


FORECASTER

A simple, yet versatile approach to prediction



BRAINPOWER

Application through learning

COMMODORE 64
CASSETTE TAPES

FORECASTER

A simple, yet versatile approach to prediction



BRAINPOWER

Application through learning

FORECASTER

**A simple, yet versatile approach
to prediction**

**Text book by Stuart Armstrong
Programs by Kevin McKeogh and
Stuart Armstrong**



Triptych Publishing Limited

Triptych Publishing Limited, Sterling House, Station Road, GERRARDS CROSS, Bucks,
SL9 8EL.

©Triptych Publishing Limited 1984

First Published 1984

All rights reserved. No part of this publication or accompanying programs may be duplicated, copied, transmitted or otherwise reproduced by any means, electronic, mechanical, photocopying, recording or otherwise without the express written permission of Triptych Publishing Limited. This book and programs are sold subject to the condition that they shall not, by way of trade or otherwise, be lent, resold, hired out, or otherwise circulated without the publisher's prior written consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

BRAINPOWER is the trademark of Triptych Publishing Limited.

ISBN 1 85050 330 3

Design and production in association with Book Production Consultants, 47 Norfolk Street, Cambridge.

Typeset by Cambridge Photosetting Services, 19-21 Sturton Street, Cambridge

Printed by Burlington Press (Cambridge) Limited, Foxton, Cambridge.

Getting Started

Forecaster has been designed to cater for people with a wide range of backgrounds and skills. Many of you will therefore not need to read through this text from cover to cover in order to use the computer programs. To accelerate your progress we suggest the following:

- A. If you already understand the principles of the various forecasting techniques, simply turn the book around, and find Chapter 9, where you will be given detailed instructions on how to use the Applications Program.
- B. If you understand the concept of forecasting, but don't know how to apply it, then turn to Chapter 1, and follow the instructions on how to use the Teaching Program.
- C. If you are starting from scratch, if you don't have your computer handy, or if you simply want to take a more leisurely approach, then please read through the Introduction before you go any further.

Note

If you are not familiar with the procedures required to load the Teaching or Application Programs into your computer, refer to Appendix 2, where you will find specific instructions for using the different versions of **Forecaster**.

Contents

The book is laid out in a way which allows you to stand it upright in the box lid alongside your computer. For this reason, the pages are in an unusual sequences. Pages 1 to 29 start from the front of the book, whilst to use pages 30 to 58, just turn the book around and start from the back.

	Page
INTRODUCTION	4
1. THE TEACHING PROGRAM	7
2. FORECASTING METHODS	10
3. BASIC TIME SERIES TECHNIQUES	14
4. ADVANCED SMOOTHING - TRENDS	19
5. CLASSICAL DECOMPOSITION	25
6. SIMPLE REGRESSION	28
7. MULTIPLE REGRESSION	32
8. THE FORECASTING PROCESS	36
9. THE APPLICATIONS PROGRAM	38
Appendix 1. APPLICATIONS PROGRAM KEYWORDS	49
Appendix 2. LOADING PROCEDURES	53
Appendix 3. USING GRAPHS	54
GLOSSARY AND BIBLIOGRAPHY	56
INDEX	58

Introduction

Welcome

Titles in the **BRAINPOWER** series are uniquely designed to harness the power of your computer to enable you to learn new skills in a simpler and more enjoyable way. The sophisticated interactive approach ensures that you can work at your own pace and, once you have mastered the topic, the Applications Program will continue to serve your needs. We have made every effort to create a course which is straightforward to use, but if you think that we could improve upon it, please write to us on the card included in the pack.

Forecaster is a complete learning and applications course based upon a number of forecasting techniques. Your purchase consists of three elements;

- 1) The Text Book which you are now reading. Please bear in mind that you will be using it continuously in conjunction with the computer, and therefore it has been designed to stand upright in the box lid so that you can refer to it more easily.
- 2) The Teaching Program, which will be used to give you a full understanding of the techniques of forecasting.
- 3) The Applications Program, which you will be able to use to solve your own forecasting problems.

You will find that the Teaching Program is not a simple tutorial on how to use the Applications Program. Once you gain an understanding of the material, you will be able to use forecasting methods to solve problems with or without your computer.

If you think that you already have a sound grasp of the principles of forecasting, then you may wish to try out the Applications Program immediately. If so, simply turn the book around, and find Chapter 9. There you will discover the detailed instructions for solving your own forecasting problems.

Forecasting Techniques

The human race has a vast history of attempting to predict the future. Many concepts have been developed to rationalise the way in which these predictions are produced. Even in this micro electronic age, many businesses base their predictions of their own futures on irrational assumptions. **Forecaster** brings together a series of analytical techniques

which will enable you to make realistic predictions of those things which can be predicted. A reasoned prediction of the outcome of future events will give you greater control over your own future. Attempting to make a prediction will direct you towards an understanding of what other influences will modify the outcome. **Forecaster** will show you how to make a forecast, and it will show you how reliable the forecast is likely to be.

It is essential that you recognise that not all events can be forecast. Many are completely unpredictable; others are so complex that the mechanism to make the prediction is too massive to construct. Some prediction situations are better handled by techniques which take account of their unpredictability, and we suggest that you consider other titles in the **BRAINPOWER** range. In particular, both **Project Planner** and **Decision Maker** have a place in analysing certain specific events.

Many people subscribe to the theory that forecasting is an intuitive skill, yet scientific studies repeatedly demonstrate that consistently successful forecasters base their predictions on the types of analysis presented in these programs. Once you have completed the Teaching Program in **Forecaster**, you will be able to use the Applications Program to direct the power of your own computer to achieve the best possible estimates of future outcomes.

Applications for Forecasting

In the world of business, the single most frequent issue to which forecasting techniques are applied is the prediction of sales figures. The specific need for the prediction varies. It may simply be to project the revenue prospects for a business, so that its profitability can be analysed. It may be to decide how much to spend on advertising to achieve the desired sales. It could be to estimate how much stock will be required to satisfy customer demand, because holding too much stock ties up funds, and too little stock leads to lost sales. The prediction could be required for the next week, next month, or next year.

Apart from forecasting into the future, the techniques of multiple regression which are explained in Chapter 7, can be applied to determine the relationships between other factors. How your car's fuel consumption is related to the length of your journey and your speed of travel; how the crime rate relates to the level of unemployment, and the expenditure on the police force; how the length of time you have to wait for a taxi depends upon how heavily it is raining. If you think that there is a measurable relationship between two or more sets of figures, then collect the data, and put it to the test using the Applications Program.

Other applications for forecasting techniques are many and varied. If you have sufficient data, you can see whether there are ways of predicting

share price changes on the stock exchange; whether there is any link between the league position of a football team, and their probable match results next Saturday; or whether you can come up with a system for winning on the horses! We do not make any promises for your success in such ventures, but with **Forecaster**, you will be better equipped to attempt it than most people who have tried before!

CHAPTER 1

The Teaching Program

1.1 Teaching Method

Before we move into the stage of actually learning anything, we will quickly review how the computer is going to be used in conjunction with this book. First of all, you will find that all written explanations of the subject will appear in the book. We don't think that you want to strain your eyes reading computer screens full of text, and anyway computer memory is a relatively expensive medium for storing the written word. Because of this principle, you will be switching back and forth between book and screen all the time, so set up the book on its stand next to the computer where you can refer from one to the other easily. You will find it useful to have a pencil and paper and a pocket calculator handy as well. The screen will be used to show you examples in operation and to present exercises to check your own understanding.

As you work your way through the book, you will be asked to operate the computer by pressing certain keys. This is so that the computer knows which point you have reached. Any key you need to press will be highlighted in the text, such as **SPACE** or **4**. Likewise, when the computer wants you to return to the book, it will direct you to your place by giving you the number of the relevant chapter sub-heading.

1.2 Subject Sequence

The subject matter in the Teaching Program is related more by common purpose than by the development of a single concept. There are a number of threads which tie it together, and it is advantageous to work through the chapters in sequence. However, you will find that the subject falls into five distinct components:

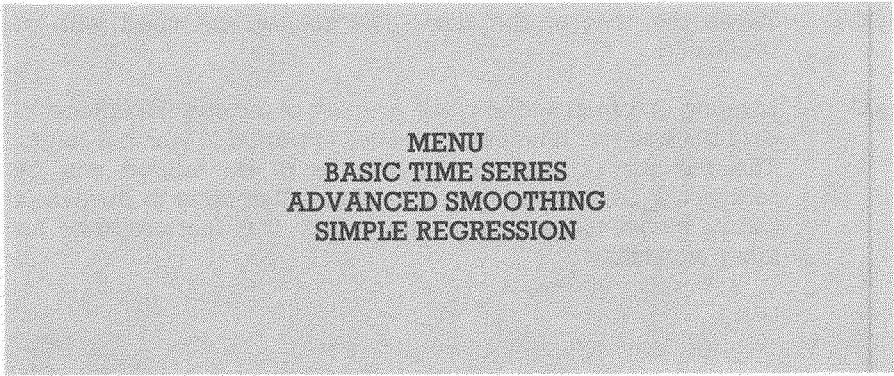
- 1) A general overview of forecasting techniques, which you will find in Chapter 2.
- 2) Time series forecasts based on weighted average techniques—moving averages and exponential smoothing. Refer to Chapters 3 and 4.
- 3) Classical Decomposition; Chapter 5.
- 4) Regression analysis methods and causal models, explained in Chapter 6 and 7.

5) Finally, Chapter 8 suggests ways of structuring a forecasting problem.

1.3 Getting Started

When you reach a chapter which uses the computer, instruct your computer to load in the Teaching Program. If you are a newcomer to computers, then you should refer to the loading instructions in Appendix 2. When the program starts, press the **SPACE** key, and it will display a list of options from which you can choose. The options relate to the chapter headings in the text book. You make your choice by using the **SPACE** and **ENTER** (or **RETURN**) keys. Each time you press **SPACE**, the bar will move one step down the list, and if it is at the bottom, it will jump to the top. When the bar is on the item you wish to select, press the **ENTER** (or **RETURN**) key, and the computer will act on your choice. This type of selection list will be referred to as a **menu** from now on.

You will notice that the options on the menu refer to Chapters 3, 4 and 6 of this Text Book.



MENU
BASIC TIME SERIES
ADVANCED SMOOTHING
SIMPLE REGRESSION

Once an option has been selected, the computer will have to load another section of the program. If you are using a cassette tape recorder, this can take quite a long time, particularly for the later topics, so please be patient.

