

Time series techniques for economists


TERENCE C. MILLS



Time series techniques for economists

TERENCE C. MILLS

Midland Montagu Centre for Financial Markets,
City University Business School



CAMBRIDGE UNIVERSITY PRESS

Cambridge
New York Port Chester
Melbourne Sydney

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40, West 20th Street, New York, NY 10011, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Cambridge University Press 1990

First published 1990

Printed in Great Britain by the University Press, Cambridge

British Library cataloguing in publication data

Mills, Terence C.

Time series techniques for economists

1. Economics. Time series. Analysis.

Econometric models

I. Title

330'.01'51955

Library of Congress cataloguing in publication data

Mills, Terence C.

Time series techniques for economists/Terence C. Mills.

p. cm.

Bibliography.

1. Economics, Mathematical. 2. Time-series analysis. I. Title.

HB135.M54 1990

330.028--dc20 89-7187 CIP

ISBN 0 521 34339 9

Preface

Although I specialised in econometrics as a student at the universities of Essex and Warwick in the early 1970s, and was appointed as a lecturer in econometrics at the University of Leeds in 1976, I first became acquainted with the theory and techniques of time series analysis while completing my doctoral thesis in the late 1970s. I suspect that this relatively late exposure to such an important area of statistics was a common occurrence amongst many econometricians of my generation.

My time series education, both in terms of theory and practice, was significantly accelerated by the time I spent in the Economics Division of the Bank of England during the period 1980–2, where I was able to benefit from the expert guidance and tuition of Peter Burman.

Convinced of the importance of time series techniques in applied econometric analysis, I returned to Leeds to introduce a significant time series component into the Applied Econometrics course in the B.A. degree in Economics and Econometrics then available but, alas, no longer offered. At the same time, a course in Time Series Analysis was also developed for the M.Sc. in Economic Forecasting, which has also fallen by the wayside in recent years. My thanks to the students of those courses, and also those of the M.Sc. in Statistics, who, unwittingly, suffered early drafts of what has turned out to be the basic material of many of the chapters of this book.

Notwithstanding these teaching ventures, and the vast intellectual debt I obviously owe to such pioneers as George Box, Gwilym Jenkins, Clive Granger and George Tiao, the major impetus to my interest and enthusiasm for time series techniques in economics has been the collaboration over the years with a number of colleagues whom I regard not just as co-research workers but as close friends: Forrest Capie, Nick Crafts, Steve Leybourne, Mike Stephenson, and Geoffrey Wood. My thanks go to all of them for making joint research such an enjoyable and fruitful experience.

At this point I must also pay tribute to the various members over the years of the 'Fenton Research Centre'; in addition to Nick, Steve and Mike mentioned above, these must include Garry Phillips, Brendan McCabe, Andy Tremayne, Karen Bevins, John and Jane Binner, Mike Hudson, and Louise Webster. Although they are now widely scattered, my thanks go to them all for helping to provide an island of sanity in an environment that was often hostile to serious economic research.

More specific thanks go to Pat Hatton for her expert typing and to Patrick McCartan, Susan Beer and Carmen Mongillo at Cambridge University Press for their input into this finished product. Finally, but of course they should be at the head of any list of acknowledgements, my thanks and love go to my family, Thea and Alexa, to whom this book is dedicated.

Contents

<i>Preface</i>	<i>page</i> ix
1 Introduction	1
PART I EXPLORATORY ANALYSIS OF ECONOMIC TIME SERIES	5
2 The graphical display of time series	7
3 Summarising time series	19
4 Transforming and smoothing time series	40
PART II THE MODELLING OF UNIVARIATE ECONOMIC TIME SERIES	61
5 Stationary stochastic time series models	63
6 Modelling nonstationary processes	92
7 Forecasting using ARIMA models	104
8 ARIMA model building	116
9 Exponential smoothing and its relationship to ARIMA modelling	153
10 Modelling seasonal time series	164
11 Further topics in univariate time series modelling	199
PART III THE MODELLING OF MULTIVARIATE ECONOMIC TIME SERIES	233
12 Intervention analysis and the detection of outliers	235
13 Transfer function–noise models	248
14 Multiple time series modelling	281

