type of calendar related behaviour is the cyclical pattern that evolves as a result of the changes in the seasons; consider the increase in energy consumption with the onset of winter for example. Other cyclical patterns, such as the upswing in retail sales in November and December as a prelude to Christmas, emerge as a result of fixed calendar related events. These cyclical patterns, referred to as seasonal cycles, are not the only patterns of behaviour resulting from regular calendar related occurrences. Other calendar related patterns of behaviour include trading or working day patterns, payday effects, and the systematic patterns of behaviour that can occur as a result of the effects of moving holidays.

Trading, or working day patterns, are purely calendar related behaviours which arise when there is greater activity occurring on some days of the week than on others. The effect arises because each day of the week can occur a different number of times from month to month and year to year (the number of times a particular day of the week occurs in a given month depending on the number of days in the month, and the way the days of the week are distributed in that month in that year). If more activity generally occurs on some days of the week than on others, then the number of times each day

2.1.1 Systematic calendar related effects continued

occurs within each month will have an impact on the behaviour of the series of reported monthly figures (for example, if more activity generally occurs on Thursdays, then a greater level of activity is likely to be reported for a month containing five Thursdays than for a month containing four Thursdays, regardless of how activity levels are changing from month to month). The result is a systematic pattern of behaviour which is based entirely on the calendar enforced manner in which the days of the week are distributed across months throughout the years.

Payday effects arise in a similar manner to trading or working day effects. Some paydays occur fortnightly, and therefore there may be either two or three paydays for a given month, depending on the month and year considered. Consequently, an original monthly payroll series can exhibit large rises and falls from month to month and year to year, based solely on the number of paydays in each month. This influence of the changing number of paydays in a month may also permeate series which are dependent on wages and salaries.

As already indicated above, fixed holidays, such as Christmas, can have an impact on the behaviour of a series. In such cases, the impact on behaviour is consistent, operating at exactly the same point in time each year, and results in the appearance of seasonal cycles. Moving holidays are holidays or events which do not occur in a fixed calendar day from year to year. In particular, they may occur in different months. Moving holidays can also impact on the behaviour of the series, but they do not necessarily affect the same month of the series every year, and consequently, do not necessarily take the form of seasonal cycles. Easter, for example, is a moving holiday which is usually observed in April, but occasionally falls in March, or straddles the end of March and the start of April. This holiday may affect activity in those months as it shifts from one month to the other over the years, resulting in the appearance of non-cyclical calendar related patterns of behaviour.

When the history of an original time series is examined, the various real world influences referred to above can all be seen to operate within the series in a sustained, predictable, and purely calendar related manner. It is these attributes which permit the identification and estimation of their effects upon the series, when sufficient data is available. As these influences have a relatively predictable impact on the behaviour of time series, some analysts have sought to remove their effects from the original data in order to more clearly determine the underlying behaviour of the series. When the systematic calendar related effects are removed from the original data, the resulting series is said to be seasonally adjusted. Henceforth, the cumbersome phrase 'systematic calendar related effects' will be replaced by 'seasonal effects', but it should be noted that the term seasonal is now used in its broadest sense to represent not only seasonal cycles, but also trading and working day effects, payday effects and the impact of moving holidays.

When the seasonal effects are removed from the original series, the seasonally adjusted series still represents the interaction of two remaining distinctly different behaviours: the irregular fluctuations and the trend behaviour.

2.1.2 Irregular/residual fluctuations

The irregular, or residual, fluctuations arise from the hiccups and transient activities that occur in socioeconomic behaviour, as well as the sampling and non-sampling errors that occur as a result of the way the data is collected. These fluctuations often contribute to a large proportion of the non-systematic volatility observed in the behaviour of a given series. Examples of transient socioeconomic effects include the impacts of certain industrial disputes and natural disasters, as well as the impact of the occasional importation of large pieces of military hardware, such as destroyers, or large pieces of capital equipment, such as oil rigs, jumbo jets or satellites. For those series derived from sample surveys, there will be the impact of sampling error, that is, the variability that occurs by chance because a sample, rather than the entire population, is surveyed. There will also be non-sampling statistical error, which represents the inaccuracies that may occur because of imperfections in reporting by respondents, errors made in collection, such as recording and coding of data, and errors made in processing the data to its final form. While attempts are made to minimise the occurrence of non-sampling statistical errors, they will nevertheless occur in any data collection, regardless of whether it is a sample survey or a full count. Sampling and non-sampling errors tend to affect the volatility of a series.

Given the nature of the influences producing this component of behaviour, the irregular fluctuations represent the effect of short-term influences upon the series. Their contribution to a series will generally change in magnitude and direction from period to period, which contrasts markedly with the regular behaviour of the seasonal effects described above, and may in large part represent the effect of 'artificial' and non-economic factors, such as statistical errors. In many situations, the irregular fluctuations are thought of as unpredictable, random like behaviour, which mask the underlying, substantive behaviour of the series. In general this behaviour is not usually regarded as relevant to business decisions or policy advice.

The irregular fluctuations can also be thought of as the residual component of behaviour remaining once the seasonal effects and trend component have been removed. As this is the way in which the components are estimated in practice, the ABS prefers to use the term residual/irregular fluctuations to refer to the component of behaviour resulting from irregular influences.

2.1.3 Trend behaviour

The trend of a series represents the fundamental, or underlying behaviour of the series. It captures the long-term behaviour of the series as well as the various medium-term business cycles. In contrast to the irregular fluctuations, the trend does not frequently change direction from period to period and in fact, trend movements are generally quite smooth and gradual in comparison. Furthermore, unlike the seasonal effects, the trend's turning points are not calendar related, for while the turning points of the 'business cycles' recur, they do not necessarily do so at fixed calendar intervals.

It should be noted that there are many interpretations commonly applied to the term trend. Throughout this paper, trend will be used to refer to the underlying path traced by the medium and long-term cyclical behaviour of the series. Equivalently, trend can be considered to represent the underlying series remaining after removing the seasonal effects from the original data and smoothing out the irregular fluctuations. Where interest centres on the direction of the underlying series behaviour at a given time the

2.1.3 Trend behaviour *continued*

paper will refer to the direction of the trend, whether the long-term pattern of behaviour is showing an increase, or a decrease.

2.1.4 Decomposition of time series

Any time series is composed of the three components of behaviour described above, the seasonal effects, the trend behaviour and the irregular fluctuations. The best method of breaking a time series into these components varies from series to series. Although there are many potential models for describing how these components fit together in a given series, only two models will be considered here. These two models allow the choice of two simple methods which can be used to approximately break any time series down into its three basic components of behaviour.

The first model is the additive model, in which a time series is thought to be the sum of the three basic components of behaviour. If O_t is used to denote the original time series, S_t to denote the seasonal effects, T_t to denote the trend behaviour and I_t to denote the irregular fluctuations, the additive model can be represented symbolically as

$$O_t = T_t + S_t + I_t \qquad (2.1)$$

This model is most appropriate for series in which the behaviour of the irregular fluctuations and the seasonal effects are independent of the trend behaviour; in particular, the seasonal effects behave in a constant manner from year to year regardless of the size of the other components.

The second model is the multiplicative model, in which a time series is thought to be the product of the three basic components of behaviour. This relationship can be represented symbolically as

$$O_t = T_t \times S_t \times I_t \qquad (2.2)$$

This model is most appropriate for series in which the behaviour of the irregular fluctuations and the seasonal effects varies in a particular manner dependent on the behaviour of the trend; in particular, the seasonal effects become more pronounced the higher the level of the trend, and less pronounced the lower the level of the trend (that is, a proportional relationship is maintained). The majority of series analysed by the ABS display the characteristics of the multiplicative time series model.

2.2 THE SPECTRUM OF A TIME SERIES

While it is natural to think of a time series as representing a collection of information across a given interval of time, there is a second, less intuitive, but very powerful way of interpreting any series. A time series can also be thought of as representing the complex interaction of many cycles of behaviour of different strengths across a wide range of frequencies. This part of the paper deals with the consideration of time series from this perspective.

The idea of breaking a time series up into cycles is not actually such a surprising concept, given the above discussion of the three components of time series behaviour. The seasonal cycles, for instance, were a substantial component of the set of systematic calendar related effects discussed earlier. These cycles were distinct patterns of behaviour which repeated throughout the year, either on an annual basis, such as when Christmas has an impact on the time series data being collected, or on a more frequent basis. The trend behaviour, also, was described in relation to cyclical behaviour, 2.2 THE SPECTRUM OF A TIME SERIES *continued*

consisting of the long-term behaviour of the series and the medium-term business cycles.

Before going on to discuss the complete cyclical breakdown of a series, it is important to clarify some terminology. The term 'cycle' will be used broadly to refer to a repeating pattern of behaviour, while the term 'cycle length' refers to the length of time it takes for the pattern of behaviour to complete exactly once. Additionally, the term 'frequency' will be used to describe the number of times a cycle repeats in a given span of time. It should be noted that a seasonal cycle of 12 months length recurs with a frequency of once a year, whereas a cycle of length 6 months occurs with a higher frequency of twice a year. On the other hand, a relatively long cycle, such as a business cycles, say roughly 5 years long, occurs with a lower frequency of 1/5 per year. Thus the terms 'cycles' and 'frequencies' are complementary, and will often be used interchangeably throughout this paper: a 'long cycle' is equivalent to a 'low frequency', and a 'short cycle' is equivalent to a 'high frequency'.

Throughout this paper, the decomposition of a time series into cycles will be considered through the use of a particular diagnostic tool called the spectrum. Broadly speaking, the spectrum of a time series indicates graphically what proportion of series variation, or behaviour, is attributable to particular cycles in the series (see graph 2.1 for an example). The horizontal axis of a spectrum portrays all the possible cycle lengths observable in the time series. If the series is collected regularly every period, whether that be every month, quarter, year, or some other regular period of time, then the range of these cycles extends from two periods long at the far right hand end to cycles approaching an infinitely long duration at the far left hand end. The vertical axis of the spectrum then reflects the contribution of each cycle, in the above range, to the variance of a particular series. This just means that the greater the 'height' a particular cycle scores on the vertical axis compared with some other cycle, the greater is its contribution to the variation in the series. This measure will be referred to as the cycle's 'contribution' or 'power'. In a sense it represents the dominance of a particular cycle or wave in a series. The total area under the spectrum graph then represents the degree of variation in a series, the greater the area, the greater the variability in the way the series behaves.

In graph 2.1, two hypothetical spectra are given. Inspection of the spectrum for series A shows that series A has a spike (i.e. displays more power) at cycles about 12 periods, than at any other range of cycles in it. On the other hand, the power of cycles in series B increases as the cycles pass from long to shorter length, that is, from low to higher frequencies. Comparing the two spectra, it can be seen that the spectrum for series B has more area under its graph overall than does the spectrum for series A, indicating that series B has the more variable behaviour of the two series.

2.2 THE SPECTRUM OF A TIME SERIES *continued*

GRAPH 2.1 STYLISED EXAMPLES OF SPECTRA

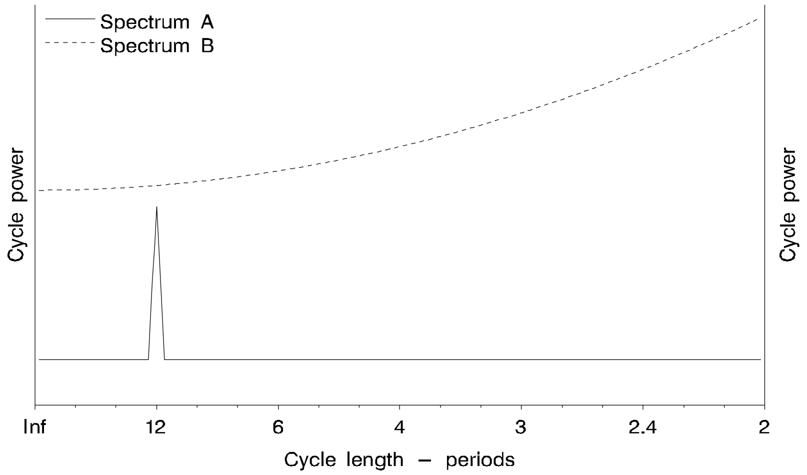

The horizontal axis in the above graph is broken up into six equal sized segments, marked off from the left by cycles of length 12, 6, 4, 3, 2.4 and 2 periods, respectively. This scale may seem a little unusual, however it is related directly to the notion of cycle frequency being complementary to cycle length. A cycle that has a length of 12 periods, for instance, is a cycle that occurs once each 12 periods, so that 1/12th of the cycle is covered over a period. For a cycle that occurs every 6 periods, 1/6th (or 2/12ths) of the cycle is covered over a period, and so on. Thus the frequency of a cycle is given by 1/(cycle length). Consequently, the values marked off on the horizontal axis from left to right represent cycles that occur with frequency 1/12th over a period, 2/12ths over a period, 3/12ths over a period and so on. Therefore, from a frequency perspective, the segments marked off represent equal distances. Cycle lengths, rather than the corresponding frequencies, are used to label the horizontal axis, because cycle lengths are easier to deal with, especially in relation to monthly data.

Examination of many spectra has shown that when seasonal cycles are present in monthly data, they are visible in the spectrum as peaks at any or all of the following cycle lengths, 12 months, 6 months, 4 months, 3 months, 2.4 months and 2 months, corresponding to cycles occurring 1, 2, 3, 4, 5, or 6 times a year, respectively (although even when such a peak is present in the spectrum of a series, the cyclical pattern of behaviour may not be readily observable in the original data). The other systematic calendar related effects (trading day, payday and moving holiday) do not exhibit such a consistent pattern of behaviour, although generally, the majority of the spectral peaks associated with these effects are concentrated in the higher frequency cycle range, particularly in the range of 2 to 6 months. Graph 2.2 shows the spectral estimates of the different components of Australian total adult male employed series.

2.2 THE SPECTRUM OF A TIME SERIES *continued*

GRAPH 2.2 THE SPECTRAL ESTIMATES OF AUSTRALIAN ADULT MALE EMPLOYED

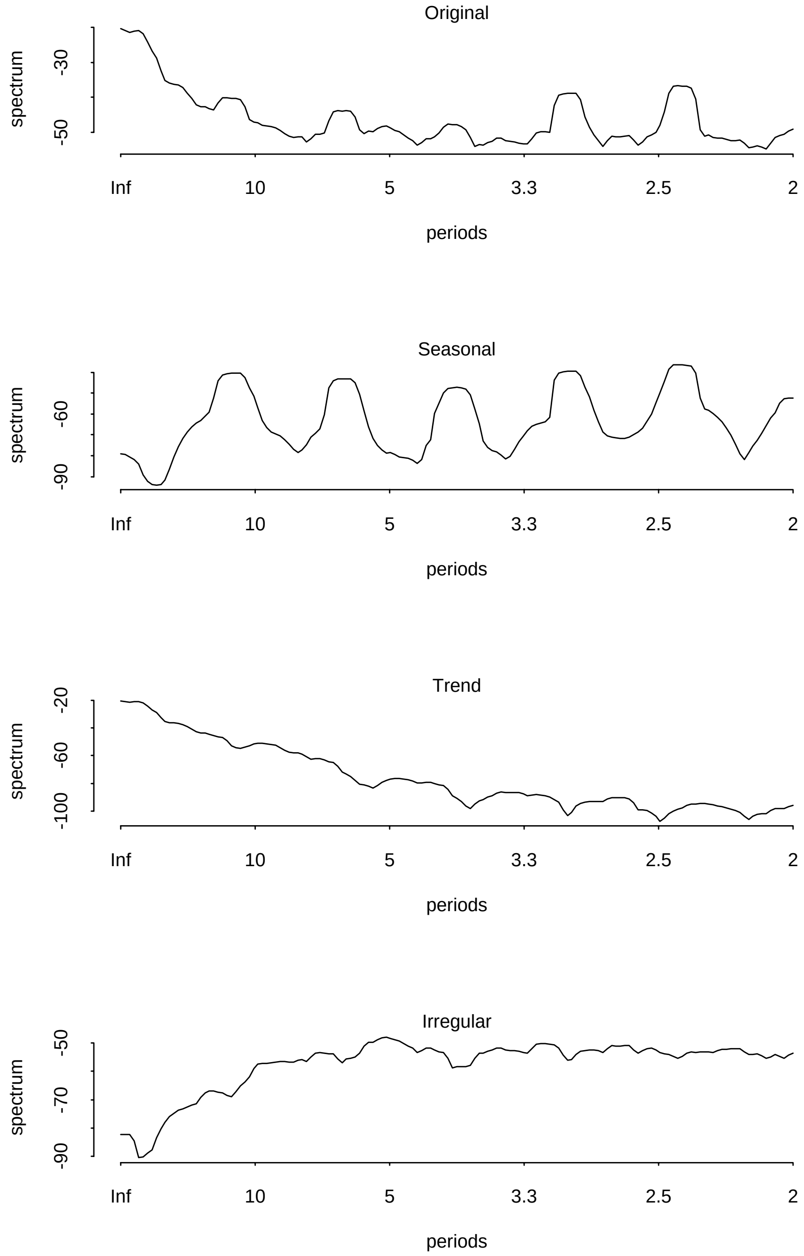

When all the seasonal effects are removed from the original series, the resulting seasonally adjusted series reflects the interaction of just two components of the original series behaviour, the trend, and the irregular fluctuations. The trend contributes to the behaviour of the spectrum around cycles of longer duration, as the trend is the long term underlying behaviour of the series, while the spectral pattern for cycles of shorter duration is dominated by the irregular fluctuations. The nature of this relationship means that considering series behaviour from a cycle structure perspective can be very useful in trying to isolate the trend behaviour of the original time series, and this relationship will be referred to frequently throughout the remainder of the paper.

CHAPTER **3**

COMMON 'NAIVE' METHODS OF SMOOTHING TIME SERIES

CONTENTS

page

SECTION

3.0 OVERVIEW

An important feature of time series is that they allow analysts to understand changes in socioeconomic activity in time. The seasonally adjusted series has had the seasonal effects removed and this helps analysts isolate the underlying behaviour of the series. However the irregular fluctuations are still present in the seasonally adjusted series and smoothing mechanisms are frequently applied to determine the underlying behaviour. Unfortunately, not all smoothing mechanisms are equally useful and some of those commonly applied actually increase the emphasis on the irregular component of series behaviour. Other indicators actually give a distorted impression of turning points in underlying activity levels, or delay their detection.

In this Chapter, some of the most commonly applied 'naive' smoothing mechanisms of time series are evaluated and some of the problems associated with them discussed. In particular, the approaches are evaluated with regard to how well they capture the underlying component of series behaviour that represents the impact of non-seasonal and non-irregular influences upon the series (this is the component of series behaviour the ABS trending procedure attempts to estimate, in order to produce an indicator of underlying socioeconomic activity). The indicators that will be considered are:

- period to period growth
- year apart growth
- moving annual sums and averages
- annual growth rates
- growth in the three months
- annualised growth rates

These forms of analysis, as commonly applied, have three major disadvantages. They can delay the changes in the behaviour of the series (phase shift), they can distort the underlying shape of the series, and some of them can actually amplify the irregularity and volatility in the series. Throughout the chapter, these deficiencies will be illustrated through the use of a typical series of data collected by the ABS.

Table 3.1 compares the performance of various popular forms of analysis at pinpointing the turning points of a major economic indicator. It reports the timing of selected turning points for Unemployed Persons as indicated by the other various possible forms of analysis relative to the turning points apparent in the trend series published by the ABS. For example, the turning point detected as occurring between June and July 1995 in the trend series, is indicated as occurring between March and April 1996 according to the yearly percentage change of the seasonally adjusted series and so on. The table containing the series and the formulae for various of common smoothing methods on which table 3.1 is based is contained in Appendix 1. The monthly percentage change of seasonally adjusted movements series which appears in column 3 of Appendix 1, has not been included in this table. This is because the constantly changing direction of the movements in this series make it difficult to detect when, if ever, the series identifies the turning points detected by the other methods.

3.0 OVERVIEW *continued*

TABLE 3.1 COMPARISONS OF VARIOUS FORMS OF ANALYSIS AT LOCATING TURNING POINTS OF THE TREND

Methods	Turning Points	Delay
1. Monthly percentage change of trend	Oct/Nov 1989	na
2. Yearly percentage change of seasonally adjusted	Feb/Mar 1990	4 months
3. Monthly percentage change of moving annual sum	Feb/Mar 1990	4 months
4. Yearly percentage change of moving annual sum	Jul/Aug 1990	9 months
5. Quarterly percentage change of 3-month	Jan/Feb 1990	3 months
6. Yearly percentage change of 3-month	Mar/Apr 1990	5 months
7. 6-month smoothed percentage change annualised	Dec/Jan 1990	2 months
8. Annualised 3-month to previous 3-months	Jan/Feb 1990	3 months
9. Annualised 6-month to previous 6-months	Feb/Mar 1990	4 months
10. Annualised half yearly change seasonally adjusted 3-month	Jan/Feb 1990	3 months
1. Monthly percentage change of trend	Jun/Jul 1995	na
2. Yearly percentage change of seasonally adjusted	Mar/Apr 1996	9 months
3. Monthly percentage change of moving annual sum	Mar/Apr 1996	9 months
4. Yearly percentage change of moving annual sum	Aug/Sep 1996	14 months
5. Quarterly percentage change of 3-month	Aug/Sep 1995	2 months
6. Yearly percentage change of 3-month	Mar/Apr 1996	9 months
7. 6-month smoothed percentage change annualised	Sep/Oct 1995	3 months
8. Annualised 3-month to previous 3-months	Aug/Sep 1995	2 months
9. Annualised 6-month to previous 6-months	Jan/Feb 1996	7 months
10. Annualised half yearly change seasonally adjusted 3-month	Oct/Nov 1995	4 months
1. Monthly percentage change of trend	Apr/May 1997	na
2. Yearly percentage change of seasonally adjusted	Jul/Aug 1997	3 months
3. Monthly percentage change of moving annual sum	Jul/Aug 1997	3 months
4. Yearly percentage change of moving annual sum	Jan/Feb 1998	9 months
5. Quarterly percentage change of 3-month	May/Jun 1997	1 month
6. Yearly percentage change of 3-month	Sep/Oct 1997	5 months
7. 6-month smoothed percentage change annualised	Jul/Aug 1997	3 months
8. Annualised 3-month to previous 3-months	May/Jun1997	1 month
9. Annualised 6-month to previous 6-months	Aug/Sep 1997	4 months
10. Annualised half yearly change seasonally adjusted 3-month	Jul/Aug 1997	3 months

The ABS trend estimates, once historically finalised, accurately identify the turning points as can be seen in graph 3.1, with both the seasonally adjusted and ABS trend series for the Unemployed Persons series. The method used by the ABS to obtain trend estimates in Chapter 4. It does have some limitations and these will be discussed in detail in Chapters 5 and 6. The remainder of Chapter 3 elaborates on the problems associated with a variety of commonly used naive methods for deriving indicators of the underlying behaviour of a time series.

3.0 OVERVIEW *continued*

GRAPH 3.1 ABS TREND INDICATOR FOR UNEMPLOYED PERSONS

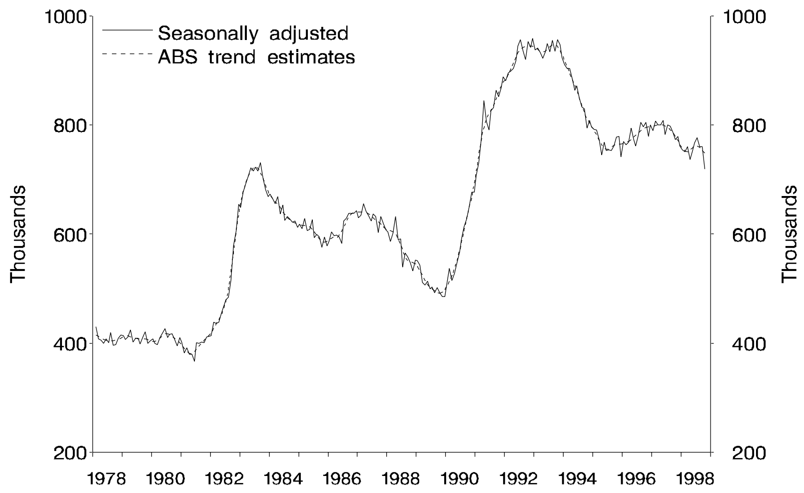

3.1 PERIOD TO PERIOD MOVEMENTS WITHOUT SMOOTHING

Analysts often focus on the period to period movements of the seasonally adjusted series to determine, in general, or at a specific time, whether socioeconomic forces are resulting in an increase in activity, a decline, or whether there is a change in behaviour. However, the problem with using this method to determine the underlying impact of socioeconomic forces lies in the fact that the seasonally adjusted series incorporates not just the underlying trend component of series behaviour, but also the irregular fluctuations. The period to period movements of the seasonally adjusted series turn out to be heavily dependent on the behaviour of these irregular fluctuations, much more so, in fact, than on the underlying behaviour of the series. To illustrate the nature of this problem, it is necessary to introduce the notion of the first-difference operator and investigate the impact of this operator on the seasonally adjusted series.

The first-difference operator, Δ ,is the mathematical operator used to generate period to period movements in a series (there are other difference operators used to generate longer term series movements). If SA_t denotes the seasonally adjusted series at time t, then the period to period movement in SA_t at time t is denoted by, ΔSA_t that is

$$\Delta SA_t = SA_t - SA_{t-1} \qquad (3.1)$$

Thus the effect of applying the first-difference operator to the series, or first-differencing the series, is just to look at the change in consecutive series values. Unfortunately, the first-difference operator amplifies any short-term irregularity, or volatility, in a time series.

As previously discussed, any time series can not only be considered to be a collection of values over time, it can also be thought of as a collection of cycles, of varying strengths, across a wide range of frequencies. Consequently, the effect of applying the first-difference operator to a series can be considered by looking at the effect of applying this operator to all possible cycles that may be observed to occur within the series. For example, consider the effect the first-difference operator has upon a cycle which repeats every two periods. Graph 3.2 shows a two period cycle and also plots the values which result when the period to period movements of this cycle are calculated. The graph

3.1 PERIOD TO PERIOD MOVEMENTS WITHOUT SMOOTHING *continued*

shows that the effect of taking the first-difference of a series of data values forming a cycle of length two periods, is to produce a more variable series, the amplitude of the resulting cycle being twice the amplitude of the original cycle. Similarly, consider the impact of the first-difference operator on a cycle of length twelve periods (a seasonal cycle in monthly data). Graph 3.3 shows that this time, the differenced values form a cycle with less strength than the original cycle, the amplitude of the cycle in fact having been halved after the process of taking the first-difference.

GRAPH 3.2 EFFECT OF FIRST-DIFFERENCING ON TWO PERIOD CYCLES

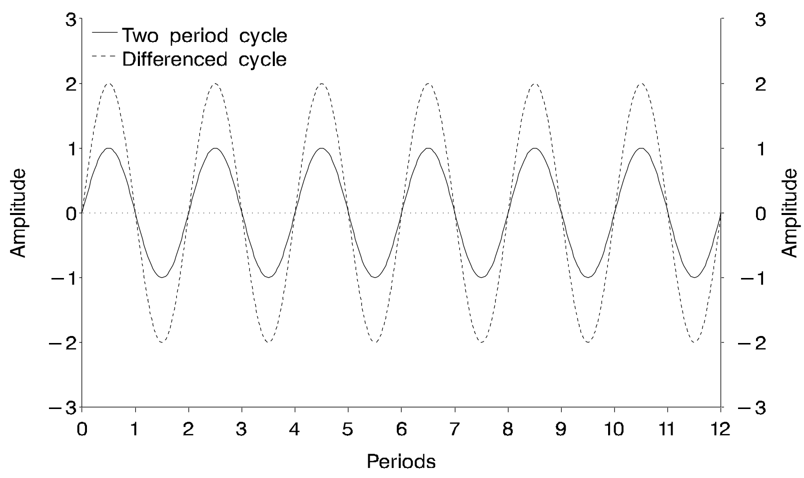

GRAPH 3.3 EFFECT OF FIRST-DIFFERENCING ON TWELVE PERIOD CYCLES

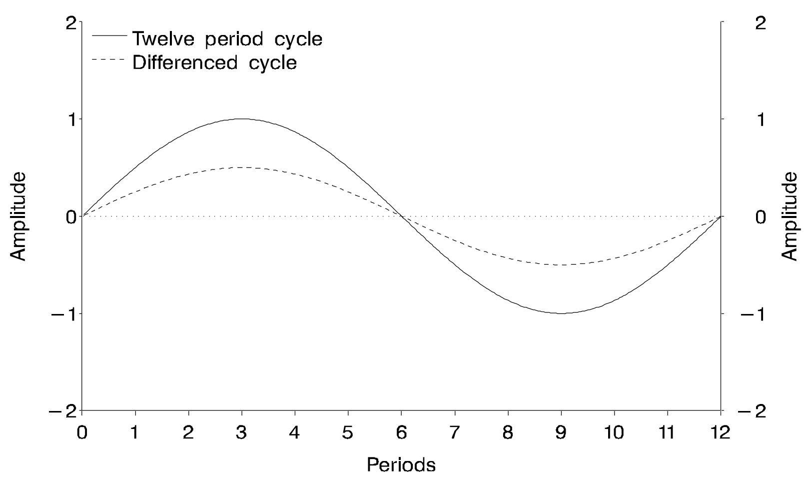

If a similar process was to be undertaken for cycles of all possible frequencies, then the effect on the strength of cycles, as measured by amplitude, would be as depicted in graph 3.4. Graph 3.4 shows to what extent the amplitudes of various cycles within a series are altered when the series is first-differenced. The vertical axis indicates the

3.1 PERIOD TO PERIOD MOVEMENTS WITHOUT SMOOTHING *continued*

percentage of a particular cycle's amplitude which remains after the first difference operator has been applied to the data. The graph shows that the only cycle with unaltered amplitude is the cycle that occurs every six periods. Cycles of shorter length have their amplitude increased, while those that are longer have their amplitude decreased. In other words cycles of shorter duration have their influence increased, while those of longer duration have their influence reduced.

GRAPH 3.4 EFFECT OF THE FIRST DIFFERENCE OPERATOR ON CYCLES

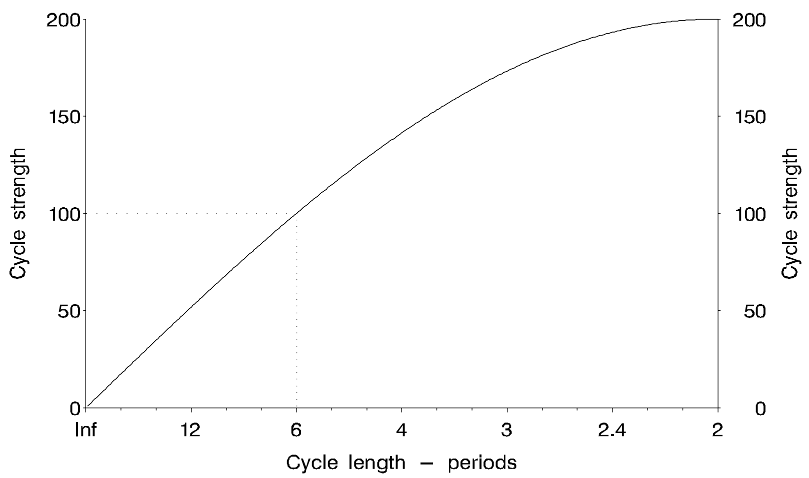

For a seasonally adjusted series, longer length cycles tend to be related to the trend behaviour, while shorter length cycles tend to be related to the irregular fluctuations. Consequently, the information displayed in graph 3.4 indicates that the series of first-differenced values from the seasonally adjusted series is even more dominated by the behaviour of the irregular fluctuations than the seasonally adjusted series itself.

Considering column 3 of Appendix 1, it can be seen that the movements in the seasonally adjusted Unemployed Persons series, expressed as a percentage, are very irregular. These movements change from being positive to negative or vice versa very frequently and as indicators of underlying economic activity, they would imply the direction of the economy was changing direction almost continuously. This can be compared with the percentage movements of the ABS trend series, a series designed to capture the timing of turning points in the underlying behaviour of the series as accurately as possible. This series signals comparatively few turning points in the economy.

Through the above discussion, it can be seen that the question of whether the movements of the seasonally adjusted series are a reasonable indicator of the underlying direction of the series, will depend on the degree of irregularity present in the seasonally adjusted series in the first place. In particular, if period to period movements are to reflect underlying behaviour rather than irregular behaviour, then these movements need to be calculated only from series containing very little irregularity, otherwise the behaviour of the irregular component will completely swamp the behaviour of the trend.

3.1 PERIOD TO PERIOD MOVEMENTS WITHOUT SMOOTHING *continued*

In real time series, the period to period movement of the irregulars often swamp/dominate changes in the underlying (trend) behaviour of a series. The following paragraphs explain (in non-technical terms) how to assess this.

There are various methods available for determining the contribution of the irregular fluctuations to the movements in a seasonally adjusted series. One approach is to calculate an average which gives the percentage contributed by the irregular component to the period to period movement as compared to the ABS estimate of trend (in which a significant proportion of the irregularity present in the seasonally adjusted series has been eliminated). The calculation is as follows. First, the amount contributed to the seasonally adjusted series by the irregular component, I_t, must be estimated. This is achieved by subtracting the series of ABS trend estimates, T_t, from the series of seasonally adjusted values, SA_t (that is $I_t = SA_t - T_t$). The period to period changes in this series, I_t, and the ABS trend series, T_t, are then computed (ie. ΔI_t and ΔT_t respectively). While these changes may be positive or negative from period to period their absolute values are used ($|\Delta I_t|$ and $|\Delta T_t|$ respectively). The contribution of the irregular fluctuations to a particular seasonally adjusted movement is then given by:

$$\frac{|\Delta I_t|}{|\Delta I_t| + |\Delta T_t|} \times 100 \qquad (3.2)$$

This calculation can be repeated for each period to period movement in a specified span of series, and the frequency with which the above measure exceeds 50 per cent (or any other value chosen by the analyst) may be noted. For instance, if this procedure is applied to the seasonally adjusted exports of goods and services data given in table 3.2, then the percentage contribution of the irregular fluctuations to each of the eleven monthly movements from February 1998 to December 1998 is: 70.6, 69.0, 45.4, 39.6, 48.4, 82.0, 98.1, 78.2, 66.3, 64.9 and 57.2 respectively. Over this interval of time, it can be seen that the irregular fluctuations have contributed more to the seasonally adjusted period to period movements than has the trend in eight of the eleven cases, or 72.7 per cent of the time. Similar calculations can be made for any time interval or any series.

TABLE 3.2 EXPORTS OF GOODS AND SERVICES ($m), 1998

	Jan	*Feb*	*Mar*	*Apr*	*May*	*Jun*	*Jul*	*Aug*	*Sep*	*Oct*	*Nov*	*Dec*
Components	$m	$m	$m	$m	$m	$m	$m	$m	$m	$m	$m	$m
SA	9 263	9 443	9 198	9 592	9 656	9 908	9 681	10 035	9 719	9 825	9 455	9 502
T	9 472	9 344	9 268	9 483	9 669	9 799	9 863	9 856	9 787	9 677	9 547	9 408

The above formula is designed to calculate the relative amount of irregular contribution to the period to period movements of the seasonally adjusted series. However analysts often focus on the period to period percentage movements of the seasonally adjusted series. A similar type of calculation can be undertaken to determine the contribution to the period to period percentage movements (the detail of this calculation can be found in the *Australian Economic Indicators, March 1992,* (cat. no. 1350.0) Feature Article 'Smarter Data Use'). In most applications, it turns out that the calculation giving the percentage contribution of the irregular fluctuations to a period to period movement in

3.1 PERIOD TO PERIOD MOVEMENTS WITHOUT SMOOTHING *continued*

the seasonally adjusted series, closely approximates the contribution to that period to period percentage movement. In the example above the irregular fluctuations contribution to the percentage movement in the seasonally adjusted Exports of goods and services from June to July 1992 is 81.9%, which is close to the 82% above.

In the ABS experience, many seasonally adjusted series exhibit a high degree of irregularity in their period to period movements, and consequently, are not in their own right useful indicators of trend behaviour. In fact the *Australian Economic Indicators, March 1992,* (cat. no. 1350.0) Feature Article 'Smarter Data Use' illustrated how poorly some of the movements in the seasonally adjusted figures for major economic indicators reflect the trend behaviour. For example, over a five-year period, the number of times the irregular fluctuations contributed more than the trend to the monthly movement in the seasonally adjusted series, is given by the following percentages:

Series	*Irregular contribution* %
Employment	50
Unemployment	70
Unemployment rate	70
Labour force participation rate	90
Motor vehicle registrations	85
Balance on goods and services	90
Exports of goods and services	95
Imports of goods and services	95
Building approvals — dwellings (number)	70
Housing finance by all lenders (number)	85
Retail trade	70
Westpac — Melbourne Institute leading index	60

Consequently, although calculating the period to period movements of the seasonally adjusted series is quick and easy, the values are frequently heavily dependent on the behaviour of short-term irregularities, rather than on the longer term underlying behaviour of the series.

3.2 YEAR APART GROWTH

Users wishing to discern underlying trends often calculate the growth between the current period and the same period a year ago. One feature of this measure is that it performs a crude seasonal adjustment when applied to the original data, because the comparison of like months reduces the impact of the constant yearly seasonal effects (although because of gradually changing seasonal patterns these effects are rarely eliminated). The problem with applying the year apart growth measure to the original series, however, is that it is unable to allow for the effects of trading day, payday or moving holidays. Nor does it cope well with the more commonly encountered form of multiplicative seasonality described in section 2.1.4, that is, when the seasonal effects evolve with the level of the series by maintaining a proportional relationship.

3.2 YEAR APART GROWTH *continued*

Applying the year apart growth measure to the seasonally adjusted series produces a measure which is free of the distortion of the calendar related influences, although it is still influenced by the irregular fluctuations. Whereas the period to period movements of the seasonally adjusted series in large part capture these short-term fluctuations, year apart movements reflect more of a long-term change in series behaviour, allowing the trend to stand out more clearly than the irregular component. This means that year apart growth values reflect more of the trend behaviour than do period to period growth values. The irregular fluctuations still play a significant role in the year apart movements however, and the contribution from the trend to the yearly change might still be less than half.

Even when the year apart growth measure reflects a substantial contribution from the trend, there is still a fundamental problem with the measure, regardless of whether it is applied to the original, seasonally adjusted, or other forms of the data. The problem is that the procedure of calculating year apart growth will rarely detect the timing of turning points in underlying series behaviour accurately, usually delaying their detection by a number of months. This is because using year apart growth to identify the timing of turning points turns out to be a very crude approximation of the standard calculus technique used for identifying the turning points of a path.

Suppose the year apart growth measure is applied to the artificially constructed series given in graph 3.5 (the ABS trend series for Unemployed Persons has followed a similar curvature in the early 1980s and early 1990s). The sign of year apart growth at time t is represented by the slope of the line that joins the point at time t and the point a year in the past, see for example lines A, B, and C. Thus a positive year apart growth value corresponds to a line with increasing slope, while a negative year apart growth value corresponds to a line with decreasing slope. If the year apart growth values were then used to indicate the direction of underlying activity, turning points would be considered to have occurred whenever the year apart growth measure changed sign.

The direction of this curved series would be determined exactly at time t by calculating the gradient of the series at time t; a positive gradient indicating the series is increasing, a negative gradient indicating the series is decreasing. In calculus terminology, the gradient is the slope of a particular straight line passing through the series, called the tangent to the curve at time t (see tangent lines A', B' and C'). It is this tangent line that is being approximated by the chord between points a year apart on the series using the year apart growth measure. As can be seen by comparing the chords A, B, and C, with the corresponding tangents A', B', and C', the slope of the chord is actually a very poor approximation to the slope of the corresponding tangent. In graph 3.5, the slope of the year apart growth measure remains positive many months after the trend peak in March has been passed (A & B), and only becomes negative when the month of December is reached (C), nine months later. Hence, the year apart growth measure would indicate that the turning point of the series occurred in December, rather than the month of March in which it actually occurred.

3.2 YEAR APART GROWTH *continued*

GRAPH 3.5 HYPOTHETICAL TREND CURVE AND YEAR APART GROWTH

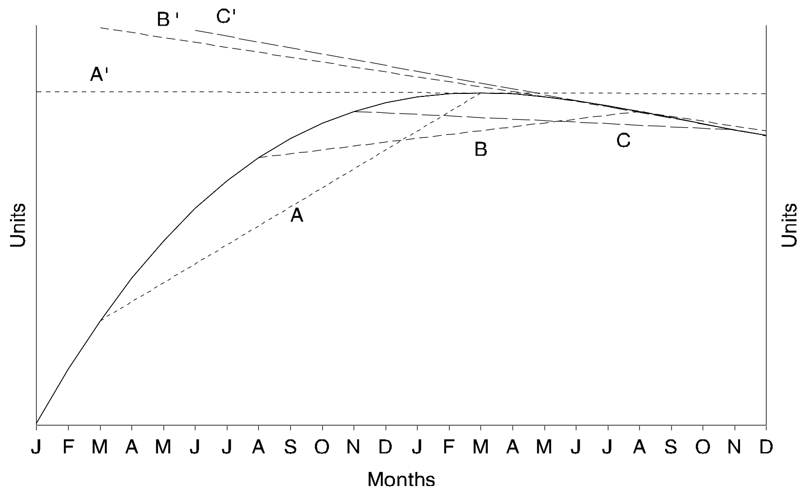

The effect of the year apart growth measure in delaying the detection of turning points can be seen in practice for the Unemployed Persons series given in table 3.1. Compared to the timing of turning points given by the percentage movements of the ABS trend series, the detection of turning points is consistently delayed, and the timing of the turning point misplaced.

In summary, year apart growth is an easy measure to calculate, but is subject to the influences of the irregular component of the series (and the seasonal component if applied to the original series). More crucially, it can delay the detection of turning points, giving a misleading indication of the timing of pivotal changes in underlying series behaviour.

3.3 MOVING ANNUAL SUMS AND AVERAGES

Another procedure commonly used to extract the underlying behaviour of the series is summing or averaging the data over a number of periods. Typically this is done over spans of twelve months, in which case the procedure is known as a moving annual sum. If the sum is divided by the number of observations aggregated, it is termed a moving annual average.

One motivation for using this procedure arises from the intuitive appeal of applying a twelve month averaging or summing process to the original series. This is done in order to even out or diminish the seasonal patterns. However, seasonally adjusted series are published and these are a better guide to the non-seasonal behaviour of the series. As the irregular fluctuations have a tendency to chop and change direction from month to month, application of a moving annual sum or average process will also tend to cancel out the irregular component of series behaviour. Consequently, applying the procedure to the seasonally adjusted series has the effect of diminishing the variation attributable to the irregular fluctuations, thereby disclosing the longer term cyclical behaviour usually associated with trend. The reduction of irregularity is borne out in practice, and moving annual sums and averages are generally much smoother series than the original or seasonally adjusted series from which they are derived. Applying such a procedure is

3.3 MOVING ANNUAL SUMS AND AVERAGES *continued*

consequently often referred to as smoothing the series, the resulting moving annual sum or average values being referred to as the smoothed series values.