When the correct section is loaded, the computer will give a message confirming the name of the unit and it will point you to the correct chapter in the book. Once any one unit is completed, the program will always give you the option of repeating the unit, stopping, or going on to the next unit. From time to time, instructions will be displayed on the screen which are not mentioned in the book. Always read and follow these instructions carefully.

Before you begin, remember to equip yourself with a pencil and paper, in order that you can make notes and sketches as you go. A calculator could also be useful at times. You should also be prepared to concentrate on a unit for quite a long period of time, for although we have made each step as simple as possible this is not a trivial subject to study. There will be plenty of opportunities to rework sections and ensure complete understanding.

CHAPTER 2

Forecasting Methods

2.1 Types of Forecasting

There are three ways of producing a forecast of the outcome of some future event, and perhaps an example will best explain the difference between them. Imagine that you are the owner of a record shop, and each month, you wish to estimate your sales volumes for the following month. How would you go about it?

Your first approach could be to look at your sales figures for all of the previous months to see whether there is any pattern in them. They could be increasing by say 2% per month, in which case, you could estimate that next month's sales will be 2% higher than those of this month. This is called a **time series** forecast; the forecast is based upon examining a series of figures collected over a period of time, and deciding whether there is any pattern in them.

Your second alternative is to attempt to work out why people buy your records. For instance, sales could appear to depend upon how much you spend on local newspaper advertising, and on the price you charge for the records. You could examine the relationship between these two elements and then calculate a mathematical formula to use to make the prediction. When you decide how much you will spend on advertising next month, and what your prices will be, just put the figures into the formula, and this will produce a forecast of your sales. This is called a **causal model**, because you have invented a mathematical model which is based on your ideas of what *causes* people to buy.

These two types of forecast are grouped under the heading **quantitative methods**, meaning that they involve a mathematical evaluation. A third alternative would be to review all of the records you stock and think about where each will be in the record charts next month; check what new records are going to be released and estimate how well they will perform; then combine your individual assessments into an overall sales forecast. This is an example of a **qualitative method**, using the term "qualitative" as the opposite of "quantitative"; i.e., not involving maths.

Each forecasting type has an appropriate role to play. A short term sales projection, such as the one we have been using as an example, is probably best handled using a time series. More complex causal models are better for longer term projections based on substantial historical data. A qualitative approach, on the other hand, is probably better suited to predicting the sales of individual records, rather than the figures for the entire shop.

2.2 Time Series

The principle on which time series forecasting is based is that all of the information required to make a forecast is contained in the previous values of the item to be forecast. That is, by analysing the way the values have changed in the past, taking nothing else into consideration, you can predict how they will change in the future. A time series is the name given to the series of measurements taken of a particular item at regular intervals over a period of time.

The actual process involves the following 5 steps:

- 1) Collect a time series of the values of the item to be forecast.
- 2) Develop a formula to predict future values from the time series.
- 3) Assess the accuracy of the prediction formula.
- 4) Modify the formula if required, and return to step (3).
- 5) Use the formula to make the prediction.

There are many different time series forecasting methods available to the forecaster, from the very simplest averaging techniques to complex computer based methods with exotic names, such as "The Box-Jenkins Univariate Autoregressive Moving Average Model". In this text, we will explain the application of a range of methods which will suit the needs of most forecasters. Chapters 3 to 5 are about these techniques, and they are arranged as follows:

CHAPTER 3	Averages Moving Averages Exponential Smoothing
CHAPTER 4	Double Moving Averages Holt's Method
CHAPTER 5	Classical Decomposition

The simple regression techniques explained in Chapter 6 can also be applied to time series data in some circumstances.

2.3 Causal Models

To develop a causal model, a great deal more data is needed than in the time series approach, and much more time and effort is required as well. There are only two more steps involved, but the additional data makes the

analysis and evaluation far more involved. The 7 steps are as follows:

- 1) Collect a time series of the values of the item to be forecast.
- 2) Decide which factors could be influencing the value of the item to be predicted.
- 3) Collect the data for those potentially-influencing factors.
- 4) Develop a formula to predict future values based upon the factors for which data has been collected.
- 5) Assess the accuracy of the prediction formula.
- 6) Modify the formula if required and return to step (4). It may be that there is evidence that additional factors could be involved, in which case, the process will have to be repeated from step (2).
- 7) Use the formula to make the prediction.

The development of causal models has only become accessible to many people through the availability of computers. This is because most effective models turn out to be relatively complex, and they tend to require extensive revision and recalculation. For this reason, the Teaching Program will demonstrate only the simpler forms of causal model. However, the Applications Program will enable you to develop far more sophisticated forecasting models. To enable you to use this facility effectively, the Teaching Program will also review the process of selecting and modifying the elements to be incorporated into your models.

You will find that Chapter 6 will explain how to use simple regression to make forecasts based on data on a single item. Chapter 7 extends into the application of computer methods to multiple regression for the development of complex causal models.

2.4 Qualitative Methods

Although many people feel that they are well qualified to make forecasts in their own particular field using a non-mathematical approach, extensive research indicates that this is not in fact the case. Quantitative methods are more reliable than a qualitative approach. However, this may well be because qualitative forecasters who have not received any formal training do not usually apply the same degree of care and structuring, even though these issues are equally important in qualitative and quantitative work.

Reliable qualitative techniques incorporate ways of structuring the

problem. This course does not intend to pursue these techniques, because it is not appropriate to use them in conjunction with micro computers. However, to give you an understanding of what they involve, the principal qualitative methods will be described:

Panel Consensus

This works on the basis that "two heads are better than one". You assemble a group of experts, and then allow them to debate the issue freely until they reach an agreement on the issue to be predicted. Unfortunately, results may suffer because of the way a group can sway the opinions of individuals within it; the prediction may reflect the views of the vocal, rather than those of the most expert.

Delphi

To overcome the disadvantages of the panel consensus approach, an alternative is to gather the group of experts, but keep them apart from each other. Each is interviewed in turn to discover his or her opinions, and a report produced which expresses their collective views. This is then distributed to each expert, who is then asked to revise his opinions in the light of the group's opinions. The process is repeated until a reasonable consensus is achieved. This tends to eliminate the group pressure problem, because the experts do not deal face-to-face with each other.

Historical Analogy

The principal involved here, is to find a previous situation similar to the one you are attempting to predict, and then make your prediction on what happened in the other case. For instance, to predict the sales of a new product, it could be reasonable to expect them to be similar to those achieved by an equivalent or similar product when it first went on the market in the past.

Market Research

Market research is the technique of systematically analysing all available information about a particular market, frequently with the intention of predicting sales volumes, but also to evaluate customers' attitudes, behaviour patterns and so on. It tends to be expensive and is usually only available to larger organisations.

CHAPTER 3

Basic Time Series Techniques

3.1 Time Series Forecasts

A **time series** is a list of values of some item or activity taken at a regular interval, or **period** over time. Time series forecasting involves examining the historical values of the series, and using that examination to predict how the series will behave in the future. A number of assumptions can be made about how the values of the series might change.

The first assumption is that the values will not change. This is a **stationary series**, and will be the subject of this chapter. The second assumption is that the series is changing at a constant rate. This is a series with a **trend**, and it will be dealt with in the next chapter. The third assumption is that the series will vary depending on the time of year, called **seasonality**. This is the subject of Chapter 5, which also deals with the situation of trend and seasonality being combined.

The items or values which we are measuring are called **variables**, because they tend to vary. The variable we are attempting to predict is called the **dependent variable**, whilst those we predict it with are called **independent variables**. In contrast, an item which always maintains a fixed value is called a **constant**.

A **stationary series** is one which, although it may have fluctuations from one period to another, remains constant, i.e., "stationary". You may think that these forms of series are simple to forecast, but unfortunately the fluctuations are frequently large enough to make the underlying values difficult to identify. All of the techniques presented in this chapter apply to stationary series.

For the purposes of demonstration in this chapter, we will be using the figures in the following situation:

Ray Jenkins runs a successful cash and carry wholesale electrical equipment shop with a wide product range. The business is only profitable because he maintains very careful control over his costs. One of his cost control objectives is to have the lowest quantity of each product in his warehouse, whilst holding sufficient stock of each to minimise the chance of running out. Keeping stocks low reduces the size of warehouse he needs, as well as reducing the size of his overdraft, both of which save him money. To help him keep control, he maintains a record of the monthly sales of each product, so that he can estimate the sales in future weeks. Two of the products, colour television sets and microwave ovens, have had the sales shown in the table on the following page over the past 20 months.

Month No.	TV Sets	Microwave Ovens
1	153	143
2	159	136
3	138	171
4	142	160
5	145	141
6	154	161
7	160	159
8	144	140
9	152	155
10	155	159
11	146	145
12	158	140
13	140	155
14	150	135
15	161	189
16	161	163
17	141	140
18	140	160
19	147	155
20	152	137

The best way of examining these sales figures is to plot them on a **scatter diagram**. A scatter diagram is a simple type of **graph** upon which each figure is marked as a point to represent its value and the week number to which it applies. The key difference between a scatter diagram and a normal line graph is that the points are not joined together. If you wish to make a further study of graphs, you will find some useful information in Appendix 3. If you now press **TS**, the computer will allow you to display scatter diagrams of both series of data.

3.2 Averages

The very simplest forecast of a time series variable is to say that the next value in the series will be the same as the one before it. For instance in the TV set case, month 21 sales could be forecast as being the same as month 20 i.e., 152 units. However, this does cause problems if the last month recorded is unusual in some way; perhaps, significantly higher or lower than average. This suggests using the average value of the series to predict the next value. An **average** is calculated by adding together all the values of the variable, and then dividing the sum by the number of variables involved.

The average of the TV set figures is thus found by adding them

together, and dividing by the number of values, 20. Press 1 to do this on the computer. The average value is 149.9, and we could use this as the forecast for month 21. Now work out the average for the microwave data given in the previous section. When you are ready, press 2 and enter your answer into the computer. How does your answer compare with using the month 20 figure as the forecast for month 21?

3.3 Moving Average

How much importance should be given to the sales in month 1 when we are attempting to forecast the sales in month 21? In the previous section, the averaging method gave equal importance to all of the previous sales figures. This is perfectly reasonable if we are confident that there are no long term changes occurring. If there are any changes, then they will not easily emerge from the average, because the inclusion of the very old values will tend to disguise them.