PART IV NONLINEAR TIME SERIES MODELS	323
15 Conditional variance models and related topics	325
16 State dependent models	342
<i>References</i>	349
<i>Index</i>	371

1 Introduction

The empirical analysis of time series has a long tradition in economics, the historical development of which has been surveyed by both Spanos (1986, chapter 1) and Nerlove et al. (1979, chapter 1). While the basis of time series analysis was laid in the inter-war period by Yule (1927), Slutsky (1937) and Wold (1938), the main thrust of econometric research was aimed at integrating the 'Fisher paradigm' – the amalgamation of descriptive statistics and the calculus of probability within the framework of the linear regression model – into econometrics (Koopmans (1937), Haavelmo (1944)). This was taken an important stage further by Haavelmo's (1943) formulation of the simultaneous equation system, the statistical analysis of which provided the research agenda for the Cowles Foundation Group. This distinguished collection of statisticians and econometricians introduced and developed new techniques of estimation and hypothesis testing, based on the concept of maximum likelihood, within the framework of the simultaneous equation model. Their results, collected in Koopmans (1950) and Hood and Koopmans (1953), provided the main research program in econometrics for the next 25 years, which was thus dominated by the linear regression model and its associated misspecification analysis, and the simultaneous equations model and its identification and estimation.

Research in time series analysis was primarily concerned, on the other hand, with developing techniques for modelling a single series. Economists were particularly concerned with the decomposition of a series into a set of unobserved components, traditionally taken to be the trend, cycle, seasonal and irregular, these being associated with the ideas of secular evolution (or long swings), the concept of the business cycle, seasonal variation, and transitory influences, respectively. Operations researchers, on the other hand, were primarily concerned with forecasting, developing techniques which were computationally feasible and based upon local trends and levels (Holt et al. (1960), Winters (1960), Brown (1963)). These

'exponential smoothing' techniques were later shown to have unobserved component representations (see, for example, Harrison (1967)), hence providing a more unified framework for analysis.

Thus, by the late 1960s the research of econometricians and time series analysts seemed to be so far apart that no synthesis of the two frameworks would have appeared possible. This dichotomy was, though, to change dramatically in the years after the publication in 1970 of the original edition of Box and Jenkins' (1976) famous book in which a feasible model building procedure for the general class of autoregressive-integrated-moving average processes was developed.

The practical manifestation of the Cowles Foundation's research program was the development of large-scale macroeconomic models designed for forecasting and policy simulation; see, for example, Klein and Burmeister (1976) for a survey and evaluations of a number of such models. The appropriateness of this form of model began to be questioned when they were found to be unable to compete with Box-Jenkins models in terms of forecasting comparisons; see Cooper (1972), Naylor et al. (1972) and Granger and Newbold (1975). Macroeconomic model builders responded by arguing that univariate time series models were irrelevant for policy considerations: that macroeconomic models could also be was forcefully argued by Lucas (1976) and Sims (1980).

The importance of a synthesis between traditional econometric and time series techniques was brought into even sharper focus by Granger and Newbold's (1974) illuminating analysis of the spurious regression problem, in which it was shown that regressing independently generated random walks would lead to a very high probability of rejecting the correct hypothesis of no relationship between the two series. The groundwork for such a synthesis was laid by Zellner and Palm (1974) and further developed in Wallis (1977) and Nerlove et al. (1979). A great deal of research in the last decade has thus been directed at producing a methodology for the econometric modelling of economic time series. One strand of this literature is probably best encapsulated in the papers by Hendry and Richard (1982, 1983), with a complete textbook treatment of this methodology being provided by Spanos (1986). Other viewpoints are provided by Sims (1980) and Leamer (1983), and current accounts of these competing views are to be found in Hendry (1987), Leamer (1987) and Sims (1987), with a review being provided by Pagan (1987). Although many arguments remain to be resolved, this general methodology does provide an important synthesis of traditional econometric and time series techniques.