The moving annual sum process is usually applied to the series by summing consecutive series values over twelve months and recording the result at the most recent time period of the twelve-month span. Similarly the moving annual average process is applied by averaging this set of consecutive series values. This means that when a new observation becomes available, the value of the moving annual sum (or average) is calculated and recorded against the new time period, the measure consequently appearing to remain 'current'. The result of applying either of these processes is a series of values which reflect more of the trend behaviour than the series from which they were derived. This resulting series can then be graphed in order to identify characteristics of socioeconomic behaviour, such as the fundamental direction in which activity is moving and the timing of turning points. Alternatively, period to period movements or percentage movements of the new series can be calculated, changes in sign indicating a change in the direction of socioeconomic activity.

Although applying a moving annual sum or average process does help to isolate the trend behaviour, particularly when applied to the seasonally adjusted series, there are a number of drawbacks associated with this method. These problems are illustrated in graph 3.6. In this example, the seasonally adjusted Unemployed Persons series has been smoothed by computing a moving annual average, that is, the moving annual sum has been calculated for each month, and then divided by twelve, so as to realign the result with the level of the seasonally adjusted series. Note that the division by twelve does not alter the shape or the timing of the smoothed curve, it just brings it into the same range as the data being smoothed. One deficiency of the procedure which is apparent from the graph is that there are eleven observations missing at the start of the series because of the practice of placing the result at the 'current' end of the moving annual average span. There are also a number of other problems associated with using the moving annual sum or average procedure and these have an impact on the conclusions that can be drawn about the underlying level of activity.

GRAPH 3.6 SIMPLE MOVING AVERAGE BASED TREND ESTIMATES FOR UNEMPLOYED PERSONS

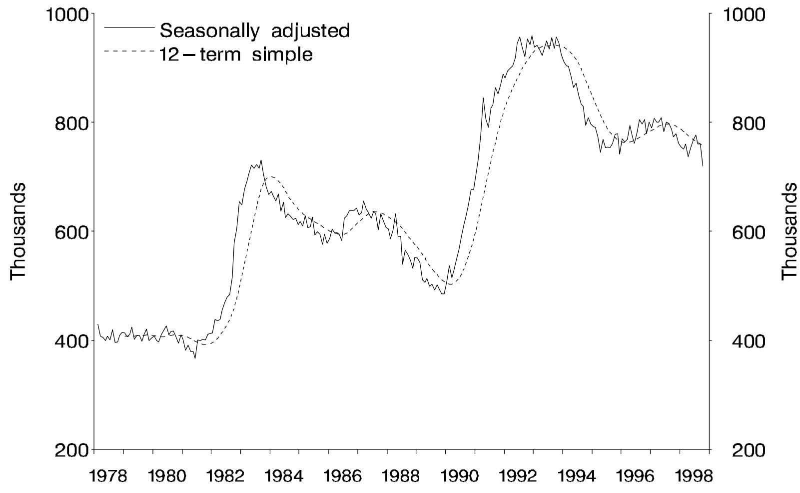

3.3 MOVING ANNUAL SUMS AND AVERAGES *continued*

First, all the turning points are disclosed about a half year after they have occurred. For example, refer to the years 1981, 1983, and 1989–90 of graph 3.6. This aspect is known technically as a 'phase shift'. Paradoxically, this significant delay in detecting turning points arises as a result of placing the results at the current end of the moving annual average span. If the results were placed in the middle of the moving annual average span, the delay in detecting turning points would disappear. This would also mean that only the first half year of results would be unavailable, rather than the first eleven months. However the results for the last half year would also be unavailable. This lack of results at the current end of the series where interest is often keenest, is referred to as the 'end-point problem', and will be discussed in Chapter 5.

Secondly, even if the phase shift of the smoother moving annual average (or sum) is corrected, on closer inspection it can be seen that the levels of the peaks and troughs are underestimated. This feature arises because this smoothing procedure is unable to track any path other than a straight line — it can only approximate the curves associated with turning points, and does so by skirting around the inside of them. Another deficiency associated with this inability, is the propensity of the moving annual average to distort the shape of the turning points, making them broader than they should be. This deficiency leads to misleading information about how quickly the economy is, for example, growing or contracting in response to policies, business decisions or other factors.

Finally, there is a problem when trying to detect points of inflection in underlying series behaviour. Points of inflection occur when the rate at which the level of activity is increasing slows down, and then at some point in time picks up again, without the level of activity itself ever actually showing a decrease (and vice versa when the level of activity is decreasing). Such a point is said to be stationary if it actually reflects zero change in the level of activity, otherwise it is said to be non-stationary, refer to graphs 3.7 and 3.8 for examples. It may be important to monitor points of inflection when attempting to determine whether business decisions, or policy actions, are moderating or stimulating activity sufficiently to achieve their goals. Moving annual sums and averages rarely detect points of inflection, for example, refer to the behaviour of the Unemployed Persons and moving annual average series in graph 3.6, around the year of 1985.

3.3 MOVING ANNUAL SUMS AND AVERAGES *continued*

GRAPH 3.7 STATIONARY POINT OF INFLECTION

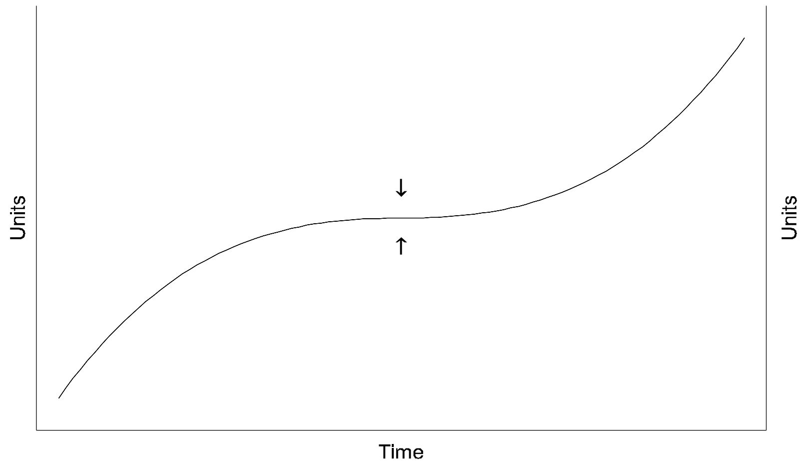

GRAPH 3.8 NON-STATIONARY POINT OF INFLECTION

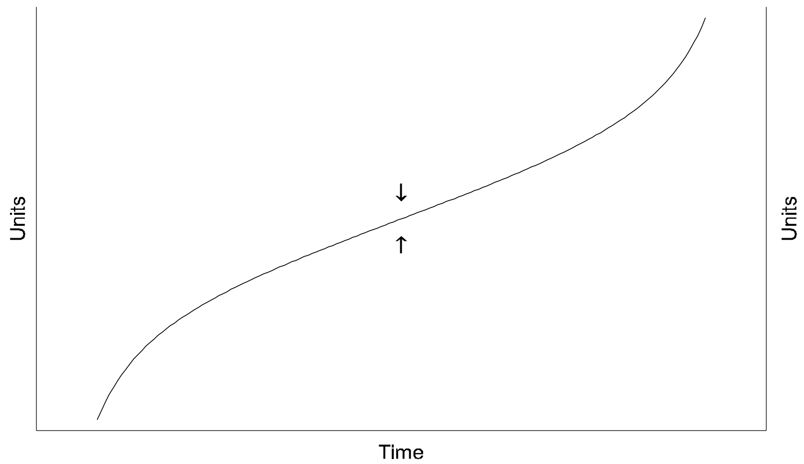

These problems are in fact a feature of the application of the above type of moving annual sum or averaging process to any series. Consequently, although the moving annual sum and average, as they are commonly used, appear to extract the underlying behaviour of the series, there are a number of major deficiencies which are invariably associated with using these methods:

- they delay the detection of turning points (phase shift)
- they underestimate the height of peaks and the depth of troughs
- they distort the shape of turning points in the series and, as a consequence, they extend the period over which a peak or trough appears to exist
- they flatten out points of inflection in the series, often resulting in their elimination.

3.3 MOVING ANNUAL SUMS AND AVERAGES *continued*

The simple sums and averages described here represents just one example of a class of procedures known as moving sums and moving averages. The ABS uses a different class of moving averages, called Henderson moving averages, to calculate its trend estimates, and the properties of moving averages in general and the Henderson moving averages in particular, will be described in detail in Chapter 4.

3.4 ANNUAL GROWTH RATES

This procedure involves computing the percentage growth between the most recently available twelve-month period and the corresponding twelve-month period a year ago. This is equivalent to amalgamating two of the measures previously discussed. First, a moving annual sum is being applied to calculate a series of values, each of which is the sum of twelve consecutive months of data. Then, the changes in behaviour of this series are being determined by applying the year apart growth measure (expressed as a percentage), that is, by comparing the difference between values one year apart in this new series.

As the annual growth rate measure is equivalent to amalgamating the year apart growth and moving annual sum procedures, it is not surprising that the deficiencies of both these procedures are compounded in the annual growth rate measure. Not only is the trend path (as represented by the moving annual sum) distorted, but the detection of turning points can also be delayed by around a year or more. The delay in detecting the turning points for the Unemployed Persons series can be seen in table 3.1. Another problem with this measure can be seen in Appendix 1 on which table 3.1 is based, namely, that there is a loss of data for about two years at the beginning of the series.

3.5 GROWTH IN THE THREE MONTHS

A variant of the above procedure involves comparing the latest three months of accumulated data, with either the immediately preceding three months, or the corresponding three months a year ago. These procedures are very similar to the annual growth rates described above, and consequently the problems associated with applying these processes are similar (although less marked). In particular, these methods can give a distorted impression of trend paths and delay the detection of turning points. Table 3.1 illustrates the deficiencies of these three-monthly growth measures with regard to detecting turning points.

3.6 ANNUALISED

Several annualised procedures exist, and each indicates what the annual percentage growth rate would be if the percentage growth during some relatively recent period was to compound or repeat over a year. One attraction of annualising data is that it transforms the latest data to the same order of magnitude as genuine annual growth rates. This enables comparisons to be made with previous time periods, or overseas data, when the annual summary data may be all that is available. The problem with annualised growth rates, however, is that they amplify the contribution of any seasonal or irregular effects present in the data being annualised, and rarely indicate reliably the annual growth rate that actually takes place during the corresponding reference period.

In this section, some of the problems associated with annualised measures will be demonstrated by considering two of the more prominent measures often employed. These are: annualised quarterly growth rate measures, and the six month smoothed percentage change at an annual rate measure.

3.6.1 Annualised quarterly growth rates

Annualised quarterly growth rates are obtained by assuming the observed growth will continue to compound without variation for the next three quarters. Thus the annualised value is obtained by raising the ratio of quarterly growth to the power of four. That is:

$$AQGR = \left(\frac{Q_t}{Q_{t-1}}\right)^4 \quad (3.3)$$

Graph 3.9 displays the results of applying this procedure to seasonally adjusted Gross Domestic Profit, GDP.

GRAPH 3.9 PERCENTAGE GROWTHS OF SEASONALLY ADJUSTED DOMESTIC PROFIT GDP

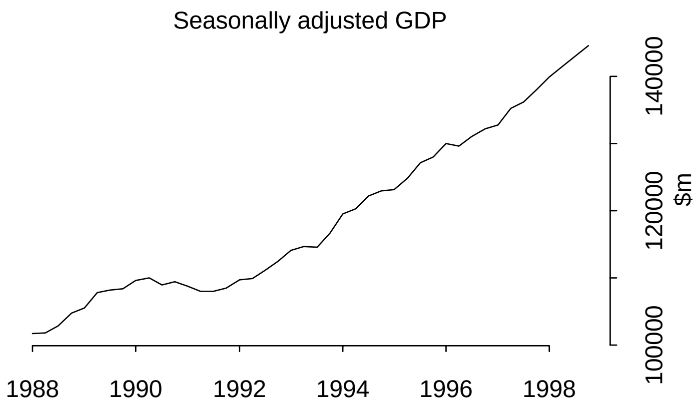

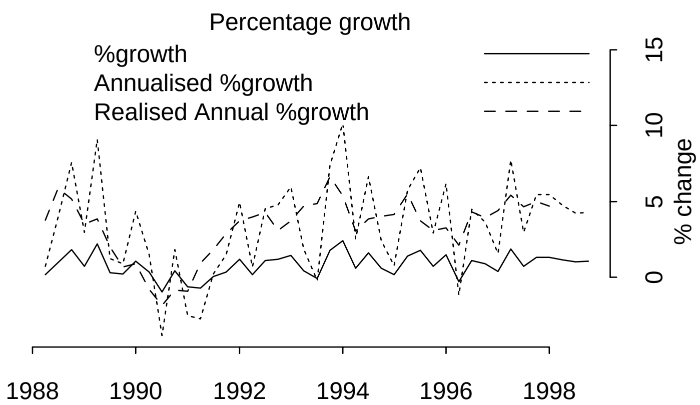

As can be seen by comparing '% growth' and 'Annualised % growth' of graph 3.9, applying this method to seasonally adjusted series can result in quite misleading estimates of the annual percentage growth. This is because this measure reflects the presence of the irregular fluctuations, whenever the trend's contribution to the quarterly growth is minor. Furthermore, since the ratio reflecting the quarterly growth is raised to the power of four, the contribution from the irregular component is amplified. Given that period to period growth will reflect even more of the irregular behaviour than the seasonally adjusted series itself (see section 3.1), the irregular fluctuations will play a major role in

3.6.1 Annualised quarterly growth rates continued

the annualised quarterly growth measure for any quarterly series with a significant degree of irregularity.

The problems associated with applying this indicator of underlying behaviour to seasonally adjusted series would be further compounded if it were to be applied to the original series since the seasonal component of series behaviour would also be amplified. This is because the ratio reflecting the quarterly growth is raised to the power of four, and therefore, the seasonal contribution to that growth would also be raised to the power of four. As a result, a significant part of each annualised value would reflect the seasonal behaviour of the original series. Furthermore, the annualised values would tend to oscillate between positive and negative growth, due to the oscillatory nature of the seasonal effects. Applying the method to original data therefore makes little sense.

The above discussion indicates that quarterly annualised growth is not a very appropriate measure of annual change in underlying series behaviour. Not only does it inflate the contribution of the irregular component of series behaviour (and the seasonal component if applied to the original series) it also assumes a constant rate of growth in the trend, an assumption which is often not valid. Where there is a need for annualised data, for example to obtain information of the same order of magnitude for comparison purposes, it is recommended that the annualised quarterly growth rate of the trend series be used, so that the underlying behaviour of the series is being examined.

3.6.2 Six month smoothed percentage change at an annual rate

The six month smoothed percentage change at an annual rate measure is another annualising procedure often applied to seasonally adjusted series in particular, in order to monitor the underlying behaviour. It is a far more complicated measure than the annualised quarterly growth measure and incorporates both a smoothing component and an annualising component. As a result of this combination of techniques, the measure is also often referred to as the smoothed annualised growth rate.

The six month smoothed percentage change at an annual rate measure is determined by first calculating the rate of growth between the current seasonally adjusted value and a smoothed value based on the preceding twelve months of data. The annualised value is then calculated by determining the amount of percentage growth represented when this ratio of growth compounds over a year. To compute the annualised measure for the month of January 1998, for example, first a smoothed value is calculated by taking the average of the previous twelve seasonally adjusted figures, that is, January to December 1997. This gives a smoothed estimate which is considered to reflect the underlying behaviour of the series exactly halfway through the twelve month span, that is, mid June 1997. Growth from this time (mid June 1997) to January 1998 is then calculated by looking at the ratio of the January seasonally adjusted value to the smoothed mid June value, a growth measure which reflects the change in behaviour over 6.5 months. An annualised value for the twelve months is then determined by calculating the percentage change represented by raising this ratio of growth to the power of 12/6.5. The result is then recorded against January 1998.

The ratio used as a basis for annualising growth contains a seasonally adjusted series value in the numerator, and a smoothed value which is an approximation to the trend in the denominator. This means that when the ratio of growth is raised to the power of 12/6.5, the contribution from the irregular component, in the numerator only, will be

3.6.2 Six month smoothed percentage change at an annual rate continued

nearly squared. Consequently, the irregularity present in the series will have an even more dominant role in the annualised growth values than in the series itself, directly influencing the growth implied by the annualised measure.

Suppose, for example, that at a particular point in time, irregularities contribute an extra 5% above the trend in the seasonally adjusted series. This can be represented symbolically as $SA = 1.05T$. The 5% contribution from the irregular component will then be almost squared in the annualised measure. As $1.05^{12/6.5} = 1.0943$, approximately 9% of the annualised growth value will actually be a direct result of the irregularity present in the series, rather than a true reflection of the change in the trend. In general, the fact that the irregular fluctuations chop and change direction from month to month, means that the annualised growth measure will oscillate, largely due to the magnified irregular variation, and the trend changes will be harder to detect.

A further concern with this particular procedure, is that it does not accurately identify the timing of turning points. According to Boehm, 1989[1]

'... in comparison with the business cycle turning points timing and in forecasting turning points, the smoothed growth rate needs to be moved three months back in order to centre it, on statistical grounds, with the corresponding turning points.'

In practice, the procedure has been observed by the ABS to delay the detection of turning points by between two and five months, depending upon the nature of the series to which the procedure is applied. The effect for the Unemployed Persons series can be seen in table 3.1. Here, a comparison of the six month smoothed percentage change at an annual rate in column 9 with the percentage movements in the ABS trend series (column 2).

Other complex ways of annualising series have also been developed, although by and large, they do not appear to be useful for understanding underlying series behaviour. In particular, when growth rates of seasonally adjusted series are annualised, the process can magnify the irregularity inherent in the series, as well as provide delayed information about turning points. Table 3.1 indicates the delays these methods introduce in the detection of turning points.

3.7 BENCHMARK SERIES FOR COMPARISON

The measures that have been discussed throughout this section are indicative of the types of methods commonly applied to series in an attempt to extract information on the underlying impact of socioeconomic forces. While some of these measures are quite easy to calculate, or have a degree of intuitive appeal, they all have some deficiencies in regard to the conclusions that can be drawn about the behaviour of the series. Some of the measures discussed have been seen to be strongly influenced by the irregular component of series behaviour (and the seasonal component when applied to the original data). Other measures have been seen to give a distorted impression of turning points and points of inflection, or to delay the timing of turning point detection.

There are many other indicators commonly applied to seasonally adjusted series in particular, in order to extract information on underlying behaviour. The reader may be aware of some of these other transformations, and be interested in comparing results of these transformations with those of the ABS trend procedure. To permit a comparison of

1 'The Usefulness and Applications of Economic Indicator Analysis', Boehm, Ernst A., (Paper presented to the Economic Society of Australia Annual Conference, 10 – 13 July 1989; page 25, footnote 15).

3.7 BENCHMARK SERIES FOR COMPARISON *continued*

performance of the alternative transformation in pinpointing turning point timing, level and shape, a benchmark seasonally adjusted and trend series is provided in Appendix 1.

The means by which the ABS trend estimates circumvent the types of problems described above will be covered in the following section. Chapter 4 describes in detail the manner in which the ABS trend estimates are derived, and the properties of those estimates.

CHAPTER 4 ABS TREND ESTIMATES

CONTENTS

4.0 OVERVIEW

Chapter 2 introduced the notion of the trend of a series, that is, the underlying component of series behaviour that remains when the seasonal and irregular influences are removed. Chapter 3 then evaluated some of the most commonly applied indicators of socioeconomic activity with regard to how well they captured this component of behaviour, and briefly discussed some of the problems associated with their use. In order to avoid the types of problems encountered when using these indicators, and to assist analysts in their search for underlying series behaviour, the ABS calculates an estimate of the trend behaviour of a series for the majority of its major economic indicators. This Chapter discusses the derivation and properties of the ABS trend series.

The seasonally adjusted series is determined by quantifying the impact of the seasonal influences on the original series and then removing this component of series behaviour. Therefore, in principle, the trend behaviour of the series could be revealed by similarly quantifying the impact of irregular influences on the series, and removing this component of behaviour from the seasonally adjusted series. Major shocks to a series can often be identified and removed in this way, however, in general, the relevant information required to both identify and quantify other irregular influences is not available, and consequently this approach is not operationally feasible.

The ABS uses a filtering procedure, based on moving averages, to remove the irregular fluctuations from the seasonally adjusted series, leaving the signal or trend. This Chapter discusses the class of moving averages used by the ABS to undertake this filtering, some of their properties, and the motivation behind using this particular group of moving averages. In particular, section 4.1 discusses moving averages and their properties in general, while section 4.2 describes the particular characteristics of the moving averages used by the ABS. Other statistical procedures for filtering series are not discussed in detail in this paper, although some alternatives are briefly explored in section 4.3.

4.1 PROPERTIES OF MOVING AVERAGES

Throughout this paper, the term average will be used to refer only to various types of arithmetic averages, unless otherwise indicated. The most familiar of these arithmetic averages is the simple average calculated by adding a set of data values together, then dividing by the number of values in that set. In this case, all the values contribute equally to the average being calculated, that is, they have an equal amount of weight in the calculation. For example, a simple average of 13 observations would be calculated by adding together 1/13th of each of the observations, so that each value contributes with a weight of 1/13. An average can also be calculated by allowing different observations to contribute to the calculation with differing weights. For instance, an average applied to 13 observations can be calculated by adding 1/24th of the first and thirteenth observations to 1/12th of the intervening 11 observations. Indeed an average for 13 observations, or a 13-term average, can be calculated using as weights any sequence of 13 proportions which sum to one.

In relation to time series data, an average may be applied to the whole time series to produce a single averaged value, or an average may be applied to successive sub-periods of the series, thereby producing a series of averaged values. In the latter case, the results are said to have been produced from a moving average. For example, when applying a 13-term moving average to a time series, the first step is to calculate the average of the first 13 consecutive time series observations. The next step is to move along the series one observation, dropping out of the calculation span the first observation and bringing

4.1 PROPERTIES OF MOVING AVERAGES *continued*

in the fourteenth observation, and calculate the average of these 13 observations. The process of moving along the time series one observation at a time and taking the 13-term average is then repeated, until there are no further time series observations to bring into the 13-term span.

A moving average can be of any length and can take on any weighting pattern in which the weights sum to one. As different weighting patterns[1] tend to give rise to moving averages with different characteristics, moving averages are often classified on the basis of their associated weighting patterns. For instance, when all observations are given the same weight, as when undertaking a simple averaging process, the moving average is said to be a simple moving average, otherwise the moving average is said to be non-simple. Thus a 13-term moving average with weighting pattern

Observation	*...*	*1st*	*2nd*	*3rd*	*4th*	*5th*	*6th*	*7th*	*8th*	*9th*	*10th*	*11th*	*12th*	*13th*	*...*
Weight	... 0	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	1/13	0 ...

is a simple moving average, whereas a moving average with the weighting pattern

Observation	*...*	*1st*	*2nd*	*3rd*	*4th*	*5th*	*6th*	*7th*	*8th*	*9th*	*10th*	*11th*	*12th*	*13th*	*...*
Weight	... 0	1/24	1/12	1/12	1/12	1/12	1/12	1/12	1/12	1/12	1/12	1/12	1/12	1/24	... 0

as mentioned earlier, is a non-simple moving average.

A moving average may also be described as being either symmetric or non-symmetric, depending on the form of its weighting pattern. In particular, a moving average is said to be symmetric if the weighting pattern that is used to calculate it is symmetric about the centre of the averaging span, otherwise it is said to be non-symmetric. All simple moving averages are symmetric, although other symmetric moving averages do exist. Consider the weighting pattern

Observation	*...*	*1st*	*2nd*	*3rd*	*4th*	*5th*	*6th*	*7th*	*8th*	*9th*	*10th*	*11th*	*12th*	*13th*	*...*
Weight	... 0	–0.02	–0.03	0.00	0.07	0.15	0.21	0.24	0.21	0.15	0.07	0.00	–0.03	–0.02	... 0

which, although more complicated than those previously mentioned, still just specifies what portion of each observation is to be added into the moving average calculation. This weighting pattern is symmetric about the centre of the 13-term averaging span, as is highlighted in graph 4.1, and defines a non-simple symmetric moving average of

1 Differences in weight values will result in different properties of the moving averaged series. An advanced design feature of a moving average filter may result negative weight values that may not have an intuitive interpretation.

4.1 PROPERTIES OF MOVING AVERAGES *continued*

13 terms. This particular moving average is called the 13-term Henderson moving average. It is the moving average applied by the ABS to seasonally adjusted monthly data in order to derive a series of trend estimates, and will be referred to frequently throughout section 4.1 (the design and properties of the family of moving averages known as Henderson moving averages will be discussed in detail in section 4.2). Non-symmetric moving averages are also used by the ABS in situations that will be described in Chapter 5. An example of a non-symmetric moving average of 13 terms is the 13-term Sutcliffe moving average, which has weighting pattern

Observation	*...*	*1st*	*2nd*	*3rd*	*4th*	*5th*	*6th*	*7th*	*8th*	*9th*	*10th*	*11th*	*12th*	*13th*	*...*
Weight	... 0	–0.06	–0.05	–0.03	–0.02	0.00	0.01	0.01	0.02	0.06	0.14	0.24	0.32	0.36	... 0

These weights are also plotted in graph 4.2, which highlights the non-symmetric nature of the weighting pattern.

GRAPH 4.1 WEIGHTING PATTERN OF 13-TERM HENDERSON MOVING AVERAGE

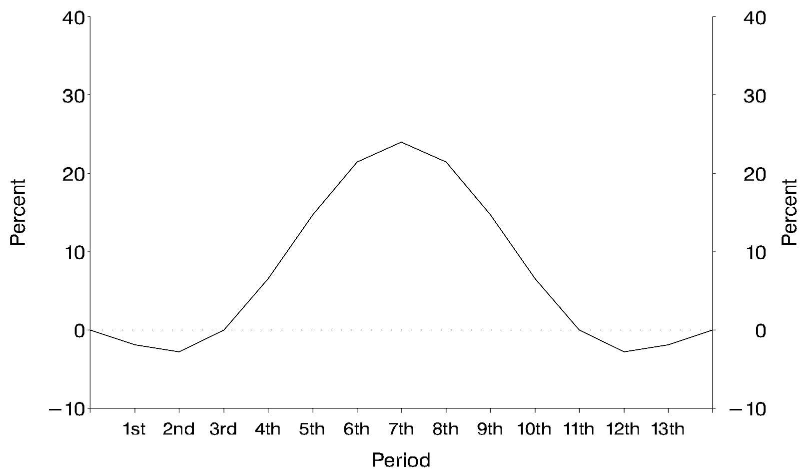

4.1 PROPERTIES OF MOVING AVERAGES *continued*

GRAPH 4.2 WEIGHTING PATTERN OF 13-TERM SUTCLIFFE MOVING AVERAGE

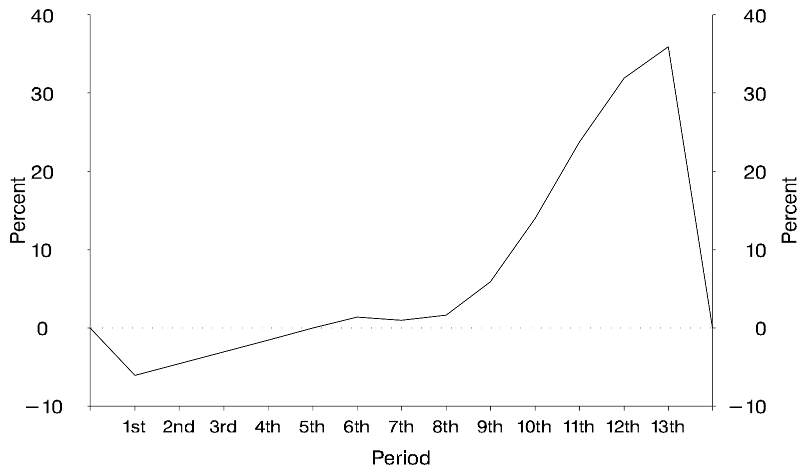

The properties of a series produced by applying a moving average are dependent on the length and weighting pattern of the moving average. Important features of moving averages that need to be considered in choosing the length and weighting pattern are:

- the degree of smoothing required
- any associated phase shift
- the end-point problem
- distortion effects.

4.1.1 Smoothing effects

The motivation behind applying a moving average to a seasonally adjusted series is to even out, or smooth out, irregularities present in the series. The irregularities have a tendency to fluctuate in magnitude and direction from period to period, and therefore the application of an averaging process will tend to cancel these fluctuations out. The result is a smoothed series, a series which is more heavily dominated by the trend component of the original series behaviour than the seasonally adjusted series from which it was derived. This smoothed series can be considered to have arisen by passing an input series through a process which filters out irregularities, and consequently, a moving average is also often referred to as a filter. The terms 'moving average' and 'filter' will be used interchangeably throughout this paper, as will the terms 'smoothed' and 'filtered', to describe the series produced by applying the moving average to an input series.

When applied to the same input series, simple moving averages generally produce smoother series the longer they become, thus, the longer the moving average, the more gradual is the difference in the change in slope between successive points in the resulting series. Hence a simple 15-term moving average will generally produce a smoother series than a simple 14-term moving average, which will in turn generally produce a smoother series than a simple 13-term moving average, and so on. The simple moving averages form a family, that is, a collection of moving averages which are defined using the same basic set of rules, irrespective of length (a simple moving average can be created for any given length by specifying a set of equal weights which sum to one). The

4.1.1 Smoothing effects continued

Henderson moving averages are also created using the same basic set of rules irrespective of length (see section 4.2) and similarly form a family. These moving averages also tend to produce smoother results the longer they become. In fact generally, within any family of moving averages specified using the same set of basic rules, longer moving averages will tend to give rise to smoother series.

For moving averages of the same length, the degree of smoothness of the series produced by applying a moving average is directly related to the degree of smoothness of the weighting pattern of the moving average, so that the smoother the weighting pattern, the smoother the resulting series (see Chapter II in Macaulay (1931)). To measure the smoothness of weighting patterns, a numerical measurement of smoothness based on the sum of the squares of the third differences is used. This measure is calculated as follows. First, an input series of weights is generated by writing down the weights of the moving average in order and appending values of zero to both ends. (Since only the terms in the moving average are involved in the calculation, it can be considered that every other observation in the time series is given a zero weight). Next, the series of third-differenced values is calculated by applying the first-difference operator to the input series of weights, and then successively applying the first-difference operator to the resulting series on two more occasions. Finally, the third-differenced series values are each squared and summed together. The smaller this sum is, the smoother is the weighting pattern of the moving average.

An example of this calculation can be seen in table 4.1, where the degree of smoothness of the weighting pattern for the simple 13-term moving average has been determined. As can be seen is this case, the weighting pattern of the moving average has a degree of smoothness of 12/169=0.0710.

TABLE 4.1 DEGREE OF SMOOTHNESS OF WEIGHTING PATTERN FOR SIMPLE 13 TERM MOVING AVERAGE

Terms	Weight	1st difference	2nd difference	3rd difference	Squared 3rd difference
...	...0	...0	...0	...0	...0
1st	1/13	1/13	1/13	1/13	1/169
2nd	1/13	0	–1/13	–2/13	4/169
3rd	1/13	0	0	1/13	1/169
4th	1/13	0	0	0	0
5th	1/13	0	0	0	0
6th	1/13	0	0	0	0
7th	1/13	0	0	0	0
8th	1/13	0	0	0	0
9th	1/13	0	0	0	0
10th	1/13	0	0	0	0
11th	1/13	0	0	0	0
12th	1/13	0	0	0	0
13th	1/13	0	0	0	0
14th	0	–1/13	–1/13	–1/13	1/169
15th	0	0	1/13	2/13	4/169
16th	0	0	0	–1/13	1/169
...	0...	0...	0...	0...	0...
Sum					12/169

4.1.1 Smoothing effects continued

As indicated above, for moving averages of the same length, the one which has the smoother weighting pattern produces the smoother result. Consider, for example, four moving averages, each 13 terms long. The simple 13-term moving average has a rectangular weight pattern, 13 units long and 1/13 units high, with zero values to the left and right. This weighting pattern is shown in graph 4.3. The weighting pattern of the non-simple 13-term moving average with end weights 1/24, intervening 11 weights 1/12, and zero weights elsewhere, is shown in graph 4.4. This weighting pattern also contains 13 positive weights, but is a little smoother, in the sense that the curve changes more gradually across the complete range of weights, including the zero weights. The non-simple moving average has a degree of smoothness of 0.0139 which is much smaller than that of the simple 13-term moving average, the smaller value indicating the smoother weighting pattern. The 13-term Henderson moving average mentioned earlier also spans 13 consecutive values and has a smoother looking weighting pattern again, as can be seen by referring back to graph 4.1. The degree of smoothness of this weighting pattern is even smaller, 0.0083. To demonstrate that this is not the smoothest possible 13-term moving average, one final example of a 13-term moving average is given in graph 4.5. The weighting pattern of this moving average has a degree of smoothness of 0.0006, making it the smoothest of the four filters considered.

GRAPH 4.3 WEIGHTING PATTERN OF SIMPLE 13-TERM MOVING AVERAGE

4.1.1 Smoothing effects continued

GRAPH 4.4 WEIGHTING PATTERN OF NON-SIMPLE 2X12 MOVING AVERAGE

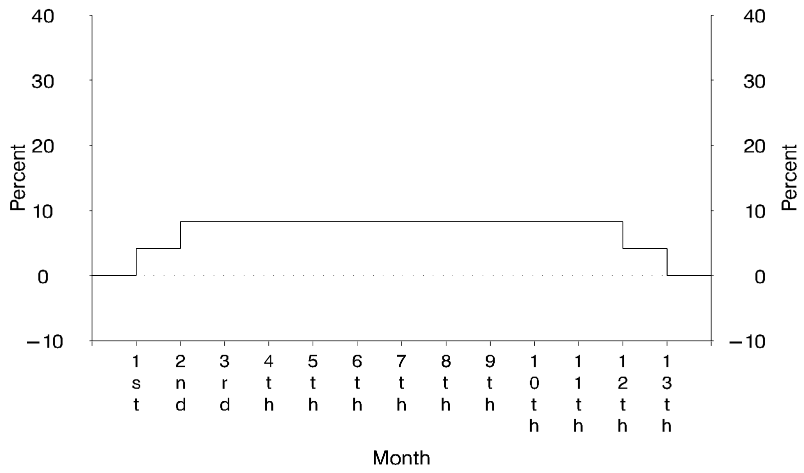

GRAPH 4.5 WEIGHTING PATTERN OF NON-SIMPLE 4X5X6 MOVING AVERAGE

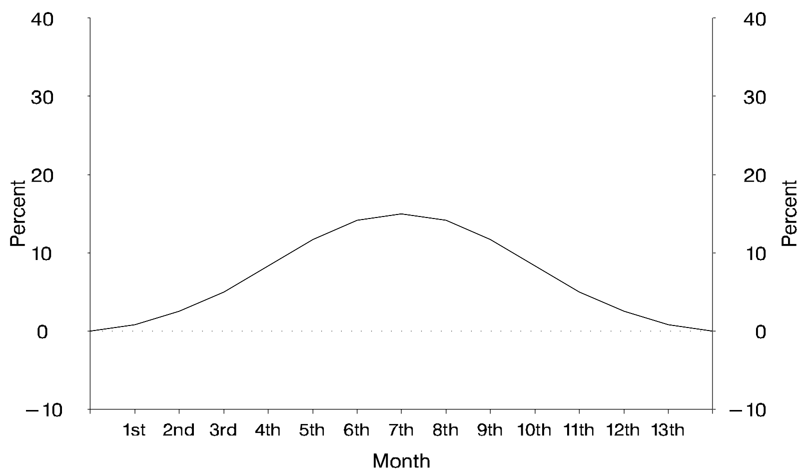

It may at first seem appropriate to seek maximal smoothness when estimating trends, but although smoothness of the weighting pattern and of the results is important, it is only one aspect of the smoothing process to consider. Other aspects are the extent to which the results are phase shifted, how much of a problem is caused in determining estimates at the end of the series, and the ability of the moving average to track the curvature of turning points in the series.

4.1.2 Phase shift

One important issue that should be considered when using a moving average, is the time point against which the result is recorded. The temptation is strong to always record the calculated average as if it referred to the end of the period over which it was calculated, in order to make the resultant time series appear as current as possible. Such an application of a moving average is said to be 'one sided', and is generally not appropriate.

4.1.2 Phase shift continued

The effect of applying a symmetric moving average by recording the results at the current end of the moving average span was illustrated in section 3.3, when the moving annual average, a simple 12-term moving average, was applied to the Unemployed Persons series (refer to graph 3.6). This example illustrated how the one sided application of the simple 12-term moving average misrepresented the timing of all the major turning points occurring in the series; it had them occurring about six months later than they actually did. This effect, of moving the turning points in the averaged series away from the turning points in the input series, is known as phase shift.

The one sided application of a symmetric moving average will always introduce phase shift at all frequencies. In fact, the only way to minimise introducing phase shift into a series when applying a symmetric moving average, is to record the moving average value in the centre of the averaged span. This is often referred to as centring the moving average (this requirement relates directly to the fact that the weighting pattern of a symmetric moving average is evenly distributed about the centre of the averaging span). Thus if a symmetric 13-term moving average is applied to a span of monthly data from January 1997 to January 1998 inclusive, then the calculated average should be recorded against July 1997, the seventh of the 13 observations spanned. Even this application will not completely avoid introducing phase shift to the series. Although the centred moving average will not introduce phase shift to the long cycles associated with the trend it will cause some of the shorter cycles associated with irregularity to experience various amounts of phase shift. However, since these shorter cycles correspond to the high frequencies targeted for removal by the filtering procedure, this effect is minimal. The remainder of this discussion will focus on how to avoid introducing phase shift to the long cycles of interest.

The effectiveness of centring a symmetric moving average in terms of phase shift is illustrated in graph 4.6. Here the seasonally adjusted Unemployed Persons series has been smoothed with both a simple 12-term moving average and a 13-term Henderson moving average. The filtered series produced by applying the 13-term Henderson moving average to the seasonally adjusted series has been obtained by placing the results of the averaging process at the centre of the 13-month averaging span, that is, at the seventh of the 13 consecutive observations being averaged. This smoothed series can be seen to reflect the timing of turning points in the seasonally adjusted series far more accurately than does the smoothed series produced by applying a simple 12-term moving average, with the averaged results recorded at the right hand end of each of the periods spanned.

4.1.2 Phase shift continued

GRAPH 4.6 HENDERSON AND SIMPLE MOVING AVERAGE BASED TREND ESTIMATES FOR THE UNEMPLOYED PERSONS

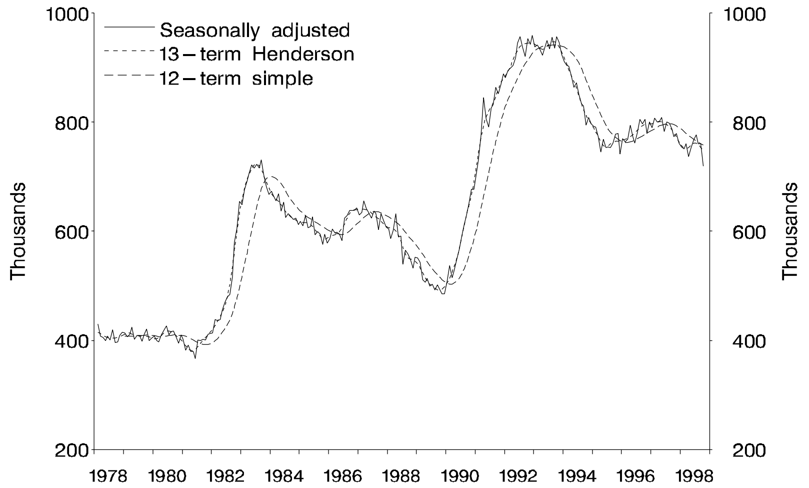

Phase shift is not only caused by recording the result of a symmetric moving average at a point other than the centre of the averaged span. Non-symmetric moving averages do not have weighting patterns evenly distributed about the centre of the averaging span, and therefore centring a non-symmetric moving average is not likely to eliminate phase shift from the longer cycles associated with the trend. In fact, it may not be possible to completely eliminate phase shift, only to minimise it. The most appropriate time period in the averaging span against which to record a non-symmetric moving average in order to minimise phase shift will depend on the weighting pattern and will always be clearly specified in conjunction with any non-symmetric weighting pattern given in this paper. In fact, all weighting patterns will be specified in such a way as to indicate the time period against which the moving average values should be recorded. Two examples are given in tables 4.2 and 4.3.

Table 4.2 gives the weighting pattern for the 13-term Henderson moving average mentioned earlier. The result of this symmetric moving average is recorded at the centre of the averaging span to avoid phase shift from the longer cycles. That is, the moving average estimate for period N is calculated by appropriately weighting the time series values for period N, and the time series values collected in the six periods before and after period N. Table 4.3 gives the weighting pattern of the 13-term Sutcliffe moving average. The result of this non-symmetric moving average is recorded at the current end of the averaging span to minimise phase shift in all cycles. That is, the moving average estimate for period N is calculated by appropriately weighting the series values for the 13 periods up to and including period N.

TABLE 4.2 WEIGHTING PATTERN OF 13-TERM HENDERSON MOVING AVERAGE

Period	*N–6*	*N–5*	*N–4*	*N–3*	*N–2*	*N–1*	*N*	*N+1*	*N+2*	*N+3*	*N+4*	*N+5*	*N+6*
Weight	–0.02	–0.03	0.00	0.07	0.15	0.21	0.24	0.21	0.15	0.07	0.00	–0.03	–0.02

TABLE 4.3 WEIGHTING PATTERN OF 13-TERM SUTCLIFFE MOVING AVERAGE

Period	*N–12*	*N–11*	*N–10*	*N–9*	*N–8*	*N–7*	*N–6*	*N–5*	*N–4*	*N–3*	*N–2*	*N–1*	*N*
Weight	–0.06	–0.05	–0.03	–0.02	0.00	0.01	0.01	0.02	0.06	0.14	0.24	0.32	0.36

4.1.2 Phase shift continued

As indicated above, application of a symmetric moving average will only avoid phase shift in the longer cycles associated with the trend if the results of the moving average are centred. A symmetric moving average of odd length can be centred by recording the result on the central observation, whereas even length averaging spans place their values between two periods, or involve the introduction of a phase shift of half a period. For instance, a 13-term symmetric moving average is centred on the seventh observation of the 13-term span it averages, whereas a 12-term simple moving average has its central value between the sixth and seventh observations, which is not observable in the data published. The value of the simple 12-term moving average could be positioned at an observable time period by placing the result on the sixth or the seventh observation, but this introduces a slight phase shift of half a period either way.

The ABS uses symmetric moving averages wherever possible, with the results centred. Also, in order to place these results meaningfully on observable time periods, the ABS uses odd length averaging spans. For smoothing monthly time series, the ABS generally uses a 13-term Henderson moving average, recording the results on the central (seventh) observation of the 13 months that are averaged. For smoothing quarterly time series, the ABS generally uses a 7-term Henderson moving average, the results being recorded against the fourth of the 7 quarters that are averaged (the weighting pattern for the 7-term Henderson moving average is given in table 4.4). The reason for using a 7 rather than 5-term Henderson moving average is discussed in section 5.2.