There is a very simple way of solving this problem – only take the average of the more recent values. For instance, we can average just the five most recent values to predict the next sales figure. In the next month, we drop the oldest of the five values, and include the new figure in the average to predict the following month. This is called a **moving average**, because the list of 5 values we use moves along the time series as we go. Because we have included 5 values, it is called a **five point moving average**. The process of using a moving average is called **smoothing** because it tends to smooth out the peaks and troughs in the data. If we take a step back in time to month 5 in Ray Jenkins' store, we can use the computer to perform five point moving average forecasts for each month up to month 21, and then we can see how the data is "smoothed" by the process – press 3 to watch this result.

To use moving averages, you must decide how many point to include. If you use a large number of points, you are getting very close to the simple average situation, and you will hide any long term changes which may be occurring. On the other hand, too few point starts to approximate the method of just taking the forecast as being the same as the current value. This places too much emphasis on short term anomalies, and does not have a smoothing effect. Typically, the number of points will be in the range 3 to 7. A 3 point moving average is only suitable for data which varies very little from one period to the next, whilst 5 and 7 point averages tend to smooth sufficiently without disguising longer term changes.

To test your understanding of this topic, press 4 and give your answers to a few problems related to the microwave oven data.

3.4 Exponential Smoothing

For organisations involved in many forecasting situations, a major problem can be the one of storing all of the data required to make the forecasts. The next method we will discuss helps to solve this problem, because all that must be recorded is the forecast itself. The forecast from the previous period is compared with the actual value for the current period, and the next forecast is based on these two figures alone. The process is called **exponential smoothing**. The exponential smoothing formula used is very simple:

$$\text{New Forecast} = (1 - A) \times \text{Old Forecast} + A \times \text{New Actual Value}$$

"A" is a constant number selected by the user in the range from 0 to 1. If you make A larger than 0.5, the effect is similar to using a moving average with only a couple of points. If A is made smaller than 0.2, this is similar to using a moving average with 5 or more points. The constant, A, is called the **smoothing constant**, because it is the constant value used in the smoothing equation.

We will now demonstrate this technique on Ray Jenkins' TV set data, starting from month 5 once again, and using an A value of 0.2. As you can see from the formula, you must have an old forecast before you can make a new one. Therefore, the very first forecast must be made in a different way. In this case, we will use the first 5 point moving average forecast from the previous section, 147.4. Now press 5 to watch the process.

The benefits of exponential smoothing are:

- 1) Only the previous forecast need be recorded. No other data is required.
- 2) The formula has the effect of giving more significance to the recent values than to older values. If you are interested in the technical reason for this, refer to the panel at the end of this chapter.

You can see just how simple this method is to apply, even though the name **exponential smoothing** sounds very sophisticated. The only judgement required is to choose the best value of A for your data. Chapter 8 will present the techniques which will help you to decide whether you have made the best choice.

Now you have the opportunity to try the technique yourself on the microwave oven data, once again. Press 6, and follow the instructions on the screen.

WEIGHTING IN EXPONENTIAL SMOOTHING

The smoothing formula has the effect of giving a progressively smaller weighting to earlier values. This can be demonstrated in algebraic form in the following way:

F_t = the forecast made in period t

V_t = the actual value which occurred in period t

Using these symbols, the exponential smoothing formula becomes:

$$F_t = (1-A) \times F_{t-1} + A \times V_t \quad - (1)$$

And thus the earlier forecasts can be expressed as:

$$F_{t-1} = (1-A) \times F_{t-2} + A \times V_{t-1} \quad - (2)$$

$$F_{t-2} = (1-A) \times F_{t-3} + A \times V_{t-2} \quad - (3)$$

Now formulae (2) and (3) can be substituted into formula (1):

$$F_t = (1-A)^3 \times F_{t-3} + A \times (1-A)^2 \times V_{t-2} + A \times (1-A) \times V_{t-1} + A \times V_t$$

This substitution can be extended to include all of the actual values. However, we can already see the weightings given to the last three values:

V_t is weighted (multiplied) by A

V_{t-1} is weighted by $A \times (1-A)$

V_{t-2} is weighted by $A \times (1-A)^2$

If A had the value of 0.2; these diminishing weightings would be 0.2, 0.16, 0.128, each one being 20% less than the previous value.

CHAPTER 4

Advanced Smoothing – Trends

4.1 Time Series with a Trend

The previous chapter was about **stationary** time series; that is, a series which does not tend to increase or decrease over time. In this chapter, we will consider data with a trend in it. A **trend** is a relatively uniform, continuous change in one direction. For instance, sales increasing steadily by 20 units per month, or decreasing steadily by say 10 units per month can be described as having a trend. If stationary forecasting methods are used on data with an increasing trend, they will tend to underestimate future figures. On the other hand, they will forecast high for series with a downward trend. If the forecasting techniques of this chapter are used on stationary data, however, they will correctly identify the absence of a trend. The only disadvantage which these methods have is that they require more data, and more judgement to make them work.

As in the previous chapter, we will start by presenting some sample data to which the techniques can be applied. Returning to Ray Jenkins store, we find the following sales volumes for stereo tape decks:

Month Number	Tape Decks	Month Number	Tape Decks
1	104	11	116
2	120	12	126
3	110	13	139
4	114	14	125
5	114	15	147
6	124	16	144
7	115	17	149
8	111	18	150
9	120	19	130
10	132	20	135

Press **AS**, and the computer will provide you with the facility to produce scatter diagrams of this data. A quick glance will reveal that there is an overall increase in sales through the period, although there are many short term dips. It is not always this simple to recognise that a set of data has a trend in it.

4.2 Double Moving Average

This method is performed in the way it is described; the moving average of the data is calculated twice. Firstly, a normal moving average is produced. Then a moving average is taken of the first moving average

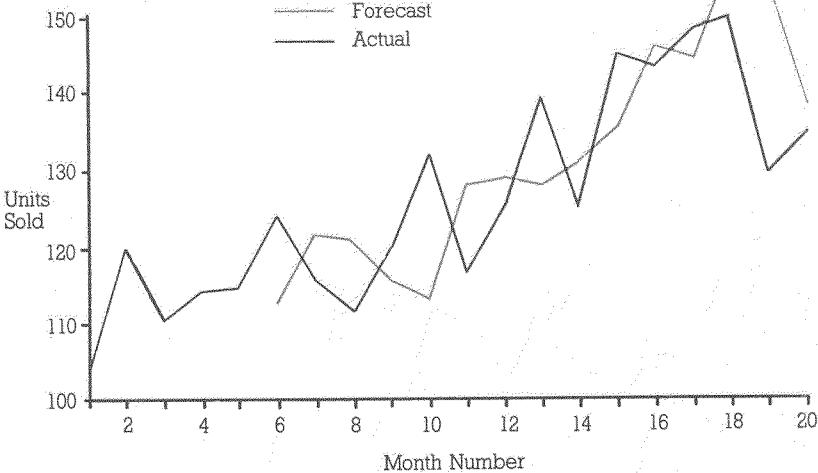
series. The difference between the two series of moving averages is a measure of the trend in the data. It is not particularly difficult to apply, but it is fairly slow to perform, and you must take care to make the correct subtraction. This is best shown by a demonstration. We will use the stereo tape deck data from the previous section and perform two 3 point moving averages. We will start by using just the first five values in the series – press 0 to display these figures. Now we will run through the following 6 steps, pressing the appropriate key to display each:

- 1) Take the average of the first three values, and enter it into a new column, opposite the middle value of the three. This alignment on the middle value is called **centering** the moving average. Press 1.
- 2) Repeat this process for data groups 2, 3, 4 and 3, 4, 5 then centre these in the second column. Press 2. You will notice that we started with five values, but there are only three values in the first moving average series.
- 3) Now take a three point average of the moving average series, and centre it in the same way. Press 3. We have now reduced the five original values to just 1. Note that this means you must have a minimum of 5 periods of data to perform a 3 point double moving average forecast.
- 4) Subtract the latest (and only) second moving average (opposite period 3) from the latest first moving average (opposite period 4), and this is the latest estimate of the trend. This is entered opposite the latest known actual value (period 5) in a new column. Press 4.
- 5) Now we are in a position to make a forecast for period 6. This is based upon the latest first moving average, adjusted for the trend. As we are forecasting a value 2 periods beyond the latest first moving average, we add twice the latest trend value. Press 5.
- 6) In the real world, we would have to wait for a month to discover the actual value for period 6, but all you need do is press 6, and the actual value will be displayed.

Well, you can see that our forecast is not too good, because of the sudden jump in sales. However, if we go on with a few more forecasts, you will see that this method does pick up the overall trend in the data, making it superior to using any of the methods demonstrated in the previous chapter. Press 7, and the computer will continue demonstrating this sequence of steps up to week 15. Follow the instructions on the screen,

which will show you how to stop when you are satisfied that you understand what is happening. To check your understanding, press 8, and you will be asked to perform the next few double moving average forecasts for yourself.

As you can see, all that you need to decide to use this method is the number of points to take in the moving average. The same considerations mentioned in single moving averages apply. A small number of points will react more quickly to changes in the data, whilst a larger number will tend to smooth the data more. In double moving average calculations, you do not risk losing the picture of the trend if you extend the number of points, rather you will tend to hide changes in the trend itself. The real problem is that this method needs a larger number of values before you can start forecasting.



MINIMUM NUMBER OF VALUES

To perform a single or double moving average, you will require a minimum number of readings before you can start. This table shows you how many readings you require:

	Number of Points in the Moving Average			
	3	5	7	9
Single Moving Average	3	5	7	9
Double Moving Average	5	9	13	17

4.3 Holt's Method

You will have noticed that double moving averages suffer even more than single moving averages from requiring the maintenance of extensive historical records of the data. This data requirement can be simplified by applying the same approach with exponential smoothing to produce **double exponential smoothing**. However we are going to go straight onto a development of this, known as **Holt's method**.