Nevertheless, even given the undoubted success of this new synthesis, economists face many problems in the analysis of time series which do not

fall comfortably, if at all, into the (single or simultaneous equation) regression framework, and this necessarily leads to the examination of a much wider class of models and techniques. Serious difficulties may be encountered in this enterprise for, since time series occur naturally in a wide range of disciplines, so the literature on time series techniques appears in a wide range of journals. Many of these may be inaccessible to economists, both physically and in terms of the notation and terminology used (this is particularly so for the time series literature in the electrical engineering field, as any economist who has attempted to read articles in control theory journals will surely testify!). Furthermore, the increase in the sheer volume of literature over the last decade has been quite phenomenal; as an indication of this, of the more than 400 citations referenced in this book, 62% have been published since 1980. Many economists may therefore be simply unaware of some of the new techniques and insights that may be of potential importance to the analysis of economic time series. The primary aim of this book is, therefore, to bring together those techniques that we feel are of importance to analysts of economic time series, to indicate by discussion and examples why they are of importance, and to provide a consistent notation and terminology by which they may be developed.

It may be important to state at this juncture what the book is *not* about. It is not an econometrics text, and therefore it does not cover regression analysis and its associated simultaneous extensions in any formal way. For such treatments see, for example, Judge et al. (1985) and Spanos (1986). Neither is it a textbook on forecasting techniques, nor does it discuss in great detail the 'Box-Jenkins' approach to analysing time series: for treatments of the former, see Abraham and Ledolter (1983) and Makridakis et al. (1983), for the latter, see Vandaele (1983) and Pankratz (1983). It is unashamedly non-rigorous; formal treatments of the theory of time series are provided by Anderson (1971) and Fuller (1976). Finally, it does not discuss spectral analysis. Our experience suggests that the frequency domain is unfamiliar territory for the majority of economists and analysis is therefore carried out entirely in the more natural time domain.

The book is structured in the following way. After this introductory chapter, the rest of the material is divided into four parts. The first part, comprising three chapters, is devoted to an area of analysis covered by few textbooks, that of the exploratory analysis of time series, a topic inspired by the work of Tukey (1977) (sometimes also referred to as initial data analysis, see Chatfield (1985)). Part II is devoted to the analysis of univariate time series and, as befits the most completely developed area of the subject, is the longest part of the book, containing seven chapters.

Multivariate time series are analysed in Part III, in which there are three chapters, while the final part, just two chapters, discusses non-linear time series models. Each part begins with an introductory discussion setting out the subsequent material.

Throughout the chapters, numerous examples are provided to illustrate the techniques discussed. Certain time series are used in a sequence of examples, thus showing how applied analysis progresses and evolves as the theoretical base of material becomes richer. Other examples draw heavily on the author's own research, both published and unpublished.

One final note of caution. Anyone familiar with the time series literature will be aware of, and perhaps irritated by, the vast number of acronyms this literature has spawned and, indeed, they have inspired a journal article (Granger (1982)). Irritating or not, such acronyms are invaluable to the textbook writer and we make no apologies for making liberal use of them throughout the book!

Part I

Exploratory analysis of economic time series

As a consequence of the highly influential book by Tukey (1977), exploratory data analysis, often referred to by its acronym EDA, has gained enormously in importance in all areas of applied statistical research over the last decade. Time series analysis is no exception and, indeed, time series data lend themselves extremely well to such exploratory analysis. It appears to be the case, however, that texts on time series and forecasting have not in the past dwelt overmuch on this aspect of data analysis, preferring to emphasise the more formal aspects of modelling techniques, although a notable exception to this is Levenbach and Cleary (1981).

This part of the book is explicitly concerned with three areas of exploratory time series analysis: (i) graphical display, both of one variable and scatter plots of the relationships between two or more variables (Chapter 2), (ii) summarising time series in terms of frequency distributions and measures of location and dispersion (Chapter 3), and (iii) the transformation and smoothing of time series to provide easier forms of analysis (Chapter 4).