TABLE 4.4 WEIGHTING PATTERN OF 7-TERM HENDERSON MOVING AVERAGE

Period	*N–3*	*N–2*	*N–1*	*N*	*N+1*	*N+2*	*N+3*
Weight	–0.06	0.06	0.29	0.41	0.29	0.06	–0.06

4.1.3 The end-point problem

As discussed above, to avoid phase shift in the long cycles in a series when applying an odd length symmetric moving average, it is necessary to place the value of each average on the central observation. This means that there are a number of time periods at the end of the series for which smoothed values cannot be determined. For example, when using a symmetric 13-term moving average, the results are recorded against the seventh observation of the 13-term span. Therefore, the most current calculation uses the last 13 terms of data available and is recorded against the seventh last time period. This means that no moving average values can be determined for the last six time periods. In general, if a symmetric moving average is of length 2m+1, then the most recent moving average estimate based on values up to period N, will be recorded against period N-m. This leaves the remaining m time periods without a value.

4.1.3 The end-point problem continued

The above situation is known as the end-point problem, and the greater the length of the moving average, the greater the number of time points affected by this problem. As interest in a time series is generally focused on the most recent data points, choosing the length of the moving average is an important issue. The use of non-symmetric moving averages provides a method for coping with this issue in part, and is discussed more fully in Chapter 5.

4.1.4 Distortion

In considering the smoothed series produced by applying the simple 12-term moving average to the seasonally adjusted Unemployed Persons series in graph 4.6, it is apparent that phase shift in the resulting series is not the only issue of concern. As discussed in section 3.3, the application of the simple 12-term moving average also has the effect of distorting the shape of the turning points and points of inflection. This is a property which does not appear to be present in the series produced by applying the 13-term Henderson moving average to the seasonally adjusted Unemployed Persons series. This distinguishing feature can be demonstrated more graphically by considering the effects of repeated application of each of the moving average processes.

If the series resulting from applying the 13-term Henderson moving average is itself smoothed using the 13-term Henderson moving average, little difference will be observed between the two series. Similarly, if the 13-term Henderson moving average is applied again, there will still be little difference observed between the various series. This feature is not replicated by repeating the smoothing process using the simple 12-term moving average, as graph 4.7 illustrates. Comparing figures 4.6 and 4.7 illustrates how the repeated application of this simple moving average slowly flattens out the whole series. The series also moves progressively further to the right as each successive sequence of one sided moving averages is calculated. This process clearly demonstrates that the simple 12-term moving average has deficiencies in tracking the series to which it is applied, particularly in relation to estimating both the timing of turning points and the shape of the curve (where the shape can be seen in the level of the series and in the sharpness of turning points and points of inflection).

GRAPH 4.7 REPEATED APPLICATION OF MOVING AVERAGE BASED TREND ESTIMATES FOR THE UNEMPLOYED PERSONS SERIES

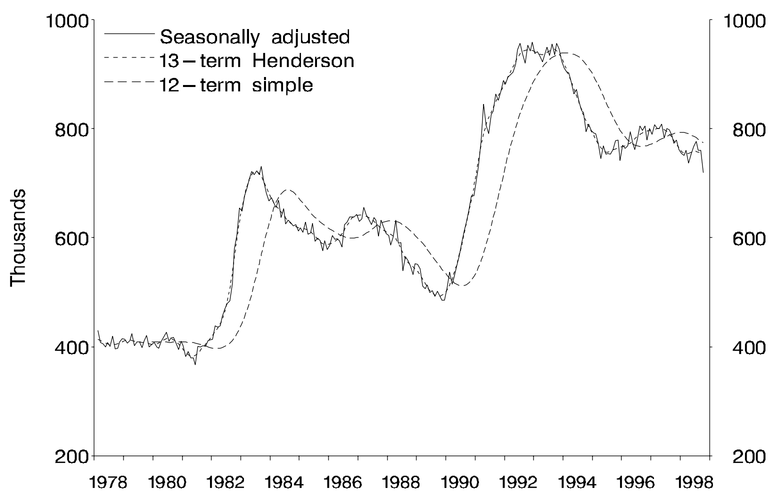

4.1.4 Distortion continued

A simple moving average can, strictly speaking, only reproduce straight line segments that appear in the data to which it is applied. When it encounters curves, corresponding to turning point peaks or troughs, or points of inflection, it can only approximate them, as seen in graph 4.6. Graph 4.8 highlights this problem when there is no irregularity in the data. Series A is the benchmark data, artificially constructed with curved and straight line segments. Series B is the data smoothed with a 13-term Henderson moving average, and Series C is the data smoothed by a simple moving average of the same length, that is, the simple 13-term moving average. From this graph it is evident that the simple moving average does not accurately track the turning points in the input series, Series A.

GRAPH 4.8 TRACKING ABILITY OF VARIOUS MOVING AVERAGES ON ARTIFICIAL DATA

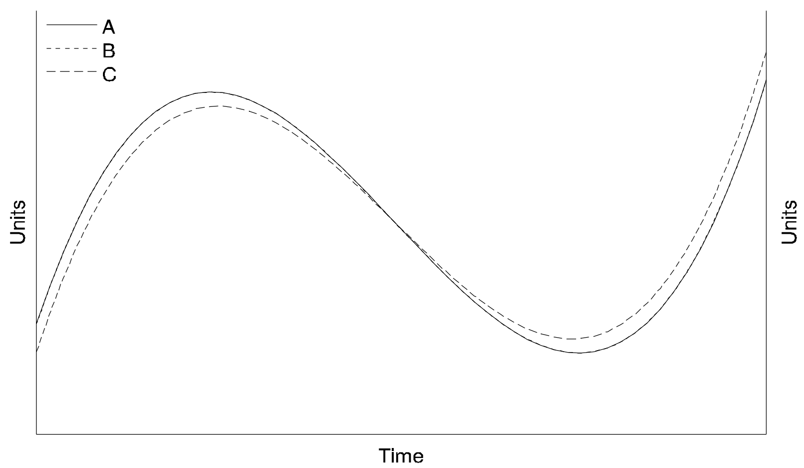

More generally, the feature of tracking only straight lines and approximating curves, arises whenever the weights of a moving average are restricted to being non-negative values. Only when some of the weights are permitted to be negative values, can moving averages be constructed to track various curvatures. There are a large number of non-simple moving averages which, by incorporating negative weights, can track specific types of curves far more efficiently than can the simple moving averages. The Henderson moving averages, for example, have been specifically designed to reproduce, not only straight line segments, but also the quadratic and cubic forms associated with peaks, troughs and points of inflection. In addition, the 43-term moving average proposed by Macaulay in 1931 as his preferred filter for smoothing monthly data, is capable of closely approximating polynomials below degree six.

As most socioeconomic time series are not well approximated by straight lines but contain curves such as peaks and troughs, points of inflection and periods of accelerating or declining growth, the ABS prefers to use moving averages that have a better chance of representing those trend characteristics than do the simple moving averages.

4.2 HENDERSON MOVING AVERAGES

The Henderson moving averages are the class of moving averages preferred by the ABS for smoothing seasonally adjusted series, in order to produce series of trend estimates. They follow a cubic polynomial trend without distortion, whilst adhering to the requirements of smoothing out a large proportion of the irregularity present in the seasonally adjusted data. In the next part of this section, the design and properties of this class of moving averages will be discussed in more detail.

4.2.1 Design of Henderson moving averages

Robert Henderson, an actuary, published a formula for the computation of the weights of his moving average in 1916. To obtain the weights, a compromise was struck between two intuitive characteristics expected of a trend series: that the trend should be able to reproduce or represent a wide range of curvatures, and that it should be as smooth as possible. The statistical criteria that Henderson used to represent these characteristics were as follows. Firstly, the weights of the moving average had to be calculated in such a way that the moving average would leave polynomials below degree four unchanged. Secondly, in conjunction with satisfying this property, the weights of the moving average had to be calculated so as to minimise the variance of the series obtained by taking the third difference of the filtered series. The first condition enables Henderson trend estimates to track curves, peaks and troughs well, and leads to the Henderson weights summing to one, while the second requirement ensures the resulting filtered series will be as smooth as possible. These conditions specified a unique weighting pattern (and consequently, a unique moving average) for each possible length considered. Henderson also demonstrated that the weights determined under these conditions gave rise to the smoothest possible weighting pattern (measured by the sum of the squares of the third differences of the weights), subject to the resulting moving average preserving cubic polynomials. Henderson's 1916 article is reproduced in Appendix 2, and is supplemented with his subsequent hand written addition of 1929, in which the exact form of the weights currently used is specified.

The Henderson moving averages can be computed for any number of odd terms, from 3 onwards. The general formula for each particular weight in the Henderson moving average of length 2m+1 is:

$$w_j = \frac{315[(m+1)^2-j^2][(m+2)^2-j^2][(m+3)^2-j^2][3(m+2)^2-11j^2-16]}{8(m+2)[(m+2)^2-1][4(m+2)^2-1][4(m+2)^2-9][4(m+2)^2-25]} \qquad (4.1)$$

for $j=-m,\ldots,m$. The weighting pattern of a 3-term Henderson moving average is:

This moving average is the shortest and simplest of the Henderson moving averages. In fact, the effect of applying and centring this particular moving average is to produce a series which is identical to the input series, that is, the application of this filter has no affect on the characteristics of the input series whatsoever. This trivial Henderson moving average is the only Henderson moving average to leave the input series unaffected and is understandably not used by the ABS.

It should be noted that all Henderson moving averages are symmetric odd length moving averages. Consequently, in addition to their smoothness and tracking properties, centred Henderson moving averages produce smoothed series which accurately reflect the timing of turning points in the series to which they are applied. The ABS uses centred Henderson moving averages for smoothing its data wherever possible.

4.2.2 Properties of Henderson moving averages

In section 2.2, the notion of the frequency decomposition of a time series was introduced as an alternative to viewing time series characteristics in the time domain. It will be recalled that, using this perspective, a time series can be considered to be the complex interaction of many cycles of behaviour of different strengths across a wide range of frequencies, with the decomposition into cycles for a given series being given by its spectrum. In section 3.1, the impact of the first-difference operator on a time series was described by considering the effect of the operator on the strength of cycles of differing frequencies. Similarly, the impact that various moving averages have upon the characteristics of time series is best appreciated using the frequency decomposition perspective.

As with the first-difference operator, the effect a given moving average has upon the strength of all possible cycles can be determined and represented graphically. An example of such a graph, or gain function, is given in graph 4.9. This graph shows how the 13-term Henderson moving average affects the strength, as measured by amplitude, of all possible cycles that can appear in a series. The horizontal axis displays the various cycle lengths ranging from two periods long at the right hand end, to very long cycles at the left hand end. The vertical axis displays what percentage of a particular cycle's amplitude remains after the moving average has been applied to the data. For example, graph 4.9 shows that the amplitude of an eight period long cycle will be reduced to around 50% of its size by the 13-term Henderson moving average. Consequently, if the 13-term Henderson moving average is used to smooth a monthly series, any eight-month long cycle present in the series will remain in the resulting series with around 50% of its strength. Cycles shorter than eight months will then remain in the filtered series with less than 50% of their strength, while longer cycles will remain with more than 50% of their strength. The feature shown in graph 4.9 is often described as the gain function of the filter.

GRAPH 4.9 EFFECT OF 13-TERM HENDERSON MOVING AVERAGE ON CYCLES

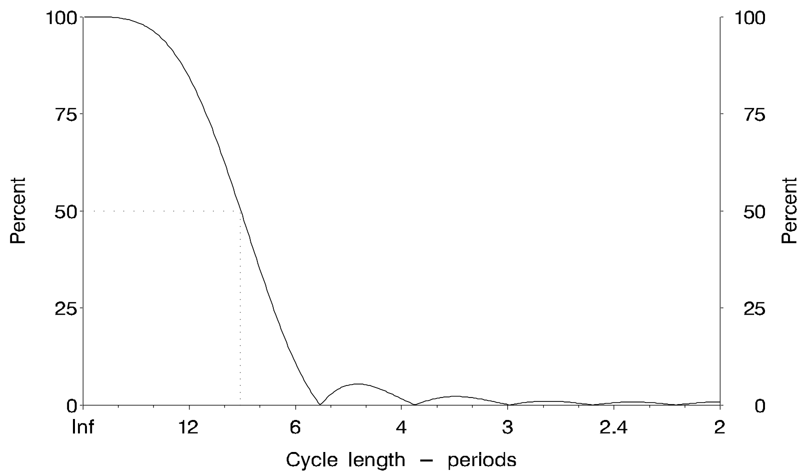

By specifying a cut-off or threshold below which cycles are considered to have been 'eliminated' from the series, graph 4.9 can be used to determine the range of cycles that

4.2.2 Properties of Henderson moving averages continued

pass through the filtering process. Taking half strength as an appropriate threshold, graph 4.9 indicates that the 13-term Henderson moving average (or 13-term Henderson filter) retains all cycles in a monthly series which are eight months or longer, and 'eliminates' all cycles in the series which are shorter. Alternatively, a threshold can be specified by using the 'power' measure. This measure indicates the contribution of a cycle to the energy or variance of the series and is given by the amplitude squared. As the 13-term Henderson filter retains about 70% of the amplitude of the ten period cycle, the ten period cycle will retain only about 50% (0.70 x 0.70 = 0.49) of its power in the filtered series. Longer cycles then retain more than 50% of their power or variance, while shorter cycles retain less. Consequently cycles longer than 10 months would remain with 50% of their power in the filtered series, while cycles shorter than 10 months would remain with less than 50% of their power in the filtered series.

The 13-term Henderson filter is used by the ABS to calculate trend estimates for monthly series. As indicated above, cycles in the range of 2 to 8 months have their strength reduced by 50% or more and consequently, the 13-term Henderson filter substantially reduces the cycles that contribute significantly to the volatility or high frequency variation of many monthly series. The 13-term Henderson filter also essentially preserves the longer term cycles associated with trend. For instance, graph 4.9 shows that cycles about 2 years long retain around 95% of their strength in the filtered series, longer cycles having even more of their strength preserved. The fact the 13-term Henderson filter barely reduces these cycles means business cycles present in even the original data will virtually remain at full strength in the trend series, as will the longer term cycles usually associated with the trend.

After the seasonal component has been removed from an original monthly series, the trend component is estimated from the resulting seasonally adjusted series by using the 13- term Henderson filter to filter out a substantial proportion of the strength of cycles contributing to irregularity. The irregular component is then estimated as the component remaining after the estimated trend component has been removed from the seasonally adjusted series. Consequently, every cycle in the seasonally adjusted series which does not pass through the filtering process into the series of trend estimates at full strength, will contribute to the estimate of the irregular component. For example, any cycle in the seasonally adjusted series which is five months long will only have around 5% of its strength passing into the estimate of trend, and therefore, around 95% of its strength will remain in the estimate of the irregular component. On the other hand, a cycle of length two years has around 95% of its strength passing into the trend estimate, while the remaining five per cent of its strength remains in the estimate of the irregular component.

It is evident from graph 4.9 that the 13-term Henderson filter captures the longer term cycles associated with trend and eliminates the shorter cycles associated with irregularity, by displaying a gradual transition in cycle attenuation properties from low to high frequencies. Thus there is a grey or twilight area for some cycle lengths in the short to medium range. In particular, cycles in the range of six months to two years contribute to the filtered series with strengths between 10 and 95%, and will consequently be split between the estimate of trend and the estimate of the irregular component. Any cycle in this range of lengths which is thought to be of interest to those monitoring trend behaviour will therefore only partially contribute to the estimate of trend while also

4.2.2 Properties of Henderson moving averages continued

partially contributing to the estimate of the irregular. On the other hand, any cycle in this range of lengths which is not thought to be of interest in monitoring trend behaviour will also contribute to the estimate of trend.

It is possible that a more desirable filtering process would be to isolate exactly those cycles with lengths of interest, while completely eliminating all others, and such a filtering process will be discussed in section 4.3.1. In practice, however, the exact choice of a cycle length cut-off differentiating between those cycles that should be passed at full strength and those that should be completely eliminated, will be dependent on the particular series being analysed and the type of analysis to be undertaken (assuming the analysis lends itself to a clear choice of cycle length cut-off at all). The 13-term Henderson filter has been found to strike an appropriate compromise between capturing those cycles with length generally of interest and filtering out those that are not, and can be used very effectively to derive a broad indicator of trend behaviour.

In considering graph 4.9, it can be seen that about 85% of the strength of the fundamental seasonal cycle of 12-months will pass through the filtering process. Consequently, the 13-term Henderson filter should not be applied to monthly data that contains fundamental seasonality, as the yearly cycle would largely remain. On the other hand, the other seasonal cycles (referred to as harmonics) of 6, 4, 3, 2.4 and 2 months long would, if not already removed, have strengths respectively of 10, 2, 0, 0 and 0 per cent. A practical advantage of the substantial reduction of the strengths of these sub-annual seasonal cycles, is that if the seasonal adjustment process has some deficiencies, such that there remain residual amounts of high frequency seasonal variation in the adjusted data, the smoothing process will greatly reduce their presence. Similarly, as the majority of the spectral peaks associated with the other systematic calendar related effects lie in the range of two to six months, any residual trading day, pay day, or moving holiday effects present in the seasonally adjusted series will also be substantially reduced.

In addition to attenuating any residual seasonal effects, the property of reducing the strength of cycles in the short cycle range also has implications for data collection methods. If a lot, or most, of the sampling and non-sampling error can be confined to the high frequency range of cycles, around two to eight months long in the original data (as has been observed to occur for some collection methodologies), they will then make minimal contribution to the Henderson based trend estimates. Conversely, if these statistical errors contribute most of their strength to cycles longer than eight months, they will have a part in determining the estimated trend behaviour. If that part is influential, the estimate of the trend component of behaviour would not necessarily reflect the real world socioeconomic activity it is intended to.

It is interesting to compare the damping properties of the 13-term Henderson moving average with the moving annual average, otherwise known as the simple 12-term moving average, described in section 3.3. The damping properties of these moving averages are given in graph 4.10. It is evident from this graph that the simple 12-term moving average reduces cycles in the medium cycle range more than the 13 term Henderson, but that it incorporates more of the strength of cycles in the high frequency range, consequently resulting in a more volatile series. The 13-term Henderson moving average also has the advantage that the contribution from the business cycles around two to five years will

4.2.2 Properties of Henderson moving averages continued

not be significantly reduced in a monthly series, whereas the simple 12-term moving average will more substantially reduce the impact of these cycles. Although not particularly significant in the analysis of series which have already been seasonally adjusted, it is interesting to note that the simple 12-term moving average removes the annual seasonal cycle of 12 months in monthly data. This is not surprising, as one of the original motivations behind using the moving annual average is to eliminate the effects of seasonality, by averaging out the behaviour of series values over the period of one year.

GRAPH 4.10 GAIN FUNCTIONS OF 12-TERM SIMPLE AND 13-TERM HENDERSON MOVING AVERAGES

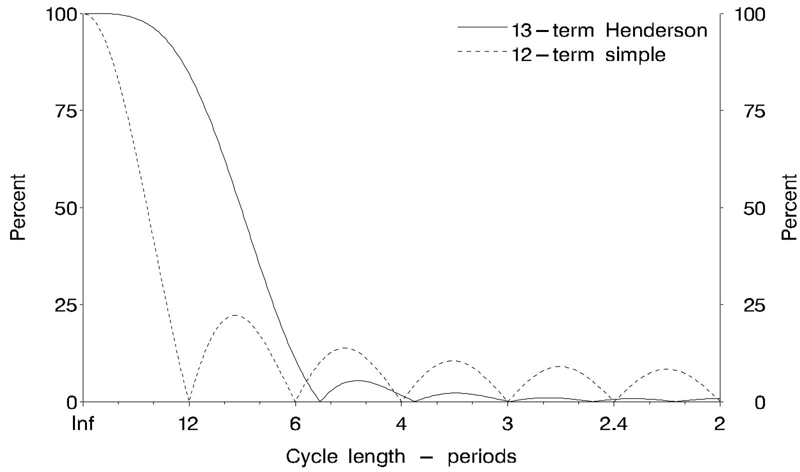

The 13-term Henderson moving average is not the only Henderson filter available for smoothing series, and graph 4.11 shows the gain functions of some of the other length moving averages used by the ABS. The impact of these moving averages on the strength of cycles is summarised in table 4.5. The weights of these filters are displayed in Appendix 3 which also includes their weight and gain calculation formulae.

4.2.2 Properties of Henderson moving averages continued

GRAPH 4.11 GAIN FUNCTION OF VARIOUS HENDERSON MOVING AVERAGES

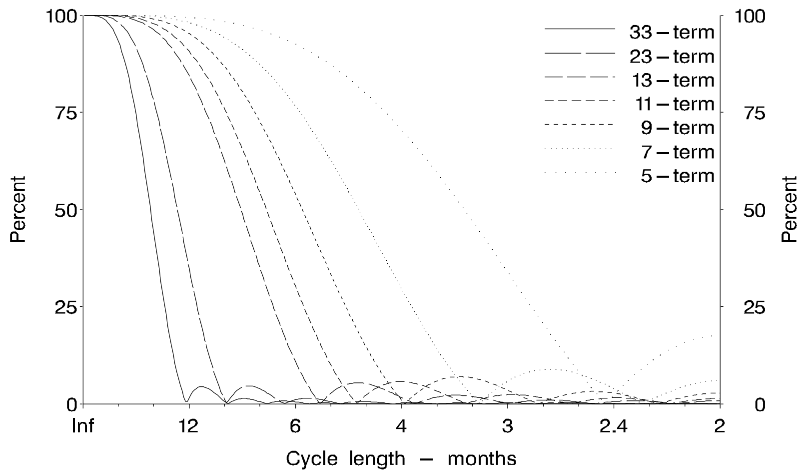

TABLE 4.5 IMPACT OF HENDERSON MOVING AVERAGES ON CYCLES

No. of terms	*10%*	*25%*	*50%*	*75%*	*90%*
5	2.60	2.84	3.34	4.21	5.51
7	3.49	3.88	4.63	5.88	7.74
9	4.33	4.84	5.81	7.41	9.78
11	5.15	5.78	6.95	8.89	11.73
13	5.95	6.69	8.06	10.32	13.64
23	9.89	11.16	13.49	17.31	22.90
33	13.77	15.56	18.84	24.18	31.99

The object of applying any filter is to preserve as much of the strength of those cycles of interest as possible, while removing as many cycles outside the range of interest as possible. For monthly series, the 13-term Henderson moving average has generally been found to provide an appropriate compromise between preserving the cycles around the length of the business cycles and longer, while removing as much of the irregularity in the series as possible. Quarterly series, on the other hand, are generally smoothed by using a 7-term Henderson moving average.

Table 4.5 shows that a 7-term Henderson moving average preserves 50% or more of the strength of cycles at least 4.63 periods long, and therefore cycles shorter than 4.63 periods are reduced to less than 50% of their strength in the filtered series. This indicates that when the 7-term Henderson filter is applied to a quarterly series cycles shorter than a year (four quarterly cycle = one year) will be largely removed from the quarterly data. As any cycle longer than 7.74 periods remains with 90% of its strength, cycles two years (eight quarters) and longer present in the quarterly series remain virtually at full strength in the corresponding smoothed series. Consequently, the 7-term Henderson filter removes a substantial part of the irregularity in quarterly series, while preserving the business cycles and the longer cycles associated with the secular trend.

4.2.2 Properties of Henderson moving averages continued

Since different length Henderson filters can be used to isolate different ranges of cycles from a particular time series, more complicated cycle extraction can be performed by combining two or more filters. An example of these filters occurs in relation to the ABS experimental composite leading indicator (CLI). The CLI is designed to detect turning points in the Australian business cycle as represented in GDP. An estimate of this business cycle is produced by smoothing the seasonally adjusted GDP series with both the 7 and 33-term Henderson moving averages, and then subtracting the 33-term smoothed series from the 7-term smoothed series. As indicated above, a 7-term Henderson moving average retains 50% or more of the strength of cycles that are at least 4.64 quarters long in quarterly data. On the other hand, a 33-term Henderson moving average retains 50% or more of the strength of cycles that are at least 18.78 quarters long. Consequently, subtracting the 33-term smoothed series from the 7-term smoothed series will result in a 'business cycle' indicator containing cycles in the range of about one to five years (4.64 quarters to 18.78 quarters), based on a 50% cycle strength threshold. (Selecting a different strength threshold will then imply a different potential range of cycles in the indicator, a 75% strength threshold, for example, indicates a range of approximately 18 months to 6 years (5.88 quarters to 24.18 quarters)). The gain of the '7 less the 33-term' Henderson filter appears in graph 4.12.

GRAPH 4.12 EFFECT OF 7-TERM LESS 33-TERM HENDERSON ON CYCLES

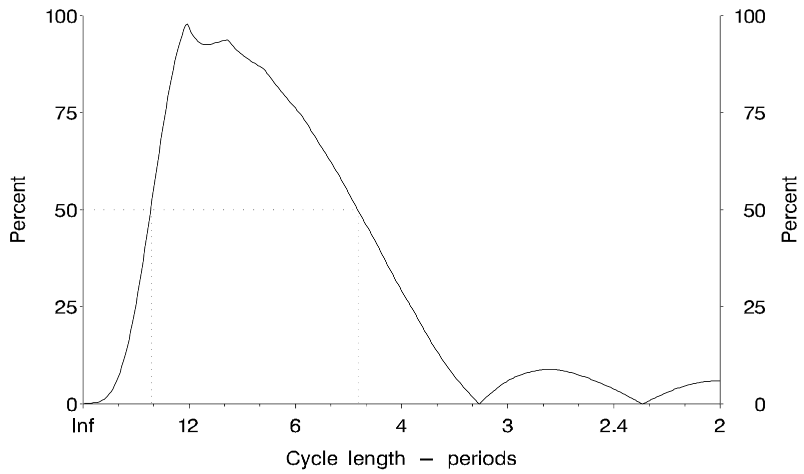

4.2.3 Impact of Henderson moving averages on noise

A seasonally adjusted series consists of the two non-seasonal components of original series behaviour, the irregular component and the trend. The irregular component represents the impact of random or irregular influences on the data being collected, and corresponds to the 'noise' or 'crackle' in the data that obscures the basic 'signal' or trend. This 'noise' is analogous to the interference that might cloud a television picture, or the hiss accompanying a radio or recorded message. The aim in applying a Henderson moving average to a seasonally adjusted series is therefore to isolate the trend by filtering out the noise present in the data. The effectiveness of the Henderson filters in achieving this goal can be considered by investigating their impact on series artificially constructed to represent the characteristics of noise.

4.2.3 Impact of Henderson moving averages on noise continued

When people talk of 'noise', they are usually referring to 'white noise'. A time series is said to consist purely of white noise when its spectrum displays equal contribution or power from all cycles, that is, equally from low, medium and high frequencies (as indicated in graph 4.13). A series produced by a random generating process tends to display the characteristics of white noise, as does a series constructed from a set of conventional 'random number tables'. Other forms of noise also exist however. For example, in a volatile seasonally adjusted series, the noise contributed by irregularity tends to be dominated by the higher frequency cycles. Such noise is here termed 'black' noise, and is also indicated in graph 4.13.

GRAPH 4.13 BLACK AND WHITE NOISE

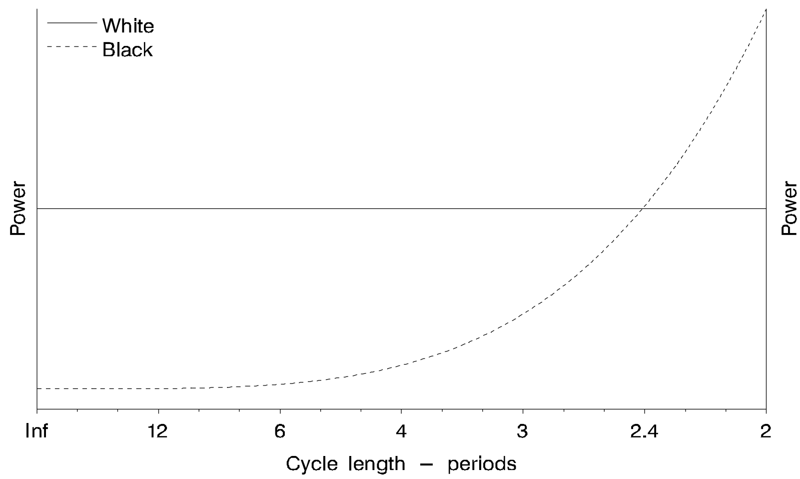

Regardless of what form the noise takes, the gain function shown in graph 4.9 indicates that applying the 13-term Henderson filter to a noise series will largely eliminate the high frequency cycles, while essentially passing those medium and low frequencies that are present in the noise series. Relating this back to the noise in the seasonally adjusted series, this indicates that applying the 13-term Henderson filter to the seasonally adjusted series will isolate not only the medium and long-term cycles associated with the trend, but also any medium and long-term cycles resulting from irregularity. As a result, there are likely to be low or medium frequency cycles arising in the original series as a result of irregularity, which will contribute to the estimate of trend. While this may seem unsatisfactory, it must be remembered that the reason for applying a Henderson filter to a seasonally adjusted series is to isolate the trend behaviour, which naturally takes the form of long-term cycles. The trade-off in ensuring this behaviour is captured is to accept the potential presence of some long cycles associated with irregularity in the trend estimates. If a series does not consist solely of irregularities (as is the case with the social and economic data collected and analysed by the ABS) that is, if a series does contain some underlying pattern of ongoing behaviour, then the contribution from the low frequency trend cycles will substantially outweigh the contribution from any irregularities

4.2.3 Impact of Henderson moving averages on noise continued

passing through the filtering process. Furthermore, the trend estimate, having had a substantial proportion of the original series irregularity removed (in the form of the high frequency cycles), will be a better indicator of trend behaviour than either the seasonally adjusted series, or the types of indicators discussed in Chapter 3.

It should be noted that Henderson filters have been designed specifically for the purpose of estimating the underlying trend component of social and economic time series behaviour. Therefore, apart from considering the impact of Henderson filters on artificially constructed noise series in order to judge their effectiveness in removing the irregular component of series behaviour, there is little meaningful information that can be gained by applying Henderson filters to series consisting solely of noise (these series have no trend to estimate). The ABS does not determine trend estimates for series with no meaningful social or economic trend growth.

4.2.4 Serial correlation in smoothed series

If the irregularity in a seasonally adjusted series is thought of as representing the random behaviour which masks the underlying substantive behaviour of the trend, then the application of a moving average should reduce the effects of the irregular influences, thereby leading to a series in which the behaviour is more predictable and less random. Returning to the time domain, and regarding each series as a collection of values over an interval of time, a means by which the strength of random influences can be measured is to look at the serial correlation of the series. Some analysts have suggested that the usefulness of ABS trend estimates is undermined by the increase in serial correlation arising from the trending procedure. This part of the paper considers this issue and how smoothing a time series affects its serial correlation.

Serial correlation is a measure, $\rho_{(i,i+k)}$, of the degree to which terms k periods apart in a series have a tendency to vary systematically with each other. If the terms k periods apart in the series vary perfectly in the same direction, then $\rho_{(i,i+k)} = 1$. If the terms k periods apart vary perfectly in the opposite direction to each other, then $\rho_{(i,i+k)} = -1$. If there is no systematic variation between the terms k periods apart, then $\rho_{(i,i+k)} = 0$.

By their very nature, both the simple and Henderson moving averages tend to produce results that have stronger serial correlation than the series to which they are applied. This is a characteristic associated with the smoothing process and is best illustrated by considering the effects of applying a moving average to a series initially displaying no serial correlation. Suppose a moving average is applied to a random series, R_i, $i = 1, 2, \ldots, t$, where it is assumed that the overall average or expected value is $E(R_i) = 0$, that the variance $E(R_i^2) = \sigma^2$, and that the covariance $E(R_i)(R_{i+k}) = 0$ for all $k \neq 0$. Also suppose that the moving average has the form $w_1 R_i + w_2 R_{i+1} + \ldots + w_n R_{i+n-1}$ where, as usual, the weights w_j, $i = 1, 2, \ldots, n$, satisfy the condition:

$$\sum_{j=1}^{n} w_j = 1.0 \qquad (4.2)$$

Then, the serial correlation of the resultant smoothed series is given by the formula

$$\rho_{(i,i+k)} = \frac{\sum_{j=k+1}^{n} w_j w_{j-k}}{\sum_{j=1}^{n} w_j^2}, \quad k < n \qquad (4.3)$$

4.2.4 Serial correlation in smoothed series continued

Firstly, consider the class of simple moving averages, moving averages in which the weights are all equal and non-negative. The above formula indicates that when applying a simple moving average of length n to a random series, the serial correlation is strongest when k=1, and declines as k increases, to the extent that there is zero serial correlation between terms n or more periods apart. This is not surprising, as the closer two values are in the filtered series, the more input series values are common to both moving average calculations, and therefore, the more likely the smoothed values are to vary systematically with each other. Another property indicated by this formula is that for a fixed value of k, the serial correlation between terms k periods apart becomes stronger as the moving average becomes longer. Again, this is not surprising, as the longer the moving average, the more terms are common to two moving average calculations k periods apart.

Unlike simple moving averages, Henderson moving averages do not have equal non-negative weights, although in general, it is still true that the above properties hold. Thus generally, the further apart two terms are in a random series smoothed by a Henderson moving average, the weaker will be the serial correlation, while for a fixed value of k, longer Henderson moving averages tend to give rise to stronger serial correlation in terms k periods apart. In addition, the serial correlation for any k is in fact generally smaller for a Henderson moving average than it is for the simple moving average of the same length.

As the serial correlation produced by applying a moving average process is not determined by the variance of the random series, the above formula can be used to calculate values of the serial correlation for any random series satisfying the properties outlined earlier. Table 4.6 gives the values of the serial correlation for the 7, 13 and 23-term Henderson moving averages, the filters used by the ABS for smoothing quarterly and monthly data respectively, as well as the simple moving averages of the corresponding lengths. This table can be used in conjunction with graph 4.14 to illustrate some of the properties described above.

4.2.4 Serial correlation in smoothed series continued

TABLE 4.6 SERIAL CORRELATION COEFFICIENT, P(I,I+K)

k	*7 Henderson*	*7 Simple*	*13 Henderson*	*13 Simple*	*23 Henderson*	*23 Simple*
1	0.76	0.86	0.91	0.92	0.97	0.96
2	0.28	0.71	0.69	0.85	0.88	0.91
3	–0.04	0.57	0.41	0.77	0.74	0.87
4	–0.09	0.42	0.16	0.69	0.58	0.83
5	–0.02	0.28	–0.01	0.62	0.41	0.78
6	0.01	0.14	–0.08	0.53	0.25	0.74
7	—	—	–0.08	0.46	0.12	0.70
8	—	—	–0.05	0.38	0.01	0.65
9	—	—	–0.01	0.31	–0.05	0.61
10	—	—	0.00	0.23	–0.08	0.56
11	—	—	0.01	0.15	–0.09	0.52
12	—	—	—	0.08	–0.07	0.48
13	—	—	—	—	–0.05	0.43
14	—	—	—	—	–0.03	0.39
15	—	—	—	—	–0.01	0.35
16	—	—	—	—	0.00	0.30
17	—	—	—	—	0.00	0.26
18	—	—	—	—	0.01	0.22
19	—	—	—	—	—	0.17
20	—	—	—	—	—	0.13
21	—	—	—	—	—	0.09
22	—	—	—	—	—	0.04
23	—	—	—	—	—	—
etc.	—	—	—	—	—	—

— nil or rounded to zero (including null cells)

GRAPH 4.14 SERIAL CORRELATION COEFFICIENTS FOR VARIOUS MOVING AVERAGES

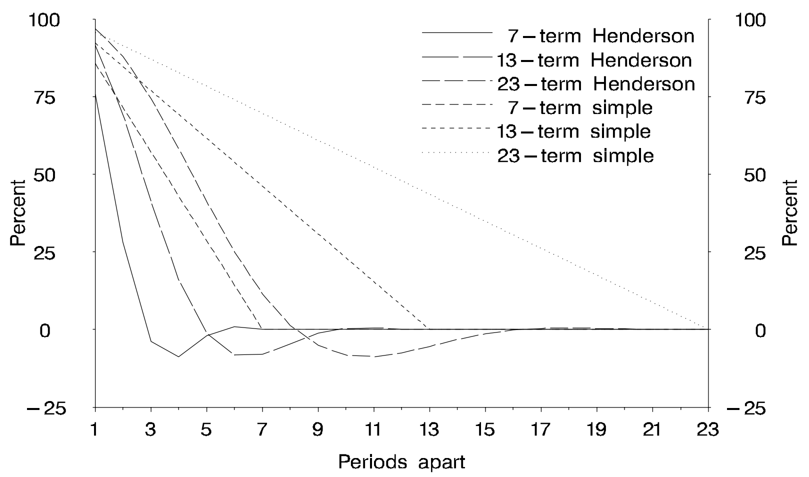

By considering the effects of applying a moving average to a random series, it can be seen that passing a seasonally adjusted series through a Henderson moving average will generally produce a series which displays stronger serial correlation than the seasonally adjusted series to which it was applied (as the filtering process itself will introduce serial correlation independently of the degree of serial correlation already present in the input series). This may seem a little strange, however in practice, most economic time series

4.2.4 Serial correlation in smoothed series continued

themselves can be considered to have high levels of underlying serial correlation (in their trend component), offset by random or irregular influences. As indicated earlier, the smoothing process seeks to strengthen the presence of this trend component by reducing the presence of the irregularities, and therefore the filtered series would naturally be expected to display stronger serial correlation than the input series. Consequently, the Henderson moving averages seek to reflect the underlying serial correlation present in the seasonally adjusted series.

4.3 ALTERNATIVES TO HENDERSON MOVING AVERAGES

4.3.1 Angularity of the gain function — cliff filters

It has already been seen in section 4.2.2 that the curvature of the gain of the 13-term Henderson filter gives rise to a degree of partial attenuation around some cycles of interest, leading to ambiguity in the analysis of some medium-term cyclical behaviour (other Henderson filters have similar properties, although the degree of curvature and the range of cycles affected by that curvature, varies for Henderson filters of different lengths). The gradual transition in the degree of attenuation for the Henderson filters can present a second problem for analysts, namely, that it is not possible to specify an exact cycle length cut-off which will differentiate between those cycles that are passed by the filtering process and those that are not. Instead, for analysis purposes, it may be necessary for the analyst to choose an appropriate threshold, based on some specified percentage reduction in the strength or power of cycles, to indicate the range of cycles that may be present in the filtered series. While the effects of the filter remain constant, different choices of threshold will give rise to a different impression of the range of cycles that are captured by the filtering process. This was particularly evident for the composite leading indicator discussed in section 4.2.2, in which two Henderson filters were combined in order to capture a range of cycles of medium length.

In order to circumvent these problems it would, ideally, be desirable to develop a set of filters which allowed cycles beyond a specified length to pass through at full strength, while simultaneously ensuring the complete elimination of all others. Combinations of such filters would then be used to isolate any particular range of interest (a filter that captured exactly those cycles longer than eight years, for example, would be combined with a filter that captured exactly those cycles longer than two years, to produce a filter which would isolate only those cycles two to eight years in length). The ideal filter for capturing cycles of length X and longer would be one that leaves the strength of all cycles of length X or longer unaffected and reduces the strength of all shorter cycles to zero (see graph 4.15 for an example), without introducing any phase shift.

In reality, it is not possible to apply such an ideal filter. An ideal filter has infinite non-zero weights and would require an infinite span of time series data. In practical terms, a truncated version of the ideal filter would have to be used. However, the truncated filter would cause some filter leakage and introduce side lobes to the gain function. Filter leakage would also contribute to the frequency aliasing problem (see section 6.3 and 6.4 in Koopmans (1974)). Therefore, ideal filters have very limited value in the production of quality trend estimates.

Although it is not possible to apply ideal filters to ABS time series, it is possible to obtain a greater degree of angularity in the gain of filters than presently possessed by the Henderson filters. As an example, graph 4.15 gives the gain function of one of the experimental angular or 'cliff' filters developed by the ABS, compared to the 13-term Henderson filter and the ideal filter it attempts to approximate.

4.3.1 Angularity of the gain function — cliff filters continued

GRAPH 4.15 EFFECT OF 13-TERM HENDERSON AND ASSOCIATED IDEAL AND CLIFF FILTERS ON CYCLES

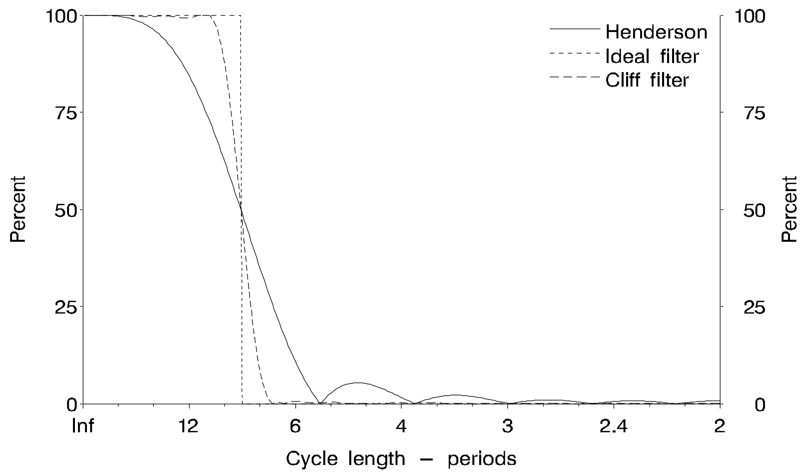

Generally, the cliff filters referred to above are obtained by averaging over relatively longer time spans than are used by the equivalent Henderson filters (the term 'equivalent' is used here in the sense of the filter having the same 50% strength cut off). As the smoothed estimates from the cliff filter are also placed in the centre of the averaging span, they cannot proceed as far to the current end of the series as can the shorter equivalent Henderson filter, and therefore there are more periods at the end of the series affected by the end-point problem. As a result, there are more periods at the end of the series for which smoothed estimates are calculated using alternative methods (how the end-point problem is dealt with for series smoothed by Henderson moving averages will be discussed in Chapter 5). As these methods give rise to smoothed estimates which are subject to revision as new data becomes available, series smoothed using a cliff filter are likely to have more estimates at the end of the series subject to revision than series smoothed using the equivalent Henderson filter, and these estimates are also likely to be subject to revision for a longer period of time. While this problem may cause cliff filters to be unsuitable for undertaking smoothing over the entire length of the series, it may be rational to use a mix of cliff and Henderson filters in order to obtain a less ambiguous indicator for at least some mid-section of the series.

4.3.2 Other filters

There is a variety of alternative filter designs based on different smoothing criteria and with different frequency properties. There are also different methods for estimating the weighting patterns in the literature. The ABS is investigating various alternative filter design approaches and the practical implications of using such approaches for its official statistics.