In this technique, the simple exponential smoothing formula is expanded into two parts; a smoothed form of the data, and a smoothed form of the trend. These are combined to produce the forecast. You will remember the smoothing constant, A, used in simple exponential smoothing. This is still required, but a second constant, B, is also used to smooth the trend. The formulae required are as follows:

$$\text{Smoothed Value} = (1-A) \times \text{Old Forecast} + A \times \text{New Actual Value}$$

$$\text{Change in Smoothed Value} = \text{Smoothed Value} - \text{Old Smoothed Value}$$

$$\text{Smoothed Trend} = (1-B) \times \text{Old Smoothed Trend} + B \times \text{Change in Smoothed Value}$$

$$\text{New Forecast} = \text{Smoothed Value} + \text{Smoothed Trend}$$

As in exponential smoothing, the formulae require some initial estimates. These include an initial forecast, an old smoothed trend, and an old smoothed value. We will demonstrate the process on the stereo tape deck data, using smoothing constants of $A=0.3$ and $B=0.6$. We will also set the first forecast and smoothed value as equal to the first actual value, and the initial trend as zero. This is a typical way of setting the first values for the formula. Press 0 to see this starting data. The actual value for month 2 is also included. We will now go through the 5 steps required to perform a forecast:

- 1) Calculate the new smoothed value by multiplying the old forecast by $(1-A)$; that is $1 - 0.3 = 0.7$; then add A times the new actual value. Press 1 to do this.
- 2) Calculate the change in the smoothed value; press 2.
- 3) Calculate the smoothed trend by multiplying the old smoothed trend by $(1-B)$ and adding B times the change in the smoothed value. Remember that $B = 0.6$ and thus $(1-B) = 0.4$; press 3.
- 4) The forecast is the sum of the results of steps 1 and 3; Press 4.

- 5) Now we can discover the actual value by pressing 5, and again we see that the forecast is pretty good.

It would not have been surprising if this forecast had not been very good, because the formula develops as more and more data points are added. As in the exponential smoothing situation, using the old forecasts as part of the formula takes into account all of the earlier data. Therefore, the longer Holt's method is applied, the better it becomes at determining the long term trend.

Press 6, and the computer will continue demonstrating this sequence of steps up to week 15. Follow the instructions on the screen, which will show you how to stop when you are satisfied that you understand Holt's method. To confirm your understanding, press 7, and you will be asked to perform some of these forecasts for yourself.

For serious trend forecasting, Holt's method offers similar benefits to exponential smoothing. The formulae have the effect of giving more weight to recent data to a greater extent than the old data, and the forecaster need only keep a record of the previous period data. Older data can be discarded.

4.4 Other Methods

Another useful method for analysing data with a trend in it is to use simple regression, discussed in Chapter 6. This has the advantage of producing a simple formula for projecting the trend into the future, but it is quite slow to apply, and the formula must be updated regularly. The significant benefit is that it can handle trends other than simple linear ones.

4.5 Reliability

How can you tell whether the forecasting technique you have chosen is sufficiently reliable for your needs? Of course, the easiest way is to make a forecast, and then see whether it is fulfilled. Unfortunately, you may not have the time to wait, and you may need to depend upon the reliability of your forecast. A simple alternative is to develop the technique without using the last couple of known values of the time series. You can then use your chosen method to forecast these known values. If the results are close to the actual values, you can have some confidence in your forecasts.

The other test is to apply some statistical measures to test how closely the model comes to the actual data you have used. The steps involved are as follows:

- 1) Use the method chosen to predict all of the past values of the time series, in just the way we have done in the examples.

- 2) Calculate the difference between each forecast and each actual value, and then ignore the sign to make them all positive numbers.
- 3) Add up all of these differences and calculate the average value of these differences.
- 4) Calculate the average value of the time series, and express the average found in step (3) as a percentage of this. This result is called the **mean absolute percentage error (MAPE)**.
- 5) As an alternative, take all of the differences found in step (2), square them, and find their average. This is called the **mean square error (MSE)**.

If you try a number of different techniques, then measure their performance in the ways mentioned here, you can then select the one with the best results. Both MAPE and MSE statistics are produced by the Applications Program.

CHAPTER 5

Classical Decomposition

5.1 Seasonality

So far, we have considered data which is **stationary** or which has a **trend**. A series which is changing, increasing or decreasing, over time is said to have a trend. We have also noted that the underlying trend can be disguised by random fluctuations around the trend. However, sometimes these apparently random fluctuations are not as random as they appear to be. This is because of another effect known as **seasonality**.

If we continue with examples based upon sales figures, it is simple to think of data which has a seasonal element. An extreme illustration can be found in the sales of Christmas cards. If monthly sales of Christmas cards in a newsagents shop were recorded from September to December, they could look like this:

September	600
October	700
November	1100
December	2000

If you applied one of the methods explained in the previous chapter to this data, you would discover a very high growth rate, and you could forecast January sales in excess of 2500. But we know that sales of Christmas cards in January will be very much lower than that. This is because of seasonality. There is a season for buying Christmas cards.

Other products experience seasonal effects, although probably not as extreme as for Christmas cards. Woolly vests, for instance, could be expected to sell in greater volumes in winter months. The seasonality of some items may be more variable. Paperback books sell well at Christmas, and also during the summer holiday period.

5.2 The Concept of Decomposition

Classical decomposition is a method of forecasting which can take into consideration both trend and seasonality, but it involves using vast quantities of data, and requires a great deal of time-consuming calculation. This is the type of thing that computers can perform very well for us, and so we will only explain the basic principals here. We can leave the actual calculations to the Applications Program.

The concept itself is very simple. It assumes that each value of the variable, which we will call "V", is made up of three components, and thus the values can be "decomposed" into these parts; hence the name. The three parts are:

A trend factor	T
A seasonality factor	S
A random, irregular factor, which cannot be predicted	I

and thus:

$$V = T + S + I$$

A series of calculations is performed to "decompose" the data into these components, and then a forecast can be produced by recombining the trend and seasonality factors, depending where in the season the period to be forecast actually falls. The random factor, "I", is ignored in the forecast, as it cannot be predicted. We will use the letter "F" to represent the forecast value, and a subscript "t" on any of the letters means "the value of that variable at time t". Thus the forecast for period "t" can be represented as:

$$F_t = T_t + S_t$$

Some models include a fourth element, called **cyclicality**. This involves the recognition that some data incorporates a longer term cycle which moves up and down over a period of years. There is some suggestion of such a cycle in the national economy which can be related to increased government expenditure around the time of general elections. Unfortunately, these cycles are very difficult to identify, and many years of historical data are required to perform this type of analysis. Therefore, we will not take cyclicality into consideration.

5.3 Using Decomposition

If you have a forecasting problem which you suspect involves a seasonal factor, then you should use the decomposition model in the Applications Program to analyse the data. Most businesses, large and small, will find that there is some seasonality in their sales, but it is essential that you have sufficient data to identify the seasonal pattern. This usually means collecting the figures for at least three full years. The data need not be in monthly form. It could be quarterly, or half yearly, for instance. A computer with larger memory capacity would even allow you to review seasonality on a weekly basis, but the Applications Program is restricted to 48 periods.

Decomposition can be used on variables which do not have any seasonal effects. However, the vast quantity of data expected by the model makes it more efficient to use the simpler techniques, such as exponential smoothing, or Holt's method.

If the data you wish to examine involves significant seasonal fluctuations, but you do not have a long enough history of the figures to use decomposition, you could consider attempting to build a causal model. The techniques involved are explained in Chapter 7, but you will also require data on all the factors which you think may be influencing the time series.

CHAPTER 6

Simple Regression

6.1 The Straight Line

If you press **SR** on the computer keyboard, the screen will display a scatter graph of the first six months of another of Ray Jenkins' products. This time, it is a new line of dishwashers, which is increasing quite quickly. The actual figures are as follows;

Month No.	Dishwasher Sales
1	3
2	8
3	13
4	15
5	30
6	35

Press 1, and the points will be joined together to form a line graph. Now you can see that the line is almost straight, apart from the kink downwards in month 4. We could go back to the scatter graph, and try to "fit" a single straight line which passes as close as possible to all the points. Press 2, and the computer will do this. Most of the points are quite close to this line. We could now use it to forecast beyond month 6. Press 3 to see how this forecast is made. The new line has a value of about 40 at month 7, so this would be our forecast for that month.

This is the basic mechanism of **linear regression**. We try to make a straight line fit the data we have, and then we project the line into the future to make a forecast. The simplest way of doing this is just as we have demonstrated; draw a scatter graph and then fit a straight line as close as possible to the data. However, the "best fit" is not always simple to pick out by eye, so mathematical methods have been developed to perform this task for us.

6.2 Preliminaries

Before examining the method used to fit a line, we need to review the way in which a straight line is defined in mathematical terms. But even before we do that, there are a couple of items to explain. First of all, take a look at the axes of the graph drawn on the computer screen. If you have any problems understanding graphs, you can turn to the additional notes in Appendix 3. The axis along the bottom of the screen is called, by tradition, the **x-axis**, and any values marked off along the x-axis are called the x values. Similarly, the vertical axis is called the **y-axis**, and the values

associated with it are the y values. In our example, the x values are the month numbers, and the y values are the sales volumes. By convention, the dependent variable, is always assigned to the y-axis.

Now press 4 to eliminate the points so that we can concentrate on the line itself. How can you describe a line in mathematical terms? There are a number of ways, but we will consider just one approach. Any straight line can be described by a formula of the form:

$$y = a + b \times x$$

where "a" and "b" are two numbers which are unique to any line. "a" is the point where the line meets the y-axis, and "b" is a measure of how steeply the line slopes. Our line does not meet the y-axis at present, but by pressing 5, the computer will extend it until it does. Notice that it actually meets the y-axis below the x-axis. This is fine; it just means that the y value at this point will be negative and thus "a" will also be negative in this case. You can see that if we can find "a" and "b" for the line which fits our data, we have a formula for making forecasts.

6.3 Fitting a Line

So how do we go about fitting a line to the data? Press 6 to add the points back to the graph. One way of measuring how well the line fits the data is to measure the distance between the data points and the line, and add all of the distances together. Press 7 to highlight these distances, which are called the **residuals**. The best fitting line is the one which reduces these residuals to a minimum. In fact, experience has shown that better results are achieved by minimising the sum of the square of the residuals. For this reason, the procedure of finding a line which fits as closely as possible to the points in the scatter graph by minimising the squares in this way is called the **least squares method**. Press 0 to clear the screen, and we will work through the least squares method on the data we plotted previously. It works like this:

- 1) Put the x and y values into two columns; press 1.
- 2) Calculate the average value of x, and the average value of y; press 2 to do this.
- 3) Subtract the average value of x from each value of x and put this in a new column; press 3.
- 4) Perform the same process for y; press 4.