Although introduced as exploratory data analysis, many of the techniques discussed here will also be used, and may recur in different guises, throughout the remaining parts of the book dealing with more formal modelling techniques. Indeed, it cannot be stressed too highly the importance of good preliminary data analysis: it will lead to better models, more efficient computing and a greater understanding of the relationships between the data at hand, the underlying economic theory, and the modelling techniques employed.

2 The graphical display of time series

2.1 Introduction

It is often the case that graphical data displays are easier to interpret than tabular representations of the same data. It is also the case that such graphical displays can be unhelpful, exaggerated and even deliberately deceiving. As Tufte (1983) says, in a highly entertaining and provocative book on the visual display of data,

Much of...statistical graphics has been preoccupied with the question of how some amateurish chart might fool a naive viewer...At the core of the preoccupation with deceptive graphics was the assumption that data graphics were mainly devices for showing the obvious to the ignorant. It is hard to imagine any doctrine more likely to stifle intellectual progress in a field.

(Tufte (1983, page 53)).

The revolution in statistical graphics came in the late 1960s, primarily through the work of John Tukey (see Tukey (1972) and Tukey and Wilk (1970)) and has since led to such masterpieces of graphical excellence as those in Tukey (1977), Chambers et al. (1983), and Cleveland (1985). Cleveland (1987) provides a survey of recent research in statistical graphics.

Graphical displays of time series data are the most frequently used form of graphic design and, as such, are subject to misuse and abuse just as any other form of visual display. They can also be one of the most powerful techniques of exploratory analysis available and, in the hands of a master such as Tukey, become a means of conveying complex and changing relationships in a simple yet flexible manner. A great deal of research has been carried out recently on the theory of graphical perception, leading to a much deeper understanding of the scientific foundations underlying graph construction: see, for example, the survey by Cleveland (1985, chapter 4) and the recent research by Cleveland and McGill (1987).




Exhibit 2.1 UK total final expenditure

Hopefully, such research findings will quickly find their way into the important area of statistical, and hence time series, graphics.

2.2 Time series plots

A time series plot is simply a graph in which the data values are arranged sequentially in time. Usually, time series are generated by taking observations at equally spaced time intervals, and hence these values must be plotted at equally spaced intervals along the time axis, conventionally taken to be the horizontal. Although the construction of a time series plot would appear to be a trivial exercise, certain commonsense guidelines concerning labelling of axes and variables, choice of scale, use of grid lines, and so on, should be adhered to. Schmid (1983, pages 17–19) and Cleveland (1985, chapter 2, particularly pages 100–1) provide convenient lists of ‘good practices’ to follow when designing time series plots.

Time series of economic data display many different and important characteristics and a plot is an effective way of quickly perceiving the evolution of a single, or a group of, time series: for an interesting case study, see Chatfield and Schimek (1987). Exhibit 2.1, for example, plots nominal and real UK Total Final Expenditure (TFE) annually for the period 1948 to 1984. It is immediately apparent that nominal TFE displays a very smooth upward trend which has steepened considerably during the observation period, a consequence, of course, of it being denominated in current prices. The real TFE series, deflated using 1980