Least squares estimation procedures can be used to generate weighting patterns that give rise to filters which have specified effects on the various cycles in time series. A more sophisticated procedure tailors the filter to a given series by using the spectral properties of the series in conjunction with least squares methods, to produce the filter which best achieves the desired goal for that particular series. To date, research by the ABS using such an approach has not provided noteworthy benefits because of two

4.3.2 Other filters continued

practical problems associated with obtaining a sustainable and reliable estimate of the spectrum of the specific time series. Firstly, the number of observations available for the particular time series may be insufficient to give the requisite precision to the spectral analysis. Secondly, due to the nature of the process involved in estimating the spectrum, a reliable estimate of the spectrum of a given series may not be attainable, regardless of the amount of data available.

There are a large number of other filters available for smoothing series. One moving average already mentioned which is of particular interest, is the non-simple symmetric 13-term moving average with weights 1/24, 1/12, ..., 1/12, 1/24 (it is often called 2x12 term moving average. Graph 4.4 shows its weight pattern). The cycle strength reduction properties of this moving average are compared to the cycle strength reduction properties of the 13-term Henderson moving average in graph 4.16. It can be seen from this graph that unlike the 13-term Henderson filter, the non-simple 13-term moving average eliminates the annual seasonal cycle of length 12 months in monthly data, as well as the other seasonal cycles of length 6, 4, 3, 2.4, and 2 months. This property of the non-simple symmetric 13-term moving average makes it very useful for removing seasonality from monthly data, and this moving average is in fact incorporated into the early stages of the seasonal adjustment process used by the ABS.

GRAPH 4.16 EFFECT OF 13-TERM HENDERSON AND 2X12 TERM MOVING AVERAGES ON CYCLES

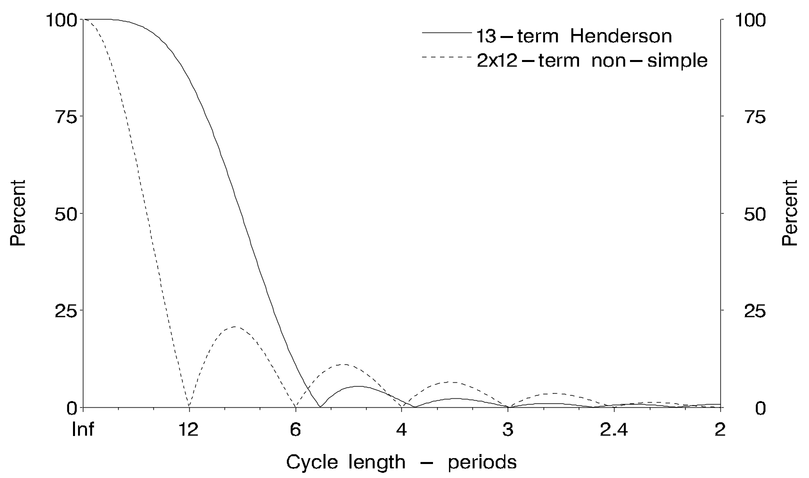

CHAPTER **5**

END-POINT ESTIMATES

CONTENTS

page

SECTION

5.0 OVERVIEW

In the last Chapter, it was established that Henderson moving averages can be used to smooth out high frequency irregularity in a series, while simultaneously tracking the types of curvatures present in the series and accurately capturing the timing of turning points. As with any centred symmetric moving average, however, this process cannot be used to directly estimate smoothed series values all the way to the end of the series. For example, consider the application of a 13-term Henderson moving average to a monthly time series that ends in September 2002. The last moving average value that can be directly calculated from the available data will be based on the last 13 consecutive time series observations, September 2001 to September 2002, and will be recorded against the seventh observation in this span, March 2002. Consequently, the 13-term Henderson filter cannot be used to calculate a smoothed value for any of the last six months, April 2002 to September 2002, as there are insufficient observations available for the moving average calculations. This absence of smoothed estimates at the end of the series is known as the end-point problem.

In general, if applying a 2m+1 term Henderson moving average to a time series, there will be m time periods at the end of the series (and indeed, also at the beginning of the series) for which smoothed values cannot be directly estimated using the Henderson filter. Thus for a 13-term Henderson moving average (m=6, 13=2x6+1), there will always be six periods at the end of the series affected, while for a 7-term Henderson moving average (m=3, 7=2x3+1), there will always be three periods at the end of the series affected. As the current end of the time series is of particular interest to those monitoring socioeconomic behaviour, it is necessary to have a mechanism in place for estimating the contribution of the trend component for the most recently collected set of original series observations. The ABS uses a collection of non-symmetric moving averages to determine estimates of trend at the current end of the series, and this chapter discusses the derivation and properties of these moving averages.

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES

One way to overcome the problem of obtaining estimates of trend at the end of the series is to explicitly extend the seasonally adjusted series by a further m values (where the Henderson filter to be applied to the series has length 2m+1). For example, suppose data for a monthly series is available up to month N, and that this data is used to project, or forecast, series values six months into the future. The 13-term Henderson moving average could then be applied to the extended series to produce trend estimates for the last six months up to and including month N. The trend estimate for month N, for instance, would be calculated by using the seasonally adjusted value for month N, the six seasonally adjusted values for the half year prior to month N, and the six forecast values for the half year following month N. As time passes and the actual real data values become available, each successive forecast value could be replaced by the actual value observed to occur at the appropriate point in time, and the Henderson moving average calculations at the end of the series repeated. If the forecast values perfectly predict the actual values that are observed, then replacing the forecast values by the actual values will not change the moving average calculations, and there will be no change to the estimates of trend at the end of the series. If any of the actual values are different from those forecast, then the recalculations will produce different 'revised' estimates of trend at the end of the series, which are more indicative of the ongoing long-term behaviour of the data being collected. In practice forecasts are rarely accurate, particularly the further

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

they look into the future, and the more inaccurate the forecast values are, the more the trend estimates will be subject to revision as real data becomes available.

Another practical approach towards addressing the end-point problem is to use a collection of progressively shorter symmetric moving averages as a set of 'surrogates' for the main Henderson filter. For example, shorter Henderson moving averages could be applied to determine the estimates of trend at the end of the seasonally adjusted series. Thus if a symmetric 13-term Henderson filter had been used to smooth the substantive part of the series, an 11-term Henderson could be used to determine an estimate of trend for the sixth last time period, a 9-term Henderson could be used to determine an estimate of trend for the fifth last, a 7-term Henderson for the fourth last, and so on. Under this methodology, however, there is still a problem with determining estimates of trend for the last two time periods. The second last estimate, for instance, would be determined by applying the 3-term Henderson filter and, as discussed in section 4.2.1, this moving average does not change the characteristics of the underlying series, that is, it does not filter out any irregularity. In addition, there is no Henderson filter available for smoothing just one observation, and consequently, no trend estimate could be determined for the most recent time period.

To combat at least part of this problem, a collection of progressively shorter simple moving averages could be used as surrogates for the main Henderson filter. Again, for the 13-term Henderson filter, a simple 11-term moving average could be used to determine the sixth last estimate of trend, a simple 9-term to determine the fifth last, and so on, up to a simple 3-term moving average for determining the second last estimate of trend. Unlike the option above, this would result in a smoothed value for the second last time period in which there was some reduction in irregularity. Again, however, no estimate of trend could be determined for the most recent time period, as there is no simple moving average which can smooth just one data point. In fact, there is no symmetric moving average which can smooth just one data point.

In addition to failing to provide trend estimates all the way to the end of the series, the use of these symmetric moving averages gives rise to a problem which is related to the smoothing characteristics of the different filters. Graph 5.1 compares the smoothing properties of the 3, 5, 7, 9 and 11-term Henderson filters with the 13-term Henderson filter. It is evident from this graph that the smoothing properties of each of these filters differ significantly from the 13-term Henderson filter, and also from each other. Consequently, the trend estimates determined for each of the last m observed time periods may capture different characteristics of the series than the Henderson based trend estimates, and each other (depending on the cycles actually present in the series). This is an important consideration, as it can result in misinterpretation of the impact of socioeconomic forces at the current topical end of the series. Furthermore, the difference in filtering characteristics can lead to a large degree of revision in the trend estimates when enough data has become available to calculate these estimates using the main Henderson filter. Graph 5.2 illustrates that the filtering properties for the simple symmetric moving averages discussed also differ from the symmetric 13-term Henderson filter and from each other. These moving averages consequently also give rise to inconsistent filtering properties at the topical end of the series.

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

GRAPH 5.1 GAIN FUNCTIONS OF VARIOUS HENDERSON MOVING AVERAGES

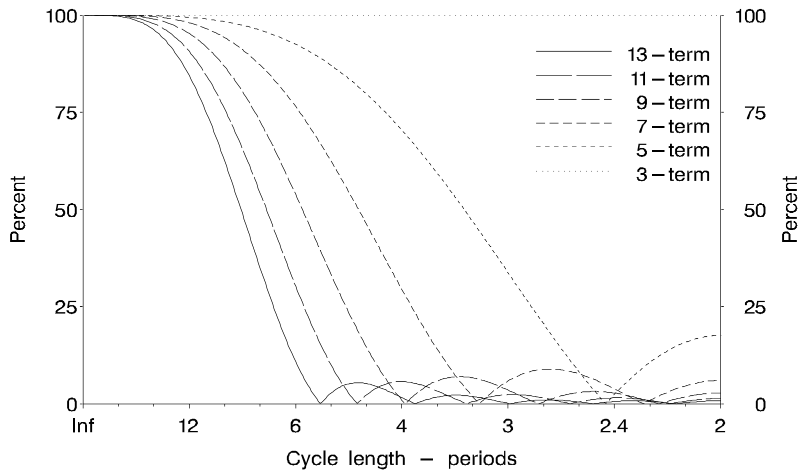

GRAPH 5.2 GAIN FUNCTIONS OF VARIOUS SIMPLE MOVING AVERAGES

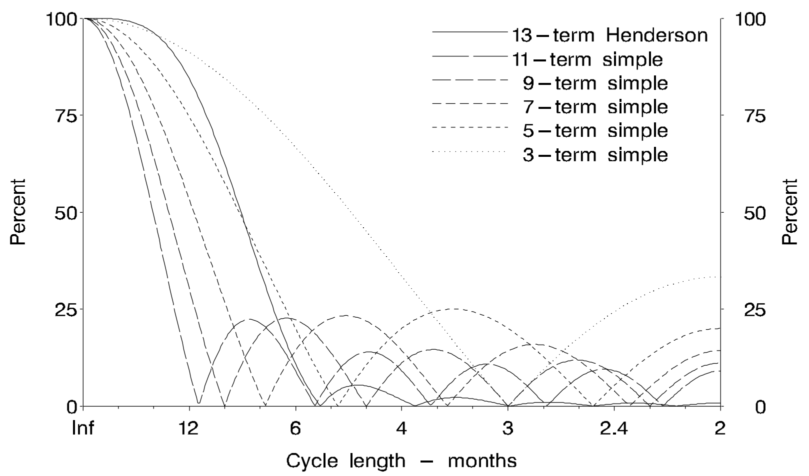

In order to develop a collection of surrogate Henderson filters (also referred to as simply surrogate filters) which can be used to calculate trend estimates all the way to the end of the seasonally adjusted series, it is necessary to consider non-symmetric moving averages. To ensure that the characteristics of the trend estimates produced by these surrogate filters closely approximate the characteristics of the trend estimates that will be produced by applying the main Henderson filter when additional data becomes available, the preferred class of non-symmetric moving averages should approximate the cycle damping properties of the main Henderson filter. The problem, however, is that non-symmetric moving averages also introduce phase shift to varying degrees over a range of cycles, including the longer cycles associated with the trend, and the non-symmetric moving averages with the best cycle damping properties tend to exhibit the worst phase shifting properties. Phase shift can also result in the surrogate trend

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

estimates having different characteristics to the final Henderson based estimates of trend (although as with cycle damping properties, this will depend on the amount of phase shift introduced to the longer cycles actually present in a given series), and can therefore also result in revision to the estimates of trend. The aim, therefore, is to find a collection of non-symmetric moving averages which strike a compromise between having similar cycle damping properties to the symmetric Henderson filter, and introducing a minimal amount of phase shift, and the revision that occurs between the surrogate and symmetric Henderson filter of trend estimates. The collection of surrogate filters used by the ABS for smoothing seasonally adjusted series to determine estimates of trend at the end of the series are designed to achieve this goal.

The basis for the weighting patterns for the surrogate filters used by the ABS were explained by Doherty (2001). They are specified under the following conditions: Firstly, the weights of each surrogate filter must sum to one (if the weights of a moving average sum to one, then the application of the moving average to a constant series will preserve the level of the series, which is a highly desirable property). The second condition upon the weights of a given surrogate filter arises by requiring that the mean square revision between the trend produced by the surrogate filter and the trend produced by the main Henderson filter is to be minimal (that is, the amount of revision should be as small as possible). These two conditions are used to generate the weights for the surrogate filters under the added assumption that the trend can be approximated by a linear path over the relatively short sub spans to which the filters are applied (this just means that the trend is assumed to be linear at the end of the series). The formula for calculating the weights for each surrogate moving average for a given 2m+1 term Henderson filter is rather complex in algebraic terms, and is consequently given in Appendix 4. This formula defines the weights for the surrogate filters for each Henderson filter of 7 terms and longer, the details for the 5-term Henderson filter will be discussed later in section 5.2.

In practical terms, the value of each weight in a surrogate filter is a function of a number of factors:

- the length of the Henderson filter for which the surrogate filter is being determined
- a subset of the weights of the Henderson filter for which the surrogate filter is being determined
- the number of terms in the surrogate filter
- a measure of the irregularity of the source data, compared to its trend behaviour.

For a Henderson filter of length 2m+1, there are m surrogate filters needed to estimate the trend at the end of the series. Consequently, although only one 13-term Henderson filter is required for smoothing a monthly series, a set of six surrogate filters is required to attend to its end-point problem. That is, a separate surrogate filter is required to determine an estimate of trend for each of the six final time periods in the seasonally adjusted series, for which no trend estimate could be calculated using the 13-term Henderson filter. Similarly, a set of three surrogates is required when using a 7-term Henderson filter to produce trends for quarterly data. When using a 33-term Henderson filter to produce trends, as is the case with the ABS composite leading indicator referred to in section 4.2.2, a set of 16 surrogates is required to determine trend estimates at the end of the seasonally adjusted series.

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

The weights for the m surrogate filters for a given 2m+1 term Henderson filter are not completely specified without knowing the relative amount of variability in the series due to irregularity, as opposed to trend behaviour. Consequently, the 'best' set of surrogate filters to be used with a particular Henderson filter depends on the series to be smoothed.

To measure the relative degree of irregular to trend variability in a series, a simple statistic known as the 'end weight parameter' is calculated (this value is also often known as the $\bar{I}/\bar{C}$ ratio). For a time series which fits the additive model, as described in section 2.1.4, the end weight parameter is computed as the average absolute period to period change of the irregular component, I, divided by the average absolute period to period change of the trend, T (where the irregular component is obtained by subtracting the trend from the seasonally adjusted series). For a time series which fits the multiplicative model, the end weight parameter is computed as the average absolute percentage period to period change of the irregular component, I, divided by the average absolute percentage period to period change of the trend, T (where the irregular component is obtained by dividing the seasonally adjusted series by the trend). The majority of series analysed by the ABS display the characteristics of the multiplicative time series model, and consequently, the latter method is generally used to calculate the end weight parameter. Regardless of which model is the most appropriate, the end weight parameter is a value which reflects how much of the variability in the seasonally adjusted series is a result of irregularity compared to trend behaviour. A numerically low value of the end weight parameter indicates that the series' period to period movements are generally driven by the trend, whereas a numerically large value of the end weight parameter indicates that the irregular fluctuations tend to determine change in the seasonally adjusted series from one period to the next.

The end weight parameter for a given series is used to specify the best set of surrogate filters to be used in conjunction with a particular length Henderson filter for that series. In practice, the end weight parameter does not play such a crucial role for the longer Henderson filters and therefore, for a particular length Henderson filter with more than 7 terms, the ABS generally uses a single 'generic' value of the end weight parameter for all series for convenience. The 'generic' values of the end weight parameter used by the ABS with some of the most commonly used Henderson filters of length 9 and above, are given in table 5.1. The surrogate filters specified by using these particular values of the end weight parameter will often be referred to as the 'conventional' surrogate filters. As long as a series specific end weight parameter ($\bar{I}/\bar{C}$ ratio) is not too different from the conventional value, the amount of revision occurring between the conventional surrogate trend estimates and the Henderson based trend estimates will not be substantially larger than that using the series specific end weight parameter.

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

TABLE 5.1 VALUES OF THE END WEIGHT PARAMETER ASSUMED BY CONVENTIONAL SURROGATE FILTERS

Number of Terms in Henderson Filter	Value of end weight parameter
9	1.0
13	3.5
15	4.5
17	4.5
23	4.5
33	6.5

The given value of the end weight parameter completely specifies the set of weights for each of the m surrogates associated with a particular length Henderson filter. In the case of the 13-term Henderson filter the set of conventional surrogate filter weights are determined by taking the end weight parameter to be 3.5. The weighting patterns for each of the conventional surrogate filters to be used in conjunction with the 13-term Henderson filter are given in table 5.2.

Each row in table 5.2 gives the weights of a conventional surrogate filter for the 13-term Henderson filter, and specifies the time period for which the estimate of trend is being calculated (where the most current time period for which data is available is period N). As with the Henderson filter, each surrogate filter calculates an estimate of trend by weighting and adding together various values from the seasonally adjusted series, and the columns specify the seasonally adjusted series value to which a given weight is applied. Thus, considering the rows in the reverse order for convenience, the bottom row indicates the weighting pattern of the filter used to determine the estimate of trend for period N–6, which is the 7th last time period for which data is available. It will be recalled that this is the last time period for which an estimate can be calculated using the 13-term Henderson filter, and indeed, the weighting pattern specified in the bottom row of the table is the weighting pattern of the 13-term Henderson filter. Here, the 13 seasonally adjusted values from period N–12 up to period N are being weighted, with the smoothed result recorded against the seventh of these 13 observations, namely, period N–6.

The second last row gives the surrogate filter weights for determining the estimate of trend at period N–5, that is, the sixth last time period for which data is available. This estimate is determined by weighting the 12 last seasonally adjusted values, from period N–11 up to period N. Similarly, the third last row gives the surrogate filter weights for determining the estimate of trend at period N–4, the fifth last time period for which data is available. This estimate is determined by weighting the 11 last seasonally adjusted values, from period N–10 up to period N. Continuing up the table, the very top row gives the surrogate filter used to determine the estimate of trend for the most current time period, period N, and is based on the last 7 seasonally adjusted series values. Whereas the 13-term Henderson filter in the bottom row can be seen to be symmetric about the time period against which the smoothed result is recorded, the surrogate filters can be seen to become progressively less symmetric about the time period to which they refer

5.1 THE ABS APPROACH TO DETERMINING END-POINT ESTIMATES *continued*

(indicated by the diagonal bar across table 5.2), as they become shorter and approach the end of the series; the shortest moving average is in fact completely one sided.

TABLE 5.2 WEIGHTS OF 13-TERM HENDERSON AND CONVENTIONAL SURROGATES

Weights for period	*N–12* Period	*N–11* Period	*N–10* Period	*N–9* Period	*N–8* Period	*N–7* Period	*N–6* Period	*N–5* Period	*N–4* Period	*N–3* Period	*N–2* Period	*N–1* Period	*N* Period
N	—	—	—	—	—	—	–0.09	–0.06	0.01	0.12	0.24	0.35	0.42
N–1	—	—	—	—	—	–0.04	–0.04	0.00	0.08	0.17	0.25	0.29	0.28
N–2	—	—	—	—	–0.02	–0.03	0.00	0.07	0.15	0.22	0.24	0.22	0.15
N–3	—	—	—	–0.01	–0.02	0.00	0.07	0.15	0.21	0.23	0.20	0.13	0.05
N–4	—	—	–0.01	–0.02	0.00	0.07	0.15	0.21	0.24	0.21	0.14	0.05	–0.02
N–5	—	–0.02	–0.03	0.00	0.07	0.15	0.21	0.24	0.21	0.14	0.06	–0.01	–0.03
N–6	–0.02	–0.03	0.00	0.07	0.15	0.21	0.24	0.21	0.15	0.07	0.00	–0.03	–0.02

— nil or rounded to zero (including null cells)

In some applications, the use of the conventional surrogate filters is not appropriate because they are based on assumed degrees of relative irregular variation that are sometimes too different from the series' actual behaviour. An example for monthly series are some of the Retail Trade series, which show a much lower degree of irregularity than is implied by an end weight parameter of 3.5. For these series, an improvement in the revision properties of the trend values estimated by the surrogate moving averages can be achieved by using a smaller end weight parameter of 2.0. These are the only monthly series for which the ABS is not currently using the conventional surrogate filters for the 13-term Henderson filter.

The choice of end weight parameter has been found to be more crucial in determining surrogate trend estimates for the shorter Henderson filters. As a result, the ABS does not use a 'generic' end weight parameter to specify the surrogate filters to be used in conjunction with the 7-term Henderson filter, which is the filter used by the ABS for smoothing quarterly series. Instead, tailored surrogate filters are used, that is, the surrogate filters for a given series are specified by using an end weight parameter which is more indicative of the degree of relative irregular to trend variability of that series. The ABS is also considering a similar approach towards smoothing at the ends of monthly series. Tailoring the surrogate filters to the specific end weight parameter of a series will result in a smaller degree of revision between the surrogate trend estimates and the main Henderson filter estimates. However, this practice has some disadvantages with regard to maintaining various accounting identities at the topical end of the series, as will be discussed in following sections.

5.2 AN ALTERNATIVE VIEW OF THE SURROGATE TREND ESTIMATES

The conditions outlined in section 5.1 are used to generate surrogate filters for calculating estimates of trend at the end of the seasonally adjusted series. However, there is an alternative view of the derivation of these surrogate trend estimates which helps to explain some of their properties. Earlier it was noted that the ABS does not explicitly forecast the future seasonally adjusted data, and then apply the full length symmetric Henderson moving averages to the combined actual and forecast data to

5.2 AN ALTERNATIVE VIEW OF THE SURROGATE TREND ESTIMATES *continued*

derive the end point trend estimates. However, given the weights of a surrogate filter are derived from an implied trend forecast based on minimising revision criteria. Thus, even though explicit forecasts are not made, they exist implicitly once the surrogate end filters have been chosen. Note however that the intention for choosing the surrogate end filters (and hence the implicitly forecast future trend values) was to minimise revisions to the preliminary trend estimates and not (necessarily) to produce good forecasts in their own right. This projection process is different for each surrogate trend estimate and the set of values used depends on the length of the surrogate filter.

In general, for a 2m+1 term Henderson filter, the first surrogate filter trend estimate is equivalent to the trend estimate that would be obtained by using the last 2m seasonally adjusted values to extend the series one period into the future, and then applying the 2m+1 term Henderson filter, the second surrogate filter trend estimate is equivalent to the trend estimate that would be obtained by using the last 2m–1 seasonally adjusted values to extend the series two periods into the future, and then applying the 2m+1 term Henderson filter, and so on. Thus the 'base' set of seasonally adjusted values used to extend the series becomes progressively smaller as trend estimates are calculated for periods progressively closer towards the end of the series. For each surrogate trend estimate, the 'base' set of seasonally adjusted values used to extend the series and apply the Henderson filter is the set of seasonally adjusted values to which the associated surrogate filter is applied. As a result, it can be seen that the length of the surrogate filters become progressively shorter (from length 2m to length m+1) as they approach the end of the series.

An advantage in associating the implicit projection process with the calculation of surrogate trend estimates is that it helps to explain why the surrogate trend estimates are subject to revision as additional data becomes available; for if the implicit projections are not perfectly accurate, then replacing the projected values with real data will change the Henderson moving average calculations. The way in which the ABS deals with revision to its trend estimates at the current end of the series will be discussed in Chapter 6.

To date, no mention has been made of the surrogate filters for the 5-term Henderson filter, and this is because these surrogate filters are derived differently to the manner described above. The 5-term Henderson surrogates can be derived (rounded to three decimal places) by assuming the 5-term Henderson filter is applied to an extended seasonally adjusted series that follows a path which is simply the average of the last two seasonally adjusted observations. Such an assumption about a series' general behaviour is inappropriate for many series, and surrogate filters so based display significant bias when the series being smoothed are growing or declining. As a result of the unsatisfactory nature of the 5-term Henderson surrogates, the ABS does not use the 5-term Henderson filter for producing trends for publication.

5.3 PROPERTIES OF SOME CONVENTIONAL SURROGATE FILTERS

The surrogate filters defined by the weights in Appendix 4 and used by the ABS as surrogates to the Henderson filters, are designed so as to have cycle damping properties which are as similar as possible to the main Henderson filter, while simultaneously introducing as little phase shift as possible. It is therefore interesting to examine the cycle reduction and phase shifting properties of some of these surrogate filters. In particular, the cycle reduction and phases shifting properties of the six conventional surrogate filters for the 13-term Henderson filter will be considered. These are the

5.3 PROPERTIES OF SOME CONVENTIONAL SURROGATE FILTERS *continued*

surrogate filters generally used in conjunction with the 13-term Henderson filter for determining estimates of trend from seasonally adjusted monthly series.

In graphs 5.3 and 5.4, the cycle strength reduction properties, or gain, of the conventional 13-term Henderson surrogate filters are compared with the symmetric 13-term Henderson filter. For clarity, only the surrogate filters used for estimating the last three smoothed values are compared to the 13-term Henderson filter in graph 5.3. The other surrogate filters, which are used to estimate the sixth, fifth and fourth last values, lie predominantly between the 13-term Henderson filter and the third last surrogate filter, and are shown in graph 5.4. It can be seen from graphs 5.3 and 5.4 that the filters that most poorly approximate the 13-term Henderson filter with regards to cycle damping are the surrogate filters used to estimate the last two trend values, while the surrogate filter used to estimate the third last trend value is a fair approximation to the 13-term Henderson filter. The surrogate filters used to estimate the sixth, fifth, and fourth last trend values approximate the 13-term Henderson filter very closely, and for most practical purposes, these filters are sufficiently good approximations of the 13-term Henderson filter to be regarded as equivalents with respect to cycle strength reduction.

GRAPH 5.3 GAIN FUNCTION OF 13-TERM HENDERSON AND LAST THREE CONVENTIONAL SURROGATES

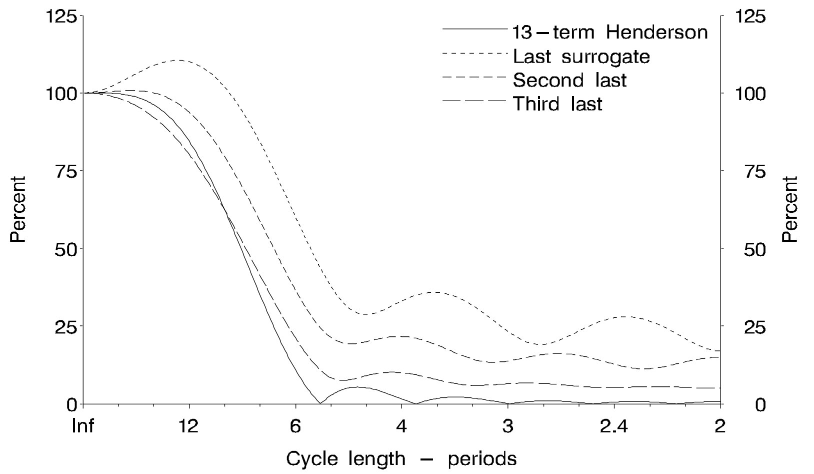

5.3 PROPERTIES OF SOME CONVENTIONAL SURROGATE FILTERS *continued*

GRAPH 5.4 EFFECT OF 13-TERM HENDERSON AND FIRST THREE CONVENTIONAL SURROGATES ON CYCLES

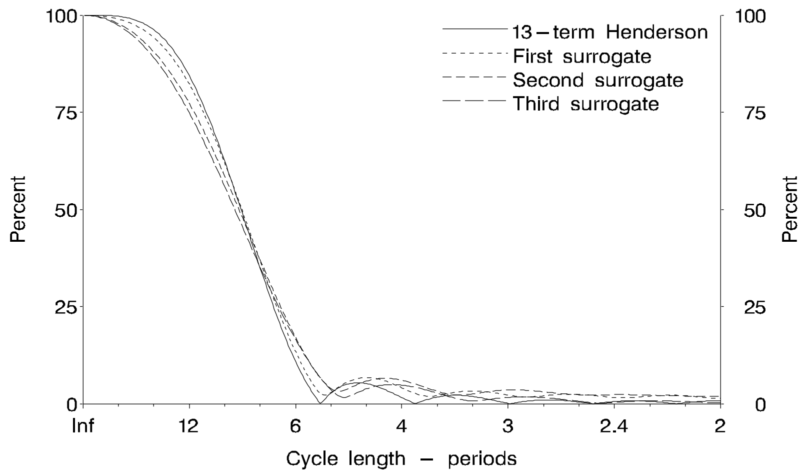

Graph 5.5 illustrates the phase shifting effects of some of the conventional surrogates for the 13-term Henderson filter. For presentational ease, only the phase shifts of cycles that remained in the filtered series with at least ten per cent of their strength are shown. This is because by reducing the strength of cycles in the short cycle range, Henderson filters also minimise the impact of any phase shifting on the filtered series, as referred to in section 4.1.2.

It can be seen that the non-symmetric moving average used for estimating the third last smoothed value does not displace cycles longer than six months by more than one week. In particular, the long cycles that make up the trend behaviour are not significantly phase shifted at all. The non-symmetric moving average used to estimate the second last smoothed value also introduces only a small amount of phase shift for the larger cycles associated with trend; the amount of phase shift introduced is less than two weeks for cycles longer than about ten months. When considering the non-symmetric moving average used for estimating the last, or most recent, smoothed value, it can be seen that cycles of five to ten months long are phase shifted by at most five weeks; shorter cycles are shifted by much lesser amounts, while cycles longer than ten months are phase shifted by about two to four weeks. The non-symmetric moving averages, used for estimating the sixth, fifth and fourth last smoothed values, have only marginal phase shifting effects in relation to cycles remaining with more than ten per cent of their strength.

When considering the phase shifting effect of a moving average, the strength that a particular phase shifted cycle would have in the smoothed data should be kept in perspective. Graph 5.5 indicates that the cycles experiencing the greatest amount of phase shift are in the higher frequency range where the greatest amount of reduction occurs.

5.3 PROPERTIES OF SOME CONVENTIONAL SURROGATE FILTERS *continued*

GRAPH 5.5 PHASE SHIFT OF CONVENTIONAL SURROGATES

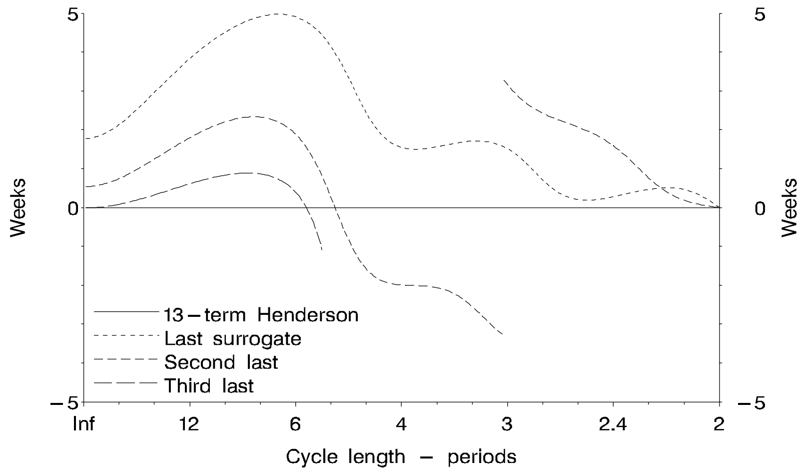

5.4 AN ALTERNATIVE SET OF SURROGATE FILTERS

While the surrogate filters discussed in section 5.1 are derived so as to minimise the amount of revision occurring in the trend estimates at the end of the series, they are biased with respect to the Henderson based estimates of trend that will be calculated for these periods when sufficient new data becomes available (thus each surrogate trend estimate is not necessarily the best available proxy for the final Henderson based estimate, although it is the least sensitive to change as new data becomes available). Furthermore, the weights of these surrogate filters are determined under the assumption that the series being filtered can be approximated by random variation about a linear path over the last part of the series. This is equivalent to saying that at the end of the series, the trend follows a linear path, offset in the seasonally adjusted series by random or irregular influences. If the trend is not following a linear path at the end of the series, but rather a quadratic or cubic polynomial path, then these estimates of trend can display even more bias with respect to the Henderson based estimates of trend. This is because the Henderson filters are generated under a different set of criteria, which include preserving the cubic and quadratic polynomial trend paths often observed to occur within socioeconomic time series. Consequently, an alternative approach towards determining surrogate filters for smoothing the end of a series, is to attempt to minimise the amount of bias with respect to the Henderson based estimate of trend. This can be achieved by using exactly the same set of criteria as are used to specify the weights of the Henderson filters.

For a given length Henderson filter, a set of surrogate filters can be determined which will each preserve a cubic polynomial trend while simultaneously smoothing out irregularity. Again, this can be achieved by calculating the weights for a given filter so as to minimise the sum of squares of the third differences of the weights, subject to the filter leaving cubic polynomial equations unaltered. The surrogate filters that arise using these conditions are called the BLUE (Best Linear Unbiased Estimate) surrogate filters (see Alistair Gray (1996a, b)), as they apply the best set of linear weights to the data at the end of the series, for producing estimates of trend which are unbiased with respect to the Henderson based trend estimates. While the BLUE filter based estimates of trend

5.4 AN ALTERNATIVE SET OF SURROGATE FILTERS *continued*

are unbiased in relation to the Henderson based trend estimates, they have, in most applications, a very high and generally unacceptable degree of revision.

For interest, table 5.3 compares the weighting pattern of the BLUE surrogate filter used to estimate the current monthly trend estimate, with the weighting pattern of the conventional surrogate filter used to determine the same trend estimate (these surrogates are both associated with the 13-term Henderson filter, which is the filter used to smooth seasonally adjusted monthly series). Note the difference in the weighting patterns, particularly at the end of the series, period N.

TABLE 5.3 COMPARISON OF WEIGHTS FOR 13-TERM HENDERSON SURROGATE FILTERS

	N–6	*N–5*	*N–4*	*N–3*	*N–2*	*N–1*	*N*
Weight	Period	Period	Period	Period	Period	Period	Period
Conventional	–0.09	–0.06	0.01	0.12	0.24	0.35	0.42
BLUE	–0.11	0.18	0.14	–0.28	–0.25	0.49	0.82

In graphs 5.6 and 5.7, the damping and phase shifting properties of these two surrogate filters are compared with the properties of the 13-term Henderson filter for which they are designed to be surrogates. Graph 5.6 shows that the BLUE surrogate filter has distinctly different cycle strength reduction properties from the 13-term Henderson filter and the conventional surrogate. On the other hand, graph 5.7 shows that it generally displays less phase shifting effects than the conventional surrogate filter, and in this sense is more like the Henderson filter. For presentational ease, only the phase shifts of cycles that could remain in the filtered series with at least 10 per cent of their strength are shown.

GRAPH 5.6 GAIN FUNCTION OF A BLUE SURROGATE MOVING AVERAGE

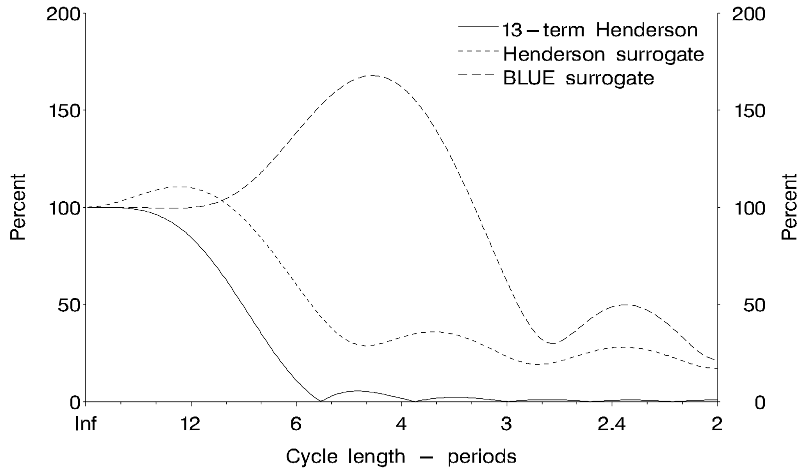

5.4 AN ALTERNATIVE SET OF SURROGATE FILTERS *continued*

GRAPH 5.7 PHASE SHIFT OF A BLUE SURROGATE MOVING AVERAGE

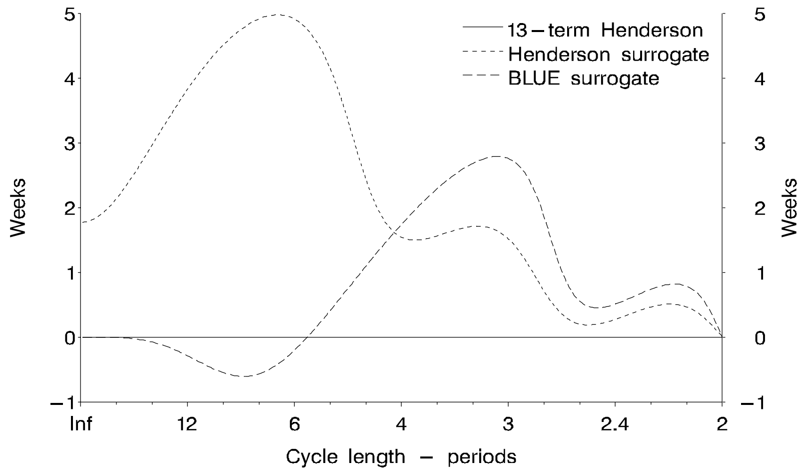

5.5 SUMMARY

Before going on to discuss how the ABS deals with revision to the trend estimates at the end of the series, it is useful to summarise the requirements considered when deciding upon a set of filters for smoothing seasonally adjusted series. When considering the appropriateness of a symmetric moving average and its attendant set of surrogate filters, the following, often conflicting criteria, must be balanced.

1. The moving average(s) must preserve the cycles expected to be in the trend series and remove or reduce other cycles. For example, if it was agreed that a trend estimate should not contain the noticeable influence of cycles shorter than nine periods, then it would not be sensible to use a simple 3-term moving average to produce estimates of trend (a simple 3-term moving average would pass large amounts of the strength of cycles shorter than nine periods, as can be verified from graph 5.2).
2. The phase shift of the longer cyclical behaviour of interest should be minimal, so that turning points and points of inflection may be detected in a reliable and timely fashion. This is the major reason for preferring the use of symmetric moving averages wherever possible, as the centring of the symmetric moving average values will result in minimal phase shift.
3. The extent of revision between the trend estimate based on the symmetric filter and the surrogate filters should be small. While the revision would be relatively small if the trend was taken to be the average path through the entire set of time series observations, this would not be a very useful estimate of trend. Alternatively, a one sided moving average could be applied along the entire length of the series, to determine estimates of trend which were subject to zero revision upon the availability of additional data. Such a procedure would, however, have significant phase shifts associated with its estimates and would therefore be misleading with respect to the timing of turning points. Consequently, the desire to minimise revision must also take into account the other criteria, as well as the suitability of the resulting trend series.
4. At the end of the series, the bias between the trend estimates calculated using the surrogate filters and the final symmetric filter based estimates of trend should

5.5 SUMMARY *continued*

be small. Again, a compromise has to be made here, because a lower bias usually means that some of the other criteria are less well met. For example, the BLUE surrogates discussed previously have zero bias, however they have a degree of revision that is generally regarded as unacceptable.

5. Finally, the smoothing procedure should be capable of tracking the types of curvatures present in the series, both historically and at the ends, so that the fundamental shape and level of the series turning points and points of inflection are preserved.

In its many years of experience in this area, the ABS has found its current trending procedure to be a robust and resource effective procedure for trending a multitude of time series. As mentioned earlier (in section 5.2), asymmetric weights of surrogate Henderson filter are derived from the implied linear prediction of future trend estimates, based on the minimising revision criteria. Recent ABS research shows that applying symmetric Henderson filters to data which has been forecasted using AutoRegressive Integrated Moving Average (ARIMA) models generally produces a smaller revision than using the current surrogate Henderson filter method does. Further research is being conducted with a view to adopt this technique in the future for series where it produces better trend estimates.

CHAPTER **6**

REVISIONS

CONTENTS

page

SECTION

6.0 OVERVIEW

As discussed in Chapter 4, the ABS estimates the trend component of original series behaviour, wherever possible, by smoothing seasonally adjusted series with Henderson moving averages. Henderson filters are centred in order to avoid phase shift in the longer cycles associated with the trend and the consequential misrepresentation of the timing of pivotal changes in trend behaviour. Therefore there are a number of time periods at the end of each series for which Henderson based trend estimates cannot be determined. As discussed in Chapter 5, the ABS overcomes this problem by using a set of non-symmetric surrogate moving averages to determine trend estimates for the most recently available original series observations. When additional original series values become available, however, it is possible to recalculate estimates of trend for these periods by applying the Henderson filter to the updated seasonally adjusted series. Consequently, the 'provisional' surrogate estimates of trend can, in time, be replaced by the Henderson based estimates they are designed to be substitutes for, giving rise to revisions in the trend estimates at the current end of the series.

The current end of the series is particularly important for those attempting to monitor the impact of socioeconomic forces. In particular, the behaviour of the trend for the most recent time periods is often used to indicate the current state of the economy, or the immediate impact of social or economic reforms. The nature of the trending process, however, means that the potential impact of revisions to the most recently calculated estimates of trend must be taken into account. This Chapter discusses some of the issues that must be considered. In particular, section 6.1 discusses the impact of revisions to the trend, ABS revision practices, and the likely extent of these revisions, while section 6.2 investigates various methods for evaluating the reliability of the provisional trend estimates.

6.1 THE NATURE OF REVISIONS

As will be recalled from the previous Chapter, trend estimates for monthly series are determined by smoothing the seasonally adjusted series with a 13-term Henderson filter in conjunction with a set of six surrogate filters. These surrogate filters are used to calculate the trend estimates for the last six periods because there is insufficient data to undertake a centred 13-term Henderson moving average calculation for these periods. When sufficient additional data has become available however, these estimates of trend can be replaced by the Henderson based estimates for which they are designed to be substitutes. For example, suppose trend estimates are initially determined on the basis of seasonally adjusted data available up to month N. The last trend estimate that can be calculated at this time using the 13-term Henderson filter is the estimate for period N–6, the remaining six estimates all being calculated using the surrogate filters. When seasonally adjusted data becomes available for month N+1 however, month N–5 becomes the seventh of the last thirteen months for which data is available, and therefore the trend estimate initially calculated for this period using the first surrogate filter, can be recalculated using the 13-term Henderson filter.

The addition of one month of data therefore gives rise to revisions in the last six trend estimates and results in estimates which are both more indicative of the final Henderson based trend estimates that will be calculated when sufficient data becomes available, and more consistent with respect to the existing historical Henderson based estimates of trend. While the estimate of trend for period N–5 is only revised once before there is enough data available to undertake the 13-term Henderson calculation for this period, it

6.1 THE NATURE OF REVISIONS *continued*

takes six months, and therefore six revisions, before the trend estimate for period N can be calculated in this way. Once sufficient data for determining the Henderson based estimate for this period has become available however, no further revisions will occur as a result of the trending process.