- 5) Form a new column by multiplying the values of the last two columns together, and add up the total value of the new column; press 5;
- 6) Form yet another column by squaring the values in the column produced in step (3), then add up the total value of the new column; press 6.
- 7) Divide the total of the step (5) column by the total of the step (6) column. This is the value of "b" for the line of best fit; press 7.
- 8) Multiply the average value of x by "b", and subtract the result from the average value of y. This will give you the value of "a"; press 8.

Now you have the formula for the best fitting line for our data:

$$a = -5.47 \text{ and } b = 6.5$$

Thus the formula becomes:

$$y = -5.47 + 6.5 \times x$$

You can now use the formula to make a prediction for any future period by setting x to the period number. For instance, to forecast month 7, replace the x with a 7:

$$y = -5.47 + 6.5 \times 7 = 40.03 \qquad \text{Forecast for month 7 is } 40.03$$

This compares closely with our guess of 40 from the first "fit" of the line in section 6.1.

6.4 Non Linear Models

Data which falls roughly in a straight line is referred to as being **linear**. If the data does not form a straight line, it is called **non linear**. The method we have described so far only works for data which is linear. However, sometimes points which lie in some form of curve can be "bent" into a straight line by modifying the x or y values or even both. For instance, taking the logarithm of the y values and plotting them again could produce a straight line. In this case, the formula would forecast the logarithm of y, and the forecaster would have to remember to change the forecast back to its original form. Modifying the data in this way is called a **transformation**.

As an example, we will look at how a transformation could be applied.

The price of a Microfruit plc share for the past four weeks have been as follows:

Week No.	Price in £
1	2.00
2	2.50
3	3.33
4	5.00

These points do not fall in a straight line. Press 0 to see them. However, if we divide each of them into the number 1, they are transformed. The result of dividing a number into 1, is called the reciprocal of that number. The transformed data looks like this:

Week Number	Original Value	Transformation		New Value
1	2.00	$\frac{1}{2.00}$	=	0.5
2	2.50	$\frac{1}{2.50}$	=	0.4
3	3.33	$\frac{1}{3.33}$	=	0.3
4	5.00	$\frac{1}{5.00}$	=	0.2

If you now look at the plot by pressing 1, you will see that they form a straight line. It does not matter that the line slopes downwards, instead of up as it did before. The important point is that the line is straight. If we project the line, we will find that the transformed forecast for week 5 is the number 0.1. To transform it back, we just need to know that the number which would turn into 0.1 if transformed in the same way as the others is 10, because:

$$\frac{1}{10} = 0.1$$

Thus, our week 5 forecast for the Microfruit share price is £10.00!

The Applications Program will permit you to perform a range of transformations, and will show the "before and after" scatter graphs. It will also perform some transformations automatically if you wish.

CHAPTER 7

Multiple Regression

7.1 More Variables

The simple regression techniques explained in the last chapter are fine if there is only one independent variable, but there are many situations where the dependent variable is influenced by more than a single variable. If more than one other variable is involved, it is impossible to draw a graph of the relationship, because each variable will need its own axis, and there can only be two axes on a flat page! However, the formula for a line can be extended to cope with more variables quite simply. For instance, if there are 4 independent variables, called u , v , w and x , the equation becomes:

$$y = a + b \times u + c \times v + d \times w + e \times x$$

The maths becomes very complex, and so solving this type of **multiple regression** is usually left to computers. The Applications Program will handle problems with up to 9 independent variables. You can even complicate the situation further by making different transformations to each of the variables to make a non-linear model into a linear model.

7.2 Causal Models

Designing a formula to make a forecast is called **modelling**, and designing a formula based upon what you think causes the changes to the forecast is called **causal modelling**. The formula you produce is thus called a **causal model**.

To develop a causal model takes common sense, and a sound understanding of the variable you are attempting to predict. What you must do is select those things which you could reasonably expect to influence the forecast. You should collect the data for those items, and run the data through the Applications Program to develop and test the model. The variables which have an insignificant effect on the outcome can be rejected, and if there are any large unexplained residuals, you should search for a new variable to explain the residuals and improve the performance of the model.

The types of causal model you will build will depend very much on your own requirements, but perhaps an example could help to give you some ideas:

Hans Fletcher is the business development manager for a chain of record stores. One of his tasks is to investigate new store sites, and he wants a model to predict sales volume for a site. He has noticed that there appears to be a link between sales, the size of

the store, the number of other shops in the street, and the distance to the nearest bus stop. Therefore, he builds a model which looks like this:

$$S = 37 \times F + 260 / D + 16 \times N$$

Where S = Monthly sales in thousands of pounds
 F = Shop floor area in thousands of square feet
 D = Distance to bus stop in yards
 N = Number of other shops in the street

This first model appears to work quite well, but the prediction for one new store turns out to be very wrong. Sales are well below the forecast. The particular store is unusual, because it is very long and narrow. He thinks that this could mean that the link could be to the size of the shopfront, rather than to the size of the store itself, so he modifies the model, and the quality of the predictions improves. The new model turned out to be:

$$S = 2 \times F + 245 / D + 17 \times N$$

where the new variable is F , the width of the shopfront in feet.

As already mentioned, the key requirements are common sense, and a sound knowledge of the variable you are attempting to model. There are no hard and fast rules to follow, but the checklist in the next section advises you on some points to bear in mind.

7.3 Modelling Checklist

- 1) There is no point in building a model which depends on data which you do not have. If you develop a model to forecast the price of grapes next month, based upon the number of grapes brought to the market, it is of no use unless you know how many grapes there are going to be.
- 2) Be certain that there is a real link between the dependent variable, and any independent variables you choose. For instance, the price of beer and the salaries of nurses may both tend to rise at the same rate, because both are related to the performance of the economy. But constructing a model to predict beer prices using nurses' salaries is doomed to failure.

- 3) Take care that you do not use two variables which effectively contain the same information. For instance, Hans removed the store size parameter from his model when he introduced the shopfront size, because the two tend to be related to each other. If you use the **STEPWISE** option in the Applications Program, the computer will tend to identify this problem, and reject one of the variables.
- 4) Check the residuals (the amount by which the formula misses the true values) to be sure that there is no pattern in them. A pattern could mean that there is an important variable missing.
- 5) In using a causal model to predict a time series, the time series itself is often a dependent variable. For instance, if a model was being built to forecast soap powder sales based upon the money spent on TV advertising, an important factor in one month's sales would be the sales of the previous month. The way to incorporate this into the model is to create an independent variable of the dependent variable figures, moved forward by one month. This is a special form of transformation called **lagging**, which is available in the Applications Program.

7.4 Testing the Model

Once you have developed a model, you will want to test it to decide how reliable it is. The simplest test is the one mentioned in section 4.5; you can make a prediction, and then wait to see how close it lies to the actual outcome. Unfortunately, the time delay would be unacceptable. The effect of this test can, however be simulated by constructing the model without using the most recent actual values, and then using it to forecast them. The performance can then be judged immediately.

Regression analysis also produces a series of statistical tests which predict the reliability of the model. These tend to have abstract names which mean little in themselves to the amateur forecaster. You will find that the most useful are:

MSE This is the same as described in section 4.5. It is calculated by finding the average of the squares of the differences between forecast and actual values. It is helpful for comparing one model with another. The model with the lower MSE is likely to be the better predictor.

R-squared This is a measure of how adequately the model fits the data. An R-squared of 1 is perfect, and R-squared values down to .5 or even less can be satisfactory, but the lower it gets, the

worse the forecasts made with the model are likely to be. If you are simply building a model to demonstrate that there is some link between two or more sets of values, then R-squared as low as 0.15 or 0.2 will suggest that the link could exist.

F-test

This is a test to decide whether the model is a better predictor than no model at all. The interpretation of the result of the F-test is not a simple matter of finding the value. It involves some further analysis. If you are not a statistician, you can accept the results of the Applications Program, which will warn you if the model fails the F-test.

CHAPTER 8

The Forecasting Process

8.1 Need for Structure

It would be very easy to jump into a forecasting problem without thinking too much about the situation into which the forecast is to fit. After all, actually producing the forecast is probably only a small part of work involved. It is essential that you develop a structure, and then relate the way you develop the forecast to the way it is to be used.

8.2 The Process

You can structure the problem by asking yourself the following five questions:

What is the forecast needed for?

Why do you want the forecast in the first place? If you know this, it will tell you what type of forecast you want, how far forward it needs to be, how often you want it, and how much accuracy you require.

What could influence the forecast?

If you were to build a causal model, what factors do you think would influence the item you are attempting to predict? In what way would they effect it?

What data is available?

Now that you have decided which factors could influence your forecast, you will need to look for data on them. How recent and how accurate is it?

What technique is appropriate?

The technique you choose will depend on your answers to the first three questions. As a guide, we suggest that the methods listed in the table on the next page can be considered in the circumstances described.

How should it be monitored?

Now that you have decided how to make the forecast, it is essential that you establish a formal procedure to review the value of the forecasts. Do they satisfy the needs expressed in answer to the first question? Can they be improved upon? Would another method be better?

TABLE OF FORECASTING TECHNIQUES

Length of Projections Required	Nature of Time Series	Historical Data on Dependent Variable	Historical Data on Influencing Factors
Short Term	Stationary Trended Seasonal Not Apparent	Moving Average Holt's Method Decomposition Simple Regression	Moving Average Holt's Method Decomposition Causal Model
Medium Term	Stationary Trended Seasonal Not Apparent	Expo. Smoothing Holt's Method Decomposition Qualitative	Causal Model or Qualitative Method
Long Term	Stationary Trended Seasonal Not Apparent	Simple Average Qualitative Decomposition Qualitative	Causal Model or Qualitative Method

The methods marked in **bold type** are available in the **Forecaster Applications Program**.

CHAPTER 9

The Applications Program

9.1 Using the Program

The **Forecaster** Applications Program has been designed to analyse up to 7 variables, with a maximum of 48 readings for each variable. If you are experienced with forecasting techniques, you will find that the program can be asked to perform an analysis to your precise requirements. If, on the other hand, you are not confident about the way to handle a problem, the program can make its own judgements about the data you provide, and help you towards your goal of developing a useful forecast.

First, you must gather the data which you intend to analyse, but before you do that, perhaps you should look back to Chapter 8 to review how you intend to handle the problem. When everything is prepared, load the Applications Program in accordance with the instructions for your type of computer. If you are uncertain how to go about this, please refer to the suggestions in Appendix 2. When the program has been loaded, the menu of options will be displayed.