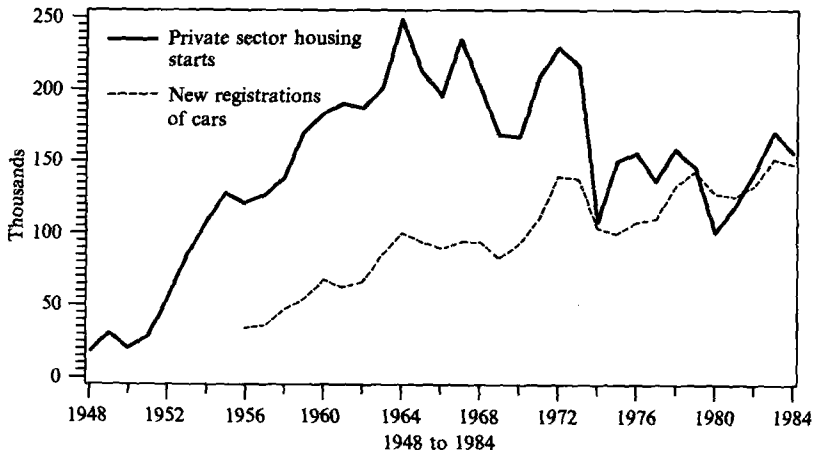


Exhibit 2.2 UK housing starts and new registrations of cars




Exhibit 2.3 UK investment in manufacturing

prices, displays a less pronounced trend, and is also rather less smooth, showing a number of cyclical movements.

Not all annual economic time series evolve so smoothly. The new registration of cars in the UK, shown in Exhibit 2.2, has a general upward trend but is dominated by short-term erratic fluctuations around this underlying movement. The UK private sector housing starts series, on the other hand, which is also shown in Exhibit 2.2, displays no general trend

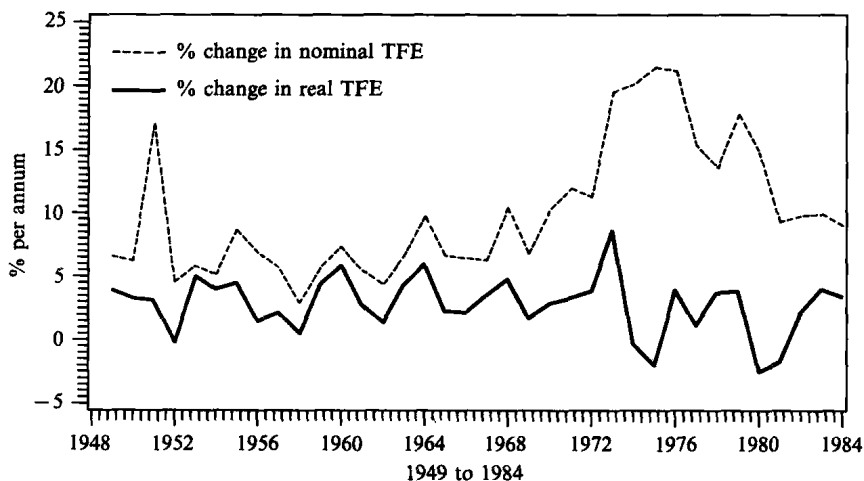


Exhibit 2.4 Annual growth in UK total final expenditure

but fluctuates widely around a mean level that itself changes fairly often. Nominal gross fixed investment in manufacturing industry, shown in Exhibit 2.3, has both a cyclical pattern and a trend that alters dramatically during the 1970s. Also shown in Exhibit 2.3 is the real version of this series, which may be best characterised as fluctuating around a constant mean level.

Transformations of time series can alter their plots dramatically. This is effectively shown by Exhibit 2.4, which plots the annual growth rates of the nominal and real TFE series whose levels were shown in Exhibit 2.1. The very smooth evolutions of these two series are, in fact, revealed to contain widely fluctuating growth rates with, for example, nominal TFE growth ranging from 2.8% in 1958 to 21.3% in 1975.

Economic time series are often observed over shorter intervals than a year and most applied research is performed using data collected at either a quarterly or a monthly time interval. Time series of this type display a further characteristic, other than trend and cyclical movements, which is not present in annual data, that of seasonality. Exhibit 2.5 displays quarterly time series from 1955 to 1984 of both the value and volume of UK retail sales. Both series are dominated by seasonal effects, a consequence of the annual Christmas 'spending spree', but whereas the volume series displays a fairly stable seasonal pattern superimposed on a slowly increasing trend, the value series has a seasonal pattern whose amplitude increases substantially as the trend steepens. When monthly time series are considered, it is also possible to observe 'double-seasonal' patterns. The bank lending to the UK private sector series shown in