Although the discussion so far has centred on the nature of revisions to monthly series, quarterly series experience a similar problem at the end of the series. As quarterly series are smoothed using a 7-term Henderson filter, only the three most recent periods are affected by the end-point problem, and therefore only three surrogate filters are employed. Consequently, the most current estimate is revised only three times before the 7-term Henderson based estimate of trend can be calculated, and therefore it only takes three-quarters for the most provisional estimate of trend to be revised into a stable form.

6.1.1 The impact of revisions

One means of gaining an appreciation of the impact of revisions in the most current trend estimates, is to simulate the performance of the smoothing procedure about an important turning point in a major socioeconomic indicator. As a case study, the smoothing of the seasonally adjusted Unemployed Persons series has been simulated through 1983. This simulation is carried out by calculating the trend estimates on the basis of seasonally adjusted data available up to January 1983, and then determining the additional trend estimates and revised estimates that can be calculated as each additional month of seasonally adjusted data becomes available. Note that seasonally adjusted series are also subject to revision, although only annually when the reanalysis of seasonal influences takes place. This simulation assumes the last seasonal reanalysis occurred in January 1983. Thus the data up to January 1983 has been seasonally adjusted, and the seasonal factors determined from this process have been used to derive the seasonally adjusted values for the rest of 1983 (as is the procedure for determining seasonally adjusted values for the months leading up to the next annual seasonal reanalysis). This guarantees that the revisions to the trend estimates are a direct result of the nature of the trending procedure at the end of the series, and not the result of subsequent changes to the seasonally adjusted values from which the trend estimates are derived.

The trend results from January 1983 are recorded in table 6.1. The value of 642.8 in the first row gives the estimate of trend for January 1983, based on the seasonally adjusted data available up to this point in time. This is the first and most provisional trend estimate for January 1983 and is calculated by using the sixth conventional surrogate filter, which is the shortest and most non-symmetric of the surrogate filters given in table 5.2. The second row gives the trend estimates that can be calculated with the addition of the February seasonally adjusted observation. The first value of 648.2 is the revised, or second estimate of trend for January, and is calculated by applying the fifth conventional surrogate filter to the seasonally adjusted data up to February. The value of 667.8 is then the first trend estimate for February 1983, and is calculated by applying the sixth conventional surrogate filter to the seasonally adjusted data up to February. The remaining rows then give the trend estimates that can be calculated as each subsequent month of seasonally adjusted data becomes available throughout 1983.

TABLE 6.1 UNEMPLOYED PERSONS SIMULATION, Trend series with all subsequent revisions to monthly trend estimates ('000)

	Jan 1983	*Feb* 1983	*Mar* 1983	*Apr* 1983	*May* 1983	*Jun* 1983	*Jul* 1983	*Aug* 1983	*Sep* 1983	*Oct* 1983	*Nov* 1983	*Dec* 1983
Jan	642.8	—	—	—	—	—	—	—	—	—	—	—
Feb	648.2	667.8	—	—	—	—	—	—	—	—	—	—
Mar	650.6	673.1	690.2	—	—	—	—	—	—	—	—	—
Apr	651.2	673.7	693.8	708.8	—	—	—	—	—	—	—	—
May	650.8	673.7	693.6	711.4	725.7	—	—	—	—	—	—	—
Jun	652.5	673.9	692.4	707.8	720.2	729.0	—	—	—	—	—	—
Jul	653.3	675.3	692.5	707.3	718.4	726.9	732.4	—	—	—	—	—
Aug	653.3	676.2	694.0	707.5	717.6	723.8	728.1	729.5	—	—	—	—
Sep	653.3	676.2	694.7	708.4	717.6	723.5	727.1	730.0	731.6	—	—	—
Oct	653.3	676.2	694.7	709.7	719.8	724.2	725.0	723.7	721.4	718.1	—	—
Nov	653.3	676.2	694.7	709.7	721.1	725.3	725.8	721.7	715.8	708.7	701.0	—
Dec	653.3	676.2	694.7	709.7	721.1	728.1	728.7	723.1	712.6	699.3	684.1	669.6

— nil or rounded to zero (including null cells)

6.1.1 The impact of revisions continued

Table 6.1 clearly illustrates the stabilising of the trend estimates for each month as more data becomes available. In particular, as discussed earlier, it takes six months from the time the most non-symmetric surrogate filter is used to determine a trend estimate, before the 13-term Henderson filter can be applied. Once there is enough data available to determine the trend estimate for a particular month using the 13-term Henderson filter, there are no further revisions to the trend estimates brought about by the trending procedure. A revision of these values will now only take place if the seasonally adjusted data used to determine the trend estimates is revised, such as when an annual seasonal reanalysis takes place, or if the original data from which the seasonally adjusted values are determined is changed.

In order to appreciate the impact of these revisions on the trend series, consider graph 6.1 in which the subsequent estimates of the smoothed series are represented pictorially. From this graph it can be seen that the revisions to the estimates are significant enough to change the apparent behaviour of the trend series at the current end. In particular, until data for October 1983 becomes available, it appears that the trend is continuing to show an increase throughout July, August and September, albeit at a much slower rate. Once the value for October has become available, however, the revised estimates have indicated the presence of a turning point in July/August, an impression which is not altered by the addition of subsequent data. It is possible, however, for the Henderson based estimate of the trend initially to identify inaccurately the timing of the turning point. In such an instance, as more data becomes available the timing of the turning point will be revised to eventually indicate the correct date.

6.1.1 The impact of revisions continued

GRAPH 6.1 REVISIONS TO TREND ESTIMATES FOR UNEMPLOYED PERSONS

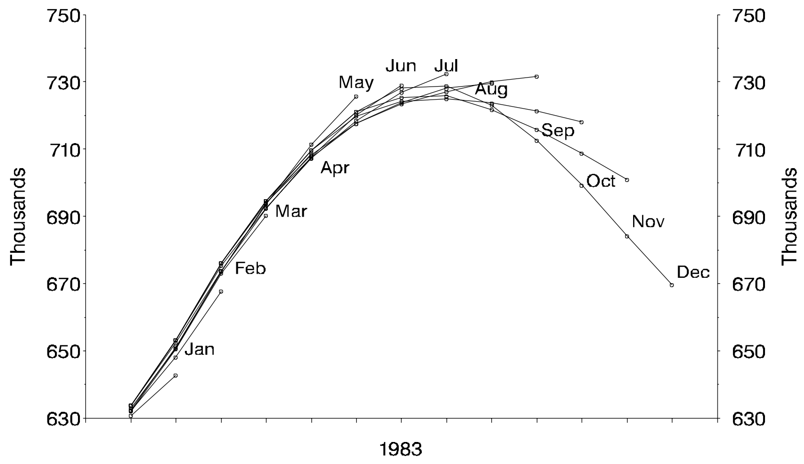

6.1.2 Revision policy and practices

As seen from table 6.1, the use of surrogate filters at the current end of the series gives rise to a number of revisions to the estimates of trend for the most recent time periods. In addition, the annual seasonal reanalysis can give rise to revisions in the seasonally adjusted data from which the trend estimates are derived, resulting in further revisions to the estimates of trend. As a result, the question arises as to how often the ABS should publish revised estimates, in particular, whether all consequential revisions should be published as they occur, or whether they should be recognised regularly but less frequently. The following discussion will consider a number of possible revision policies and the impact these policies could have on the accurate detection of changes in the behaviour of the trend.

One possible policy is not to revise any of the initial trend estimates until the annual seasonal reanalysis takes place. Thus, for the period of a year, the only change to a monthly trend series would be the addition of each month to the current trend estimate. In terms of the above example, this policy would involve publishing only the estimates that appear on the first diagonal of table 6.1. Table 6.2 shows the trend series (as of the annual seasonal reanalysis based on data to January 1983) that would be published each month. Each of these trend estimates is determined at the given point in time using the sixth conventional surrogate filter, which is the conventional surrogate filter that most poorly approximates the 13-term Henderson filter with regard to cycle damping and phase shift. Furthermore, these trend estimates will not be revised until the next annual seasonal reanalysis takes place, which in this case study was scheduled for January 1984. Consequently, under this policy, those monitoring behaviour at the current topical end of the series would have to rely on the provisional trend estimates that most poorly approximate the Henderson based estimates of trend for up to a year.

TABLE 6.2 UNEMPLOYED PERSONS SIMULATION, Trend series with NO subsequent revisions to monthly trend estimates ('000)

Month in which trend series calculated	*Jan* 1983	*Feb* 1983	*Mar* 1983	*Apr* 1983	*May* 1983	*Jun* 1983	*Jul* 1983	*Aug* 1983	*Sep* 1983	*Oct* 1983	*Nov* 1983	*Dec* 1983
Jan	642.8	—	—	—	—	—	—	—	—	—	—	—
Feb	642.8	667.8	—	—	—	—	—	—	—	—	—	—
Mar	642.8	667.8	690.2	—	—	—	—	—	—	—	—	—
Apr	642.8	667.8	690.2	708.8	—	—	—	—	—	—	—	—
May	642.8	667.8	690.2	708.8	725.7	—	—	—	—	—	—	—
Jun	642.8	667.8	690.2	708.8	725.7	729.0	—	—	—	—	—	—
Jul	642.8	667.8	690.2	708.8	725.7	729.0	732.4	—	—	—	—	—
Aug	642.8	667.8	690.2	708.8	725.7	729.0	732.4	729.5	—	—	—	—
Sep	642.8	667.8	690.2	708.8	725.7	729.0	732.4	729.5	731.6	—	—	—
Oct	642.8	667.8	690.2	708.8	725.7	729.0	732.4	729.5	731.6	718.1	—	—
Nov	642.8	667.8	690.2	708.8	725.7	729.0	732.4	729.5	731.6	718.1	701.0	—
Dec	642.8	667.8	690.2	708.8	725.7	729.0	732.4	729.5	731.6	718.1	701.0	669.6

— nil or rounded to zero (including null cells)

6.1.2 Revision policy and practices continued

Relying on these provisional trend estimates can substantially delay the accurate detection of turning points in the trend. Consider the simulation for the Unemployed Persons series again. Table 6.3 gives the month to month movements of the trend series published under this policy, as they would be calculated each month. To provide a benchmark series of movements for comparison, table 6.3 also contains the month to month movements of the stable 1983 trend estimates calculated as of February 1998. Whereas the benchmark movements indicate that the turning point actually occurred in July/August of 1983, the movements of the trend estimates published under this policy would indicate that the turning point occurred between September and October 1983, two months later than it actually did. This perception would not have been amended until revised trend estimates became available, shortly after the next annual seasonal reanalysis. Therefore, under this policy, users would not have been provided with an indication that there was a turning point between July and August 1983, until around February 1984. In other circumstances, such as when the next seasonal reanalysis has been scheduled to take place later in the year, it is possible for this error period to be longer.

TABLE 6.3 UNEMPLOYED PERSONS SIMULATION, Trend series with NO subsequent revisions to monthly trend estimates (per cent)

Month in which trend series of movements calculated	*Feb*	*Mar*	*Apr*	*May*	*Jun*	*Jul*	*Aug*	*Sep*	*Oct*	*Nov*	*Dec*
Jan	25.0	—	—	—	—	—	—	—	—	—	—
Feb	25.0	—	—	—	—	—	—	—	—	—	—
Mar	25.0	22.4	—	—	—	—	—	—	—	—	—
Apr	25.0	22.4	18.6	—	—	—	—	—	—	—	—
May	25.0	22.4	18.6	16.9	—	—	—	—	—	—	—
Jun	25.0	22.4	18.6	16.9	3.3	—	—	—	—	—	—
Jul	25.0	22.4	18.6	16.9	3.3	3.4	—	—	—	—	—
Aug	25.0	22.4	18.6	16.9	3.3	3.4	–2.8	—	—	—	—
Sep	25.0	22.4	18.6	16.9	3.3	3.4	–2.8	2.1	—	—	—
Oct	25.0	22.4	18.6	16.9	3.3	3.4	–2.8	2.1	–13.5	—	—
Nov	25.0	22.4	18.6	16.9	3.3	3.4	–2.8	2.1	–13.5	–17.1	—
Dec	25.0	22.4	18.6	16.9	3.3	3.4	–2.8	2.1	–13.5	–17.1	–31.3
Benchmark Series	19.8	14.5	11.5	9.5	6.5	2.0	–3.4	–7.6	–11.0	–11.9	–11.1

— nil or rounded to zero (including null cells)

6.1.2 Revision policy and practices continued

Another possible revision policy is to publish only the first revision of each initial trend estimate, and no more, until the next annual seasonal reanalysis takes place. Thus, for the period of a year, the only change to a monthly trend series would be the addition each month of the current trend estimate, and the revision of last month's trend estimate. Referring again to the simulation for the Unemployed Persons series, this would involve publishing only the estimates appearing on the first-off diagonal on table 6.1. Table 6.4 shows the trend series (as of the annual seasonal reanalysis based on data to January 1983) that would be published each month. This policy is similar to the previous policy, in that for a period of about a year, it still involves reliance upon the poorest approximating surrogate filters at the end of the series, where interest in reliable estimates of trend behaviour is keenest.

TABLE 6.4 UNEMPLOYED PERSONS SIMULATION, Trend series with only one subsequent revision to monthly trend estimates ('000)

Month in which trend series of movements calculated	*Jan* 1983	*Feb* 1983	*Mar* 1983	*Apr* 1983	*May* 1983	*Jun* 1983	*Jul* 1983	*Aug* 1983	*Sep* 1983	*Oct* 1983	*Nov* 1983	*Dec* 1983
Jan	642.8	—	—	—	—	—	—	—	—	—	—	—
Feb	648.2	667.8	—	—	—	—	—	—	—	—	—	—
Mar	648.2	673.1	690.2	—	—	—	—	—	—	—	—	—
Apr	648.2	673.1	693.8	708.8	—	—	—	—	—	—	—	—
May	648.2	673.1	693.8	711.4	725.7	—	—	—	—	—	—	—
Jun	648.2	673.1	693.8	711.4	720.2	729.0	—	—	—	—	—	—
Jul	648.2	673.1	693.8	711.4	720.2	726.9	732.4	—	—	—	—	—
Aug	648.2	673.1	693.8	711.4	720.2	726.9	728.1	729.5	—	—	—	—
Sep	648.2	673.1	693.8	711.4	720.2	726.9	728.1	730.0	731.6	—	—	—
Oct	648.2	673.1	693.8	711.4	720.2	726.9	728.1	730.0	721.4	718.1	—	—
Nov	648.2	673.1	693.8	711.4	720.2	726.9	728.1	730.0	721.4	708.7	701.0	—
Dec	648.2	673.1	693.8	711.4	720.2	726.9	728.1	730.0	721.4	708.7	684.1	669.6

— nil or rounded to zero (including null cells)

6.1.2 Revision policy and practices continued

Table 6.5 gives the month to month movements of the trend estimates published under this policy, as they would be calculated each month. Note that each of these estimates of movement is also revised only once as the revision of the initial trend estimate for a given month also gives rise to revision in the movement for that month. As can be seen from table 6.5, these movements would lead users to perceive the timing of the turning point as being between August and September, rather than July and August. Once again, this impression would not have been corrected until revised trend estimates became available around February 1984. Thus this policy can also substantially delay the detection of turning points in the trend.

TABLE 6.5 UNEMPLOYED PERSONS SIMULATION, Trend series with all subsequent revisions to monthly trend estimates (per cent)

Month in which trend series of movements calculated	*Feb* 1983	*Mar* 1983	*Apr* 1983	*May* 1983	*Jun* 1983	*Jul* 1983	*Aug* 1983	*Sep* 1983	*Oct* 1983	*Nov* 1983	*Dec* 1983
Feb	19.6	—	—	—	—	—	—	—	—	—	—
Mar	24.9	17.1	—	—	—	—	—	—	—	—	—
Apr	24.9	20.7	15.0	—	—	—	—	—	—	—	—
May	24.9	20.7	17.6	14.3	—	—	—	—	—	—	—
Jun	24.9	20.7	17.6	8.8	8.8	—	—	—	—	—	—
Jul	24.9	20.7	17.6	8.8	6.7	5.4	—	—	—	—	—
Aug	24.9	20.7	17.6	8.8	6.7	1.2	1.4	—	—	—	—
Sep	24.9	20.7	17.6	8.8	6.7	1.2	1.9	1.6	—	—	—
Oct	24.9	20.7	17.6	8.8	6.7	1.2	1.9	–8.6	–3.3	—	—
Nov	24.9	20.7	17.6	8.8	6.7	1.2	1.9	–8.6	–12.8	–7.7	—
Dec	24.9	20.7	17.6	8.8	6.7	1.2	1.9	–8.6	–12.8	–24.5	–14.5
Benchmark Series	19.8	14.5	11.5	9.5	6.5	2.0	–3.4	–7.6	–11.0	–11.9	–11.1

— nil or rounded to zero (including null cells)

6.1.2 Revision policy and practices continued

A further revision policy is to publish all revisions to the trend estimates each month as they occur. Thus each month, the current trend estimate is calculated, and the previously calculated surrogate estimates are all revised by incorporating the additional month of seasonally adjusted data. The trend series, as of January 1983, that would be published each month for the Unemployed Persons series, is exactly the series given in table 6.1. Table 6.6 gives the month to month movements of these trend estimates as they would be calculated each month. Note that as the surrogate trend estimates are all revised each month, so are all the movements that rely on surrogate trend estimates. It can be seen that if this policy is followed, then the turning point will be correctly identified as having occurred in July/August of 1983 when the October 1983 figure is first published, and that subsequent revisions to the surrogate trend estimates will not alter this timing. Thus this policy does not require users to wait until the next seasonal reanalysis before revised trend estimates can be viewed, and consequently represents a shorter delay in the detection of trend turning points.

TABLE 6.6 UNEMPLOYED PERSONS SIMULATION, Trend series with all subsequent revisions to monthly trend estimates (per cent)

Month in which trend series of movements calculated	*Feb* 1983	*Mar* 1983	*Apr* 1983	*May* 1983	*Jun* 1983	*Jul* 1983	*Aug* 1983	*Sep* 1983	*Oct* 1983	*Nov* 1983	*Dec* 1983
Feb	19.6	—	—	—	—	—	—	—	—	—	—
Mar	22.5	17.1	—	—	—	—	—	—	—	—	—
Apr	22.5	20.2	15.0	—	—	—	—	—	—	—	—
May	23.0	19.9	17.8	14.3	—	—	—	—	—	—	—
Jun	21.4	18.6	15.3	12.5	8.8	—	—	—	—	—	—
Jul	22.0	17.3	14.8	11.2	8.5	5.4	—	—	—	—	—
Aug	23.0	17.8	13.5	10.1	6.3	4.3	1.4	—	—	—	—
Sep	23.0	18.4	13.7	9.2	5.9	3.6	3.0	1.6	—	—	—
Oct	23.0	18.4	15.0	10.0	4.4	0.8	–1.2	–2.3	–3.3	—	—
Nov	23.0	18.4	15.0	11.4	4.3	0.5	–4.1	–5.9	–7.1	–7.7	—
Dec	23.0	18.4	15.0	11.4	7.1	0.6	–5.6	–10.5	–13.4	–15.1	–14.5
Benchmark Series	19.8	14.5	11.5	9.5	6.5	2.0	–3.4	–7.6	–11.0	–11.9	–11.1

— nil or rounded to zero (including null cells)

To minimise potential delays in the accurate detection of turning points, or other important trend changes, the ABS publishes all revisions to its trend estimates as they occur. This practice is consistent with the revision policies generally applied to both the original and seasonally adjusted series. Experience indicates that while the last three estimates of a monthly trend series can be misleading, the remaining estimates are a sufficient guide to the behaviour of the trend. This is because the surrogate filters used to calculate the last three trend estimates are the surrogates that most poorly approximate the 13-term Henderson filter, whereas the remaining surrogate filters generally approximate the 13-term Henderson filter quite well.

When using the series of trend movements to detect turning points, all but the last two estimates of movement are a sufficient guide to turning point timing. Again, this is because the last two movements are based on the trend estimates that most poorly approximate the Henderson based estimates for which they are designed to be

6.1.2 Revision policy and practices continued

substitutes. For example, when October 1983 data becomes available for the Unemployed Persons series, the movements for October and September will rely entirely on trend estimates for August, September, and October, which are expected to be poor approximations of the Henderson based estimates for those months. On the other hand, the movement for August is calculated on the basis of trend estimates for July and August which are calculated using the third and fourth conventional surrogate filters respectively. As the third conventional surrogate filter is a good approximation of the symmetric 13-term Henderson filter, and the fourth conventional surrogate filter is generally a fair approximation, this movement can be taken as a reasonable indication of the movement in the final Henderson based estimates for that month. Similarly, the monthly movements for July and earlier will be calculated on the basis of trend estimates which are either Henderson based estimates, or very good approximations of Henderson based estimates, and therefore they will also be reliable estimates of movement.

6.1.3 Guidelines for the extent of revisions

The majority of monthly series are smoothed using the 13-term Henderson filter and the set of conventional surrogate filters, and as indicated in section 5.3, it is the last three of these conventional surrogate filters that most poorly approximate the 13-term Henderson filter with regard to cycle damping and phase shift (refer to graphs 5.4, 5.5 and 5.6). Therefore, the trend estimate for a particular month should experience the most significant amount of revision in the first three months, with subsequent revisions becoming progressively smaller.

The more a particular series is dominated by cycles in the high frequency domain, where the damping and phase shifting properties of the conventional surrogate filters diverge most from the properties of the 13-term Henderson filter, the greater the amount of revision that will occur. Thus for series that are relatively less erratic, the extent of the revisions is expected to be smaller. Whether any of these revisions are 'significant' however, depends on the particular circumstances in which these indicators are being used.

As a result of examining many series, a rule of thumb has been developed for evaluating the expected amount of revision to the provisional trend estimates calculated for monthly series using the conventional surrogate filters. The rule states:

'Generally, the expected average percentage revision of the last, second last and third last trend estimates, relative to their 'final' estimates, will be respectively about one half, one quarter and one tenth of the series' star value. Other trend estimates are subject to minimal revision.'

The 'star value' (defined as the average absolute percentage changes of the irregular component of a time series) is a particular measure of the irregularity present in a series and is calculated as follows. First, a series of values giving the multiplicative contribution of the irregular component to the time series is determined, by dividing each seasonally adjusted value by the corresponding trend estimate. The star value is then taken to be the average percentage period to period change, without regard to sign, of the most recent ten year span of this series of irregular values. Table 6.7 gives the star value for a selection of monthly and quarterly series analysed by the ABS.

6.1.3 Guidelines for the extent of revisions continued

TABLE 6.7 STAR VALUE FOR SELECTED INDICATORS

Series	STAR
Monthly indicator	
Unemployment rate	1.72
Unemployed persons	1.86
Employed persons	0.26
Labour force	0.29
Participation rate	0.29
New motor vehicle registrations	
— Total vehicles	7.08
— Cars and wagons	7.24
— Other vehicles	7.25
Residential building approvals	8.16
Retail trade	0.69
Housing finance	3.6
Industrial Disputes	69.72
Quarterly indicator series	
Building commencements	3.1
Residential completions	1.89
Goods cebits	2.34
Goods credits	2.8
Current account deficit	11.03
GOS — Trading enterprises	2.06
GOS — Mfg companies	1.72
GDP	0.55
CPI	0.24

As an example of the rule in practice, consider the monthly series of new registrations of cars and station wagons. This series has a star value of 7.24, and therefore the rule of thumb indicates that the average percentage error of the trend is generally expected to be about:

- 1/2 x 7.24 = 3.62% for the current trend estimate
- 1/4 x 7.24 = 1.81% for the second most recent trend estimate
- 1/10 x 7.24 = 0.72% for the third most recent trend estimate.

Thus if data is available up to month N, the estimate of trend calculated for month N using the most non-symmetric conventional surrogate filter, would be expected to differ by around 3.62% from the Henderson based estimate eventually calculated for month N. Similarly, the estimate calculated for month N–1 at month N using the second most non-symmetric conventional surrogate filter, would be expected to differ by around 1.81% from the Henderson based estimate eventually calculated for month N–1. Finally, the estimate calculated for month N–2 at month N using the third most non-symmetric conventional surrogate filter, would be expected to differ by around 0.72% from the Henderson based estimate eventually calculated for month N–2.

The above rule gives the overall amount of revision expected between each provisional estimate and the final Henderson based estimate, and does not address the amount of revision that may be expected to occur from month to month. Thus, each surrogate trend estimate may be revised up and/or down from month to month, with additional revisions partially cancelling out earlier ones, before they converge to the Henderson

6.1.3 Guidelines for the extent of revisions continued

based estimate. Furthermore, this rule only indicates the average expected revision, and therefore individual revisions may be above or below these levels. For instance, the star value for Unemployed persons is 1.86, indicating that the expected average amount of revision for the provisional trend estimates for a given month, to the final Henderson based trend estimate for that month, are 0.93% for the first estimate of trend, 0.46% for the second estimate of trend, and 0.19% for the third estimate of trend. However, in the simulated trend series for 1983, the actual revisions for January were 0.67%, 0.23% and 0.32% respectively.

The rule indicates that very irregular series can experience a large amount of revision to the last three trend estimates. Such large revisions can give rise to large swings in apparent trend behaviour at the end of the series and have the potential to cause confusion to those using the data. For this reason, the ABS implemented a practice called clipping for many years, whereby the trend estimates for the last three months were not published for particularly irregular monthly series. While this led to fewer problems in dealing with provisional trend estimates, it did result in the loss of some useful information about the likely behaviour of the trend, and consequently, the practice of clipping ceased around 1993. At present, all provisional trend estimates are published for monthly series, although the possibility of significant revisions to the last three estimates, particularly for highly irregular series, means the most current trend estimates should be de-emphasised in the decision making process. Methods by which the reliability of the provisional trend estimates can be investigated will be discussed in section 6.2.

It should be noted that the above rule of thumb relates specifically to monthly series for which the provisional trend estimates are calculated using the conventional surrogate filters. It will be recalled from Chapter 5 that these surrogates to the 13-term Henderson filter are determined by taking the end weight parameter to be 3.5. For a number of Retail Trade series that show significantly less irregularity than is implied by an end weight parameter of 3.5, surrogate filters are used based on an end weight parameter of 2.0. The provisional trend estimates calculated from these surrogate filters are also subject to revision, with the last three trend estimates again expected to experience the most significant amount of revision each month. As these surrogate filters are more specifically tailored to the degree of irregularity in the series, the amount of revision is generally expected to be smaller than for the conventional surrogate filters, although there may be greater revision to the estimates whenever there are large irregular movements in the seasonally adjusted series.

A rule of thumb has also been developed for the average amount of revision expected to the three provisional trend estimates associated with quarterly series. This rule of thumb is based on a set of surrogate filters defined by a single typical end weight parameter, and as the surrogate filters for quarterly series are tailored to the degree of irregularity of the series being smoothed, the amount of revision would in fact generally be expected to be better than indicated. The rule states:

'Generally, the expected average percentage revision of the last, second last and third last trend estimates, relative to their 'final' estimates, will be respectively about one half, one fifth and one fifth of the series' star value.'

6.1.4 Revisions to the seasonally adjusted series

To those regularly using ABS data, it may appear initially that seasonally adjusted series are subject to less revision than trend series, and are therefore more suitable for analysis. This is not the case however, because an end-point problem similar to that for the trend affects the estimation of seasonal factors, which in turn influences the stability of the seasonally adjusted series as additional data is brought into the adjustment process. Instead of noticeably affecting just the last few periods (months/quarters) as is the case for revisions to the trend, it is the estimates of the last few years that are subject to revision in the seasonally adjusted series. Although they are more widespread, the revisions to ABS seasonally adjusted series are often not as obvious to users as revisions to the trend series because they are generally made annually rather than every period.

The ABS has been releasing seasonally adjusted series since 1967. Its experience in doing so has lead to an empirically based guideline which relates revisions of both monthly and quarterly seasonally adjusted series to their degrees of irregularity. The rule of thumb states:

'Generally, the average percentage revision (without regard to sign) for the most recent year, will be approximately one half of the series' star value. This degree of revision can be expected to fall gradually to about one fifth of the star value for years four or more earlier.'

It should be noted that the above rule of thumb examines the gradual revisions that generally occur from one year to the next as additional data is incorporated into the seasonal analysis process. Later revisions can, and sometimes do, partially cancel out earlier ones. Thus the rule does not address the question of how much the first seasonally adjusted estimates differ from the stable or 'final' estimates that are obtained some years later. Furthermore, the above rule indicates the average revision expected, and some revisions for particular months may once again be considerably larger or smaller than suggested. For instance, the star value for industrial disputes is about 70, indicating that the expected average revision of the most recent year of seasonally adjusted data is about 35%, and that these revisions should gradually decline to about 14% for years four or more earlier. One month, however, had a revision of 45%, while another had the relatively small revision of 2%.

Given that the trend is obtained directly from the seasonally adjusted series, it might appear that 'final' trend estimates are not possible until the seasonally adjusted figures themselves have ceased being revised. As discussed, it will take some years for the seasonally adjusted figures to become final, however, it turns out that the trending procedure quickly dampens the revisions generally experienced by the seasonally adjusted series, thereby leading to an effectively 'final' trend series being obtained very much sooner than the seasonally adjusted series.

It should be noted that the rules discussed above are only guidelines and as with all global and summary indicators, there will be occasional exceptions to what they indicate. With respect to revisions, there can also be other factors to take into account besides incorporation of additional data. For instance, there is the nature of revisions to the original data — whether they are seasonal in character, trend related, or more irregular in behaviour. Revisions to the seasonally adjusted series may also occur because allowances for trading day, payday factors and systematic moving holiday effects may be improved with hindsight and additional subject matter information. On occasions there

6.1.4 Revisions to the seasonally adjusted series continued

will also be abrupt dislocations to the trend or seasonal patterns, that give rise to an additional degree of revision as estimates adapt or are made to reflect the new behaviour. To date, the ABS has not developed rules of thumb in relation to these revision factors.

6.2 EVALUATING PROVISIONAL TREND ESTIMATES

The trending procedure gives rise to revision in the last three trend estimates for quarterly series and the last six trend estimates for monthly series (although as discussed above, noticeable degrees of revision are only experienced for the last three of these), with the revision expected to be the most significant for the most recent time period. As a result of these revisions, the trend estimates for the last three time periods, in particular, have the potential to give a misleading impression of recent trend behaviour. This was evident in the case of the simulation for the Unemployed Persons series discussed above, where the provisional trend estimates for March, April and May 1997 indicated that the Unemployed persons trend estimate was continuing to increase throughout these months. It was not until data for June 1997 had become available, and the March, April and May 1997 estimates had been further revised, that the March/April (later revised to April/May) turning point was revealed.

As a result of the potential impact of revisions to the last three trend estimates, for both quarterly and monthly series, these estimates should not be taken as conclusive indicators of recent trend behaviour on their own. This is especially so for seasonally adjusted series that are highly irregular, as the amount of revision can be quite significant for these series. In particular, this means that emphatic statements should not be made about the presence of a current trend turning point developing, based on the behaviour of the current trend estimates, without considering how reliable these provisional estimates are likely to be. One way to evaluate the reliability of the provisional trend estimates is to consider the reliability of the corresponding estimate of the irregular component of series behaviour. This approach will be discussed in section 6.2.1. Another approach is to determine how sensitive the current trend estimates are to incorporation of additional data. Such analysis is referred to here as 'sensitivity analysis', and will be the subject of discussion in sections 6.2.2 and 6.2.3. Section 6.2.4 briefly discusses some alternative methods that are being investigated by the ABS for evaluating the reliability of provisional trend estimates.

6.2.1 Assessment against implied irregularity

To assist analysis of recent and current trend behaviour, the ABS generally graphs and tabulates the trend estimates with the seasonally adjusted series from which they were derived. The reason for this practice is to enable the users to determine the relative magnitude and direction of the differences between the two series, the difference being an estimate (in the additive sense) of the irregular component of series behaviour. The user is then able to contrast these estimates of recent irregularity with their own knowledge or anecdotal information of the erratic influences in the series. If the user believes the irregular contribution to the series in a particular period is significantly larger or smaller than the value determined from the series provided by the ABS, they may amend the provisional estimates of the trend series accordingly. For example, suppose the user believes that the irregular for the current period is larger than the value calculated from the published seasonally adjusted and trend series. The user could then reduce the value of the provisional trend estimate for the current period, thereby increasing the corresponding irregular to a level considered to be more reasonable.

6.2.1 Assessment against implied irregularity continued

It should be noted that while the original series is also available, contrasting the trend series with the original series is not particularly helpful. As discussed earlier, the difference between the trend and the original series is an estimate of the combined effect of evolving seasonality, trading day patterns, moving holiday influences, and irregularity. It is unlikely that many users of the data will have detailed knowledge of the magnitude and directional role of all these factors at each point in time.

6.2.2 Sensitivity analysis I

Another means by which the recent provisional trend estimates can be evaluated, is to determine what the seasonally adjusted figure would have to be next period in order for the current trend estimate to be revised by a specified amount (for instance, one might consider what seasonally adjusted figure is required next period, all other factors constant, for the current trend estimate not to be revised). This scenario may be helpful because most series become available for analysis five months or so after the period to which they refer. Consequently, there may be other indicators and anecdotal information available to judge whether the seasonally adjusted figure for the period in question is likely to be above or below the calculated value. Users may wish to adjust the provisional trend estimates themselves using their own additional information about the seasonally adjusted estimate next period.

In order to calculate the seasonally adjusted value required to give rise to a particular degree of revision in the current trend estimate, it is necessary to refer back to the weighting process. Each provisional trend estimate is determined, at the given point in time, by using the weights of the appropriate filter to weight and add together a number of seasonally adjusted values. That is, each trend estimate is determined by substituting various seasonally adjusted values into a particular weighting formula. For instance, suppose seasonally adjusted data for a monthly series is available up to month N and that the trend estimates at the end of the series are determined using the six conventional surrogate filters to the 13-term Henderson filter. The estimate of trend for month N, $T_{N,N}$ is then calculated using the formula

$$\begin{aligned} T_{N,N} = & -0.092A_{N-6} - 0.058A_{N-5} + 0.012A_{N-4} + 0.12A_{N-3} \\ & + 0.244A_{N-2} + 0.353A_{N-1} + 0.421A_N \end{aligned} \quad (6.1)$$

where A_i denotes the seasonally adjusted value at month i, and where –0.092, –0.058, 0.012, 0.120, 0.244, 0.353 and 0.421 are the weights of the sixth conventional surrogate filter, as given in table 5.2. When seasonally adjusted data becomes available for month N+1, the trend estimate for month N will be recalculated by applying the fifth conventional surrogate filter to the updated seasonally adjusted series. Therefore, the trend estimate for month N, calculated at month N+1, is determined from the formula

$$\begin{aligned} T_{N,N+1} = & -0.043A_{N-6} - 0.038A_{N-5} + 0.002A_{N-4} + 0.080A_{N-3} \\ & + 0.174A_{N-2} + 0.254A_{N-1} + 0.292A_N + 0.279A_{N+1} \end{aligned} \quad (6.2)$$

where –0.043, –0.038, 0.002, 0.080, 0.174, 0.254, 0.292, and 0.279 are the weights of the fifth conventional surrogate filter, as also given in table 5.2. Subtracting equation (6.1) from equation (6.2) then gives the amount of revision that occurs in the provisional trend estimate for month N, upon the addition of the seasonally adjusted value for month N+1. This equation is

6.2.2 Sensitivity analysis I continued

$$Revision = 0.049A_{N-6} + 0.020A_{N-5} - 0.010A_{N-4} - 0.040A_{N-3} - 0.070A_{N-2} - 0.099A_{N-1} - 0.129A_N + 0.279A_{N+1} \quad (6.3)$$

Rearranging this equation gives equation (6.4),

$$A_{N+1} = \frac{Revision - 0.049A_{N-6} - 0.020A_{N-5} + 0.010A_{N-4} + 0.040A_{N-2} + 0.070A_{N-2} + 0.099A_{N-1} + 0.129A_N}{0.279} \quad (6.4)$$

which provides the mechanism for determining the seasonally adjusted value required next month, in order for the current trend estimate to be revised by a particular amount. For example, consider the Unemployed Persons series for 1997. Table 6.8 gives the seasonally adjusted values (SA) and trend values (T) that were published in December 1997.

TABLE 6.8 UNEMPLOYED PERSONS ('000), 1997

Month	*Jan*	*Feb*	*Mar*	*Apr*	*May*	*Jun*	*Jul*	*Aug*	*Sep*	*Oct*	*Nov*	*Dec*
SA	790.9	810.8	803.0	802.3	809.8	777.8	801.2	798.7	791.5	772.7	778.9	753.2
T	797.0	799.6	801.7	802.4	801.1	798.7	795.6	791.5	786.1	779.6	772.7	767.1

Using equation (6.4), the January 1998 seasonally adjusted value required in order for the December trend value to remain unchanged (that is, for zero revision to occur in the December trend estimate) is given by

$$\begin{aligned} A_{N+1} &= [0 - 0.049 \cdot (777.8) - 0.020 \cdot (801.2) + 0.010 \cdot (798.7) + 0.040 \cdot (791.5) \\ &\quad + 0.070 \cdot (772.7) + 0.099 \cdot (778.9) + 0.129 \cdot (753.2)]/0.279 \\ &= 766.6 \end{aligned}$$

Thus, in order for the December trend value to stay at 767.1 when the surrogate trend estimates are revised in January 1998, the seasonally adjusted value for January would have to be 766.6. Similarly, if the trend was to be revised up by 10 over the initial December figure of 767.1, the January seasonally adjusted value required would have to be

$$\begin{aligned} A_{N+1} &= [10 - 0.049 \cdot (777.8) - 0.020 \cdot (801.2) + 0.010 \cdot (798.7) + 0.040 \cdot (791.5) \\ &\quad + 0.070 \cdot (772.7) + 0.099 \cdot (778.9) + 0.129 \cdot (753.2)]/0.279 \\ &= 802.4 \end{aligned}$$

If this value seemed reasonable on the basis of all other available information, then the January seasonally adjusted figure of 802.4 could be used to revise the other provisional trend estimates calculated in December, particularly the October and November values. The analyst would then be able to determine whether such a significant revision to the December estimate, having been regarded as plausible, was indicative of a turning point having occurred some time in the last few months.

6.2.2 Sensitivity analysis I continued

By referring back to the weighting process and the equations used to derive the provisional trend estimates at various points in time, it is possible for the analyst to derive formulae for considering other types of scenarios that may be of interest. For instance, formulae can be derived to determine the seasonally adjusted figure required next month in order to produce a particular percentage growth for that month, or the present month. These scenarios may be useful because they allow analysts to consider the conditions required for the current growth rate to be maintained, or for there to be a quickening or slowing of growth. In addition, consideration of zero growth can provide information about the possibility of a turning point, point of inflection, or plateau developing in the trend. Regardless of the scenario considered, the analyst is left to judge whether the implied value for next month's seasonally adjusted figure is likely to be achieved, given present economic conditions and the history of the seasonally adjusted series.

It should be noted that calculations of the type discussed above only consider revisions that are brought about by the trending procedure, and do not take into account the effect of revisions to the original or seasonally adjusted series. A further point to note is that while such calculations can also be made for quarterly series, it would be necessary to determine the weighting formulae for the surrogate filters used for the series in question (bearing in mind that the weights of the surrogate filters are tailored to the specific series on the basis of it's relative degree of irregular to trend variability), before appropriate formulas could be derived. Similarly, specific formulae would have to be derived for any monthly series, such as some of the Retail Trade series, for which the trend estimates at the end of the series were not determined by the conventional surrogate filters.

6.2.3 Sensitivity analysis II

Another type of sensitivity analysis involves considering the alternative trend paths the series would follow if the seasonally adjusted series increased or decreased by a typical amount in the next month. This is the type of sensitivity analysis that regularly appears in ABS publications. As an example, consider the analysis undertaken for the Non-residential Building Approvals series in graph 6.2. In this example, the average monthly percentage movement, without regard to sign, of the seasonally adjusted series was about 20 per cent for the ten years to October 1998. The alternative trend paths were therefore determined by revising the trend series on the assumption that the seasonally adjusted value for November rose (or fell) by 20 per cent over the seasonally adjusted value for October.

6.2.3 Sensitivity analysis II continued

GRAPH 6.2 RELIABILITY OF TREND ESTIMATES FOR NON-RESIDENTIAL BUILDING APPROVALS

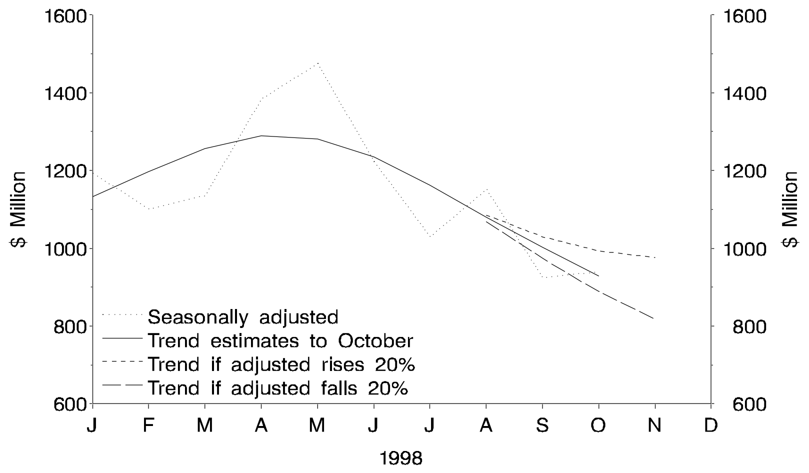

In this case, the figure of 20% was chosen because it represented the average monthly percentage movement in the seasonally adjusted series over the past ten years. However, the trend paths resulting from any particular level of growth could be considered. Furthermore, such paths could be determined by using the analysis already undertaken, without requiring the revised trend estimates at the end of the series to be recalculated under every scenario of interest. Suppose, for instance, that the trend series is available up to period N, and that revised trend paths have been determined for the cases where the seasonally adjusted value for period N+1 rises or falls by X per cent over the value for period N. Let $U_{N,X}$ denote the (upper) path traced when the seasonally adjusted value grows by X per cent, and let $L_{N,X}$ denote the (lower) path traced when the seasonally adjusted value falls by X per cent. Then the path, Z_N, traced by the mid-points of the paths $U_{N,X}$ and $L_{N,X}$, represents the trend outcome obtained if there is zero percentage growth in the seasonally adjusted series in period N+1 (see graph 6.3). As the path Z_N corresponds to zero growth and the path corresponds to X per cent growth, the path for any particular positive level of growth, say W per cent, can now be determined by tracing the path W/X from Z_N to $U_{N,X}$, as indicated in graph 6.3. That is, the trend path $U_{N,W}$ corresponding to W per cent growth is given by

$$U_{N,W} = Z_N + \frac{W}{X}(U_{N,X} - Z_N) \quad (6.5)$$

For instance, in the example above, where the upper trend path is determined by growth of 20 per cent in the seasonally adjusted value for the next period, the trend path corresponding to 1 per cent growth is given by

$$U_{N,1} = Z_N + \frac{1}{20}(U_{N,20} - Z_N) \quad (6.6)$$

Similarly, the trend path corresponding to 40% growth in the seasonally adjusted value for the next period is given by

$$U_{N,40} = Z_N + \frac{40}{20}(U_{N,20} - Z_N) \qquad (6.7)$$

6.2.3 Sensitivity analysis II continued

that is, the path followed by the trend in this case is given by the mid-point path plus twice the difference between this path and the upper alternative path. The trend path for any particular rate of decrease, say W per cent, can similarly be determined by considering the line W/X from Z_N to $L_{N,X}$. That is, the trend path $L_{N,W}$ corresponding to a fall of W per cent is given by

$$L_{N,W} = Z_N - \frac{W}{X}(Z_N - L_{N,X}) \quad (6.8)$$

From these examples it can be seen that the user can examine the trend path resulting from any particular rise or fall in the seasonally adjusted value next period, regardless of the initial growth scenarios published.