9.2 The Main Menu

The main menu will be displayed on the screen as follows:

```
MAIN MENU
LOAD DATA
SAVE DATA
ENTER DATA
EDIT DATA
TRANSFORM DATA
TIME SERIES
CAUSAL MODEL
FORECAST
```

You will see that the **MENU** heading is highlighted by a bar on the screen. If you press the **SPACE** key, the bar will move down to the next item. If you continue to press **SPACE**, the bar will move down the list, and then it will return to the top item below the word **MENU**. To choose any item on the menu, simply use **SPACE** to move the bar onto that item, then press **ENTER** (or **RETURN**), and it will be selected. Each item on the menu works as follows:

- LOAD DATA** Loads in a data file from tape or disk, depending which version of the program you have purchased. This will only be of use to you once you have entered some data, and used the next option to save it.
- SAVE DATA** Saves a data file to tape or disk. You will find a full explanation of these first two options in section 9.4.
- ENTER DATA** Allows you to enter a new set of data into the computer. This is the option you will require to enter your first forecasting problem. A full explanation is given in section 9.3.
- EDIT DATA** Enables you to make changes, additions, deletions and print out the data you have already entered. Refer to section 9.5 for the full details of the editing options available.
- TRANSFORM DATA** Allows you to alter any of the variables by using a series of functions to transform the data. An explanation of the use of transformations is provided in section 9.6.
- TIME SERIES** Performs a time series analysis on the dependent variable using Holt's method, or Classical Decomposition. Please refer to section 9.7, for an explanation.
- CAUSAL MODEL** Performs a regression analysis. The full details of this facility are provided in section 9.8.
- FORECAST** Produces a forecast based upon the time series or regression model created by one of the two previous commands. Section 9.11 will show you how this is done.

We recommend that the first time user should select the **ENTER DATA** option, and move on to the next section to see how to continue.

9.3 Entering New Data

Entering your data into the computer is a simple process. When you select **ENTER DATA** on the main menu, the screen will clear, and you will be

asked by the computer to provide the name for your dependent variable. Any name up to 10 characters long will be accepted. You will notice that the request for the name will appear in a box at the bottom of the screen. In the data entry routine, this box will always tell you what you can or must do next. When you have typed in the dependent variable name, press **ENTER** (or **RETURN**), and it will be printed at the top of the screen. The numbers 1 to 10 will appear in a column, and a white box will appear adjacent to the first number. This is where the first value of the dependent variable will appear when you type it in. After you type in each value, press **ENTER** (or **RETURN**), and go on with the next. If you make a mistake, you cannot correct it immediately, but do not be concerned. Finish typing in all the data, and then use the **EDIT DATA** option when you get back to the main menu.

When you have typed in all of the values for the dependent variable (there can be up to 48 altogether), simply press **ENTER** (or **RETURN**) twice. The computer will then ask for a name for the next variable in the same way as before. The only difference from this point is that you must type in the same number of readings for each variable as you did for the dependent variable. When you have entered all of the variables (up to a maximum of 7), just press **ENTER** (or **RETURN**) when asked for another variable name, and you will be returned to the main menu.

9.4 File Handling

Once you have keyed in all of your data, it would be wise to save it on tape or disk before continuing. Simply select **SAVE DATA** on the main menu to do this. The program will ask you for a name for the data file, and will then start the save operation. The exact details of this operation will vary from one computer type to another, so follow the instructions on the screen very carefully. When the data file has been saved, keep a record of the name used so that you can recall it when required.

When you go on from this point, you may make changes to the data, in which case you may wish to make another copy of the file. If you save the data after a forecasting model has been produced, the file will also contain details of the model.

Loading the data back into the computer is simply the reverse procedure. Select the appropriate menu option, and follow the instructions on the screen. Note that disk versions of this program will only produce files on disk, and tape versions only produce tape files.

9.5 Editing the Data

When all the data has been entered, you can review and alter it by selecting the **EDIT** option. You can change individual entries, and add or

delete rows and columns. In addition, you use **EDIT** to select the variables to include in the forecasting model, and which is to be the dependent variable. This enables you to enter all of the data and then just work with part of it. The specific commands, listed below, are selected by pressing the key representing the first letter of the word you want. For example, if you wish to **ERASE**, press the **E** key:

INSERT To insert a new line of data below the current cursor line, or a new column to the left of the current cursor column. On selection, you will be asked whether you want a **ROW** or **COLUMN**. Type **R** or **C** to select.

ERASE This erases the current cursor row or column. Use this command with care, as the data in an erased line cannot be recovered. On selection, you will be asked whether you want a **ROW** or **COLUMN** erased.

DEPENDENT Makes the current cursor column variable the dependent variable. There can only be one dependent variable, so this will automatically de-select the previous dependent. The current dependent variable is the one with all values shown in inverse (dark on a light background).

SELECT Selects the current cursor column variable or current cursor row for inclusion in the time series or causal analysis, or if it has already been selected, it de-selects it. As many or as few variables and readings as you wish can be included. The currently selected variables are identified by having the name at the top of the column in inverse, and the currently selected readings are likewise shown in inverse in the left hand column. Press **R** or **C** to indicate **ROW** or **COLUMN** when requested.

ALTER Allows you to alter the value at the current cursor position. You will be requested to enter a new value at the top of the screen.

TIME Inserts a "time" variable in the first column. This is required if a regression (causal) analysis of a time series is to be attempted. The variable takes the values 1 to n, where n is the number of readings. Refer to sections 9.8 and 9.9 to see how this facility is used.

PRINT Prints a copy of all the data you have entered, as well as showing any transformations you have applied.

MENU Returns you to the main menu (see section A1.1).

Once you have completed your manipulation of the data, select **MENU** by pressing **M** to return to the main menu. If you have made any important changes to the data, it may be appropriate to save them onto tape or disk once again before you continue.

9.6 Transformations

A **transformation** is the modification of all the readings of one variable in a consistent way. For instance, multiplying all of the readings of a variable by 2 would be a transformation. Transformations are often used in the production of more complex causal models, and the way they are used is described in Chapters 6 and 7. An infinite range of transformations are possible, but space limitations mean that only a restricted number are provided in **Forecaster**. When you select this option on the main menu, you will be presented with a list of all of the variables you have entered. You can select one from the list using **SPACE** and **ENTER** (or **RETURN**) in the usual way, or you can select **RETURN** to go back to the main menu. Once a variable has been selected, you will be shown a scatter diagram of that variable plotted against the dependent variable. If you select the dependent variable, it will be plotted against time as the independent variable, assuming that the readings are in time order.

If you decide to perform a transformation, you have the choice of replacing the original data with the transformed data, or entering the transformation as a new variable. In this latter case, you will be asked to provide the new variable with a name. If you ask for the transformation to replace the original data, the type of transformation performed will be shown against the variable name when it is displayed on the edit or transformation menu screens. Remember that you can always recover the original values by selecting the **NORMAL** option on the transformation menu.

Only one transformation can be performed on a variable if the transformation has replaced the original variable. If you attempt a second transformation, this will replace the first. Therefore, to perform multiple transformations on a single variable, you must create a new variable for each stage of the transformation. The way each transformation operates is explained in the list which follows on the next page. The abbreviation " V_t " means "the value of the variable in time period 't'."

- RETURN** Takes you back to the main menu.
- NORMAL** Returns the data to its original form, but only if a previous transformation has superceded the original data.
- LOG** Takes the natural logarithms of the data; New $V_t = \ln(\text{Old } V_t)$
- EXP** Takes the exponential values of the data; New $V_t = e^{\text{Old } V_t}$
- POWER** Raises the data to a power. The power to be used will be requested. Using a power of minus one will produce the inverse values of the data; New $V_t = \text{Old } V_t^n$
- ROOT** Takes the root of the data. The root to be used will be requested; New $V_t = \sqrt[n]{\text{Old } V_t}$
- LAG** Lagging of a time series is the process of moving the column of data back so that it no longer relates to period to which it refers. The number of periods the data is to be lagged will be requested. Lagging is only used when attempting to develop a causal model to predict a time series. It is not appropriate in any other situation. The lag can be by any number of periods, from 1 upwards. The transformation for a lag of n periods is;
New $V_t = \text{Old } V_{t+n}$

9.7 Time Series Analysis

Select **TIME SERIES** on the main menu to use this function. The program will perform one of two types of analysis on the dependent variable alone, depending upon how many readings you have provided. When you select this function, you will be asked:

Do you suspect a seasonal factor? (Y/N)

If you answer **N** to this, the program will perform a Holt's method analysis. If you answer **Y**, you will be asked:

No. of periods in the seasonal pattern

If you are using monthly periods, and you suspect annual seasonality, the answer is **12**. Similarly, for quarterly data, you would answer **4**, as there are 4 quarters in one year. You may be studying daily data in which you think there is seasonality from day to day. In this case, answer **5**, **6**, or **7**, depending on the number of days in your working week.

When you have answered, the program will check the total number of readings you have provided. If there are enough to cover at least three times the number of periods in a season, it will perform Classical Decomposition. If there are less than the required number, it will warn that there is insufficient data, and will default to Holt's method.

To start on Holt's method, the program will ask for initial values for the smoothing constants, A and B. If you have a reasonable idea of what they may be, enter the values. If in doubt, we suggest that you use low values, such as 0.1 for both. The program will re-perform Holt's method with various combinations of smoothing parameters until the pair which give the lowest mean squared error are found. The results displayed will be the chosen formula, and the mean squared error and mean absolute percentage error which it gives.

The results produced by a Classical Decomposition analysis will be the trend equation, a table of seasonal factors, the mean squared error, and the mean absolute percentage error.

Another way of analysing a time series is to perform a simple regression, using time as the dependent variable. The Applications Program can be asked to do this by using the causal modelling option, and the appropriate procedures are outlined in section 9.9.

9.8 Causal Modelling

The **Forecaster** multiple regression analysis routine is designed for use either by people who are unsure of the form of the causal model they are investigating or by others who are confident that they do understand it. The needs of both groups can be satisfied. When you select **CAUSAL MODEL**, you will be presented with a six option menu. Each of the options is described below:

PRESELECTED All of the variables selected by the user, as transformed, will be included in the regression model without question. This is the fastest method, and can be used for preliminary investigations, and whenever the user is confident that he is able to make his own selection of the appropriate variables.

STEPWISE The variables selected and transformed by the user will be examined, and only those which have a statistically significant effect on the dependent variable will be included in the model. All of the other variables will be rejected. Indeed, it is possible that the program will reject all variables, in which case the message **NONE OF THE VARIABLES ARE SIGNIFICANT** will be displayed.