GRAPH 6.3 ALTERNATIVE TREND PATHS BASED ON EQUAL TYPICAL SEASONALLY ADJUSTED RISE AND FALL

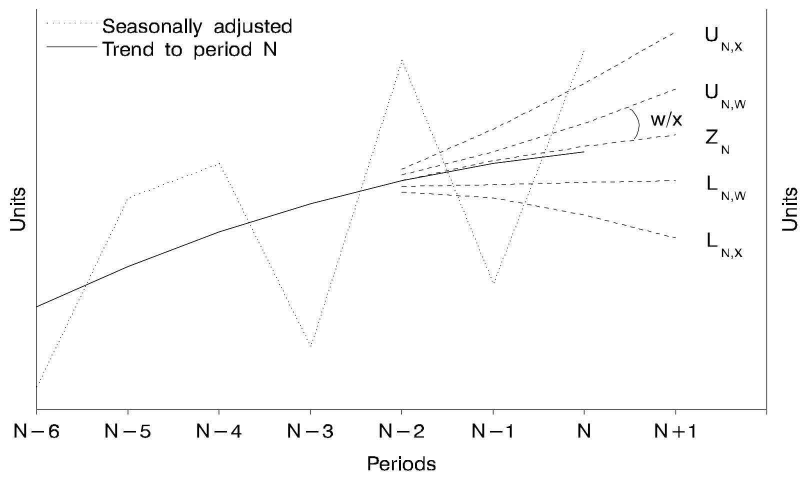

In this type of sensitivity analysis, a number of aspects should be noted. Firstly, the average percentage growth from one period to the next may not be the most appropriate 'typical' measure to consider. In some applications the median or modal percentage growth may be more relevant, and these measures of growth are easily accommodated for in this type of analysis.

Secondly, for some series, the typical seasonally adjusted increase may not be the same magnitude as the typical decrease. This means that the upper and lower trend paths initially determined are not spaced equally about the trend path corresponding to zero growth. For example, consider the situation illustrated in graph 6.4. In this case the upper alternative trend path is based upon a representative seasonally adjusted increase of 5%, and the lower trend path on a representative seasonally adjusted decline of 1%. The range of these seasonally adjusted movements is six units (+5 – (–1) = 6) and therefore, logic indicates that the path of zero growth will lie 5/6ths of the way from the upper path corresponding to a five per cent increase, to the lower path corresponding to a one per cent decrease. That is, if $U_{X,5}$ denotes the upper alternative trend path based on a rise of five per cent, and $L_{N,1}$ denotes the lower alternative trend path based on a fall of one per cent, then the trend path, Z_N, associated with zero growth, will be

$$Z_N = U_{N,5} - \frac{5}{6}(U_{N,5} - L_{N,1}) \quad (6.9)$$

6.2.3 Sensitivity analysis II continued

This rule can be extended to the general case, so that if $U_{N,X}$ is the upper alternative trend path based on a rise of X per cent, and $L_{N,Y}$ is the lower alternative trend path based on a fall of Y per cent, then the trend path, Z_N, corresponding to zero growth, will be

$$Z_N = U_{N,X} - \frac{X}{X+Y}(U_{N,X} - L_{N,Y}) \qquad (6.10)$$

Once the path corresponding to zero growth has been determined, the trend path associated with any other specific level of growth can be calculated as above.

GRAPH 6.4 ALTERNATIVE TREND PATHS BASED ON UNEQUAL TYPICAL RISE AND FALL

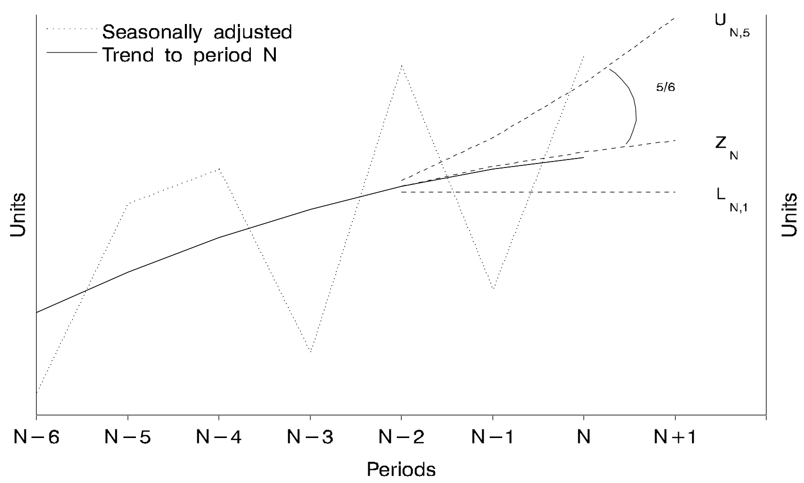

The third aspect of this type of sensitivity analysis that should be noted, is that the alternative trend paths are not necessarily symmetrically distributed about the current trend path, even when they are based upon future seasonally adjusted rises and falls of equal magnitude. In fact, it is possible for both of these paths to be either above or below the current trend estimates. Such situations will arise whenever the most recent seasonally adjusted movement is exceptionally large and the typical seasonally adjusted rise or fall next period is insufficient to counteract the effect that level of growth has had on the trend estimates.

The fourth aspect to note is that this type of sensitivity analysis is carried out on the basis that all other factors are held constant. That is, it assumes that both the original and seasonally adjusted series are not themselves revised, so that the only revisions occurring to the trend estimates are a direct result of the fact that an additional period of seasonally adjusted data is being incorporated into the smoothing process.

Finally, it should be noted that this analysis only looks one period into the future, and ignores the fact that the present trend estimates will be revised not only by the addition of next period's seasonally adjusted value, but by the addition of subsequent periods' as well. Consequently, considering the impact that next period's seasonally adjusted value has on topical trend estimates is only part of the overall revision story, albeit an important part.

6.2.4 Alternative methods of analysis

In the above sensitivity analyses, each of the scenarios considered involved conditional statements such as, 'This is what the trend series would be if the seasonally adjusted series was such and such next period, all other factors constant.' In such analysis, no likelihood of the particular scenario becoming an outcome is explicitly stated. That aspect is left to the analyst to determine, given that additional information or hypotheses may be available to the analyst about the state of the economy, or the activity the series represents and is related to. However, general probability statements concerning the trend's revisions may be useful to analysts, and the ABS is investigating ways in which this may be achieved. For example, Retail Turnover is published with a 90% confidence interval about the trend estimate. This confidence interval is based on the trend revision history since 1962.

CHAPTER 7 PRACTICAL CONSIDERATIONS

CONTENTS

page

SECTION

7.0 OVERVIEW

While it is important to consider the impact of revisions when using and interpreting trend estimates, there are a number of other practical considerations that need to be taken into account. For instance, it is important to consider how sudden changes in the behaviour of any of the original series components may affect the trend estimates, and whether corrections should be made in accordance with these effects. In addition, it is valuable to consider the compatibility of related trend estimates. For example, how do the trend estimates for an aggregated series relate to the trend estimates for the components? How do trend estimates for quarterly series relate to trend estimates for monthly series? This Chapter deals with such considerations. In particular, section 7.1 considers the impact of disturbances in the components of the original data, section 7.2 considers the compatibility of various trend series, and section 7.3 discusses some other general issues that should be taken into account when using Henderson filters to derive estimates of trend.

7.1 DISTURBANCES IN THE DATA

ABS trend estimates are derived directly from the seasonally adjusted series by applying an averaging process in the form of the Henderson moving averages and their surrogates. Therefore, any disturbances in the original series data which pass through to the seasonally adjusted series, also have the potential to pass through to the estimates of trend. Examples of such phenomena include the impact of extreme one-off irregular influences, abrupt rises or falls in the level of the trend path, and disruption of seasonal and trading day patterns. The general nature and impact of these problems is considered below.

7.1.1 Extreme irregulars

As previously indicated, the original series data is affected by real world socioeconomic irregular effects and imperfections in the statistical collection process, such as sampling and non-sampling error. These influences give rise to the irregular fluctuations which are present in the original and seasonally adjusted series, and which mask the underlying trend component of series behaviour. As the irregularities have a tendency to chop and change in magnitude and direction from period to period, they tend to cancel each other out on average. This is the motivation behind using an averaging process to smooth the irregularity from the seasonally adjusted series. If a particularly large irregular resides in the seasonally adjusted series, it may not be counteracted over the length of the averaging span, thereby contaminating the trend estimates for a number of periods. The effect of this contamination is indicated in graph 7.1, where the impact of the Reserve Bank of Australia's bulk sell-off of gold stocks in June 1997 is evident in the data for exports of other non-rural goods.

7.1.1 Extreme irregulars continued

GRAPH 7.1 IMPACT OF EXTREME IRREGULARS ON TREND ESTIMATES

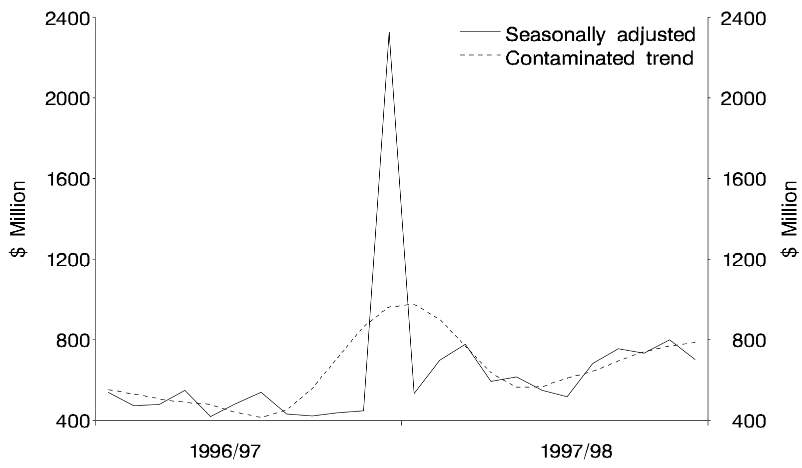

By considering the weighting patterns of the filters used to derive the trend estimates for a particular series, it is possible to determine what proportion of a single large irregular, or 'contaminant' flows through to various estimates of trend. For instance, consider a monthly series which is smoothed using the 13-term Henderson filter and its six conventional surrogate filters (see table 5.2 for the weighting patterns of these filters). If a contaminant resides in the seasonally adjusted value for month M, then a proportion of this contaminant will pass through to the trend estimates that use the seasonally adjusted value for month M in their moving average calculations. These proportions are given in table 7.1. Each row of this table indicates the effect of the contaminant on the estimates of trend calculated at a given point in time; thus the first row indicates the effect on the trend estimates calculated at month M, the second row indicates the effect on the estimates calculated at month M+1, and so on. Each column of table 7.1 then gives the percentage of contaminant contributed to a particular trend estimate as it is revised over time.

TABLE 7.1 MONTHLY CONTAMINATION, Percentage of a specific contaminant in month M passing to trend estimates

As at month	M–6	M–5	M–4	M–3	M–2	M–1	M	M+1	M+2	M+3	M+4	M+5	M+6
M	–1.9	–3.4	–1.8	4.6	14.8	27.9	42.1	—	—	—	—	—	—
M+1	–1.9	–2.8	–0.6	5.0	13.1	21.6	29.2	35.3	—	—	—	—	—
M+2	–1.9	–2.8	0.1	6.1	13.6	20.1	24.1	25.4	24.4	—	—	—	—
M+3	–1.9	–2.8	0.1	6.6	14.4	20.5	23.0	21.6	17.4	12.0	—	—	—
M+4	–1.9	–2.8	0.1	6.6	14.7	21.2	23.5	20.8	14.9	8.0	1.2	—	—
M+5	–1.9	–2.8	0.1	6.6	14.7	21.4	23.8	21.0	14.5	6.8	0.2	–5.8	—
M+6	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.3	14.5	6.6	0.3	–3.8	–9.2
M+7	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.7	0.4	–2.5	–4.3
M+8	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.3	–2.2	–1.6
M+9	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.1	–2.2	–0.9
M+10	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.1	–2.5	–1.1
M+11	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.1	–2.8	–1.7
M+12	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.1	–2.8	–1.9
etc	–1.9	–2.8	0.1	6.6	14.7	21.4	24.0	21.4	14.7	6.6	0.1	–2.8	–1.9

— nil or rounded to zero (including null cells)

7.1.1 Extreme irregulars continued

Table 7.1 indicates that when a contaminant occurs in the seasonally adjusted value for month M, then initially, 42.1 per cent of the contaminant passes through to the trend estimate for that month. Progressing down the column for month M, it can be seen that as additional data is brought into the smoothing process, and the trend estimate for month M is revised using longer surrogate filters, the contribution of the contaminant declines. Eventually, once enough data is available to determine the trend estimate for month M using the 13-term Henderson filter, the contribution of the contaminant stabilises, the contaminant remaining with only 24% of its strength.

As the last six trend estimates are revised each month by incorporating the new seasonally adjusted value into the smoothing procedure, estimates prior to month M will be affected retrospectively by the contaminant. For instance, when the seasonally adjusted value becomes available for month M, the revised trend estimate calculated for month M–1 will contain about 28% of the contaminant. In the following five months, as this estimate is further revised using more appropriate filters, the contribution of the contaminant will decline to about 21%, where it stabilises. Similarly, the trend estimate for month M–2 contains about 15% of the contaminant, both at month M and in later months, the estimate for month M–3 contains between five and seven per cent, and only minor amounts of the contaminant pass into the estimates for months M–4, M–5 and M–6 each month. The trend estimates for periods more than six months before month M are then unaffected by the contaminant as they are all determined by applying the 13-term Henderson filter to a 13-term span that does not include month M.

Trend estimates for months following month M will also be influenced by the contaminant. Again, this is because the seasonally adjusted value for month M, and therefore the contaminant in month M, is incorporated into the moving average calculations for future months. Table 7.1 shows that the trend estimate for month M+1 will initially contain about 35 of the contaminant, and this reduces to about 21% as the estimate is revised over the next six months. In the month after that (month M+2), about 24% of the contaminant will pass to the trend estimate initially, although

7.1.1 Extreme irregulars continued

eventually this is reduced to about 15%, and so on. The estimate for month M+6 is then the last trend estimate incorporating the seasonally adjusted value for month M in its calculation, and is therefore the last trend estimate affected by this contaminant (the contribution of the contaminant varies for the six months following month M+6 as the trend estimate is revised using the various surrogate filters, and eventually, the 13-term Henderson filter).

From the above discussion, it can be seen that if a contaminant occurs in a monthly series being smoothed by the 13-term Henderson filter, then not only will the trend estimate for the affected month respond to the presence of the contaminant, the trend estimates for the preceding six months and the following six months will also respond. Furthermore, the impact of the contaminant on a particular trend estimate will vary over time as the estimate is revised, the proportion of contaminant passing into the trend estimate depending on the filter used to calculate it at a given point in time.

The effect of a contaminant occurring at a particular point in time can also be considered for the trend estimates derived from a quarterly series. In this case the seasonally adjusted series is smoothed by the 7-term Henderson filter, in conjunction with a set of three surrogate filters which are tailored to the degree of irregular to trend variability of the series being smoothed. Table 7.2 gives some examples of the impact of a contaminant occurring in quarter Q, for a few values of the end weight parameter measure of irregular to trend variability. Case 1 in table 7.2 examines the situation in which the smoothing procedure is applied to a seasonally adjusted series in which the period to period movements are driven predominantly by trend behaviour rather than the irregular fluctuations (an end weight parameter of 0.4). Cases 2 and 3, on the other hand, consider circumstances in which the irregular fluctuations are generally more dominant in the series' period to period movements (end weight parameters of 2.0 and 3.5 respectively).

TABLE 7.2 QUARTERLY CONTAMINATION, Percentage of a specific contaminant in quarter Q passing to trend estimates

As at quarter	*Q–3*	*Q–2*	*Q–1*	*Q*	*Q+1*	*Q+2*	*Q+3*
Case 1: End weight=0.4 (very stable series)							
Q	–5.9	2.0	28.2	68.4	—	—	—
Q+1	–5.9	5.9	26.6	40.7	43.3	—	—
Q+2	–5.9	5.9	29.4	39.7	29.4	6.6	—
Q+3	–5.9	5.9	29.4	41.3	29.0	6.5	–18.3
Q+4	–5.9	5.9	29.4	41.3	29.4	6.6	–4.7
Q+5	–5.9	5.9	29.4	41.3	29.4	5.9	–3.9
Q+6	–5.9	5.9	29.4	41.3	29.4	5.9	–5.9
Case 2: End weight=2.0 (moderately variable series)							
Q	–5.9	2.4	28.5	61.0	—	—	—
Q+1	–5.9	5.9	26.9	40.8	40.9	—	—
Q+2	–5.9	5.9	29.4	39.8	29.4	9.1	—
Q+3	–5.9	5.9	29.4	41.3	28.9	6.3	–11.0
Q+4	–5.9	5.9	29.4	41.3	29.4	6.4	–5.0
Q+5	–5.9	5.9	29.4	41.3	29.4	5.9	–4.4
Q+6	–5.9	5.9	29.4	41.3	29.4	5.9	–5.9
Case 3: End weight=3.5 (quite volatile series)							
Q	–5.9	3.0	28.8	55.5	—	—	—
Q+1	–5.9	5.9	27.2	41.0	39.0	—	—
Q+2	–5.9	5.9	29.4	39.9	29.4	10.9	—
Q+3	–5.9	5.9	29.4	41.3	28.8	6.2	–5.5
Q+4	–5.9	5.9	29.4	41.3	29.4	6.0	–5.3
Q+5	–5.9	5.9	29.4	41.3	29.4	5.9	–5.0
Q+6	–5.9	5.9	29.4	41.3	29.4	5.9	–5.9

— nil or rounded to zero (including null cells)

7.1.1 Extreme irregulars continued

From table 7.2 it can be seen that the effect of the contaminant varies depending on the relative amount of irregular to trend variability of the series being smoothed. For example, consider the trend estimate for the quarter in which the contaminant occurs, quarter Q. The greater the amount of irregularity assumed by the trending procedure, the less weight is assigned to the contaminant in the trend estimate for this quarter. In Case 1, 68.4% of the contaminant initially passes through to the trend estimate for quarter Q, with the degree of contamination varying over the next three quarters as the estimate is revised, stabilising at 41.3%. On the other hand, in Case 3 it can be seen that 55.5% of the contaminant initially passes to the trend estimate for quarter Q, before stabilising at 41.3%.

Table 7.2 also illustrates how other trend estimates are affected both before and after the period in which the contaminant occurs, as subsequent observations become available. In particular, it can be seen that not only does the contaminant affect the trend estimate for the quarter in which it occurs, it also retrospectively affects the trend estimates for the preceding three quarters to varying degrees over time. In addition, the trend estimates for the three quarters following the occurrence of the contaminant are affected, the contribution for each trend estimate varying over four quarters as the estimate is revised.

7.1.1 Extreme irregulars continued

The above tables indicate the proportion of a particular contaminant that passes through to the various trend estimates, but they do not indicate the overall impact of this contaminant. The actual impact will in fact depend on how much of the trend estimate is determined by the proportion of contaminant passing into it. A measure of this impact is defined as follows. Suppose a seasonally adjusted series is smoothed using a Henderson filter of length 2m+1. Let $C_{N,t}$ be the size of the contaminant occurring at period N, as it has been assessed at time t (the initial size of a contaminant may be revised over time with the benefit of hindsight and additional subject matter information). Let $\Pi_{p,t}$ be the proportion of this contaminant that passes into the trend estimate calculated for period p at time t, as given in Tables 7.1 and 7.2. The amount of contaminant passing into the trend estimate is then $\Pi_{p,t} \times C_{N,t}$. Hence, if $T_{p,t}$ is used to denote the trend estimate calculated for period p at time t, then a measure, $Y_{p,t}$, of the relative contamination of this trend estimate is given by

$$Y_{p,t} = \frac{\Pi_{p,t} \times C_{N,t}}{T_{p,t}} \times 100 \qquad (7.1)$$

for $p = N-m, \ldots, N, \ldots, N+m$, and $t = N, N+1, \ldots, N+2m$.

To understand the operation of this formula, consider the following example in which there is a contaminant in the December 1992 seasonally adjusted figure of a monthly series. Suppose the contaminant is thought to have accounted for 52 units of the seasonally adjusted December 1992 figure, and that the December trend estimate is 5,097. The relative contamination of this trend estimate is then given by

$$Y_{M,M} = \frac{0.421 \times 52}{5097} \times 100 = 0.43\%$$

$$M = Dec\ \ 1992$$

The impact of this contaminant on the December trend estimate will change over time as the estimate is revised. Suppose the trend estimate for December has been revised to 5,129 in June 1993, both as a result of incorporating the additional seasonally adjusted values into the smoothing procedure, and as a result of any revisions to the seasonally adjusted values from which the December trend estimate is derived. Assuming the contaminant is still thought to contribute 52 units to the seasonally adjusted December figure, the relative contamination of the December trend estimate calculated in June 1993 will be

$$Y_{M,M+6} = \frac{0.240 \times 52}{5129} \times 100 = 0.24\%$$

$$M = Dec\ \ 1992$$

Table 7.1 indicates that the proportion of the contaminant passing from the December seasonally adjusted figure to the December trend estimate, stays at 24% from June 1993 onwards (M+6, etc.). Furthermore, as the December trend estimate is calculated using the 13-term Henderson filter from June 1993, the trend estimate will only be revised in later months if the seasonally adjusted values are revised. Consequently, the level of relative contamination of the December trend estimate will only change in later months if the size of the contaminant is revised, or revisions in the seasonally adjusted series give rise to revisions in the December trend estimate.

7.1.1 Extreme irregulars continued

As discussed above, months prior to December 1992 will also be contaminated retrospectively. Consider the month of September 1992 for instance. Once seasonally adjusted data for December has become available, the revised estimate calculated for September from the updated seasonally adjusted series will contain the effect of the December contaminant. Suppose the revised trend estimate for September as at December is 4,908. The relative contamination of the September trend estimate at this time is then

$$Y_{M-3,M} = \frac{0.046 \times 52}{4908} \times 100 = 0.05\%$$
$$M = Dec \;\; 1992$$

The September 1992 trend estimate will be contaminated in following months as well, although the level of contamination will vary as the trend estimate is revised. For instance, consider the impact of the contaminant on the September 1992 trend estimate calculated in March 1993. If the trend estimate for September has been revised to 5,006, and the contaminant hasn't varied, then the relative contamination will be

$$Y_{M-3,M+3} = \frac{0.066 \times 52}{5006} \times 100 = 0.07\%$$
$$M = Dec \;\; 1992$$

Once again, the level of relative contamination of this trend estimate will only change in later months if the size of the contaminant is revised, or the trend estimate for September is revised as a result of revisions to the seasonally adjusted values from which it is derived.

As previously mentioned, months after December 1992 will also be affected by the contaminant in the data. For instance, consider the June 1993 trend estimate, which will be the last estimate affected by the contaminant in the seasonally adjusted figure for December 1992. If the initial trend estimate for June 1993 is 4,975, and the contaminant is still thought to account for 52 units of the seasonally adjusted December figure, then the relative contamination of the trend estimate for June 1993 will initially be

$$Y_{M+6,M+6} = \frac{-0.092 \times 52}{4975} \times 100 = -0.10\%$$
$$M = Dec \;\; 1992$$

This negative sign just indicates that the contaminant has been responsible for reducing the size of the trend estimate rather than increasing it (as in the previous cases) and thus the contaminant has had the effect of reducing the June trend estimate by 0.10%. As indicated earlier, the contribution of the contaminant then changes over the next six months as the June estimate is revised. In December 1993, for example, when the June 1993 trend estimate has been revised to 5,318 units, the relative contamination of the June 1993 trend estimate will be

$$Y_{M+6,M+12} = \frac{-0.019 \times 52}{5318} \times 100 = -0.02\%$$
$$M = Dec \;\; 1992$$

7.1.1 Extreme irregulars continued

This level of relative contamination will then represent the impact of the contaminant on the June estimate in future months, unless the size of the contaminant is revised or the seasonally adjusted series undergoes revision.

While this example illustrates how the level of relative contamination can be determined for monthly trend estimates, equation (7.1) can similarly be used to determine the level of relative contamination for quarterly trend estimates. In the latter case, the only difference is that the proportions, $\Pi_{p,t}$, depend on the relative degree of irregular to trend variability of the series being smoothed.

The above discussion indicates that a particularly large irregular can contaminate the trend estimates for a number of periods. One way to prevent the trend series from responding to the presence of these extreme irregulars, would be to remove them from the seasonally adjusted series before the series is smoothed. The main problem with this strategy, however, is the difficulty of identifying what constitutes an extreme irregular. Furthermore, the implementation of such a strategy would mean that users could not reproduce the trend estimates themselves by applying the Henderson and surrogate filters to the published seasonally adjusted series (the published seasonally adjusted series does not have any irregularities removed, extreme or otherwise, as it is designed to be a measure of the non-seasonal behaviour of the original series, that is, it is designed to contain the trend and irregular components of original series behaviour).

The problems mentioned above were among those considered when the ABS first began publishing trend estimates in 1984. The ABS has recently developed an algorithm which uses a cut-off based on a statistical measure of variability, to identify time series observations which are unusually high or low. When available, subject matter knowledge is then used to determine whether each observation thus identified has arisen as a result of an identifiable random, or one-off influence. In such cases, the impact of the irregular influence is then estimated and removed from the extreme seasonally adjusted value before smoothing is undertaken. This procedure is generally applied to all seasonally adjusted series smoothed by the ABS and produces trend estimates which are influenced less by irregularity and which are therefore more representative of the underlying behaviour of the series.

While extreme irregulars that have been quantified are removed from seasonally adjusted series, there are nevertheless situations in which irregular influences can distort the trend estimates. For instance, the identification of extreme irregulars usually takes place when the annual seasonal reanalysis occurs, unless a specific investigation is requested. Consequently, extreme irregulars may affect the original and seasonally adjusted series values that become available between seasonal reanalyses, thus distorting the most recent trend estimates. Furthermore, even when the seasonal reanalysis takes place, there may be extreme observations for which a specific cause cannot be found. Such extremes are not modified before smoothing unless their specific causes are found and/or the resulting trend estimates are suspicious because they unlikely reflect the underlying movement of the series. The most recent couple of seasonally adjusted values are also not generally modified at a reanalysis, because there is usually insufficient information at the very end of the series to determine if the series is responding to a one-off occurrence, or undergoing a more sustained change in behaviour. Because the ABS publishes the trend series in conjunction with the seasonally adjusted series, the

7.1.1 Extreme irregulars continued

user may be able to determine whether the trend series has been left to respond to the presence of a large irregular influence. In such cases, a large spike in the seasonally adjusted series (either up or down) will be associated with a smaller and smoother bump in the trend series (as indicated in graph 7.1), which signifies the contribution of the irregular to the surrounding estimates of trend. If the analyst has an idea of the size of the extreme irregular from other sources, then equation (7.1) can be used to quantify the impact of the contaminant on the surrounding trend estimates in these situations.

When the trend is distorted by the presence of an extreme irregular, it may be tempting to emphasise the seasonally adjusted or original series in the analysis of underlying series behaviour instead. However, any such irregular will remain at full strength in both the original and seasonally adjusted series, as these series contain all the elements of the irregular component of behaviour. As a result, it may be more difficult to interpret the movements of these series about the affected period than it is to interpret the movements in the trend series. Furthermore, it should be noted that an unmodified extreme irregular in the original series can also result in distortion of the seasonally adjusted series. Unlike the trend series, this distortion tends to affect the seasonally adjusted series for a number of years before and after the occurrence of the extreme irregular, rather than just a relatively small number of periods.

A final point to note is that the above discussion has only considered the impact of a single, large contaminant on the trend estimates about it. In some situations, large extreme movements are followed immediately by swings in the opposite direction, of a similar order of magnitude. This can occur, for example, when processing of data is disrupted in one month, leading to a low figure, and a catch-up is made in the following month, leading to a high figure. In cases such as this, the adverse impact on the trend series is relatively short-lived because the Henderson filter and some of its surrogates dampen out much of the gyration as they pass through these episodes. The impact of these sudden abrupt swings in behaviour will remain fully in both the original and seasonally adjusted series, however, making interpretation of their period to period movements more difficult.

7.1.2 Trend breaks

On occasions there may be abrupt and appreciable changes in the level of a series between one period and the next, which reflect sudden and sustained changes in the underlying behaviour of the series. An example of this type of behaviour can be seen in graph 7.2, where the original data has undergone a sudden shift in level from July 1992. Such occurrences are here termed 'trend breaks' and can occur either because socioeconomic change has occurred abruptly and in a sustained fashion, or because of a change to collection or compilation methodology. An example of the former behaviour is a building society becoming a bank, thereby leading to a trend break in the deposit series for both the bank and non-bank sectors (although the series representing the deposits of all financial institutions will be unaffected by this change of status, because both banks and building societies are included). An example of the latter behaviour is the trend break that occurred in the Employed Persons series in 1986, when the reclassification of Unpaid Family Helpers resulted in additional persons being classified as employed.

7.1.2 Trend breaks continued

GRAPH 7.2 IMPACT OF TREND BREAK ON NON-BANK FINANCIAL INSTITUTIONS DEPOSITS

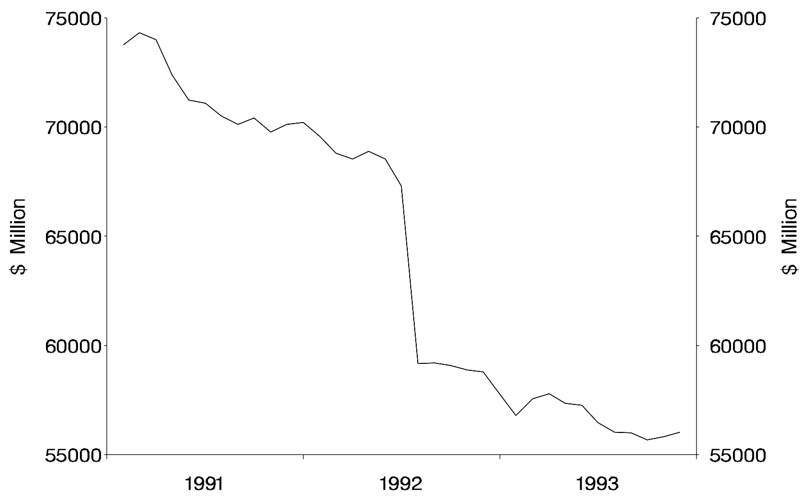

If a trend break is present in a seasonally adjusted series, then directly smoothing the series across the break with the Henderson and surrogate filters can give rise to distortion in the trend estimates for a number of periods around the break. An example of this distortion is indicated in graph 7.3. This graph shows the effect of applying the 13-term Henderson filter to a hypothetical monthly series that has experienced a sudden rise in level between December 1997 and January 1998. Note that the seasonally adjusted series contains no irregularity and apart from the abrupt change in level, displays a constantly increasing linear trend. The graph shows the distorted trend path that would be obtained by smoothing across the affected periods with the 13-term Henderson filter. If the trend break is considered to have occurred exactly midway between December 1997 and January 1998 (the two periods that reflect the sudden change in the level of the series), then it can be seen that the trend estimates are distorted for six months on either side of the trend break, with the two closest Henderson based trend estimates on either side experiencing the most significant amount of distortion.

7.1.2 Trend breaks continued

GRAPH 7.3 SMOOTHING ACROSS A TREND BREAK WITH 13-TERM HENDERSON FILTER

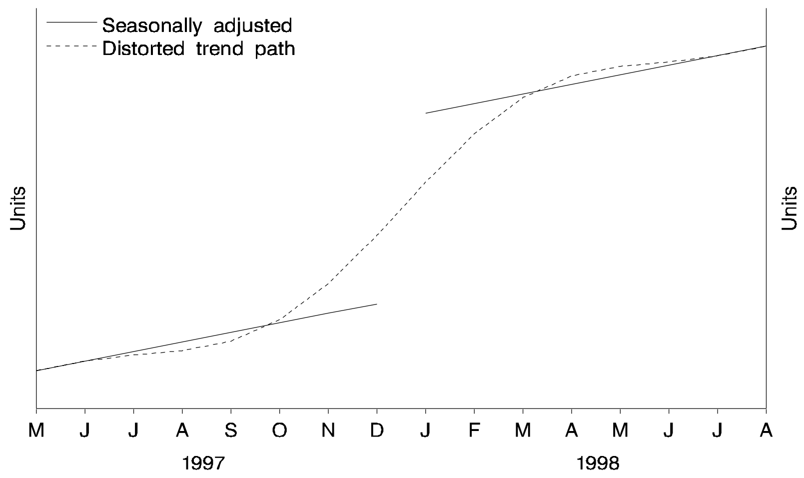

While graph 7.3 shows the eventual impact of smoothing across a historical trend break with the 13-term Henderson filter, graph 7.4 indicates the 'real time' impact of smoothing across a recent trend break using the six conventional surrogate filters to the 13-term Henderson filter. Series 1 is the trend path estimated immediately after the trend break occurs, Series 2 is the trend path estimated one month later by using an additional seasonally adjusted value to revise the last six trend estimates and calculate the most current estimate, and so on. It can be seen from this graph that smoothing directly across a recent trend break with the conventional surrogate filters can give a distorted impression of the level of the series at the current end, and a distorted impression of the current gradient or slope of the series. While the trend estimates start to reflect the level of the series fairly accurately three months after the occurrence of the trend break, it takes an additional two or three months before the trend estimates also start to accurately reflect the gradient of the trend at the current end.

7.1.2 Trend breaks
continued

GRAPH 7.4 SMOOTHING ACROSS A TREND BREAK WITH 13-TERM HENDERSON CONVENTIONAL SURROGATES

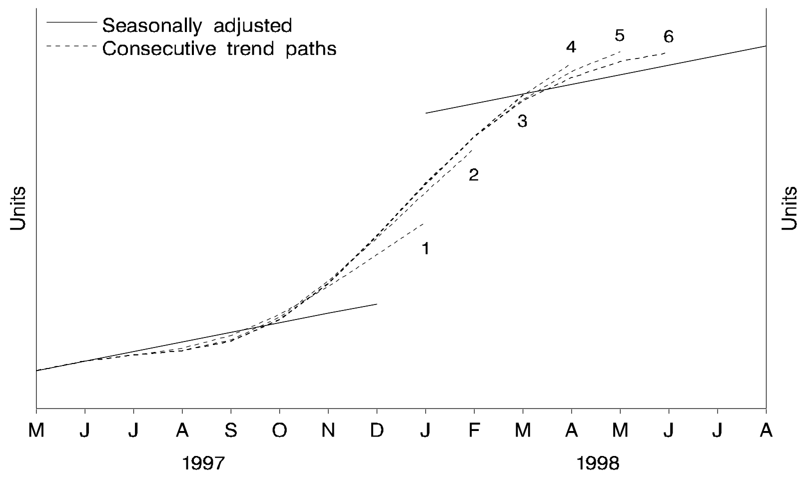

Given that the presence of a trend break can substantially distort a number of the trend estimates, the ABS takes such breaks into account, wherever possible, when smoothing seasonally adjusted series. Suppose the seasonally adjusted series is considered to consist of two parts, the part of the series before the trend break (up to December 1997 in graphs 7.3 and 7.4) and the part of the series after the trend break (from January 1998 in graphs 7.3 and 7.4). To correct the trend break for smoothing purposes, the ABS calculates a 'dislocation factor' indicating the extent of the change in trend level. This factor is then applied to the first part of the seasonally adjusted series to bring its level into line with the level of the second part of the series. The resulting modified seasonally adjusted series is then smoothed with the appropriate filters to derive a series of trend estimates. So that the trend estimates reflect the change in the level of the trend path as they should, the first part of the trend series is then moved back to the original level of the seasonally adjusted series over those periods, by dividing the estimates by the same dislocation factor. The result is a discontinuous series of trend estimates which reflect the abrupt change in the fundamental level of the series. If this procedure is applied to the pure trend series being smoothed in graphs 7.3 and 7.4, then the series of trend estimates obtained is exactly the same as the input series, that is, the above procedure captures the increasing linear trend of the seasonally adjusted series with its abrupt shift in level.

As with extreme irregulars, there are situations in which trend breaks are either not detected or not corrected. In particular, trend breaks are generally only detected during an annual seasonal reanalysis unless a specific investigation is requested, and are then only corrected (i.e. the smoothing procedure is modified in the above manner) when there is sufficient prior information about the occurrence and size or data available to estimate the extent of the break. To determine how the presence of an untreated trend break may affect trend estimates, the ABS investigated the impact of a wide variety of trend breaks on trend estimation for a number of different series (Bend 1993, Report on the Investigation Into Smoothing Across Trend Breaks). These empirical investigations used the method described above to determine a series of trend estimates that

7.1.2 Trend breaks continued

accounted for the trend break, as a basic guide to the 'true' trend behaviour, and then compared this series to the trend series obtained by applying the Henderson filters and their surrogates directly across the trend break without making a correction. The effects of varying the size of the trend break, the direction of the trend break (up or down), the change in the slope of the series' trend after the break, and the amount of irregularity in the series about the break, were then considered.

The amount of distortion observed for the Henderson based estimates was found to be directly related to the size of the trend break, while the direction, slope and amount of irregularity about the trend break, were found to have very little influence. As these last three factors had little effect on the distortion to the trend estimates, the results obtained under all the conditions considered were used to determine very robust average values indicating the relative amount of distortion to the Henderson based estimates of trend for trend breaks of varying sizes. The results for series smoothed by a 13 and 7-term Henderson filter are given in tables 7.3 and 7.4 respectively.

These tables indicate the percentage distortion without regard for sign of the trend estimates in the periods leading up to the occurrence of the trend break. As an example, table 7.3 indicates that if the level of the trend in graph 7.3 has risen by 5% between December 1997 and January 1998, then the Henderson based trend estimate for December will overestimate the trend level by 1.92%, the Henderson based trend estimate for November will overestimate the trend level by 0.84%, October will overestimate it by 0.10%, September will underestimate it by 0.25%, August will underestimate it by 0.25 per cent, and July will underestimate it by 0.10%. These earlier underestimates of the trend level are a characteristic of the weighting pattern of the Henderson moving average and relates specifically to the presence of negative weights.

It can be seen from these tables that the trend estimates are only distorted for six periods leading up to the trend break when smoothing with the 13-term Henderson filter, while only three periods are affected when using the 7-term Henderson filter. Furthermore, it is interesting to note that for both monthly and quarterly series, the relationship between the size of the trend break and the distortion of the Henderson based trend estimates around that break is a proportional one. Thus if the size of the trend break is twice as large, the distortion to the trend estimates will also be twice as large.

TABLE 7.3 TREND BREAK DISTORTION, MONTHLY

Percentage trend break	*M–5*	*M–4*	*M–3*	*M–2*	*M–1*	*M*
5	0.10	0.25	0.25	0.10	0.84	1.92
10	0.20	0.50	0.50	0.20	1.68	3.84
20	0.40	1.00	1.00	0.40	3.35	7.68

7.1.2 Trend breaks continued

TABLE 7.4: TREND BREAK DISTORTION, QUARTERLY

Percentage trend break	*Q–2*	*Q–1*	*Q*
5	0.30	0.07	1.40
10	0.60	0.14	2.80
20	1.20	0.28	5.60

It should be noted that, ignoring the direction of distortion, the actual level of distortion to the Henderson based estimates is nearly symmetric about the trend break. That is, if the trend break is considered to have occurred exactly midway between the periods reflecting the sudden change in the trend, then the absolute difference between the untreated and treated Henderson based trend estimates n periods after the trend break, will be approximately the same as the absolute difference between the untreated and treated Henderson based trend estimates n periods before the break (note that the values in tables 7.3 and 7.4 indicate the difference over the periods leading up to the trend break relative to the level of the series over those periods, and do not reflect the relative amount of distortion in the periods after the break as well because the trend experiences a sudden change in level on the other side of the trend break). The symmetric nature of the distortion to the Henderson based estimates can be seen in graph 7.3.

The above tables indicate the degree of distortion to estimates of trend based upon the Henderson filter when an appropriate allowance is not made for a historical trend break. The investigation also attempted to determine the impact of a trend break at or near the end of the series, where the trend estimates are calculated using the surrogate filters. Unlike the situation above, these investigations did not uncover a direct relationship between the size of a trend break and the percentage distortion to the surrogate based trend estimates. However, some general results were found to hold for monthly trend estimates calculated using the conventional surrogate filters, as will be described below.

When considering the effect of a recent trend break on the trend estimates at the current end of a series, there are two factors to take into account; whether the surrogate based trend estimates accurately reflect the true level of the trend component of the series, that is, whether the level of the trend has been distorted, and whether the surrogate based trend estimates accurately reflect the true growth rate of the trend component of the series, that is, whether the slope of the trend has been distorted (this factor is particularly important as distortion to the trend gradient can give a misleading impression of the way the behaviour of the trend is changing at the current end of the series). With regard to level, the investigations found that as a general rule, the surrogate based trend estimates approached the true trend level from the third period after the break onwards. This is evident in graph 7.4, as the surrogate trend estimates calculated for series 3, 4, 5 and 6 for the third period after the break and later all lie very close to the true trend level of series A. On the other hand, it was found that the true trend gradient is not reflected until a few periods later, and that how quickly it is reflected depends on how much the true trend gradient changes after the trend break. In general, the appropriate trend gradient is estimated more quickly when the true trend gradient

7.1.2 Trend breaks continued

either stays the same after the trend break or becomes slightly steeper, while the distortion of the estimated trend gradient is greatest whenever the true trend gradient adopts an opposite direction after the break, or maintains the same direction but becomes less steep. In the case illustrated in graph 7.4 (where the gradient or slope of the true trend path has not changed after the trend break) it can be seen that while the provisional trend estimates in series three to six lie very close to series A from the third period after the break onwards (in terms of trend level), the slope of series 3, 4, and to some extent five, are not truly indicative of the true trend slope of series A at the current end.

As indicated earlier, there are situations in which trend breaks are either not detected or not corrected. In these situations, where some of the trend estimates can be distorted, there may be a temptation to emphasise the original or seasonally adjusted series in the analysis of underlying series behaviour, rather than the trend series. The trend break will be present in both the original and seasonally adjusted series however, and therefore movements of the original and seasonally adjusted series over the break will be very difficult to interpret, especially when they also contain the complicating influences of seasonal and irregular factors, respectively. Furthermore, when a trend break cannot be corrected, its presence in the original series also has the potential to distort the seasonally adjusted series, as will be described below.