AUTOMATED The variables selected by the user will be examined; a range of transformations will be applied by the computer, and the model will be constructed which best predicts the dependent variable. If the user has transformed any of the selected variables, the user's chosen transformation will be accepted. Otherwise, the transformations that the computer will apply will be; natural logarithm, exponential, reciprocal and a lag of 1. This method is extremely useful if you have no idea about how to start transforming the data, but it can be a very slow process.

STATISTICS This option can only be selected once a regression has been performed. It will display a full analysis for those people who have statistical training. For the lay forecaster, you can use this display to discover the R-squared value, MSE and the difference between forecast and actual values for the data used. This difference is known as the **residual**.

PRINTOUT Again, this can only be selected if a regression has been performed. The computer will print a copy of all of the statistical data, together with a description of the model produced by the regression.

RETURN Takes you back to the main menu.

The actual regression analysis will be very slow, particularly if the third option is selected. This is because of the vast range of alternatives which will be tested. Sometimes you will certainly have time to go away and have a cup of tea. Whenever the computer is working on lengthy calculations, a flashing **CALCULATING** message will be displayed on the screen.

When the processing is completed, the screen will display the selected regression model, together with a qualitative statement of how reliable it is likely to be, and a suggestion on what to do next. These screen messages are explained in section 9.10. If you want a more detailed statistical analysis, select **STATISTICS** on the menu described above.

9.9 Time Series Regression

If you have read Chapter 6 of the teaching text, you will recall that it is possible to use a simple regression model to forecast a time series. This is particularly useful if the time series is obviously linear, or following an

exponential, or power curve. To perform a time series regression, you set the time series as the dependent variable, and time as the independent variable. If you wish, you can perform a transformation on the time variable. The precise steps you must take to use this program to run a time series regression are as follows:

- 1) Enter the time series data using the **ENTER DATA** option. You may also wish to enter data for other variables at this time if you intend to test some form of causal model as well.
- 2) Select the **EDIT DATA** option. Use this to ensure that the time series is marked as the dependent variable, and that any other variables are de-selected.
- 3) Remaining in the **EDIT DATA** option, key **T** for **TIME**. This will create a new independent variable at the left of the screen, named **PERIOD**, which simply numbers the data from 1 onwards. This is automatically selected as the independent variable.
- 4) Return to the main menu. If you wish, you can use the **TRANSFORM** option to apply any of the transformations to the **PERIOD** variable.
- 5) Use the **CAUSAL MODEL** option to perform the regression. You can use any of the three options available, but we suggest that the **AUTOMATED** facility offers the greatest benefit.

The results will be displayed in the way described in section 9.8.

9.10 Analysing the Regression Model

Once a regression analysis is complete, one of three messages will appear on the screen, as follows:

NONE OF THE VARIABLES ARE SIGNIFICANT
– try another selection

In this case, no useful model can be found. If you used the **AUTOMATED** option, then there is probably no way to make these variables fit a model. If you ran the **STEPWISE** facility, then it could be that transforming the variables in some way could produce a result. If you used **PRESELECTED**, then try an **AUTOMATED** or **STEPWISE** analysis before you discard these variables.

THIS MODEL HAS SOME VALUE
– there could be a factor missing

The model produced is better for making predictions than no model at all. It may be that a transformation of one or more of the variables will yield a better result, or you could consider adding another variable. If you did not use **STEPWISE**, try this, because the model could then be improved if the computer chose to discard one of the variables you have included.

THIS APPEARS TO BE A GOOD MODEL
– refer to checklist in 7.2

You have found a very good model, but you should confirm that you have taken all of the items in section 7.2 into account.

9.11 Making a Forecast

The actual process of producing a forecast is very simple, once the model has been developed. Select the **FORECAST** option on the main menu. If you have used the **TIME SERIES** option to produce the forecasting model, you will simply be asked for the number of periods ahead for which you wish to produce a forecast. The forecast figures will then be presented on the screen. You can then press **ENTER** (or **RETURN**) to go back to the main menu.

If a causal model has been developed, you will be required to enter values for all of the independent variables for each prediction you wish to make. Follow the instructions on the screen to do this. When you have performed all of the required forecasting, simply press **ENTER** (or **RETURN**) to go back to the main menu.

Remember that, before you use this routine, you must have developed a model first, using the **TIME SERIES** or **CAUSAL MODEL** option. If you enter data, develop a model, and then save the data as a file, the last model you produced will be saved with it. Therefore if you load a file of this type back into the program, you will be able to use the **FORECAST** option immediately if you wish.

A FINAL TIP

Although the error data supplied with the two types of model can give you some confidence in the final method you select, there is no better test than comparing your prediction with the actual outcome. You may be happy to wait for the next period, be it a day, week, month or year away, to see how good the forecast really is. On the other hand, you may not. If you fall into this latter category, we suggest the following solution. Develop your model without using the last two or three readings in the series of data you have collected; you can de-select them on the **EDIT DATA** option. Then use the model to make predictions for these last few readings. If these forecasts fall within the tolerance that you require, you can have added confidence in the model. But remember to add back the readings you deleted before you make your real forecasts.

APPENDIX 1

Applications Program Keywords

A1.1 The Main Menu

- LOAD DATA** Loads in a data file from tape or disk. The specific procedures vary from one version of the program to another, so follow the instructions as they appear on the screen.
- SAVE DATA** Saves a data file to tape or disk. Again, the procedures will vary from one program version to another. Follow the on-screen instructions.
- ENTER DATA** Allows you to enter a new set of data into the computer. If you already have some data in the computer, use the **EDIT DATA** option to add more variables or readings.
- EDIT DATA** Enables you to make changes, additions and deletions to the data you have already entered. Refer to A1.2 for the full details of the editing options available once you have made this selection.
- TRANSFORM DATA** Allows you to alter any of the variables by using a series of functions to transform the data. Full details of the transformations available are given in section A1.3.
- TIME SERIES** Performs a time series analysis on the dependent variable using Holt's method or, if there is sufficient data, Classical Decomposition.
- CAUSAL MODEL** Performs a regression analysis on the variables selected in the **EDIT** option. Full and stepwise options are available, as well as an automatic transformation facility – refer to section A1.4 for details.
- FORECAST** Produces a forecast based upon the time series or regression model created by one of the two previous commands. You will be asked for the period number or variable data required to undertake the forecast.

A1.2 Editing

INSERT To insert a new line of data below the current cursor line, or a new column to the left of the current cursor column. On selection, you will be asked whether you want a **ROW** or **COLUMN**.

ERASE This erases the current cursor row or column. Use this command with care, as the data in an erased line cannot be recovered. On selection, you will be asked whether you want a **ROW** or **COLUMN** erased.

DEPENDENT Makes the current cursor column variable the dependent variable. There can only be one dependent variable, so this will automatically de-select the previous dependent. The current dependent variable is the one with all values shown in inverse (red on a white background).

SELECT Selects the current cursor column variable or current cursor row for inclusion in the time series or causal analysis, or if it has already been selected, it de-selects it. As many or as few variables and readings as you wish can be included. The currently selected variables are identified by having the name at the top of the column in inverse (blue on a white background), and the currently selected readings are likewise shown in inverse in the left hand column. Press **R** or **C** to indicate **ROW** or **COLUMN** when requested.

ALTER Allows you to alter the value at the current cursor position. You will be requested to enter a new value at the top of the screen.

TIME Inserts a "time" variable in the first column. This is required if a regression (causal) analysis of a time series is to be attempted. The variable takes the values 1 to n, where n is the number of readings.

PRINT Prints a copy of all the data you have entered, as well as showing any transformations you have applied.

MENU Returns you to the main menu (see section A1.1).

A1.3 Transformations

Performing transformations can be a complex task. Therefore, if you have any doubt about the way this process operates, we suggest that you refer to sections 6.4 and 9.6 for a full explanation. The options available on the transformation menu are as follows:

RETURN Takes you back to the main menu.

NORMAL Returns the data to its original form, but only if a previous transformation has superceded the original data.

LOG Takes the natural logarithms of the data.

EXP Takes the exponential values of the data.

POWER Raises the data to a power. The power to be used will be requested. Using a power of minus one will produce the inverse values of the data.

ROOT Takes the root of the data. The root to be used will be requested.

LAG Lagging of a time series is the process of moving the column of data back so that it no longer relates to period to which it refers. The number of periods the data is to be lagged will be requested. Lagging is only used when attempting to develop a causal model to predict a time series. It is inappropriate for other situations.

A1.4 Causal Models

Three alternative regression processes are available as follows:

PRESELECTED All of the variables selected by the user, as transformed, will be included in the regression model without question. This is the fastest method, and should be used for preliminary investigations, and whenever the user is confident that he can select the appropriate variables.

STEPWISE The variables selected and transformed by the user will be examined, and only those which have a statistically significant effect on the dependent variable will be included in the model.

AUTOMATED The variables selected by the user will be examined; a range of transformations will be applied, and the model will be constructed which best predicts the dependent variable. If the user has transformed any of the selected variables, the user's chosen transformation will be accepted. **** WARNING ** - this can be a very slow process.**

STATISTICS This option can only be selected once a regression has been performed. It will display a full analysis for those people who have statistical training. For the lay forecaster, you can use this display to discover the R-squared value, MSE and the difference between forecast and actual values for the data used. This difference is known as the **residual**.

PRINTOUT Again, this can only be selected if a regression has been performed. The computer will print a copy of all of the statistical data, together with a description of the model produced by the regression.

RETURN Takes you back to the main menu.

APPENDIX 2

Loading Procedures

A2.1 Spectrum Cassette Version

The Teaching Program is on Cassette No. 1 and the Applications Program is on Cassette No. 2. Both programs load and auto run by typing **LOAD ""** and **ENTER**.

A2.2 Spectrum Microdrive Version

The Teaching and Applications Programs both start on Cartridge No. 1. Simply insert Cartridge No. 1 in the Microdrive and key **NEW ENTER RUN ENTER**. An option to run the Teaching or Applications Program will then soon appear on the screen. However, the Teaching Program menu routine may ask you to insert Cartridge No. 2 in the Microdrive to load certain parts of the course. Please obey the instructions which you see on the screen.

A2.3 Commodore 64 Cassette Version

The Teaching Program is on Cassette No. 1 and the Applications Program is on Cassette No. 2. Both programs load and auto run simply by pressing **SHIFT** and **RUN/STOP** at the same time.