As previously indicated, the original series is analysed for seasonal influences approximately once a year, that is, there is an annual seasonal reanalysis. At this point in time, seasonal factors are determined for the time periods for which data is available and these factors are then applied to the original series data to remove the seasonal effects, thus giving rise to the seasonally adjusted series. Seasonal factors for forthcoming periods (forward factors) are also determined on the basis of the available data, and these are used to remove the seasonal influences from forthcoming original series values as they become available. When the next seasonal reanalysis takes place, the entire history of the original series is then used to re-evaluate the impact and extent of seasonal influences and revise the seasonally adjusted series.

When a trend break occurs before the re-estimation of seasonal factors, the current set of seasonal factors will be used to remove the seasonality from the original series values about the trend break. As these factors have not been influenced by the presence of the break, the full impact of the trend break will pass directly from the original series into the seasonally adjusted series, as it should. On the other hand, once the original series data containing the trend break has been included in the next annual reanalysis, the sudden change of level in the original series can distort the estimates of the seasonal factors, which in turn gives rise to distortion of the seasonally adjusted series. Unlike the trend estimates, which are only affected for a few periods either side of the trend break, the seasonally adjusted figures can be affected for a number of years before and after the occurrence of the trend break.

It should be noted that when contemporary trend breaks are detected but cannot be corrected, users are warned by the ABS of the presence of a trend break, and the consequences of the trend break remaining untreated.

7.1.3 Seasonal breaks

In addition to experiencing dislocations in the trend component of time series behaviour, a given series may also experience abrupt and sustained changes in the behaviour of the seasonal component. Such changes are here termed seasonal breaks. Seasonal patterns in time series appear to evolve slowly. There is, however, a distinction between gradual changes and abrupt changes to seasonal patterns. The seasonal adjustment procedure utilised by the ABS allows for some movement in time series' seasonal patterns. However, when an abrupt and permanent change in a seasonal pattern occurs, the ABS allows for a discontinuity in estimates of the seasonal pattern. This is done so that (much like the concept of a trend break — refer to section 7.1.2) estimates of a seasonal pattern are not allowed to 'smooth across' a break thus distorting estimates of the seasonal pattern on either side of the seasonal break. While in principal these abrupt changes should not affect the estimation of the trend (the trend being the underlying component of non-seasonal and non-irregular series behaviour) they do have the potential to distort the seasonally adjusted series from which the trend estimates are derived, thereby potentially leading to some distortion of the trend. In this respect, it is valuable to briefly consider the impact of seasonal breaks in relation to the smoothing procedure.

Seasonal breaks occur when seasonal patterns, trading day effects, payday effects, or moving holiday effects change abruptly. For example, the seasonal pattern of recorded export trade changed abruptly as a result of special reminder action requesting exporters to lodge their returns on time. In this case it was found that some of the activity previously attributed to January actually occurred in December, and thus when exporters lodged their returns on time, the December figures rose abruptly while the January figures correspondingly suddenly decreased by a substantial amount. Trading day patterns in series such as the Housing Finance series changed when the major banks switched from recording information on the last Wednesday of the month to the last working day of the month. Finally, the systematic moving holiday effect associated with the Australia Day holiday changed in 1988 when the national convention for its observance was varied between states and territories. Instead of this holiday falling on the first Monday after January 26 and impacting on January and February figures in the same way Australia-wide, the day of observance varied for the different states and territories and consequently affected state and territory figures differently.

Just as with trend breaks, untreated seasonal breaks in the original series can distort the seasonal factors derived from the series, and hence distort the seasonally adjusted series. In particular, if a seasonal break occurs after an annual reanalysis, then the current seasonal factors applied to the series will be determined solely on the basis of the old pattern of systematic calendar related behaviour, and thus will not reflect the new pattern of systematic calendar related behaviour. Consequently, applying the forward factors to the original series values containing the new pattern of behaviour can result in some residual calendar related effects remaining in the seasonally adjusted series for those periods. On the other hand, once the data affected by the seasonal break has been incorporated into the next reanalysis, the presence of two distinct patterns of calendar related behaviour will distort the seasonal factors around the break, and they will not completely capture either pattern. The result is that the seasonally adjusted series can contain some residual calendar related effects for a number of years either side of the seasonal break.

7.1.3 Seasonal breaks continued

The amount of distortion a seasonal break has on the trend estimates depends on the nature of the break. For instance, suppose a seasonal break results in residual trading day, pay day or moving holiday effects remaining in a seasonally adjusted monthly series. As the cycles associated with these behaviours lie predominantly in the range of two to six months, they will generally be filtered from the seasonally adjusted series by the 13-term Henderson filter and all but it's last two conventional surrogates (refer to graphs 5.4 and 5.5). On the other hand, the cycles associated with 'true' seasonality lie at the seasonal frequencies of, 12, 6, 4, 3, 2.4, and 2 months. As the 13-term Henderson and its conventional surrogates do not substantially attenuate the annual seasonal cycle of twelve months, the presence of residual seasonality in a monthly seasonally adjusted series may also lead to the presence of some residual seasonality in the trend series.

It should be noted that seasonal breaks are corrected wherever possible before determining seasonal factors, however, it takes at least three years before there is enough information available about the new pattern to recognise, reliably, the presence of a seasonal break, and correspondingly make a correction for it.

7.1.4 Compound breaks

In some instances a series will have a number of aspects of its character changed abruptly and simultaneously, such as when significant changes are made to economic policies, or the way data is collected. In such cases, the presence of both trend and seasonal breaks, in combination with changed irregular variation, create a confusing set of period movements in what has effectively become a new series. While these situations can be resolved, it generally takes a number of years before adequate corrections can be incorporated into the seasonally adjusted and trend series.

7.2 COMPATIBILITY OF TREND ESTIMATES

Another practical issue that is very important to consider when deriving and interpreting trend estimates is how compatible are the trend estimates for various series. Section 7.2.1 considers the trend estimates derived for a series composed of component series, and compares them to the trend estimates for the components from which they are derived. Section 7.2.2 considers the compatibility of trend estimates derived from quarterly and monthly series.

7.2.1 Component series

Many ABS series can be considered as being built up from, or comprised of, a number of component series. For instance, original series figures are available for the number of employed and unemployed persons, and adding these original series values together gives the original series figures for labour force. Similarly, the original series figures for the unemployment rate are derived directly by dividing the unemployed persons figures by the labour force figures. As a further example, the original series Balance of Trade figures are derived by subtracting the original series figures for imports from those for exports. When an original series, O_t, is a particular combination of a number of component series, and the trend series are available for both O_t and its components, the question arises as to whether the trend series for O_t can be obtained by combining the trend series for the components in the same way. In order to consider this question, two basic scenarios will be discussed; the first is the situation in which O_t is the sum or difference of a number of component series, and the second is the situation in which O_t is derived by multiplying or dividing a number of component series together.

7.2.1 Component series continued

ADDITIVE RELATIONSHIPS

The situation in which an original series can be obtained by summing or differencing other available original series arises frequently in ABS publications, as the majority of ABS time series are published at a number of different levels known to be particularly relevant to the decision making process. Thus the employed persons figures, for instance, can be found broken up by state, by gender, by full-time and part-time persons, or by combinations of these and other variables. These published original series figures are consistent across levels, that is, the original figures across all the states and territories will sum together to give the original series figures for the total number of employed persons in Australia, as will the figures for full-time and part-time persons, while the difference between the original figures for the total number of employed persons and male employed persons will be the figures for female employed persons, and so on. The same is not necessarily true for the seasonally adjusted and trend estimates, however. The difference is often called statistical discrepancy.

If a particular original series, O_t, is the sum (or difference) of a number of component series and there are seasonally adjusted series available for both O_t and the component series, then there are two separate approaches that can be taken towards producing trend estimates for O_t. One approach is to directly smooth the seasonally adjusted series for O_t, while another option is to smooth each of the component seasonally adjusted series, and then combine them together to produce a trend series. For instance, as the original figures for the total number of employed persons can be derived by adding the male employed persons figures to the female employed persons figures, a trend series for the total number of employed persons could be derived either by smoothing the seasonally adjusted series for total employed persons directly, or by smoothing each of the seasonally adjusted series for males and females, and adding these two smoothed series together. Generally, in order for these two methods to give the same results, three conditions must be satisfied:

- The seasonally adjusted series for O_t equals the sum (difference) of the seasonally adjusted components.
- None of the seasonally adjusted series are modified for extreme irregulars or trend breaks (in a non-linear fashion) during the smoothing process.
- All of the seasonally adjusted series are smoothed with the same set of filters, that is, the same length Henderson filter and the same set of surrogate filters are applied to the components as well as the composite series.

The first condition holds for a large number of seasonally adjusted series published by the ABS, as seasonally adjusted series published at a broad level are generally derived by aggregating a particular set of seasonally adjusted components that have been published at a finer level. This practice is implemented because summing the original series components and then seasonally adjusting the sum rarely gives exactly the same result as seasonally adjusting the components and adding the component seasonally adjusted series together (although the results are generally very similar). By choosing to publish the seasonally adjusted series derived by aggregating the seasonally adjusted components, the ABS minimises the number of explicit measures of the same quantity that are available, and retains consistency with results obtained at a finer level.

7.2.1 Component series continued

ADDITIVE RELATIONSHIPS *continued*

The second condition indicates that the modification of extreme irregulars or correction of trend breaks in either the component or composite seasonally adjusted series, can give rise to differences between the trend estimates derived from the two methods. This occurs because different component series may react to a particular extreme irregular or trend break in a different way (or not at all) compared to the composite series. Consequently, the corrections made at the component level and the composite level will generally not have a consistent impact on the trend estimates derived under the two methods. For example, the employed persons figures for the Northern Territory are often quite volatile, and are more susceptible to the influence of irregular influences than the figures for the larger states and territories. Consequently, the Northern Territory figures may contain extreme irregulars that are not present in either the other state and territory figures, or at the Australian level. In this case, the Northern Territory figures may be modified for these irregular influences before smoothing, thus affecting the trend series derived by aggregating the state level trend estimates. On the other hand, no equivalent modification will be made to the Australian level figures, as the same irregular influences have not given rise to the presence of extreme irregulars at the Australian level.

The third condition indicates that whenever different filters are used to filter the component series and the composite series, the trend estimates derived through the two processes will not necessarily be identical. This is particularly important for series in which tailored surrogate filters are used, for while the substantive part of the series is smoothed using a 2m+1 term Henderson filter, the first and last m trend estimates determined for each series are calculated using surrogate filters that depend on the relative amount of irregular to trend variability of the series being smoothed. Thus for quarterly series, where the ABS generally uses tailored surrogates in conjunction with the 7-term Henderson filter, this lack of consistency can affect the trend estimates for the first and last three quarters. While the differences are not generally significant, they tend to be greatest when the component series are substantially more irregular than the composite series, or when the composite series displays appreciable period to period growth, and are more noticeable for the trend estimates derived from the shortest surrogate filters.

To minimise the number of explicit measures of the same quantity that are available, the ABS generally prefers to derive trend estimates at a broad level by adding together the trend estimates for a particular set of published finer level components. While this prevents the published trend for the composite series from being different from the aggregation of those component trend series, there may still be combinations of other trend components which will give rise to different results. The difference is called statistical discrepancy. For example, the trend series for the total number of employed persons is determined by aggregating the trend series for full-time males, full-time females, part-time males and part-time females at Australia level. Consequently, any combination of these trend components to the total persons level will give rise to the trend series published for the total number of employed persons. However, there is no guarantee that the aggregation of the trend series at the state and territory level will give rise to the trend series published for employed persons at the total Australian level although it is true for original estimates.

7.2.1 Component series continued

ADDITIVE RELATIONSHIPS *continued*

It should be noted that there are occasions in which a composite series will be smoothed directly. Such situations generally occur when the direct seasonal adjustment of the composite series yields significantly higher quality in term of less irregularity and residual seasonality than the aggregating approach, or when there is more historical information available at the composite level. As an example, the retail trade figures for Australia are smoothed directly rather than summing the smoothed state and territory components[1].

MULTIPLICATIVE RELATIONSHIPS

Original series can also be derived by multiplying or dividing together various component series. For instance, the original series figures for the unemployment rate are obtained by dividing the figures for unemployed persons by the figures for labour force. Again it is possible to either smooth the composite series directly, or smooth the individual components and then multiply or divide the smoothed component series in order to derive a series of trend estimates. While these two procedures generally approximate one another very well, the differences tend to be most significant when:

- the component series are highly irregular and the resultant series is much less irregular due to cancelling effects
- a component series contains sharp turning points or steep regions of inflection which affect its smoothing and which become reduced or enlarged in the composite series.

Once again, where component trend series are available, the ABS generally manipulates those smoothed components arithmetically to obtain products or ratios, thereby avoiding two explicit measures of the same quantity.

7.2.2 Quarterly and monthly trend indicators

In order to make informed decisions about the state of the economy, or the impact of socioeconomic reforms, it is often necessary to consider series for a number of different collections. As will be seen below, whether the trend indicators for all these series are directly comparable will depend not only on the smoothing procedure employed for each series, but also on the type of data being smoothed.

Original, seasonally adjusted, and trend series can represent two distinct types of time series, namely, stock and flow series. The values in a stock series each measure activity at a particular point in time, while the values in a flow series measure accumulated activity over a particular period of time. Thus for a monthly stock series, for instance, the data value for a particular month is determined at some specific point in time during that month and does not involve any information about activity at any other time during the month. For a monthly flow series on the other hand, the data value for a particular month represents the total accumulated amount of activity that has taken place over the month. As the values of the Employed Persons series each indicate the number of employed persons at a particular point in time each month, the Employed Persons series is an example of a stock series. On the other hand, the Retail Trade series is an example

1 Additional reason is that the state and territory figures have only been collected since 1982, whereas Australian figures are available from 1962. In these situations, more reliable or informative trend estimates can be obtained by trending the composite series directly.

7.2.2 Quarterly and monthly trend indicators continued

of a flow series, because the retail trade figures each give the level of accumulated daily retail trade over the period of a month.

The following discussion considers the trend estimates derived for both quarterly and monthly data for stock and flow series, and considers how comparable these trend indicators are. The discussion begins with the most commonly encountered type of series published by the ABS, flow series.

FLOW SERIES

In order to illustrate the differences between trend estimates derived from quarterly flow series using the 7-term Henderson filter, and trend estimates derived from monthly flow series using the 13-term Henderson filter, it is necessary to consider flow data that is available in the form of both monthly and quarterly series. Suppose a monthly flow series is available. As the values in a flow series each indicate the accumulated level of activity over the period to which they refer, a quarterly flow series can be derived from this series by adding together the three monthly values in each quarter. Thus the June quarter figure for a particular year, Q_{June}, will be the sum of the April, May, and June monthly figures for that year, M_{April}, M_{May}, and M_{June}, respectively (in fact, any quarterly flow series can be considered as having arisen from a particular underlying monthly series in this way, the monthly series being the series that would have arisen by accumulating the flow data every month instead of every quarter).

Suppose the above quarterly flow series is smoothed by the 7-term Henderson filter, so that each quarterly figure in a particular 7-term span is weighted by the appropriate 7-term Henderson weight, and summed with the other weighted values in the 7-term span. Each quarterly figure is the sum of three monthly figures, and therefore, applying a weight to a particular quarterly figure is the same as weighting and adding together each of the three monthly values in that quarter. For example, if weight w is applied to the June quarterly figure in a 7-term Henderson moving average calculation, then

$$\begin{aligned} w \times (Q_{June}) &= w \times (M_{April} + M_{May} + M_{June}) \\ &= w \times M_{April} + w \times M_{May} + w \times M_{June} \end{aligned} \qquad (7.2)$$

Consequently, weighting and adding together the quarterly values in a 7-term span with the 7-term Henderson weights is the same as weighting and adding together the values in the corresponding 21-term monthly span with a related set of weights. (In fact, given that the weights of the 7-term Henderson filter are –0.059, 0.059, 0.294, 0.412, 0.294, 0.059, and –0.059, the first three weights in the 21-term weighting pattern are all –0.059, the second three weights are all 0.059, the third three are all 0.294, and so on). This 21-term weighting pattern then defines a moving average which can be moved along the monthly series in quarterly steps to give rise to the exact same quarterly estimates as are derived by smoothing the quarterly series with the 7-term Henderson filter.

Now suppose the monthly flow series is smoothed using the 13-term Henderson filter. These smoothed estimates will certainly not reflect quarterly behaviour, as they are only based on accumulated month to month activity rather than accumulated quarter to quarter activity. However, a quarterly indicator can be obtained by adding the monthly smoothed estimates in each quarter together. Again, as each monthly smoothed estimate is a weighted sum of monthly figures, and adding three of these weighted sums

7.2.2 Quarterly and monthly trend indicators continued

FLOW SERIES *continued*

together is the same as weighting the underlying monthly figures with a different set of weights, this process is equivalent to applying a particular moving average directly to the monthly figures. As 15 consecutive terms are required to determine three consecutive 13-term Henderson moving average calculations, this moving average has 15 terms, where each weight is derived by adding together up to three particular weights from the 13-term Henderson filter.

As indicated above, applying the 7-term Henderson filter to the quarterly flow series is the same as applying a particular 21-term moving average to the underlying monthly series in quarterly steps, while applying a 13-term Henderson filter to the monthly flow series and then adding the smoothed estimates together in quarters, is the same as applying a particular 15-term moving average to the underlying monthly series in quarterly steps. These two moving averages have different damping properties, as can be seen from graph 7.5. For example, it can be seen that the 50% strength cut-off for the 21-term moving average is about 14 months (or 4.67 quarters) while the 50% cut-off for the 15-term moving average is about 9 months (or 3 quarters). Consequently, the two filters are likely to derive different smoothed estimates (and therefore different estimates of movement) and they are also likely to derive smoothed estimates which capture different behavioural characteristics of the underlying data (although how different these characteristics are will depend upon what cycle strength exists in the series being smoothed).

GRAPH 7.5 FILTERS GIVING RISE TO QUARTERLY FLOW TREND INDICATORS

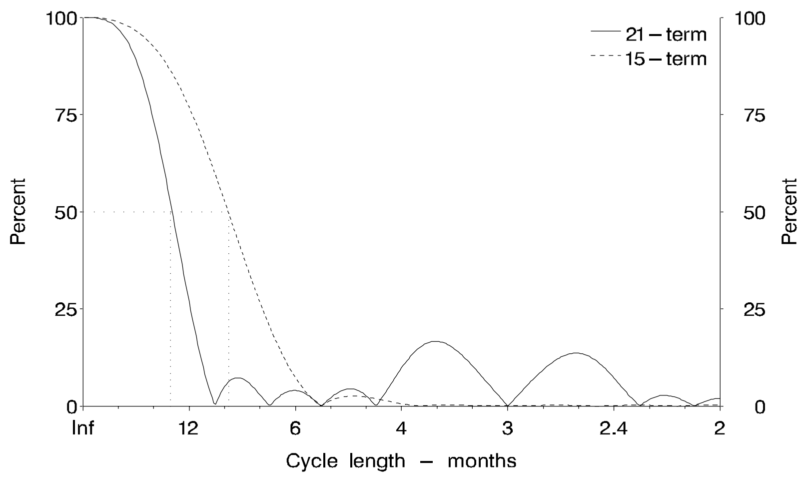

The above discussion indicates that deriving a quarterly trend indicator from a monthly flow series by applying the 13-term Henderson filter to the monthly series and then adding the trend estimates together over quarters, can result in trend estimates with different properties to those derived directly from quarterly series using the 7-term Henderson filter. Hence, if strict comparability of trend indicators is desired between monthly and quarterly flow data, original monthly flow data should be aggregated to

7.2.2 Quarterly and monthly trend indicators continued

FLOW SERIES *continued*

form a quarterly series, and this in turn seasonally adjusted and then filtered using the 7-term Henderson filter.

STOCK SERIES

Again, in order to illustrate the differences between trend estimates derived from monthly and quarterly stock series, it is necessary to consider stock data that is available in the form of both monthly and quarterly series. Suppose a monthly stock series is available, then a quarterly stock series can be derived from this series by taking every third monthly figure. Thus, for example, the March, June, September and December monthly figures for each year can be used to form a quarterly stock series. Note that the values in this quarterly series then represent a regular 'stocktake' of activity once every quarter, and give no information about activity at any other stage during the quarter. Smoothing this quarterly series with the 7-term Henderson filter is then equivalent to applying the 7-term Henderson weights to each of the monthly values that were used to define the quarterly series, while taking no account of the other monthly values at all. Thus, given that the weights of the 7-term Henderson filter are –0.059, 0.059, 0.294, 0.412, 0.294, 0.059, and –0.059, smoothing the quarterly series with the 7-term Henderson filter is equivalent to moving along the monthly series in quarterly steps and applying a 19-term filter with weights –0.059, 0, 0, 0.059, 0, 0, 0.294, 0, 0, 0.412, 0, 0, 0.294, 0, 0, 0.059, 0, 0, –0.059.

Now suppose the monthly stock series referred to above has been smoothed using the 13-term Henderson filter. A quarterly indicator can then be derived by taking every third estimate of this smoothed monthly series, such as the March, June, September and December smoothed estimates. These estimates are not the same as the estimates determined directly from the quarterly stock series using the 7-term Henderson filter however, as they are based on different monthly observations and have been derived using different weights. For example, the March smoothed estimate obtained from the monthly series will be a function of the 13 monthly values centred around March, which involves all the monthly values from the previous September to the following September. In contrast, the smoothed estimate obtained for March from the quarterly series will be based solely upon the 7 monthly values used to derive the quarterly series about March, namely, the June, September, and December values immediately before and after March. Whereas March would receive a weight of 0.24 in the 13-term Henderson calculation, it would receive a weight of 0.412 in the 7-term Henderson calculation, the previous December would receive a weight of 0.066 compared to 0.294, the previous September a weight of –0.019 compared to 0.059, and so on. As a result, it can be seen that the two smoothing procedures correspond to applying different filters to the underlying monthly series, and may therefore give rise to smoothed series which capture different behavioural characteristics. The damping properties of these two filters are shown in graph 7.6.

7.2.2 Quarterly and monthly trend indicators continued

STOCK SERIES *continued*

GRAPH 7.6 FILTERS GIVING RISE TO QUARTERLY STOCK TREND INDICATORS

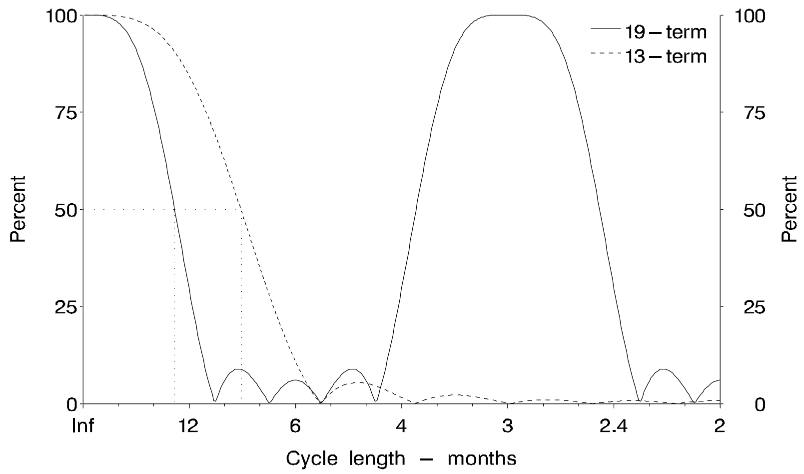

Again, the above discussion illustrates how trend indicators derived from quarterly stock series using the 7-term Henderson filter may have different properties from those obtained from monthly stock series by applying the 13-term Henderson filter. Consequently, if strict comparability is required between the trend indicators for monthly and quarterly stock series, then the monthly stock data should be used to derive a quarterly stock series, which should in turn be seasonally adjusted and then smoothed with the 7-term Henderson filter.

7.3 FILTERING REQUIREMENTS

7.3.1 Filter length

As previously stated, the ABS uses the 13-term Henderson filter to determine trend estimates for monthly series, and the 7-term Henderson filter to determine trend estimates for quarterly series. Applied to monthly and quarterly series respectively, these filters capture the medium-term business cycles and the long-term cycles associated with the trend, while simultaneously eliminating a substantial amount of the short-term irregularity present in time series. While this approach satisfies general user needs, there may be situations in which it is desirable to capture different behavioural characteristics from a particular series. In these situations, it may be useful to consider choosing a Henderson filter with different filtering characteristics to smooth the series. The ABS has developed procedures that permit it to use Henderson filters of length 5–33 terms, and the impact of each of these filters on cycle strength is summarised in table 4.5.

When choosing the length of Henderson filter to be used to smooth a particular series, a number of issues should be taken into account. Firstly, it is important to consider what behavioural characteristics the filtering process should capture, so that a filter can be chosen which will effectively reduce or eliminate the range of unwanted cycles. A second issue to consider is the extent to which revisions can be tolerated in the most recent trend estimates. This is important when choosing filter length, because the longer the

7.3.1 Filter length continued

Henderson filter employed, the more surrogate filters are required to determine trend estimates at the end of the series, and therefore the more recent trend estimates are subject to revision. As revisions can affect the timely detection of stable turning points near the end of a series, this may be particularly relevant to those focusing on behaviour at the current end.

Filter selection may also be related to what cycles are present in the data to be smoothed and to what extent. For instance, suppose in a monthly series a user wishes to dampen all cycles whose length is shorter than a year. As the 50% cycle strength threshold for the 21-term Henderson filter is about 12 periods (refer to table 7.5) this may seem an appropriate filter to apply. However, if the short-term cyclical variations in the data are only in the range of 5 months or less, then as the 50% threshold for the 7-term Henderson filter is about 5 periods, this filter could be used to achieve the same goal. The 7-term Henderson filter would then have a less severe end-point problem than the 21-term Henderson filter, and only the last three smoothed estimates would require the use of surrogate filters, rather than the last ten.

In some situations it may seem that a longer Henderson filter is required to achieve the damping objective. However in these cases the problem of having additional revisions at the end of the series may be such that it is not considered practical to use a filter of this length, and a shorter filter may be selected as a compromise. For example, suppose a user is not interested in capturing short-term cyclical behaviour in the region of one year or less, and these cycles are thought to contribute significantly to the behaviour of a monthly series. In order to sufficiently remove the influences of these cycles, it may seem appropriate to choose a Henderson filter with particularly strong reduction characteristics in the short-term range. This would require the use of a 31-term Henderson filter, which will leave less than 10% of the strength of cycles shorter than 13 months remaining in the series (see table 7.5). While this filter might seem the most desirable to achieve the dampening objective, the last 15 months would have to be estimated using surrogate filters and would therefore be subject to revision. As a result, it may be more practical to use the 21-term Henderson filter, which still eliminates 50% or more of the strength of cycles shorter than a year, and which will have a less severe end-point problem.

TABLE 7.5 HENDERSON MOVING AVERAGE GAIN, At indicated cycle length (measured in periods)

No. terms	10%	25%	50%	75%	90%
7	3.49	3.88	4.63	5.88	7.74
21	9.10	10.27	12.42	15.93	21.06
31	13.00	14.68	17.77	22.81	30.18

7.3.2 Other considerations

There are a number of other practical considerations to note when filtering series to obtain trends. Firstly, the 13 and 7-term Henderson filters do not eliminate the seasonal cycle in monthly and quarterly data respectively. Consequently, these filters should not be applied to respective monthly or quarterly series that contain seasonal patterns, as some of the strength of these patterns of behaviour will pass into the smoothed series, *7.3.2 Other considerations continued*

thereby producing a misleading trend indicator. As a result, these filters should only be used to determine trend estimates from data that has been seasonally adjusted or is non-seasonal.

Another consideration is the amount of data that must be available before trend estimates can be determined for a new series. As indicated above, if a new series contains the classic annual seasonal cycle, then the series will have to be seasonally adjusted before smoothing can be undertaken. This requires a significant amount of data to be available, as a minimum of three years of data is required to seasonally adjust a series, and usually more than five years of data is required before the seasonal adjustment can be considered reliable. If a new monthly series is non-seasonal, then the shortest period of data that could be smoothed by a symmetric 13-term Henderson filter and its surrogates is thirteen months — the first and last six trend values being obtained by using the surrogate filters. Similarly, when using the 7-term Henderson filter and its surrogates to smooth a non-seasonal quarterly series, a minimum of seven observations would be required.

CHAPTER 8 CONCLUSION

SUMMARY

This Information Paper has discussed the distinguishing characteristics of the original, seasonally adjusted and trend series, and has explained why it is difficult to identify the underlying behaviour of the series from either the original or seasonally adjusted data. In the case of the original data, the seasonal and irregular factors obscure the relatively subtle changes of the trend. In the case of the seasonally adjusted data, the irregular variation generally dominates the behaviour of the period to period movements, obscuring the trend. Given that the irregular variation is attributable to statistical errors and very short lived socioeconomic events, it is important to recognise the adverse impact reliance on these movements may have for business decisions, policy advice, commentary and analysis.

This paper has also illustrated how a number of other indicators of socioeconomic activity commonly derived from original and seasonally adjusted series can provide misleading information when taken to represent trend behaviour. Some of these indicators are strongly influenced by the irregular component. Other indicators give a distorted impression of the shape and level of turning points in the trend, or delay their detection. Such deficiencies may lead to incorrect assessments being made about leads and lags in economic systems or result in the misinterpretation of cause and effect relationships.

To minimise these common difficulties, the ABS is producing trend estimates for a widening range of subject areas. The statistical procedure used smoothes out a significant proportion of the high frequency irregularity in a series, while simultaneously tracking the types of curvatures present in the series and accurately capturing the timing of trend turning points. These features allow recent turning points and rates of increase or decrease to be detected quickly, facilitating the responsive setting of appropriate policies and commercial actions.

ABS trend estimates generally provide a reliable indicator of trend behaviour, however their usefulness as a 'current' indicator can be reduced whenever there is a high degree of irregular variation present in the original and seasonally adjusted data. Abrupt changes in the time series characteristics of the original series components may also impact on the reliability of the trend estimates, although such occurrences also affect the original and seasonally adjusted series and generally to a larger extent. Particular care must therefore be exercised when attempting to infer trend behaviour from either the original, seasonally adjusted or trend series in these situations.

Given the nature of time series, as discussed in this paper, the ABS generally recommends that those using its statistics for the purpose of monitoring underlying trends should pay attention to its explicit trend estimates. These non-seasonal and non-irregular series generally provide a superior guide to the underlying impact of

SUMMARY *continued*

socioeconomic forces, than do the alternative measures of trend behaviour commonly derived from original and seasonally adjusted data.

This paper describes the most commonly used method in ABS to produce the trend estimates - smoothing seasonally adjusted estimates by applying an appropriate Henderson filter (and its surrogates).

As mentioned in Chapter 6 both seasonally adjusted series and trend estimates are subject to revisions as new observations may result in changes in the estimated seasonal pattern and in using better smoothing filters for the latest part of the series. However, the revisions of trend estimates are relatively larger and converge quicker to the final estimates than the revisions of seasonally adjusted series do.

To reduce the revisions in size and length, the ABS will continue conducting its own research and development, and evaluating new methods and informing its users once the new methods are proved to be valuable in improving the quality of ABS statistics.

APPENDIX 1 BENCHMARK SERIES

UNEMPLOYED PERSON: COMPARISON OF COMMONLY USED INDICATORS

Table A1[1] compares the performance of various popular forms of analysis, in order to determine how well they pinpoint the turning points of the Unemployed Persons series, a major economic indicator. Column 1 contains the seasonally adjusted (s.a.) Unemployed Persons series as published by the ABS. This series has been smoothed by the procedure discussed in this paper, to produce a series of trend estimates, also as published by the ABS. The month to month percentage changes of the ABS trend series appear in column 2. This series is the benchmark series, disclosing turning points whenever the growth rate in column 2 changes sign, for example, in Jul/Aug 1983, Nov/Dec 1985, Mar/Apr 1987, Oct/ Nov 1989, Nov/ Dec 1992, Apr/ May 1993, Sep/Oct 1993, Jun/Jul 1995 Apr/May 1997, Mar/Apr 1998, etc. Columns 3 to 12 show to what extent alternative forms of analysis of the seasonally adjusted Unemployed persons series delay the detection of the turning points evident in column 2.

Col. 1 Benchmark series — seasonally adjusted
$= SA_t$

Col. 2 Monthly percentage change of trend
$= (T_t / T_{t-1} - 1) \times 100\%$

Col. 3 Monthly percentage change of s.a.
$= (SA_t / SA_{t-1} - 1) \times 100\%$

Col. 4 Yearly percentage change of s.a.
$= (SA_t / SA_{t-12} - 1) \times 100\%$

Col. 5 Monthly percentage change of s.a. 12-month

$$= \left(\frac{\sum_{i=0}^{11} SA_{t-i}}{\sum_{j=0}^{11} SA_{t-j-1}} - 1 \right) \times 100\%$$

Col. 6 Yearly percentage change of s.a. 12-month

$$= \left(\frac{\sum_{i=0}^{11} SA_{t-i}}{\sum_{j=0}^{11} SA_{t-j-12}} - 1 \right) \times 100\%$$

Col. 7 Quarterly percentage change of s.a. 3-month

$$= \left(\frac{\sum_{i=0}^{2} SA_{t-i}}{\sum_{j=0}^{2} SA_{t-j-3}} - 1 \right) \times 100\%$$

Col. 8 Yearly percentage change of s.a. 3-month

$$= \left(\frac{\sum_{i=0}^{2} SA_{t-i}}{\sum_{j=0}^{2} SA_{t-j-12}} - 1 \right) \times 100\%$$

1 The table only lists part of time series for illustration purpose. The full time span time series are available on request.

UNEMPLOYED PERSON: COMPARISON OF COMMONLY USED INDICATORS *continued*

Col. 9 6-month smoothed percentage change annualised

$$= \left(SA_t \Big/ \frac{1}{12}\sum_{j=0}^{11} SA_{t-j-1} - 1 \right)^{\frac{12}{6.5}} \times 100\%$$

Col. 10 Annualised s.a. 3-month to previous 3-months

$$= \left(\left(\frac{\sum_{i=0}^{2} SA_{t-i}}{\sum_{j=0}^{2} SA_{t-j-3}} \right)^{4} - 1 \right) \times 100\%$$

Col. 11 Annualised s.a. 6-month to previous 6-months

$$= \left(\left(\frac{\sum_{i=0}^{5} SA_{t-i}}{\sum_{j=0}^{5} SA_{t-j-6}} \right)^{2} - 1 \right) \times 100\%$$

Col. 12 Annualised half yearly change sa 3-month

$$= \left(\left(\frac{\sum_{i=0}^{2} SA_{t-i}}{\sum_{j=0}^{2} SA_{t-j-6}} \right)^{2} - 1 \right) \times 100\%$$

TABLE A1 COMPARISON OF COMMONLY USED INDICATORS — UNEMPLOYED PERSON SERIES

Date	Col. 1	Col. 2	Col. 3	Col. 4	Col. 5	Col. 6	Col. 7	Col. 8	Col. 9	Col. 10	Col. 11	Col. 12
Feb–78	429.60	—	—	—	—	—	—	—	—	—	—	—
Mar–78	408.40	–0.69	–4.94	—	—	—	—	—	—	—	—	—
Apr–78	405.60	–0.77	–0.69	—	—	—	—	—	—	—	—	—
May–78	400.20	–0.59	–1.33	—	—	—	—	—	—	—	—	—
Jun–78	408.20	–0.41	2.00	—	—	—	—	—	—	—	—	—
Jul–78	401.00	–0.24	–1.75	—	—	—	–2.76	—	—	–10.58	—	—
Aug–78	419.20	0.01	4.53	—	—	—	1.17	—	—	4.75	—	—
Sep–78	396.00	0.19	–5.52	—	—	—	0.19	—	—	0.76	—	—
Oct–78	397.00	0.16	0.25	—	—	—	0.24	—	—	0.96	—	–4.98
Nov–78	409.50	0.35	3.14	—	—	—	–2.10	—	—	–8.14	—	–1.90
Dec–78	414.80	0.53	1.31	—	—	—	0.42	—	—	1.71	—	1.23
Jan–79	413.20	0.56	–0.40	—	—	—	2.09	—	—	8.61	–0.26	4.72
Feb–79	406.90	0.41	–1.51	–5.28	–0.46	—	2.70	—	–0.73	11.24	–0.40	1.09
Mar–79	411.70	0.07	1.17	0.81	0.07	—	0.86	—	2.29	3.47	1.91	2.59
Apr–79	424.50	–0.39	3.10	4.66	0.39	—	0.45	–0.04	8.10	1.81	4.94	5.16
May–79	402.30	–0.44	–5.21	0.54	0.04	—	0.28	2.01	–2.78	1.14	3.54	6.07
Jun–79	408.80	–0.31	1.62	0.17	0.01	—	0.31	1.79	0.06	1.24	2.47	2.35
Jul–79	409.20	–0.20	0.08	2.04	0.17	—	–1.83	0.91	0.20	–7.12	1.12	–2.76
Aug–79	398.30	–0.07	–2.67	–4.99	–0.43	—	–1.79	–0.98	–4.98	–6.98	1.42	–3.01
Sep–79	410.70	0.05	3.12	3.70	0.30	—	–1.41	0.16	1.37	–5.54	0.05	–2.21
Oct–79	420.30	—	2.34	5.87	0.47	—	0.73	1.41	5.21	2.96	–2.48	–2.21
Nov–79	399.30	–0.27	–5.00	–2.48	–0.21	—	1.16	2.31	–5.12	4.70	–2.16	–1.31
Dec–79	403.80	–0.55	1.12	–2.67	–0.23	—	0.43	0.17	–2.78	1.74	–2.09	–1.97
Jan–80	407.90	–0.36	1.03	–1.27	–0.11	0.02	–1.48	–2.14	–0.51	–5.80	–1.87	–1.52
Feb–80	399.70	0.06	–2.02	–1.78	–0.15	0.34	–1.54	–1.91	–4.01	–6.01	–1.06	–0.80
Mar–80	396.70	0.57	–0.75	–3.64	–0.31	–0.03	–1.56	–2.23	–5.07	–6.09	–2.11	–2.25
Apr–80	409.70	0.99	3.29	–3.47	–0.30	–0.72	–0.40	–2.97	1.35	–1.60	–2.63	–3.73
May–80	418.40	1.05	2.10	3.99	0.33	–0.44	1.11	–1.10	5.91	4.50	–0.84	–0.89
Jun–80	426.10	0.66	1.86	4.23	0.35	–0.10	4.15	1.51	8.90	17.64	1.40	5.11
Jul–80	409.70	0.13	–3.87	0.12	0.01	–0.25	3.98	2.77	0.59	16.90	1.65	7.25
Aug–80	415.60	–0.37	1.45	4.35	0.35	0.53	2.17	2.89	3.29	8.96	2.84	6.71
Sep–80	417.00	–0.78	0.33	1.53	0.13	0.35	–0.96	1.98	3.24	–3.78	5.74	6.39

— nil or rounded to zero (including null cells)

TABLE A1 COMPARISON OF COMMONLY USED INDICATORS — UNEMPLOYED PERSON SERIES *continued*

Date	Col. 1	Col. 2	Col. 3	Col. 4	Col. 5	Col. 6	Col. 7	Col. 8	Col. 9	Col. 10	Col. 11	Col. 12
Oct–80	406.90	–0.91	–2.43	–3.20	–0.27	–0.39	–1.18	0.83	–1.57	–4.62	6.42	5.60
Nov–80	394.80	–0.92	–2.96	–1.13	–0.09	–0.28	–2.62	–0.95	–6.40	–10.06	2.79	–1.01
Dec–80	409.70	–0.95	3.77	1.47	0.12	0.07	–2.48	–0.98	0.39	–9.57	–0.40	–6.72
Jan–81	401.80	–1.23	–1.93	–1.50	–0.12	0.05	–2.67	–0.39	–3.38	–10.26	–1.18	–7.48
Feb–81	382.10	–1.48	–4.89	–4.39	–0.36	–0.16	–2.05	–1.46	–11.72	–7.95	–5.09	–9.01
Mar–81	391.50	–1.37	2.44	–1.32	–0.11	0.04	–2.97	–2.40	–7.09	–11.35	–8.59	–10.47
Apr–81	379.10	–0.91	–3.16	–7.48	–0.63	–0.29	–4.44	–4.43	–12.26	–16.63	–10.50	–13.50
May–81	379.60	–0.17	0.14	–9.26	–0.80	–1.41	–3.64	–6.09	–11.01	–13.79	–9.96	–10.92
Jun–81	366.80	0.70	–3.38	–13.93	–1.23	–2.97	–4.25	–10.27	–15.25	–15.94	–12.06	–13.68
Jul–81	400.20	1.38	9.12	–2.30	–0.20	–3.17	–0.53	–8.58	1.87	–2.10	–11.62	–9.66
Aug–81	399.80	1.49	–0.11	–3.81	–0.33	–3.84	1.44	–6.76	2.04	5.89	–7.75	–4.46
Sep–81	401.40	1.37	0.41	–3.74	–0.33	–4.28	6.74	–3.29	3.44	29.82	–4.96	4.46
Oct–81	401.20	1.27	–0.05	–1.40	–0.12	–4.13	4.86	–2.99	3.97	20.90	–0.85	8.79
Nov–81	410.90	1.15	2.44	4.09	0.34	–3.71	4.01	–0.42	8.94	17.02	3.13	11.32
Dec–81	412.80	1.20	0.45	0.76	0.07	–3.77	1.96	1.12	9.16	8.08	11.20	18.45
Jan–82	414.00	1.44	0.29	3.03	0.26	–3.40	2.95	2.61	9.60	12.32	12.62	16.53
Feb–82	438.10	1.79	5.81	14.63	1.18	–1.91	4.23	5.97	21.08	18.04	14.42	17.53
Mar–82	436.50	1.88	–0.35	11.51	0.94	–0.88	5.20	9.63	17.71	22.46	16.69	15.04
Apr–82	438.70	1.74	0.51	15.73	1.23	0.98	6.10	13.93	16.78	26.74	17.95	19.31
May–82	454.60	1.81	3.61	19.75	1.53	3.35	5.13	15.62	21.91	22.17	18.83	20.09
Jun–82	468.50	2.30	3.05	27.72	2.04	6.78	5.68	20.99	25.30	24.73	19.32	23.59
Jul–82	479.50	3.47	2.36	19.81	1.56	8.66	6.79	22.32	26.02	30.08	23.88	28.40
Aug–82	484.40	4.95	1.01	21.16	1.64	10.81	7.71	22.76	24.77	34.59	24.21	28.23
Sep–82	516.00	6.13	6.53	28.55	2.19	13.61	8.67	23.18	36.07	39.47	27.81	31.89
Oct–82	579.70	6.52	12.34	44.49	3.33	17.54	12.65	31.41	62.06	61.06	36.69	44.74
Nov–82	605.40	6.16	4.45	47.33	3.51	21.25	18.76	40.18	65.31	98.94	45.83	63.63
Dec–82	654.80	5.28	8.16	58.63	4.23	26.29	24.33	50.21	79.27	138.96	56.90	82.56
Jan–83	648.40	4.10	–0.99	56.62	3.93	30.91	20.80	54.20	63.07	112.95	65.01	85.20
Feb–83	676.80	3.00	4.38	54.49	3.85	34.36	16.40	56.54	64.38	83.56	77.61	91.10
Mar–83	692.10	2.13	2.27	58.56	3.97	38.40	9.64	56.55	59.80	44.51	84.26	85.83
Apr–83	706.40	1.66	2.06	61.01	4.00	42.17	8.73	58.02	54.42	39.77	78.43	72.52
May–83	721.90	1.34	2.20	58.81	3.84	45.41	7.09	59.46	49.54	31.54	71.25	55.39
Jun–83	715.40	0.91	–0.91	52.71	3.41	47.36	6.27	57.42	37.17	27.52	57.10	35.75
Jul–83	723.10	0.29	1.09	50.81	3.26	49.82	4.10	54.04	31.52	17.45	47.42	28.12
Aug–83	715.60	–0.46	–1.04	47.75	2.99	51.81	1.59	50.40	21.60	6.51	34.85	18.36
Sep–83	730.70	–1.05	2.10	41.61	2.70	52.57	1.20	46.60	19.66	4.89	25.04	15.65
Oct–83	702.70	–1.54	–3.83	21.23	1.51	49.87	–0.53	36.01	6.00	–2.10	17.01	7.23
Nov–83	681.30	–1.69	–3.05	12.52	0.91	46.11	–1.83	24.31	–2.62	–7.14	8.38	–0.55
Dec–83	668.20	–1.61	–1.92	2.03	0.16	40.41	–5.41	11.53	–7.61	–19.94	2.93	–8.36
Jan–84	673.30	–1.28	0.77	3.84	0.30	35.51	–5.88	5.98	–6.57	–21.51	–3.00	–12.34
Feb–84	665.50	–0.89	–1.16	–1.67	–0.13	30.31	–5.09	1.36	–9.06	–18.86	–7.03	–13.20
Mar–84	655.50	–0.68	–1.50	–5.29	–0.44	24.79	–2.82	–1.14	–11.34	–10.80	–11.99	–15.49
Apr–84	669.20	–0.71	2.09	–5.26	–0.44	19.46	–1.61	–4.10	–7.15	–6.28	–13.29	–14.23
May–84	637.10	–0.94	–4.80	–11.75	–1.02	13.88	–2.25	–7.48	–14.50	–8.69	–13.56	–13.93
Jun–84	653.10	–1.04	2.51	–8.71	–0.76	9.28	–1.75	–8.60	–8.79	–6.82	–12.29	–8.83
Jul–84	625.00	–0.99	–4.29	–13.57	–1.20	4.57	–3.77	–11.35	–14.70	–14.24	–12.36	–10.35
Aug–84	631.80	–0.83	1.08	–11.72	–1.04	0.47	–2.65	–11.34	–11.03	–10.17	–11.76	–9.44
Sep–84	627.20	–0.72	–0.73	–14.16	–1.29	–3.43	–3.85	–13.16	–10.52	–14.52	–9.78	–10.75
Oct–84	622.30	–0.54	–0.78	–11.44	–1.02	–5.83	–1.77	–12.46	–9.65	–6.90	–10.50	–10.65
Nov–84	624.40	–0.36	0.33	–8.35	–0.73	–7.37	–1.89	–11.38	–7.36	–7.33	–9.10	–8.76
Dec–84	612.70	–0.34	–1.87	–8.30	–0.72	–8.18	–1.30	–9.39	–9.32	–5.12	–10.35	–9.94
Jan–85	618.30	–0.18	0.91	–8.18	–0.72	–9.10	–1.38	–8.27	–6.56	–5.40	–8.46	–6.15

UNEMPLOYED PERSON: COMPARISON OF COMMONLY USED INDICATORS *continued*

Graph A1 shows that a turning point (indicated by a vertical line) at July/August 1983 is correctly detected by the trend movement (changing from postive to negative). All the rest of the methods detect the turning point much later or are problematic in terms of certainty.