A2.4 Commodore 64 Disk Version

The Teaching Program operates by typing **LOAD "TEACH", 8 RETURN**. When the program has loaded, a **READY** message will appear on the screen. You must then type **RUN** and **RETURN** to operate the program. The Applications Program is loaded in the same way. Type **LOAD "APPLY", 8 RETURN** and then **RUN RETURN** when the **READY** message appears.

A2.5 BBC Model B Cassette Version

The Teaching Program is on Cassette No. 1 and the Applications Program is on Cassette No. 2. Both programs load and auto run simply by typing **CHAIN ""** and **RETURN**.

A2.6 BBC Model B Disk Version

Load the Teaching Program by typing **CHAIN "TEACH" RETURN**, and the program will run automatically. Similarly, load the Applications Program by typing **CHAIN "APPLY" RETURN**.


APPENDIX 3

Using Graphs

A Graph is a way of representing numerical information in a diagrammatic form. There are a number of types of graph, but all work on the same basic principals. We will consider the simpler forms of graph. Basically a graph can be used to display the relationship between two groups of information. One of the groups may be some form of time scale, such as the date, and the other group in this case will be a time series.


The basic framework of a graph consists of two lines, drawn at right angles, each of which is called an **axis**. The plural form of axis is **axes** (pronounced "axeese"). In most business applications the axes are drawn along the bottom and up the left hand side of the work area. The point where they cross is called the **origin**. By tradition, the axis along the bottom of the work area is called the "x-axis", and the one up the left hand side is called the "y-axis". One axis is then divided into parts to represent one set of information, and the other axis is divided to represent the other set. If one axis is being used to represent time, such as dates, weeks, months, etc., then it is normal to select the horizontal axis (or x-axis) for this. The way the axis is divided is called the **scale**, and the scale must be carefully selected to include all values of the information it is intended to represent.

When the scales have been marked on the axes, the information can be marked or **plotted** on the graph. We will assume that you are plotting a graph of sales of Ray Jenkins' dishwashers over a period of weeks. The **data** you have and the graph produced are as follows:



If you look at how each point is plotted, the process involved can be seen to be very simple. First the value for the first piece of data is selected. This value is found on the horizontal scale. Then the corresponding value for the second piece of data is found on the vertical scale. An imaginary line is drawn up from the value on the horizontal scale, and across from the vertical scale. Where the two lines cross, the point is marked. When all the points have been marked in this way, they are connected by a series of straight lines to form the graph, known as a **line graph**. If the points are left unconnected, the graph is called a **scatter graph**, or **scatter diagram**. The scatter diagram is useful for making a preliminary examination of the data. Joining the points into a line graph is helpful for attempting to identify any pattern.

Another form of graph frequently used in business is called the **bar graph**. To generate a bar graph, plot the points in the usual way. Instead of joining them to each other, draw a vertical bar from the bottom axis up to the point. Bar graphs create a much stronger image than a simple line graph, and are thus very useful for making presentations of data to other people. The data from the previous page would look like this if presented in bar graph form:



Glossary

Average, the "typical" value of a **Variable**, calculated by adding all of the values together and dividing by the number of values.

Causal Modelling, developing a mathematical formula which predicts the value of something based upon the factors which are believed to cause it to happen.

Classical Decomposition, a **Time Series** forecasting method which "decomposes" the time series into its component parts; a **Trend**, a **Seasonal** effect and a **Cycle**.

Constant, an item which has an unchanging value; the number of inches in a foot is a constant, with a value of 12.

Cycle, a regular repeating pattern of up and down movements in the values of a **Time Series**.

Data, a collection of information about a particular object or activity.

Dependent Variable, the **Variable** which is being predicted or forecast. Its value is "dependent" upon the values of the other variables in the forecasting **Model**.

Double Moving Average, a **Moving Average** of a moving average.

Exponential Smoothing, a **Time Series** forecasting method which produces a new forecast on the basis of the previous forecast and the most recent actual value of the item being forecast.

F Test, a measure of whether a **Model** produced by **Regression Analysis** is of any value for prediction purposes.

Graph, a diagrammatic display of the relationship between sets of **data**.

Holt's Method, a form of **Exponential Smoothing** which can be applied to a **Time Series** with a **Trend**.

Independent Variable, a **Variable** used in a forecasting **Model** to calculate the **Dependent Variable**. Its value does not "depend" upon the other variables involved.

Linear Regression, a simple **Regression Analysis** of the relationship between two variables which produces a straight line when plotted on a **Graph**.

Mean Squared Error, a measure of the accuracy of the forecast. A good forecast should have an MSE which is low when compared with the typical values of the **Dependent Variable**.

Model, in mathematical terms, a "model" is simply a formula (or sometimes a series of formulae) which explains how a group of **Variables** relate to each other.

Moving Average, an **Average** of a specific number of the most recent values of a **Variable**. As a new value occurs, it is included in the average, and the oldest value is dropped out. Usually the average is of the 3 or 5 most recent values.

Multiple Regression, a **Regression Analysis** which takes into account three or more **Variables**.

Qualitative Methods, an analytical approach in which the result is achieved by reasoning, rather than by using a mathematical formula.

Quantitative Methods, an analytical approach in which the result is achieved by using mathematical techniques.

Regression Analysis, a method of determining a formula to explain the relationship between two or more **Variables**.

R Squared, a measure of how well a **Regression Analysis** fits the data you have used. An R Squared of 1 is perfect, and an R Squared of 0 means that the **Model** is worthless.

Scatter Diagram, a simple form of **Graph** on which the data points are not connected together. It is useful for identifying broad relationships between variables.

Seasonality, a (possibly complex) pattern of up and down movements in the values of a **Time Series** which repeats at regular intervals; usually annually, but it could be a shorter period such as monthly or weekly. Patterns lasting more than a year are normally described as **Cycles**.

Smoothing, modifying a **Time Series** to remove random fluctuations, and thus to reveal the underlying **Trend**.

Stationary Series, a **Time Series** which, allowing for some random fluctuations is not increasing or decreasing.

Time Series, a list of values of a **Variable** recorded at regular intervals over a period of time.

Transformation, modifying the values of a **Variable** by applying the same mathematical to every value.

Trend, a relatively uniform, continuous change in a **Time Series**.

Variable, a measurable activity, event or outcome, which does not remain the same.

Bibliography

1. Coping With Numbers, by David Targett, Martin Robertson 1983.
2. Forecasting Methods in Business and Management, by Michael Firth, Edward Arnold 1977.

Index

- Analysing the Regression Model, 46
- Applications Program Keywords, 49
- Average, 15, 56
- Axis, 28, 54
- Bar Graph, 55
- Causal Modelling, 10, 11, 32, 44, 52, 56
- Centering, 20
- Classical Decomposition, 25, 56
- Constant, 14, 56
- Cycle, 26, 56
- Data, 56
- Delphi, 13
- Dependent Variable, 14, 56
- Double Exponential Smoothing, 22
- Double Moving Average, 19, 56
- Editing Data, 40, 50
- Entering Data, 39
- Exponential Smoothing, 17, 56
- File Handling, 40
- Fitting a Line, 29
- Forecasting, 47
- Forecasting Methods, 10, 37
- Forecasting Process, 36
- F Test, 35, 56
- Graph, 54, 56
- Historical Analogy, 13
- Holt's Method, 22, 56
- Independent Variable, 14, 56
- Lagging, 34
- Least Squares Method, 29
- Linear Regression, 23, 28, 56
- Line Fitting, 29
- Line Graph, 55
- Loading Procedures, 53
- Market Research, 13
- Mean Absolute Percentage Error, 24
- Mean Squared Error, 24, 34, 56
- Menu
 - Applications, 38, 49
 - Teaching, 8
- Model, 32, 56
- Modelling Checklist, 33
- Multiple Regression, 32, 57
- Non Linear Regression Models, 30
- Origin, 54
- Panel Consensus, 13
- Qualitative Methods, 10, 12, 57
- Quantitative Methods, 10, 57
- Regression Analysis, 23, 28, 32, 57
- Residuals, 29
- R Squared, 34, 57
- Scale, 54
- Scatter Diagram, 15, 55, 57
- Seasonality, 25, 57
- Simple Regression, 23, 28
- Smoothing, 16, 57
- Smoothing Constant, 17
- Stationary Series, 14, 57
- Straight Line, 29
- Teaching Method, 7
- Testing the Model, 34
- Time Series, 10, 11, 14, 43, 57
- Time Series Regression, 45
- Transformation, 30, 42, 51, 57
- Trend, 19, 57
- Variable, 14, 57
- Weighting, 18



TRIPTYCH PUBLISHING LIMITED. Sterling House, Station Road, GERRARDS CROSS, Bucks SL9 8EL.

FORECASTER

It has always been man's dream to predict the future and many have tried. They would have fared better with FORECASTER.

FORECASTER takes most of the guesswork out of forecasting because it accepts that tomorrow's events are best predicted on the basis of today's facts. To do this FORECASTER shows you how to assimilate and express your current knowledge ready for entry into your micro. Then by following one of FORECASTER'S special routines, you will arrive at the most logical and accurate extension of the available data. Armed with FORECASTER the assessment of business trends becomes that much easier. Key issues such as sales targets, growth trends, population figures and store traffic can be accurately and quickly predicted. All of which can only result in your profit forecasts looking even brighter.

THE BRAINPOWER RANGE

With the BRAINPOWER range you have, for the first time, the intelligent application of your micro. And for the first time, involved and even difficult subjects are made easy and stimulating.

The key to this is the three vital elements to be found in this pack.

The TEACHING PROGRAM and TEXT BOOK which together provide the simplest but most comprehensive introduction to the subject, and the APPLICATIONS PROGRAM itself which unlocks the power of your micro.

The entire BRAINPOWER range has been conceived by experts with the aim of making it simple to understand and easy to apply.

© 1984. TRIPTYCH PUBLISHING LIMITED. All rights reserved. No part of this program or accompanying literature may be duplicated, copied, transmitted or otherwise reproduced by any means, electronic, mechanical, photocopying, recording or otherwise without the express written permission of Triptych Publishing Limited. This program and book are sold subject to the condition that they shall not, by way of trade or otherwise, be lent, resold, hired out, or otherwise circulated without the publisher's prior written consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser. BRAINPOWER is the trade mark of Triptych Publishing Limited.

ISBN 1-85050-260-9



9 781850 502609

BRAINPOWER

Application through learning