UNEMPLOYED PERSON: COMPARISON OF COMMONLY USED INDICATORS *continued*

GRAPH A1 COMPARSION OF DETECTED TURNING POINTS

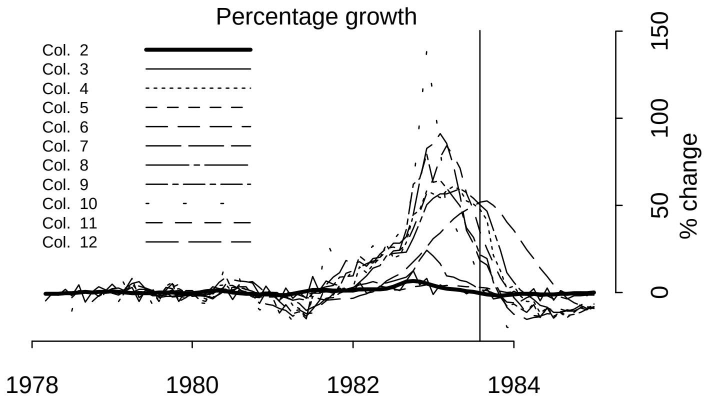

NOTE ON GRADUATION BY ADJUSTED AVERAGE[1]

BY ROBERT HENDERSON

Paper

The expression 'graduation by adjusted average' is used to include those methods by which the graduated value of a function is determined by adding together a number of adjacent terms each multiplied by a numerical factor. Summation formulas are thus included as particular cases but other formulas are intended to be included which cannot be expressed in the form of successive summations. We will first investigate the formula which will arise when the graduated value is determined by fitting an algebraic function of the third degree to the ungraduated values. It is known that in the case of an algebraic function the method of least squares and the method of moments give identical results.

Let U_0 denote the ungraduated value of the function corresponding to the graduated value to be determined and suppose that only terms from U_{-n} to U_n are to be taken into account and that W_x is the weight to be assigned to the term U_x for all values of x.

Then if $a+bx+cx^2+dx^3$ is the algebraic function determined and consequently, a the graduated value of U_0 we have the following four equations in a, b, c and d:

$$\sum_{x=-n}^{+n}(a+bx+cx^2+dx^3)W_x = \sum_{x=-n}^{+n} W_x U_x,$$

$$\sum_{x=-n}^{+n}(ax+bx^2+cx^3+dx^4)W_x = \sum_{x=-n}^{+n} xW_x U_x,$$

$$\sum_{k=-n}^{+n}(ax^2+bx^3+cx^4+dx^5)W_x = \sum_{k=-n}^{+n} x^2 W_x U_x,$$

$$\sum_{k=-n}^{+n}(ax^3+bx^4+cx^5+dx^6)W_x = \sum_{k=-n}^{+n} x^3 W_x U_x.$$

Or if we designate $\sum_{x=-n}^{+n} x^r W_x$ by s_r,

$$s_0 a+s_1 b+s_2 c+s_3 d = \sum_{x=-n}^{+n} W_x U_x,$$

$$s_1 a+s_2 b+s_3 c+s_4 d = \sum_{x=-n}^{+n} xW_x U_x,$$

$$s_1 a+s_2 b+s_3 c+s_4 d = \sum_{x=-n}^{+n} x^2 W_x U_x,$$

$$s_1 a+s_2 b+s_3 c+s_4 d = \sum_{x=-n}^{+n} x^3 W_x U_x,$$

These are linear equations in a,b,c and d, therefore we have:

$$a = h\sum_{x=-n}^{+n} W_x U_x + j\sum_{x=-n}^{+n} xW_x U_x + k\sum_{x=-n}^{+n} x^2 W_x U_x + l\sum_{x=-n}^{+n} x^3 W_x U_x$$

$$= \sum_{x=-n}^{+n}(h+jx+kx^2+lx^3)W_x U_x$$

provided h, j, k and l are determined so that

1 Henderson, R., Note on Graduation by Adjusted Average, Trans. Amer. Soc. Actuaries, 17 (1916), 43–48.

Paper continued

$$hs_0+js_1+ks_2+ls_3 \text{ or } \sum_{x=-n}^{+n}(h+jx+kx^2+lx^3)W_x=1$$

$$hs_1+js_2+ks_3+ls_4 \text{ or } \sum_{x=-n}^{+n}x(h+jx+kx^2+lx^3)W_x=0$$

$$hs_2+js_3+ks_4+ls_5 \text{ or } \sum_{x=-n}^{+n}x^2(h+jx+kx^2+lx^3)W_x=0$$

$$hs_3+js_4+ks_5+ls_6 \text{ or } \sum_{x=-n}^{+n}x^3(h+jx+kx^2+lx^3)W_x=0$$

If then we write V_x for $(h+jx+kx^2+lx^3)W_x$ and put U'_0 for the graduated value of U_0 we have

$$U'_0=\sum_{x=-n}^{+n}V_xU_x$$

where

$$\sum_{x=-n}^{+n}V_x=1,$$

$$\sum_{x=-n}^{+n}xV_x=0,$$

$$\sum_{x=-n}^{+n}x^2V_x=0,$$

$$\sum_{x=-n}^{+n}x^3V_x=0.$$

Thus far no limitations have been placed on the values of W_x so that none are placed on those of V_x except the four conditions expressed above. If, however, we suppose that the values of W_x are all necessarily positive then the values of V_x will change sign along with $(h+jx+kx^2+lx^3)$. Conversely any series of values of V_x satisfying these four conditions and changing sign not more than three times, may be considered as derived in this manner from a series of values of W_x of the form $V_x/(h+jx+kx^2+lx^3)$, where the function $h+jx+kx^2+lx^3$ changes sign in the same intervals as V_x.

Suppose now that the values of W_x are symmetrical about W_0, so that $W_x=W_{-x}$. Then, where r is odd, we have $s_r=0$ since positive and negative terms will cancel one another, so that the four equations in a, b, c and d reduce to:

$$s_0a+s_2c=\sum_{x=-n}^{+n}W_xU_x,$$

$$s_2b+s_4d=\sum_{x=-n}^{+n}xW_xU_x,$$

$$s_2a+s_4c=\sum_{x=-n}^{+n}x^2W_xU_x,$$

$$s_4b+s_6d=\sum_{x=-n}^{+n}x^3W_xU_x.$$

Since a appears only in the first and third of these equations we have

$$a=h\sum_{x=-n}^{+n}W_xU_x+k\sum_{x=-n}^{+n}x^2W_xU_x$$

$$=\sum_{x=-n}^{+n}(h+kx^2)W_xU_x$$

where

$$hs_0+ks_2 \text{ or } \sum_{x=-n}^{+n}(h+kx^2)W_x=1,$$

$$hs_2+ks_4 \text{ or } \sum_{x=-n}^{+n}x^2(h+kx^2)W_x=0.$$

Paper continued

Here
$V_x = (h + kx^2)W_x$

Hence the values of V_x are also symmetrical. Conversely any symmetrical graduation formula such that $V_x = V_{-x}$ with two changes of sign may be considered as derived from a symmetrical set of weights W_x of the form $V_x/(h+kx^2)$ where $h+kx^2$ changes sign in the same intervals as V_x .

It is interesting to determine the relative weights which would produce certain well known summation formulas.

In Woolhouse's formula V_0 to V_4 inclusive are positive, V_0 being equal to $\frac{1}{5}$. V_5 is zero and V_6 and V_7 are negative. This suggests $(25-x^2)/125$ as the transforming factor. This gives the following series of values for W_0 to W_7 inclusive, W_{-x} being, it will be remembered, equal to W_x for all values of x.
$1; 1; 1; \frac{7}{16}; \frac{6}{18}; ?; \frac{4}{22}; \frac{3}{24}$

The value of W_5 is indeterminate but the analogy of the neighbouring terms suggests the value $\frac{5}{20}$ or $\frac{1}{4}$.

Similarly in Higham's formula we have for the values of W_0 to W_8 inclusive,
$1; 1; \frac{6}{7}; \frac{5}{8}; \frac{1}{3}; \frac{1}{4}; \frac{2}{11}; \frac{1}{12}; \frac{1}{39}$

And in Karup's formula we have for the values of W_0 to W_9 inclusive,
$1; \frac{19}{20}; \frac{29}{35}; \frac{53}{80}; \frac{7}{15}; \frac{1}{4}; \frac{8}{55}; \frac{3}{40}; \frac{2}{65}; \frac{1}{140}$

In all of these cases the vanishing of V_5 gives a suggestion of the transforming factor, but in the case of Spencer's 21-term formula none of the terms vanish. The change of sign, however, takes place between $x = 5$ and $x = 6$ and investigation shows that in order that W_x should always decrease as x increases numerically the change of sign in the transforming factor must occur between $x^2 = 29.5$ and $x^2 = 33.25$

Taking $(30-x^2)/175$ as the factor the series of weights is
$1; \frac{57}{58}; \frac{47}{52}; \frac{11}{14}; \frac{9}{14}; \frac{3}{5}; \frac{1}{6}; \frac{5}{38}; \frac{5}{68}; \frac{1}{34}; \frac{1}{140}$

and taking $2(33-x^2)/385$ as the factor it is
$1; \frac{627}{640}; \frac{517}{580}; \frac{121}{160}; \frac{99}{170}; \frac{33}{80}; \frac{11}{30}; \frac{11}{64}; \frac{11}{124}; \frac{11}{320}; \frac{11}{1340}$

This result illustrates the fact that the same series of values of V_x may be derived from many different series for values of W_x .

Let us now revert to the case where the values of V_x are not necessarily symmetrical and suppose that graduated values are determined for a number of successive terms, the same values of V_x being used for each. Then generally we have
$U'_y = \sum_{x=-\infty}^{+\infty} V_x U_{y+x} = \sum_{z=-\infty}^{+\infty} V_{z-y} U_z$ where z = y + x

While the above summation is indicated as extending infinitely in both directions the terms under the summation will vanish where $V_x = 0$.

Let us now examine the conditions for the smoothest possible graduated series. In doing so we shall tentatively adopt the usual criterion of smoothness, namely the smallness of the mean square of the error in the third difference. Then
$\Delta^3 U'_y = \sum_{z=-\infty}^{+\infty} U_x \Delta^3 V_{z-y}$

where the variable on both sides is y. Then the mean square of the error in $\Delta^3 U'_y$ is

Paper *continued*

$$\sum_{z=-\infty}^{+\infty} e_z^2(\Delta^3 V_{z-y})^2$$

where e_z^2 is the mean square of the error in U_Z. Summing this for all values of y and assuming that, except for a finite range of values of z, e_z^2 vanishes, we have for the sum

$$\sum_{y=-\infty}^{+\infty}\sum_{z=-\infty}^{+\infty} e_z^2(\Delta^3 V_{z-y})^2 = \sum_{z=-\infty}^{+\infty} e_z^2 \sum_{y=-\infty}^{+\infty}(\Delta^3 V_{z-y})^2$$

$$= \sum_{z=-\infty}^{+\infty} e_z^2 \sum_{x=-\infty}^{+\infty}(\Delta^3 V_{x-3})^2$$

since$\Delta^3 V_{z-y}$(y variable) = - $\Delta^3\ V_{x-3}$

But $\sum_{x=-\infty}^{+\infty}(\Delta^3 v_{x-3})^2$ is independent of z so that the expression reduces to the product of two factors, one of which, $\sum_{z=-\infty}^{+\infty} e_z^2$, is independent of the graduation formula. It is only necessary therefore to consider the conditions for a minimum value of $\sum_{x=-\infty}^{+\infty}(\Delta^3 v_{x-3})^2$.

The differential coefficient of this sum with respect to V_x is

$$2(\Delta^3 v_{x-3} - 3\Delta^3 v_{x-2} + 3\Delta^3 v_{x-1} - \Delta^3 v_x) = -2\Delta^6 v_{x-3}$$

and therefore the complete differential is

$$\sum_{x=-\infty}^{+\infty} -2\Delta^6 v_{x-3}\delta v_x$$

But owing to the four conditions limiting V_x we have, where all values of V_x except those from V_{-n} to V_n inclusive vanish:

$$\sum_{x=-n}^{+n} \delta V_x = 0,$$

$$\sum_{x=-n}^{+n} x\delta V_x = 0,$$

$$\sum_{x=-n}^{+n} x^2\delta V_x = 0,$$

$$\sum_{x=-n}^{+n} x^3\delta V_x = 0.$$

For the smoothest possible series with a graduation formula extending from $-n$ to n inclusive we must have

$$\sum_{x=-n}^{+n} \Delta^6 v_{x-3}\delta v_x = 0$$

for all values of the differentials satisfying these four equations, or in other words we must have a relation of the form

$$\Delta^6 v_{x-3} = \kappa + \lambda x + \mu x^2 + \nu x^3$$

for values of x from $-n$ to n inclusive.

But $\Delta^6 v_{x-3}$ includes terms from v_{x-3} to v_{x+3} inclusive. Therefore v_x must be an algebraic function of not more than the ninth degree for values of x from $-(n+3)$ to $n+3$ inclusive.

But we know that V_x vanishes for six values of x, viz., $-(n+3)$; $-(n+2)$; $-(n+1)$; $n+1$; $n+2$; and $n+3$; therefore $\{(n+1)^2 - x^2\}\{(n+2)^2 - x^2\}\{(n+3)^2 - x^2\}$ must be a factor of V_x which must therefore take the form

$$V_x = [(n+1)^2 - x^2][(n+2)^2 - x^2][(n+3)^2 - x^2][h + jx + kx^2 + lx^3]$$

where h, j, k and l are to be determined from the four conditions limiting V_x .

Paper continued

This is seen to be exactly the formula which would be arrived at by assigning weights $W_x = \{(n+1)^2 - x^2\}\{(n+2)^2 - x^2\}\{(n+3)^2 - x^2\}$ and as the values of W_x are symmetrical we see that j and l vanish and the expression reduces to
$V_x = (h + kx^2)W_x$

Suppose now, to obtain another view of the matter we approach it from the standpoint of a smooth series of weights and determine the condition for a minimum value of $\sum_{x=-\infty}^{+\infty} (\Delta^3 W_k)^2$ subject to the condition that all values of W_x except those from W_{-n} to W_n inclusive should vanish and that $\sum_{x=-n}^{+n} W_x = K$. In this case we have by reasoning similar to the above a relation $\Delta^6 W_{x-3} = k$ for values of x from $-n$ to n inclusive. Hence W_x is of the sixth degree from $-(n+3)$ to $n+3$ inclusive. But W_x vanishes for the six values of x, $-(n+3)$; $-(n+2)$; $-(n+1)$; $(n+1)$; $(n+2)$; $(n+3)$; hence W_x takes the form:
$W_x = W[(n+1)^2 - x^2][(n+2)^2 - x^2][(n+3)^2 - x^2]$

The same formula is thus obtained as before, since only relative weights affect the final result. We thus see that the smoothest possible graduated series from a formula of given range is obtained by assigning the smoothest possible series of weights to the successive terms.

In the preceding investigation it has been assumed that the weight assigned to a term depends only on its distance from the term for which the graduated value is being determined. In view, however, of the fact that the intrinsic weights of the terms of the ungraduated series are sometimes known to vary from term to term it may be desirable to take these variations into account and to assign to each term a weight calculated as the product of two factors one $W_z = W_{y+x}$, peculiar to the term itself and the other $W_x = W_{z-y}$ dependent only on the distance. We must then substitute $W_x W_{y+x}$ for W_x and consider the series as probably not symmetrical. Analogy would suggest that the smoothest series of graduated values would be obtained by giving to W_z values inversely proportional to e_z^2 and to W_x the same values as before. The problem of testing this suggestion is left to the readers.

Handwritten Note

See addendum, Vol. XIX, p.29, where it is stated that putting $m = n + 2$, so that number of terms $= 2m - 3$.

$$V_x = \frac{315[3m^2-(16+11x^2)][m^2-(x-1)^2][m^2-x^2][m^2-(x+1)^2]}{8m(m^2-1)(4m^2-1)(4m^2-9)(4m^2-25)}$$

Proof,

$$V_x = (h + kx^2)[m^2(m^2-1)^2 - (3m^4+1)x^2 + (3m^2+2)x^4 - x^6]$$

$$= (h + kx^2)[(m^2 - x^2)^3 - 2(m^4 - x^4) + (m^2 - x^2)] \text{ ; m = n + 2}$$

Let $R_t = \sum_{x=-(m-1)}^{+m} \frac{x^t}{2m^{t+1}}$

therefore

$$R_0 = 1$$

$$R_2 = \frac{1}{3} + \frac{1}{6m^2}$$

$$R_4 = \frac{1}{5} + \frac{1}{3m^2} - \frac{1}{30m^4}$$

$$R_6 = \frac{1}{7} + \frac{1}{2m^2} - \frac{1}{6m^4} + \frac{1}{42m^6}$$

$$R_8 = \frac{1}{9} + \frac{2}{3m^2} - \frac{7}{15m^4} + \frac{2}{9m^6} - \frac{1}{30m^8}$$

$$R_{10} = \frac{1}{11} + \frac{5}{6m^2} - \frac{1}{m^4} + \frac{1}{m^6} - \frac{1}{2m^8} + \frac{5}{66m^{10}}$$

Handwritten Note continued

Then since $V_{-(m-1)}$, V_{m-1} and V_m all vanish

$$\frac{1}{2}\sum_{x=-(m-1)}^{+m} V_x = h[m^7(R_0 - 3R_2 + 3R_4 - R_6) - 2m^5(R_0 - R_4) + m^3(R_0 - R_2)]$$

$$= +k[m^9(R_2 - 3R_4 + 3R_6 - R_8) - 2m^7(R_2 - R_6) + m^5(R_2 - R_4)]$$

$$= \frac{1}{2};$$

$$\frac{1}{2}\sum_{x=-(m-1)}^{+m} x^2 V_x = h[m^9(R_2 - 3R_4 + 3R_6 - R_8) - 2m^7(R_2 - R_6) + m^5(R_2 - R_4)]$$

$$= +k[m^{11}(R_4 - 3R_6 + 3R_8 - R_{10}) - 2m^9(R_4 - R_8) + m^7(R_4 - R_6)]$$

$$= 0$$

or

$$\frac{m(m^2-1)(4m^2-1)(4m^2-9)}{315}\{9h + (m^2-4)k\} = \frac{1}{2}$$

$$\frac{m(m^2-1)(m^2-4)(4m^2-1)(4m^2-9)}{3465}\{11h + (3m^2-16)k\} = 0.$$

Let $h = (3m^2 - 16)l$ and $k = -11l$. Then $9h + (m^2 - 4)k = (16m^2 - 100)l = 4(4m^2 - 25)l$,

HENDERSON MOVING AVERAGE WEIGHTS

TABLE A2 SYMMETRIC HENDERSON FILTER WEIGHTS

Period/ Terms	*5*	*7*	*9*	*11*	*13*	*23*	*33*
–16	—	—	—	—	—	—	–0.001
–15	—	—	—	—	—	—	–0.004
–14	—	—	—	—	—	—	–0.008
–13	—	—	—	—	—	—	–0.010
–12	—	—	—	—	—	—	–0.011
–11	—	—	—	—	—	–0.004	–0.010
–10	—	—	—	—	—	–0.011	–0.005
–9	—	—	—	—	—	–0.016	0.003
–8	—	—	—	—	—	–0.015	0.014
–7	—	—	—	—	—	–0.005	0.027
–6	—	—	—	—	–0.019	0.013	0.042
–5	—	—	—	–0.028	–0.028	0.039	0.058
–4	—	—	–0.041	–0.027	0.000	0.068	0.072
–3	—	–0.059	–0.010	0.036	0.065	0.097	0.085
–2	–0.073	0.059	0.118	0.141	0.147	0.122	0.095
–1	0.294	0.294	0.267	0.239	0.214	0.138	0.101
0	0.559	0.413	0.331	0.278	0.240	0.144	0.103
+1	0.294	0.294	0.267	0.239	0.214	0.138	0.101
+2	–0.073	0.059	0.118	0.141	0.147	0.122	0.095
+3	—	–0.059	–0.010	0.036	0.065	0.097	0.085
+4	—	—	–0.041	–0.027	0.000	0.068	0.072
+5	—	—	—	–0.028	–0.028	0.039	0.058
+6	—	—	—	—	–0.019	0.013	0.042
+7	—	—	—	—	—	–0.005	0.027
+8	—	—	—	—	—	–0.015	0.014
+9	—	—	—	—	—	–0.016	0.003
+10	—	—	—	—	—	–0.011	–0.005
+11	—	—	—	—	—	–0.004	–0.010
+12	—	—	—	—	—	—	–0.011
+13	—	—	—	—	—	—	–0.010
+14	—	—	—	—	—	—	–0.008
+15	—	—	—	—	—	—	–0.004
+16	—	—	—	—	—	—	–0.001

— nil or rounded to zero (including null cells)

The symmetric Henderson moving averages can be computed for any number of odd terms, from 3 onwards. The general formula for each particular weight in the Henderson moving average of length 2m+1 is:

$$w_j = \frac{315[(m+1)^2-j^2][(m+2)^2-j^2][(m+3)^2-j^2][3(m+2)^2-11j^2-16]}{8(m+2)[(m+2)^2-1][4(m+2)^2-1][4(m+2)^2-9][4(m+2)^2-25]}$$

where $-m \leq j \leq m$.

The Gain function can be calculated by the following formula:

$$G(\omega) = \left| \sum_{i=-m}^{m} W_j \cos(\omega j) \right|$$

where ω is the frequency.

SURROGATE FILTERS

SURROGATE FILTERS

This Appendix gives the formula for the derivation of the weights of the surrogate filters, under the conditions specified in section 5.1. The filter design features are explained in Doherty (2001).

If the most recent time period for which data has been observed is period N, then the last smoothed observation that can be determined using a 2m+1 term Henderson filter is the (N–m)th. Thus for each period N-i with $0 \le i \le m-1$, a surrogate filter is required to determine a smoothed estimate. Consequently, a symmetric Henderson moving average of length 2m+1 has m surrogate filters, surrogate filter i ($0 \le i \le m-1$) being used to calculate the smoothed estimate for period N–i. Surrogate filter i has m+i+1 weights $W_j^{(i)}$, which are applied to the last m+i+1 series observations, that is, weight $W_j^{(i)}$ is applied to the observation from period N-i+j, where $-m \le j \le i$. Thus, the trend T at time point N-i is given by

$$T_{N-i} = \sum_{j=-m}^{i} W_j^{(i)} A_{N-i+j}$$

for $i = 0, 1, \ldots, m-1$, where A_{N-i+j} is the seasonally adjusted series value at time N-i+j.

The weights are computed as follows:

$$W_j^{(i)} = a + bj + W_j$$

where $-m \le j \le i$ and W_j = weight j of the main Henderson filter.

$$a = \overline{W} + \frac{m-i}{2} b$$

$$b = \frac{\lambda}{1 + \left[\lambda \frac{(m+i+2)(m+i+1)(m+i)}{12} \right]} \sum_{j=-i}^{m} j \left[W_j + \overline{W} \right]$$

and

$$\overline{W} = \frac{1}{m+i+1} \sum_{j=i+1}^{m} W_j$$

$$\lambda = \frac{4}{\pi (EWP)^2}$$

where EWP stands for the End Weight Parameter. For example, for a 13-term Henderson filter and an end weight parameter of 3.5, the surrogate weights are computed as follows:

EWP=3.5, and therefore

$$\lambda = \frac{4}{3.1415926 \times 3.5^2} = 0.1039379$$

Then, for i=0,

$$\overline{W} = \frac{1}{7} \sum_{j=1}^{6} W_j = \frac{0.3799714}{7} = 0.05428163$$

$$b = \frac{0.1039379}{1 + 0.1039379 \times \frac{8 \times 7 \times 6}{12}} \sum_{j=0}^{6} j \left[W_j + 0.05428163 \right] = 0.04226406$$

$$a = 0.05428163 + \frac{6}{2} \times 0.04226406 = 0.1810783$$

and therefore

SURROGATE FILTERS *continued*

$$W_{-6}^{(0)} = 0.1810738 - 6\times 0.04226406 - 0.0193499 = -0.0918604$$
$$W_{-5}^{(0)} = 0.1810738 - 5\times 0.04226406 - 0.0278638 = -0.0581103$$
$$W_{-4}^{(0)} = 0.1810738 - 4\times 0.04226406 + 0.0000000 = +0.0120176$$
$$W_{-3}^{(0)} = 0.1810738 - 3\times 0.04226406 + 0.0654918 = +0.1197734$$
$$W_{-2}^{(0)} = 0.1810738 - 2\times 0.04226406 + 0.1473565 = +0.2439022$$
$$W_{-1}^{(0)} = 0.1810738 - 1\times 0.04226406 + 0.2143368 = +0.3531465$$
$$W_{0}^{(0)} = 0.1810738 - 0\times 0.04226406 + 0.2400572 = +0.4211310$$

For i=1,

$$\overline{W} = \frac{1}{8}\sum_{j=2}^{6} W_j = \frac{1}{8}\times 0.1656347 = 0.02070433$$
$$b = \frac{0.1039379}{1+0.1039379\times\frac{9\times 8\times 7}{12}}\sum_{j=-1}^{6} j\left[W_j + 0.02070433\right] = 0.01258897$$
$$a = 0.02070433 + \frac{5}{2}\times b = 0.05217677$$

and therefore

$$W_{-6}^{(1)} = 0.05217677 - 6\times 0.01258897 - 0.0193499 = -0.0427069$$
$$W_{-5}^{(1)} = 0.05217677 - 5\times 0.01258897 - 0.0278638 = -0.0386319$$
$$W_{-4}^{(1)} = 0.05217677 - 4\times 0.01258897 + 0.0000000 = +0.0018209$$
$$W_{-3}^{(1)} = 0.05217677 - 3\times 0.01258897 + 0.0654918 = +0.0799016$$
$$W_{-2}^{(1)} = 0.05217677 - 2\times 0.01258897 + 0.1473565 = +0.1743553$$
$$W_{-1}^{(1)} = 0.05217677 - 1\times 0.01258897 + 0.2143368 = +0.2539245$$
$$W_{0}^{(1)} = 0.05217677 - 0\times 0.01258897 + 0.2400572 = +0.2922339$$
$$W_{1}^{(1)} = 0.05217677 + 1\times 0.01258897 + 0.2143368 = +0.2791025$$

The surrogate weights for i>1 can be obtained in a similar fashion.

It should be noted that in relation to the discussion given in section 5.2, the weights derived by extending the series k periods into the future are the weights $W_j^{(i)}$ specified by the equation above with i = m–k. Furthermore, the parameter a, used to define the slope of the extended path associated with the kth surrogate, is given by

$$a = \frac{G^2}{1+G^2}$$

where

$$G^2 = \frac{12}{(2m+1-k)(2m-k)(2m+2-k)\lambda}$$

The Gain function of an asymmetric Henderson filter can be calculated by

$$G(\omega) = \sqrt{\left(\sum_{j=-m}^{i} W_j^{(i)}\cos(\omega j)\right)^2 + \left(\sum_{j=-m}^{i} W_j^{(i)}\sin(\omega j)\right)^2}$$

GLOSSARY

Absolute value The value of a data point, or movement, ignoring its sign.

Amplitude Amplitude is a non-negative number measuring the size of the peak (and/or trough) of a cycle.

Autocorrelation The autocorrelation coefficient at lag k is the autocovariance at lag k, divided by the variance of the series. Denoting the autovariance by $\gamma(k)$, and the variance by $\gamma(0)$, the autocorrelation $\rho(k)$ is given by

$$\rho(k) = \frac{\gamma(k)}{\gamma(0)}$$

estimated as

$$\hat{\rho}(k) = r(k) = \frac{c(k)}{c(0)}$$

Autocovariance The autocovariance of a time series X at lag k is the covariance of observations k time units apart. Denoted by $\gamma(t, t+k)$ the autocovariance is given by

$$\gamma(t, t+k) = cov(X_t, X_{t+k}) = E[(X_t - M_t)(X_{t+k} - M_{t+k})]$$

where $M_t = E(X_t)$ is the mean or expected value of the series at time t. For a stationary series neither the mean nor the autocovariances depend on t and we may write $\gamma(k) = E[(X_t - M)(X_{t+k} - M)]$. Using this notation, $\gamma(0)$ is the variance of the time series. Given a sample realisation of a time series of length L, we may estimate the autocovariance using the formula

$$\hat{\gamma}(k) = c(k) = \frac{1}{L}\sum_{t=1}^{L-k}(x_t - \bar{x})(x_{t+k} - \bar{x})$$

Sometimes the divisor L – k –1 is used instead of L to give an unbiased estimate.

Business cycle The cycle of recession and recovery apparent in some economic time series.

Cubic polynomial Refer to polynomial.

Cycle A regularly repeating fluctuation in a time series. Strictly speaking, the length of the cycle should be fixed, for example, 1 year for the fundamental seasonal cycle. However the term is also used for less regular repetitions, for instance, the business cycle.

Cycle length The elapsed time, or number of time units, for each repetition of a cycle.

Filter A transformation of time series data to isolate certain frequency components, for example trend or seasonal filters.

Frequency A cyclical fluctuation which repeats itself once every p time units (ie. has cycle length p) is said to have frequency = 1/p. The frequency is the fraction or multiple of the cycle's length that passes by in one time unit.

Frequency domain analysis The properties and characteristics of the frequency set (or fluctuations) exhibited by a time series. For a continuous time series this set ranges from zero frequency (for a straight line trend) to infinity. For a discrete time series the frequency range is from 0 to 1/2 (cycle per unit time), as the highest frequency observable (called the Nyquist frequency — Jenkins, 1968) in this case is a cycle of length 2 time units. All higher frequencies alias, or are indistinguishable from, other cycles in the stated range. The frequency domain is an alternative to the time domain, for examining the behaviour of a time series. Analysis of a time series in the frequency domain usually involves the Spectral analysis.

Gain The change in amplitude of a cycle caused by the application of a filter. If the gain for cycles of a particular frequency is greater than 1, cycles of this frequency are said to be amplified, if the gain is less than 1, attenuation is said to occur. The gain of a linear filter is given by

$$G(\omega) = \sqrt{\sum_{j=-\infty}^{\infty} \left\{ \left[a_j \cos(\omega j) \right]^2 + \left[a_j \sin(\omega j) \right]^2 \right\}}$$

where a_j is the jth weight of a linear filter with weights $\{ a_{-\infty}, \ldots, a_{-1}, a_0, a_1, \ldots, a_{\infty} \}$, ω is the frequency usually expressed in radians and ω j is the product of the frequency and the integer j.

Irregular or Residual/irregular The remaining component of a series after taking account of seasonal and trend components. This variation is caused by chance events such as unusual weather conditions, industrial disturbances, political pronouncements, statistical errors, etc. affecting particular values of the series.

Linear filter A filter consisting of a linear combination of the observations of a time series. Thus applying a linear filter with weights $\{a_t\}$ to a time series $\{X_t\}$ gives a new series $\{Y_t\}$ defined by

$$Y_t = \sum_{j=-\infty}^{\infty} a_j x_{t+j}$$

The properties of the linear filter are given by the function

$$A(\omega) = \sum_{j=-\infty}^{\infty} a_j e^{-i\omega j} = |A(\omega)| e^{i\, Arg[A(\omega)]}$$

$A(\omega)$ is the transfer function of the linear filter whose absolute value $|A(\omega)|$ gives the gain caused by the filter at frequency ω and whose argument gives the phase shift. If the weights a_j sum to 1 the linear filter is sometimes called a moving average. In this case the filtered series has the same level as a constant input series. If the weights are symmetric $(a_j = a_{-j})$ the phase shift of the filter is negligible for those cycles remaining with appreciable strength.

Moving average A linear filter the weights of which sum to 1. A simple moving average has equal weights.

Peak The maximum point within a repetition of a cycle.

Phase A measure of the position of a cycle relative to the time frame of observation.

Phase shift The time lag that exists between the filter output and the filter input. When a filter is applied to a time series, as well as changing the amplitude of cycles of a particular frequency, it may also alter their phase. For example, differentiating a sine curve gives a cosine curve, a shift of – $2\pi/4$ radians from the original cycle. The phase shift of a linear filter is given by

$$Arg[A(\omega)] = Arc\tan\left\{\frac{b(\omega)}{a(\omega)}\right) = \left\{ \begin{array}{rl} \arctan\left(\frac{b(\omega)}{a(\omega)}\right), & a(\omega) > 0 \\ \arctan\left(\frac{b(\omega)}{a(\omega)}\right) + \pi, & a(\omega) < 0, b(\omega) \geqslant 0 \\ \arctan\left(\frac{b(\omega)}{a(\omega)}\right) - \pi, & a(\omega) < 0, b(\omega) < 0 \\ \frac{\pi}{2}, & a(\omega) = 0, b(\omega) \geqslant 0 \\ \frac{\pi}{2}, & a(\omega) = 0, b(\omega) < 0 \end{array} \right\}$$

the angle of the transfer function $A(\omega)$ at frequency ω, where

$$a(\omega) = \frac{[A(\omega) + \overline{A(\omega)}]}{2}$$

$$b(\omega) = \frac{[A(\omega) - \overline{A(\omega)}]}{2i}$$

is the conjugate of $A(\omega)$.

A symmetric linear filter produces no appreciable phase shift.

Polynomial A polynomial of degree n is an algebraic expression of the form:

$$a_o + a_1 x + a_2 x^2 + a_3 x^3 + \ldots\ldots + a_n x^n$$

Polynomial *continued*

A cubic polynomial has the form:
$a_o + a_1 x + a_2 x^2 + a_3 x^3$

Seasonal

The component of behaviour arising from seasonal influences and other systematic calendar related effects.

Series

See Time series.

Spectral analysis

Analysis of the variation of a time series in the frequency domain employing the spectral density function or the spectrum.

Spectral density function

A Fourier transformation of the autocorrelation function. The area under the spectral density function is equal to unity. It is a function, f, of frequency ω , denoted by $f(\omega)$, given by

$$f(\omega) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \rho(k) \cdot e^{-i\omega k},\ 0 < \omega < \pi$$

which, because cos(k) = cos(-k), is the same as

$$f(\omega) = \frac{1}{2\pi} [\rho(0) + 2 \sum_{k=1}^{\infty} \rho(k) \cos(\omega k)]$$

Spectrum

The Fourier transformation of the autocovariance function of a stationary series. Substituting the autocorrelation function in the transformation instead of the autocovariance function gives the spectral density function. The area under the spectrum is equal to the variance of the series.

Stationary

A series is said to be stationary if neither its mean nor its autocovariance depend on time. This is weak or second-order stationarity. Strong stationarity requires time independence of the entire distribution of the observations of the series, not just the first and second order moments.

Time domain analysis

The properties and characteristics a time series in relation with the timing of observations. Analysis of a time series in the time domain usually involves the autocovariance (or autocorrelation) function.

Time series

A set of observations made sequentially or chronologically in the time domain.

Trend

The longer term movement of a time series. For an economic time series the trend is usually a movement due to population and economic growth. In discussions of seasonal adjustment, trend is often used to mean trend and business cycle.

Trough

The minimum point attained within a repetition of a cycle.

BIBLIOGRAPHY

Balance of Payments, Australia (cat. no. 5301.0) January 1993.

Benda, O 1993, Report on the Investigation Into Smoothing Across Trend Breaks, internal paper, Australian Bureau of Statistics, Canberra.

Boehm, EA 1989, The Usefulness and Applications of Economic Indicator Analysis, paper presented to the Economic Society of Australia Annual Conference, 10 – 13 July.

Doherty, M 2001, The Surrogate Henderson Filter in X11, Australian & New Zealand Journal of Statistics, Vol. 43, No. 4.

Dagum, EB, Chhab N & Chiu, K 1996, Derivation and Properties of the X11ARIMA and Census X11 Linear Filters, Journal of Official Statistics, Vol. 12, No. 4, 329–347.

Findley, DF, Monsell, BC, Bell, WR, Otto, MC & Chen B 1998, New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program, Journal of Business and Economic Statistics, Vol. 16, No. 2, 127–177.

Gray, A & Thomson, P 1996a, Design of moving-average trend filters using fidelity, smoothness and minimum revisions criteria, Bureau of the Census, RR96/01

Gray, A & Thomson, P 1996b, Design of moving-average trend filters using fidelity and smoothness criteria, Vol. 2: Time Series Analysis in Memory of EJ Hannan ed. P Robinson and M Rosenblatt. Springer Lecture Notes in Statistics 115, 205–219.

Gomez, V & Maravall, A 2000, Automatic modelling methods for univariate series, Working Paper 9808, Research Department, Bank of Spain. (Appeared as Ch. 7 in Pena, D, Tioa, GC, & Tsay, RS (eds.) 2001, A Course in Time Series Analysis, New York, Wiley.

Henderson, R 1916, Note on Graduation by Adjusted Average, Trans. Amer. Soc. Actuaries, No. 17, 43–48.

Information Paper: A Guide to Smoothing Time Series — Estimates of 'Trend' (cat. no. 1316.0) 1987a.

Information Paper: Introduction of Concurrent Seasonal Adjustment into the Retail Trade Series (cat. no. 8514.0) 1993.

Information Paper: Time Series Decomposition — An Overview (cat. no. 1311.0) 1987b.

Jenkins, GM and Watts, DG 1968, Spectral Analysis and its applications, Holden-Day.

Koopmans, LH 1974, The Spectral Analysis of Time Series, Academic Press.

Macaulay, FR 1931, 'The Smoothing of Time Series', National Bureau of Economic Research.

Smarter Data Use, Australian Economic Indicators (cat. no. 1350.0) March 1992.

FOR MORE INFORMATION...

INTERNET	**www.abs.gov.au** the ABS web site is the best place to start for access to summary data from our latest publications, information about the ABS, advice about upcoming releases, our catalogue, and Australia Now—a statistical profile.
LIBRARY	A range of ABS publications is available from public and tertiary libraries Australia-wide. Contact your nearest library to determine whether it has the ABS statistics you require, or visit our web site for a list of libraries.
CPI INFOLINE	For current and historical Consumer Price Index data, call 1902 981 074 (call cost 77c per minute).
DIAL-A-STATISTIC	For the latest figures for National Accounts, Balance of Payments, Labour Force, Average Weekly Earnings, Estimated Resident Population and the Consumer Price Index call 1900 986 400 (call cost 77c per minute).

INFORMATION SERVICE

Data which have been published and can be provided within five minutes are free of charge. Our information consultants can also help you to access the full range of ABS information—ABS user-pays services can be tailored to your needs, time frame and budget. Publications may be purchased. Specialists are on hand to help you with analytical or methodological advice.

PHONE	**1300 135 070**
EMAIL	**client.services@abs.gov.au**
FAX	1300 135 211
POST	Client Services, ABS, GPO Box 796, Sydney 2001

WHY NOT SUBSCRIBE?

ABS subscription services provide regular, convenient and prompt deliveries of ABS publications and products as they are released. Email delivery of monthly and quarterly publications is available.

PHONE	1300 366 323
EMAIL	subscriptions@abs.gov.au
FAX	03 9615 7848
POST	Subscription Services, ABS, GPO Box 2796Y, Melbourne 3001

2134900001008
ISBN 0 642 54286 4

RRP $10.00

Produced by the Australian Bureau of Statistics